

LIMITATIONS ON GEOTECHNICAL RISK MANAGEMENT: DESIGNING FOR RESILIENCE

Tim Davies

Geological Sciences, University of Canterbury, NZ

ABSTRACT

The imprecisions inherent and unavoidable in probabilistically-based risk analysis and management are addressed and shown to be a significant limitation on reliable infrastructure design against the loads resulting from hazard events. The nature of hazard events is outlined and the concept of resilience defined in the context of poorly-quantified and unknown hazards. It is suggested that resilience might be improved by purposefully designing infrastructure to match complex system attributes such as decentralisation, heterogeneity and redundancy. The use of event and effects scenarios in choice of a design event in geotechnical engineering is outlined, and the implications of this strategy explored.

1 INTRODUCTION

The increasing technical ability of geotechnical engineers to characterise and model the nature of both natural and made ground, its response to loads applied to it, and ways of improving that response, are helping to increase the ability of the built environment to withstand the loads generated by specific hazard events. However, in terms of the need for the built environment to function satisfactorily, these abilities are only as good as the knowledge of what events will affect a specific piece of ground during the period of interest – that is, during the design life, which is more or less the next hundred or so years in most cases. Herein I address the inevitable imprecision inherent in estimates of the magnitudes and numbers of events a structure might be exposed to in its design life, and the limitations this imprecision places on the realism of probabilistically-based design loads. This leads to consideration of a different way to decide how an engineering structure might be designed and built to resist geotechnically-derived loads resulting from hazard events.

The term „resilience“ is much used nowadays to refer to the ability of something (in our context a unit or system of infrastructure) to survive a major shock (in our context a geotechnical failure of some type, such as foundation collapse or landslide) and to maintain/regain its ability to function effectively during the event and afterwards. Engineering design can obviously enhance the ability of a structure to withstand the loads that a specific anticipated geotechnical event will apply if and when it occurs, but choosing the design event is the tricky bit in this contribution to resilience. Again, if we knew what was going to happen in the design life of a structure we'd be able to design for that – but we don't.

2 RISK MANAGEMENT

The usual solution to this conundrum is to design a structure so that it will fail in an event of such intensity that it has an acceptably low probability of occurring during the design life; or, increasingly, so that the risk associated with the failure (where risk is probability multiplied by cost of failure) is acceptable. This process is called Risk Management, and it has developed well-defined procedures that result in risk being reduced to such a level that either it is deemed acceptable (e.g. Finlay & Fell, 1997), or that the utility of investment in risk management is optimised in some way. This process obviously requires that the relation between event magnitude and frequency is quantified to an adequate degree.

While for some events (e.g. river floods) substantial quantities of this type of data are available, this is less often the case for geotechnically-related events such as liquefaction, landslides or earthquakes, so the statistical basis for probabilistic design has intrinsic errors. A less well appreciated issue, however, is that even with infinitely accurate statistical data, design based on probability is unavoidably imprecise because – *by definition* – the number of near-design-level events that occur during a design life is small. This means that the infinitely accurate statistics are used to predict the magnitudes of a small number of events, and the magnitudes of events that actually occur will in that case be likely to differ significantly from these predictions. This effect is illustrated in Figure 1, which shows the maximum error in ten averages of various-sized random samples of all the integers from 0 to 99 (whose average is exactly 50). Obviously the smaller the sample, the greater is the deviation of the sample average from 50.0; in particular, when sample size is very small (1 or 2) the error (defined in this way) approaches $\pm 100\%$. Applying this to a structure with a design life of 100 years, the

average magnitudes of the ensemble of 1-year events can be predicted to $\pm 10\%$, but that of 50-year events has an error of $\pm 95\%$. Thus any design based on predicting the occurrence of small numbers of events is liable to extreme errors, even with infinitely accurate probabilities of future event occurrence.

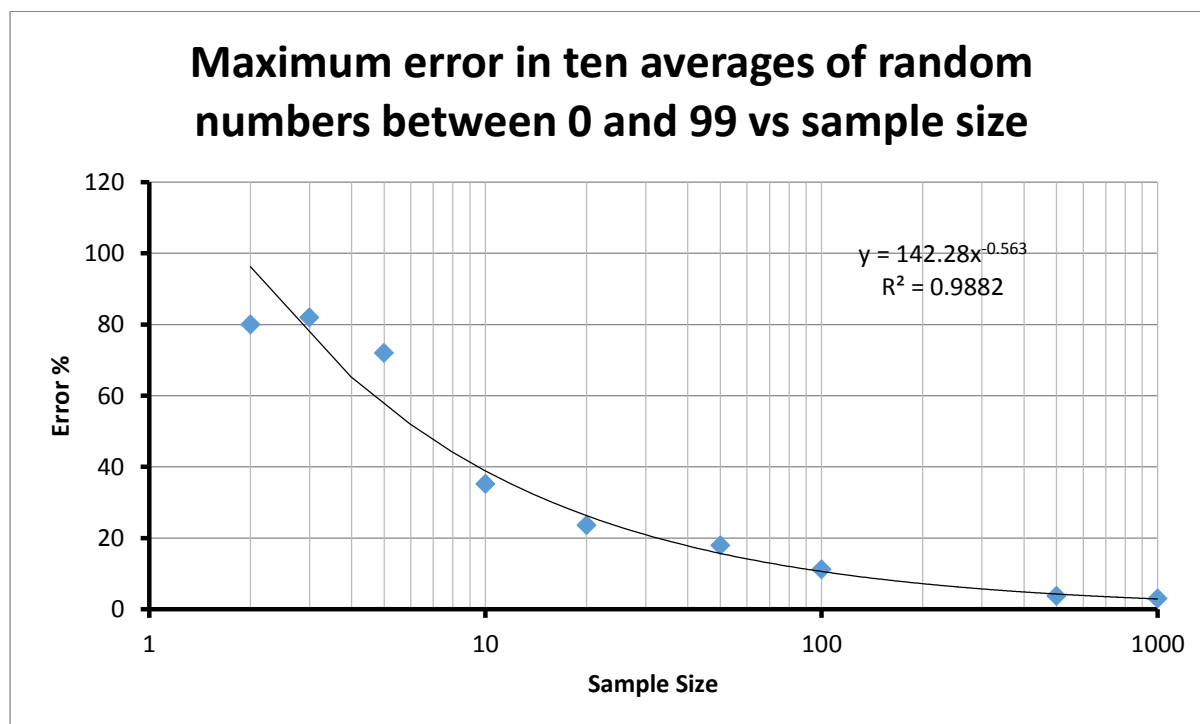


Figure 1: Error in predicting averages of various-sized samples of numbers from 0 to 99.

This situation becomes worse when some utility optimisation procedure is involved (e.g. cost-benefit analysis of event mitigation measures), because in addition to quantified event magnitudes, and thus costs and benefits, being imprecise for the reasons set out in the previous paragraph, the net benefit which is optimised is the difference between two large and imprecise quantities – cost and benefit. Thus the net benefit, which is a much smaller number, has a very much higher imprecision still and its optimisation seems likely to be unrealistic as a design procedure.

A final drawback to the use of event probabilities in design is that very large events, which have very low probabilities, are treated as if they will never occur. In fact, there is absolutely no reason why the next event that occurs cannot be a very low-frequency event. For example, the September 4 2010 Greendale earthquake in New Zealand had a probability of occurrence in 2010 of about 10^{-4} (Stirling et al., 2007), and no rational person would have made a statistically-based decision to prepare for it in that year (or, in fact, in any given year) with a higher priority than preparation for a whole suite of much more likely, if less intense, events. Nevertheless, the 1-in-10,000-year event occurred in 2010. In reality the low probability of the event caused it to be unexpected when it occurred, and this will always be the case when intense events are associated with probabilities. Events become unexpected when they are assigned low probabilities.

Probabilistic design is often justified on the grounds that it defines the most likely event, and that it is rational to design for this rather than for any other event. What is not stated, however, is that in most cases the likelihood of this event is small in absolute terms, and that it is much *more* likely that some other event – that has not been designed for - will occur. For example, if the probability of the most likely event is 10%, then there is only one chance in ten that it will occur, and it is ten times more likely that something else will happen. There is a 90% probability that the design event will not occur.

Thus, reliably designing to avoid ground failure on the basis of probabilistic analysis is very difficult, and there is a strong case for thinking in addition about non-probabilistically-based ways of developing resilient infrastructure.

3 DESIGN FOR RESILIENCE

3.1 SAFE-FAIL DESIGN

One example of a way of preventing ground failure turning into disaster is to *design the failure*, so that the failure mode is known and the consequences foreseeable, even if the time of occurrence and the magnitude of the event that causes failure scale are unpredictable – measures can then be put in place to reduce the impact of the consequences. This is “safe- fail” instead of “fail-safe” design (Park et al., 2013). It renders the failure itself expectable, even if the trigger event is not. This is equivalent having a “fuseplug” in a levee – a length of lower bank that will overtop and wash out before any other part of the bank does, so that the location of failure can be foreseen and the consequences mitigated. It also defines very closely the magnitude of the smallest event that will cause failure; but most important of all, its very existence means that the failure is *expected*, and people will not be taken by surprise when it occurs.

3.2 COMPLEXITY AND RESILIENCE

Most major disasters, by contrast, result from events that are unexpected and take people by surprise, and obviously these cannot be designed for. One of the reasons for unexpectedness has been mentioned above – the event is expectable but is assigned a low probability. Other reasons are that no event of a particular type has ever been experienced or imagined (officially, at least) in the location at which it occurs (an example is the Mt St Helens debris avalanche in 1980: Glicken, 1994); and that while the event type was expected, the event magnitude was not (examples are the 2004 M9.1 Andaman, 2010 M7.0 Haiti, 2010 M8.8 Maule and 2011 M9.0 Tohoku earthquakes, all of which exceeded their maximum credible magnitudes). However, one doesn’t need to know what sort or size of disaster will befall a ship before deciding that lifeboats are a good idea, because the main vulnerability of a ship is that it can sink. It is more useful to think of things that can be done – equivalent to installing lifeboats on a ship – that, whatever the nature of a catastrophe, will reduce society’s vulnerability to the event.

Alternatively, coming from the opposite direction, what characteristics of the way society is organised today lead to vulnerability of infrastructure to unexpected shocks? Some examples include:

- spatial concentration and centralisation, so that all one’s infrastructural eggs are in one basket in a single location;
- short-term economic optimisation and planning, which ignore the possibility of expected but rare events;
- lack of redundancy, so that failure of one component is disastrous;
- component uniformity, so that a failure of one component can trigger cascading failure of equally-strong adjacent components. Component diversity, by contrast, is like biodiversity, it reduces the “brittleness” that leads to sudden and widespread failure;
- top-down governance and communication in order to be better organised, with the result that information about developing crises cannot easily move upwards to decision-makers;
- the operation of systems to eliminate “troublesome” small-scale oscillations in fluxes; this can lead to catastrophic instability on a much larger scale. Steadiness, like uniformity, is potentially destabilising; and
- increasing focus on data-informed decision-making, as a consequence of increased accountability and legal paranoia; as emphasised above, the information carried by data is often spurious, and in addition this diverts attention from events on which data are sparse or absent – usually big events.

The opposite characteristics, which by implication impart resilience to systems, are obvious: decentralisation, long-term planning, redundancy, heterogeneity, unsteadiness, isotropic communication and qualitative decision-making. Interestingly, most of these are also characteristics of the structure or behaviour of *complex dynamic systems*; these are a type of system that is increasingly demonstrated to represent the essential functions and behaviours of many-component, strongly-linked systems such as are found to have evolved in ecology, physiology, finance, geomorphology, sociology and many other fields (e.g. Sornette, 2002). Note that these are the very systems that interact to generate disasters – natural systems and societal systems. Thus there is reason to hypothesise that *complex dynamic system processes are fundamental to disaster processes*. If this is indeed the case, it has some serious implications for how disasters can be managed or reduced, because one of the major characteristics of complex dynamic systems

is the *emergence* of behaviour patterns at a given scale that cannot be predicted or expected from knowledge or understanding of smaller-scale behaviour; in other words, complex system behaviour is fundamentally and irreducibly unpredictable. Thus, future disasters are fundamentally unexpected. This corresponds closely to the common experience that major disasters at a given location are unexpected, either in type or in magnitude, and clearly you can't design to manage the risk of an unknown disaster. Nevertheless, as suggested above, designs that purposefully incorporate characteristics associated with resilience are likely to prove to be more resilient.

But is there any *evidence* that these characteristics do in fact confer resilience? If we think about evolution as the basis of natural system development, the answer is yes. The natural systems of the planet (in particular, species of life-forms, but it is becoming apparent that non-living systems such as rivers behave similarly) - are accepted to have evolved over billions of years as a result of competition with other systems in their environment, and in a context of scarcity of resources. When severely stressed by their environmental systems, species lose those members that are least equipped to survive those stresses, and retain as breeding stock those members better equipped - thus over time a species evolves to become more resilient to the shocks it experiences. Since life began it is accepted that species complexity has increased (albeit punctuated and altered by occasional megashocks, such as mass extinctions), and systems of species have evolved the characteristics and behaviours of complex dynamic systems. It appears that as the systems of Earth have co-evolved over æons, they have all become complex dynamic systems - indicating that such systems are optimally resilient in some sense. The same should apply to societal systems, which have also evolved in competition with neighbouring systems - so how has society developed the less resilient traits bulleted above?

To answer this question authoritatively needs specialised expertise that I lack, but one suggestion is that since humans developed a high level of intelligence and ability to utilise natural resources, human society has increasingly insulated itself from natural systems and is thus decreasingly affected by them - or at least, by the ones it can manage, which are the smaller ones. In this way other priorities than resilience have been able to develop and survive - such as the priority to make as much money as quickly as possible; to exercise maximum power over fellow humans; to minimise personal feelings of insecurity - and these underlie many of the traits listed above. While they may confer resilience to the behaviours of the evolving societal environment, they do not confer resilience to natural system behaviours.

4 SCENARIO-BASED DESIGN

These general considerations are perhaps interesting, but they do not provide the engineer with an alternative to magnitude and probability in designing infrastructure to be resilient to future events. They leave us instead with the uncomfortable feeling that the future is unknowable, and therefore cannot be designed for. While this is true, the situation is more hopeful in reality.

While the future events that will trigger disasters will be of unknown type, unknown magnitude, unknown time of occurrence and unknown location, the societal effects that can result are less unknown. Irrespective of whether the triggering *event* is a tsunami, earthquake, eruption or tornado, the resulting *effects* on society are relatively common: destruction of assets, death and displacement of people, interruption of commerce. How a disaster affects society depends more on the characteristics of society than it does on the characteristics of the trigger event. For example, an identical earthquake to that which hit Tokyo in 1923 will have very different effects when it next occurs in the same location. Again, we cannot know where major landslides will be caused by earthquakes in mountain ranges, but we do know they will block rivers, and since we know where the towns are located that are vulnerable to landslide dambreak floods on rivers, we can figure out what the effects can be. Thus the unknown nature, size, location and time of a trigger *event* is not crucial to usefully characterising its potential *effects*.

This indicates that useful information about what can happen to society in future disasters can be derived from knowledge of past disasters, combined with physical and social science, by developing disaster *scenarios* as a basis for future planning and design. Science, community knowledge and historical/paleo-data provide information on what types of event are possible at a given location, and how big they can be; selecting one of these, the scenario of effects on society can be developed relatively easily, including the longer-term effects of interruption to commerce. By generating a limited suite of scenarios, the likely range

of societal effects can be estimated; this knowledge can then form the basis of community planning to reduce vulnerability to the scenario effects.

Again, this is all very well when expressed in general terms. How can the idea be applied to specific geotechnical design considerations? Let's think about landslide risk management...

5 GEOTECHNICAL RESILIENCE; EXAMPLE

Consider a typical landslide hazard situation. An existing development at the base of a hillslope is affected by a small landslide during heavy rain some years after the development was completed. No other information is available on the history of landsliding in this area, but some geomorphic indications of past slope instability are found during a post-landslide investigation. The conventional response to this situation would be to carry out limited geotechnical investigation and analysis of the slope, develop a numerical model of its stability under heavy rainfall (calibrated, note, by one data point representing the one known landslide), and establish what rainfall intensity and precursory moisture content are required to cause a landslide sufficiently large to be considered unacceptable. If these conditions occur, according to meteorological data, at an unacceptable frequency or probability, then action is required; if the frequency is acceptable, then nothing needs to be done.

Note, first, the sources of imprecision in this procedure. How well do the geotechnical data represent the full suite of ground conditions across the whole slope? What assumptions underlie the numerical model? Is one data point adequate to calibrate the model? What is the length of rainfall data record, and the corresponding statistical precision? What is the design life of a building? These call into question the robustness of a "do nothing" decision, or in other words the extent to which the people living at the base of the slope are in fact safe. The model may indicate a 10^{-4} per annum chance of damage, but, even if it is accurate this does not mean that no damage can occur in any given year, or in the 100-year expected life of the development. The model may be wrong, and the true probability may be 10^{-2} ; and even though the site is then 100 times more dangerous than calculated, it may well be the case that no further landslide occurs within the 100-year design life. Where is the reality in all this? In addition, who decides what probability of damage or death is acceptable, and acceptable to whom?

If, on the other hand, the frequency appears unacceptable, then action is required either to engineer the slope to reduce the risk; to strengthen or protect the buildings so that a landslide does little or no damage; or to abandon the site, immediately (if feasible) or gradually. Any of these is in reality a useful contribution to risk reduction, but whether the residual risk attending any strategy except abandonment is acceptable is, as noted in the previous paragraph, not simple to determine. Thus, basing a design to counter a geotechnical hazard on magnitude-frequency data is not simple, and is in fact fundamentally unsatisfactory (Wong, 2014). Although based on accepted procedures, much of the rationale is spurious.

A resilience-based approach to the situation, by contrast, starts from the simple and irrefutable observation that an unexpected hazard event has occurred. It can be reliably inferred from this that landslides will occur again on this slope, and thus present further hazards to the development - the hazard is now expected. When, where and how big these landslides will be is unknown. Some scenario landslides can be generated, based on the known occurrence and the geomorphic indications, together with perhaps an idea of the largest conceivable event so that scenarios can be aligned with respect to this; the effects of these on the development can be estimated, and the consequent local and regional impacts on building values, community sociology, commercial activities and regional development can be estimated (note that developing scenarios for the latter impacts will require considerable collaborative work with the community). With these impacts in mind, the community and its officials together can then work out how they wish to respond to the hazard threat. This may, as before, involve slope engineering, and/or building protection and/or partial or wholesale site abandonment; the significant point is that the scenarios provide much more useable information on landslide effects than did the numerical/probabilistic data. No probabilities are attached to the scenarios, because it cannot be known which of them will occur first - all we know is that they all can, and in fact will, occur on some unknown days in the future. In addition, because the scenarios are understandable by everyone involved - from lay people to government department heads and politicians - the community has available a much sounder and broader basis for making realistic and acceptable decisions.

Certainly, in this latter case, there is still a need for detailed geotechnical input for slope engineering, and for building protection and strengthening, but again the point is that this design is now based on a realistic picture of what will happen on the site one day, rather than a probabilistic view of what will happen on average every few thousand years. In essence, the possibility of a poorly-constrained probabilistic decision resulting in unacceptable damage and deaths is avoided. If damage and deaths do occur – and this cannot be ruled out, except by abandonment – then the situation is ameliorated by the fact that the community has taken a genuine part in the decision that led to this, and will therefore accept responsibility along with their professional advisers; a witch-hunt is much less likely because the decision process has been inclusive, open and transparent. This should be comforting to engineers designing structures that might kill people one day.

6 DISCUSSION

6.1 ACCOUNTABILITY

One significant criterion for assessing the feasibility of a design methodology is to think about what will happen at a court of enquiry following a disastrous failure. An enquiry tends to put a lot of weight on demonstrating that approved conventional procedures have been properly followed. Thus suggesting that a conventional procedure is irrational, as herein, tends to generate understandable scepticism and reluctance to adopt any different procedure, no matter how rational. This is a serious disincentive to seriously consider arguments such as those made herein. In this context, it is possibly useful to think of design for resilience as a procedure that can complement probabilistic design, rather than replacing it. In reality the imprecisions that accumulate in probabilistic design allow plenty of scope to implement resilience concepts, and involvement of a community in designing for its future is (or should be) in any case an ethical requirement of most professional disciplines. Thus pragmatic considerations suggest incorporating resilience considerations into conventional design rather than abandoning the latter in the short term, irrespective of its shortcomings.

6.2 GEOMORPHOLOGY

In the hypothetical example above, it is clear in retrospect that the hazard could perhaps have been anticipated had a competent geomorphologist been consulted with respect to the initial hazard assessment of the development, and that might have led to more thorough investigation of the landslide hazard at the outset. Regrettably, geomorphologists are still under-utilised in everyday geotechnical engineering, and indeed generally in engineering site investigations. The latter usually focus on ground conditions from a foundation perspective, perhaps gleaning a deeper geological perspective from engineering geology if the situation appears to warrant it (a decision usually left to the often inadequate geological knowledge of the engineering consultant). However, where hazards may be subtly-indicated, if at all, a geomorphologist can often bring valuable insights on future landform events such as river aggradation and slope instability from an understanding of the geomorphological setting and history of the site.

7 CONCLUSIONS

Examination of the assumptions underlying probabilistic risk management leads to the conclusion that, where the target sample size is small (as in hazard and design event contexts), the technique is of little value in selecting a rational design event to reliably reduce the occurrence of future damaging events in a specific location.

Scenario-based design, by contrast, can avoid probabilities entirely and deal only with event- and effects-scenarios, by whose selection and use communities and experts can work together to decide on the nature and level of protection required against future hazard events and their consequences.

For reasons of professional accountability, however, it is probably more realistic to suggest that scenario-based design is implemented initially to complement standard probabilistic design procedures.

8 REFERENCES

- P.J. Finlay and R. Fell, 1997. Landslides: Risk perception and acceptance *Canadian Geotechnical Journal* 34: 169-188
- J. Park, T. P. Seager, P. S. C. Rao, M. Convertino and I. Linkov 2013. Integrating Risk and Resilience Approaches to Catastrophe Management in Engineering Systems. *Risk Analysis*, Vol. 33, No. 3, 2013 DOI: 10.1111/j.1539-6924.2012.01885.x.
- Sornette D. 2002. Predictability of catastrophic events: material rupture, earthquakes, turbulence, financial crashes and human birth. *Proceedings of the National Academy of Science* 9: SUPP1 2522-2529.
- Stirling M., N. Litchfield, W. Smith, P. Barnes, M. Gerstenberger, G. McVerry, J. Pettinga 2007. Updated Probabilistic Seismic Hazard Assessment for the Canterbury Region, *GNS Science Consultancy Report 2007/232 - ECan Report U06/6 August 2007*
- Wong, H.N. 2014. Is landslide risk quantifiable and manageable? In: Zhang et al. (ed) *Geotechnical Safety and Risk IV* Taylor & Francis Group, London, ISBN 978-1-138-00163-3. 101-111.