

Young Writers Competition

GUESTimation

Breaking the deadlock on wedding guest lists

Earlier this year, together with the Young Statisticians Group of the RSS, *Significance* ran its second Young Writers Competition. A selection of runners-up will be published on these pages and on our website, www.significancemagazine.com, over the coming months. Here we present the winner. **Damjan Vukcevic** talks us through independent variables, probability distributions, and robustness of models, all in the context of that nightmare that most of us encounter at some time in our lives – whom should we invite to our wedding?

It is truth universally acknowledged that one of the most stressful parts of organising a wedding is deciding the guest list. My partner and I faced a classic dilemma. There were many people we wanted to invite, but we also had to stay within the constraints of our venue and budget. Should we invite all of our loved ones and risk breaking the bank? Play it safe and end up regretting not filling out the venue? Or resort to that tempting but socially risky manoeuvre of having multiple invitation rounds? (B-list friend to A-list friend: “You got the invitation in March? Mine just arrived.”)

If all our guests were in town, the guest numbers would be easy. You can count on people showing up to weddings. But our family and friends were not so conveniently located. Many of them were in Europe and Asia. How many would make the trip Down Under?

Fortunately, this was one type of uncertainty that my education equipped us to deal with. We created a probability model to work out how many people we could invite in a single round of invitations, while taking on a comfortable level of overbooking risk. Statistics allowed us to sleep easy on this particular wedding issue.

Predicting our guest attendance

Our venue could host 110 people comfortably, and 120 at a squeeze. For guest comfort, we wanted to aim for 100–110 attendees. To emphasise that our counts include the bride and

groom (since we are interested in the total number of people), we refer to *attendees* rather than guests.

Our first draft guest list far exceeded this, which would have been a concern except that of course there would be people who would decline. In theory, this natural attrition would allow us to “overbook” the wedding. In fact, overbooking is likely to be necessary in order to hit our target range. (Airline companies routinely overbook

**Too many guests on the list?
Too few? Statistics can make it
just right**

their flights, to avoid flying with part-empty planes.)

The big question then is how much could we overbook? What is a safe level before it becomes high risk? Anyone who has been bounced off a pre-booked aircraft flight will know how annoying it can be – and we did not want start our married life by upsetting half our friends.

A simple approach would be to guess an average response rate and multiply it by the number of guests to estimate the numbers. Surely this

approach is too crude. We know something about the circumstances of each guest and should be able to use this information to estimate – or “GUESTimate” – their attendance.

The most important factor that differentiated our friends and family was where they live. Some are in our city, Melbourne, while a significant number live overseas. Of those overseas, we knew some were avid travellers, some less able to travel, and others had already been talking up the idea of visiting us. Therefore, we split our guests into the following groups:

- **Definitely** – locals and overseas Melbournians who generally return for the Christmas break (which is when we held our wedding);
- **Likely** – overseas guests who seemed enthusiastic about a holiday in Melbourne;
- **Maybe** – overseas guests who love travelling and do it readily;
- **Unlikely** – overseas guests who generally do not travel far from home.

For each group we assigned a probability of attendance based on our subjective judgement: 100%, 80%, 50% and 10%, respectively. Choosing a probability of 100% for the first group is clearly an overestimate, but without a better basis to judge we decided to be conservative. We could now calculate an expected number of guests by multiplying the number of guests in each group by their probability of attendance, then summing these.

This approach to guest number estimation is probably the most sophisticated that people with access to Microsoft Excel get to. However, what are the chances of actually getting close to that number? Are the final numbers likely to be higher or lower than the point estimate? How confident are we that we will end up with fewer than 110 guests?

What we need is a probability distribution. As well as the four groups defined above, our guests could be naturally grouped on a smaller scale. Most attendees came in pairs (couples), and some formed larger family units. Attendance of people within units is unlikely to be independent. We wanted our model to capture this.

We assumed that guests within a unit either attend or do not attend all together. To illustrate, if we invited a couple and their two children, we assumed either all four of them would attend, or none of them.

Across units, we assumed independence. For example, if we invite two friends and their respective partners, we assume that if one couple attends, it has no bearing on whether the other couple does.

These assumptions together define a probability distribution for the total attendance. We can then examine this distribution to help us decide on the list of people to invite. In particular, we can vary the invitation list until we obtain a distribution whose properties we are happy with. (The calculation of the distribution can be done fairly easily, either by iteratively adding in each guest unit and updating the distribution, or by running a simulation based on the given probabilities; but we shall spare you the mathematics.)

We began with a “wish list” of all people we would invite if we had no constraints. This consisted of 185 attendees. For this list, our model



Stockbyte/Thinkstock

Table 1. Summary counts for our final invitation list

Group	Units					Total per group
	1	2	3	4	5	
Definitely	29	25		4	1	100
Likely		1				2
Maybe	4					4
Unlikely	6	10	1	1		33
Total						139




Figure 1. Our estimated probability distribution. The bars shaded in green form a 95% prediction interval

gave a 95% prediction interval of 137–152, well beyond our target range of 100–110. Perhaps we could tolerate a small probability of discomfort (more than 110 attendees), but we wanted the probability of outgrowing our venue (more than 120 attendees) to be close to zero.

We gradually reduced the number of guests until we met our criteria. Our final list had 139 attendees. A summary of it is shown in Table 1, as counts split by group and unit size. Applying our model gave a 95% prediction interval of 102–113.

Figure 1 shows a plot of the distribution, with the 95% interval highlighted in green. It shows that we were most likely to get 106 attendees, but it wouldn’t be at all surprising to get values a bit higher or lower than this. The probability of discomfort was 11%, and that for outgrowing the venue was almost 0. With our questions now thoroughly answered, we sent out our invitations and slept easy.

Assumptions, assumptions

At this point, some of you might have reservations about our modelling efforts. After all, haven’t we just invented the probabilities for each group?

This is a genuine concern. Nevertheless, the fact that we even have a model is substantial progress. Most people faced with the same dilemma must rely on intuition alone. Without a handle on the uncertainty, they might choose to play it safe and invite fewer guests, possibly leading to disappointment and hurt feelings. Others might pay for a larger venue or forgo their dream outdoor marquee wedding because they are worried about overshooting the mark. Still others might accede to family pressures to invite more relatives, quietly hoping that it will be all right in the end.

What we have done is used probability to help us make a calculated “stab in the dark”. Our

model is a tool to convert intuition and experience into a form that can be easily interpreted and acted on.

Any statistical modelling effort invariably requires assumptions. The strength of the conclusions will depend on how reasonable these are (both before and after seeing any data), and how robust the conclusions are to varying these assumptions.

In our case, we did not have any external data, and thus any assessment of our assumptions will be subjective. Using the model as a personal decision tool, we judged that the assumptions were reasonable since they accorded with our personal expectations. The model formalised the information that we would use to make the decision in any case. Furthermore, it has the additional benefit of making our assumptions transparent.

If we wished to go beyond our personal context, then some more scrutiny of the assumptions is warranted. In fact we have repeated the exercise on behalf of our friends (see below).

What about the effect of varying the assumed attendance probabilities? If we reduced the probabilities all by 10 percentage points, the prediction interval goes down to 83–101. If we increase them all by 10 percentage points (up to a maximum of 100%), then the interval shifts up to 105–118. These are both substantial shifts, which we could accommodate in our venue and budget.

The other important assumptions are that people act in small units, and that the units act independently of each other. It turns out that the first of these assumptions has a minor impact, while we believe the latter to be close to the reality of wedding behaviour (see box).

Independence assumptions

Assuming independence is the main reason why our model gives prediction intervals that are not impractically wide.

You can see this intuitively by imagining that you invite an extremely large family, say of size 20, and that either they can all come or none of them. This gives them a strong influence on the final number: it will vary by ± 20 based on one decision alone.

As we assume more independence, we reduce the uncertainty implied by our model. Thus, to assume that each guest acts independently is an extreme case. The only way we deviated from this was to assume that people act in small units (couples and families). We think this is a realistic set of assumptions for most weddings. It implies that the different units do not greatly influence each other.

How might such influence arise? One way would be if a guest is more likely to attend if they know their friends will be there as well. Alternatively, they might be less likely to attend if they know that, for example, their ex-partner will be there. For our wedding, we believed this effect would be negligible.

In practice, the impact of not assuming complete independence is minor because unit sizes are generally small relative to the total number of guests. To give you a sense of this, we can rerun the calculations for our wedding assuming complete independence. The resulting 95% prediction interval is 103–111. This is only slightly narrower than the original interval. We could probably have gotten by with the simpler independence model if we had wanted to.



Damjan Vukcevic and Joan Ko at their perfectly-filled wedding. Photo: Teddy Tan (<http://teddytan.com.au>)

Other weddings

Our friends were enthusiastic about our probabilistic model for wedding attendance. Soon we had requests to help break the guest list deadlock on other weddings. We noted earlier that our model was designed for our own expectations, so it may not be completely transferable to our friends. However, our friends had mostly settled on a guest list and simply wanted an estimate of the final numbers. They happily adopted our assumptions.

Here we present results for three other weddings. We will refer to them as weddings B, C and D, and to our wedding as A. Our friends gave

us counts split into groups and units, just as we had for our wedding. We used this information to calculate 95% prediction intervals using the same model. Table 2 shows the summary data for all weddings, and Table 3 the predictions.

How well did we do?

Of the people on our original list, 97 celebrated with us on that happy day. This is fewer than we predicted. It turned out that very few of our overseas guests could come to the wedding.

This was fortunate in the end because we had to add unexpected visitors and partners in the run-up to the event. The actual final total was 105 attendees, an ideal number.

How accurate was our GUESTimation for our friends' weddings? The outcomes are shown in Table 3. For each wedding, our model tended to overestimate. The overestimation was mild for weddings A–C. For wedding D the discrepancy was substantial.

Now that we have actual data, we can evaluate our assumptions more directly by examining the outcomes by each of the four guest groups. Figure 2 shows 95% confidence intervals for the attendance probability for each group–wedding combination. It shows that the main deviation is for the *definitely* group, which covers the majority of invited guests. Our conservative estimate of 100% attendance was always going to be too high. Over these four weddings, attendance by the *definitely* group varied from as low as 84% (for wedding D) up to 99% (wedding B). This suggests that using 100% attendance is playing it safe, with lower values often being more realistic.

For the *likely* and *definitely* groups, our assumed attendance probability was also high, and for *unlikely* it seems about right.

A natural next step would be to use these data to re-estimate the group probabilities, either as single values or summarised as a prior distribution to allow for variation between weddings. This would give a more accurate model to use for future weddings predictions.

Lessons learnt

Estimating wedding attendees in the absence of real data put me personally in an uneasy situation. As a scientist, every day I seek the comfort of objectivity, data, facts and conclusion. In contrast, as first-time wedding planners, we could draw only on gut feel and the stories we tell each other.

Tips for modelling your own guest list

Step into your guests’ shoes. How strongly do you wish to attend this wedding? What personal circumstances might prevent you from coming?

Get data if you can. Use ours as a starting point.

If you have a fixed guest limit, be conservative with your probability assumptions.

Try varying your assumptions.

Leave some buffer for unexpected late invitations and “wedding-crashers”.

We are convinced, though, that this project has real value. The process of modelling led us to more clearly think about our assumptions and enabled us to evaluate implications of our decisions. We also have friends clamouring for GUESTimation services for their own weddings. The weddings industry in Australia is worth a few billion dollars annually. Anything that puts statistics at the heart of a billion-dollar industry is surely worthwhile.

Most importantly, we have a greater appreciation that sometimes decisions must be made without data. Non-scientists do it all the time. This project shows that even without data, statisticians can help improve decision-making by taking a systematic and probabilistic approach to assumptions.

Not long after our wedding, I read David Palmer’s *Significance* article about risk management for an engineering project¹. His challenge was to estimate how long it would take, and how much it would cost, to build and transport a special drilling machine. I was struck by how similar in spirit his approach to the challenge was to our wedding invitation problem. We both broke the problem down into small parts, modelled these with simple distributions based on experience, and summed it all up in the right way.

It shows that even without the rigour of a randomised controlled trial, we can go a long way. That is good news for those of us making decisions every day in the face of uncertainty.

Table 2. Summary data for our friends’ weddings

Group	Invited				Attended			
	A	B	C	D	A	B	C	D
Definitely	100	201	67	183	96	199	61	154
Likely	2	3	5	22		3	4	10
Maybe	4	2	12	13	1		3	1
Unlikely	33		16	11			3	2

Table 3. Predictions and outcomes for our friends’ weddings (PI: prediction interval)

	A	B	C	D
Invited	139	206	100	229
95% PI	102–113	202–206	72–85	200–216
Attended	97	202	71	167

References

1. Palmer, D. (2011) Will the big machine arrive? Engineering the Miami Tunnel. *Significance*, 8(4), 148–153.

Damjan Vukcevic is a postdoctoral researcher at the Murdoch Childrens Research Institute in Melbourne, Australia. His work focuses on statistical genetics and biostatistics.




Figure 2. Estimates and 95% confidence intervals for the attendance probability for each group for each wedding. The green horizontal dashed lines show our assumed values