

THE KEG INDEX AND A MATHEMATICAL THEORY OF DRUNKENNESS

CHRISTOPHER TUFFLEY

Consider the following problem:

Problem O. There are n people sitting in a ring, one of whom takes a swig from a keg of beer, and then passes it right or left with a 50% probability independently of what has happened before. The process repeats until everyone has had at least one swig, then stops. Show that the probability that the keg stops at a particular (non-starting) person is independent of that person's position.

I'm not going to discuss this problem until the end, because to my mind there's a much bigger issue here, and that is: just how drunk are these people??

After some thought, we might suppose that completely sober people would realise that the most efficient (not to mention fairest!) method of having everyone get a swig from the keg is to have the keg travel round the circle in either a clockwise or anti-clockwise direction, and that they would pass the keg accordingly: on receiving the keg from their left they would pass it to their right, and vice-versa. Completely drunk people, on the other hand, would be capable of little more than shoving the keg back where it came from; while someone somewhere between these two states might pass it back where it came from with some probability p (increasing with drunkenness) or pass it on with probability $1 - p$. This leads us to define the keg index:

*Someone is **p -drunk** (or a **p -drunk**) if he or she passes the keg back with probability p , and passes it on with probability $1 - p$.*

Recapping, “sober” corresponds to $p = 0$, “drunk” corresponds to $p = 1$. So according to the keg index, the people in Problem O are half drunk! Well! That answers the question we set out to answer, but now that we have a method of measuring drunkenness (and one so readily estimated by anyone, anywhere, without any need for fancy equipment—simply use the observed frequencies with which the subject passes the keg on, and back) we may have a new-found predictive power. Maybe now we can answer some of those questions you've always wanted to ask but never knew how, such as: just how drunk is the combination of a p -drunk person and a q -drunk person?

Suppose a p -drunk person P is standing with a q -drunk person Q on the right and is passed the keg from the left (see Figure 1). P might pass it straight back with probability p , or pass it to Q with probability $1 - p$, who might pass it on with probability $1 - q$, or back to P with probability q . P in turn may pass it out

Mathematical Entertainments column, Mathematical Intelligencer 24(3):34–36,67, 2002.

Figure 1 is an `xfig` reproduction of the original hand-drawn figure appearing in the published article. I would like to thank Chris Godsil of the University of Waterloo for directing me to the Lovász and Winkler paper.

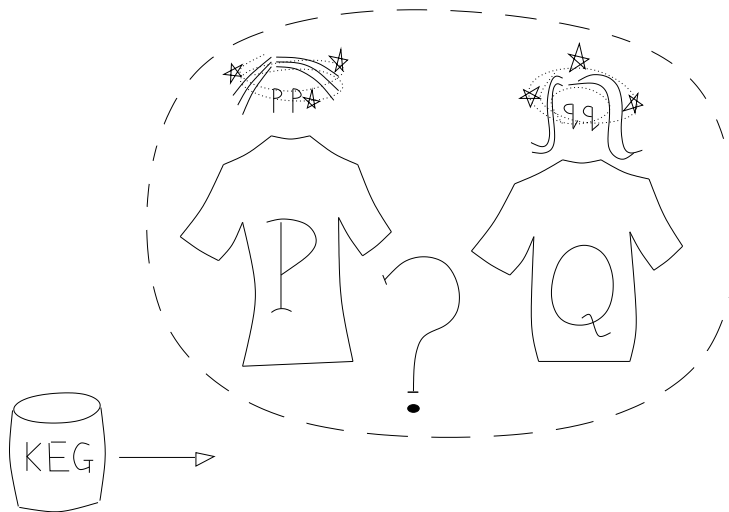


FIGURE 1. Just how drunk is a p -drunk person with a q -drunk person?

(with probability $1 - p$) or back (probability p). In this way it might shuttle back and forth between P and Q for some time before eventually emerging from one side or the other. Each round trip from Q to P and back again (or from P to Q and back, after the first) occurs with probability pq , so we get a geometric series. The probability that P passes it out is

$$\begin{aligned} p + (1 - p)^2 q \sum_{k=0}^{\infty} p^k q^k &= p + \frac{(1 - p)^2 q}{1 - pq} \\ &= \frac{p - 2pq + q}{1 - pq}, \end{aligned}$$

while the probability that Q passes it out is

$$(1 - p)(1 - q) \sum_{k=0}^{\infty} p^k q^k = \frac{1 - p + pq - q}{1 - pq}.$$

I make two observations about these two probabilities: first, that they sum to one, and second, that they're symmetric in p and q . That they sum to one is unsurprising, but nevertheless comforting if you want a swig from the keg too, while the fact they're symmetric in p and q is quite unexpected. Unexpected, and fortuitous, saving us from having to introduce such concepts as left- and right-drunkenness: together these two facts imply that the (p, q) -drunk combination behaves as a single person with the net keg index

$$(1) \quad \text{P[the keg is passed back]} = \frac{p - 2pq + q}{1 - pq}.$$

Since

$$\frac{p - 2pq + q}{1 - pq} - q = \frac{p(1 - q)^2}{1 - pq} \geq 0,$$

it follows that the (p, q) -drunk combination is drunker than $\max\{p, q\}$ —at last a mathematical proof of something readily confirmed experimentally.

Before going on let's take another look at equation (1). At first sight there seems to be a problem here when $p = q = 1$, because the denominator is zero. However, a closer inspection shows that everything's okay after all, as the numerator is zero too and the limit as $p, q \rightarrow 1^-$ is 1, which is what we expect. In spite of this, to make things nicer later on, I will assume that complete drunks don't exist. This is mathematics: we're allowed to do this kind of thing. And if sober people start being a problem, I just might assume them out of existence too.

Some people may object to the calculation leading to equation (1) on the grounds that the effect of the swigs on the individual keg indices is not taken into account when calculating the net keg index. This is clearly an issue that will have to be addressed, but as we shall see it is perhaps not as serious a flaw as it may at first appear. The effect of the swigs will depend on the number taken—the fewer swigs, the better the swig-free approximation will be. We therefore calculate the expected number of swigs before P and Q pass out the keg, obtaining

$$(2) \quad E[\text{swigs}] = p + (1-p)(1-q) \sum_{k=1}^{\infty} 2k(pq)^{k-1} + (1-p)^2 q \sum_{k=2}^{\infty} (2k-1)(pq)^{k-2}.$$

Using the series

$$\frac{1}{(1-x)^2} = \sum_{k=0}^{\infty} kx^{k-1},$$

we obtain the expressions

$$\frac{2x}{(1-x^2)^2} = \frac{1}{2} \left(\frac{1}{(1-x)^2} - \frac{1}{(1+x)^2} \right) = \sum_{k=0}^{\infty} 2kx^{2k-1}$$

and

$$\frac{1+x^2}{(1-x^2)^2} = \frac{1}{2} \left(\frac{1}{(1-x)^2} + \frac{1}{(1+x)^2} \right) = \sum_{k=1}^{\infty} (2k-1)x^{2(k-1)}.$$

Applying these to (2) with $x = (pq)^{1/2}$ we find

$$\begin{aligned} E[\text{swigs}] &= p + (1-p)(1-q) \frac{2}{(1-pq)^2} + (1-p)^2 q \frac{3-pq}{(1-pq)^2} \\ &= 2 + \frac{q-p}{1-pq}, \end{aligned}$$

and because $pq \leq p, q < 1$ we have

$$\left| \frac{q-p}{1-pq} \right| < 1$$

on $[0, 1) \times [0, 1)$. Therefore the expected number of swigs lies strictly between 1 and 3, and we may sleep peacefully in the knowledge that in the low swig-alcohol-content limit, the correction to the swig-free approximation isn't too large.

Now suppose we have three people: a p -drunk, a q -drunk and an r -drunk, named P , Q and R . To calculate their net keg index we may simply use the above result to treat P and Q as a single person S , and then apply the result again to S and R . Of course, we could equally well apply the result first to Q and R and then include P , since either calculation gives the probability the keg is passed back; carrying out

the first we get

$$\frac{\frac{p-2pq+q}{1-pq} - 2 \frac{p-2pq+q}{1-pq} r + r}{1 - \frac{p-2pq+q}{1-pq} r} = \frac{p+q+r-2qr-2pr-2pq+3pqr}{1-qr-pr-pq+2pqr}.$$

The result is symmetric in p , q and r , as we now expect in light of the p , q -symmetry of equation (1): interchanging two neighbours doesn't change the net keg index, and such transpositions generate the symmetric group. It follows that we may use equation (1) to assign a well defined net keg index to any group of people, without first having to line them up in a row, and we may thereby extend the keg index to less orderly arrangements. As an application, it's easy to show inductively that n half-drunks are equivalent to a single $n/(n+1)$ -drunk. The rapid convergence of this figure to 1 as $n \rightarrow \infty$ may go a long way towards explaining crowd behaviour in pubs.

In fact, if we define $++ : [0, 1) \times [0, 1) \rightarrow [0, 1)$ (pronounced double-vision-plus) by

$$p ++ q = \frac{p-2pq+q}{1-pq},$$

then $([0, 1), ++)$ forms an abelian, associative semigroup with identity: the keg semigroup. So the keg index forms a structure with a number of nice properties—in fact, about the only nice property we'd like to have but don't is the existence of inverses. What's more, inverses would seem to have a natural interpretation in “sober-up” pills: pills or potions that, when taken, sober you up, or at the very least, make you less drunk. Can the keg semigroup be embedded in a group?

Yes! Indeed, there is a natural candidate for the target group: as an addition law defined on a half open interval and satisfying $p++q \geq \max\{p, q\}$, the keg semigroup is reminiscent of the non-negative real numbers under addition. Furthermore, the fact that n half-drunks are equivalent to a single $n/(n+1)$ -drunk suggests an explicit map: the function $\psi : \mathbb{R}_{\geq 0} \rightarrow [0, 1)$ given by

$$\psi(t) = \frac{t}{t+1}.$$

A simple check shows that ψ is indeed an isomorphism, with inverse

$$\phi(p) = \frac{p}{1-p},$$

so mathematically speaking, sober-up pills exist! To construct them simply extend ψ to all of \mathbb{R} . The fact that $\psi(-1) = \infty$ appears to present a problem, but this is resolved by defining

$$p ++ \infty = \lim_{q \rightarrow \infty} \frac{p-2pq+q}{1-pq} = \frac{2p-1}{p} = \psi(\phi(p) - 1)$$

for $p \neq \infty$ and (taking a limit once more) $\infty ++ \infty = 2 = \psi(-2)$. We then have $\psi(t+u) = \psi(t) ++ \psi(u)$ for all $t, u \in \mathbb{R}$ (in fact, we may regard $p = 1$ as $\psi(\infty)$, in which case this holds for all $t, u \in \mathbb{R} \cup \{\infty\}$) so ψ maps $(\mathbb{R}, +)$ isomorphically onto $(\mathbb{R} \cup \{\infty\} \setminus \{1\}, ++)$. There! Sober-up pills exist, and correspond to things that pass the keg back with probabilities outside $[0, 1]$ —so perhaps “pills” is not the correct word, but in any case, as mathematicians our part is done: we've proven the theoretical existence of sober-up *some things* and we may leave the details of physically realising them to disciplines better suited to the task.

More sober reflections. Now that the keg index has been developed to such a satisfactory conclusion, let's turn our attention back to the original problem. How do you solve Problem O, and what leads to such a surprising result? Our intuition suggests that the keg should be more likely to finish further from rather than closer to the starting point. Where is it going wrong?

To answer the first question, suppose you're sitting somewhere in the circle and did not start with the keg. The keg will stop at you if and only if both your neighbours have swigs before you do, and for this to happen the keg, having visited one of them, must make it all the way around the circle to the other without ever being passed to you. The probabilities of the two events "your left neighbour gets the keg before your right" and "your right neighbour gets the keg before your left" *do* depend on where you're sitting, but their sum does not and equals one. The probability of the keg stopping at you is then the probability of it getting from one of your neighbours to the other without ever being passed to you. But this is just the probability of it stopping at you given that it started at your neighbour, and so does not depend on where you are in the circle.

More generally, consider a random walk on a connected graph G that begins at some vertex x , moves at each step with equal probability to any neighbour of the the current vertex, and stops as soon as it has visited every vertex. Such a walk is called a *cover tour*, and in these terms, the result of Problem O is that a cover tour from any vertex on a cycle is equally likely to end at any other vertex. This is true of complete graphs too, by symmetry. In a paper with acknowledgments "to Ron Graham for extra incentive, and to the Hunan Palace, Atlanta GA, for providing the napkins" Lovász and Winkler [1] show that complete graphs and cycles are the only graphs with this property. In doing so they show that our intuition is correct in general and give insight into where it is failing us for the cycle. Namely, letting $L(x, y)$ be the event that a cover tour beginning at x ends at y , they show the following:

Theorem (Lovász and Winkler [1]). *Let u and v be nonadjacent vertices of a connected graph G . Then there is a neighbour x of u such that $\mathbb{P}[L(x, v)] \leq \mathbb{P}[L(u, v)]$; further, the inequality can be taken to be strict if the subgraph induced by $V(G) \setminus \{u, v\}$ is connected.*

The theorem is proved by showing that $\mathbb{P}[L(u, v)]$ is equal to the average of $\mathbb{P}[L(x, v)]$ at its neighbours, plus a nonnegative correction term that is positive if $G \setminus \{u, v\}$ is connected. This implies that for a fixed finishing vertex y the minimum of $\mathbb{P}[L(x, y)]$ occurs at a neighbour of y , as we expect; but it is not a strict minimum for the cycle, a cycle being disconnected by the removal of any two nonadjacent vertices. However, there is a surprise: they give an example to show that for fixed initial vertex x the minimum of $\mathbb{P}[L(x, y)]$ need not occur at a neighbour of x . The example is a complete graph K_n with an extra path u, x, y, z, v joining two of its vertices u and v (see Figure 2 for the case $n = 6$), for which they claim $\mathbb{P}[L(y, x)] = \mathbb{P}[L(y, z)] \rightarrow 1/3$ as $n \rightarrow \infty$ while $\mathbb{P}[L(y, w)] \rightarrow 0$ for any given w in the complete graph. Heuristically, we may understand the limiting behaviour of $\mathbb{P}[L(y, z)]$ as follows: For n large we expect a random walk to become "trapped" in the complete graph once it enters it, only emerging long after it has visited each vertex; thus as $n \rightarrow \infty$ we expect $\mathbb{P}[L(y, z)]$ to tend to $1/3$, the probability of the walk getting from y to u without going to z .




FIGURE 2. Lovász and Winkler's graph when $n = 6$. For n sufficiently large a cover tour beginning at y is more likely to end at z than at some given w in the complete part of the graph.

To classify the graphs G such that $P[L(x, y)]$ is the same for all distinct x and y they observe that a cover tour cannot end at a vertex v such that $G \setminus \{v\}$ is disconnected. Thus G cannot be disconnected by the removal of any single vertex (that is, it must be 2-connected), but must nevertheless be disconnected by the removal of any two nonadjacent vertices, by the theorem; such a graph is necessarily a complete graph or a cycle.

The extension of the keg index to this more general situation could perhaps lead to further insights. I leave this open avenue of research to others, together with the development of the keg index in the non-swifree case.

REFERENCES

- [1] L. Lovász and P. Winkler, A note on the last new vertex visited by a random walk, *Journal of Graph Theory* 17, No. 5 (1993), 593–596.