

What are the odds of that?
Practical Stats for Astrologers v1.0
by Kyle Pierce

The purpose of this document is to give insights into the meaning of probability in astrology. Our guide is Kyle Pierce, who graduated as a biostatistician. The document will expand as more questions and answers are added, with each major new question being reflected by an increment of 1.0 in the version number.

Q: What are the chances of several planets being in the same house?

A: Even though it sounds simple enough, this kind of problem doesn't have an obvious solution -- or at least, one obvious solution turns out to be wrong. We can't simply take the probability of one point showing up in one house (1/12) and multiply this by the same probability for another point, and so on.

This method will tell us the probability that a given planet (Uranus) is in a given house (3rd) and another given planet (Moon) is in another given house (9th), which would be about $1/12 * 1/12 = 1/144$. And so on, with each specific planet in its house. The probability of finding a given combination of specific planets in specific houses is known as the "joint probability" and is found by multiplying the individual probabilities together.

This is not what we're looking for, however. We want to know the probability that *any* planet (or any number of planets) is in a given house, which is something else entirely.

I found a more promising solution by resorting to my favourite method, which is to simulate the problem by using Monte Carlo techniques. It works like so: you have 12 bins (houses) and 11 balls (10 planets, and the nodal axis has only one independent point -- the other is always opposite it). You have some procedure whereby each ball will end up in one bin or another with equal probability. You run through this procedure and then you count the number of balls that fell into a particular bin. You repeat this same procedure thousands of times, keeping track of the count of balls in that bin. Fortunately, I have software that can do this gruelling work for me.

In this way you can compute a very good estimate of the true probability that, say, 2 or more balls will fall into a given bin. This is a rough analogy to the original problem, and we could of course work with 360 bins (degrees) and measure house cusps precisely, but in this case we are simply roughing it out and making the houses all of uniform size. So it turns out that for any chosen house, and using 11 points:

- over 60% of charts will have at least one point in a given house
- about 23% will have more than one point in that house
- 5 to 6% will have more than two points in that house
- about 1% will have more than three points in that house

When it comes to looking at a house axis, the probabilities are the same as for any pair of houses, all other things being equal. And these probabilities turn out like so for any chosen pair of houses, using 11 points:

- over 85% of charts will have at least one point in that house pair
- about 57% will have more than one point in that house pair
- about 27% will have more than two points in that house pair
- about 10% will have more than three points in that house pair
- at least 2% will have more than four points in that house pair

So for example, a chart which has four points in the 3rd/9th house pair, will occur about 10% of the time -- or 8%, *if* we subtract the 2% that have more than *four* points in a given house pair. Of course, this kind of analysis will leave out any other striking features of that chart, so it may be unfair to say that the chart is not one in a million, but one in ten.

I hope I have managed to make some kind of sense out of this explanation. It may well seem counterintuitive, but the experiment has been run and these are the findings. I find it quite satisfying that there is in fact a fairly accessible solution that doesn't require laying out any huge theoretical edifice -- also, I don't know how else one could do it.

One might wonder whether, if you actually computed the charts, things could turn out quite differently. Well, I have tried this, not on this particular problem but with a number of other problems, and in fact estimates based on a similar kind of rough analogy agree surprisingly well with the results when computing the charts. This analogy would fall on its face, no doubt, with problems that involve counting aspect contacts or zodiacal positions within individual charts, since as we know, planetary positions in the zodiac are not distributed very uniformly. So at least we can anticipate where this approach won't work.

Thanks to Jessica Adams for asking the question that instigated this discussion.