A really dumb old joke about which I hadn't had a single thought in many years rushed back into my head a few days ago. It goes like this:

> A man walks into a bookstore and says to the clerk, "I'm looking for a book whose title I can't remember exactly. It's called 'The Red Boat', or 'The Scarlet Ship',... it's something like that." The clerk answers, "Here's the one you're thinking of, sir—*The Rubaiyat*!"

Now that you're primed (you don't have to wait to stop groaning, that's optional), be sure to pay a visit to the website www.therubaiyat.com. The person or people who put up this site have included several translations of Omar's quatrains, not just the famous ones by Edward Fitzgerald. Pay particular attention to Richard Brodie's version, and let me know if you're as impressed by what he did as I was. He has produced an anagram of each one of Fitzgerald's verses, rearranging the exact same letters into new words without needing a single additional letter or having a single letter left over, and the anagrammed verses also scan and rhyme, and have basically the same meaning as the originals but with completely different wording! Apparently this was not the first time he has undertaken a project of this type. He published a book of anagrams of verses from the King James bible, where again the meanings of the verses were left intact while the letters were rearranged into completely different wordings.

If I had been asked whether this feat were possible before I saw it black and white, I probably would have said that in general it isn't. For certain special cases, yes, why not, but why should it be possible for an arbitrary sequence of words? But of course a moment's clear thought, i.e., once you know what the correct answer is, shows that this intuitive response has it completely backward. Any sequence of N letters can be permuted in $N!/\prod_{i=1}^{26} N_i!$ different ways, where N_i is the number of times that the i^{th} letter of the alphabet occurs in the sequence. The symbol 'Π' in that formula (for the non-mathematicians in the audience, and my apologies to others for the insult to your intelligence) means to compute the product of all 26 numbers $N_i!$.

If you use that formula to compute the number of permutations of the letters in your name, or some other such short sequence, you'll get a number that it tractable, i.e., you could actually write down every single permutation on a piece of paper. In such a case, there clearly is a non-trivial probability that none of them except for the original will be sensical. But if you compute the number of permutations possible for a letter sequence the length of a bible verse or a Fitzgerald quatrain, you will almost certainly end up with a number that literally is beyond astronomical. (That is, larger than the number of stars in the universe. In fact, probably larger than the number of elementary particles in the universe. [Including the photons and neutrinos? Don't know, I hadn't actually considered them. Has any of you guys with physics background seen any estimates of how many of these there are at any given moment?] And so large that if you had started at the moment of the Big Bang and had written down a permutation per second ever since, by the present time you would have gotten through only a tiny fraction of the possibilities.) That being the case, it actually is almost a certainty that at least one besides the

original has all the letters in legitimate words, and scans and rhymes in the desired manner, and has basically the same meaning as the original.

The real problem is how does one find the perfect permutation that is buried within the beyond-astronomical set of possibilities. More than approximately a decade ago, there really was no way to do it. The problem has become solvable at all only with the availability of fast computers and specialized software applications (such as Anagram Genius, which is highly entertaining and I recommend it to any of you who don't already have it). Modern PCs can run through many permutations per second, but that alone isn't enough, so the applications work in a stepwise manner—first using a large word database to find all the 1-letter, 2-letter, 3-letter, etc. words that can be made from those in the original sequence, then searching for groupings of these words that use each letter of the original sequence exactly once—to weed out permutations that are guaranteed to be nonsensical and to run through only the ones that have a prayer of satisfying the requirements.