This is a section of a paper I wrote several years ago:

> Great Expectations. in Tim Lewens, ed. *Risk: A Philosophical View* Routledge, 2007, 84-98, also on this site (follow the 'research' link.)

The paper is about risk-taking virtues, but in the section below I explain the connection between risk-taking and social inequality, how in a society of individuals gaining from and losing from risks, increasingly greater inequality will result, without any diversity in skills or effort, simply from the accumulated effects of chance. (In the absence of interventions such as redistributive taxation schemes or limits of inheritance.) I have put the relevant conclusion in **blue**, so that you can scroll down to it to see if it interests you. The wider context is important, though, as it shows that this is a part of a very general mathematical fact.

**variability and expectation: three facts** An agent is facing a choice between two options. The options could be taking the right or left fork in the road. How things will turn out after her choice depends on some facts that she does not know. (Perhaps whether there are still bandits in this territory, who will attack travellers who take the short-cut path.) Suppose that one of the options is riskier than the other, in that it might turn out much better facts are to be true. (She knows that there have been no attacks for years, though not all of the known bandits have been caught.) Then she can ask 'do the potential benefits of taking the first option outweigh its potential costs?' She has to ask herself this, because this is the essence of the situation she faces. But the concept of outweighing demands a lot. It asks her to count for an action its benefits and the probability that these will follow, to count against it its costs and the probability that these will follow, and to combine these probabilities and values in a way that allows them to be compared. The result of the combination is the *expectation* of the action: its benefits weighted by the probability they will occur reduced by its costs weighted by the probability *they* will occur.

How can we calculate expectations? The standard model of a situation in which the calculations are unproblematic is given by games of chance. Suppose that instead of a fork in the road our agent faces a choice between two gambles, $o_1$ and $o_2$. In $o_1$ a fair coin will be tossed: if it lands Heads she wins $200, and if it lands Tails she looses $100. In o2 she gets $30 whatever. (She might have paid to be in the situation where this choice is open to her. The person offering the gambles does not have to be benevolent.) If the gamble were repeated many times and she took $o_1$ every time her gains would approach $ 50 times the number of repetitions, since she would win roughly half the time and loose roughly half the time. And if she took o2 every time her gains would approach $30 times the number of number of repetitions. (Or, equivalently, if a zillion duplicates of her were to take o1 they would end up with 50 zillion dollars, and if they took o2 they would end up with 30 zillion.) So this aspect of the gamble, the average amount that it would yield if repeated indefinitely, is clear. For each option it is the probability of Heads times the benefit from Heads for that gamble plus the probability of Tails times the benefit from Tails. Then it is just a small step to taking this expected or average value to be 'what the gamble is worth', to specify how the option should be ranked in comparison with other options. (Standard expositions of this idea are in chapter 1 of Jeffrey 1983, chapter 2 of Raiffa 1968, and chapter 3 of Resnik 1987. For some history see Hacking 1975, especially chapter 11.)

Suppose that the agent takes this small step, and evaluates the options by their expected value. Then the ‘average’ amount she will gain from the risky option will be $50, which is better than the $30 she will get if she takes the less risky one. ‘Average’ here means average over some indefinitely large set of possible occasions. It could be all the ways things might turn out if she takes the option, or it could be the way things would turn out if she (impossibly) were to repeat the choice over and over again. After the risky choice her possible selves will fall into two classes. Half of them will be $200 better off and half $100 poorer, and after the riskless choice all her possible selves will have the same profit of $30. So the downside for the greater expected outcome is the greater variability of actual outcomes.

This fact is completely general. Take one gamble to be riskier than another when its possible outcomes are more varied. (There are several ways of making this precise, and their differences do not matter here. For simplicity take a riskier gamble to have a greater variance of distribution of outcomes.) It will follow that people who make riskier choices will experience more varied outcomes than those who make safer ones. Consider the effects of this on a population of people in a gamble very similar to the one just described. A coin is tossed: if it lands Heads players gain $1, if Tails they lose $1. The game can continue, and then after the second toss each player may have $2 (after two successive Heads), $0 (after Heads-Tails, or Tails-Heads) or $-2 (after two successive Tails). And so on. Call this game g-risky and compare it to an alternative g-safe in which the players win or lose nothing. The two games have the same expected value, 0. Consider two sub-populations of players, one playing each game. After one round of g-two rounds about one quarter will be richer by $2, one half will have returned to zero, and one half will be poorer by $2. And of course the whole of the second sub-population, playing (or perhaps not-playing) g-safe will have remained at zero. And if we continue to play the game with more and more rounds then though the average wealth of the two populations is the same, it is distributed very differently. In the first population there are eventually some extremely wealthy people and some grotesquely indebted ones, while in the second population no one has changed their wealth more than anyone else.

Some very general facts are beginning to emerge. Fact number one: *when two gambles have the same expected value the riskier one will also produce a wider distribution of results*: more or greater winners and also more or greater losers.

**There are social consequences of this fact, though they are not the focus of this chapter. A society in which people are free to take risks for their own benefit will often end up with a higher average wealth, but it will also end up with a greater variation in wealth, so that it is quite easy for many people to be worse off as a result of the choices that lead to an increase in the average well-being. (And the people who emerge well off will compliment themselves on their wise choices and their sense of opportunity, when often the fact will be that they are**

**the few for whom the coin came down Heads many times.)** More to the present point is the consequence that a riskier gamble can make it more likely that one does badly. Or, to put it more carefully, given two gambles with the same expected value, the riskier one will sometimes present a larger probability of emerging with less than the expected value of the gamble. And, more generally, sometimes though one gamble has a higher expected value than another it also makes it more likely that you will do worse than you will if you had taken the other. (Two ways: it can be more likely that you will do worse than the expected value of the other, and it can also be more likely that you will do worse than the most likely outcome of the other gamble.) This can be illustrated by variants of the g-risky game just considered or by cases like the shifty character story at the beginning of this chapter. Fact number two: *by choosing a gamble which has a higher expected value but also a greater risk, one can often increase the probability of doing badly.*