Dynamics of Death in the Lee-Carter model

Shripad Tuljapurkar

The Lee-Carter model is widely used, both in its original form and newer versions with fancier formulations. The model builds on a remarkable constant exponential rate of decline in mortality rates. It has often been remarked that life expectancy under this model seems to increase linearly over multi-year projections but it is not clear why. We also know that the variance in age at death has fallen over time in many countries but we don't know if and why the Lee-Carter model shows this trend. but has not been examined for this trend.

We study the dynamics of life expectancy and variance in age at death when mortality follows a Lee-Carter model. We present analytical results to show under what circumstances the model predicts linear increases in life expectancy. Our results exploit the fact that the probability distribution of period age at death is sharply peaked in modern human experience. Using this fact we derive excellent approximations to the rate of change of the mode and variance of the distribution of adult age at death. These in turn tell us most of what we need to know about changes in e0 and the variance. We show that the rates of change of the mode and variance depend on the covariance between the age-response function in the Lee-Carter model and the current age distribution of log-mortality. We explain why this has predicted a nearly linear increase in e0 in the past. We also describe conditions under which the variance in age at death will decrease and/or increase overtime. We show that these results illuminate what happens in recent data, especially in terms of possible increases in the variance in age at death. We also argue that a study of forecast variance in relation to the data provide a sensitive test of projection models and their assumptions.

Finally we consider a difference of opinion that has been around for some time about fitting the model to observed death rates. The original work and many subsequent papers have used a two-stage fitting process to make the model hew more closely to the data. Others have proposed using higher-order variants (in the sense of singular value projections). We argue that the structure and assumptions imply that there should be no second-stage fitting, because we are identifying a linear subspace that supports the main Lee-Carter age response. Second-stage methods attempt to load information from orthogonal subspaces onto this main subspace and this process must surely fail. Higher-order variants avoid this error.