Example 1. Assume that you want to estimate the mean $\mathbb{E}_{\mu^\dagger}[X]$ of some random variable X with respect to some unknown distribution $\mu^\dagger$ on the interval $[0,1]$ based on the observation of n i.i.d. samples, given to finite resolution δ, from the unknown distribution $\mu^\dagger$. The Bayesian answer to this problem is to assume that $\mu^\dagger$ is the realization of some random measure distributed according to some prior π (i.e. $\mu \sim \pi$) and then compute the posterior value of the mean by conditioning on the data. Now to specify the prior π you need to specify the distribution of all the moments of μ (i.e. the distribution of the infinite dimensional vector $(\mathbb{E}_\mu[X], \mathbb{E}_\mu[X^2], \mathbb{E}_\mu[X^3], \ldots)$). So a natural way to assess the sensitivity of the Bayesian answer with respect to the choice of prior is to specify the distribution $\mathbb{Q}$ of only a large, but finite, number of moments of μ (i.e. specify the distribution of $(\mathbb{E}_\mu[X], \mathbb{E}_\mu[X^2], \ldots, \mathbb{E}_\mu[X^k])$ where k can be arbitrarily large). This defines a class of priors Π and our results show that no matter how large k is, no matter how large the number of samples n is, for any $\mathbb{Q}$ that has a density with respect to the uniform distribution on the first k moments, if you observe the data at a fine enough resolution then the minimum and maximum of the posterior value of the mean over the class of priors Π are 0 and 1.