government treasury bills – as a hedge against low-consumption states. Consequently, such relatively safe assets bear very low, even negative, rates of return.

I do not believe that it will be easy to dismiss such type of Bayesian statistical explanation for the equity premium puzzle. After all, the *qualitative* fact that $E[r_e]-r_f$ is positive comes as no surprise, just from first principles of risk aversion. The equity premium puzzle is the *quantitative* paradox that the observed value of $E[r_e]-r_f$ is too big to be reconciled with the standard neoclassical stochastic growth paradigm. But *compared with what* is the observed value of $E[r_e]-r_f$ "too big"? The answer given in the equity-premium literature is: "compared with the right hand side of formula (19)." Unfortunately for this logic, the right hand side of (19) is in practice a very bad estimate of the true value of $E[r_e]-r_f$ as given by equations (13) or (16). Anyone wishing to downplay this line of reasoning in favor of the *status quo ante* would be hard pressed to come up with their own Bayesian rationale for calibrating variances of non-observable subjectively-distributed future growth rates by point estimates equal to past sample averages.

In effect, the frequentist-inspired literature that produces the family of equity puzzles avoids the consequences on expected utility of non-uniform convergence (in n, for any positive δ) only by imposing the pointwise-convergent extreme case $m=\infty$ right from the beginning. Given any model of utility, it is well known that in principle there exist subjective probabilities that can produce the necessary marginal-utility state shadow-weights to "explain" the observed prices of traded financial assets. The interesting question then becomes: are these subjective probabilities sufficiently close to objective probabilities to be plausible? The existing literature errs by attempting to address this question in the non-relevant space of observed past consumption, where the answer is negative, instead of in the relevant space of non-observed subjectively-expected utility of future consumption, where the answer is positive.

We are witnessing growth data from the past that look as if they are normally distributed with mean $\hat{g}$ and variance $\hat{V}$. But the corresponding Bayesian posterior distribution, which is required to evaluate properly the true impact of risk aversion embodied in Theorem 1, indicates that the all-important difference is an unnoticeable (for large $m+n$) upward adjustment in the probabilities of the higher-variance scenarios. Theorem 1 says that once we compute the Bayesian equity premium rigorously, then the paradox recedes. The underlying statistical reason is that with