less than an infinite amount of data, the t-like part of the probability density function (33) has a sufficiently plump tail to make very dramatic the Bayesian expected-utility implications of model uncertainty, as captured here by uncertain structural parameter values. The thesis of this paper is that such kind of model uncertainty drives the entire family of equity 'puzzles,' thereby giving rise to a very different world view than what emerges, e.g., from just plugging into the equity-premium formula (13) a normal distribution whose variance is point-calibrated to past values.

In an early attempt to explain the equity premium puzzle, Rietz (1988) argues that we cannot exclude the possibility that our sample size is not large enough to describe adequately the full macro-risk of unknown future growth rates. The impact on financial equilibrium of a situation where there is a tiny probability of a catastrophic out-of-sample event has been dubbed the "peso problem." In a peso problem, the small probability of a disastrous future happening (such as a collapse of the presumed structure from a natural or socio-economic catastrophe) is taken into account by real-world investors (in the form of a "peso premium") but not by the model, because such an event is not in the sample. The fault is not really with the model, so to speak, but rather the fault is that the modeler is forcing the rational-expectations sample variance of past growth rates to stand in for the overall risk-adjusted effects on expected utility of a future growth process whose basic structure is unknown.

I think Theorem 1 is trying to tell us that a statistical analogue of the peso problem may be generically ingrained in the "deep structure" of how Bayesian inferences about exponential processes (of future economic growth, at unknown rates) interact with a curved utility function. Bayesian inferences from finite data fatten the posterior tails of probability density functions, as the example of replacing the workhorse normal distribution by its t-like posterior distribution demonstrates with dramatic consequences (when expressed in units of expected marginal utility). This "Bayesian-statistical peso problem" means that it is not so absurd to believe that *no* finite sample size is large enough to capture all of the relevant structural model uncertainty concerning future economic growth. I think the Bayesian peso problem is trying to tell us that to calibrate an exponential process having an unknown growth rate, which is essentially intended to describe future worldwide economic prospects, by plugging the sample variance of observed growth rates from the past into a "very bad" approximation of the subjectively-distributed future growth rate, is