

element, which would capture the characteristic that appears to make stocks comparatively so uncertain, seems to be missing from the standard model. At least for asset pricing applications, a consensus has developed among economists that the standard model is seriously flawed.

Not surprisingly therefore, this family of equity macro-puzzles has stimulated a lot of economic research. In attempting to explain the paradoxes, an impressive literature has developed, which is filled with some imaginatively fruitful variations on the standard model. To overcome one or another equity puzzle, many new models feature exotic reverse-engineered formal (or behavioral informal) preferences having aggregated coefficients of relative risk aversion that are typically very high, time-varying, and correlated with the real economy. Some valuable insights have come from these new models, but it still seems fair to say that no new consensus has yet emerged from within the economics profession as a whole that the puzzles have been satisfactorily resolved.

The point of departure for this paper is to note that, throughout the existing literature, the risk premium and the risk-free rate are effectively calibrated by plugging into the relevant formulas the sample mean and sample variance of past growth rates. But strictly speaking, the correct procedure requires inserting the full subjective probability distributions of uncertain structural parameters of the model, not just their point estimates. Missing from the framework is a formal incorporation of the decision-theoretic specification required to make a rigorous statistical-economic general-equilibrium growth model. In effect, the implicit statistical methodology assumes that the time series are long enough that the law of large numbers allows substituting the sample moments of past growth rates for the population moments of future growth rates. This intuitive methodology may well be justified in some economic applications. Unfortunately, as the paper shows, such point calibration is a fatally flawed procedure for the particular application of analyzing aversion to model-structure uncertainty, which underlies (or, more accurately, *should* underlie) all asset-pricing calculations. The core problem is that calibrating population moments to sample frequencies understates significantly the researcher’s (and the investor’s) utility-weighted predictive uncertainty, which spills over into biased expectation formulas that produce dramatically incorrect asset valuations.