

tractable, garden-variety specifications in order to be able to derive a relatively transparent expression for the family of equity discrepancies, it will become apparent that the basic insights have much broader applicability.

This paper is far from being the first to investigate the effects of Bayesian statistical uncertainty on asset pricing. Earlier examples include Barsky and DeLong (1993), Timmerman (1993), Bossaerts (1995), Cechetti, Lam and Mark (2000), Veronesi (2000), Brennan and Xia (2001), Abel (2002), Brav and Heaton (2002), Lewellan and Shanken (2002), and several others. Broadly speaking, these papers indicate or hint, either explicitly or implicitly, that the need for Bayesian learning about structural parameters tends to reduce the degree of one or another equity anomaly. What has been missing from this literature, however, is a generic appreciation of the incredible *strong force* that tail-fattening structural parameter uncertainty brings to bear on asset pricing by its utterly dominating influence over the outcome of any calculation involving expected marginal utility. In effect, the direction of this Bayesian force is appreciated in the literature, but not its strength.

The one noteworthy exception is an important paper by John Geweke (2001), who applies a Bayesian framework to the most standard model prototypically used to analyze behavior towards risk and then notes the extraordinary fragility of the existence of finite expected utility itself.<sup>1</sup> In a sense the present paper begins by accepting this non-robustness insight, but pushes it further to argue that the inherent fragility of the standard prototype formulation constitutes an important clue for unraveling what may be causing the equity puzzles and for giving them a unified general-equilibrium interpretation that simultaneously fits the stylized time-series facts.

This paper will end up arguing that there are no equity ‘puzzles’ as such arising from within a Bayesian framework. Instead, the arrow of causality in a unified Bayesian explanation is reversed: the ‘puzzling’ numbers being observed empirically are trying to tell us something important about the implicit background prior distribution of structural model-parameter uncertainty that is generating such data. In the final section of the paper the three ‘puzzling’ time-

---

<sup>1</sup> I wish to express my gratitude to two readers of a previous version of this paper for making me aware of Geweke’s earlier article after noticing that I had derived a similar result.