plus infinity. (It is essentially in order to make this moment-generating integral converge that the condition $\delta>0$ is imposed in the first place.)

This entire preliminary discussion of the future consequences of what people now think that people long ago "might have been thinking" about such things as an upper bound on V (of $1/\delta$) or a degree of prior confidence in $\hat{V}$ (of m) has an unreal tone about it. In practice this issue ought to be non-operational – and therefore not worth contemplating – because the intervening n observations should have bleached the prior parameters out of the posterior distribution. Thus, if the number of data points n is large enough, it "should not matter" what values of δ or m we select now to represent past beliefs. This "should not matter" intuition is true, it turns out, for the parameter m, whose effects on expected utility converge *uniformly* in n for all $m>0$. However, the parameter δ behaves fundamentally differently, because its effects on expected utility *do not* converge uniformly in n for all $\delta>0$. In this sense there is a critical distinction, which is crucial for all expected-utility asset-pricing implications, between not knowing what value to assign now to the prior parameter m and not knowing what value to assign now to the prior parameter δ.

The fact that expected utility is *not* uniformly convergent in n for all positive δ has great significance for the interpretation of this paper. A prior distribution is *our* imputation *now* of what "they might have" imposed n years ago during the pre-data past. It is essentially a mental artifice for framing a subjective thought-experimental dialogue between the present and the past about what to expect from the future. In such a setting, *pointwise* convergence of expected utility in n for a given δ is not nearly enough to guarantee a robust prior, because the prior is a subjective creature of *our* imagination *now*, not an objective unchangeable reality that a real person carved in stone n years ago to represent some intrinsic characteristic of the then-observable world.

To have faith in the standard practice of calibrating means and variances of normal distributions to past historical averages presupposes a robustness in the interpretation of observable data with respect to whatever values of δ or m are chosen. Therefore, a necessary precondition for the validity of the classical statistical idea to just "let the data speak for themselves" is that the effects of δ or m should be negligible for sufficiently large n. This condition holds (in the space of expected utility) for m, but such a robustness condition *does not* hold (in the space of expected utility) for δ . The value of δ that has now been chosen to