Example 2. Assume that you want to estimate the mean $\mathbb{E}_{\mu^\dagger}[X]$ of some random variable X with respect to some unknown distribution $\mu^\dagger$ on the interval $[0,1]$ based on the observation single data point $d_1 \approx 0.5$ up to resolution δ (i.e. we observe $d_1 \in B_\delta(x_1)$ with $x_1 = 0.5$). Consider the following two Bayesian models $\mu^a(\theta)$ and $\mu^b(\theta)$ on the unit interval $[0,1]$, parameterized by $\theta \in (0,1)$, with densities $f^a(\theta)$ and $f^b(\theta)$ given by

$$f^a(x,\theta) = (1-\theta)\big(1+\tfrac{1}{\theta}\big)(1-x)^{\frac{1}{\theta}} + \theta\big(1+\tfrac{1}{1-\theta}\big)x^{\frac{1}{1-\theta}},$$

$$f^b(x,\theta) = \begin{cases} f^a(x,\theta)\frac{1}{Z(\theta)}\big(\mathbb{1}_{\{x\notin(x_1-\frac{\delta_c}{2},x_1+\frac{\delta_c}{2})\}} + 10^{-9}\mathbb{1}_{\{x\in(x_1-\frac{\delta_c}{2},x_1+\frac{\delta_c}{2})\}}\big), & \text{if } \theta < 0.999,\\ f^a(x,\theta), & \text{if } \theta \geq 0.999,\end{cases}$$

where $Z(\theta)$ is a normalization constant close to one, chosen so that $\int_{[0,1]} f^b(x,\theta)\,dx = 1$. See Figure 1 for an illustration of these densities.

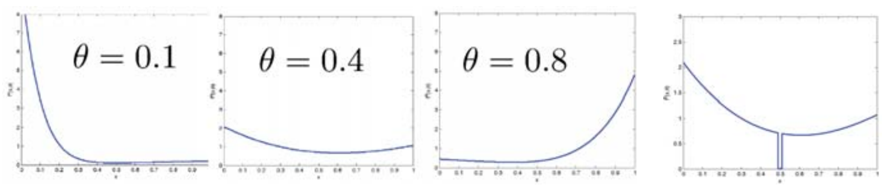

(a) $f^a(x,\theta)$ (b) $f^b(x,\theta)$

Figure 1: Illustration of the density $f^a(x,\theta)$ of model a and $f^b(x,\theta)$ of model b.

Observe that the density of model b is that of model a besides the small gap of width $\delta_c > 0$ created around the data point when $\theta < 0.999$ (see Figure 1); since the data point is fixed at $x_1 = \frac{1}{2}$, the total variation distance $d_{\mathrm{TV}}\big(\mu^a(\theta),\mu^b(\theta)\big)$ between the two models is, uniformly over $\theta \in (0,1)$, less than a constant times δ_c. Assuming that the prior distribution on θ is the uniform distribution on $(0,1)$, observe that the prior value of the quantity of interest $\mathbb{E}_\mu[X]$ under both models (a and b) is approximately $\frac{1}{2}$. Now, when θ is close to one (zero) then the density of model a puts most of its mass towards one (zero). Observe also that the density of model b behaves in a similar way, with the important exception that the probability of observing the data under model b is infinitesimally small for $\theta < 0.999$. Therefore, for $\delta < \delta_c$, the posterior value of the quantity of interest $\mathbb{E}_\mu[X]$ under model a is $\frac{1}{2}$ whereas it is close to one under model b (because, under the density $f^b(x,\theta)$, the probability of observing the data, up to resolution δ, is extremely unlikely if $\theta < 0.999$ when compared to the alternative $\theta \geq 0.999$). Observe also that a perturbed model c analogous to b can easily be constructed which leads to a posterior value close to zero. This simple example of brittleness under infinitesimal model perturbations is derived from the proof of theorem 6.4 of http://arxiv.org/abs/1304.6772, which shows that Bayesian posterior values are generally brittle under infinitesimal perturbations of Bayesian models in TV and in Prokhorov metrics.