Example 5. Figure 3 gives an illustration when the parameter space Θ is the interval $[0,1]$ and the data generating distribution is $P(\theta^*)$. The probability density functions of π and π' are p and p' with respect to the uniform distribution on $[0,1]$. In sub-figure (a), π is an arbitrarily small perturbation of π' of sub-figure (c) in total variation. π_n converges towards δ_{θ^*} while the distance between the support of π'_n and θ^* remains bounded from below by $a > 0$. Note that the mechanism illustrated in Figures 2 and 3 does not generate non qualitative robustness at *all priors* but instead for the full class of *consistency priors*, defined by the assumption of having positive mass on every Kullback-Leibler neighborhood of θ^*. One may wonder whether this non *qualitative robustness* can be avoided by selecting the prior π to satisfy Cromwell's rule (that is, the assumption that π gives strictly positive mass to every nontrivial open subset of the parameter space Θ), as in sub-figure (b). The answer is that this is not the case if the

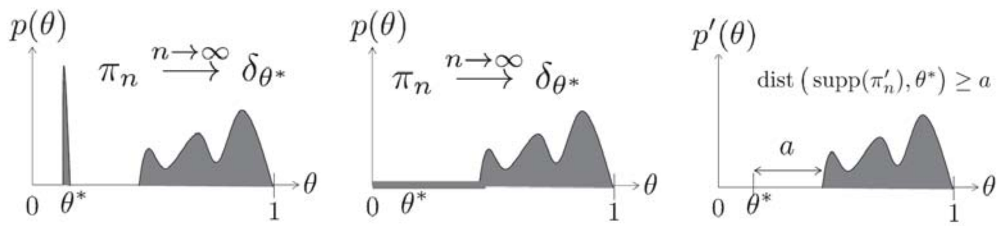

(a) Density of π (b) Density of π (c) Density of π'

Figure 3:

parameter space Θ is not totally bounded. For example, when $\Theta = \mathbb{R}$, for all $\delta > 0$ one can find $\theta \in \mathbb{R}$ such that the mass that π places on the ball of center θ and radius one is smaller than δ, and by displacing this small amount of mass one obtains a perturbed prior π' whose posterior distribution remains asymptotically bounded away from that of π when the data-generating distribution is $P(\theta)$. Similarly if Θ is totally bounded then one can place an upper bound on the size of the perturbation of the prior π that would be required as a function of the covering complexity of the parameter space Θ. Note that these observations suggest that a maximally *qualitatively robust* prior should place as much mass as possible near all possible candidates θ for the parameter θ^* of the data generating distribution, thereby reinforcing the notion that a maximally robust prior should have its mass spread as uniformly as possible over the parameter space.