

Example 7. Figure 5 represents a non-parametric generalization of the mechanism of Figure 4. Assume that the data generating distribution is $\mu \notin P(\Theta)$, so that the model is misspecified. Let $\pi \in \mathcal{M}(\Theta)$ be an arbitrary prior distribution and $P_*\pi \in \mathcal{M}^2(X)$

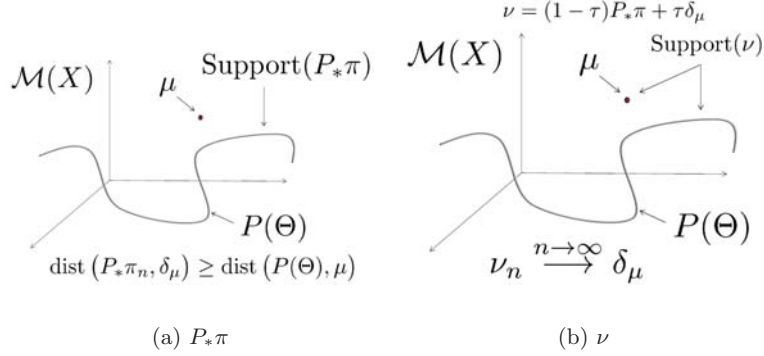


Figure 5: Non-robustness caused by misspecification

its corresponding non-parametric prior. By removing an arbitrarily small amount of mass from $P_*\pi$ and placing it on μ one obtains an arbitrarily close prior distribution ν that is consistent with respect to the data generating distribution μ . Therefore although $P_*\pi$ and ν may be made arbitrarily close, their posterior distributions would remain asymptotically separated by a distance corresponding to the degree of misspecification of the model (the distance from μ to $P(\Theta)$).