INCLUDE	**EXCLUDE**
Include if it is important to predictions	We might just be fooling ourselves (relying on something we don't really understand)
Once we have the data it is easier to incorporate if the feature is represented	Exclude if we think it is not important to our desired predictions
Enables sensitivity analysis	Exclude if it decreases calibration quality. If it calibrates well without it, that may not be evidence that it doesn't exist, but it is not evidence that we have to have it in the model.
If you don't have it in the model, you can't see the uncertainty associated with it. The uncertainty is still there, you just can't see how it affects the prediction.	Including features we know little about opens us to criticism
	It may cause numerical instability or model convergence issues