Centring costs $O(Np)$ once. It removes one unpenalized variable from each subsequent fit and simplifies the normal equations. It does not change the asymptotic cost of a dense slope solve. Reusing centring across a cross-validation split requires the same affine model and unpenalized intercept; data-dependent scaling has separate requirements.