A Mean Field Approximation for LDA

$$Q(\Phi, \mathbf{Z}, \Theta \mid \lambda, \pi, \gamma) = \prod_{i=1}^{K} Dir(\phi_i \mid \lambda_i) \prod_{j=1}^{M} q_j(\theta_j \mid \gamma_j) \prod_{t=1}^{N_j} q_j(z_{j,t} \mid \pi_{j,t})$$

- Uses a **fully factorized** surrogate distribution
- λ, γ, π are the **variational parameters**
- These are adjusted to minimize $KL(Q \mid\mid P)$