



Figure S.7: Trace plot of the latent binary values obtained from BAYES (top) and HESS (bottom) in SIM1 for the 10 true positive associations simulated in the third hot-spot ($j = 225, k = 91, \dots, 100$). For the 25 replicates, the output (γ_{kj}) of each algorithm was piled up giving rise to a vector of 50,000 ($2,000 \times 25$) and 125,000 ($5,000 \times 25$) sweeps, respectively. Red dot and blue cross indicate $\gamma_{kj} = 0$ and $\gamma_{kj} = 1$, respectively. HESS correctly identifies the 10 transcript-marker associations as indicated by a large majority of blue crosses. Good MCMC mixing is clear from the sequence of blue crosses interrupted by red dots and *vice versa*. On the contrary, BAYES misses the simulated associations (false negative) and gets stuck in $\gamma_{kj} = 0$ producing long stripes of consecutive red dots. Overall, the different efficiency in the MCMC mixing between BAYES and HESS is apparent from the diverse coloured stripe patterns.