applied in its entirety (see [25, 37]). Thus, any symmetric preconditioning matrix may be easily integrated into the conjugate gradients algorithm.

Although CG computes the desired solution $A^{-1}b$, it does not explicitly calculate any entries of A^{-1} in the process, and thus does not directly provide error covariance information. For inference problems defined by the standard normal equations (2.1, 2.2), Schneider [65, 67] has developed a method for extracting the error variances directly from the search directions generated by the CG iteration. Similar methods may also be applied to the information form of the normal equations (2.6, 2.7) considered in this thesis [9, 57]. However, accurate error variances cannot be obtained without reorthogonalization of the search directions, which may greatly increase the computational cost of the CG iteration. In addition, the error variances often converge quite slowly relative to the conditional mean estimates.

2.4.3 Stopping Criteria

When applying any of the iterative methods considered in the previous sections, it is necessary to determine when to stop the iteration, i.e. to decide that x^n is a sufficiently close approximation to $A^{-1}b$. Given a tolerance parameter ϵ, we would ideally like to iterate until

$$||A^{-1}b - x^n||_2 \leq \epsilon \tag{2.57}$$

Unfortunately, because $A^{-1}b$ is unavailable, equation (2.57) cannot be evaluated. However, the residual $r^n = b - Ax^n$ *is* available. Therefore, as suggested in [7], we may instead iterate until

$$\frac{||r^n||_2}{||b||_2} \leq \epsilon \tag{2.58}$$

Using the identity $(A^{-1}b - x^n) = A^{-1}r^n$, equation (2.58) is easily shown to yield the following bound on the final error:

$$||A^{-1}b - x^n||_2 \leq \epsilon \sigma_{\min}(A) ||b||_2 \tag{2.59}$$

All of the simulations in this thesis use this residual–based stopping criterion.