Table 21.3 Naive and CV Estimates of the MSE of $\hat{\beta}$

β_0	n	Naive	One Control Variate	Two Control Variates
0.1	25	.03739 ($.510 \times 10^{-3}$)	.03720 ($.317 \times 10^{-3}$)	.03728 ($.272 \times 10^{-3}$)
0.1	100	.00959 ($.134 \times 10^{-3}$)	.00973 ($.468 \times 10^{-4}$)	.00970 ($.390 \times 10^{-4}$)
0.1	400	.00252 ($.351 \times 10^{-4}$)	.00247 ($.650 \times 10^{-5}$)	.00246 ($.524 \times 10^{-5}$)
0.5	25	.03161 ($.522 \times 10^{-3}$)	.03171 ($.454 \times 10^{-3}$)	.03139 ($.384 \times 10^{-3}$)
0.5	100	.00777 ($.734 \times 10^{-4}$)	.00768 ($.696 \times 10^{-4}$)	.00767 ($.542 \times 10^{-4}$)
0.5	400	.00193 ($.281 \times 10^{-4}$)	.00187 ($.976 \times 10^{-5}$)	.00188 ($.756 \times 10^{-5}$)
0.9	25	.01725 ($.413 \times 10^{-3}$)	.01725 ($.413 \times 10^{-3}$)	.01731 ($.377 \times 10^{-3}$)
0.9	100	.00277 ($.563 \times 10^{-4}$)	.00276 ($.548 \times 10^{-4}$)	.00274 ($.439 \times 10^{-4}$)
0.9	400	.00054 ($.922 \times 10^{-5}$)	.00053 ($.748 \times 10^{-5}$)	.00053 ($.534 \times 10^{-5}$)

Table 21.3 shows naive estimates and two sets of CV estimates of the mean squared error of $\hat{\beta}$, for the same nine cases as Table 21.2. Using only one control variate, (21.19), generally yields more accurate estimates than using no control variates, and using two control variates, (21.19) and (21.20), always works better than using only one. However, the gains relative to the naive estimator are always less than those achieved when estimating the mean; compare Table 21.1. This illustrates the general result that control variates tend to be most helpful for estimating means and progressively less helpful for estimating higher moments; see Davidson and MacKinnon (1992b).

Given the highly variable gains from using control variates, it may be advisable in cases for which computational costs are large to determine the number of replications N adaptively. One could decide in advance the acceptable level of precision for the various quantities to be estimated, then calculate those quantities for an initial fairly small value of N (perhaps 500 or so), and use those initial results to estimate how many replications would be needed to obtain standard errors that are sufficiently small. Alternatively, one could calculate standard errors of the quantities of interest after every few hundred replications, stopping when they are sufficiently small. In practice, few Monte Carlo experiments have been designed this way; N is generally just fixed in advance, and the precision of the estimates is whatever it turns out to be.

21.7 Response Surfaces

As we have stressed above, one of the most difficult aspects of any Monte Carlo experiment is presenting the results in a fashion that makes them easy to comprehend. One approach that is sometimes very useful is to estimate a **response surface**. This is simply a regression model in which each observa-