

neglected. However, if we had included an additional error, we would find the weights much more comparable. Of course, we can choose any values for  $w_{ij}$  that we wish- the only effect would be to increase (or decrease) the error in the final result. Therefore, we attempt to choose weights which increase the final error by a negligible amount and also help compensate for any non-formal measurement errors. Since a star with a very high signal-to-noise does not improve the final QSO light curve much more than a star with a signal-to-noise roughly twice that of the QSO, we “clip” the error in the stars used for the weightings such that percentage error is not less than half that of the QSO error. This will allow more stars to contribute to the calculation of the final light curve and yet not substantially increase the final errors. Note that for all computations other than these weights, we use the unclipped formal errors.

Before computation, the values  $s_{ij}$  are rescaled to minimize numerical roundoff problems. First, we rescale all the values of  $s_{ij}$  for each image  $i$  based on the values of  $s_{i'j}$  in one selected image  $i'$ . This eliminates most of the large intensity differences due to airmass and exposure time. Next, we rescale  $s_{ij}$  for each star  $j$  such that the mean of  $s_{ij}$  for a given star  $j$  is one. As long as we use the same rescaling factor for all the stars in a given image, and the same rescaling factor for a given star in all the images, we have not affected the final result.

The computation of the values  $c_i$  are done iteratively. The scaling factor for each image is computed in turn, with the cycle repeating until the weighted sum of the errors in  $c_i$  converges. At this point, values which deviate by more than  $3\sigma$  are thrown out ( $w_{ij} = 0$ ), and the fit is repeated. This process is repeated until no stars are rejected at the end of a cycle. A larger rejection threshold,  $6\sigma$ , is used when throwing out the last few stars in any given image (to avoid having any value of  $c_i$  be based on a single comparison star).

Note that the calculation of  $c_i$  for a given image  $i$ , uses the data from all the other comparison stars in all the images for a given field. Each image effectively gives us information on the relative intensities between the comparison stars, which are assumed