is already close to the converged solution. This strategy may fail, however, in the MIT coexistence region. For illustration, let us assume that the target discretization is $\Delta\tau = 0.1$ and that we are studying a phase point which corresponds to the black cross in Fig. 3.13 a, where both metal and insulator are stable (for $\Delta\tau = 0.1$). If the convergence process is started using $\Delta\tau = 0.25$, the intermediate result is a metallic solution, since the insulator is, then, unstable. Obviously, subsequent iterations for smaller $\Delta\tau$ can only increase the precision of the metallic solution, but are unlikely to cause "jumps" towards the insulating solution. A second, comparatively minor aspect of the explanation is that Jarrell initially accelerates the convergence of the iteration scheme using overrelaxation. At least in regions of the phase space where the insulator is thermodynamically metastable, this might also render insulating solutions numerically unstable and, eventually, unobservable.

3.4.3 Estimation of Errors

Generally, QMC estimates for any observable have both statistical and systematic errors. Systematic errors which arise from using finite $\Delta\tau$ such as the remaining error in the Fourier transformation or the inevitable Trotter error will be explicitly considered in Sec. 3.5. More precisely, most results will first be presented for a range of finite $\Delta\tau$. Error bars will indicate the statistical errors as determined empirically from a set of measurements (taking autocorrelation into account). Only in a second step, conclusions are drawn for the physical limit $\Delta\tau = 0$.

The statistical error bars computed for finite $\Delta\tau$ are only a good approximation for the true error when other sources of errors are removed. One possible source is the thermalization of the auxiliary Ising field used for the MC solution of the impurity problem. In our code, this field is initialized randomly for each iteration. Generically, these initial configurations have a vanishingly small true probability of occurring in an infinitely long Monte Carlo run.[18] Consequently, they would be overweighted in any run of finite length. As a resolution, a certain number of Monte Carlo "warm-up" sweeps after initialization is discarded before starting measurements. Schlipf (1998) also introduced the possibility of using multiple walkers in each iteration, i.e., of averaging over completely independent solutions of the impurity problem (with identical $\mathcal{G}(\tau)$). Thereby, the result of an iteration is less influenced by a particularly "bad" initial Ising configuration. We typically used 2000-4000 warm-up sweeps and 6 or 8 walkers, in parallel runs up to 32 walkers (i.e., starting points). A study on the influence of the number of warm-up sweeps is shown in Fig. 3.14. Here, the double occupancy is computed for a relatively high temperature ($\beta = 10$) using between 0 and 10000 warm-up sweeps. Since the curves are practically on top of each other in the main panel, the differences with respect to the unweighted average of the results (for 2000, 3000, 5000, and 10000 warm-up sweeps) are plotted in the inset for a magnified view. Runs without any warm-up sweeps show a significant systematic error within the metallic phase; all other runs agree within error bars. Here, the

[18] Note that this effect is already strongly reduced by random initialization compared with using an ordered initial Ising field. After all, the usage of only the initial Ising fields would correspond to a simple Monte Carlo procedure without importance sampling.