Chapter 9

Single-mode interacting Bose gas thermodynamics

In Section 5.6 it was shown that the gauge P representation can also be used to calculate properties of the grand canonical ensemble of an interacting Bose gas described by the lattice Hamiltonian (2.12). A convenient feature of calculations made with this method (rather than, say path integral Monte Carlo, or variational Monte Carlo methods) is that a single simulation can estimate any observables, and gives results for a wide range of temperatures $T \geq 1/k_B \max[\tau]$.

Before attempting full many-mode simulations (see Chapter 11), it is pertinent to make sure that the simplest special case of a single-mode is correctly simulated. This is a good test case because it can easily be solved exactly for a unambiguous comparison, and because it is simple enough that a broad analysis of the statistical behavior of the simulation is possible.

The single-mode model is particularly relevant to a locally-interacting many-mode Hamiltonian (2.17). Here the stochastic equations to simulate are given by (5.50), and it can be seen that the nonlinearity and noise in the evolution of local amplitudes $\alpha_{\mathbf{n}}$ for each mode $\mathbf{n}$ depend only on the local variables $\alpha_{\mathbf{n}}$ and $\beta_{\mathbf{n}}$. In such a situation the nonlinear and stochastic features of the many-mode behavior should appear in their entirety in the single-mode toy model.

9.1 The single-mode model

9.1.1 Quantum description

Isolating a single mode from the description of Section 2.6, the master equation for the evolution of the (un-normalized) density matrix is

$$\frac{\partial \widehat{\rho}_u}{\partial \tau} = \left[\mu_e(\tau)\widehat{n} - \widehat{H}\right]\widehat{\rho}_u \tag{9.1}$$

with $\widehat{n} = \widehat{a}^\dagger \widehat{a}$ (the number operator) defined in the usual way in terms of boson creation and annihilation operators. The "imaginary time" parameter is $\tau = 1/k_B T$, and the "effective chemical potential" μ_e is defined in (2.30). That expression can be integrated to give

$$\mu(\tau) = \frac{1}{\tau}\left[-\lambda_n + \int_0^\tau \mu_e(t)dt\right] \tag{9.2}$$

with constant λ_n, and it is then seen that at low temperatures $\mu(T) \to \mu_e(T)$. The single-mode analogue of the interacting Bose gas Hamiltonian is

$$\widehat{H} = \hbar\chi\widehat{n}\left(\widehat{n} - 1\right), \tag{9.3}$$

where any linear contribution $\propto \widehat{n}$ has been amalgamated into the chemical potential. The initial state is (as per (2.31))

$$\widehat{\rho}_u(0) = \exp\left[-\lambda_n\widehat{n}\right]. \tag{9.4}$$

Physically, this single-mode model is an approximation of e.g. a boson orbital in a heat bath at temperature T, and chemical potential μ, where this chemical potential includes any kinetic and external potential effects.

9.1.2 Gauge P stochastic equations

Defining

$$\breve{n} = \alpha\beta = n' + in'' \tag{9.5}$$

as in previous chapters, the Ito stochastic equations in a gauge P representation are obtained from (5.50):

$$d\alpha = (\mu_e - 2\hbar\chi\breve{n})\,\alpha\,d\tau + i\alpha\sqrt{2\hbar\chi}\,(\,dW - \mathcal{G}\,d\tau) \tag{9.6a}$$

$$d\beta = 0 \tag{9.6b}$$

$$d\Omega = \Omega\,[(\mu_e - \breve{n})\,n\,d\tau + \mathcal{G}\,dW]\,. \tag{9.6c}$$

The Wiener increments dW can be implemented as independent real Gaussian noises of mean zero, and variance $d\tau$ at each time step of the simulation. The single complex drift gauge is $\mathcal{G}$.

Because only a single non-weight variable α experiences diffusion, then the ungauged diffusion matrix is $\underline{D} = -2\hbar\chi\alpha^2$, a single complex function. There are then no standard diffusion gauges (as described in Section 4.4.2) because the orthogonal 1×1 matrix (4.68) is just $O = 1$, apart from a phase factor.

Initial conditions on the variables α and β are (from (5.59))

$$P_G(\alpha,\beta,\Omega) = \delta^2(\Omega - 1)\delta^2(\beta - \alpha^*)\frac{1}{\pi\overline{n}_0}\exp\left(\frac{-|\alpha|^2}{\overline{n_0}}\right), \tag{9.7}$$

where

$$\overline{n}_0 = \frac{1}{e^{\lambda_n} - 1}, \tag{9.8}$$

is the mean particle number at high temperature. α is easily sampled using Gaussian random variables.

The mean number of particles is $\langle\widehat{n}\rangle = \mathrm{Tr}\,[\widehat{\rho}_u\widehat{n}]\,/\mathrm{Tr}\,[\widehat{\rho}_u]$, and is estimated by

$$\overline{n} = \frac{\langle \mathrm{Re}\,\{\breve{n}\Omega\}\rangle_{\mathrm{stoch}}}{\langle \mathrm{Re}\,\{\Omega\}\rangle_{\mathrm{stoch}}}. \tag{9.9}$$

These stochastic equations exhibit completely different behavior to the dynamics of this system (7.3). Rather they are more similar to the two-boson absorption of Section 6.4.

9.1.3 Exact solution

The quantum evolution of this model is easily evaluated exactly in a Fock number state basis $|n\rangle$. This is useful to make a definitive check of the correctness of the

stochastic simulations. Using $\widehat{\rho}_u = \exp\left[\mu\tau\widehat{n} - \widehat{H}\tau\right]$, the density matrix elements in this basis $\rho_{n\widetilde{n}} = \langle n|\,\widehat{\rho}_u\,|\widetilde{n}\rangle$ are found to be:

$$\rho_{n\widetilde{n}} \;=\; \delta_{n\widetilde{n}} \exp\left\{\left[\mu - \hbar\chi(n-1)\right] n\tau\right\}. \tag{9.10}$$

A long times $\tau \to \infty$, i.e. low temperatures, the density matrix will be dominated by the populations for which the exponent is largest. This maximum of the exponent occurs when $n_\infty = [\lim_{\tau\to\infty}\mu(\tau) + \hbar\chi]/2\hbar\chi$. Since n takes on only integer values, then the dominant mode occupation will be at the nearest integer to n_∞ (or the two nearest components if n_∞ takes on half-integer values). Using (9.2),

$$n_\infty = \frac{1}{2} + \frac{1}{2\hbar\chi}\lim_{\tau\to\infty}\left[\frac{1}{\tau}\int_0^\tau \mu_e(t)dt\right], \tag{9.11}$$

and the $T \to 0$ state $\widehat{\rho}(\tau \to \infty)$ is:

- A vacuum if $n_\infty < 0.5$.
- Otherwise, if n_∞ is a half integer, then $\widehat{\rho}(\tau \to \infty)$ is an equal mixture of the $|n_\infty \pm \frac{1}{2}\rangle$ states.
- Otherwise, $\widehat{\rho}(\tau \to \infty)$ is the Fock number state with the nearest whole occupation number to n_∞.

9.2 Moving singularities and removal with gauges

9.2.1 Moving singularity

Consider, for now, the case of constant effective chemical potential μ_e. From (9.6), and using (B.7), the Stratonovich stochastic equation for the complex occupation variable $\breve{n}$ is

$$d\breve{n} \;=\; \breve{n}\left[\mu_e\,d\tau - \hbar\chi(2\breve{n}-1)\,d\tau + i\sqrt{2\hbar\chi}\,(\,dW - \mathcal{G}\,d\tau)\right]. \tag{9.12}$$

In a positive-P-like simulation[1], $\mathcal{G} = 0$, and the deterministic part of the $d\breve{n}$ equation has the phase-space structure shown in Figure 9.1**(a)**. There are stationary points

[1]But with a weighting gauge to allow for complex weight evolution $d\Omega$.

at vacuum and at

$$\breve{n} = n_\infty = \frac{1}{2} + \frac{\mu_e}{2\hbar\chi}, \tag{9.13}$$

with the more positive stationary point being an attractor, and the more negative a repellor. The deterministic evolution is easily solved, and with initial condition $\breve{n}(0) = n_0$ gives

$$\breve{n}(\tau) = \frac{n_\infty n_0}{n_0 + (n_\infty - n_0)e^{-2\hbar\chi n_\infty \tau}}. \tag{9.14}$$

A moving singularity appears along the negative real axis, and if one starts with a negative real $n_0 = -|n_0'|$, then $\breve{n} \to \infty$ at time

$$\tau_{\text{sing}} = \frac{1}{2\hbar\chi n_\infty} \log\left(1 - \frac{n_\infty}{n_0}\right). \tag{9.15}$$

Since initial conditions (9.7) lead to samples with $n_0 \in (0, \infty)$, then some low-n_0 trajectories can be expected to rapidly diffuse into the negative real part of phase space where the super-exponential growth occurs. This will tend to cause either misleading systematic boundary term errors or uncontrolled spiking, which renders the simulation useless after a short time. As is seen in Figure 9.5 in this case it is spiking.

9.2.2 Minimal drift gauge

To correct the problem one has to change the phase-space topology in some way to prevent the occurrence of such moving singularities. The radial evolution $d|\breve{n}| = -2\hbar\chi|\breve{n}|n'\,d\tau + \dots$ is at fault, and the offending term can be removed with the gauge

$$\mathcal{G} = i\sqrt{2\hbar\chi}\,(\mathrm{Re}\,\{n\} - |\breve{n}|)\,. \tag{9.16}$$

This gauge has the effect of replacing the $-2\hbar\chi|\breve{n}|\mathrm{Re}\,\{n\}\,d\tau$ term in $d|\breve{n}|$ with $-2\hbar\chi|\breve{n}|^2\,d\tau$, which always attracts trajectories into the phase space in the vicinity of the origin. Polar phase evolution is unchanged, and the resulting deterministic phase space is shown in Figure 9.1**(c)**. If $n_\infty > 0$ the attractor and repellor remain, while for $n_\infty \leq 0$, there is just an attractor at vacuum. For near-classical trajectories having real $\breve{n}$, phase-space evolution is unchanged, and the gauge is zero.

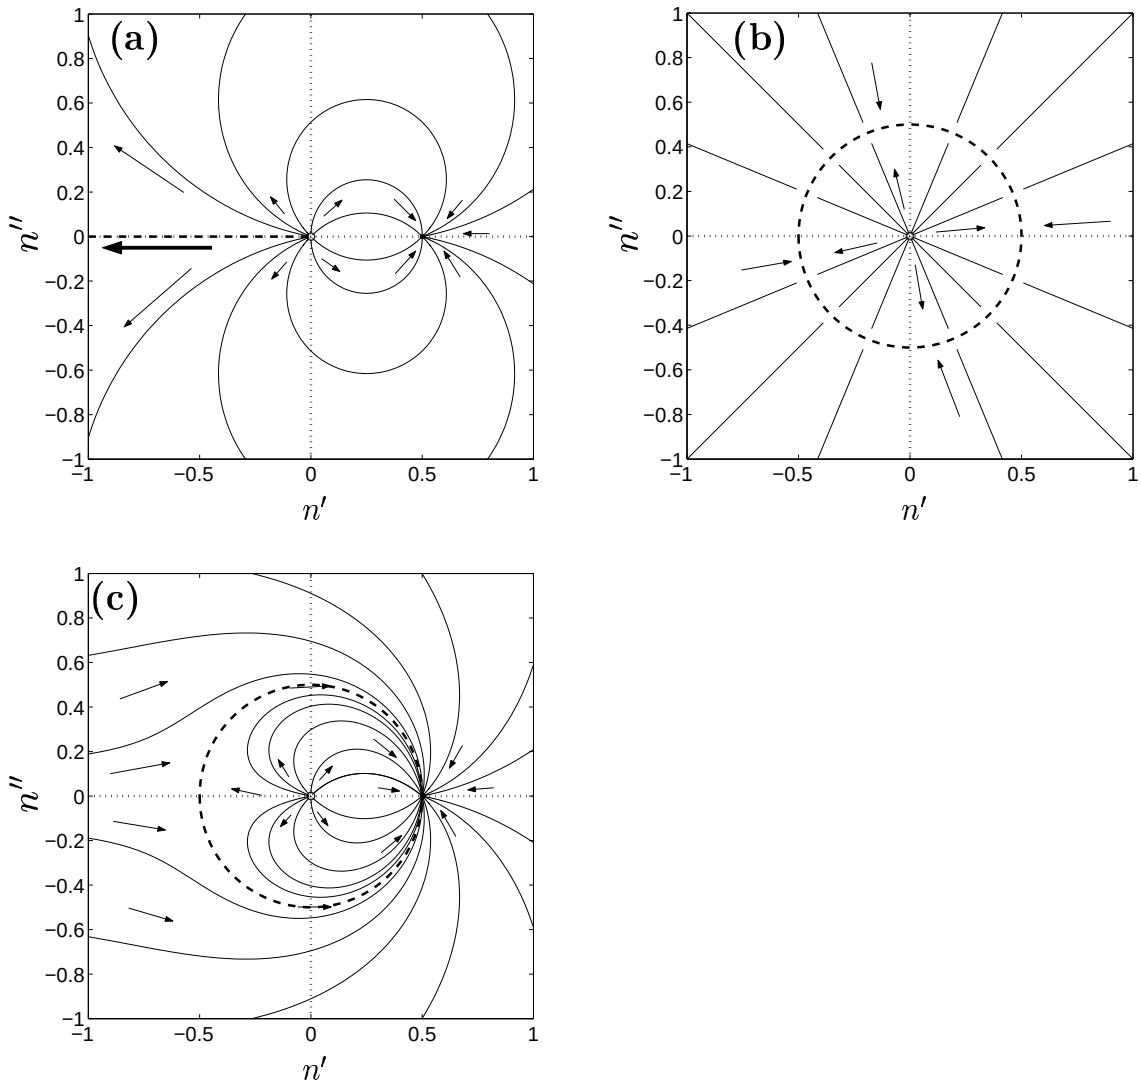

Figure 9.1: **Deterministic phase space** for Stratonovich form of the $d\check{n}$ equation. The case of $n_\infty = 0.5$ is shown. **(a)**: positive-P-like simulation $\mathcal{G} = 0$; **(b)**: radial gauge (9.17); **(c)**: minimal gauge (9.16). The moving singularity in **(a)** is shown with large arrow, the attractors at $|\check{n}| = n_\infty$ with a thick dashed line.

How do the modified equations measure up to the gauge choice criteria of Section 6.3.2? In order:

1. The moving singularity in $d\check{n}$ has been removed (at high $|\check{n}|$ the deterministic behavior is restorative towards the origin), and no other moving singularities are present. Note that if $\check{n}$ remains finite, then so does Ω, which just experiences exponential growth with an exponent that is always finite, since it depends only on $\check{n}$.

2. No new moving singularities have been introduced.

3. No noise divergences (new or old) are present since all noise terms satisfy (6.3).

4. No discontinuities in the drift or diffusion coefficients are present provided the time-dependence of $\mu_e(\tau)$ is not singular.

5. Is the weight spread minimized?

   (a) Gauge is zero when $\breve{n}$ is positive real, and small in its neighborhood, where gauge corrections are unnecessary.

   (b) Explicit variational minimization of $\mathcal{G}$ has not been carried out, but rather criteria 6 have been applied.

   (c) Gauge is zero at deterministic attractor $\breve{n} = n_\infty$.

6. No unwanted gauge behaviors are present:

   (a) Gauge is nonzero in a large part of phase space, so most trajectories contribute to removal of bias.

   (b) Gauge changes smoothly in phase space.

   (c) Gauge is autonomous.

   (d) Gauge breaks the analyticity of the equations.

9.2.3 Radial drift gauge

The gauge (9.16) looks good, but it was found that a more severe phase-space modification usually gives smaller statistical uncertainties. This better gauge form is

$$\mathcal{G} = i\sqrt{2\hbar\chi}\,(\breve{n} - |\breve{n}|)\,, \tag{9.17}$$

which will be called here the "radial" gauge, due to elimination of any deterministic polar-angle evolution as shown in Figure 9.1**(b)**. The deterministic attractor is the entire $|\breve{n}| = n_\infty$ circle (or the vacuum if $n_\infty \leq 0$), with some phase diffusion in the phase of $\breve{n}$.

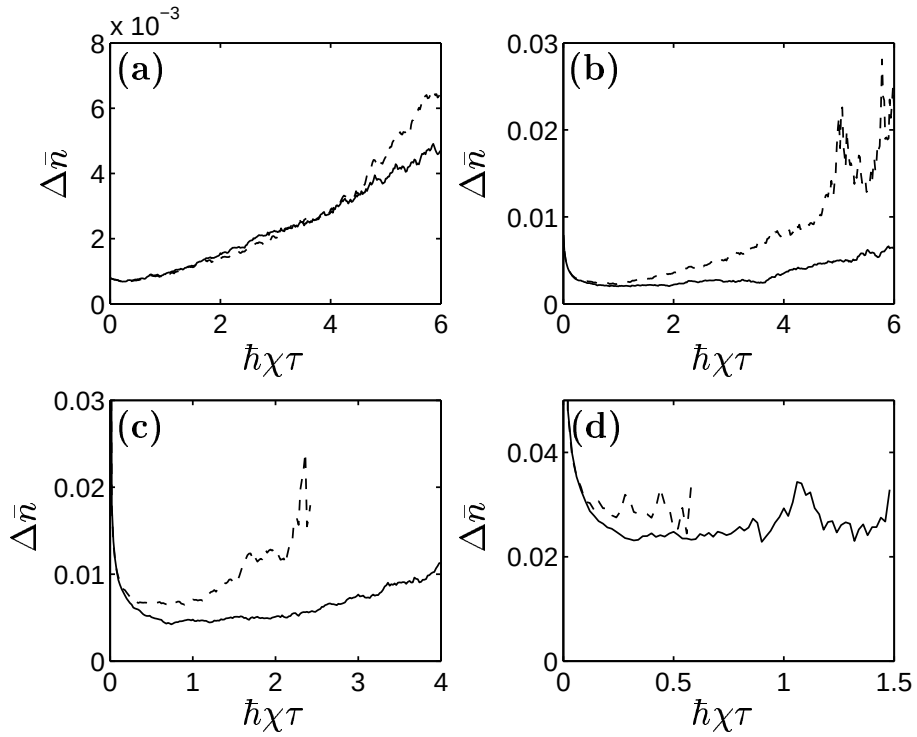

Figure 9.2: **Uncertainty in** $\overline{n}$, the estimate of $\langle \widehat{n} \rangle$ for the $\mu_e = 0$ system. **(a)**: $\overline{n}_0 = 0.1$; **(b)**: $\overline{n}_0 = 1$; **(c)**: $\overline{n}_0 = 10$; **(d)**: $\overline{n}_0 = 100$; SOLID LINE: radial gauge (9.17); DASHED LINE: gauge (9.16). Simulations were carried out with $\mathcal{S} = 2 \times 10^4$ trajectories.

The efficiency of the two gauges is compared in Figure 9.2 for $\mu_e = 0$ and a variety of initial $\overline{n}_0$. It is not immediately clear why the radial gauge is better, but the improvement is significant whenever $\overline{n}_0 \geq \mathcal{O}(1)$. Similar behavior was seen for other μ_e values. In particular also for $\mu_e = \hbar\chi$, which corresponds to a different low temperature state (i.e. a pure Fock number state — see Figure 9.3).

Comparison to the heuristic gauge choice criteria of Section 6.3.2 follows through identically as for the previous gauge (9.16), apart from the gauge not being zero on the entire $|\breve{n}| = n_\infty$ attractor but only at $\breve{n} = n_\infty$.

9.3 Numerical simulations

The results of some example simulations using the radial gauge (9.17) are shown in Figures 9.3 and 9.4. These are compared to exact results obtained from (9.10) and also to mean field semiclassical calculations. As outlined in Section 5.4, mean

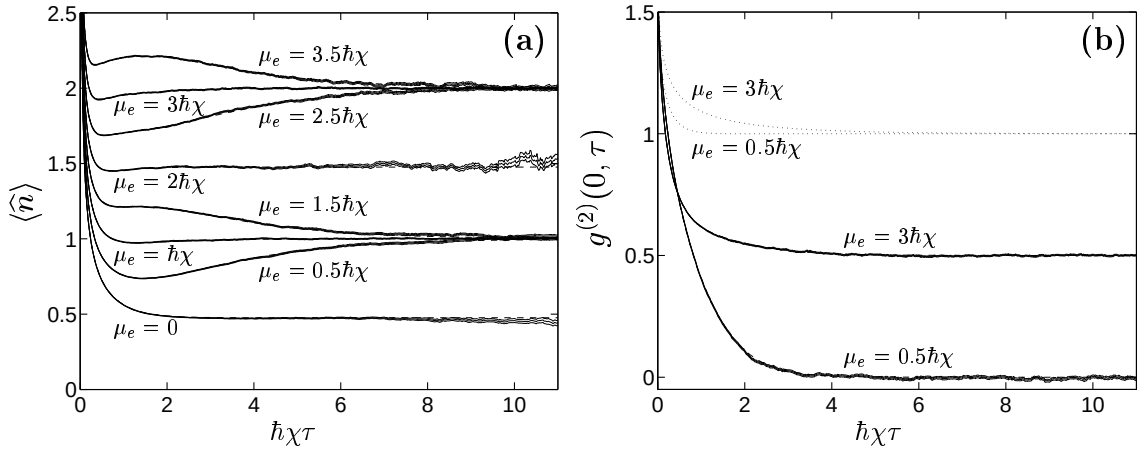

Figure 9.3: **Temperature dependent observables** given constant μ_e and $\overline{n}_0 = 10$ at high temperature. SOLID LINES: simulation with $\mathcal{S} = 2 \times 10^5$ trajectories. Triple lines indicate error bars (often not resolved at this scale). DASHED LINES: exact values (mostly obscured by simulation results) DOTTED LINES: mean field calculation (subplot **(b)**). The second order correlation function $g^{(2)}(0,\tau)$ is given by (8.8).

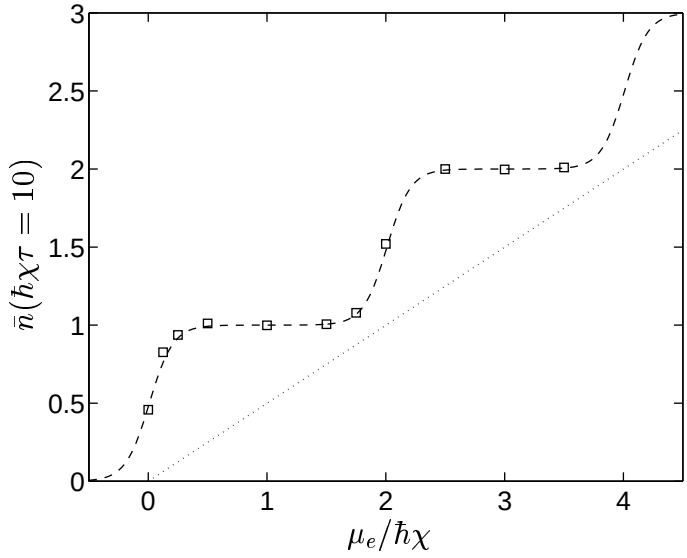

Figure 9.4: **Mean mode occupations at low temperature** $T = \hbar\chi/10k_B$ as a function of constant μ_e. DASHED: exact result — note the quantization, DOTTED: mean field result, SQUARES: simulation results with radial gauge (9.17). Uncertainty was of symbol size or smaller.

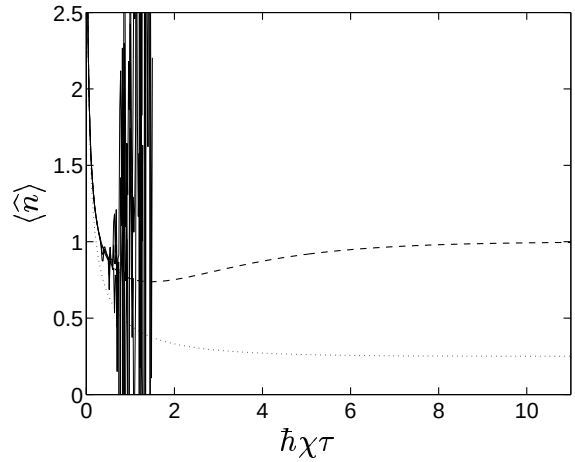

Figure 9.5: Performance of the mean field and un-gauged simulations for comparison with Figure 9.3**(a)**. **Mean mode occupation** with $\mu_e = 0.5\hbar\chi$ and $\overline{n}_0 = 10$. DASHED: exact result, DOTTED: mean field result, SOLID: un-gauged simulation.

field calculations such as solution of Gross-Pitaevskii equations are equivalent to simulation of only the deterministic part of the Ito gauge P equations.

Additionally, a positive-P-like simulation with no gauge ($\mathcal{G} = 0$) is shown in Figure 9.5 for comparison with the gauged technique.

A subtlety to keep in mind in these thermodynamics simulations is that (in contrast to dynamics) the normalization $\langle \mathrm{Re}\{\Omega\}\rangle$, which appears in the denominator of observable estimates (3.14), is not unity in the $\mathcal{S} \to \infty$ limit. This then requires a sufficient number of trajectories per subensemble so that the denominator is always positive and not too close to zero for any of these, as discussed in detail in Appendix C.

Features seen in these figures include:

- Convergence of the gauged simulation is excellent, and precisely reproduces the exact quantum behavior.
- At low temperatures $\tau \gtrsim 1/\hbar\chi$ the semiclassical approximation gives completely wrong results for $g^{(2)}$ and is out by $\mathcal{O}(1)$ in $\langle \widehat{n} \rangle$, which is also completely wrong if mode occupation is $\leq \mathcal{O}(1)$.
- The gauge P method, on the other hand, reproduces the exact quantum behavior precisely despite using a semiclassical coherent-state basis.

- The un-gauged simulation breaks down while still at high temperatures due to boundary term induced spiking.

- Low temperature observable estimates are consistent with the Fock state low temperature ground state described in Section 9.1.3. ($\langle\widehat{n}\rangle$ approaches integer values[2] and the two-particle correlation $g^{(2)}$ approaches $1 - 1/\langle\widehat{n}\rangle$.)

- Mid-temperature behavior is also simulated precisely.

- At low temperature (long times) uncertainty in the observable estimates is greatest for μ_e even integer multiples of $\hbar\chi$ (and hence even integer multiple chemical potential μ at low temperatures), and lowest for μ_e odd integer multiples of $\hbar\chi$. These cases correspond to the low temperature state being a mixture of two Fock number states, or just a single pure Fock state, respectively.

9.4 Chemical potential as a free gauge parameter

Suppose one is interested in obtaining properties at a given temperature T and chemical potential μ. This determines the final simulation target time τ_T and target chemical potential $\mu(\tau_T)$, but the chemical potential at intermediate times $\tau < \tau_T$ is not specified. The only conditions are 1) that the quantity $\mu(\tau)\tau$ be (piecewise) differentiable so that $\mu_e(\tau)$ can be calculated, and 2) that $\lambda_n = -\lim_{\tau\to 0}(\mu\tau)$ be positive so that the initial distribution (9.7) is a normalizable probability. Apart from these, $\mu(\tau)$ is formally an arbitrary function (and λ_n an arbitrary positive number).

This is all in the limit of many trajectories, however. The form of $\mu(\tau)$ at intermediate τ has no effect on observable estimate means in the limit $\mathcal{S} \to \infty$, but can have a strong effect on the broadness of the trajectory distribution. And, hence — on the precision of a finite sample estimate. In this sense then, $\mu_e(\tau)$ (given by (2.30)), acts as effectively an additional gauge function.

[2]Apart from when μ_e is an even integer multiple of $\hbar\chi$, since then $\langle\widehat{n}\rangle$ approaches half-integer values.

An indication of the characteristics of an efficient form of $\mu_e(\tau)$ can be gained by comparing calculated observable uncertainties for several basic cases. Let us look at the situation when $\tau_T = 1/\hbar\chi$, and $\mu(\tau_T) = \hbar\chi$. (This gives mean occupation $\overline{n}(\tau_T) = 0.5628$).

- **Case 1:** Variation in $\overline{n}_0(\lambda_n)$, i.e. in the high-temperature starting mode occupation, while $\mu_e(\tau)$ is a constant chosen appropriately using (9.2). Calculated uncertainties in $\langle\widehat{n}\rangle$ and $g^{(2)}(0,\tau)$ are compared in Figure 9.6. It is seen that initial occupation $\mathcal{O}(1) - \mathcal{O}(10)$ gives the best performance, while too large or too small initial occupation lead to excessive uncertainty in the observables.

- **Case 2a:** Time-varying μ_e, while $\overline{n}_0 = 1$ is held constant. μ_e is chosen to be a nonzero constant up to a time τ_c, and zero for times $\tau > \tau_c \leq \tau_T$. Note that as $\tau_c \to 0$, the size of μ_e grows to make up for the shorter time over which it acts. Uncertainties are compared in Figure 9.7**(a)**. It is seen that if μ_e is varied too strongly, an inefficient simulation results, although some relatively small variation may be beneficial for the $g^{(2)}$ calculation.

- **Case 2b:** Again time-varying μ_e, while $\overline{n}_0 = 1$ is held constant, but this time μ_e is held zero at high temperatures (low times $\tau < \tau_c$), and makes up for this by becoming a nonzero constant for $\tau_c < \tau < \tau_T$. As $\tau_c \to \tau_T$, the size of the required μ_e grows. Uncertainties are compared in Figure 9.7**(b)**. Again, it is seen that strong variation in μ_e leads to an inefficient simulation.

From the simulation data it appears that, other things being equal,

1. $\overline{n}_0 \approx \mathcal{O}(\overline{n}(\tau_T))$ or slightly greater appears to give the best performance

2. A constant (or nearly so) $\mu_e(\tau)$, which keeps the phase-space behavior time independent appears to give better performance than simulations for which μ_e strongly varies with τ.

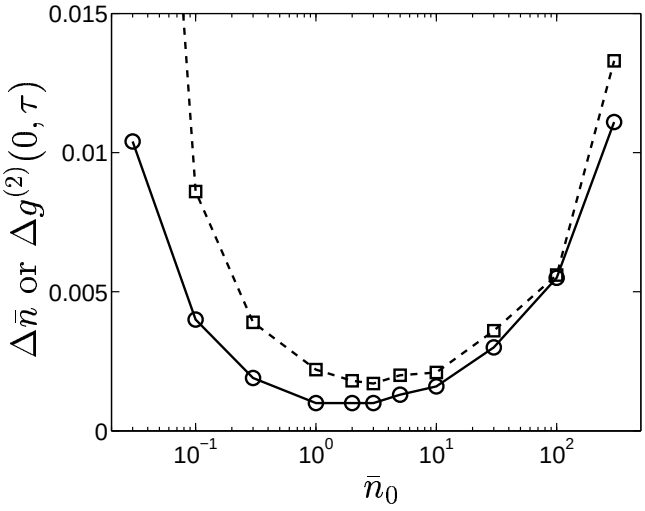

Figure 9.6: **Dependence of observable estimate uncertainties on the starting occupation** $\overline{n}_0$. Shown are calculated uncertainties in CIRCLES: $\overline{n} = \langle \widehat{n} \rangle$ SQUARES: $g^{(2)}(0,\tau)$, for $\hbar\chi\tau = 1$ and $\mu = \hbar\chi$. Calculations were made with the radial gauge (9.17) varying the starting occupation $\overline{n}_0$ and choosing a constant μ_e according to (9.2). Each calculation was with 2×10^4 trajectories.

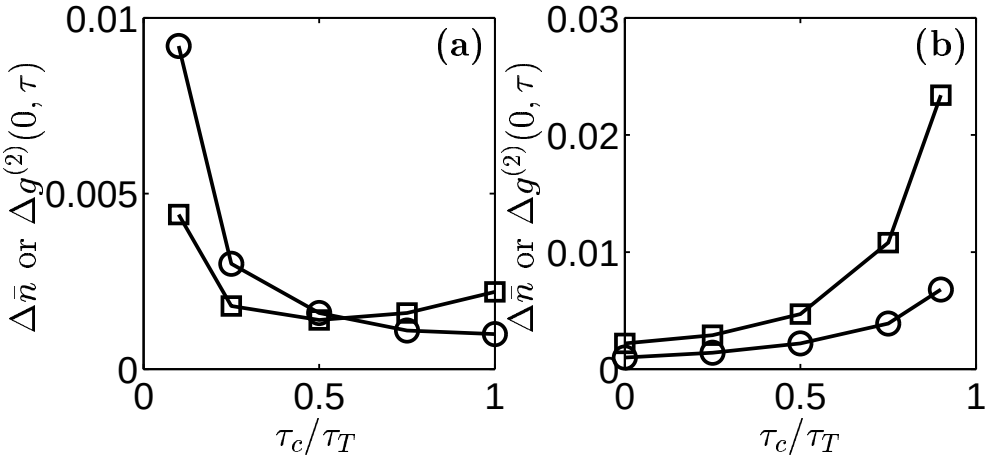

Figure 9.7: **Dependence of observable estimate uncertainties on the form of** μ_e. Shown are calculated uncertainties in CIRCLES: $\overline{n} = \langle \widehat{n} \rangle$ SQUARES: $g^{(2)}(0,\tau)$, for $\hbar\chi\tau = 1$ and $\mu = \hbar\chi$. Calculations were made with the radial gauge (9.17), starting with $\overline{n}_0 = 1$ initially, but varying the form of μ_e. In **(a)**, this was taken to be a nonzero constant for $\tau \leq \tau_c$ and zero for $\tau > \tau_c$. In **(b)**, μ_e was zero for $\tau < \tau_c$, and nonzero constant for $\tau \geq \tau_c$. Each calculation was with 2×10^4 trajectories.

9.5 Summary

The drift gauges (9.16) and (9.17) have been developed to overcome shortcomings of the un-gauged positive-P-like technique. Using these gauges it has been shown that a stochastic simulation can precisely simulate the full quantum grand canonical thermodynamics of a two-body interacting Bose gas mode or orbital. Features include

1. Simulation with good precision down to the lowest temperatures $\hbar\chi\tau \gg 1$ where discretization of mode occupation takes place (See e.g. Figures 9.3 and 9.4).

2. Simulation of a range of temperatures down to a minimum $1/k_B\tau_T$ in one run.

3. Calculation of any desired set of observables from one run.

The uncontrollable spiking (which was present in an un-gauged simulation) is removed.

Numerical simulations indicate that the radial gauge (9.17) gives superior efficiency for a wide range of parameters. Also, it has been found that if one is aiming for a particular target temperature and chemical potential, it appears advantageous to choose the intermediate-time chemical potentials in the simulation such that: 1) μ_e is slowly varying with τ (or constant), and 2) The starting occupation $\overline{n}_0(\lambda_n)$ is of the same order as the mean occupation at the target time.

The gauges developed should carry over in a straightforward manner to multi-mode simulations when the interparticle scattering is local to the lattice sites as described by the Hamiltonian (2.17). This is implemented in Chapter 11.