




**Universidade do Minho**

Escola de Engenharia

Alexandra Francisco Ramôa da Costa Alves

**Learning the physics of open quantum systems from experiments**





**Universidade do Minho**

Escola de Engenharia

Alexandra Francisco Ramôa da Costa Alves

**Learning the physics of open quantum systems from experiments**

Master Thesis  
Engineering Physics

Work developed under the supervision of:

**Ernesto Galvão**  
**Luís Soares Barbosa**


## **COPYRIGHT AND TERMS OF USE OF THIS WORK BY A THIRD PARTY**

This is academic work that can be used by third parties as long as internationally accepted rules and good practices regarding copyright and related rights are respected.

Accordingly, this work may be used under the license provided below.

If the user needs permission to make use of the work under conditions not provided for in the indicated licensing, they should contact the author through the RepositoriUM of Universidade do Minho.

### ***License granted to the users of this work***



**Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International  
CC BY-NC-SA 4.0**

<https://creativecommons.org/licenses/by-nc-sa/4.0/deed.en>



# Acknowledgements

First of all I would like to express my gratitude to the people who made this learning experience possible. In particular, to Raffaele Santagati, for the opportunity to undertake this project and for his guidance throughout the year; to Ernesto Galvão, for his invaluable help; to Nathan Wiebe, for his insightful comments; and to professor Luís Barbosa, for being a great teacher.

I would also like to thank my family and friends, who supported me in other ways - especially my sister Mafalda, for being my biggest companion throughout this journey.

Finally, I acknowledge the use of the use of IBM Quantum services [61], Google Colab [47], and the NOVAthesis L<sup>A</sup>T<sub>E</sub>X template [75].

### **STATEMENT OF INTEGRITY**

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the Universidade do Minho.

*"[...] all knowledge degenerates into probability ; and this probability is greater or less, according to our experience of the veracity or deceitfulness of our understanding, and according to the simplicity or intricacy of the question."*

*– David Hume*

# Abstract

## Learning the physics of open quantum systems from experiments

The ability to efficiently determine the dynamics to which a quantum device conforms is vital for its reliable operation. Thus, as quantum machines evolve, the means for their characterization must evolve alongside them, especially as they reach the limits of classical tractability. Bayesian inference has been proposed as a solution, as it offers a flexible way of using experimental data to learn the dynamical parameters - or even models - governing the evolution of quantum systems. It gives rise to noise-resilient protocols, which are capable of quantifying their own uncertainty, of learning from scarce information, and of real-time estimation. Importantly, and apart from the obvious applications in sensing devices, online processing enables adaptivity, a stepping stone for achieving fundamental limits of metrology via what is called *quantum-enhanced* estimation. Like so, the exploitation of quantum control as a resource within the Bayesian paradigm opens the door to uncertainties scaling at the Heisenberg limit.

Most work on the subject has relied on simple methods for portraying the Bayesian posterior, but they quickly become a bottleneck in realistic scenarios. The difficulty of the task grows with the complexity of the devices to be verified, which adds to the fact that they are ultimately open systems and as such undergo uncontrollable interactions with their surroundings. While capturing unwanted processes like decoherence is by itself a valuable endeavor, this poses extra challenges, and requires careful computational treatment. In general, when scaling up, the biggest difficulty of Bayesian learning is numerical integration. In this context, pairing it with advanced Monte Carlo methods makes for remarkably robust algorithms, which can succeed under complex features and control a multitude of trade-offs.

This dissertation aims to overview both of these methodologies, and to apply them to the characterization of noisy quantum computers. Special emphasis is placed on state-of-the-art methods for statistical simulation - namely Hamiltonian and sequential Monte Carlo, variants thereof, and subsampling approaches. Using these strategies to post-process 150 single-shot measurements on a Ramsey sequence, we learn an oscillation frequency and coherence time to 0.4% and 33% uncertainty respectively, compared to 15% and 333% using curve fits. With an echoed pulse set-up and 75 shots only, we then achieve an uncertainty of 0.3% for the qubit's precession frequency, whereas standard fitting methods either do worse or fail completely. Switching to online processing and further lowering the total shot count to 15, we use adaptive experimental design to infer the frequency to 5% uncertainty, compared to 44% using offline processing under otherwise identical conditions.

**Keywords:** Bayesian inference, Monte Carlo integration, quantum metrology

## **Caraterização de sistemas quânticos abertos através de dados experimentais**

A capacidade de determinar de forma eficiente a dinâmica que rege um dispositivo quântico é fulcral para a viabilidade da sua operação. Logo, à medida que as máquinas quânticas evoluem, os meios para a sua caraterização devem acompanhá-las, sobretudo quando se atingem os limites da tratabilidade clássica. A inferência Bayesiana foi proposta como solução, já que oferece uma forma flexível de usar dados experimentais para estimar os parâmetros - ou até modelos - que governam a evolução de sistemas quânticos. Esta técnica produz protocolos resilientes a ruído, capazes de quantificar a sua própria incerteza, de aprender com base em informação escassa, e de estimar em tempo real. Crucialmente, além das óbvias aplicações em sensores, o processamento em tempo real possibilita adaptatividade, um ponto de partida para alcançar limites fundamentais da metrologia no regime quântico. Assim, o uso de controlo quântico dentro do paradigma Bayesiano proporciona a progressão de incertezas no limite de Heisenberg.

A literatura tem recorrido maioritariamente a métodos simples para retratar distribuições Bayesianas, mas estes são limitadores em cenários realistas. A dificuldade da tarefa aumenta com a complexidade dos dispositivos a verificar, o que é agravado pelo facto de estes serem sistemas abertos, que interagem de forma incontrollável com o ambiente. Apesar do interesse inerente à captura de fenómenos indesejáveis como a decoerência, estes criam desafios adicionais, e exige um tratatamento computacional cuidado. Em geral, quando a complexidade aumenta, a maior dificuldade da aprendizagem Bayesiana é a integração numérica. Neste contexto, combiná-la com métodos de Monte Carlo avançados produz algoritmos de uma robustez notável, capazes de lidar com atributos complexos e de gerir vários compromissos.

Esta dissertação pretende abordar estas duas metodologias, e aplicá-las à caraterização de computadores quânticos ruidosos. Colocar-se-á especial ênfase no estado da arte da simulação estatística - nomeadamente Monte Carlo Hamiltoniano e sequencial, variantes destes, e técnicas de subamostragem. Usando estas estratégias para pós-processar 150 medições (isoladas) numa sequência de Ramsey, inferimos uma frequência de oscilação e tempo de coerência com incertezas de 0.4% e 33% respetivamente, comparando com 15% e 333% com ajuste de curvas. Acrescentando um pulso de inversão, obtemos com apenas 75 medições uma incerteza de 0.3% para a frequência; métodos padrão deram resultados piores ou nulos. Com processamento em tempo real e reduzindo as medições para 15, usamos medições adaptativas para inferir a frequência com uma incerteza de 5%, em contraste com 44% usando pós-processamento e mantendo as restantes condições.

**Palavras-chave:** inferência Bayesiana, integração de Monte Carlo, metrologia quântica

# Contents

|   |              |
|---|--------------|
| <b>List of Figures</b>  | <b>xii</b>   |
| <b>Acronyms</b>   | <b>xviii</b> |
| <b>1 Introduction</b>   | <b>1</b>     |
| 1.1 Characterization of quantum systems . . . . .                                       | 2            |
| 1.2 Research questions, literature overview and objectives . . . . .                    | 2            |
| 1.3 Related topics . . . . .  | 5            |
| 1.4 Document structure . . . . .  | 6            |
| <b>2 Quantum Parameter Estimation</b>   | <b>8</b>     |
| 2.1 Bayes' rule . . . . .   | 8            |
| 2.2 Bayesian inference . . . . .  | 14           |
| 2.2.1 Estimators and expressions of uncertainty . . . . .                               | 17           |
| 2.2.2 Estimation in practice: approximate inference and conjugate priors . . . . .      | 24           |
| 2.3 Application to quantum systems and general likelihood models . . . . .              | 32           |
| 2.3.1 Brief detour into model selection . . . . .                                       | 35           |
| 2.3.2 Likelihood evaluations in the quantum case . . . . .                              | 39           |
| 2.3.3 Quantum characterization examples: spin precession and phase estimation . . . . . | 42           |
| 2.4 Bayesian experimental design . . . . .  | 50           |
| 2.4.1 Adaptive measurements for a greedy strategy . . . . .                             | 54           |
| 2.4.2 Beating shot noise: quantum-enhanced, Heisenberg-limited estimation . . . . .     | 57           |
| 2.4.3 Empirical strategies for a precessing qubit . . . . .                             | 59           |
| <b>3 Monte Carlo methods for posterior sampling</b>                                     | <b>66</b>    |
| 3.1 Monte Carlo and importance sampling . . . . .                                       | 67           |
| 3.2 Sequential Monte Carlo . . . . .  | 71           |
| 3.2.1 Particle propagation and the Liu-West filter . . . . .                            | 76           |
| 3.2.2 Sequential importance resampling . . . . .  | 81           |
| 3.2.3 Tempered likelihood estimation . . . . .  | 85           |
| 3.2.4 General SMC schemes and estimation of the model evidence . . . . .                | 88           |
| 3.3 Markov Chain Monte Carlo and the Metropolis-Hastings algorithm . . . . .            | 91           |

|          |  |            |
|----------|--|------------|
| 3.3.1    | Random walk Metropolis . . . . .   | 95         |
| 3.3.2    | Hamiltonian Monte Carlo . . . . .  | 96         |
| 3.3.3    | Markov kernels within SMC . . . . .  | 103        |
| 3.4      | Implementations of Hamiltonian Monte Carlo . . . . .                       | 105        |
| 3.4.1    | Conditions for applicability . . . . .                                     | 106        |
| 3.4.2    | Sampling a trajectory and state . . . . .                                  | 107        |
| 3.4.3    | Dynamic implementations: the no U-turn sampler . . . . .                   | 109        |
| 3.5      | Subsampling strategies . . . . .   | 112        |
| 3.5.1    | Unbiased estimators and variance reduction with control variates . . . . . | 112        |
| 3.5.2    | Stochastic gradient Hamiltonian Monte Carlo . . . . .                      | 115        |
| 3.5.3    | Hamiltonian Monte Carlo with energy conserving subsampling . . . . .       | 116        |
| <b>4</b> | <b>Applied sampling and inference examples</b>                             | <b>118</b> |
| 4.1      | Sampling problems . . . . .  | 118        |
| 4.1.1    | 6-dimensional normal distribution . . . . .                                | 119        |
| 4.1.2    | Rosenbrock function . . . . .  | 120        |
| 4.1.3    | Smiley kernel density estimate . . . . .                                   | 122        |
| 4.2      | Inference problems . . . . .   | 123        |
| 4.2.1    | Phase estimation . . . . .   | 123        |
| 4.2.2    | Precession and bimodality . . . . .  | 125        |
| 4.2.3    | Multivariate function (2 and 4-dimensional estimation) . . . . .           | 127        |
| <b>5</b> | <b>Experiments on quantum hardware</b>                                     | <b>131</b> |
| 5.1      | Models for open quantum systems . . . . .                                  | 131        |
| 5.1.1    | Inside the Bloch sphere: visualizing mixed states . . . . .                | 131        |
| 5.1.2    | The Kraus operator sum representation . . . . .                            | 133        |
| 5.1.3    | Energy relaxation: the amplitude damping channel . . . . .                 | 135        |
| 5.1.4    | Dephasing: the phase damping channel . . . . .                             | 138        |
| 5.1.5    | Spin precession . . . . .  | 140        |
| 5.2      | Characterization of quantum devices . . . . .                              | 144        |
| 5.2.1    | Experimental setup . . . . .   | 145        |
| 5.2.2    | Hahn echo experiment (T2 estimation) . . . . .                             | 147        |
| 5.2.3    | Quantifying energy loss (T1 estimation) . . . . .                          | 150        |
| 5.2.4    | Ramsey experiment (2-dimensional estimation) . . . . .                     | 150        |
| 5.2.5    | Echoed Ramsey experiment (1-dimensional estimation) . . . . .              | 153        |
| <b>6</b> | <b>Discussion and conclusions</b>  | <b>158</b> |
|          | <b>Bibliography</b>  | <b>160</b> |

# List of Figures

|    |   |    |
|----|---|----|
| 1  | Diagram of the topics to be covered in this dissertation. Dashed boxes denote specific algorithms. . . . .  | 7  |
| 2  | Bayesian posterior probability density for a binomial likelihood. The dashed and full vertical lines mark the real parameter used to generate the data and its frequentist estimate respectively. The same data vector contemplating 100 trials was used for the Bayesian and frequentist calculations. . . . .           | 16 |
| 3  | Bayesian posterior probability density for a binomial likelihood. The dashed and full vertical lines mark the real parameter used to generate the data and its frequentist estimate respectively. The same data vector contemplating 1000 trials was used for the Bayesian and frequentist calculations. . . . .          | 21 |
| 4  | Bayesian posterior probability evaluations for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector. . . . .  | 25 |
| 5  | Bayesian posterior probability distribution for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector. . . . .   | 26 |
| 6  | Bayesian posterior probability mass function for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector. . . . .  | 28 |
| 7  | Bayesian conjugate posterior density for a binomial likelihood and a flat prior. The dashed line marks the real parameter used to generate the data, and the 'x's mark single point evaluations of the posterior. 100 trials were included in the data vector. . . . .  | 31 |
| 8  | Illustration of the model $M_1$ from equation 2.67 for different values of $W$ . The Bayes factor of $M_1^{(ii)}$ relative to $M_1^{(i)}$ is $\mathcal{B} = 1/2$ for $D = 1$ . . . . .  | 38 |
| 9  | Illustration of the model $M_2$ from equation 2.70 as compared to its one-dimensional projections. Two such projections can be viewed on a scale in figure 8 for different values of $W$ . The space forms a square with side $W$ . . . . .   | 38 |
| 10 | Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs. The tilde $\sim$ means "is distributed as" and the dot $\cdot$ is a stand-in for the object of the sampling. . . . . | 40 |
| 11 | Bloch sphere representations of spin precession for a specific initial state and for a general one. . . . .   | 44 |



|    |   |    |
|----|---|----|
| 12 | Bayesian posterior probability distribution for an oscillatory likelihood. The dashed vertical line mark the real parameter used to generate the data. 100 trials were contemplated in the data vector, using evenly spaced times on $[0, 5]$ . . . . .   | 45 |
| 13 | Iterative phase estimation circuit diagram. . . . .   | 46 |
| 14 | Impact of the magnitude of the parameter's coefficient(s) in the shape of the likelihood for a simple binomial model (14a) and a damped oscillator model (14b). . . . .   | 51 |
| 15 | Binary tree illustrating the computation of the expected utility for a set of experiments $\vec{E} = \{E_1, E_2, E_3\}$ . To each layer of branches corresponds one of the individual experiments, in chronological order from top to bottom. Each branch represents the occurrence of a specific datum, and is tagged with its likelihood. The leaves are marked with the utility of the data sequence traced out by the path of branches from the root to themselves. Using these calculations, the expected utility of $\vec{E}$ can be derived from 2.109. . . . .  | 53 |
| 16 | Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference, based on iterative updates. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs. . . . .  | 54 |
| 17 | Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference with local utility optimization. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs. . . . .  | 55 |
| 18 | Binary trees illustrating the computation of the locally expected utility for two followed experiments, $E_1$ and $E_2$ . To each of the height-1 trees corresponds one experiment, in chronological order from left to right. Each branch represents the occurrence of a specific datum, and is tagged with its likelihood. The leaves are marked with the utility of the respective branch's datum. Using these calculations, the expected utilities of $E_1$ and $E_2$ can be derived from 2.109. The equality annotated bottom right evidences that the latter calculation is dependent on $E_1$ and the effective $D_1$ . . . . .            | 56 |
| 19 | Diagrammatic representation of a local optimization routine in the context of Bayesian experimental design. . . . .   | 56 |
| 20 | Posterior distributions for a time-dependent binomial model, for 2 sets of measurement times. The likelihoods take the undamped squared-sinusoidal form of 2.84. Both posteriors were evaluated on the same grid comprising 10 points (the 'x' markers), and are juxtaposed with smooth curves of matching color representing their respective continuous posterior distributions. The latter were approximated using 1000 points. The dashed line marks the real parameter used to generate the data. 100 measurements were included in each data vector, and the times were picked in constant increments up to the maximum $t_{max}$ . . . . . | 64 |

|    |  |     |
|----|--|-----|
| 21 | Posterior distributions for a time-dependent binomial model, for 2 sets of measurement times. The likelihoods take the undamped squared-sinusoidal form of 2.84. The dashed line marks the real parameter used to generate the data. 10 measurements were included in each data vector, and the times were picked in constant increments up to the maximum $t_{max}$ . . . . .   | 65  |
| 22 | Evolution of the posterior probability distribution with the iteration number. The dashed vertical line marks the real parameter used to generate the data. The posterior is depicted for $k = 5, 10,$ and $25$ accumulated data, according to equation 3.9. The $D_k$ were obtained using exponentially increasing times $t_k = (9/8)^k$ . . . . .  | 72  |
| 23 | Illustration of the particle degeneracy/weight collapse problem, prompted by the particle cloud's evolution as the iterations (and with them the data) accumulate. The distributions are the same as in figure 22. The 'x' markers stand for particles, their sizes representing their weights. . . . .  | 74  |
| 24 | Illustration of the obtention of a weighted particle set and the necessary treatment to generate an efficient unweighted sample. This corresponds a segment of 23, with a change of the sampling techniques. The bottom clipping displays the following posterior, to be targeted by the darkest particle set. The 'x' markers stand for particles, their sizes representing their weights. . . . .                            | 77  |
| 25 | Illustration of the process behind a HMC proposal (for a gaussian distribution). . . . .   | 98  |
| 26 | Posterior probability (crimson) and log-probability (blue) for two different likelihood models and flat priors. The dashed vertical lines mark the real parameter used to generate the data $D$ . . . . .  | 106 |
| 27 | Effect of noise on the phase space trajectory for a harmonic oscillator model ( $U = \theta^2/2$ ). The trajectory with correct dynamics is plotted in black, that with a noisy gradient in green, and that with a noisy gradient compensated by a friction term in a gradation from blue to red (by ascending order of elapsed time). This example is based on an identical one from [20].                                    | 115 |
| 28 | Trajectories generated using RWM (top) and HMC (bottom) to sample from a multivariate normal distribution with 6 dimensions/parameters. Both chains started at the origin and were evolved for $10^3$ states. These states are plotted as points gradating from blue to red (from earlier to later iterations), whereas the starting point is marked with an 'x'. The dimensions are plotted pairwise for convenience. . . . . | 119 |
| 29 | 3-dimensional plot of the Rosenbrock function from equation 4.1. . . . .   | 120 |
| 30 | Trajectories generated using RWM (30a) and HMC (30b) to sample from a (bivariate) Rosenbrock function. Both chains started at the origin and were evolved for $10^4$ states. These states are plotted as points gradating from blue to red (from earlier to later iterations). . . . .   | 121 |

|    |  |     |
|----|--|-----|
| 31 | Trajectories generated using different variants of HMC to sample from a (bivariate) Rosenbrock function. They differ (both between themselves and from the standard case of 30b) in how the trajectories are generated and their samples picked. The sample selection methods are: uniform sampling (31a), biased progressive sampling (31b) and NUTS with partly biased progressive sampling (31c). Figures 31a and 31b used static implementations with $L = 2^6$ , whereas 31c's dynamic termination produced on average $L = 2^{5.4}$ . All chains started at the origin and were evolved for $10^4$ states. These states are plotted as points gradating from blue to red (from earlier to later iterations). | 122 |
| 32 | Target bivariate smiley function (32a), generated using 768 gaussian kernels, and 3-d scatter of the particle cloud produced by SHMC (32b) plotted over a surface plot of the target.  | 122 |
| 33 | Trajectories generated using different SMC/MCMC sampling methods to sample from a (bivariate) smiley KDE. Results are shown for 4 methods: parallel RWM chains (33a), SMC with RWM moves (33b), parallel HMC chains (33c), and SHMC (33d).   | 124 |
| 34 | Evolution of the sampling process through the iterations for MCMC with RWM and HMC move steps (34a) and for GRF (34b). Both consider 100 experiments total, but their design and use are different. The horizontal coordinate is the phase (angle) value. The points are positioned vertically and upward according to the iteration number, e.g. the topmost point corresponds to the final result. The likelihoods based on the full datasets are plotted in blue in a second axis for visualization.  | 124 |
| 35 | Evolution of the sampling process through the iterations for for MCMC with RWM and HMC move steps (35a) and SMC with LWF (35b). Both consider 100 experiments total, and the same set of pre-determined controls. The points are positioned vertically and upward according to the iteration number, e.g. the topmost point is the final result; the horizontal coordinate is the phase (angle) value. The likelihood is plotted in blue on a second axis for visualization.   | 126 |
| 36 | Cumulative density functions obtained via MCMC with RWM and HMC move steps (36a) and SMC with LWF (36b), for a prior with support over negative frequencies. The modes are marked with dashed vertical red lines.  | 126 |
| 37 | Kernel density estimates of the posterior resulting from full data SHMC and from subsampling 50/400 observations (SHMC-ECS).   | 127 |
| 38 | Results of 2-dimensional inference using SIR with RWM propagation for $15^2$ particles (38a), a grid for $15^2$ particles (38b), and a grid for $50^2$ particles (38c). The modes are marked with red 'x' markers.   | 128 |
| 39 | Results of using stochastic gradients in HMC within SIR while introducing friction (39b) and not (39a). The modes are marked with red 'x' markers, but are covered by the particles.   | 128 |
| 40 | Achieved results using random (40a), adaptive (40b) and random but increasing (40c) measurement times. The modes are marked with red 'x' markers.  | 129 |

|    |   |     |
|----|---|-----|
| 41 | Median successful results using random (40a), adaptive (40b) and random but increasing (40c) measurement times. The modes are marked with red 'x' markers. . . . .  | 130 |
| 42 | Results of 4-parameter estimation (with dimensions plotted pairwise) for two variations of SMC: SIR with the prior as importance function (42a and 42b), and TLE (42c and 42d). The modes are marked with red 'x' markers. . . . .  | 130 |
| 43 | Bloch representations for pure and general states. . . . .  | 133 |
| 44 | Pulse schedule for measuring $T_1$ . . . . .  | 138 |
| 45 | Pulse schedule for measuring $T_2^*$ . . . . .  | 139 |
| 46 | Pulse schedule for measuring $T_2$ . This is known as a <i>Hahn echo</i> or <i>spin echo</i> sequence. . . . .  | 140 |
| 47 | Pulse schedule for estimating the detuning $\delta_{\text{meas}}$ in a Ramsey experiment. This is known as a <i>Ramsey sequence</i> . Note that $\delta$ is an experimental control determining the pulse frequency. The frequency estimate can then be corrected to higher precision as $\Omega_0 = \tilde{\Omega}_0 + \delta_{\text{meas}}$ . . . . .   | 143 |
| 48 | Pulse schedule for measuring $\Omega_0 = \tilde{\Omega}_0 + \delta_{\text{meas}}$ in an echoed Ramsey experiment. Note that $\delta$ is an experimental control determining the pulse frequency. This is a simplification of the ideas in [112], who called this type of setup a <i>Hahn-Ramsey sequence</i> . . . . .  | 144 |
| 49 | Estimated time fluctuations of two quantities of interest for the IBMQ backend <code>ibmq_armonk</code> . These are the parameters to be targeted in the Ramsey experiment of section 5.2.4. The $x$ coordinates are the starting times of data collection. The error bars represent standard deviations. . . . .   | 146 |
| 50 | Evolution of the standard deviation of the dephasing time $T_2$ during the inference process, for the IBMQ device <code>ibmq_rome</code> and two different orders of data addition. Both results were based on the same dataset containing 1500 observations. The shaded region represents the interquartile range. . . . .   | 147 |
| 51 | Curves produced by $T_2$ estimates for the IBMQ device <code>ibmq_rome</code> , using 75 different measurement times. In figure 51a the model of equation 5.21 was used, whereas 51b relaxes it to include two constants $A$ and $B$ . . . . .  | 148 |
| 52 | Curves produced by $T_2$ estimates for the IBMQ device <code>ibmq_casablanca</code> , using also estimates for $A, B \neq 1/2$ . The results were obtained using Bayesian inference (solid red line) and a curve fit (dashed black line), superposed on the data points used for the latter (gray dots). The curve fit relied on 512 shots per point, including the extremes to which it was adapted. For the inference, 20 shots per point were used, and 100 for each of the hyperparameter-adjusting ( $A$ and $B$ ) measurements. . . . . | 148 |
| 53 | Curve produced by the Qiskit coherence fitter for $T_2$ for the IBMQ device <code>ibmq_casablanca</code> , using 15 different measurement times for 169 shots each (a total of 2 535 shots). . . . .  | 149 |
| 54 | Evolution of the standard deviation during the estimation of the energy relaxation time constant $T_1$ of IBMQ device <code>ibmq_guadalupe</code> . The dataset contemplated 1500 observations. The shaded region represents the interquartile range. . . . .   | 150 |

|    |  |     |
|----|--|-----|
| 55 | Evolution of the standard deviations during the estimation of the detuning frequency $\delta$ and coherence factor $\gamma_2^*$ of IBMQ device <code>ibmq_armonk</code> . The dataset contemplated 150 observations. The solid curves are the median results for 300 runs split by 100 datasets. The shaded regions represent interquartile ranges. . . . .  | 151 |
| 56 | Curves produced by $\delta$ (labeled $f$ ) and $T_2^*$ estimates for the IBMQ device <code>ibmq_guadalupe</code> , using 75 different measurement times. . . . .   | 152 |
| 57 | Curve produced by the Qiskit coherence fitter for $T_2^*$ for the IBMQ device <code>ibmq_guadalupe</code> , using 10 different measurement times for 15 shots each (a total of 150 shots). . . . .   | 153 |
| 58 | Ramsey experiment inference results obtained using 100 runs over the same dataset, or distributed by successive datasets spanning a longer period of time. The inference results (dashed red lines) are plotted against the same set of data points (gray dots), from which the data for the inference was extracted in the case of figure 58a. A curve fit (dashed black lines) tailored to these points is included for reference. Each of the data points used 512 shots, whereas for the inference only 2 were used for each time. The targeted device was <code>ibmq_armonk</code> , with 75 different measurement times. . . . . | 154 |
| 59 | Effect of inserting a refocusing pulse in a Ramsey sequence for <code>ibmq_armonk</code> . Each of the 75 data points (black dots) was obtained using 512 shots. To aid visualization, a curve fit (dashed blue lines) performed using SciPy for the function of equation 5.36 is plotted along. . . . .   | 154 |
| 60 | Evolution of the standard deviation during the inference process for adaptive (blue line) and offline (red line) estimation. The shaded regions represent interquartile ranges. . . . .  | 155 |
| 61 | Results of using TLE with HMC move steps to learn a detuning frequency on the IBMQ device <code>ibmq_armonk</code> . The dataset contemplated 75 observations, and ten tempering coefficients were used. . . . .   | 156 |
| 62 | Results of using TLE with RWM move steps and data subsampling with control variates to learn a detuning frequency on the IBMQ device <code>ibmq_armonk</code> . A close up of the distribution is presented, though the prior covered $]0, 10]MHz$ . The dataset contemplated 375 observations, and 10 tempering coefficients were used, 3 of which were employed in the warm up in the subsampling case. When warming up or subsampling, each likelihood evaluation relied on 150/375 data, in the former case the 150 ones with lower evolution times. . . . .   | 157 |

# Acronyms

|             |   |
|-------------|---|
| <b>bPM</b>  | block pseudo-marginal <a href="#">117</a> , <a href="#">126</a>   |
| <b>CLT</b>  | central limit theorem <a href="#">57</a> , <a href="#">58</a> , <a href="#">103</a> , <a href="#">113</a> , <a href="#">115</a> , <a href="#">119</a>   |
| <b>ECS</b>  | energy conserving subsampling <a href="#">112</a> , <a href="#">117</a> , <a href="#">118</a> , <a href="#">126</a>   |
| <b>ESS</b>  | effective sample size <a href="#">82</a>  |
| <b>GRF</b>  | Gaussian rejection filtering <a href="#">80</a> , <a href="#">123</a> , <a href="#">124</a> , <a href="#">125</a>   |
| <b>HL</b>   | Heisenberg limit <a href="#">59</a>   |
| <b>HMC</b>  | Hamiltonian Monte Carlo <a href="#">4</a> , <a href="#">5</a> , <a href="#">7</a> , <a href="#">96</a> , <a href="#">98</a> , <a href="#">101</a> , <a href="#">103</a> , <a href="#">104</a> , <a href="#">105</a> , <a href="#">106</a> , <a href="#">107</a> , <a href="#">109</a> , <a href="#">110</a> , <a href="#">112</a> , <a href="#">113</a> , <a href="#">114</a> , <a href="#">115</a> , <a href="#">116</a> , <a href="#">117</a> , <a href="#">118</a> , <a href="#">119</a> , <a href="#">120</a> , <a href="#">121</a> , <a href="#">123</a> , <a href="#">125</a> , <a href="#">126</a> , <a href="#">128</a> , <a href="#">130</a> , <a href="#">155</a> , <a href="#">159</a> |
| <b>IPE</b>  | iterative phase estimation <a href="#">45</a> , <a href="#">46</a> , <a href="#">62</a>   |
| <b>IQFT</b> | inverse quantum Fourier transform <a href="#">44</a> , <a href="#">45</a>   |
| <b>KDE</b>  | kernel density estimate <a href="#">84</a> , <a href="#">122</a> , <a href="#">123</a> , <a href="#">127</a>  |
| <b>LWF</b>  | Liu-West filtering <a href="#">3</a> , <a href="#">4</a> , <a href="#">5</a> , <a href="#">6</a> , <a href="#">79</a> , <a href="#">87</a> , <a href="#">104</a> , <a href="#">118</a> , <a href="#">125</a> , <a href="#">126</a> , <a href="#">159</a>  |
| <b>MC</b>   | Monte Carlo <a href="#">3</a> , <a href="#">6</a>   |

- MCMC** Markov Chain Monte Carlo [4](#), [5](#), [7](#), [66](#), [91](#), [92](#), [93](#), [96](#), [98](#), [103](#), [104](#), [113](#), [114](#), [116](#), [118](#), [123](#), [125](#), [126](#), [131](#), [145](#), [159](#)
- MH** Metropolis-Hastings [4](#), [5](#), [93](#), [116](#), [117](#), [118](#), [119](#), [121](#), [126](#)
- NISQ** noisy intermediate scale quantum [45](#), [48](#), [158](#)
- NUTS** no U-turn sampler [5](#), [7](#), [110](#), [111](#), [121](#)
- OSR** operator-sum representation [7](#), [135](#), [137](#)
- pCN** preconditioned Crank-Nicolson [4](#)
- PM** pseudo-marginal [117](#)
- QAOA** quantum approximate optimization algorithm [45](#)
- RWM** random walk Metropolis [5](#), [7](#), [95](#), [96](#), [102](#), [103](#), [104](#), [105](#), [106](#), [107](#), [118](#), [119](#), [120](#), [123](#), [125](#), [126](#), [145](#), [149](#), [155](#)
- SG** stochastic gradient [112](#), [118](#), [128](#)
- SHMC** sequential Hamiltonian Monte Carlo [104](#), [118](#), [123](#), [126](#)
- SIR** sequential importance resampling [3](#), [4](#), [5](#), [7](#), [66](#), [82](#), [83](#), [84](#), [88](#), [104](#), [118](#), [125](#), [126](#), [128](#), [145](#), [159](#)
- SMC** sequential Monte Carlo [3](#), [4](#), [5](#), [6](#), [7](#), [66](#), [71](#), [72](#), [78](#), [79](#), [81](#), [88](#), [89](#), [91](#), [92](#), [93](#), [103](#), [104](#), [105](#), [113](#), [114](#), [116](#), [117](#), [118](#), [123](#), [125](#), [130](#), [131](#), [145](#), [146](#), [147](#), [149](#), [151](#), [154](#), [155](#), [159](#)
- SQL** standard quantum limit [58](#)
- TLE** tempered likelihood estimation [4](#), [5](#), [7](#), [86](#), [87](#), [89](#), [117](#), [118](#), [126](#), [130](#), [144](#), [155](#), [156](#), [159](#)





# Introduction

Speaking about quantum mechanics, Albert Einstein famously expressed distaste at the idea that God would play dice with the universe. Yet, as of today, this renders our best means of describing reality. What then is the object of science if not the movement of these dice?

If measurements are dice play, metrology is a game of chance; its ultimate goal is to put the odds in our favor. It should then come as no surprise that methods borrowed from statistics may be especially handy when dealing with quantum systems. In particular, statistical inference is well tailored to the characterization of quantum systems - the pillar of tasks as vital as certification and simulation.

Inference paradigms sketch a path between experimental data and system properties, these data consisting of measurements performed on the system. This work focuses on *Bayesian* inference, a method that relies on the Bayes' theorem to update one's *confidence* in a hypothesis given observations - which we call the *probability* of this hypothesis. This stands in contrast with frequentist statistics, where probabilities are relative frequencies. A downside of the Bayesian method is requiring that we consider multiple hypotheses instead of a single one; a benefit is *allowing us* to consider multiple hypotheses instead of a single one.

In other words, as much as this statistical paradigm may bring added complexity, it is not in vain that it does so. Improved descriptions can be had by averaging over hypothesis, be they different parameters or even different models. What this lacks in intuitiveness, it makes up for in predictive power; and science is ultimately concerned with reality, not some immaterial definition of truth. Projecting *truth* into a stiff mold is no more natural - as in frequentist statistics, where doubt is artificially sifted out of the main theory and left up to meta-statistics.

Furthermore, acknowledged doubt paves the way for shaping strong convictions. In practice, systems do not always behave as our idealized version of themselves. This is doubly true of quantum systems, where noise is joined by fundamental randomness that amplifies our uncertainty. By embracing that uncertainty and welcoming it into our framework, we can reduce it from within - which makes for remarkably reliable and informative methods. Reliable, because they can recover from errors and estimate their own uncertainty; informative, because they can pick out the questions whose answers are the most meaningful.

Another point in favour of the Bayesian approach is the fact that it is quite general, forgoing the need for a new solution in the face of a new problem. What is more, allows for incorporating previous knowledge into the estimation process, something which is often available in physically interesting scenarios. Lastly, it naturally allows for leveraging quantum resources, and for exploiting them beyond what is requisite. The characterization of a quantum nature calls for quantum resources, much like its simulation.

## 1.1 Characterization of quantum systems

If building a quantum device is a challenging task, so too is certifying its behaviour. A closed quantum system, such as an analogue quantum simulator, abides by a Hamiltonian. Being given the intended Hamiltonian, how can one check that it does? The knee-jerk solution is to compare its behaviour with classical predictions. However, this defeats the purpose of a quantum computer, which is of little interest if it can so easily be replicated classically.

As quantum devices approach the frontier of classical tractability, more apt verification methods are in order [51]. Bringing together quantum resources and Bayesian statistics creates a powerful characterization tool for working in the quantum regime, which is crucial for approaching fundamental estimation limits [57, 33].

Bayesian inference systematizes the use of quantum resources for quantum characterization tasks; this could mean state or process tomography, parameter learning, establishing noise models, or others. It relies on comparing the expected behaviour of a comprehensive set of quantum simulators with the behaviour of the quantum system of interest, to obtain the configuration most similar to it. Owing to the broad scope of this analysis, the final product is a probability distribution over a continuum of hypotheses, called the *posterior* distribution. As compared to a single point estimate, this provides an extensive insight rather than a lone guess.

Apart from being a useful assessment of the confidence the results should inspire, this insight arms us with the tools to perform *quantum-enhanced* estimation. Although a probabilistic system will never unequivocally behave like itself, it will under some circumstances more than in others. We can choose those circumstances mindfully if we construct a holistic understanding of it, and allow it to interact with the way we observe it - a technique termed Bayesian experimental design. It leverages the Bayesian insight to forecast the outcomes resulting from the possible choices, aiming to choose the experimental controls that are expected to bring the best results in light of current knowledge. This offers the means to make the most out of each experiment, bringing the greatest possible returns out of the fewest possible measurements.

## 1.2 Research questions, literature overview and objectives

The main open problems pertaining to Bayesian Hamiltonian learning arise in two different fronts. One goal is to balance the quality of the experiments with the computational cost of finding them; and another is to choose the aforementioned *comprehensive set* of quantum simulators, in order to capture relevant information while economizing simulation runs.

The first of these pillars is intrinsic to the inference process, whereas the second one concerns its computational treatment. Together they dictate the usefulness of the inference-produced distribution, being responsible for shaping it and extracting information from it respectively. In short, the two main focuses are:

- **Inference:** choose the experiment controls to maximize the utility of the experiments. Another relevant topic in this category is the choice of models, but this is more context-dependent. Should exact inference be within reach, these and kindred concerns would be the only ones, as they alone shape the posterior distribution.
- **Representation:** discretize the probability distributions so as to best capture the inference results. There is generally a trade-off between resources and quality of the representation, where the resources are simulation runs. They provide the Bayesian *likelihoods*, which tell how likely the simulator would have been to replicate the observed behaviour.

Both problems become harder when scaling up complexity. In particular, when dealing with open quantum systems - which realistically are all quantum devices -, a greater number of parameters is expected to be at play. Their descriptions are given by more intricate formalisms, and learning all of their underlying variables commands both informative data and robust processing. Higher dimensional models are much less forgiving of careless measurements and brute-force algorithms. Further, this difference is not merely quantitative, as they tend to give rise to more complex features - such as multi-modality, the embodiment of redundancy in the probabilistic domain.

These problems require provisions of their own, namely more refined representation methods. Representing the posterior distribution appropriately is key for the correction of the relayed system behaviour, given in terms of expectation integrals over the posterior. [Monte Carlo \(MC\)](#) methods propose to estimate these expectations numerically, by probabilistically choosing samples to be used in numerical integration. Interestingly, they just like inference itself exploit the relation between *simulation* and *characterization*, but now in a statistical (rather than quantum-mechanical) sense. They offer a vast collection of solutions, which can differ greatly in how they produce samples and thus in how efficient they are. Unlike deterministic approaches such as grid-based quadrature, their required resources don't necessarily scale exponentially with the dimension, which is a first step toward escaping what is dramatically called *the curse of dimensionality*.

We now consider the main sources for this project, most of which concern Bayesian learning applied to quantum systems, and their associated discretization strategies. Arguably the most popular representation is via [sequential importance resampling \(SIR\)](#) - a type of [sequential Monte Carlo \(SMC\)](#)/particle filter - with [Liu-West filtering \(LWF\)](#), adopted in [51, 19, 116, 117, 113, 48, 58] as well as by the available software [50]. This is a quite simple and lightweight strategy, especially due to [LWF](#). Its main advantage is keeping quantum simulation to a minimum; the drawback is that it preserves only the two first statistical moments of the distribution. In other words, normal distributions are kept invariant, but others are only assured to maintain their mean and variance. In the latter case, a sensitive trade-off between structure preservation and efficiency is imposed, endangering the viability of the method. This is a glaring defect; not only does it impose severe restrictions on accuracy, it may even lead to absolute failure. In particular, this happens in the case of multimodality, be it native or due to redundancy in the early stages of the inference process (when the information isn't yet decisive). Still less structure preserving approaches are the rejection filtering

of [115] and Gaussian conjugate posteriors of [32], which work under the assumption of absolute normality.

In [49], the **SIR-LWF** protocol is adapted to multimodal distributions by means of clustering, in addition to other sophistications. This provides a working solution, but solves only the most blatant symptom and not the root cause - inaccuracy may still be an issue. Some physically motivated problems are concerned only with approximations and their uncertainty, rejecting the need for more complex or accurate expectations. In this case, **SIR** and **LWF** may suffice, and they are quite enticing due to their sparing use of quantum resources. Even still, simplistic and/or mode-seeking methods are likely to fail in more complex applications (namely in high dimensional parameter spaces). Further, this fix employs costly classical processing, which could be scraped entirely simply by replacing either **LWF** or **SIR** altogether with a more conservative approach. **Markov Chain Monte Carlo (MCMC)** can stand in for either, and is the most obvious solution. The downside is the increase in the computational cost, and if dispensing with **SMC** a likely decrease in the overall convergence speed.

Simple **MCMC** methods have more recently been applied to Bayesian learning for quantum problems. Naive **Metropolis-Hastings (MH)** is applied to quantum tomography in [81]; [118] points out the drawbacks of that type of strategy and uses slice sampling instead, but acknowledges that e.g. multimodal distributions require more advanced methods, mentioning **SMC** as a possible solution. It is remarked that **SMC** may incur inaccuracies, but this strongly depends on the implementation - while it may be true of **SIR** with **LWF**, this is not necessarily so for more robust variants. Another source using slice sampling is [76].

Still more recently, **preconditioned Crank-Nicolson (pCN)** has been applied to pseudo-Bayesian quantum tomography [77, 78, 80], improving upon vanilla **MH**. This type of **MCMC** is rather recent, having been introduced in [21]. By adapting the stepsize, it offers better acceptance probabilities for high dimensions as compared to standard **MH** algorithms. However, it is not informed by the geometry of the target distribution, and so is still expected to produce diffusive transitions. Another more popular and widely studied variation that scales favourably with the dimension is **Hamiltonian Monte Carlo (HMC)**<sup>1</sup> [29, 88, 15], which brings two added benefits: good mixing, and robustness under complex distribution features [10]. This is one of the methods tested here. Also, none of the above consider embedding **MCMC** in an **SMC** scheme, which may be very advantageous. **SIR** is more reliable than isolated **MCMC** and suitable for online estimation, enabling adaptivity. If post-processing, a still more reliable option is **tempered likelihood estimation (TLE)** [89], a well-established method for complex statistical simulation problems.

Other publications on this subject typically consider less general approaches, or focus on other points instead. In [100], an exact representation (a Fourier series) makes a simple single parameter problem analytically tractable, but this method doesn't generalize well to more complex inference problems. Similarly, rough approximations such as the Gaussian mixtures of [22] can be quite agile for few parameters, but are generally not scalable. Some papers focus mostly on the experimental design and barely [60, 38] or not at all [33] in the discretization; [36] tackles the same problem but makes a normality assumption to make it at once general and analytically tractable. Likewise, many attempt to improve yet other aspects,

---

<sup>1</sup>Note that the ideas behind **pCN** and **HMC** are not mutually exclusive, and **pCN**-like improvements can be adopted for **HMC** too. References [21, 10, 1] elaborate on the relation between these methods.

while briefly mentioning [SMC](#) strategies and referring to the sources above. Some notable examples are [\[35\]](#), which relates the effect of relaxing the simulation requirements, [\[34\]](#), which considers embodying model uncertainty to improve robustness, and [\[44\]](#), which presents a more general algorithm that loosens the parametrization assumptions.

This dissertation intends to overview the main foundational topics of Bayesian learning, and to combine them with more advanced numerical methods that enjoy widespread acceptance by the statistical community. The final objective is to test the performance of these algorithms with an emphasis on quantum applications, while also implementing some of the standard strategies for comparison. The central statistical methods to be approached are all Monte Carlo algorithms, and fall into three categories:

- [MCMC \(MH\)](#): [HMC](#) and [random walk Metropolis \(RWM\)](#). This includes enhancements of these methods, such as subsampling and extensions of [HMC](#) (namely the [no U-turn sampler \(NUTS\)](#)).
- Improvements to the [SIR-LWF](#) algorithm. More specifically, replacements for [LWF](#) with stronger assurances are considered. They precisely match the [MCMC](#) methods mentioned above.
- [TLE](#), an [SMC](#) algorithm distinct from [SIR](#).

[HMC/NUTS](#) and [SMC](#) with [MCMC](#) kernels (namely [TLE](#)) are techniques used in leading software, such as state-of-the-art probabilistic packages/programming languages for Bayesian inference and statistical simulation: [PyMC3](#) [\[98\]](#) supports both, and [Stan](#) [\[42, 105\]](#) intends to despite currently leaving [SMC](#) out. Another technique they both use is variational inference, perhaps the most interesting lower precision and lower cost alternative to Monte Carlo algorithms - but this dissertation will focus on the latter.

Overall, the aims of this dissertation fall into four main topics, each corresponding to a chapter. First, to overview Bayesian inference, with an emphasis on the problem of learning the dynamical parameters of quantum systems. This includes experimental design, namely adaptive experimental design. Second, to explore numerical methods well-tailored to this paradigm and evaluate their relative efficiency. These are mostly Monte Carlo methods. Third, to apply the considered strategies to representative test cases in order to assess their merits. And lastly, to use this knowledge to characterize quantum devices, considering in particular quantum channels capturing decoherence processes or other open system effects. For this end, the theoretical foundations necessary to realize simple open system descriptions are considered as well.

## 1.3 Related topics

We will now allude to some well-known subjects that are connected with the ones to be explored in this dissertation, and the differences that separate them.

Quantum state tomography is the most related concept; however, it targets an instantaneous quantum state and not the entirety of the systems' dynamics. Quantum process tomography, on the other hand, means to characterize processes in the sense of operations or maps, and not continuous time evolution. By contrast, here we consider the dynamics undergone by a determined initial quantum state; furthermore, this is realized under some practical assumptions, resulting a full description of the system's time evolution.

This can be given by operator(s), or a function; either way, the entity to be characterized is time-dependent, and so too are the measurements. As an example, this can mean determining the evolution of the excited state population as opposed to an instantaneous density matrix. More generally, any dynamical parameters could be targeted in the learning process. And finally, unlike in tomography, all the experiments here use measurements on a single basis - though this is a choice and not the method's own limitation.

Another related topic is quantum machine learning. While Bayesian learning can broadly speaking be considered machine learning, quantum machine learning in its usual sense is outside of the scope of this work. It generally refers to quantum-*enhanced* machine learning, i.e. machine learning algorithms which seek to extract an advantage from quantum resources. In this case, the learning techniques are entirely classical, only they're *applied* to quantum systems. In a few words, in our case quantum systems are a *subject* and not a tool.

Likewise, quantum Monte Carlo - which deals directly with quantum problems - is quite distant from the Monte Carlo techniques approached here, inasmuch as the latter are not quantum by construction. Rather, they are structurally application-blind, their reliance on quantum simulation notwithstanding. They simply require resources that match the problem; should the problem be quantum, so too must they.

Finally, quantum Bayesianism is associated with both quantum systems and Bayesian statistics, but on an epistemic level. It gives rise to a multitude of interpretations of quantum mechanics, the most notable being QBism. This dissertation focuses merely on the mathematical aspects of Bayesian statistics, and how it can serve quantum-mechanical systems (*quantitatively*).

## 1.4 Document structure

For a clearer picture, figure 1 schematically represents the topics to be covered throughout the document. The left half of the tree corresponds to chapter 2, while the right half respects chapter 3. Chapters 4 and 5 showcase the results of applying this collection of methods to a selection of test cases; they consider statistical simulation and characterization of quantum devices from IBMQ [61] respectively.

More specifically, the dissertation is structured as follows. Chapter 2 begins with the basics of Bayesian statistics and inference (sections 2.1 and 2.2), including basic methods for exact and approximate inference (subsection 2.2.2). After that, section 2.3 discusses how this applies to quantum systems and the resources involved (2.3.2), and provides illustrative examples (2.3.3). It also touches upon model selection (2.3.1), which opens the door to more general algorithms. Finally, section 2.4 respects Bayesian experimental design strategies. It covers both global optimization strategies and local adaptive ones (2.4.1), as well as how they relate to the fundamental limits on measurement precision (2.4.2). This section finishes with some cost-effective heuristics for a typical quantum example (2.4.3).

After this overview of Bayesian inference, chapter 3 focuses on complementary numerical methods whose purpose is to characterize the resulting distribution. It starts with the motivation and idea behind MC methods in section 3.1, and follows with SMC in 3.2. Contained therein is LWF (3.2.1), as well

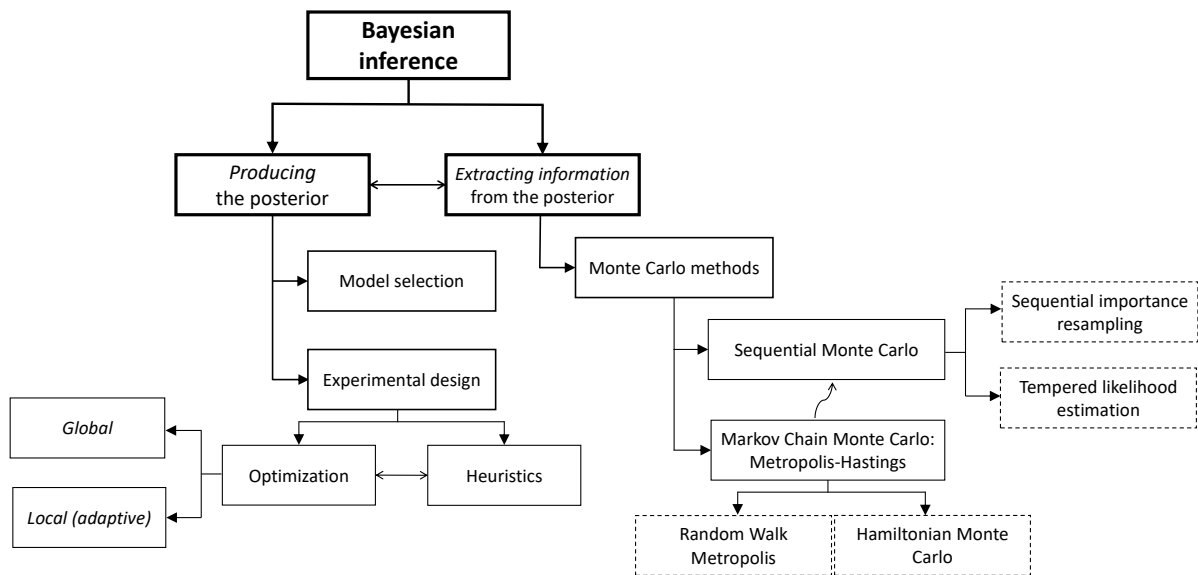


Figure 1: Diagram of the topics to be covered in this dissertation. Dashed boxes denote specific algorithms.

as [SIR](#) (3.2.2) and [TLE](#) (3.2.3). Then, section 3.3 introduces [MCMC](#), exemplifying the algorithm with the elementary [RWM](#) variant (3.3.1) before discussing [HMC](#) (3.3.2) and the usefulness of [MCMC](#) for [SMC](#) (and vice versa; subsection 3.3.3). Finally, section 3.4 delves deeper into [HMC](#), namely its applicability (3.4.1), progressive biased variants (3.4.2), and dynamic implementations such as the [NUTS](#) (3.4.3). Section 3.5 finishes with an overview of subsampling strategies, with an emphasis on [HMC](#): control variate-enhanced estimators (3.5.1), stochastic gradients (3.5.2) and energy conserving subsampling (3.5.3).

Chapter 4 considers a few applications of statistical simulation. Three simple sampling problems are considered first in section 4.1 (6-dimensional Gaussian in subsection 4.1.1, Rosenbrock function in 4.1.2, a smiley face in 4.1.3), followed by two inference problems using generated data in section 4.2 (phase estimation in subsection 4.2.1, multivariate sum of squared sinusoids in 4.2.3).

As for chapter 5, it starts with some rudiments on open quantum systems (mixed states in the Bloch picture in subsection 5.1.1, and the Kraus [operator-sum representation \(OSR\)](#) in 5.1.2) and the grounds for the experiments to be performed in section 5.1 (subsections 5.1.3 and 5.1.4 present simple noise channels depicting amplitude and phase damping respectively, and 5.1.5 discusses Larmor/Rabi/Ramsey oscillations). The results of these experiments follow on section 5.2. Relevant considerations, including an overview of the methods, are laid out in subsection 5.2.1, while the following sections present and compare the results of using measurements to characterize qubit dynamics for a few physical processes: Hahn echo sequences (subsection 5.2.2), energy relaxation (5.2.3), Ramsey sequences (5.2.4), and *Hahn-Ramsey* sequences (5.2.5).

Lastly, chapter 6 briefly discusses the results and recapitulates the main takeaways.

The code used for this project can be found online at the [GitHub repository](#) [3].

# Quantum Parameter Estimation

The attribution of limited predictive power to reality itself has been a source of controversy in quantum mechanics, and one whose implications reach well beyond conceptual perplexity. Ultimately, to reject predictability is to resign to ambiguity; that truth leaves the future to chance precludes this future from conveying its past.

Less abstractly, the fact that quantum measurements are inherently probabilistic bars any single one from disclosing the complete state of the system. This indetermination may be overcome by repeating different measurements multiple times on an ensemble of states to be characterized. But evidently, this adds a layer of complexity to metrology, effectively placing data collection on the same level as the usual experimental limitations (such as the precision of the measuring apparatus).

In any case, should we want to describe the physics of quantum systems, we must honor their non-deterministic essence. This can only be achieved by adopting characterization techniques appropriate to phenomena of probabilistic nature. Hence, methods borrowed from statistical inference may be of use when one means to work with technology in the quantum regime.

One such method is Bayesian inference. Interestingly, Bayesianism is another theory which has been subjected to criticism due to holding probabilities a little too close to heart. However, it is also one more which came to enjoy widespread acceptance on account of its practical successes, once they were made possible by the available technological means.

This chapter will present basic concepts of Bayesian statistics and inference, and how they can be of use to quantum science. We will use a few basic examples to illustrate these techniques and contrast them with their more traditional counterparts. The formal justification for using such techniques will also be briefly touched on, as well as some refinements that the framework enables.

## 2.1 Bayes' rule

If we repeatedly toss a coin and it always lands tails up, we might strongly suspect that it is biased. But what if the trend is more subtle than that, or multi-factorial? The outcome distribution may depend on its temperature, or the throwing angle. That doesn't make it any less of a tendency - simply one we struggle more to understand.

From determining whether a coin is fair to getting a thorough description of its behaviour, the leap is an monotonous one, and relies mostly on uninspired processing. Of course, *uninspired* doesn't necessarily mean easy. While the human mind may excel at many things, analysing large amounts of data is not one



of them.

Fortunately, there exist machines to do that for us; we have but to instruct them on how to. This brings the additional complication of doing so - but ever since the advent of computers, it has often been rewarding to translate intuition into reason, and reason into mechanical work.

Bayes' rule renders the former transition. It arises from what is essentially a *common sense* observation, and translates empirical intuition into a powerful mathematical expression.

In order to present it, it is only natural that we start with its constituents. Firstly, we have an entity  $A$ , whose *probability* we write as:

$$P(A)$$

As basic as this expression may seem, its interpretation captures the essence of Bayesian thinking. In frequentist statistics, the probability of an event  $P(A)$  is simply a frequency: it asymptotically matches the relative frequency of occurrence of  $A$  (in the scope of an experiment). In Bayesian statistics, on the other hand,  $P(A)$  is viewed as a *degree of belief* in  $A$ . In this interpretation, a probability describes the knowledge held by some *agent* (such as ourselves), rather than a property of the phenomenon.

For instance,  $P(A)$  could represent how likely we initially think it is that the aforementioned coin is fair (i.e. that it yields equal probability of flipping heads or tails). Let's say somebody told us that it was, but we only moderately trust them. Then we might attribute the fairness a confidence of 50%:

$$P(\text{fair coin}) = 0.5 \tag{2.1}$$

This notion of *confidence* or *degree of belief* may not seem very *mathematical*, insofar as it appears to hold some subjectivity - but it lies at the heart of Bayesian thinking. For all the things that set it apart from a frequentist approach, it is this this shift from *probabilities as frequencies* to *probabilities as quantifications of confidence* what best illustrates their fundamental divergence.

One may argue that this stands in the way of impartiality, but such is not necessarily the case. We may choose to rely essentially on experimental data alone to establish these *beliefs*, thereby removing the practitioner's own partisanship from the process.

A more concrete take on what exactly this uncontentious choice entails is left for section 2.2, but the key point is that we leave it up to the data to shape our convictions, as well as the confidence we have in them. Because these data are produced precisely by the reality we mean to study, one would be hard pressed to find a more unbiased source of insight.

Next up are *conditional* probabilities, which play a central role in Bayesian statistics. We write the probability of  $A$ , conditional on  $B$ , as:

$$P(A \mid B)$$

This can be read as the probability of  $A$ , given  $B$ . Note that this *probability* continues to be interpreted in the Bayesian sense. For example, it could indicate how strongly we are convinced that the coin is fair, given that we have tossed it three times and always got a tails outcome:

$$P(\text{fair coin} \mid \{\text{tails, tails, tails}\}) \tag{2.2}$$

We may be skeptical that such is the case, since a fair coin would have been unlikely to only produce tails outcomes - but it would not be impossible, so there's still a chance. Of course, how plausible it is depends on the total amount of coin tosses. In this case, they were indeed not many.

At this point, we may start wondering whether these *beliefs* actually belong to the realm of science. *How exactly* are we supposed to quantify how much we trust that the coin is fair based on some outcomes?

The answer is, *we're not* - that's precisely where Bayes' rule steps in.

We can straightforwardly derive it by remarking on the fact that the joint probability of two events, namely  $A$  and  $B$ , cannot depend on the arbitrary order according to which we choose to write them.

$$\mathbf{P}(A \cap B) = \mathbf{P}(B \cap A) \quad (2.3)$$

Then we move on to *conditional* probabilities. Naturally, the joint probability of  $A$  and  $B$ ,  $\mathbf{P}(A \cap B)$ , is the conjunction of  $B$  occurring with  $A$  occurring given that  $B$  occurs. Since the conjunction of two events translates into the product of their probabilities, this means that we write equation 2.3 as:

$$\begin{aligned} \mathbf{P}(A | B)\mathbf{P}(B) &= \mathbf{P}(B | A)\mathbf{P}(A) \\ \leftrightarrow \mathbf{P}(A | B) &= \frac{\mathbf{P}(B | A)\mathbf{P}(A)}{\mathbf{P}(B)} \end{aligned} \quad (2.4)$$

, where we have assumed  $\mathbf{P}(B)$  to be non-zero (with little to no loss of generality, as it's rather unlikely we would need to compute the aftermath of an impossible event).

The last form of equation 2.4 is the statement of Bayes' theorem. It provides a quantitative measure of how we should change our degree of belief in  $A$ , should  $B$  happen.

Note that this brings somewhat of an asymmetry between  $A$  and  $B$ , as is evident in the rearrangement of equation 2.4. And as a matter of fact, these variables assume quite different roles.  $A$  is an underlying factor that influences an observation;  $B$  is the object of that observation.  $A$  is or is not true;  $B$  does or does not happen.  $A$  is a hypothesis (e.g. " $\mathbf{P}(\text{tails}) = 0.5$ ");  $B$  is an event (e.g. "tails landed up").

As such, with the purpose of making notation clearer,  $A$  is often replaced by  $\theta$  or  $x$  (denoting a possible value of an unknown parameter), and  $B$  by  $D$  (denoting data, or a single datum). With this choice of notation, equation 2.4 becomes:

$$\mathbf{P}(\theta | D) = \frac{\mathbf{P}(D | \theta)\mathbf{P}(\theta)}{\mathbf{P}(D)} \quad (2.5)$$

Relatedly, there is a more straightforward manner of computing  $\mathbf{P}(D | \theta)$  than there is its antipode  $\mathbf{P}(\theta | D)$ . Reverting equation 2.2, we get:

$$\mathbf{P}(\{\text{tails, tails, tails}\} | \text{fair coin}) \quad (2.6)$$

Here we simply wish to know how likely these outcomes are, given that the coin is fair. This can be readily assessed, without much in the way of a philosophical discussion:

$$\mathbf{P}(\{\text{tails, tails, tails}\} | \text{fair coin}) = \prod_{i=1}^3 \mathbf{P}(\text{tails} | \text{fair coin}) = \left(\frac{1}{2}\right)^3 = \frac{1}{8} \quad (2.7)$$

Note the strategic avoidance of the word *probability*. Our previous interpretation of probability as strength of belief is somewhat uncalled for here: having flipped tails or not is hardly a matter of conviction. Are the relative frequencies of random events demoted from their status as probabilities in Bayesian statistics? As one might expect - or hope -, they are not.

The Bayesian concept of probability encompasses those too, and their probabilities may still be read as frequencies; the Bayesian perspective is simply more flexible than the frequentist one. That flexibility is made apparent by the fact that Bayesian strategies can be (and are) gainfully employed in otherwise frequentist paradigms. We will return to this fact later.

Even still, we typically distinguish the roles of the parameter  $\theta$  and the event  $D$ , and call  $\mathbf{P}(D \mid \theta)$  the *likelihood* of  $\theta$ , given  $D$ . Again, it still identifies a probability; it *is* a probability by definition.

$$\mathbf{L}(D \mid \theta) \equiv \mathbf{P}(D \mid \theta) \quad (2.8)$$

Regardless, the likelihood formulation is not only correct, but also more descriptive. Considering how likely the parameter was to have generated the data, rather than how probable the data was given the parameter, emphasizes the position of each element. The order reversal of the arguments in 2.8 suits this purpose as well: it evidences the central standing of  $\theta$ , while placing  $D$  in a supporting role. Conveniently, we also avoid overburdening the term *probability*, which streamlines any discussions relating 2.5.

Plugging this into equation 2.5, we get our final statement of the Bayes' rule:

$$\mathbf{P}(\theta \mid D) = \frac{\mathbf{L}(\theta \mid D)\mathbf{P}(\theta)}{\mathbf{P}(D)} \quad (2.9)$$

Now that we have our final formula, let's try to make it a bit clearer by getting back to our coin example. In that case,  $\theta$  would be the probability of tails (or equivalently, of heads):

$$\theta \equiv p_{\text{tails}} \quad (2.10)$$

In more complex cases where we have not a binomial distribution, which is described by a single variable - the probability of success -, but one which requires multiple parameters,  $\theta$  will be a vector as long as the number of parameters. For instance, if its components are real numbers:

$$\vec{\theta} \in \mathbb{R}^d \quad (2.11)$$

, where  $d$  denotes the dimension of the parameter space where  $\vec{\theta}$  lives.

On the other hand,  $D$  would be an outcome:

$$D \equiv \text{outcome} \in \{0, 1\} \quad (2.12)$$

, where for the sake of conciseness we have mapped to 0 and 1 the two possible outcomes (heads and tails respectively).

$$\{\text{heads, tails}\} \equiv \{0, 1\} \quad (2.13)$$

Naturally, other contexts could give rise to different sets of outcomes. In our case of a binomial distribution, only two exist; the likelihood is then given by:

$$\mathbf{L}(\theta | D) = \theta^D (1 - \theta)^{1-D} \quad (2.14)$$

, as it represents the probability that the specific parameter value  $\theta$  would have produced outcome  $D$ .

Returning to the general case given by equation 2.9, we can identify as the central element of the right-hand side  $\mathbf{P}(\theta | D)$ , the *likelihood* that the parameter(s)  $\theta$  would result in  $D$  happening. On the other hand,  $\mathbf{P}(\theta)$  represents our belief in  $\theta$ . Simply put, it stands for the probability of its being true, to the best of our knowledge - which doesn't have to be much, as we shall get to shortly.

Regardless, once we get experimental data, such as heads/tails outcomes, this belief is obsolete. If we thought that the probability of tails was 100%, but flip the coin and get heads, then that conviction no longer stands. Likewise, any outcome should slightly tweak our knowledge about  $\theta$ , hopefully getting it a bit closer to the truth.

This fine tuning process is achieved by applying equation 2.9. Schematically:

$$\mathbf{P}(\theta) \xrightarrow{D} \mathbf{P}(\theta | D) \quad (2.15)$$

, where the arrow with superscript  $D$  denotes evolution contingent on the occurrence of  $D$  (i.e. on having observed the datum  $D$ ).

The pivotal idea is that we get a fresh stance, based on data collection. This stance reflects both our *prior* beliefs, as evinced by the term  $\mathbf{P}(\theta)$ , and the experimental results we observed, as evinced by the term  $\mathbf{L}(\theta | D)$ . This produces our *posterior* beliefs, which have been revised in light of the new evidence:

$$\mathbf{P}(\theta | D) \propto \mathbf{L}(\theta | D)\mathbf{P}(\theta) \quad (2.16)$$

If we gather a lot of experimental data that our  $\theta$  hypothesis would have been very likely to generate, we have more of a reason to believe that it *did* generate these data. That is portrayed in the factor  $\mathbf{L}(\theta | D)$ . If on the other hand our hypothesis parameter  $\theta$  would have been unlikely to replicate the events that took place, the very same factor will penalize its probability.

Lastly,  $\mathbf{P}(D)$  is the prior probability of  $D$ , also called the marginal probability, marginal likelihood, or the model *evidence* (this refers to the statistical model, which provides the likelihood and the prior probability). It represents the *a priori* probability of  $D$  happening for any possible  $\theta$  (i.e. integrated or marginalized over the prior for  $\theta$ ). It has been left out of equation 2.16 because it doesn't take as much of a conceptual significance as the other elements. In fact, for most of our purposes it works as a normalization factor, assuring that  $\mathbf{P}(\theta | D)$  integrates to one as any probability should. It is given by:

$$\mathbf{P}(D) = \int \mathbf{P}(D | \theta)\mathbf{P}(\theta)d\theta \quad (2.17)$$

By this we don't mean to disregard its importance; it may be of relevance when considering other applications, such as Bayesian model comparison (subsection 2.3.1). This divisor will again be briefly

touched upon in section 2.2, as will the practical details regarding its computation (or that of an equivalent counterpart within the scope of our framework).

With all the elements at hand, one key aspect is perhaps worth stressing. As conveyed by the representation in 2.16, there is an implicit *evolution* from  $\mathbf{P}(\theta)$  to  $\mathbf{P}(\theta | D)$ . This isn't just a separate calculation which depends on  $\mathbf{P}(\theta)$  - rather, it provides an improved version of it. Hence, there's a hierarchical relation of *substitution* between these mathematical entities.

So if we register a datum  $D_1$ ,  $\mathbf{P}(\theta | D_1)$  is our *updated*  $\mathbf{P}(\theta)$ . The « $| D_1$ » annotation is purely denoting that this calculation takes  $D_1$  into consideration, whereas in  $\mathbf{P}(\theta)$  no such explicit indications were present. If we subsequently want to add yet another datum, namely  $D_2$ , to the estimate, then  $\mathbf{P}(\theta | D)$  takes the place of  $\mathbf{P}(\theta)$  in equation 2.9, and so on for any additional data.

This is in accordance with the fact that adding the data one by one or all at once should have no bearing in our findings. In equation 2.12, we defined  $D$  as an outcome. Alternatively, we can define it as a *list* of  $N$  outcomes:

$$\vec{D} \equiv \langle \text{outcome} \rangle \in \{0, 1\}^N \quad (2.18)$$

But this is in truth no more general, because the probability given a vector of independent data is simply the product of the probabilities of its individual constituents.

$$\mathbf{P}(\theta | \vec{D}) = \mathbf{P}\left(\theta | \bigwedge_{j=1}^N D_j\right) = \prod_{j=1}^N \mathbf{P}(\theta | D_j) \quad (2.19)$$

This is true if both the likelihood and the marginal probability are themselves products over all the data (which they clearly are, by virtue of their respective definitions as joint probabilities), and  $\mathbf{P}(\theta)$  is replaced consecutively by

$$\begin{aligned} \mathbf{P}'(\theta) &\equiv \mathbf{P}(\theta | D_1) \\ \mathbf{P}''(\theta) &\equiv \mathbf{P}((\theta | D_1) | D_2) \equiv \mathbf{P}(\theta | (D_1 \wedge D_2)) \equiv \mathbf{P}(\theta | D_1, D_2) \\ &\dots \end{aligned}$$

in expression 2.9 (when adding data one by one). So whether we consider multiple data or a data vector is of no consequence, as should be. Should we wish to do so, we can consider a series of outcomes at once when updating our distribution.

$$\mathbf{P}(\theta | \vec{D}) = \frac{\mathbf{P}(\vec{D} | \theta)\mathbf{P}(\theta)}{\mathbf{P}(\vec{D})} = \frac{\mathbf{L}(\theta | \vec{D})\mathbf{P}(\theta)}{\mathbf{P}(\vec{D})} \quad (2.20)$$

In short, we have constructed a framework that allows us to adjust our knowledge so that it better matches our observations - which ultimately means that it better matches reality. This provides a way of correcting our expectations about probabilistic phenomena, such as coin flipping, by proceeding in accordance with the results they produce. In a nutshell, we characterize a random phenomenon by sampling from it.

This approach is in order when a phenomenon we are interested in is governed by an unknown - some hidden variable(s) that dictate its (probabilistic) behaviour, such as the probability of flipping tails. Our expectations then take the form of degrees of belief, which we assign to specific instances of this unknown. But what good is that by itself? And what if we don't have any knowledge to begin with? We do not merely want to evaluate how good an *isolated* hypothesis is - we want to secure a reliable one.

Clearly, a broader approach is due. We must build on this framework and devise a systematic procedure - one which will allow us to retrieve more serviceable information. In other words, we want a protocol for inference.

## 2.2 Bayesian inference

In frequentist statistics, the very definition of probabilities equipped us with a direct way of computing their estimates; unfortunately, the same cannot be said for its Bayesian analogue.

In the former case, we could effortlessly derive an estimate for a coin's probability of tails by flipping it several times. That's because we would identify the true probability of the *tails up* event with the limiting value of its relative frequency, i.e. with its propensity of occurrence as the number of trials (flips) approaches infinity:

$$\theta \equiv \lim_{\text{coin flips} \rightarrow \infty} \left( \frac{\text{number of tails up outcomes}}{\text{coin flips}} \right) \equiv \lim_{N \rightarrow \infty} \frac{\sum_{j=1}^N D_j}{N} \quad (2.21)$$

, where the quotient has been written more compactly by defining  $N$  as the total number of coin flips, recalling convention 2.13, and representing with  $D_j$  the  $j$ th datum/outcome (so  $j \in \{1, \dots, N\}$ ).

Of course, this is unachievable, so we would truncate the expression after some finite number of experiments.

$$\tilde{\theta} \approx \frac{\text{number of tails up outcomes}}{\text{coin flips}} \quad (2.22)$$

By contrast, Bayes' theorem gives us solely  $\mathbf{P}(\theta)$ , our degree of trust in a given hypothesis  $\theta$  (conditioned on some information). In the previous case, at no point did we even attribute probabilities to different values of  $\theta$ , or contemplate these values at all - but now this is all we have. As for the expression being general with respect to  $\theta$ , that can be dealt with by sweeping over the possible  $\theta$ s. But how can we procure our  $\tilde{\theta}$  while relying on the same information as before? And how will the result differ?

Let us start by focusing on equation 2.16. The likelihood  $\mathbf{L}(\theta | D)$  can be obtained for any  $\theta$ , using only the vector of outcomes. We will use  $\vec{D}$  to emphasize that this vector encompasses the results  $D_j$  from multiple trials (coin flips). Then:

$$\mathbf{L}(\theta | \vec{D}) = \prod_{j=1}^N \theta^{D_j} (1 - \theta)^{1-D_j} \quad (2.23)$$

, where we have applied 2.14 and used the same conventions as in equation 2.21. This highlights their shared dependence on  $\vec{D}$ , the list of outcomes, and nothing else ( $N$  is simply the cardinality of  $\vec{D}$ ).

Here we have chosen to consider an ordered vector, in anticipation for other applications where our data won't all be equivalent (due to additional flexibility in the measurements). If we disregard the order and consider a set of outcomes instead, we get the more familiar form of a binomial distribution,

$$\mathbf{L}(\theta \mid \vec{D}) = C(N, s) \cdot \theta^s (1 - \theta)^{N-s} \quad (2.24)$$

, where  $C(N, s) = \binom{N}{s}$  is the binomial coefficient and  $s \equiv \sum_{j=1}^N D_j$  is the number of successes (1 outcomes) among the  $N$  trials.

At this point, we already have something interesting: a general formula for the likelihood of specific instances of the variable  $\theta$ , which is provided by the data alone (along with the description, or model, of the phenomenon - in this case, that's just a binomial distribution). Now all we're missing in proportionality 2.16 is  $\mathbf{P}(\theta)$ .

As said before, this  $\mathbf{P}(\theta)$  consists of the initial probabilities we assign to possible values of  $\theta$  - that is, of the probability distribution choose to we start with. If we were inclined to believe it is nearly fair, we could start out with a Gaussian with some spread centered at  $\theta = 1/2$ . Of course, we should be cautious in making our choice. We don't want to bias our estimate or slow down convergence by making unwarranted assumptions. Yet, we might not know much about our coin, nor about its fairness, and as such not be confident enough to include any such inclination in our reasoning. In that scenario, the probability of tails could be anywhere between 0 and 1, and we wouldn't have any sound reasons to favor some values over others. How then could we define this prior distribution?

We already have. We have established that it:

- Must treat all possible  $\theta$ s equally, i.e. cannot depend on  $\theta$ ;
- Must integrate to one, because it is a probability distribution;
- Must be 0 for  $\theta \notin [0, 1]$ , because  $\theta$  itself is a probability.

These three characteristics uniquely define a probability distribution. It consists of a flat distribution covering all possible values:

$$\mathbf{P}(\theta) = \begin{cases} 1, & \text{if } 0 \leq \theta \leq 1 \\ 0, & \text{otherwise} \end{cases} \quad (2.25)$$

While in more complex cases plenty can be said about how a *representation of ignorance* can be best achieved, for many applications of interest a so-called *objective* prior distribution can be found - that is, one which conveys little to no information, and so barely affects the inference results. This is the case of our example above, where the prior simply expressed problem constraints and otherwise followed the *principle of indifference* by assuming a uniform distribution over the plausible region (the one delimited by these constraints).

What is more, often the prior beliefs barely affect the posterior ones at all, so long as enough data is considered (although this defense has in some cases come to be criticized [30], it stands for many others). As put by [31], even if two individuals initially disagree, the inference process is expected to

reconcile their positions. That means that by collecting enough evidence, we could expect to bring their viewpoints together up to unbounded like-mindedness (assuming they're not Bayesian skeptics). For this reason, it may be unnecessary to overly stress the choice of a prior; resources could be more gainfully employed in gathering the data that should cause its relevance to subside.

This is more formally expressed by the Bernstein–von Mises theorem [109, chapter 10], which proves under some fairly inclusive assumptions that the posterior distribution can be made not to depend on the prior by means of data collection. These assumptions often hold for physically motivated estimation problems, under the typical circumstances in which these problems occur (namely the class of priors that they prompt).

It may be worth mentioning that this matter brings forth the major divergence within Bayesianism, which is split into what are called objective and subjective interpretations. The former is more temperate in how it adopts the idea of probabilities as beliefs, trying to stay by *rational* beliefs; the latter interprets it quite literally (as in personal beliefs). In practice, this distinction mostly concerns the prior.

Finally, having both a prior and a likelihood, we can compute the right-hand side of proportionality 2.16 in full. By evaluating it on a grid and interpolating a continuous curve, we can visualize the result. Figure 2 represents the posterior probability density - which is proportional to  $\mathbf{P}(\theta \mid D)$  - that it yields for a randomly generated data vector.




Figure 2: Bayesian posterior probability density for a binomial likelihood. The dashed and full vertical lines mark the real parameter used to generate the data and its frequentist estimate respectively. The same data vector contemplating 100 trials was used for the Bayesian and frequentist calculations.

The omission of the y-axis scale on figure 2 is analogous to the omission of the divisor in proportionality 2.16. Clearly, these pieces of information are not strictly necessary for all purposes: even visually, we can



approximate the maximum just by looking at the probability density graph. For many applications, relative probabilities are sufficient.

But even if not, we don't need a *separate* or very precise computation of the marginal  $\mathbf{P}(D)$ . What we need is given by our distribution already, because it is a normalization factor. In effect, that is *the best possible* estimate of  $\mathbf{P}(D)$ , because it is the only one which guarantees that our probability distribution behaves like one.

This discussion may be somewhat unclear without discussing the practical representation of the distributions, or how the integrals are to be evaluated. Figure 2 shows a smooth curve achieved through interpolation, but this does not have much bearing in practice because in general a continuous representation is not viable. Further clarification is left for subsection 2.2.2, which will cover the most conceptually simple implementations. But before that, we will look into Bayesian inference a little bit more.

### 2.2.1 Estimators and expressions of uncertainty

The furthest we got so far was a probability density function given by equation 2.16 and represented in figure 2. This function quantifies how likely each of the values that the parameter  $\theta$  could take would have been to produce the outcomes we observed<sup>1</sup>. From there, we must still extract an estimate  $\hat{\theta}$ .

An estimator, usually denoted by a circumflex - as in  $\hat{\theta}$  -, is a construction which maps the posterior (or equivalently, the conjunction of the prior and the data) to an estimate. Looking at the graph from figure 2, the answer seems obvious: we should just take the value of  $\theta$  corresponding to the posterior peak. In this case, the choice is made painless by the fact that the values are approximately normally distributed: for normal distributions, the mean, mode and median exactly match. We can pick any - all - of them as our estimator.

By contrast, we will generally have to pick a single one when looking at less regular posteriors. Which to select is then an implementation choice. When discussing this selection, we generally speak in terms of *loss functions* (or *utility functions*, a symmetric construction). They typically penalize inexactness (i.e. deviations of the estimators with respect to the real values) in different ways:

$$\mathcal{L}(\theta, \hat{\theta}(\vec{D})) = f(|\theta - \hat{\theta}(\vec{D})|) \quad (2.26)$$

Note the calligraphic  $\mathcal{L}$ , which means to avoid confusion with the  $L$  symbol we use for the likelihood.

We will mention some popular choices of loss functions shortly, as well as the estimators they prompt. But the process of determining these estimators requires further consideration: as is, equation 2.26 is unusable, by virtue of its dependence on  $\vec{D}$  and  $\theta$ . We would like to somehow keep the loss to a minimum; this expression doesn't allow us to evaluate it at all, still less minimize it.

Of course, the actual loss incurred by an estimator will depend on the collected data; this is made clear by the fact that we don't get our estimate until the inference process has terminated. We can however

<sup>1</sup>Recalling that a flat prior was used, we know that the posterior is maximized precisely at the likelihood maximum. The prior probabilities  $\mathbf{P}(\theta)$  - and of course  $\mathbf{P}(D)$ , the marginal - exactly match for all points.

evaluate the *a priori* merit of an estimator by considering the expected loss over all possible outcomes. This expectation over the data yields what is called the *risk* of the estimator,  $R(\theta, \hat{\theta})$ :

$$R(\theta, \hat{\theta}) = \sum_{\vec{D}} \mathbf{P}(\vec{D} | \theta) \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \equiv \mathbb{E}_{\mathbf{P}(\vec{D}|\theta)} \left[ \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \right] \quad (2.27)$$

, where we have introduced the  $\mathbb{E}_{\mathbf{P}(-)}$  notation for the expected value (average) over a distribution  $\mathbf{P}(-)$ . We will discuss expectations in more detail in the next subsection, but for now the notation will mostly be used to symbolize ideas that are described in the text.

We have removed the dependence on the data by using the likelihood to transfer it to  $\theta$ , the real parameter. This  $\theta$  is also an argument to the loss function itself. We don't have a clear-cut estimate of this parameter in advance, so we can't exactly pinpoint the values of these quantities (if we did, we would choose the estimator  $\hat{\theta} = \theta$ , making the problem trivial and removing the need for this discussion entirely). But we *do* have our best-effort estimate, embodied in our prior. So to remove this unwanted variable, we just integrate it out, getting the expectation of 2.27 over  $\theta$ :

$$r(\mathbf{P}(\theta), \hat{\theta}) = \mathbb{E}_{\mathbf{P}(\theta)} \left[ R(\theta, \hat{\theta}) \right] = \mathbb{E}_{\mathbf{P}(\theta), \mathbf{P}(\vec{D}|\theta)} \left[ \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \right] \quad (2.28)$$

$$= \int \sum_{\vec{D}} \mathbf{P}(\vec{D} | \theta) \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \mathbf{P}(\theta) d\theta \quad (2.29)$$

, where the innermost distribution in the subscript is the distribution over which to take the expectation first.

We will henceforth omit the dependence of this quantity on the prior and simply write  $r(\hat{\theta})$ , to avoid overcrowding equations; the use of a flat prior distribution within some bounds (which dictate the integration domain) can be assumed, as this strategy will be adopted throughout the document.

The removal of our last unwanted variable  $\theta$  in equation 2.29 by averaging over the prior then yields  $r(\hat{\theta})$ , which is called the *Bayes risk*. It depends on how the estimator function  $\hat{\theta}$  maps data to estimates. Once a loss function is chosen, the *Bayes estimator* is the one which minimizes the the Bayes risk, or equivalently the posterior expected loss; this expectation is supplied by the prior (along with the likelihood function).

To say that the risk represents the posterior expected loss may be somewhat confusing. It is intuitive that it should be so, since the posterior effectively encodes our results, and it is these results' expected performance we are concerned with. However, at no point did the posterior come up during the discussion. As a matter of fact, we took the expectation of the risk over the prior. So where exactly does the posterior come in, and what do we really mean by *posterior expected loss*?

The key point resides in the step represented by equation 2.27, which takes place before the integration of 2.29. After taking the expectation of the loss over the data, when performing the integration the distribution already resembles our posterior more than our prior. This is made clear by some algebraic

manipulation:

$$r(\hat{\theta}) = \int \sum_{\vec{D}} \mathbf{P}(\vec{D} | \theta) \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \mathbf{P}(\theta) d\theta = \sum_{\vec{D}} \int \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \cdot \mathbf{P}(\vec{D} | \theta) \mathbf{P}(\theta) d\theta \quad (2.30)$$

$$= \sum_{\vec{D}} \int \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \cdot \mathbf{P}(\vec{D}) \left( \frac{\mathbf{P}(\vec{D} | \theta) \mathbf{P}(\theta)}{\mathbf{P}(\vec{D})} \right) d\theta \quad (2.31)$$

$$= \sum_{\vec{D}} \mathbf{P}(\vec{D}) \int \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \cdot \mathbf{P}(\theta | \vec{D}) d\theta \quad (2.32)$$

, where we have recognized equation 2.20.

More compactly:

$$r(\hat{\theta}) = \sum_{\vec{D}} \mathbf{P}(\vec{D}) \mathbb{E}_{\mathbf{P}(\theta|\vec{D})} \left[ \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \right] = \mathbb{E}_{\mathbf{P}(\vec{D}), \mathbf{P}(\theta|\vec{D})} \left[ \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \right] \quad (2.33)$$

, where again the expectations are taken in right-to-left order to remove all dependence in both  $\theta$  and  $\vec{D}$ .

That is, the Bayes' risk of an estimator is sum of its posterior expected loss conditional on every possible dataset, weighted by the respective datasets' probabilities. We can obtain the posterior expectation as long as we condition it on a specific list of outcomes: we proceed just as we would if we had actually gathered these hypothetical data. As for the probabilities of the data, they are the likelihoods averaged over the prior;  $\mathbf{P}(\vec{D})$  is by definition of the marginal distribution (equation 2.17).

In hindsight, we can see that we did in equation 2.29, when we took the conditional (on the parameter) expectation over the data and then averaged over the prior, was equivalent to taking the conditional (on the data) posterior average and then the expectation over the data.

Comparing the subscripts for the expectations in 2.29 and 2.33, they look like they make up a Bayes' theorem of their own; and the proof of their likeness (2.30) was indeed achieved by means of that theorem. This ties into the relation between the the probabilities at hand (prior, posterior, likelihood, and marginal). The equivalence can be seen by taking

$$\mathbf{P}(\theta) \mathbf{P}(\vec{D} | \theta) = \mathbf{P}(\vec{D}) \mathbf{P}(\theta | \vec{D}) \quad (2.34)$$

, multiplying by the loss, integrating over  $\theta$ , and summing over  $\vec{D}$ . By linearity of integration, the order of the two last operations is irrelevant.

Having clarified our objective and terms, we can finally discuss specific estimators. A popular choice is the posterior mean. It minimizes a *quadratic* loss function - that is, large errors are more penalized than small ones, which tends to be appropriate in scientific applications.

$$\mathcal{L}^{(\text{quadratic})}(\theta, \hat{\theta}(\vec{D})) = \left( \theta - \hat{\theta}(\vec{D}) \right)^2 \quad (2.35)$$

In the one-dimensional (single parameter) case, the quadratic loss function reduces to the mean squared error. Clearly, for each possible dataset, the expectation value, or mean, minimizes the expected

mean squared difference (this can readily be checked by differentiation), which is then the posterior variance (or more generally, the covariance).

$$\mathbb{E}_{\mathbf{P}(\theta|\vec{D})} \left[ (\theta - \hat{\theta})^2 \right] = \mathbb{V} \left[ \mathbf{P}(\theta | \vec{D}) \right] \triangleq \mathcal{L}^{(\text{quadratic})}(\theta, \hat{\theta}(\vec{D})) \quad (2.36)$$

, where  $A \hat{=} B$  denotes  $A$  estimates  $B$ . Note the fixed  $\vec{D}$ , which is assuming an after data collection point of view. The fact that the Bayes' estimator is the mean is unchanged after marginalizing over *a priori* possible data records, because for all of them the mean minimizes the variance.

In general, these are the estimators we will use (the mean estimates the real parameter, and the variance its mean squared error); these tend to be popular choices in research.

A different possibility would be the median. It minimizes a *linear* loss function.

$$\mathcal{L}^{(\text{linear})}(\theta, \hat{\theta}(\vec{D})) = \left| \theta - \hat{\theta}(\vec{D}) \right| \quad (2.37)$$

The median divides the distribution into two even parts containing equal probability mass, i.e. it falls in the middle of the distribution. That's where the weighted average of the absolute distances to all other parameter locations (which is in this case the expected loss) is smallest.

Finally, yet another option would be the mode, which minimizes a *binary* (0/1) loss function.

$$\mathcal{L}^{(\text{binary})}(\theta, \hat{\theta}(\vec{D})) = \begin{cases} 0, & \text{if } \hat{\theta}(\vec{D}) = \theta \\ 1, & \text{otherwise} \end{cases} \quad (2.38)$$

It attempts to maximize the chance of correctness by taking the posterior maximum, i.e. the highest probability point. The expected loss will simply consist of the probability mass assigned to the region where the loss is non-zero,  $\hat{\theta}(\vec{D}) \neq \theta$ . For continuous distributions one can relax the equality condition by defining a tolerance (otherwise the contribution of a point estimate to an integral is always null, as the corresponding volume is infinitesimal), or use an alternative definition involving the Dirac delta construction.

When the mode is chosen as estimator, the estimation method is generally termed *maximum a posteriori probability estimation* (or *MAP* for short). If additionally a flat prior is used, the designation further simplifies to *maximum likelihood estimation* (*MLE*), for plain reasons. Clearly, these techniques are less general than their common framework, because they automatically produce a single point estimate rather than a whole distribution to pick from (and from where to retrieve additional information).

Of these three estimators, the mode is perhaps the one which most easily escapes the criticism directed at Bayesian inference, or at least much of it. This is owing to the fact that it converges to correct values even in cases where other estimates fail due to (for some choices of a prior) not meeting the premises for asymptotic correctness [39, 40]. However, it may also display inherently less robust behaviour [101], namely in the presence of noise. Finding the mode is finding the posterior maximizer, and optimization based strategies are prone to overfitting.

Since irregularities invalidating asymptotic correctness won't be of concern in our case studies, we will generally - as said before - use the posterior mean. In any case, this choice is not very critical if the posterior is uni-modal and relatively regular, since the three estimators we mentioned will fall within close proximity of each other.

Something that stands out is the fact that when we took a frequentist stance, *we never had this discussion*. We didn't need an estimator, because we automatically produced an estimate,  $\tilde{\theta}$ . Initially, it seemed like somewhat of an unnecessary complication not to directly get one. Arguably, it is a nuisance to have to extract an estimate, and that too after having meticulously built up a thorough state of knowledge. But also, *we meticulously built up a thorough state of knowledge*.

This contrast is effectively illustrated by figure 2, where a single line (identified by a a number) characterizes the frequentist's knowledge, but a full distribution is required to describe the Bayesian viewpoint. This exhaustive characterization may seem more fastidious, but it also provides a more holistic overview.

For example, let's see what happens if we add data to the distribution in figure 2. The aftermath is represented in figure 3, along with a line marking the frequentist guess based on the same data.




Figure 3: Bayesian posterior probability density for a binomial likelihood. The dashed and full vertical lines mark the real parameter used to generate the data and its frequentist estimate respectively. The same data vector contemplating 1000 trials was used for the Bayesian and frequentist calculations.

Comparing figures 2 and 3, we can see that the posterior has gotten *sharper* with the inclusion of extra trials. We can tell that our knowledge has been refined since the previous checkpoint, and that would not change even if we were unaware of how many trials each of the posteriors actually contemplated. Contrastingly, the immediate frequentist result continues to be reported by the same unsavoury single line; its position has solely been readjusted.

Put differently, the Bayesian framework inherently provides not only an estimate of the real parameter, but also one of our trust in it. Furthermore, the Bernstein–von Mises theorem (which we already referred to in connection with the prior) proves that both these results are asymptotically correct from a frequentist point of view, under some conditions. This is convenient inasmuch as it allows us to perform estimation without subscribing to an entire new philosophical doctrine.

More specifically, the aforementioned theorem states that the posterior converges to a normal distribution centered at an efficient estimator<sup>2</sup>. We can notice that in the simple case of figures 2 and 3 the posterior is approximately a Gaussian curve whose center coincides with the frequentist estimate.

Furthermore, this asymptotic form's variance is known, and it matches its frequentist analogue [17]. This holds provided that the prior fulfils some criteria; note that this statement is otherwise independent of the prior, in agreement with our previous claims.

This effectively harmonizes Bayesian and frequentist statistics in many contexts. In particular, it enables the approximation of frequentist quantifications of uncertainty by means of the Bayesian posterior, via the legitimization of its frequentist interpretation. It does so by proving that the two frameworks' measurements of uncertainty - expressed in terms of intervals or, more generally, regions -, mathematically match in the asymptotic limit (in spite of their conceptual disagreement). Summarily, it sustains the use of Bayesian means for non-Bayesian purposes.

Credible regions delimit a section of the parameter space where we believe the real parameter is contained to some level of certainty. This *belief* denotes Bayesian probability. So more specifically, a 95% credible interval would be one where our  $\theta$  is 95% likely to fall (given our state of knowledge).

In figure 3, this interval would be one over which the posterior probability (which, contrary to the probability density, is normalized by definition) integrates to 0.95<sup>3</sup>. Put another way, the confidence level is the fraction of probability mass bounded within the confidence interval.

As is evident, this definition is implausible from a frequentist perspective. The *distribution* over possible real values could be but a pulse on a single one, regardless of whether we know it or not. To speak of the *probability* that the true parameter lies within a designated region is nonsense: as a static entity, it either does or it doesn't. We can at most be more or less certain that one of these cases is true.

Hence in this paradigm the *credibility* of the parameters is replaced by the *confidence* in the estimate. The (so to say) *active* role shifts from the true variable to the interval. Instead of speaking of the probability that *the parameter is contained* in the interval, we speak of the probability that *the interval would have contained* the parameter. This *confidence* concerns the intervals we might have obtained *in general*, whereas the Bayesian credible region pertains to the particular one we got.

This key divergence is brought on by the fact that in the latter case we regard the interval as probabilistic, whereas in the former it is the parameter itself which we consider to be so. The origin of the observed

---

<sup>2</sup>This efficiency is, as before, relative to a loss function.

<sup>3</sup>We can observe that this condition does not uniquely define an interval. Several methods for picking one can be selected; this selection is akin to the designation of an estimator.

data's randomness is ascribed to either the nature of the measurement operation isolatedly, or also that of the parameter that generated it.

Such a parting is attributable to the fact that unlike frequentism, what Bayesianism seeks to describe is a subjective reality. This *subjectivity* - or lack thereof - dictates whether uncertainty is assimilated into the theory, or is simply an immaterial assessment pertaining to the reliability of its findings. Accordingly, in frequentist statistics our 95% confidence interval would be defined as one that would have contained the parameter with 95% probability. Now this is a frequentist probability: it is the limit of a relative frequency.

More specifically, by proceeding the way we proceeded to obtain our 95% confidence interval estimate, we could have arrived at a *correct* interval (i.e. one holding the true parameter) or not. It follows from definition of confidence intervals that the former scenario would occur 95% of the times if we continued to repeat the attempts eternally. In this case, the intervals are determined under a normality assumption (for the sample means). The desired confidence level then determines - along with some summary statistics - the width of the region, which is centered at the normal curve's peak.

For instance, a 95% confidence interval would cover 1.96 standard deviations to each side of the Gaussian's maximum. This is a property of the Gaussian, and can be calculated or looked up in a table. The only extra details needed to identify the interval are the the population variance  $\sigma^2$  (or an estimator) and sample size  $N$ .

We should bear in mind that these statistics respect the sample means, *not* a single sample. This is in line with our discussion about how confidence regards all of the possible end products our estimation could have yielded. That's why the sample size is required: the standard deviation of the sample means is approximated by  $\sigma/\sqrt{N}$  (in subsection 2.4.2 we will come across this again, when discussing shot noise). We need this factor to convert the interval width from units of standard deviations to our actual units.

These techniques occur within the scope of *interval estimation*, which improves upon the single point frequentist estimates discussed above to provide a more descriptive conclusion. One may then question whether any benefit is to be had from the Bayesian strategy. Wouldn't the discrepancy in thoroughness be redeemed by this stepping up of frequentist statistics?

That is not quite the case. First of all, we can observe that while it does allow for uncertainty estimation, it is not entirely as general as the Bayesian approach. A posterior distribution allows for computing the expectation value of nearly *any* function of the parameters.

In practice we can do this simply by sampling from the posterior<sup>4</sup>. Statistical functions, such as the mean or variance, are just particular cases of expectations. While distributions are completely described by these parameters when they are normal, in general they are not.

Secondly, the method itself is more universal: it requires only a parametric model and the ability to replicate it (for the likelihoods; this requirement will be presented more clearly in subsection 2.3.2). Moreover, its degrees of freedom are built into the algorithm by default, independently of the model's complexity. Contrarily, the frequentist process demands more structure. These characteristics again

<sup>4</sup>An overview of what this *simply* understates is left for chapter 3.

betray an overarching context of unmatched flexibilities.

As a consequence, it is often more viable to demarcate Bayesian credible regions than it is frequentist confidence regions [17]. It should then be clear why we would be interested in a theorem which substantiates their interchangeability. By stating that the credible intervals will converge to their frequentist counterpart, the Bernstein-von Mises theorem produces extra paths toward the computation of the latter.

## 2.2.2 Estimation in practice: approximate inference and conjugate priors

Just before subsection 2.2.1, we downplayed the role of the marginal probability, on the grounds of its being an unimportant normalization factor as far as parameter estimation is concerned. We will now return to pick up on that statement.

Let's begin by considering the computation of this constant,  $\mathbf{P}(D)$ . Recall that we defined it as:

$$\mathbf{P}(D) = \int \mathbf{P}(D | \theta)\mathbf{P}(\theta)d\theta \quad (2.17)$$

, and said that it represented the probability of the data  $D$  occurring independent of  $\theta$ .

In fact, it is the integral of the denominator in equation 2.5 over the parameter space. As the term marginal probability suggests, the variable  $\theta$  is integrated out (by marginalizing the conditional probability over the prior). For a flat prior (re-scaled to a normalized domain  $\theta \in [0, 1]$ ), it is also called the marginal likelihood.

More often than not, the most feasible way of computing integrals like 2.17 is through numerical integration. That is, we replace the infinite summation over infinitesimal volume units with a finite one, by selecting some set of points at which to evaluate the integral.

These points can be chosen deterministically (e.g. by arranging them on a dense enough grid) or randomly (as in Monte Carlo integration). Then we simply take the average:

$$\mathbf{P}(D) \approx \frac{V}{M} \sum_{i=1}^M \mathbf{P}(D | \theta_i)\mathbf{P}(\theta_i) \quad (2.39)$$

, where we have considered  $M$  evaluation sites, and  $V$  is the volume covered by the original integral. For instance, in our coin case we have a single-variable integral, so  $V$  is a length. Since  $\theta \in [0, 1]$ , we have that  $V = 1$ . Clearly,  $V/M$  is the finite unit of volume that we attribute to each point, which comes to replace the infinitesimal one  $d\theta$ .

We should note approximations such as 2.39 can be improved upon. This discussion is left for chapter 3, as it is extraneous to our current purposes; the analysis that follows holds regardless.

All that's left is to choose the  $M$  points ( $\theta$ s) at which to evaluate the function. This is one of the aspects that strongly affect the cost-to-performance ratio, but without much of a thought, we can simply choose them to be evenly spaced and as many as we're willing to spend computational resources with.

Actually, we already did something similar back in section 2.2, where we evaluated Bayes' rule on a grid to get the smooth curve of figure 2. In reality, our immediate result would look more like figure 4,



since there's no way to plot a continuous function directly (as much as we can bring neighbouring points close up to any desired precision).




Figure 4: Bayesian posterior probability evaluations for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector.

Let's again opt for that strategy, and evaluate approximation 2.39 using those same sites. With that, all details are fixed: we can finally normalize our distribution. That can be done by dividing the  $\mathbf{L}(\theta | D)\mathbf{P}(\theta)$  calculations by the approximation of  $\mathbf{P}(D)$  given by 2.39 to get the complete form of the Bayes' rule as per equation 2.9. Figure 5 shows the result, again represented with a smoothed line.

So finally, we have a normalized probability distribution, which integrates to one (asymptotically). Which seems like a good time to ask ourselves, *what exactly do we want that for?*

The short answer is, because that's how probability distributions work - they *have* to be normalized. But it is also not a very satisfactory one, and somewhat of a lapalissade - why should we particularly care that they are?

In the beginning of this section, we remarked that we could point out an estimate just from looking at the unnormalized probability density, and further noted that this construct sufficed for some applications. Where is it that this falls apart? We can do better by looking deeper into our motives.

In section 2.2.1, we talked about estimators of the real value and of the uncertainties, and suggested we'd adopt the mean and the variance respectively. These are two particular cases of expectation values that we can extract from a posterior distribution, much like we extracted  $\mathbf{P}(D)$  from its unnormalized version. In general, we can calculate the expected value (or equivalently, the mean) of any function of the




Figure 5: Bayesian posterior probability distribution for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector.

parameters by integrating over the posterior distribution:

$$\mathbb{E}_{\mathbf{P}(\theta|D)} [f(\theta)] = \int f(\theta) \mathbf{P}(\theta | D) d\theta \quad (2.40)$$

Of course, this holds for *any* probability distribution, not just the posterior. We could replace the occurrences of  $\mathbf{P}(\theta | D)$  in equation 2.40 by some other distribution, such as our prior, to get the corresponding expectation values.

Ultimately, *this* is what characterizes a probability distribution: the behaviour that it induces, described in terms of expectations. For the mean  $\bar{\theta}$  we simply have  $f(\theta) = \theta$  (hence the interchangeability of the terms *expectation* and *mean*: this definition can be generalized to functions other than the identity), and for the variance  $f(\theta) = (\theta - \bar{\theta})^2$ . The fact that the two metrics that are most commonly used to describe a distribution consist of a particular case of expression 2.40 only comes to underline its centrality<sup>5</sup>.

Thus, the foremost goal of the inference process is to enable the computation of integral 2.40. This unlocks all the information we may want, which takes the form of estimates. And interestingly, *we already know* what that computation entails in practice, because all that was said in the beginning of this subsection about approximating the marginal probability 2.17 holds here as well. So as before, we approximate the integral in 2.40 as:

$$\mathbb{E}_{\mathbf{P}(\theta|D)} [f(\theta)] \approx \frac{V}{M} \sum_{i=1}^M f(\theta_i) \cdot \mathbf{P}(\theta_i | D) \quad (2.41)$$

<sup>5</sup>Of course, if one seeks to find the posterior maximizer, single point evaluations are to be employed in numerical optimization rather than in computing expectations. This is one of the cases where normalization is not required. However, averages enjoy a theoretical soundness that optimization does not. This ties into the discussion of chapter 3.

Once more, we have replaced the distribution by single point evaluations; these can be the same grid points as before. Which raises the question, what do we need *the rest* of the distribution for?

If our goal is to compute equation 2.40 and we do so through equation 2.41, we don't need it at all. For this reason, the distribution can be more conveniently expressed discretely, as a *probability mass function* with support over a finite set of points.

In this case, normalization demands that the points *sum* to one, as opposed to *integrating* to one. So instead of using the exact equation 2.9, we will apply a discretized version of it to each of our  $M$  grid points  $\theta_i$  ( $i \in \{1..M\}$ ).

First we sweep of all these points and attribute them unnormalized *probabilities* (this is clearly an abuse of terminology, because there is no such thing). These tell of their *relative* likelihoods given the data. This simply amounts to repeatedly applying the formula:

$$W_i(D) \equiv \mathbf{L}(\theta_i | D)\mathbf{P}(\theta_i) \quad (2.42)$$

We have represented with  $W_i(D)$  the unnormalized *probability* of  $\theta_i$ , in anticipation of the notation to be used in the following chapter.

$$W_i(D) \propto \mathbf{P}(\theta_i | D) \quad (2.43)$$

Equations 2.42 and 2.43 essentially make up proportionality 2.16, but with added emphasis on the discretization.

As we do this for all the  $M$  grid points, we successively sum the results that we get for each. This normalization constant is to be used for all the points, so one need only calculate it once. Finally, we go back to get the normalized probabilities for all points, by dividing their respective values of equation 2.42 by this common factor:

$$w_i(D) \equiv \frac{W_i(D)}{\sum_{j=1}^M W_j(D)} = \frac{\mathbf{L}(\theta_i | D)\mathbf{P}(\theta_i)}{\sum_{j=1}^M \mathbf{L}(\theta_j | D)\mathbf{P}(\theta_j)} \quad (2.44)$$

At this point, it is clear why normalization substitutes the calculation of the marginal probability: *because it precisely amounts to calculating it*, apart from a constant factor  $V/M$ . This can readily be seen by comparing equation 2.39 with 2.44's denominator.

$$\mathbf{P}(D) \approx \frac{V}{M} \sum_{i=1}^M W_i(D) \quad (2.45)$$

A single issue remains: the omission of this constant factor. But if one might be tempted to ask why we would leave it out, a more appropriate question would be, *why would we not?*

As it turns out, this constant also appears in expression 2.41, which we have deemed *our ultimate goal*. So if we were to include it, we would constantly be *dividing* equation 2.44 by it (multiplying its quotient), only to then *multiply* it by the same factor in equation 2.41.

In the end, this *oversight* is rather opportune, in the sense that assimilating the factor into the formalism simplifies the expectations we draw from it. Coincidentally (or not), it does so all the while imposing the familiar condition of normalization.

With this, when working with a a probability mass function, expectations take the form:

$$\mathbb{E}_{\mathbf{P}(\theta|D)} [f(\theta)] = \int f(\theta)\mathbf{P}(\theta | D)d\theta \approx \sum_{i=1}^M f(\theta_i) \cdot w_i(D) \quad (2.46)$$

This is in essence a discretized version of equation 2.40, which formalizes the approximation we considered in 2.41 by transitioning into an intrinsically discrete setting. Note that this is correct provided that the  $w_i(D)$  are calculated according to 2.44, whereas in the continuous domain  $\mathbf{P}(\theta_i | D)$  represented single point evaluations of equation 2.9. The difference is that now our normalization constant is an estimate of

$$\frac{1}{P(D) \cdot M/V} = \frac{V}{M} \cdot \frac{1}{P(D)} \quad (2.47)$$

, which as we've seen takes care of the approximation in full, leaving only the target function evaluations.

Equivalently, we can describe the distribution using Dirac deltas to place *impulses* at the grid points. These will map the expectation integrand  $f(\theta)\mathbf{P}(\theta | D)$  at each chosen location  $\theta_i$  to the corresponding evaluation at that point  $f(\theta_i)w_i(D)$ , while disregarding all others - those (infinite)  $\theta_X$ s for which  $\forall i \in \{1..M\} : \theta_X \neq \theta_i$ . This converts 2.40 to 2.46 directly.

$$\mathbf{P}(\theta | D) \approx \sum_{i=1}^M w_i(D) \cdot \delta(\theta - \theta_i) \quad (2.48)$$

When we are working in the discrete domain, we typically represent probabilities not with scattered points as in figure 4, but rather as the visually clearer alternative of a stem graph, as in figure 6.




Figure 6: Bayesian posterior probability mass function for a binomial likelihood. The dashed line marks the real parameter used to generate the data. 100 trials were included in the data vector.

In contrast with figure 5, the probability mass function in 6 has been normalized in accordance with 2.44, which prompts a re-scaling of the y-axis by  $V/M$ . As said before, in our case  $V = 1$ , and  $M = 100$  points were used; these features materialize as a compression by a factor of 100.

With this we conclude our probe into approximate inference for now. In chapter 3 we will take a more in-depth look at these methods, and elaborate on how their foundations can be built upon so as to outperform the basic grid arrangements we considered.

We will shortly proceed to discuss more realistic applications for the matters of this exposition, but before that one last topic merits a mention.

We assumed from the outset that we *needed* to approximate the distributions, due to the obvious memory and power constraints. This was a sensible line of reasoning, and it *is* the most general one. Nonetheless, it isn't necessarily the *only* one: in some particular cases, the posterior is available analytically, dispensing with the need for numerical calculations entirely.

This happens for some classes of likelihood models, when paired with specific priors. For a few combinations of these two distributions, the posterior distribution is known to belong to the same family as the prior. In these circumstances, we call the latter a *conjugate prior* for the likelihood function under consideration.

Such requirements may seem restrictive, because they are. Even so, a considerable amount of commonly occurring combinations of functions enable this approach, and its perks are enticing: we may abandon the resource-intensive techniques outlined above altogether, and simply take a known closed-form solution for a fraction of the effort.

This is made possible by the fact that we know in advance we can describe the posterior using the same parametric model as the prior. As such, we can generally compute the parameters that fully determine it efficiently. What is more, *these parameters alone* contain all the information we could possibly want.

Clearly, these new parameters are distinct from the one(s) we are targeting with the learning process. They're just auxiliary constructs, used to characterize the consecutive distributions that depict our knowledge of the parameter that governs the probabilistic phenomenon under investigation (the  $\theta$  from before). For this reason, and to avoid confusion, these somewhat tangential parameters are commonly termed *hyperparameters*. For example, if we had a gaussian prior distribution, the hyperparameters would be the mean and variance; the parameter would simply be the independent variable.

We can think of the prior distribution of being molded by experimental results, yielding the posterior, whose shape has been altered to accommodate the data. This is done by tweaking the existing degrees of freedom; the extra structure that conjugate distributions provide reduces these degrees of freedom into a few variables. This enormously facilitates information storage and processing, because it is much easier to characterize a finite set of variables than it is a continuous distribution.

A well-known example is the normal distribution, which is self-conjugate. So if we start with a Gaussian prior, and the likelihood is Gaussian itself, then so too is the posterior.

Another case is when we have binomial likelihood as in our coin example, and the prior is a beta distribution. A beta distribution is described in the interval  $\theta \in [0, 1]$  by a probability density function

which depends on two (hyper)parameters,  $\alpha$  and  $\beta$ :

$$\mathbf{Beta}(\alpha, \beta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha, \beta)} \quad (2.49)$$

, where  $B(\alpha, \beta)$  is the beta *function*, which consists of the integral of the denominator and acts as a normalizing constant.

Let's say we start with some hyperparameters  $\alpha_0, \beta_0$ . It can be checked by direct calculation (by plugging equations 2.24 and 2.49 into the Bayes' rule 2.9 and rearranging) that this yields a posterior given succinctly by:

$$\mathbf{P}(\theta | D) = \mathbf{Beta}(\alpha_0 + s, \beta_0 + (N - s)) \quad (2.50)$$

It is convenient to remember from the probability density function in 2.24 that  $s$  is the number of successes (1 outcomes), and  $N$  is the total number of trials (so  $(N - s)$  corresponds to the number of failures, or 0 outcomes).

That means that all we have to do is count the number of 1 (or tails) outcomes and keep track of the total number of coin tosses. Here our data can be summed up by two constants:  $D = \{N, s\}$ . In the end, we will have perfect understanding of our posterior, just by updating the hyperparameters according to these figures.

Of course, not all is perfect, and we find ourselves constricted in the choice of a prior, which must be a beta distribution itself (aside from the restriction we already placed on the likelihood). But often, and invoking if necessary the Bernstein–von Mises theorem, we can adjust the prior into taking the desired form; if so, this is a small price to pay considering the unparalleled convenience that it begets us.

In the case of the coin, we had picked a uniform prior, given by equation 2.25. As luck would have it, that happens to be precisely a beta distribution, with parameters  $\alpha = 1, \beta = 1$ . This cognizance completes the necessary conditions for us to be two integer sums and a subtraction away from full possession of the exact posterior distribution.

The posterior beta distribution is plotted in figure 7, along with superposed discrete evaluations of the posterior using Bayes' equation directly (these are the same points as in figures 4 and 6).

So not only does this allow us to forgo all the calculations laid out in section 2.2.2, but also yields an improved accuracy. Instead of performing  $M$  likelihood evaluations and subsequent normalizations to get an approximate result, we perform a couple of elementary operations *and* get a faultless one.

As a consequence, instead of the expectation values being approximated by equation 2.46, they have closed forms. Most quantities of interest, namely all the estimators we mentioned before, can be looked up in a table and depend on the hyperparameters  $\alpha$  and  $\beta$  alone.

Naturally, more instances of conjugate distributions exist. Furthermore, strategies reminiscent of this one are still an option where conjugate distributions are not applicable: the posterior - and consequently expectations - can be approximated variationally. This is realized by assuming that it can be described by a parametrized distribution of choice, and optimizing its (hyper)parameters with respect to some measure of




Figure 7: Bayesian conjugate posterior density for a binomial likelihood and a flat prior. The dashed line marks the real parameter used to generate the data, and the 'x's mark single point evaluations of the posterior. 100 trials were included in the data vector.

likeness to the actual posterior [16]. A similar selection was carried out implicitly in the case of conjugate priors, where the result was exact and determined by a quick data based calculation.

In general such an assumption does not hold, wherefore the choice of a model introduces a bias. The problem can be alleviated by introducing enough variational freedom through additional hyperparameters, which yield a more customizable family of distributions and thus enhanced expressibility. This further demonstrates the ever present tradeoff between exactness and computational load (or tractability).

In the sample-based approaches we covered earlier, asymptotic exactness was generally assured, but the finite approximation error could be appreciable if not enough resources were employed. With variational inference, the error is mostly due to the limited descriptive power of the imposed model, whose complexity in turn determines the required resources.

Should the main goal be to lighten the computational cost, the latter option is a befitting choice. However, it can heavily weigh down the achievable accuracy. For this reason, we shall content ourselves with this brief overview of these alternative approaches, and instead focus on the *general* approach.

Working with the latter will bring us the benefits of a broad-spectrum solution, which preserves in full the flexibility that makes Bayesian inference attractive in higher complexity scenarios. Admittedly, it also presents more diverse intricacies. We will cover some of them in chapter 3; but still before that, we ought to examine one of the contexts that give rise to such scenarios.

## 2.3 Application to quantum systems and general likelihood models

Finally, we would like to apply this inference technique to something a bit more useful than coin examination. By itself, determining the fairness of a coin is not the most rewarding endeavour. But this is where we benefit from the generality that we have claimed Bayesian inference enjoys: owing to it, measuring a qubit is essentially glorified coin tossing. On that account, we can characterize a quantum system much like we did a simpler one in section 2.2.

This idea has been explored both theoretically and experimentally; here we will mostly follow [51, 33, 36], but similar techniques have been approached in many other references [35, 116, 117, 100, 113].

The rules are now set by the Born rule, which yields the probability of finding a qubit in either basis state given its wavefunction. It states that the probabilities of these states are given by the squared modulus of their respective (complex) amplitudes. So if the wavefunction is given by

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad (2.51)$$

, then when measuring we will get each of the two outcomes with probabilities

$$P(0) = |\alpha|^2 \quad (2.52)$$

$$P(1) = |\beta|^2 \quad (2.53)$$

As one might expect, imposing that  $\Psi$  be normalized renders these two values redundant.

$$P(1) = 1 - |\alpha|^2 \quad (2.54)$$

This observation alone gives us a model quite alike the one of equation 2.14.

$$\mathbf{L}(\theta | D; M_{\text{coin}}) = \theta^D (1 - \theta)^{1-D} \quad (2.14)$$

$$\mathbf{L}(\alpha | D; M_{\text{qubit}}) = |\alpha|^{2D} (1 - |\alpha|^2)^{1-D} \quad (2.55)$$

, where the annotations  $M_{\text{coin}}$  and  $M_{\text{qubit}}$  clarify which system each likelihood model pertains to.

We can point that one retrieves not the underlying parameter (amplitude)  $\alpha$  from this information, but only its norm; this is of course a consequence of the projective measurement postulate of quantum mechanics. Aside from that, to shift from determining a coin's fairness to characterizing a pure quantum state, we need only identify the parameter  $\theta$  with the squared norm of one of the wavefunction coefficients. This is because performing projective measurements on these systems is very much like sampling from a binomial distribution.

It is worth noting that quantum systems do *not* in general behave like classical probabilistic phenomena. In fact, that they do not is ultimately attributable to the amplitude-versus-probability matter mentioned above. That the amplitude's phases at the time of measurement bear no impact the outcome does not



mean that they aren't of consequence - which they are, due to the interference they give rise to as long as we're working in the quantum regime.

In the instant we measure in the computational basis, writing the wavefunction as in 2.51 tells the whole story as far as measurements are concerned. But leading up to that, the amplitudes can interfere with one another, be it constructively or destructively, to produce very different results than a weighted sum of the isolated basis states would have.

If our formulation seems unaffected by this, that's simply because we are implicitly describing the end result through the *final* coefficients  $\alpha$  and  $\beta$ , whose phases don't affect the outcome. Notwithstanding, all the ones coming before play a role in the algebra that produces it.

Moreover, this isn't the only subtlety that these coefficients hide. When speaking of coins we were targeting a constant variable, but that's no longer the case. The wavefunction could depend, in addition to the parameter(s) we're trying to find, on some external controls, such as the time of measurement; we could even change the measurement basis, which would condition the outcomes as well. In general, we could have some set of experiment controls  $E$ , such that:

$$\mathbf{P}(0) = f(\theta; E) \quad (2.56)$$

We will from now separate random variables (here  $\theta$ ,  $D$ ) and non-random ones (here  $E$ ,  $M$ ) with a semicolon to avoid confusion; commas will be used to split arguments within the same *class*. We prefer not to use the form  $f(\theta | E, M)$ , or in general  $f(\theta | V)$  for  $V$  a deterministic entity. While in English we might say that probabilities are *conditioned* (dependent) on the experiment and model, we reserve the idea of conditioning to denote a dependence on *random* variables, which are held fixed when taking conditional probabilities. Accordingly, we destine the vertical bar to standing between two random variables/events.

If the distribution of outcomes is given as in 2.56, we must accommodate for it by adjusting the structure of the likelihood function we use to model the outcome distribution.

$$\mathbf{L}(\theta | D; E, M) = f(\theta; E)^D (1 - f(\theta; E))^{1-D} \quad (2.57)$$

The function  $f$  is linked to the particular model  $M$ , mapping parameters to their associated conditional probabilities, given the experiment(s) and under  $M$ . We again separate random and non-random parameters using a semicolon, remembering that  $\theta$  itself is treated as a random parameter.

Now, equation 2.9 too must be modified, to account for the likelihood's dependence on some experiment controls.

$$\mathbf{P}(\theta | D; E) = \frac{\mathbf{L}(\theta | D; E) \mathbf{P}(\theta)}{\mathbf{P}(D; E)} \quad (2.58)$$

We can also take the occasion to in addition make the model explicit, as we did to distinguish 2.55 from 2.14 and in equation 2.57 to mark the use of the function  $f$  in the left-hand side.

$$\mathbf{P}(\theta | D; E, M) = \frac{\mathbf{L}(\theta | D; E, M) \mathbf{P}(\theta; M)}{\mathbf{P}(D; E, M)} \quad (2.59)$$

The prior too signals dependence on the model, as it is itself an implementation choice; plus, the starting point for the parameters depends on the model's structure. Two different models may work on different scales, or even have different dimensions. The prior could be considered to depend on the controls too, but we assume that dependence to be absorbed into the model.

Finally, the marginal probability, being the integral of the nominator, also manifests a dependence on these new entrants. It is the prior probability of having obtained the datum  $D$  given a choice of experiment  $E$  and underlying model  $M$ .

$$\mathbf{P}(D; E, M) = \int \mathbf{P}(D \mid \theta; E, M) \mathbf{P}(\theta; M) d\theta \quad (2.60)$$

Of course, we don't necessarily have to consider data records of unit length, and we can condense longer data records into a single expression as we did in 2.20.

$$\mathbf{P}(\theta \mid \vec{D}; \vec{E}) = \frac{\mathbf{L}(\theta \mid \vec{D}; \vec{E}) \mathbf{P}(\theta)}{\mathbf{P}(\vec{D}; \vec{E})} \quad (2.61)$$

The data and experiment controls must have the same length and be written in the same order, that is,

$$\vec{D} = D_1, D_2, \dots, D_N \quad \vec{E} = E_1, E_2, \dots, E_N$$

means that experiment  $E_1$  was performed and yielded datum  $D_1$ , etc. (We could of course implicitly incorporate everything into the data vector, making it a vector of tuples.)

When estimating quantum parameters, likelihood evaluations as per 2.57 amount to finding the probability distribution of outcomes for an experiment, given a parametrized quantum system. The experiment  $E_j$  is a constant for each datum  $D_j$ , and we bind the parameters to a target value  $\theta_i$  at every evaluation to get  $\mathbf{L}(\theta_i \mid D_j; E_j)$  - or, if we prefer, we put the data together and consider  $\mathbf{L}(\theta_i \mid \vec{D}; \vec{E})$ . Either way, we do this for multiple  $\theta_i$  to approximate expectations as in equation 2.46.

Such a parametrized description of the target system's behaviour is indispensable for the inference process, which cannot take place without a likelihood model; the only degrees of freedom bestowed on the solution are the parameters. As such, the existence of a fixed model has been assumed all throughout this discussion, to the point of being implicit in our early version of Bayes' rule (e.g. equation 2.9).

However, such an assumption places severe restrictions on the scope of our analysis. In general we won't have one single, perfect model. What - if anything - can be done if we don't know the exact shape of the likelihood? Some alternative function  $g(\theta; E)$  could just as well replace the  $f$  in equation 2.57. Ideally, we would be able to tell which does best. Better yet, we would want to *refine* the best one by increasing its complexity and evaluating whether its aptness to fit to the data is enhanced. All of these possibilities rely on the ability to compare models.

So the question is, can this framework be extended to provide guidelines for model comparison?

The answer is, it doesn't need to be.

### 2.3.1 Brief detour into model selection

Not only is it possible to derive further generalization still within the Bayesian framework, but all the tools are already in our hands; we need only recognize and systematically apply them. More specifically, broader algorithms can be achieved by branching the analysis so as to encompass several models. We can then rely on model selection rules based on constructions we've already discussed to choose between them.

The usual approach is to develop evolutionary strategies that construct and compare increasingly complex models, and it has recently been applied to quantum systems [44]. An earlier such application, albeit somewhat *less quantum* in nature (in the sense that model selection takes place within a classical clustering problem) is presented in [49]

How can we perform systematic model comparison? Bayes' theorem again provides an answer. This selection criterion is used and briefly discussed in e.g. [44, 49]; a short but broader overview can be found in e.g. [34].

It should come as no surprise that Bayes' rule can be used to reason about models also, bringing them from the sidelines to occupy the central role that in equation 2.5 belonged to the parameters.

$$\mathbf{P}(M | D) = \frac{\mathbf{P}(D | M)\mathbf{P}(M)}{\mathbf{P}(D)} \quad (2.62)$$

We can then quantify the relative merits of any two models  $M_A$  and  $M_B$  by means of the ratio of their posterior probabilities - in the words of [64, chapter 5], the proportionality of their chances:

$$\frac{\mathbf{P}(M_B | D)}{\mathbf{P}(M_A | D)} = \frac{\mathbf{P}(D | M_B)\mathbf{P}(M_B)}{\mathbf{P}(D | M_A)\mathbf{P}(M_A)} \quad (2.63)$$

For an objective assessment of their comparative predictive power, we choose  $\mathbf{P}(M_A) = \mathbf{P}(M_B)$ , getting:

$$\mathcal{B}_{B>A} \equiv \frac{\mathbf{P}(D | M_B)}{\mathbf{P}(D | M_A)} \quad (2.64)$$

, where ">" denotes superiority.

The ratio in 2.64 is termed a Bayes factor. It is called a factor because it transforms the prior odds into posterior odds. Notice that this quantity is entirely determined by the data: we make no *a priori* distinction between models.

The magnitude of the ratio 2.64 can be interpreted as how strongly the evidence supports one model (here  $M_B$ ) over another (here  $M_A$ ). Its significance can be gauged in a similar way to statistical hypothesis testing; scales have been proposed in [64, appendix B] and later [65]. In [49] this is applied to a quantum characterization example.

As to how we can compute it in practice, that's a discussion we've already had back in subsection 2.2.2. The elements of 2.64 are simply the marginal probabilities of the data over each models' parameter space, as in equations 2.17 (where the model was implicit) and 2.60 (where a single model was considered, and we also explicitly marked the experiment  $E$ ):

$$\mathbf{P}(D; M_X) = \int \mathbf{P}(D | \theta_X; M_X)\mathbf{P}(\theta_X; M_X)d\theta_X \quad (2.65)$$

This is where Bayes factors differ from traditional likelihood ratio tests, often used in hypothesis testing: maximization is replaced by marginalization over model parameters. This is convenient for two reasons. First, the marginal is at hand by virtue of the inference process (this has been illustrated in 2.2.2 and will be discussed further in chapter 3). Second, we are exempted from optimization, which in addition to the cost is more susceptible to problems affecting generalization; integration tends to produce a more robust metric, in particular if noise is an issue.

Lastly, it naturally penalizes superfluous model complexity, privileging models with less structure if all else is equal [65, 110]. This is opportune, given that model selection shouldn't unreservedly aim for the best goodness-of-fit, but rather factor in parsimony as well, seeking a happy medium. If we don't enforce the principle of parsimony (also known as Occam's razor), we risk overfitting owing to undue flexibility. This is related to the bias variance tradeoff and the idea of *capacity* that is prevalent in machine learning.

Other approaches, such as the Akaike or Bayesian information criteria, artificially work in a penalty to disfavor an increase in parameters [34]. By-default retention of the null hypothesis in frequentist hypothesis testing is another example of imposed bias against excessive structure [110].

In light of this, Bayes factors' inherent propensity towards simpler models has earned them the epithet of *automatic Occam's razor* [65]. But where exactly does it come from?

It stems, of course, from the nature of the integral in 2.65. We can shed some light on the reason why by taking a step back and looking at what it stands for: instead of taking isolated likelihoods as a metric, we take their expectation value over the full prior.

This doesn't favor high likelihood densities per se, instead weighing them against space occupation. Should the parameter space be unnecessarily voluminous, a large part of it will have very low likelihoods associated. If so, the corresponding model will be penalized when taking the expectation in 2.65 as compared with a more frugal one which uses volume more sparingly. The gratuitous expansion of model space size for little or no benefit in terms of explanatory power provokes a penalty.

Put differently, the data tend to become more sparse relative to the working region as the latter expands. So if we add one parameter to our model, thereby increasing its dimension, this extra degree of freedom must actively contribute towards the interpretation of the data. If it does not, and the data fit is no better than it was before, its addition is unwarranted. This fact will be penalized in the Bayes factor analysis, where we take the average over a more *diluted* prior. As a consequence, we will be led to keeping the simpler one, which does as good a job .

We can illustrate this with a simple example. Let's say we have a black-box that discards any inputs (i.e. the controls are irrelevant) and produces 0 or 1 outcomes. We start by proposing a model  $M_1$ , with a likelihood of the form

$$\mathbf{L}(\theta \mid D; M_1) = \begin{cases} D, & \text{if } a \leq \theta \leq b \\ 1 - D, & \text{otherwise} \end{cases} \quad (2.66)$$

, with  $a, b$  fixed and again  $D \in \{0, 1\}$ .

We can interpret this as the description of a black-box controlled by some secret parameter. If the parameter lies between  $a$  and  $b$ , it always outputs 1; if it doesn't, it always outputs 0.

Still, this doesn't fully determine our model; we need to define a domain for the parameter. Instead, we write:

$$\mathbf{L}(\theta \mid D; M_1) = \begin{cases} D, & \text{if } \theta \in [a, b] \\ 1 - D, & \text{if } \theta \in [0, a[\cup]b, W[ \end{cases} \quad (2.67)$$

We furthermore assume a flat prior, as usual.

$$\mathbf{P}(\theta; M_1) = \frac{1}{W} \quad (2.68)$$

Here  $a, b, W$  are constants with  $0 \leq a \leq b \leq W$ . The total prior width  $W$  can be regarded as a hyperparameter that concretizes  $M_1$ .

Let's suppose our first call to the black-box produced a 1, and we want to tune our model so that it reproduces this behaviour. In this case, the support of the posterior is confined to the interval  $[a, b]$ , as this is the only region with non-null likelihood for  $D = 1$ . This arrangement imitates what usually happens when performing inference: most of the posterior probability mass is contained within a region (a credible region, as discussed in subsection 2.2.1), whose volume ideally decreases with the data collection if we fix the probability it should amass.

Recalling equation 2.65, we can easily compute the marginal, or the *evidence*, for this model:

$$\mathbf{P}(D = 1; M_1) = \frac{b - a}{W} \quad (2.69)$$

We can see that if we increase  $W$ , unnecessarily expanding the parameter space for no net increase in descriptive power for the data (in this case, datum), the marginal will decrease, effectively marking the model as *less worthy* (figure 8).

What if go further, and needlessly add *an entire dimension*? We can lift the one-dimensional model up to two dimensions by taking the natural inclusion map of the likelihood 2.67 along each dimension and intersecting both (figure 9).

More formally:

$$\mathbf{L}(\vec{\theta} \mid D; M_2) = \begin{cases} D, & \text{if } \{\theta_1, \theta_2\} \in [a, b]^2 \\ 1 - D, & \text{if } \{\theta_1, \theta_2\} \in [0, W[^2 \setminus [a, b]^2 \end{cases} \quad (2.70)$$

This is at least as expressive as the previous model. The posterior is symmetric relative to a  $45^\circ$  inclination line, and we could recover the previous posterior by projecting onto either of the axis the directional maxima orthogonal to it.

However, despite our having doubled the number of parameters, this doesn't fit the datum any better. As a matter of fact, the one-dimensional model already rendered a perfect description of what transpired,




Figure 8: Illustration of the model  $M_1$  from equation 2.67 for different values of  $W$ . The Bayes factor of  $M_1^{(ii)}$  relative to  $M_1^{(i)}$  is  $\mathcal{B} = 1/2$  for  $D = 1$ .




Figure 9: Illustration of the model  $M_2$  from equation 2.70 as compared to its one-dimensional projections. Two such projections can be viewed on a scale in figure 8 for different values of  $W$ . The space forms a square with side  $W$ .

because any point with non-zero posterior probability would have reproduced it with a 100% chance. Thus, increased dimensionality could never be of benefit to us.

We can see how this impacts the model's worth as measured by Bayes' factors. The same choice of a uniform prior now gives us

$$\mathbf{P}(\vec{\theta}; M_2) = \frac{1}{W^2} \tag{2.71}$$

, and the evidence becomes:

$$\mathbf{P}(D = 1; M_2) = \left( \frac{b - a}{W} \right)^2 \quad (2.72)$$

Other than in the scenario where all possible parameter specifications behave in the exact same manner (in which case all models are equally bad, because inference is useless), the two-fold increase in parameters brings a quadratic decrease in the evidence in favor of the model. A 2-dimensional model with a projection matching 2.67 could only have as large a marginal probability as the latter if the likelihood was independent of the additional parameter (with support on a vertical or horizontal strip, as the lighter colored sections of figure 9).

In the case of 2.70, the ratio is:

$$\mathcal{B}_{M_2 > M_1} = \frac{b - a}{W} \leq 1 \quad (2.73)$$

Under these circumstances, we would reject the redundant parameter. This is due to the fact that the prior probability density is stretched thinner relative to the data-appointed plausible region, as a result of the tension between probability mass and parameter space volume. If we further generalize the likelihood to a cube and then hypercubes, the fraction of volume contained within the region will shrink exponentially with the dimension.

Expectably, the decisions are more nuanced in realistic scenarios; this is an extreme example. Notwithstanding, it mimics the inner workings of Bayesian model comparison, and shows how it attributes higher scores to models which make better use of the space they inhabit. Encouraging simplicity in this manner is likely to improve generalization, making the winning model more apt to predict unseen data.

### 2.3.2 Likelihood evaluations in the quantum case

Once models are out of the way, we are left with equation 2.58. The prior is predetermined, the marginal probability can be regarded as a normalization factor - and lastly, we have the likelihood function.

Likelihood evaluations are the key element when performing Bayes' updates, and their most challenging requirement. It is the experimentally determined likelihoods that will shape our final beliefs, by reworking the prior ones in consonance with our measurement results.

Conceptually, they consist of the probabilities that some hypothetical parameters  $\theta_i$  would have responded like our target system did to our observations. That is, they correspond to the probabilities of sampling the conditional outcomes  $D_j \mid E_j$  that we did (or the joint probability of having obtained the full data record  $\vec{D} \mid \vec{E}$ ) had the (unknown) real parameters been  $\theta_i$ . This describes the task of quantum simulation.

Before proceeding to discuss this point, we observe that it settles the key differences between characterizing quantum system and characterizing a coin. A schematic overview of the inference process applied to the quantum case is presented in figure 10. It illustrates the context that the application of equation

2.58 in a quantum parameter learning problem demands. Similar diagrams (on which this one was based) can be found in [51] and [33].

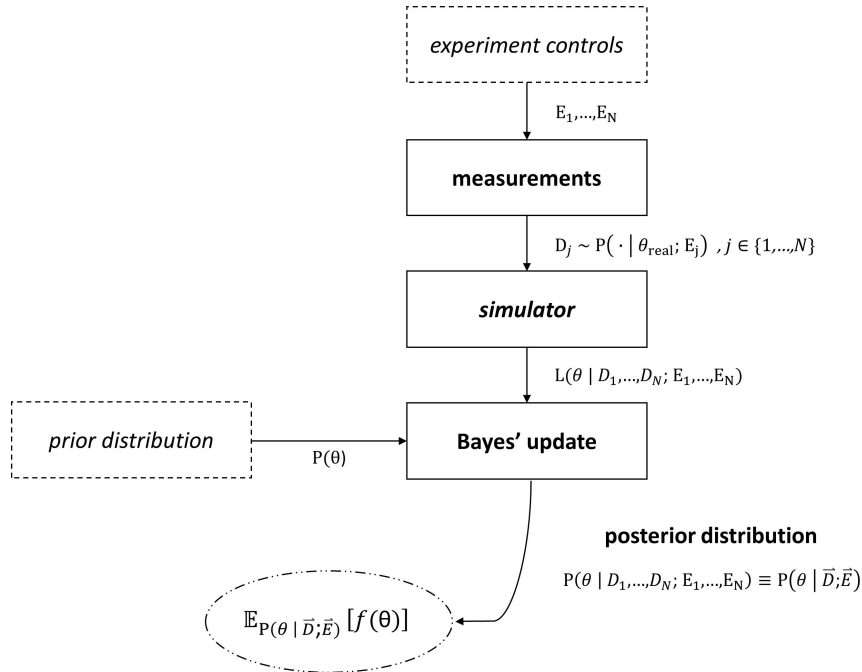


Figure 10: Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs. The tilde  $\sim$  means "is distributed as" and the dot  $\cdot$  is a stand-in for the object of the sampling.

It is sometimes possible to estimate the probabilities by means of analytic expressions, derived from quantum mechanics. If so, apart from the measurements on the system under study, the inference protocol is as straightforward to implement as it was in the case of our coin. This is feasible for small quantum systems which are simple to analyse; such is the case of the examples we will approach here. But retrieving measurement statistics from a generic (parametrized) Hamiltonian is more often than not a rather challenging task.

Still more generally, if what we mean to characterize is an open quantum system - a realistic assumption -, the Schrödinger equation is substituted by a master equation, whose operators and constants we must discover to characterize the system. Evaluating likelihoods at some instant (this instant being an experimental control) requires simulating the time evolution produced by our formalism until then.

One possibility is to simulate this evolution classically. However, classical resources are known not to excel at simulating systems which manifest their non-classical nature. For this reason, the calculations soon become intractable, as the size of the quantum system of interest increases and with it the size of the classical simulator required to replicate it - it being that the latter scales exponentially with the former.

Such circumstances prompt the use of quantum devices, which deliver our best known description of reality. This solution has been increasingly entertained in literature, ever since it was first advocated and



discussed [37, 74]. For our particular purposes, they are interesting in that they supply efficient estimates of the likelihoods.

Let us briefly recap where we stand. We have only now stated that we would like to reconstruct the likelihood function for some parameter specification and experiment. It should be remembered that this is preceded by the measurement of a *target* quantum system, which we have no control over. We mean to characterize that system; more specifically, we have a parametrized model describing it, and want to estimate the actual system parameters.

This can mean for instance that we are set on a fixed form Hamiltonian  $\hat{H}(\vec{\theta})$ , which we believe describes the system's dynamics and whose real parameters  $\vec{\theta}$  we would like to find. The description is to be concretized by establishing the numerical parameters  $\vec{\theta}_{\text{real}}$  up to some tolerance. This settles our problem statement. In this context, quantum parameter estimation is often called Hamiltonian learning.

All we need from the system is the ability to measure it (we would of course also appreciate it if it conformed to our model, but that responsibility falls on us). These measurements are described by experimental controls, such as the time, which we know give rise to a distribution of outcomes. We would like to know what distribution, because it is determined by the underlying parameters which we are attempting to estimate.

Equivalently, knowing the exact distribution would inform us about the relevant parameters, as illustrated by equation 2.56 or 2.57. Yet, all we get are individual samples - our data, which are comprised of measurement outcomes (paired with their respective controls). So we pick those samples, and we examine many possible configurations of our quantum model, expecting that the ones most likely to have behaved similarly under the same conditions correspond to the most accurate parametrizations. Conformity to the data *is* accuracy for all we care, because the ultimate goal is to replicate the system's behaviour. Such an achievement enables the estimation of its properties by means of a well-tuned simulator, including the anticipation of future observations based on past ones.

Of course, if we only perform *one* measurement, this doesn't convey very useful information. But as we gather more and more experimental data, we anticipate that it will. It also tends to perform better on limited data than just inverting expression 2.56 for a rough estimate of  $\mathbf{P}(0)$  would, and is a more flexible approach (further exemplification will follow later in the document).

It is when determining the merits of parametrization hypotheses that we need access to a quantum simulator. We want to mimic the real system; for that, we fill in the unknown (the parameters  $\vec{\theta}$ ), and instruct our simulator to evolve accordingly. If dealing with a closed system, the evolution is dictated by the Schrödinger equation: we must approximate the dynamics induced by  $\hat{H}(\vec{\theta})$ , e.g. by trotterizing it. The evolution time for the Hamiltonian dynamics must match how long the target system was allowed to evolve after some initialization (generally a projection).

Lastly, we replicate the exact same measurements we performed on our system, and assess how likely the outcomes would have been to match the registered data by repeating each and taking the relative frequencies. These constitute our likelihoods.

Note that, in spite of their likeness, there's an asymmetry between what is required of the target system

as compared to the simulator; from the former we get binary samples, and from the latter we are assuming we get exact (or close enough, but certainly not binary) probabilities.

These ideas match in the limit where we take many samples (measurements), because we may then reconstruct the full probabilities of an outcome as the ratio of its occurrence to the total measurements. Or, meeting at the other end, if the simulator produces a single outcome. For this reason, the terms *sampling* and *simulation* are often used interchangeably in the domain of statistics. If they are not, a too strict view of simulation is to blame, for their difference is but a matter of degree.

Likewise, when speaking of quantum simulation one generally distinguishes between *strong* and *weak* simulation; the latter describes precisely the drawing of samples from the probability distribution, as opposed to its full specification. Their difference is quantitative and not qualitative, in spite of this quantitiveness' being rather formidable still (and quite possibly the difference between the simulation being computable and not). This emboldens one to ask, could weak simulation be used in the context of Bayesian inference? It can, and it has been [35].

Still within quantum device based likelihood computations, more sophisticated strategies have been proposed, exploiting the structure of Bayesian learning to tune the protocol's degrees of freedom [116, 117]. Their underlying idea of adaptivity is related to the strategies to be covered in section 2.4.

This concludes our fundamental considerations on this topic. With them in hand, we can finally consider some real examples, corresponding to specific instances of the diagram in figure 10.

### 2.3.3 Quantum characterization examples: spin precession and phase estimation

The main references for this section are [51, 115], for the first and second illustrative cases respectively.

One simple yet elucidatory example is that of a single qubit oscillatory wavefunction whose frequency we mean to discover. This application has been widely studied, namely in [100, 33, 51, 36, 113].

We consider a qubit under the action of the following Hamiltonian:

$$\hat{H}(\omega) = \frac{\omega}{2} \hat{\sigma}_z \quad (2.74)$$

This class of Hamiltonians in particular is well-suited for the description of multiple physical phenomena of relevance, such as Larmor precession, Rabi flopping and Ramsey oscillations. We will elaborate further on chapter 5, section 5.1; for now, we will be content with understanding the behaviour it induces.

That can be obtained by expanding the time evolution operator corresponding to 2.74 in its Taylor series, grouping the terms into even and odd functions, recognizing each as the cosine and sine series and using the identity  $\hat{\sigma}_i^2 = \hat{I}$  for  $\hat{\sigma}_i$  in the Pauli group to get:

$$\hat{U}_\omega(t) = e^{-i\hat{H}t} = \cos\left(\frac{\omega}{2}t\right) \hat{I} + i \sin\left(\frac{\omega}{2}t\right) \hat{\sigma}_z \quad (2.75)$$

Clearly, such an operator has no effect on a z-eigenstate, leaving basis states  $|0\rangle$  and  $|1\rangle$  unchanged apart from a global phase. For other initial states however, it will impact the  $\hat{\sigma}_x$  and  $\hat{\sigma}_y$  expectations.

In general, we can apply this to some initial state  $|\psi(0)\rangle$ ,

$$|\psi(t)\rangle = \hat{U}_\omega(t) |\psi(t=0)\rangle \quad (2.76)$$

, and use the Born postulate to extract outcome probabilities:

$$\mathbf{P}(D | t) = \|\langle D | \psi(t) \rangle\|^2 \quad (2.77)$$

for  $|D\rangle$  some basis state corresponding to an observable's eigenstate.

A straightforward choice of a starting state is the x-eigenstate.

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad (2.78)$$

Since the identity maps it to itself and  $\hat{\sigma}_z$  is an inverter in the x-basis (mapping  $|+\rangle$  to  $|-\rangle$  and vice-versa), it is easy to see that:

$$|\psi(t)\rangle = \hat{U}_\omega(t) |+\rangle = \cos\left(\frac{\omega}{2}t\right) |+\rangle + i \sin\left(\frac{\omega}{2}t\right) |-\rangle \quad (2.79)$$

This describes a spin rotating about the z-axis. It can be visualized using the Bloch sphere (figure 11a): the Bloch vector lies on the equatorial plane, and precesses about the vertical axis. Naturally, there is nothing unique about the z direction; we could modify 2.74 to make the Bloch vector rotate around other axes.

The result is that the  $\hat{\sigma}_z$  expectation is constant (0), whereas those of  $\hat{\sigma}_x$  and  $\hat{\sigma}_y$  oscillate.

$$\mathbf{P}(+ | t) = \|\langle + | \psi(t) \rangle\|^2 = \cos^2\left(\frac{\omega}{2}t\right) \quad (2.80)$$

$$\mathbf{P}(- | t) = \|\langle - | \psi(t) \rangle\|^2 = \sin^2\left(\frac{\omega}{2}t\right) \quad (2.81)$$

Thus by measuring on the x-basis, we can infer  $\omega$ . (Equivalently, we could measure in the y-basis or others, preferably orthogonal to the z-axis.)

For more general states, the operator preserves the  $\hat{\sigma}_z$  expectation while inducing oscillations in the  $\hat{\sigma}_x$  and  $\hat{\sigma}_y$  components. This can most easily be seen by applying it to a general pure state in the Bloch form:


$$|\psi(0)\rangle = \cos(\theta_0/2) |0\rangle + e^{i\varphi_0} \sin(\theta_0/2) |1\rangle \quad (2.82)$$

$$|\psi(t)\rangle = \hat{U}_\omega(t) |\psi(0)\rangle = \cos(\theta_0/2) |0\rangle + e^{i\varphi_0 + \omega t} \sin(\theta_0/2) |1\rangle \quad (2.83)$$

The initial angle with the z axis is fixed at  $\theta_0$ , whereas  $\varphi(t) = \varphi_0 + \omega t$  (figure 11b). For simplicity, we will consider the simpler case of  $|+\rangle$  initialization.

All ingredients for inference are now set. We can identify the notation with the one used throughout the chapter:

$$\boxed{\mathbf{P}(\theta | D; E) \leftrightarrow \mathbf{P}(\omega | +; t) = \cos^2(\omega t/2)} \quad (2.84)$$



(a) Spin precession starting at the  $|+\rangle$  state. This illustrates the motion described by equation 2.79. Here  $0 < t_1 < t_2$ .

(b) Spin precession in the general case. The angle  $\theta_0$  is a constant, whereas the phase  $\varphi(t)$  is time dependent.

Figure 11: Bloch sphere representations of spin precession for a specific initial state and for a general one.

Here the parameter is the precession frequency, the data are whether the measurement found the qubit at state  $|+\rangle$  or  $|-\rangle$ , and the only experimental control is the evolution time that elapsed since the wave-function's collapse into the initial state. The likelihood in 2.84 contains all information, because the probability of  $|-\rangle$  is complementary to that of  $|+\rangle$ .

With this, everything is in place to learn  $\omega$ . An illustrative posterior for the likelihood in 2.84 is presented in figure 12. The times were chosen arbitrarily; we will get back to this choice later. For now, we'll be content with noting that the posterior looks similar to what it did in section 2.2 (e.g. figure 2)

Another interesting example is phase estimation. The goal of a phase estimation is to estimate the phase  $\phi$  of the eigenvalue  $\lambda$  of a unitary  $\hat{U}$ 's eigenstate  $|\psi\rangle$ . The two last elements are given.

$$\hat{U}|\psi\rangle = \lambda|\psi\rangle \quad (2.85)$$

$$\lambda = e^{i\phi} \quad (2.86)$$

The most well-known approach is to construct a circuit where we create superposition in an auxiliary register, apply to a second register initialized to  $|\psi\rangle$  a sequence of multiples of  $\hat{U}$  controlled on the said register ( $C\text{-}\hat{U}^m$ ), and process the resulting state with an inverse Fourier transform (inverse quantum Fourier transform (IQFT)). Then we measure to retrieve the phase information.

However, such a circuit is of prohibitive depth, considering the decoherence problems that plague today's (and tomorrow's) quantum devices. The main culprit is the Fourier transform, which requires a sequence of controlled rotations. This is doubly inconvenient, being a *sequence of controlled rotations*.




Figure 12: Bayesian posterior probability distribution for an oscillatory likelihood. The dashed vertical line mark the real parameter used to generate the data. 100 trials were contemplated in the data vector, using evenly spaced times on  $[0, 5]$ .

Each of these rotations must be decomposed into a sequence of elementary gates, whose length scales oppositely to the required precision [66]. These groups of decompositions are to be executed mostly in a series, for the duration of which the system must not lose coherence.

Moreover, the precision is determined by the number of available qubits - a scant resource in quantum devices, and one which brings additional troubles (yet one more source of error). The result is hardly a reasonable thing to ask of near-term quantum devices, whose very limited coherence times will join a plethora of other sources of experimental error to produce essentially noise at the output.

For this reason, alternative approaches to phase estimation have emerged, namely [iterative phase estimation \(IPE\)](#) algorithms. They are typically based on a shorter quantum sub-routine, trusting on some level of classical processing to replace the troublesome quantum processing (the [IQFT](#)).

This transfer of responsibility is achieved by shrinking the quantum circuit from a full algorithm to a fundamental measurement operation, which we plug into a classically controlled feedback loop. The typical quantum circuit used for this purpose is depicted in figure 13. For a single qubit, it matches the traditional phase estimation circuit.

Unsurprisingly, the circuit for iterative phase estimation resembles the key motif of its textbook version. We detach it from the circuit, alter it to our convenience - namely adding some degrees of freedom -, and feed it to a classical processor. This has a similar flavor to other [NISQ](#) era algorithms (such as the [QAOA](#)), where we take a shallow parametrized circuit that loosely reproduces the pattern of an exact quantum algorithm and throw in classical resources to (hopefully) compensate. The goal is to somehow extract




Figure 13: Iterative phase estimation circuit diagram.

the quantum advantage of the fundamental building blocks while placing most of the computational load elsewhere.

More specifically, the attractiveness of these iterative phase estimation protocols stems from three points. First, they are based on several more shallow circuits rather than one long one, which is beneficial when decoherence is a problem. Second, their precision is generally a function of the number of iterations rather than that of ancilla qubits. And lastly, they allow for the reduction of the ancillary register at the expense of extra rounds. This helps remedy the issues due to large circuit dimensions (both vertical and horizontal).

Several routes are possible within IPE. A common approach is to obtain bits of the phase from last to first, using successive calls to the circuit in figure 13 for a pre-determined sequence of integers  $m^{(k)}$  and adaptive angles  $\theta^{(k)}$  [68, 7, 67, 104, 26]. In this subsection we will prefer to use a superscript  $(k)$  to mark the iteration number, instead of a subscript which may cause confusion with binary notation and other subscripts (namely position indices).

It is straightforward to see that the wavefunction of the upper qubit just before the Hadamard gate is:

$$\frac{|0\rangle + e^{i(m^{(k)}\phi^{(k)} + \theta)}|1\rangle}{\sqrt{2}} \quad (2.87)$$

, i.e. the eigenphase has been kicked back into the upper qubit  $M$  times, as has the offset  $\theta$ , producing a relative phase. Clearly, the protocol exploits equations 2.85 and 2.86's implication that  $|\psi\rangle$  is also an eigenstate of  $\hat{U}^m$ .

$$\hat{U}^m |\psi\rangle = \lambda^m |\psi\rangle \quad (2.88)$$

$$\lambda^m = e^{im\phi} \quad (2.89)$$

The final state will be, in the computational basis:

$$\frac{1}{2} \begin{bmatrix} 1 + e^{i(m^{(k)}\phi + \theta^{(k)})} \\ 1 - e^{i(m^{(k)}\phi + \theta^{(k)})} \end{bmatrix} \quad (2.90)$$

If the exponent's norm,

$$\xi \equiv m^{(k)}\phi + \theta^{(k)} \quad (2.91)$$

, can be written in terms of a single digit binary fraction

$$\xi = -2\pi(0.\xi_1)_2 = 2\pi\frac{\xi_1}{2^1} \quad (2.92)$$

, the probability of measuring 0 is

$$\mathbf{P}(0 \mid \xi) = \begin{cases} 1, & \text{if } \xi_1 = 0 \\ 0, & \text{if } \xi_1 = 1 \end{cases} \quad (2.93)$$

In this case, the bit obtained by measuring precisely matches  $\xi_1$ , giving us  $\phi$ . These results are exact, apart from experimental error.

For longer  $\phi$ , still written as a binary fraction,

$$\phi = -2\pi(0.\phi_1 \dots \phi_M)_2 = 2\pi \left( \sum_{k=1}^m \frac{\phi_k}{2^k} \right), \phi_k \in \{0, 1\} \quad (2.94)$$

, we can get any  $\phi_k$  as the only digit in a fraction of the form 2.92, and thus determine them all. For that, we leverage the degrees of freedom in equation 2.91. We start by using  $m^{(k)}$  to push the digits that are more significant than itself to the left of the radix point, rendering them an inconsequential integer factor for  $2\pi$ . After that, we erase the ones that are *less* significant by performing a corrective rotation on the phase using  $\theta^{(k)}$ .

We don't know in advance what these corrections are; however, we do know that, for any digit, they are determined by the digits to the right of it. So we start with the rightmost digit  $\phi_M$ , which has none of those. In this case no correction is required. Then we move on to  $\phi_{M-1}$ , leveraging the freshly learned  $\phi_{M-1}$  to tune our circuit, and we proceed similarly for all others as we accumulate more and more digits.

In short, we run the circuit from figure 13  $M$  times, getting all the  $\phi_k$  in reverse order (from  $\phi_M$  to  $\phi_1$ ). We do so by iterating  $k$  backwards from  $M$  to 1 and choosing at each iteration the controls that will carry out the corrections described above.

$$m^{(k)} = 2^{k-1} \quad (2.95)$$

$$\theta^{(k)} = -2\pi(0.0\phi_{k+1}\phi_{k+2} \dots \phi_M)_2 \quad (2.96)$$

The iterator  $k$ , in addition to determining these controls, tags the digit being computed in the current iteration,  $\phi_k$ . Note that the phase correction  $\theta^{(k)}$  depends on the digits inferred in all previous iterations  $k + 1, \dots, M$ . That is, it considers as many digits as have been inferred so far (so if there are none,  $\theta^{(k)} = 0$ ; we've seen this happens for the edge case  $k = M$ ).

It is this feedback mechanism what enables the direct extraction of the binary  $\phi_k$  from the respective circuit's measurement outcome. At  $k = M$  we infer  $\phi_M$  with no prior knowledge; at  $k = M - 1$  we infer  $\phi_{M-1}$  after using  $\phi_M$  to rotate the phase; and so on. We can see that

$$\theta^{(k)} = -m^{(k)}\phi_{\text{curr}} \quad (2.97)$$

, where  $\phi_{\text{curr}}$  contains the binary digits determined up to the previous iteration  $k + 1$ , which are the  $M - k$  least significant ones, in their proper positions.

$$\phi_{\text{curr}} = (0.\underbrace{0 \dots 0}_{k \text{ zeros}} \phi_{k+1} \dots \phi_M)_2 \quad (2.98)$$

That leaves  $k$  placeholder zeroes corresponding to the digits to be filled in, all but one of which we eliminate when multiplying by  $m^{(k)} = 2^{k-1}$ , retrieving 2.96. This intuition will be alluded to later on.

Clearly, this strategy relies on accurate measurements of deterministic outcomes, as did the original protocol. Additionally, the maximum number of sequential applications of  $\hat{U}$  - and with it the necessary coherence time - scales rigidly and exponentially<sup>6</sup> with the total number of iterations (as the original one did with the number of ancilla qubits). Thus, it is quite susceptible to experimental errors, especially when error corrected devices are not a reality (which they aren't).

A more flexible approach is to regard circuit 13 not so much as a building block for a lineup that will output a solution, but rather as an entity on its own. It consists on no more than a system - one whose outcome distribution has a known dependence on the eigenphase  $\phi$ , in addition to some customizable parameters. This is the description of an inference problem.

Not unexpectedly, this can mean Bayesian inference. The resolution is up to the classical processing; we can apply whatever techniques we wish, namely machine learning ones. The Bayesian approach in particular has been proposed in [115] and subsequently implemented experimentally in [92]. Finally, [79] compares a few strategies, Bayesian and not, and in particular their error resilience.

This case differs from the precession example in that the target is a small circuit rather than a single qubit; but as before, its role is simply to relay the outcomes to a classical processing unit. Because the application of  $\hat{U}$  on  $|\psi\rangle$  is of no material consequence, we must necessarily fall back on some auxiliary structure. We thus assemble a circuit that will impart some relevant information; this can be the again that of figure 13.

We know that the structure of that circuit is sufficient, because it was sufficient when neatly arranged into a a specific layout. We will simply control and process it a less orderly manner. This is even more reminiscent of other NISQ algorithms where optimization, feedback and repetition are joined together in an attempt to counter experimental limitations. The pairing of shallow circuits with more mature (classical) machine learning techniques is by nature quite robust. In particular, it is much more error resilient than routines ran entirely on non fault-tolerant quantum devices.

With this choice, we need only apply the Born rule to the final state 2.90 to get our likelihood function

<sup>6</sup>Except if some additional structure can be exploited to get more efficient implementations of the required powers of  $\hat{U}$ .



corresponding to either of the basis states:

$$\mathbf{P}(\phi \mid 0; m^{(k)}, \theta^{(k)}) = \frac{1 + \cos\left(m^{(k)}\phi + \theta^{(k)}\right)}{2} \quad (2.99)$$

$$\mathbf{P}(\phi \mid 1; m^{(k)}, \theta^{(k)}) = \frac{1 - \cos\left(m^{(k)}\phi + \theta^{(k)}\right)}{2} \quad (2.100)$$

Actually,

$$\mathbf{P}(\phi \mid 0; m^{(k)}, \theta^{(k)}) = \cos^2\left(\frac{m^{(k)}\phi + \theta^{(k)}}{2}\right) \quad (2.101)$$

$$\mathbf{P}(\phi \mid 1; m^{(k)}, \theta^{(k)}) = \sin^2\left(\frac{m^{(k)}\phi + \theta^{(k)}}{2}\right) \quad (2.102)$$

(using half angle formulas).

The likelihoods 2.101 and 2.101 and those from the precessing qubit example (2.80 and 2.81) look very quite alike. As a matter of fact, *they are identical* if we identify the spanned angle  $\omega t$  with  $m^{(k)}\phi + \theta^{(k)}$  and convert between x and z basis.

This is, of course, because the two processes are equivalent. The effect of circuit 13 on the upper qubit is to pulse it into the x-basis, rotate it at the equatorial plane around the z axis, and then do an x-basis measurement (Hadamard transform + computational basis projection). The phase gate and the controlled  $\hat{U}$  both change the relative phase between basis states, affecting the Bloch coordinate  $\phi$  additively such that the accumulated phase is  $m\phi + \theta$ .

Predictably, the problem summary too nearly matches 2.84.

$$\boxed{\mathbf{P}(\theta \mid D; E) \leftrightarrow \mathbf{P}(\phi \mid 0; m, \theta) = \cos^2\left(\frac{m\phi + \theta}{2}\right)} \quad (2.103)$$

Here the parameter is the eigenphase, the data are the classical bits obtained upon measuring (which tell whether the qubit was found at state  $|0\rangle$  or  $|1\rangle$ ), and the experimental controls are two: the number of times  $\hat{U}$  is applied, plus the offset angle. The likelihood in 2.103 contains all information, because the probability of  $|0\rangle$  is complementary to that of  $D = |1\rangle$ .

The nuance is that here our *time* has discrete units  $m$  (we would recover continuity if we took  $\hat{U}^m = e^{-i\hat{H}m}$ <sup>7</sup>); also, there is now a controllable offset  $\theta$ , a freedom we could have added to the previous example by choosing a different initial state still within the x-y plane.

Whatever they be, the existence of these variable controls is opportune: they allow us to tweak the outcome distribution. Generally, degrees of freedom in the likelihood function are welcome, albeit not strictly necessary. With them, we can make more out of each measurement by choosing appropriate controls. We are not simply passive agents who register and process data: we have the power to design the experiments that produce those data. But how do we use such a power?

<sup>7</sup>A similar example is characterization of coherent gate error. For instance, applying an imperfect  $\hat{\sigma}_x$  gate twice will rotate the Bloch vector by  $(2\pi + 2\epsilon)$  radians around the z axis. This error  $\epsilon$  will accumulate to cause an angle  $2m\epsilon$  if the gate is applied an even number of times  $2m$ , which produces a similar oscillation of the probability of  $|+\rangle$  as a function of  $m$ .

## 2.4 Bayesian experimental design

When performing inference, not all measurements are made equal. The structure of the likelihood function itself impacts the achievable utility of each experiment: if it varies more slowly with the parameter(s), it will produce less precise parameter estimates. As an example, it is not hard to imagine that an exponential model exhibits reduced sensitivity as compared to a linear one.

Having chosen a model, the shape of the likelihood still changes depending on the controls. As a simple example, let's go back to our coin tosses. We had a coin, say coin A, with some probability of tails  $p = \theta$ , where  $\theta$  was the target parameter. We can now imagine that we had access to a *second* coin, coin B, in addition to the original one. We are told coin B's probability of tails  $p = 0.1\theta$  is a tenth of coin A's. That is,

$$p_A = \mathbf{L}(\theta \mid 1; A) = \theta \quad (2.104)$$

$$p_B = \mathbf{L}(\theta \mid 1; B) = 0.1\theta \quad (2.105)$$

This introduces a new control: the coin pick. We can choose which coin to toss a few times before taking a guess at  $\theta_{\text{real}}$ . Which *do* we choose? We may expect that coin B will do worse: it is harder to resolve  $0.1\theta$  than it is  $\theta$ , as the latter manifests a stronger dependence on  $\theta$ . That is indeed the case. Figure 14a shows the posterior distributions resulting from 100 coin flips for each coin; the one with  $p = 0.1\theta$  has a flatter curve, with larger variance (signifying more uncertainty).

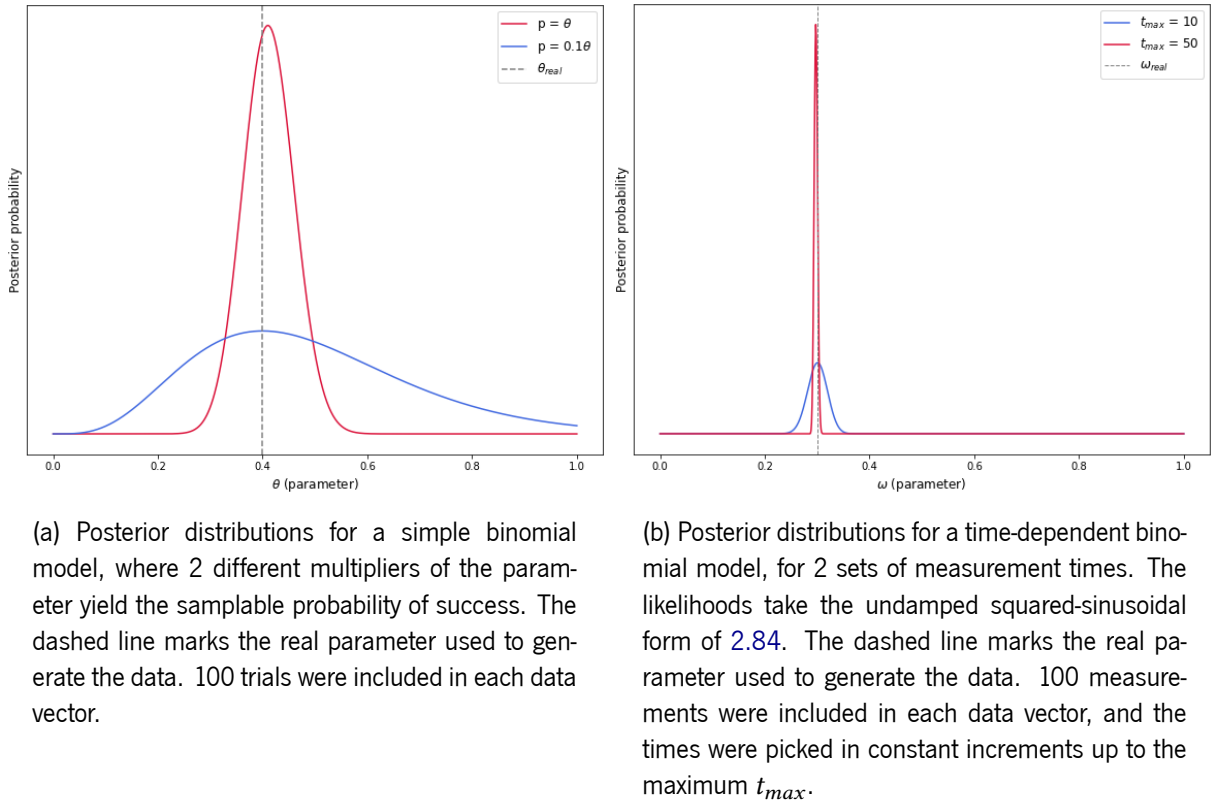
How about the oscillatory likelihoods of 2.84 and 2.103 - how should we pick the times? They show inside the sinusoidal function as a multiplying factor for the target parameter. For a fixed parameter value, longer times mean faster oscillations. We can anticipate that a Bayes' update will be more informative, and the posterior sharper, if the likelihood varies rapidly with the parameter.

Yet again, we have a point. The posterior distributions for two different sets of measurement times (spaced evenly on  $[0, t_{\text{max}}]$  for two values of  $t_{\text{max}}$ ) are depicted in figure 14b. The longer times yield a visibly sharper curve<sup>8</sup>.

One more thing to notice is that both curves of 14a are narrower as compared to those of 14b, testifying to the inherent differences between models mentioned earlier: the sinusoidal wave varies more rapidly than the proportionality (though naturally this depends on their frequency and slope). But once again, we will be holding the model fixed for each test case; it is ultimately determined by the target system, and we have already discussed how to select it. The experiments, on the contrary, can be chosen completely at our whim.

This is very promising: what we are observing is similar to what happened back in section 2.2 when we added more data (from figure 2 to 3). But here we are *not* adding data, nor are we making any more Bayes' updates. We are simply changing the experiments we perform, and that alone has the power to change our results for the better - despite there being no extra processing whatsoever.

<sup>8</sup>Although equation 2.84 is the one directly considered here, these results can be straightforwardly generalized to the case of 2.103, where the *times* are integers due to experimental constraints.



(a) Posterior distributions for a simple binomial model, where 2 different multipliers of the parameter yield the samplable probability of success. The dashed line marks the real parameter used to generate the data. 100 trials were included in each data vector.

(b) Posterior distributions for a time-dependent binomial model, for 2 sets of measurement times. The likelihoods take the undamped squared-sinusoidal form of 2.84. The dashed line marks the real parameter used to generate the data. 100 measurements were included in each data vector, and the times were picked in constant increments up to the maximum  $t_{max}$ .

Figure 14: Impact of the magnitude of the parameter's coefficient(s) in the shape of the likelihood for a simple binomial model (14a) and a damped oscillator model (14b).

To make use of this fact, we need a way to identify the best experiments *before* we perform them. How can we do that, if we don't know the posterior they would give rise to?

While it is true that we don't know the posterior in advance, we are from the outset in possession of the next best thing: our prediction of where it will stand, based on our current knowledge (the prior). We also don't know or how sharp a set of experiments will actually make it; but by the same token, we can compute how sharp a set of experiments makes it *on average*. This is sufficient to make an informed choice, and stack the odds in our favour.

Interestingly, this discussion seems somewhat familiar. Back in subsection 2.2.1, when talking about estimators, we spoke of their associated Bayes risk, which was the expectation over the data of the posterior averaged loss. We suggested we'd pick the estimator that we expected to minimize this risk.

$$r(\hat{\theta}) = \mathbb{E}_{\mathbf{P}(\vec{D}), \mathbf{P}(\theta|\vec{D})} \left[ \mathcal{L}(\theta, \hat{\theta}(\vec{D})) \right] \quad (2.33)$$

With the choice of an estimator, the Bayes risk becomes a fixed quantity.

$$r = \mathbb{E}_{\mathbf{P}(\vec{D}), \mathbf{P}(\theta|\vec{D})} \left[ \mathcal{L}(\theta, \hat{\theta}_X(\vec{D})) \right] \quad (2.106)$$

Or does it? In the meantime, we have decorated Bayes' equation with some extra dependencies. As

we are considering a fixed model, we'll refer back to equation 2.58. Adjusting 2.106 accordingly, we get:

$$r(\vec{E}) = \mathbb{E}_{\mathbf{P}(\vec{D}; \vec{E}), \mathbf{P}(\theta | \vec{D}; \vec{E})} \left[ \mathcal{L}(\theta, \hat{\theta}_X(\vec{D}; \vec{E})) \right] \quad (2.107)$$

The risk now depends *on the experiments*. This is a formalization of what we saw in figure 14. In subsection 2.2.1, we talked about several loss functions, and how they were typically a function of the deviation from the real value. When taking expectations over the posterior, they become measures of the spread of the distribution. The loss function we picked, the variance, clearly demonstrates this. If a set of experiment controls has a smaller risk relative to some reference, it is likely to make the posterior sharper.

When talking about experiments, we don't generally talk about *risk*, but rather about *utility*. Unlike an estimator, an experiment doesn't incur error; it is simply more or less informative. We can choose our utility function to be symmetric to the risk:

$$U(\vec{E}) = -r(\vec{E}) \quad (2.108)$$

We will do just that and take the negative variance as our utility<sup>9</sup>, which was the result of our choice of a loss function in subsection 2.2.1. Of course, alternative functions can be considered; another popular choice for estimation problems in physics is the information, measured by the entropy [33, 51]. Penalties can also be factored in, using e.g. the total evolution time if the cost of simulation is the main concern.

All tools are now at our disposal. For any group of experiments  $\vec{E}$ , we can obtain the posterior expected utility of a list outcomes by mock-updating the distribution, and then marginalizing the loss over the resulting posterior distribution  $P(\theta | \vec{D})$ . At the same time, we can obtain the current expected probabilities of the outcome list by marginalizing the likelihood over the prior.

This is done for all datasets, after which we multiply each dataset's probability by the corresponding utility (both calculations depend on the experiments  $\vec{E}$ ). Finally, we sum over all possible data vectors, which are  $2^N$  for  $N$  experiments with binary outcomes. With this, we get the expected utility of the vector of experiments:

$$U(\vec{E}) = \sum_{\vec{D}} \mathbf{P}(\vec{D}; \vec{E}) \cdot \mathbb{E}_{\mathbf{P}(\theta | \vec{D}; \vec{E})} \left[ U(\theta, \vec{D}; \vec{E}) \right] = \sum_{\vec{D}} \mathbf{P}(\vec{D}; \vec{E}) \cdot U(\vec{D}; \vec{E}) \quad (2.109)$$

Figure 15 shows the operations required to compute the utility of a set of experiment controls. To make an optimal choice, this must of course be evaluated for multiple possibilities of  $\vec{E}$ .

Clearly, utility maximization is not all roses. In addition to computing the marginal probabilities of all  $2^N$  data-vectors, we must compute their average posterior loss (e.g. the final variance). Each of these computations requires as much processing as the entire inference process did, because they are precisely a data-conditional *look-ahead* into the posterior. And this is just for a single evaluation of the utility at a

<sup>9</sup>The variance is a property of a specific posterior, so to get the final utility we must still average the variance over all possible ones (i.e. take the expectation over the data). The term utility is often used to refer both to the data-conditional utility and the expected utility. The subsequent double occurrence of the function  $U$  in 2.109 with different arguments further attests to this fact.




Figure 15: Binary tree illustrating the computation of the expected utility for a set of experiments  $\vec{E} = \{E_1, E_2, E_3\}$ . To each layer of branches corresponds one of the individual experiments, in chronological order from top to bottom. Each branch represents the occurrence of a specific datum, and is tagged with its likelihood. The leaves are marked with the utility of the data sequence traced out by the path of branches from the root to themselves. Using these calculations, the expected utility of  $\vec{E}$  can be derived from 2.109.

specified  $\vec{E}$  - if we are to take on optimization, these the calculations need to be done all over again for many other possibilities, as they're evidently not re-usable for different choices of  $\vec{E}$ .

This exponential scaling with the total number of experiments is prohibitive, and soon becomes unfeasible. We're in the presence of a trade-off between the number of experiments and the processing cost. If measurements are an expensive resource, it may be interesting to take full advantage of each by being selective in the choice of controls. If not, however, we may be better off rejecting the processing overhead and running a larger number of arbitrary experiments. Between these two options, could there be hope for a more balanced compromise?

The issue is brought on by the construction of the annotated binary tree in figure 15, which has height  $N$ . To a degree, it unavoidable that we undertake it; but a more lightweight strategy is possible. We can exploit the parallel between the central Bayes' updates and the accessory ones by building the tree as we go, effectively intertwining its construction with the mandatory processing.

Such a coupling can be achieved by processing the data real-time, and forgoing the quest for a global maximum in favour of local optimality in what is called a *greedy* algorithm. Equipped with adaptivity, we can move along the linear path that the data conduct us through.

### 2.4.1 Adaptive measurements for a greedy strategy

In sections 2.2.2 and 2.3, we claimed that whether we processed the data one by one or all at once was indifferent. And it is - as long as everything is equal. This *everything* includes, of course, the experiments. Figure 16 unfolds the process of 10, representing the iterative version of the inference algorithm. These diagrams are again based on [51, 33], as is the analysis itself. Under these conditions, the two diagrams are equivalent.

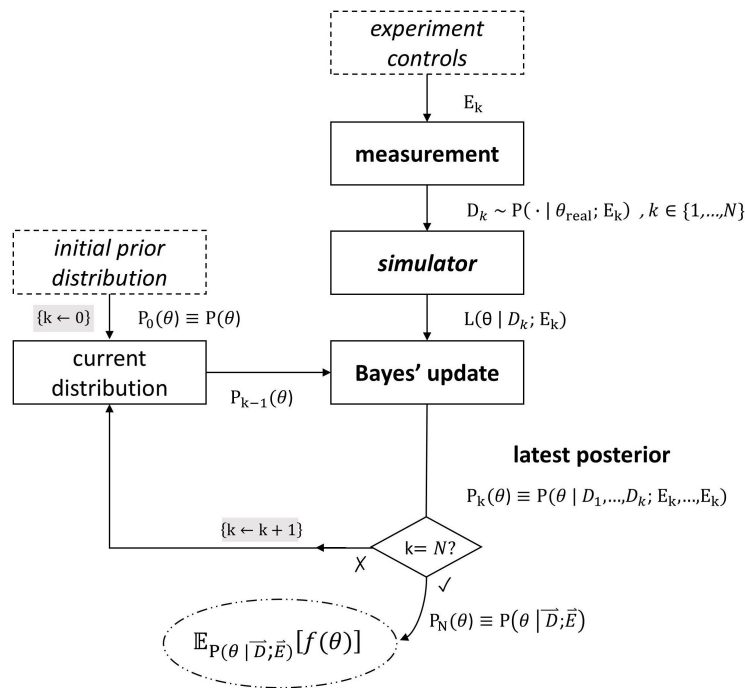


Figure 16: Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference, based on iterative updates. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs.

In spite of this equivalence, when going from the single datum updates of 2.58 to the batch updating of 2.61, we actually did lose something: the passage through intermediate stages of knowledge. While it's fine to skip over them to our final conclusion if that's all we care about, we are now in a different scenario. After the collection of the first datum  $D_1$ , we already know a little bit more than we did to begin with; we have a first-iteration posterior,  $\mathbf{P}(\theta | D_1; E_1)$ . Then why not use this knowledge, instead of following a pre-determined route?

This can be done by introducing a feedback mechanism, as shown in figure 17. We can delay deciding which second experiment  $E_2$  to perform until the moment we collect  $D_1$  and update the distribution according to it. *That is* the full extent of our *a priori* knowledge at the time we must perform the experiment  $E_2$  - and we can put ourselves in a position to use it.

This is entirely applicable to all subsequent experiments too; for each, we use all the data gathered before it. The datum obtained with each experiment precipitates an update from the latest *prior* to a

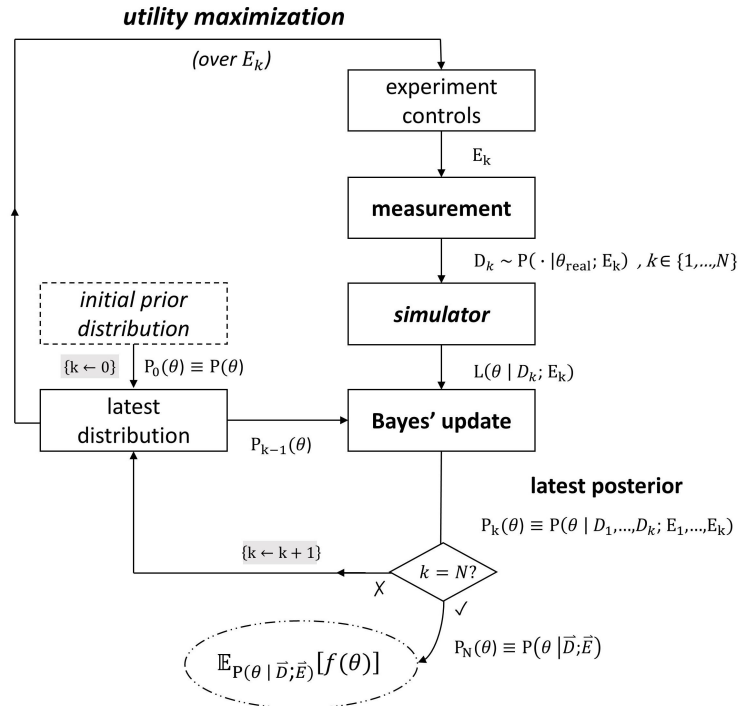


Figure 17: Diagrammatic representation of a quantum parameter estimation algorithm using Bayesian inference with local utility optimization. Italics signal dependence on implementation choices, and dashed boxes denote one time only inputs.

transitory *posterior*, which becomes the prior for the ensuing iteration. We had already mentioned that the posterior can be seen as an evolved version of the prior in equation 2.15.

The terms prior and posterior denote precisely the preexisting knowledge and its after-the-fact state, where the order is relative to the data-led updates. We are free to decide how phased the assimilation of the data should be; it can be broken up into chunks or regarded as a whole. Should we want to use it before finalizing the process, we can *compile* it more often.

The only new imposition is that we switch to an online setting, and process the data as we go: in between measurements, we must update the distribution using the previous datum  $D_{k-1}$  and employ local optimization to choose the next control  $E_k$ . The update is something we'd have to do eventually anyway, and the cumulative cost of the optimization will be smaller than it would if we were to seek a global optimum.

Looking at the binary tree of 15, it is as if we travel down the branches as the inference advances, only ever using a single level look-ahead. Instead of concerning ourselves with the likes of the formidable tree in 15, we consider a sequence of mini-trees of height 1, only ever thinking in terms of the experiment coming next and its two possible outcomes. Figure 18 shows two of these smaller trees. Their construction is necessarily sequential.

The feedback loop in 17 is oversimplified, and presupposes further processing, and access to a simulator (just as the main algorithm does). Figure 19 shows the required set up explicitly. The dashed section

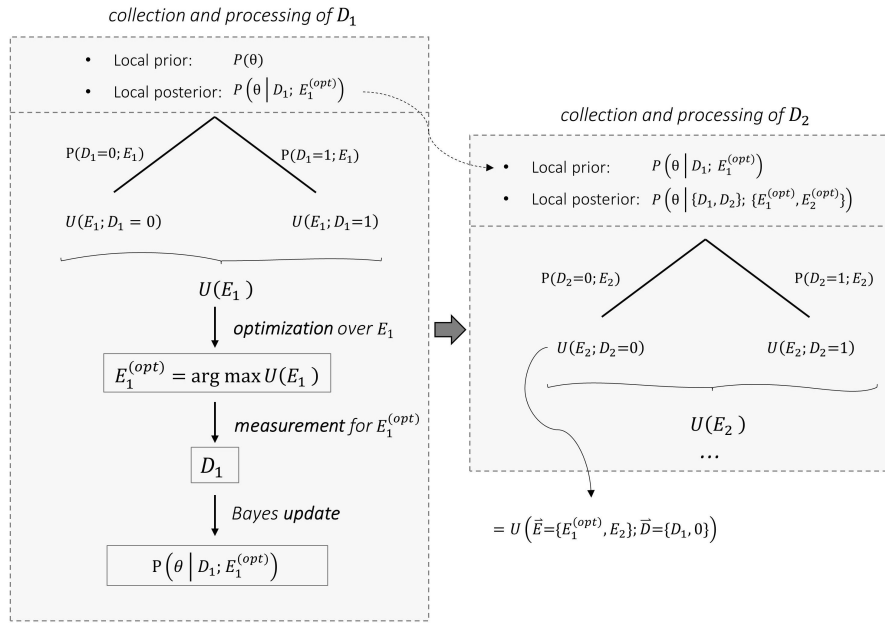


Figure 18: Binary trees illustrating the computation of the locally expected utility for two followed experiments,  $E_1$  and  $E_2$ . To each of the height-1 trees corresponds one experiment, in chronological order from left to right. Each branch represents the occurrence of a specific datum, and is tagged with its likelihood. The leaves are marked with the utility of the respective branch's datum. Using these calculations, the expected utilities of  $E_1$  and  $E_2$  can be derived from 2.109. The equality annotated bottom right evidences that the latter calculation is dependent on  $E_1$  and the effective  $D_1$ .

encloses the computation of the utility as per equation 2.109.

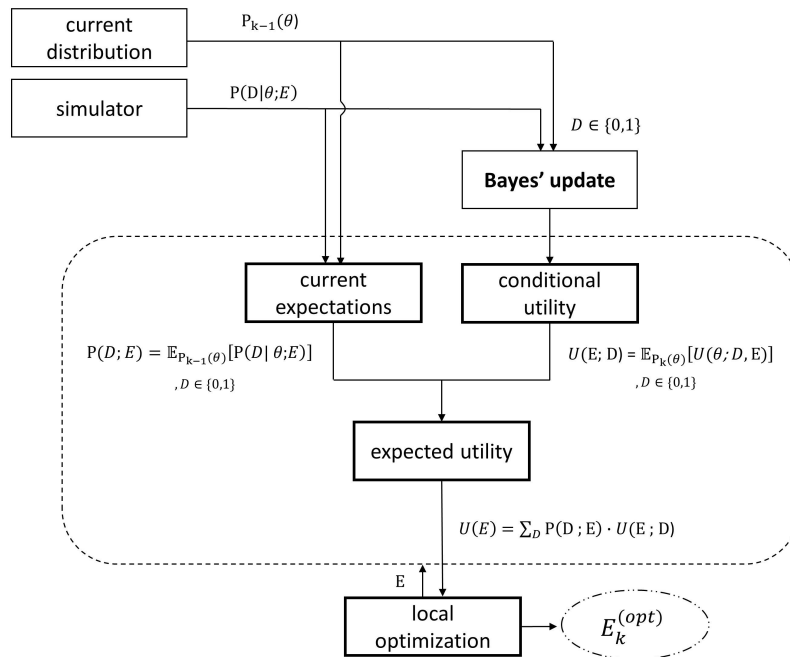


Figure 19: Diagrammatic representation of a local optimization routine in the context of Bayesian experimental design.



In general, this greedy strategy will not prompt the same choices that the globally optimal one would have. However, it is a quite less demanding alternative, and has been shown (numerically) to produce good results in the case of a flat prior [33]. It may also be particularly interesting in the case of such an uninformative prior, since it particularly benefits from the inference process as it moves forward.

This type of measurement interdependence has been used to perform estimation with a precision scaling at the best limit allowed by quantum mechanics. As a matter of fact, adaptivity underpinned the first experimental realization of Heisenberg-limited phase estimation - filling in for the laborious state preparation at times thought to be irreplaceable [57]. It may then be interesting to look at *why* it allows reaching such a high precision regime.

## 2.4.2 Beating shot noise: quantum-enhanced, Heisenberg-limited estimation

In metrology, one aims to achieve the best possible measurement tactics; in the case of a quantum system, it is to expect that they're be inescapably tied to a quantum description. For any estimation strategy, we don't arrive at the limits of describability until we overcome the shortcomings of our description. For this reason, a quantum problem begs a quantum solution, which promises the best known achievable accuracy.

Here we will present a quick overview<sup>10</sup> of the limits quantifying the attainable bounds of this advantage, as they are quite linked to adaptive inference strategies.

For that end, we first consider an ensemble of  $N$  independent probabilistic measurements on a qubit; as before considering experimental design, we would perform them at once, without much concern for their controls. They could be separated in space (different qubits) or in time (different passes through the same circuit). Based on these experiments, we get an estimate for some target quantity; for instance, the probability of finding the qubit in state  $|1\rangle$ . Writing each measurement outcome as  $D_j$ , that estimate is:

$$\bar{D} \equiv \mathbb{E} [D_j] = \frac{\sum_{j=1}^N D_j}{N} \quad (2.110)$$

This could be applied for instance to the frequency/phase estimation problems summarized in 2.84 and 2.84; for some fixed experiment, we could determine the likelihood and then solve with respect to the parameter to get a ratio estimate.

Taking a finite number of random samples imposes a limit on the precision of this estimate; limited resolution is tantamount to observing the target quantity through noise. For our independent stochastic variables in particular, we consider gaussian noise: we can invoke the [central limit theorem \(CLT\)](#) to say that the distribution of outcomes  $D_j$  is approximately normal with some mean  $\mu$  and variance  $\sigma^2$ .

$$D_j \sim \mathcal{N}(\mu, \sigma^2) \quad (2.111)$$

<sup>10</sup>The discussion is quite simplified; rigorous analysis requires more careful definitions, in particular of the resource count  $N$  [120, 57].

Equation 2.110 is the specific mean we get from a sample of  $N$  results, and we are equating its obtention with sampling from a normal distribution 2.111. This is despite the fact that the original variables are not normally distributed (this is justified by the CLT under our suppositions).

In some cases, this sample mean might align perfectly with the exact value; but more often than not, it won't. To quantify the expected magnitude of this variability, we consider meta-statistics regarding this distribution of sample means.

$$\bar{D} \sim \mathcal{N}(M, \Sigma^2) \quad (2.112)$$

This distribution has itself a mean, which by linearity of expectation and 2.111 is simply  $\mu$ .

$$M = \mathbb{E}[\bar{D}] = \mathbb{E}\left[\frac{\sum_{j=1}^N D_j}{N}\right] = \frac{1}{N} \sum_{j=1}^N \mathbb{E}[D_j] = \frac{1}{N} \sum_{j=1}^N \mu = \mu \quad (2.113)$$

Centrally, we quantify its spread around the mean  $M$  by means of its variance  $\Delta^2$ :

$$\Delta^2 = \mathbb{V}[\bar{D}] = \mathbb{V}\left[\frac{\sum_{j=1}^N D_j}{N}\right] \quad (2.114)$$

The reason why we concern ourselves with the first two moments of the distribution only is that the mean is trivial and the variance suffices for our purposes; hence the representation as a Gaussian in 2.112.

By the properties of the variance, 2.114 becomes:

$$\mathbb{V}[\bar{D}] = \mathbb{V}\left[\frac{\sum_{j=1}^N D_j}{N}\right] = \frac{1}{N^2} \mathbb{V}\left[\sum_{j=1}^N D_j\right] \quad (2.115)$$

We can appeal to the fact that the measurements are independent, and so the variance of their sum is the sum of their variances.

$$\frac{1}{N^2} \mathbb{V}\left[\sum_{j=1}^N D_j\right] = \frac{1}{N^2} \sum_{j=1}^N \mathbb{V}[D_j] \quad (2.116)$$

At which point we invoke 2.111 to get:

$$\Delta^2 = \frac{1}{N^2} \sum_{j=1}^N \sigma^2 = \frac{\sigma^2}{N} \quad (2.117)$$

This error arising from statistical fluctuations represents the so-called **standard quantum limit (SQL)**, or shot noise limit (a term borrowed from electronics, where it arises from discrete charges rather than discrete samples). The standard deviation of the sample means  $\Delta$  quantifies our uncertainty, and the SQL describes how it scales with the relevant resource (the number of shots,  $N$ ).

$$\Delta^{(SQL)} \propto \frac{1}{\sqrt{N}} \quad (2.118)$$

The outcome variance  $\sigma^2$  is of course independent of  $N$ : the total number of experiments could never influence their individual results.

Finally, we can dissect our assumptions. This is valid for *independent* measurements; the calculations do not hold in other cases. In particular, better scaling can be achieved by exploiting state preparation and measurements in more thought-out schemes. If arbitrary quantum effects are used, the fundamental limit is Heisenberg's (HL):

$$\Delta^{(HL)} \propto \frac{1}{N} \quad (2.119)$$

The scaling of 2.119 represents the regime of *quantum-enhanced* estimation, which offers the fundamentally optimal resolution [120]. Many of the proposals to achieve it consider introducing correlation by means of entanglement. The quantum Fisher information quantifies the extractable information for a state-observable pair, and can be used to directly determine the maximum attainable measurement accuracy (through the quantum Cramér-Rao bound). For separable states it is linearly bounded by the number of particles, whereas for general states this bound is quadratic [107].

This has motivated the use of entangled states for parameter estimation [120]. However, the complexity required in the preparation of such states has often shown to be prohibitively challenging for the available technology [45], and the presence of noise sorely undermines their longevity [107]; such problems are further exacerbated by high qubit counts. These points have painted sub-shot noise estimation as something to be left for the future, especially in the case of large numbers of qubits.

A significantly more manageable alternative was proposed in [57], and relied precisely on a locally optimal Bayesian algorithm. The procedure sought to minimize the expected posterior variance at each step - effectively replacing intricate input states by sequential optimized reruns of a parametrized circuit. This greatly relieved the requirements, enabling the experimental demonstration of Heisenberg scaling.

### 2.4.3 Empirical strategies for a precessing qubit

Even if so far we have spoken only of general estimation strategies, looking into the concrete structure of the likelihoods at hand may often prove advantageous. The adaptive measurements we surveyed produce good results, but are still quite demanding in view of the optimization to be carried out between measurements, especially when considering simulation based likelihoods. The processing overhead can be alleviated by problem-specific heuristics, which may afford favorable solutions at dramatically lower cost.

In this subsection, we will take a closer look at the particular case of our examples of section 2.3.3. Based on [36, 115, 51, 116], we will go over some adaptive heuristics that can be a proxy for full-on optimization or more invested approaches, posing a less demanding alternative. While at it, we will take the opportunity to get a more intuitive understanding of the mechanisms driving measurement optimality, and bring up some issues that become relevant if considering approximate inference.

The strategies to be considered here are *handcrafted*, in the sense that they are grounded in observed behaviour, be it theoretical or practical. More general (and powerful) approaches can be achieved by automating the reasoning; in some cases, pre-existing heuristics can even be used as a starting point, should

any be available. As an example, [79] relies on particle swarm optimization to create an *offline adaptive*<sup>11</sup> phase estimation protocol. More recently, reinforcement learning has been used to create neural-network based heuristics; details can be found in [38], which also presents an interesting performance comparison encompassing the approaches to be addressed here. However, these machine learning based methods are left out of this overview.

Having noted this, we will now introduce a few relevant schemes, starting with analytical considerations. For the squared-sinusoidal likelihoods in 2.84, [36] analyses the behaviour of the risk under a normality assumption for the prior. Such an assumption enables the analytical treatment of the problem, which may assist in designing practical strategies for improving inference outcomes. A flat prior can be made approximately Gaussian by an update based on a small initial set of sensible predetermined measurements.

Assuming the warm up phase has been realized, we seek optimality through adaptivity; the analytically tractable prior can be used to arrive at an asymptotic expression for the Bayes risk (again for quadratic loss) as a function of the experiment control. This control is the measurement time, a continuous variable.

The authors find that this quantity oscillates between a constant maximum and a lower envelope, with a frequency that increases as the inference advances (i.e. with more measurements, as the spread of the distribution is reduced).

We consider a counter  $k \in \{1..N\}$  starting out at the first adaptive measurement, after the prior has been warmed up to

$$\mathbf{P}(\omega) = \mathcal{N}(\mu_0, \sigma_0) \quad (2.120)$$

At any stage, the minimum of the envelope is given in terms of the latest posterior's variance as

$$r_k^{(opt)} \equiv r(t_k^{(opt)}) = (1 - e^{-1})\sigma_{k-1}^2 \quad (2.121)$$

, and occurs for

$$t_k = 1/\sigma_{k-1} \quad (2.122)$$

,  $\sigma_k$  being the standard deviation of the  $k$ th iteration posterior,  $(k + 1)$ th iteration prior.

As the distribution becomes sharper, the minimum of the risk approaches this lower bound location, due to the oscillations' rapid variation with respect to the bottom of the envelope [36, Appendix B]. Thus, the optimal scaling of the uncertainty with the iteration number is asymptotically described by the successive envelope minima:

$$\sigma_k = \left(\sqrt{1 - e^{-1}}\right)^k \sigma_0 \quad (2.123)$$

This gives the rate at which we expect the uncertainty to shrink. These minima are associated with the times

$$t_k = \frac{1}{\left(\sqrt{1 - e^{-1}}\right)^{k-1} \sigma_0} \approx \frac{1.26^{k-1}}{\sigma_0}, k \in \{1, \dots, N\} \quad (2.124)$$

---

<sup>11</sup>This is in the sense that each measurement relies on the previous result, but virtually no online processing is required. The optimization is performed beforehand, after which the feedback is determined by a quick calculation.

, i.e.  $t_1 = 1/\sigma_0, t_2 = 1.26/\sigma_0, \dots$

This result is however only valid asymptotically, and does not hold for a finite number of experiments. Nonetheless, it may still be a useful insight; for instance, we could concentrate the search for an optimal solution around this lower bound.

We can even drop full-blown optimization altogether, and instead just browse around a vicinity of the envelope's minimum by comparing the utility of a few possible  $t$ s. We expect that the risk's minimum will be nearby, an expectation that only becomes stronger as we move forward.

In light of this, a possible solution is to compute the current standard deviation to obtain the time given by 2.124, and then perturb its location for a number of attempts (e.g. by sampling from a Gaussian centered at it) to get a few tentative  $t_j$ . Then we could determine the expected utility of these selected  $t_j$ , following equation 2.109, and choose the one among them which yielded the maximal value.

$$t_k^{(\text{chosen})} = \arg \min [U(t_j)] \quad (2.125)$$

But we don't necessarily need to calculate the current standard deviation. We could pick pairs of values  $(\omega_j, \omega'_j)$  at random by sampling from the current distribution:

$$\omega_j, \omega'_j \sim \mathbf{P}(\cdot | \vec{D}_{curr}) \quad (2.126)$$

, after which we take the reciprocal of their distance  $t_j = 1/|\omega_j - \omega'_j|$ . This replaces  $1/\sigma_{curr}$  while organically introducing variability. Also, the intuition of 2.124 can be generalized to higher dimensions [116], where this surrogate heuristic conveniently eschews the inversion of the covariance matrix.

This type of heuristic based guesses have been used in e.g. [51, 50]. In subsection 5.2.5 they will be applied to a quantum computer, and 4.2.3 will attempt multi-modal generalizations.

With this, we may significantly reduce the complexity of choosing good experimental controls, opting for an in-between solution where we neither choose them arbitrarily nor extensively optimize them. Computing the utilities still demands a number of likelihood evaluations, but we are in a position to decide how many (by adjusting the total number of attempts). Further, the utility can be made to account for experimental issues, namely by being partial to evolution times well within the coherence time of the system. This is pertinent when dealing with real devices, possibly more than idealized utility maximization; decoherence may significantly weigh down the accuracy of otherwise promising guesses.

What is more, concessions can easily be made, as experimental design can tolerate rougher calculations than the nucleus of the inference process. It is not of the utmost importance that we choose the absolute optimal experiments, so precision may be better left for where it matters most. The positioning of approximations in the inference process is discussed in [51].

As a matter of fact, even settling on some apt guesses without ever considering the utility explicitly may bring better results than random measurements, for little to no added processing cost. By picking the asymptotic optimum 2.124, we can hope to land in proximity to the actual optimum more than if we used indiscriminate measurements. Since as mentioned before the upper bound for the risk is flat, we never increase the maximum risk.

For these reasons, times inversely proportional to the standard deviation have been experimentally applied to the precession example in e.g. [113], and also to phase estimation in [92] after being proposed in [115]. In the latter case the times are rounded to an integer  $m$ , and there's one more parameter to choose: the feedback angle  $\theta$ , which is picked directly from the prior/latest posterior (apart from a multiplying constant dependent on the chosen  $m$ )<sup>12</sup>. Summarizing, the phase estimation controls are obtained as (by this order):

$$m_k \propto \left[ \frac{1}{\sigma_{k-1}} \right] \quad (2.127)$$

$$-\frac{\theta_k}{m_k} \sim \mathbf{P}_{\text{curr}}(\cdot) \quad (2.128)$$

Since these results are not guaranteed to be optimal, the constant of proportionality of 2.127 can be optimized (e.g. by picking a set of hypothetical real values at random from the prior, performing a hyperparameter sweep, and choosing the hyperparameter instance that yielded the lowest loss or expected loss on average). The second equation 2.128 just means that we sample a value from the current distribution (of the parameter, the phase  $\phi$ ), which considers the data record up to the present iteration,  $\mathbf{P}_{\text{curr}}(\phi) \equiv \mathbf{P}_{\text{curr}}(\phi \mid \vec{D}_{\text{curr}})$ , and then multiply by the symmetric of the previously obtained constant  $m_k$ . Refer to figure 13 and/or box 2.103 for the meaning of these controls.

As an alternative, two different strategies are presented in [36]. One is based on locating the risk minimum with a correction to the asymptotic solution, while at the same time extending the Gaussian assumption to the successive posteriors (rather than restricting it to utility related considerations). This reduces the Bayes' update to a pair of rules that produce the mean and variance directly, thereby lessening the online processing load. This type of strategy is sometimes called *offline adaptive* [79].

The second one is to choose the times to be exponentially spaced out (loosely emulating those of 2.124, aside from a multiplying constant), i.e.

$$t_k = C^k \quad (2.129)$$

for some  $C$  chosen beforehand (*offline*, without making use of the data). The article proposes  $C = 9/8$  (for a prior domain normalized to  $\omega \in ]0, 1[$ ).

We started out with rigorous considerations, and gradually relaxed them to the point they're just some reasonable bids. If one is to only barely ground them in reason, why didn't we start with guesswork from the outset? Let us backtrack to the beginning of this section to try to get some insight on what useful experiments look like in practice, and why we can't identify them so easily. Along the way, we will consider the ways applied practice means can interfere with this usefulness, and how approximate inference may result in a more delicate interplay.

<sup>12</sup>This seems inspired in the exact IPE case of 2.97 (note that there the iterations are in superscript to avoid ambiguous notation, whereas in other sections, namely this one, we opt for subscripts), of which 2.128 is reminiscent. However, [79] shows it is not optimal for inference, at least not under a (fairly applicable normality assumption).

Considering the dependence illustrated by figure 14, the choice may have seemed easy enough: we can simply choose the controls whose gradient with respect to the parameters is larger in magnitude. They should be particularly useful, because they'll be quicker to tell us where in parameter space we should focus our attention.

So in the oscillator case, we might be tempted to snub shorter evolution times entirely. Of course, there's a limit to this; after a certain point decoherence and noise will take over, and we don't want our data collection process to take forever. But ideally, weren't it for those trade-offs, we may be inclined to make the intervals as long as possible. This idea seems deceptively simple. Should we?

That's not really the case. Spectacularly sharp likelihoods aren't ideal so long as we can't perform exact inference, which we saw in section 2.2.2 happens often. In the limit, our likelihood function could take the form of an infinitely tall and infinitesimally wide pulse at an exact value. It would be completely unambiguous, but that wouldn't do much for us, since we'd have vanishing probability of pinpointing its location in parameter space.

In less extreme cases, it is still a point to consider: we have to select a finite set of points at which to evaluate the posterior, and if the function is very sharp we need a very dense grid to ever find non-zero probabilities. If we simply want a rough estimate, we may be better off relying on a flatter curve.

This is illustrated in figure 20, where posteriors analogous to those of figure 14b were evaluated on the same grid. The single point evaluations fail to provide any insight for the sharper curve, whereas in the case of the flatter curve one of the points succeeds in signaling mode proximity.

It is clear that sharp likelihoods aren't the be-all and end-all of *approximate* inference, in particular if resource economy is a concern. But save for that, should our idea inspire no more misgivings?

That is still not the case. Elusive distribution peaks aren't the only caveat that it holds. As much as higher frequencies mean more information in a way, they also tend to bring ambiguity. This is demonstrated in 21, where as compared to 20 we have reduced the number of measurements while preserving the intervals over which we space their corresponding times.

As the likelihoods associated to the data oscillate faster, their periodicity becomes more densely packed in parameter space. This gives rise to equivocal posteriors, where we are likely to observe multi-modality. This leads to irresolution: which mode do we pick if they have equal or similar predictive power for the results that we obtained? Our credible region is now much more disperse. The issue with *high frequency information* is not unique to this problem or strategy; similar drawbacks appear when considering entangled states for phase estimation [9].

In our case, incorporating lower times should help resolve this; clearly, there aren't as many modes as there could be in figure 21, and that's because we are considering a multitude of experiments. Each experimental time yields a datum, which produces an individual likelihood; only when we multiplying them all together do we get our final results.

In other words, our understanding doesn't spring from lone observations, but rather from their juxtaposition. If we add in a datum associated to a very long measurement time to the blue posterior, we don't




Figure 20: Posterior distributions for a time-dependent binomial model, for 2 sets of measurement times. The likelihoods take the undamped squared-sinusoidal form of 2.84. Both posteriors were evaluated on the same grid comprising 10 points (the 'x' markers), and are juxtaposed with smooth curves of matching color representing their respective continuous posterior distributions. The latter were approximated using 1000 points. The dashed line marks the real parameter used to generate the data. 100 measurements were included in each data vector, and the times were picked in constant increments up to the maximum  $t_{max}$ .

expect that it will be immediately destabilized - most of the peaks will be cancelled out by lower probability events at the same parameter site.

This is evidenced by comparing the crimson curves of figures 20 and 21; in the former we don't face the same problem as in the latter, despite using as long times. We strengthen the result through increasing the times more slowly; of course, this also means more data were necessary, and with them a higher processing cost. With this in mind, it is pivotal that we find a balance between reliable but shallow likelihoods, and sharp but undependable ones.

The problems represented in figures 20 and 21 are in a way related. They both suggest that we shouldn't unreservedly aim for long evolution times, even if dealing with a perfectly isolated system; and they both could be managed by progressive management of the search space.

In other words, we could profit from a fluid discretization of the posterior - one which deploys its grid<sup>13</sup> points at strategic locations. This choice can draw on the previous one, using information that is already available by virtue of the inference process. In fact, each single-point evaluation of the Bayes rule quantifies the worth of the corresponding point (the coordinate(s) of a location in parameter space).

Shifting the disposition of evaluation sites via a serialization of the posterior sampling process would allow us to spare resources, by removing support from the unlikely regions. That way, we could concentrate

<sup>13</sup>The term *grid* is used loosely here.






Figure 21: Posterior distributions for a time-dependent binomial model, for 2 sets of measurement times. The likelihoods take the undamped squared-sinusoidal form of 2.84. The dashed line marks the real parameter used to generate the data. 10 measurements were included in each data vector, and the times were picked in constant increments up to the maximum  $t_{max}$ .

them where they can actually be of benefit. As a bonus, we may to a degree evade ambiguity through preemptive action.

Ideally, we would start with a likelihood like the shallow one in figure 20, and then gradually transition (by dint of data collection) to the sharper one, while re-focusing our grid on the pertinent section of parameter space, where the posterior is non-null. Such a migration could populate this key section more densely, while abandoning its irrelevant surroundings. With that, a likelihood with ambiguity at the scale of 21 should pose no issue: we'd have no coverage for the second peak, having eliminated it from the search space earlier on.

This reflection seems somewhat distant from the previous deliberation on utility. It stands closer to the details of how we represent distributions, in particular when dealing with analytically intractable inference and finite memory resources. While multi-modality or (faux) flatness *would* be penalized if minimizing the expected posterior variance<sup>14</sup>, these problems may be more easily remediable via a tighter integration with the discretization methods. This is the subject of the next chapter.

<sup>14</sup>Excessive sharpness in particular is technically associated with a low variance, but it is fair to assume optimization will operate our representation of the posterior. Since that's where the problem lies, this will reflect in the estimated variance.

## Monte Carlo methods for posterior sampling

Up until now, we've been focusing on the core of the inference process. We were especially interested in how the allotted resources could be used to devise the best possible action plan. But all that is worth nothing, if we fail to seize the legacy of our perfect plan.

This was evident in figure 20, where a particularly good posterior produced particularly bad results. A chimerical dream of optimality is hardly any consolation when we've let it slip between our fingers.

Granted, we could just reinforce the grid we started out with, increasing the number of evaluations from 5 to 10, or 100, or 1000. At some point, we would catch the mode; and we could pack the space as densely as our target precision demanded. However, this is somewhat precipitate: after all, we have four out of five points standing idly on the sidelines, relaying near zero probability. Why should we devote any more computational power to the problem, when we're squandering 80% of it on undeserving calculations?

On this basis, this chapter concerns more resource-effective ways of sampling from the posterior. This is a determining factor for the success of Bayesian learning: it is these samples that provide a way of computing quantities of interest, as the sole carriers of the collected information.

An alternative to the grid based quadrature discussed in subsection 2.2.2 is Monte Carlo integration, a simulation based approach which tends to scale better with dimension and complexity. This chapter overviews a selection Monte Carlo methods, starting with its simplest versions in section 3.1.

Because these algorithms are still unreliable for the task of characterizing a complex posterior, we will then move on to more advanced ones. These are grouped into two classes: [sequential Monte Carlo \(SMC\)](#) algorithms (section 3.2), which adaptively construct a dynamic grid, and [Markov Chain Monte Carlo \(MCMC\)](#) algorithms (section 3.3), which sample from the posterior directly. Within each of these groups, a few structural variations are considered.

For [SMC](#), the main degrees of freedom are the the source of dynamism (the sequence of distributions to be targeted) and the source of adaptivity (the mechanism for the strategic relocation of *grid* points when transitioning from one target to the next). For the former we consider two possibilities: [sequential importance resampling \(SIR\)](#) with the prior as importance function, and [SMC](#) with tempered likelihood estimation. For the latter we consider another two: Gaussian kernel smoothing and Markov kernels.

In the case of [MCMC](#), the variability lies with the construction of the transition kernels. Two variants of Metropolis-Hastings [MCMC](#) are presented: random walk Metropolis and Hamiltonian Monte Carlo. Section 3.4 then goes into some implementation aspects concerning the latter: applicability, anti-correlation measures, and dynamic implementations (namely the no U-turn sampler extension). Finally, section 3.5 discusses subsampling, overviewing variance reduction and pseudo-marginal methods with an emphasis

on Hamiltonian Monte Carlo.

### 3.1 Monte Carlo and importance sampling

In section 2.2.2, we resorted to an elementary grid arrangement to perform numerical integration. For simple problems, a grid may actually be a viable choice - possibly with some improvements, such as the trapezoid rule and the like. But when working with large numbers of parameters and/or expensive likelihood evaluations, such a solution is a waste of resources in the best-case scenario, and a lost cause in the worst. In particular, creating a grid dense enough to capture features of high dimensional spaces is usually quite the thankless task; in these circumstances, a more efficient representation of the distribution is of the essence.

This is due to that fact that as the parameter space's dimension grows, contributions to the expectation integral depend on a more complicated interplay between integrand (or probability) density and space volume. More specifically, the volume tends to accumulate away from the highest probability neighbourhood (or any neighbourhood for that matter). As the parameters grow in number, the imbalance of volume grows exponentially with it. Indeed, this is the very reason why expectation integrals are far superior to optimization in high dimensional spaces: the latter blindly seeks probability density only, which means nothing by itself. These points are rooted in measure theory [12], which underpins statistics.

This discussion isn't unlike that of subsection 2.3.1, where we considered the effect of the dimension of an integral in its net value. In this context, this means that the region over which the integrand is to be accumulated must be selected carefully to achieve optimal balance between accuracy and evaluations. This is essential if one means to preserve the robustness of the expected value framework.

To see how this fits the overarching narrative, we can refer back to our central equation. If we have collected a vector of data  $\vec{D}$ , we can compute the posterior expected value of a function as:

$$\mathbb{E}_{\mathbf{P}(\theta|\vec{D})} [f(\theta)] = \int f(\theta)\mathbf{P}(\theta | \vec{D})d\theta \approx \sum_{i=1}^M f(\theta_i) \cdot w_i \quad (3.1)$$

, where the  $w_i = w_i(\vec{D})$  are our discrete-domain version of evaluations of the posterior  $\mathbf{P}(\theta | \vec{D})$ . We suppress the dependence on the data to lighten the notation.

$$w_i \equiv \frac{W_i}{\sum_{j=1}^M W_j} = \frac{\mathbf{L}(\theta_i | \vec{D})\mathbf{P}(\theta_i)}{\sum_{j=1}^M \mathbf{L}(\theta_j | \vec{D})\mathbf{P}(\theta_j)} \quad (2.44)$$

We have been trying to get a posterior with favorable properties to come by, so we could be confident in the expectations we compute (such as the mean, which we chose as our estimator). But our right-hand side calculation in 3.1 reflects not this posterior, but rather its values at a discrete set of  $M$  points.

$$\{\theta_i\}_{i=1}^M \quad (3.2)$$

Clearly, it is not sufficient that the posterior boasts good properties: it is also vital that these properties be relayed adequately by the  $\theta_i$  of our choosing. Poorly chosen positions will produce negligible contributions, leading to inaccurate expectations.

Before proceeding to analyse and develop our algorithms, we will dub two quantities to facilitate discussion. First, the parameter space points  $\theta_i$  from  $\{\theta_i\}_{i=1}^M$ , which we use to discretize the distribution, are often called *particles* (especially in the context of sequential Monte Carlo algorithms, which we will get to shortly). And second, their associated multipliers in 3.1 are called their *weights*, or importance weights.

Apart from this detail, we had already used these quantities in subsection 2.2.2, where we described a basic procedure to approximate expectations. Algorithm 1 outlines this first approach. It assumes single-parameter inference, though the generalization is straightforward.

**Inputs:** prior distribution boundaries  $\theta_L$  (left) and  $\theta_R$  (right), data vector  $\vec{D}$ , number of particles  $M$ .

**Computes:** list of pairs  $\{\theta_i, w_i\}_{i=1}^M$  to be used in 3.1.

**Assumes:** ability to evaluate the prior distribution and the likelihoods.

1. Space out  $M$  particles  $\theta_i$  evenly in the interval  $\theta \in [\theta_L, \theta_R]$  to get the fixed set  $\{\theta_i\}_{i=1}^M$  (a grid).
2. For each particle  $\theta_i \in \{\theta_i\}_{i=1}^M$ , compute the unnormalized weight given the full dataset.

$$W_i = \mathbf{L}(\theta_i | \vec{D})\mathbf{P}(\theta_i) \quad (2.42)$$

Additionally, accumulate these quantities to get the normalization constant  $C$ .

$$C = \sum_{i=1}^M W_i$$

3. For each  $\theta_i \in \{\theta_i\}_{i=1}^M$ , compute the normalized weights by dividing the former ones by the previously calculated normalization constant.

$$w_i = \frac{W_i}{C} \quad (2.44)$$

Algorithm 1: Grid-based algorithm for computing expected values.

For one-dimensional integrals, improved efficiency can be achieved via interpolation, by using the well-known Newton-Cotes quadrature rules or other more sophisticated formulas. This can also be applied to higher dimensions (and so to multiple parameter estimation problems), through sequential application justified by Fubini's theorem.

However, this quickly becomes infeasible, due to exponential scaling on the number of dimensions: repeated application of 3.1 along  $d$  dimensions will require  $M^d$  function evaluations. In that case, Monte Carlo (non-deterministic) methods tend to be preferable. They rely on random evaluations of the integrand, and their absolute error scales approximately with  $M^{1/2}$ , which is dimension-independent.

Their simplest version could be described by 3.1 still, only the set  $\{\theta_i\}_{i=1}^M$  would now be sampled from a uniform distribution. For the sake of clarity, it should be said that any integral (not just the ones

corresponding to expected values) can be computed like so:

$$\int g(x)dx \approx \frac{V}{M} \sum_{i=1}^M g(x_i) \quad (3.3)$$

In our case, we have seen in section 2.2.2 that normalizing the weights instead of dividing them by the marginal likelihood took care of the  $\frac{V}{M}$  factor, because we were precisely dividing by  $\frac{M}{V}\mathbf{P}(\vec{D})$ .

Unfortunately, the idea of dimension-independence doesn't tell the whole story; these methods still suffer from increased dimensionality. In high-dimensional spaces, we expect that less of our evaluations will recount non-negligible probability, since the probability mass will tend to be more concentrated with respect to the total volume we must cover. This is likely to happen for our *expectation* integrals, where we consider some parametrized distribution. Typically, vast sections of the parameter space will have near-zero probabilities; the more dimensions, the truer this is. The result is an inflation of the total space to explore; should particles fall in these regions, their weights will be insignificant, barely affecting the integral. Thus, the relative error will be amplified by the number of model variables, still frustrating our attempted estimates in complex cases.

The problem can be alleviated by using *importance* sampling. Instead of sampling points from a uniform distribution, we sample them from one that is tailored to our specific integrand. Ideally, we would focus the  $\theta_i$  in the regions that actually contribute to the integral; accumulating the integrand over null probability regions is, of course, futile. A sensible choice will reduce the variability in our estimate, making it more robust for finite  $M$ .

With that in mind, we would introduce a proposal distribution, also called an importance function, that we could sample from:

$$\pi(\theta) \quad (3.4)$$

, and use it to obtain our parameter locations:

$$\theta_i \sim \pi(\cdot) \quad (3.5)$$

In that case, equation 3.1 must be altered to compensate for the points' relative importance, which is a trick to make the samples land around a critical zone and not literal *importance*. That is, this importance is ascribed by our strategy, and not by the integrand. Consequently, if a point is very likely to occur under our choice, that likeliness should be discounted from its contribution; the intended value comes instead from the evaluation of the integrand, and is larger or smaller according to it alone.

For the general case of 3.3, this means:

$$\int g(x)dx \approx \frac{1}{M} \sum_{i=1}^M \frac{g(x_i)}{\pi(x_i)} \quad , \quad x_i \sim \pi(\cdot) \quad (3.6)$$

In hindsight, we can see that uniform sampling was just a particular case of this, with  $\pi(x) = \frac{1}{V}$  (for  $V$  the volume of the domain of integration, which should match that of the uniform distribution) yielding the fixed multiplicative constant  $\frac{V}{M}$ , which in that case was independent of the effective set of samples.

As for our case, 3.1 becomes:

$$W_i = \frac{\mathbf{P}(\theta_i | \vec{D})}{\pi(\theta_i)}, \quad \theta_i \sim \pi(\cdot) \quad (3.7)$$

The normalization to get  $w_i$  is as usual (i.e. divide each  $W_i$  by the sum  $\sum_{i=1}^M W_i$ ). Again, it is easy to see that for the particular case of a uniform distribution these contributions are as before, because the  $\pi(\theta_i) = 1/V$  factors cancel out when normalizing.

In algorithm 1, importance sampling would change two points:

- In step 1, we would sample  $\theta_i \sim \pi(\cdot)$  instead of arranging the particles on a grid.
- In step 2, we would compute the weights  $W_i$  according to 3.7 (for both occurrences).

As an example of importance sampling, we can think of estimating the fairness of a nearly balanced coin: we could sample the points from a Gaussian centered at  $1/2$ . Evaluating the posterior at those sites should be more fruitful than spreading them uniformly between 0 and 1, since we expect the probability to be low away from the center. Had we used uniform samples and small  $M$ , we would expect a significant discrepancy between any two different sets of samples, whereas a Gaussian should produce more consistent results; this testifies to the robustness of the estimation strategy.

How can we develop a *general* approach to importance sampling, that will allow us to compute posterior expectations efficiently? This calls for some insight, but we have long decided to be neutral; our flat prior would suggest precisely uniform sampling.

Overall, a nice importance function would resemble the integrand,  $f(\theta)\mathbf{P}(\theta | \vec{D})$ . We should remember that  $f(\theta)$  can be any function of interest, such as the identity or the mean squared distance from the mean (whose expectations, as seen in subsection 2.2.1, comprise our main targets, the mean and variance). What we want is to have a fairly generic group of  $\theta_i$  - having registered these locations and their associated posterior weights  $w_i$ , we are in a position to retrospectively compute the expectation of whatever function of the parameters.

As such, the *ideal* choice is actually objective: we would like to sample directly from  $\mathbf{P}(\theta | \vec{D})$ , in which case 3.7 would reduce to perfectly distributed evaluations of  $f$  with no multipliers at all. This manifests a dependence on the data-vector  $\vec{D}$ , and so assumes batch-processing. In such a context, we could use the results of our  $N$  experiments to a *posteriori* condition our proposal distribution:

$$\pi(\theta | \vec{D}) = \pi(\theta | \{D_1, \dots, D_N\}) \quad (3.8)$$

This means that we can choose the posterior as the proposal distribution - though not quite. While we can resort to ready-made implementations to sample from some well-known families of functions, we don't know how to *sample* from an arbitrary function just by evaluating it<sup>1</sup>.

What we *can* do is to settle for a different strategy - one grounded on the data still, but aided by some extra contrivance. Expectably, this contrivance is again sequential processing.

<sup>1</sup>Yet; that's left for section 3.3.

## 3.2 Sequential Monte Carlo

*Sequential Monte Carlo* ([sequential Monte Carlo \(SMC\)](#)) is an umbrella term for methods concerned with characterizing the succession of states in a stochastic process, based on consecutive partial observations. A review of such methods is presented in [27], which is the reference we'll follow for this section. They find use in a wide array of applications, notably Bayesian inference.

In that case, the target process easily finds a formalization in the theory of hidden Markov models. Such a model considers two concurrent processes: one is *hidden*, and one is directly observable. The latter depending on the former, its direct observation amounts to making *partial* observations regarding the hidden states<sup>2</sup>. These hidden states are of course the target parameter(s), while the *visible* ones are the measurement outcomes; the *emission probability* that functionally relates them is none other than the likelihood. The discrete time is set by the data collection, e.g. one datum corresponds to a time step.

Put simply, the purpose of [SMC](#) is to sample from a sequence of distributions, in our case the successive posteriors. This is also called *simulating* the successive posteriors. Just like quantum simulation provides the means for characterizing and predicting the behaviour of observable quantum systems, statistical simulation does so for observable statistical phenomena.

One possible application for sequential simulation is obvious: *live* parameter learning, where we want either to capture the evolution of dynamic parameters or to track the estimates real-time. Clearly, this is our case, as we need to follow the data as they arrive in order to adapt the experiments accordingly.

However, these methods are as often used *offline*, purely to aid at the post-processing stage. In that case, the goal is generally to travel across a sequence of slowly changing, increasingly complex distributions and arrive at the intended one, when it would be hard to do so directly. This is *also* our case.

Our target sequence of distributions is comprised of the successive posterior distributions, with single-datum granularity. The  $k$ -th iteration posterior is:

$$\mathbf{P}_k(\theta) \equiv \mathbf{P}(\theta \mid \{D_1, \dots, D_k\}) \quad (3.9)$$

A selection of such local posteriors is presented in figure 22 for the same case as in the previous chapter (the precession example).

This local posterior should be available with no more information (i.e. not rely on  $D_j$  for  $j > k$ ). Again, one possible motivation for this is clear: we may want to choose the next experiment adaptively, or mean to get the up-to-date information as soon as possible. With it, we can get the current expectation for some function:

$$\mathbb{E}_{\mathbf{P}_k(\theta)} [f(\theta)] = \int f(\theta) \mathbf{P}_k(\theta) d\theta \approx \sum_{i=1}^M f(\theta_i) \cdot w_i^{(k)} \quad (3.10)$$

Another rationale in favor of adding data gradually is the fact that more attainable adjustments in the discretization should be due when transitioning between consecutive distributions, since a single datum

<sup>2</sup>This *partialness* could also be due to observing variables through some noise.




Figure 22: Evolution of the posterior probability distribution with the iteration number. The dashed vertical line marks the real parameter used to generate the data. The posterior is depicted for  $k = 5, 10,$  and  $25$  accumulated data, according to equation 3.9. The  $D_k$  were obtained using exponentially increasing times  $t_k = (9/8)^k$ .

isn't expected to bring about massive changes. This is made clear by figure 22, and will shortly be made still clearer.

To add to this fact, in the case of the previous chapter and figure 22 extra flexibility is introduced by experimental freedom. That is, we can sort the data in increasing order of measurement times, to make the distribution evolve more smoothly; and we can increase them as fast or as slowly as we wish, in order to control the spread of the distribution through the iterations. If we maximize some utility function, we will be using our particle cloud to evaluate it, and so numerical precision will be minded by default.

Having outlined the appeal of a sequential sampling scheme, algorithm 2 lays out the associated procedure. It uses the same number of likelihood evaluations as 1, but it processes the data sequentially.

We now want to modify this to use sequential *importance* sampling, in order to improve the positioning of the particles. Through such improvements, we hope to soothe our previous concerns relating the discretization of the distribution.

In the context of SMC, the problem we seek to treat is generally called the *particle degeneracy* problem. As the inference process advances, the posterior probability becomes more and more concentrated with respect to the parameter space volume, leading to increasing disparity between the particle weights. As such, if we start out with  $M$  uniform uncorrelated samples, our *effective* number of samples will shrink down to a small few with time, due to particle collapse: most particles will cease to provide any useful information after a point.



**Inputs:** data vector  $\vec{D}$  (to be used incrementally:  $\{D_1\}, \{D_1, D_2\}, \dots, \{D_1, D_2, \dots, D_N\}$ ), number of particles  $M$ .

**Computes:** successive lists of pairs  $\{\theta_i, w_i^{(k)}\}_{i=1}^M$  for (by order)  $k \in \{1, \dots, N\}$  to be used in 3.10.

**Assumes:** ability to sample from the prior and to evaluate the likelihoods.

**Assures:** The  $k$ -th list is computed in the  $k$ -th iteration, using the datum  $D_k$  and before processing  $D_{k+1}$ .

1. Sample  $\{\theta_i\}_{i=1}^M$  from the prior distribution and set  $w_i^{(0)} = 1/M$  for all  $i \in \{1, \dots, M\}$ .
2. For  $k \in \{1, \dots, N\}$ :

- a) For each particle  $\theta_i \in \{\theta_i\}_{i=1}^M$ , update the weight using the latest datum  $D_k$ .

$$W_i^{(k)} = w_i^{(k-1)} \cdot \mathbf{L}(\theta_i | D_k) \quad (3.11)$$

Additionally, accumulate these quantities to get the normalization constant  $C$ .

$$C^{(k)} = \sum_{i=1}^M W_i^{(k)}$$

- b) For each  $\theta_i \in \{\theta_i\}_{i=1}^M$ , compute the normalized weights by dividing by the previously calculated normalization constant.

$$w_i^{(k)} = \frac{W_i^{(k)}}{C^{(k)}} \quad (2.44)$$

Algorithm 2: Sequential algorithm for computing expected values. A  $(k)$  superscript signals dependence of the marked quantity on the first  $k$  data and no more.

This is due to the nature of inference: if all goes well, as the data pile up, the true posterior *should* converge to an each time sharper peak around the real value. It is this fact what allows us to get results, and to trust them. But if we don't take heed when approximating the distributions, this same fact can be our undoing.

The achievable resolution will be determined by the disposition of the particles, as will how sharp a distribution we can capture. In general, it's safe to say that a coarse grid is a severe performance bottleneck, and undermines the credibility of the results - if they surface at all. What is more, this coarseness is relative, and bound to become a reality sooner or later.

In the case where we work with a fixed grid, after some precision threshold, the particles that fall closest to the peak are bound to amass all the weight, causing all others to yield negligible contributions (remembering that the weights are made to sum to one). Should we proceed further, we are likely to have a single particle be responsible for our estimates, as the others get sidetracked with (rather expensive) computations of zero.

Finally, if we add in still more data, even this resistant particle will find zero probability eventually. By then, we will encounter a zero-division indetermination when normalizing it, and be forced to stop. Even if one particle is lucky enough to fall very close to the asymptotic mode by chance - which we clearly should

not rely on  $\cdot$ , any estimates derived from it will be defective. In the end, if we take the inference process far enough, we're unlikely to be left with much more than the scant information that the particles we started out with are not plausible parameter specifications.

Figure 23 illustrates this problem, using the posteriors displayed in figure 22. After the first curve, the weights are still fairly even; however, we can see that after the second the two leftmost particles have grown disproportionately, while the others wither away. This disparity foreshadows what will happen next: the grid is unfit to capture further information, because at all of its points the posterior probability is nearly zero (often it is for all purposes zero, due to arithmetic underflow).




Figure 23: Illustration of the particle degeneracy/weight collapse problem, prompted by the particle cloud's evolution as the iterations (and with them the data) accumulate. The distributions are the same as in figure 22. The 'x' markers stand for particles, their sizes representing their weights.

It being that we proposed importance sampling as a possible remedy for this type of issue, we will try to incorporate it into algorithm 2. For each iteration  $k$ , we can importance sample *new* particles instead of recycling the previous ones. We can do so just before step 2a of iteration  $k$ ; and naturally, we can utilize the data supplied so far.

$$\theta_i^{(k)} \sim \pi(\cdot \mid \{D_1, \dots, D_k\}) \quad (3.12)$$

As a consequence, the weight calculations will be done from scratch at the new particle locations, and must be divided by the importance function instead of just being updated according to the latest datum

(as before in step 2a).

$$w_i^{(k)} \propto \frac{\mathbf{P}_k(\theta_i^{(k)} \mid \{D_1, \dots, D_k\})}{\pi(\theta_i^{(k)} \mid \{D_1, \dots, D_k\})} \quad (3.13)$$

The values of the final weights are considered up to normalization for simplicity. Any scaling factor will be eliminated when normalizing, so determining the weights apart from a common normalizing constant is sufficient.

Even if we have some convenient importance function in mind, this is very demanding in terms of resources. We get the sequence of distributions we want and can use the past data to inform the importance distribution, but the price to pay is an immense cost as compared to our previous implementation. If before we took on the calculations for  $N$  data, now we will do the same for 1 datum, then for 2, and so on until we *finally* get to our last iteration - which alone is as costly as the full-length procedure previously was.

At this point, it is also not very clear why the situation that motivated our switch to a *sequential* routine should be any better with it. Our problem was the importance function, or lack thereof. We started out with the purpose of ameliorating the case of 3.8, after observing that using the posterior as the importance distribution was infeasible and we had no other clues upon which to base it. What has changed as compared to then?

Two things have. At the time of sampling for the  $k$ -th iteration:

- The available dataset at the time of importance sampling is not necessarily the full data-vector  $\vec{D}$ , but rather its  $k$  first elements;
- Before computing the set  $\{\theta_i^{(k)}, w_i^{(k)}\}_{i=1}^M$ , we have computed its precursor  $\{\theta_i^{(k-1)}, w_i^{(k-1)}\}_{i=1}^M$ .

The first point means that the best importance function we could wish for at the  $k$ -th iteration is the posterior  $\mathbf{P}_k(\theta)$  considering the first  $k$  data, and not the final posterior. Here we encounter the same problem as before: we can't sample from the posterior.

However, thanks to the second point, we are *much* closer. Equation 3.12, doesn't consider all the available resources at the time of its use. We can work in a dependence in all the preceding lists of particle-weight tuples, which will be all set by then. In particular, we can use the latest one, as we pointed out. Having just completed some iteration  $k$ , we have

$$\{\theta_i^{(k)}, w_i^{(k)}\}_{i=1}^M \quad (3.14)$$

In fact, this particle cloud is a weighted sample from  $\mathbf{P}_k(\theta)$ , as made clear by equation 3.10. If we were in possession of uniform samples from  $\mathbf{P}_k(\theta)$  to use as importance samples in the  $(k+1)$ th iteration,  $\{\theta_i^{(k+1)}\}_{i=1}^M$ , we could choose this  $\mathbf{P}_k(\theta)$  as our importance distribution, and the weights to be calculated in 3.13 for the  $(k+1)$ th iteration would become simply:

$$w_i^{(k+1)} \propto \frac{\mathbf{P}_{k+1}(\theta_i^{(k+1)})}{\mathbf{P}_k(\theta_i^{(k+1)})} \propto \mathbf{L}(\theta_i^{(k+1)} \mid D_{k+1}) \quad (3.15)$$

That is, we would be a single datum away from the target, in which case this datum's likelihood is all that is needed for the iteration in question.

Were it not for the issues with reusing particle locations, we would just take the weighted sample and multiply the old weights by the likelihood, recovering algorithm 2; this is enough to actualize the transition from  $\mathbf{P}_k(\theta)$  to  $\mathbf{P}_{k+1}(\theta)$ . But we have seen the problem with that, so we want to proceed differently, somehow getting different - better - importance samples for the following operations.

Achieving this via the latest particle cloud would be doubly beneficial: on the one hand we would avoid repeating likelihood evaluations, and on the other the changes in the discretization would consist of smaller *corrections* based on the information gathered since the previous alterations.

Essentially, we would want to base these smaller-scale corrections on 3.14, transferring the dependence of 3.12 on the data to the particles in hand (except for the last datum). This leaves us with the last piece of the puzzle - we must figure out how to use these particles.

### 3.2.1 Particle propagation and the Liu-West filter

At this point, we are in possession of a particle cloud representing a weighted sample from the latest posterior, and from it mean to get a fresh sample for the following one.

$$\theta_i^{(k+1)} \sim \pi(\cdot \mid \{\theta_i^{(k)}, w_i^{(k)}\}_{i=1}^M) \quad (3.16)$$

More specifically, we want a new and improved set  $\{\theta_i^{(k+1)}\}_{i=1}^M$ , whose reweighting will be more fruitful than if we had reused the old locations. For that, we can *recycle* poorly situated particles, summoning them to critical regions.

Equation 3.16 raises two questions. First, the samples we have access to are weighted, whereas our importance samples up to this point were simple parameter vectors with no associated weights. We expect our importance density to place more particles where it is higher, but we don't want it to give us a high probability point and say it counts for 10.

Which takes us to our second point - we want *variability*. We want to mobilize our particles away from the barren land that most witnessed, in order to sift through the interesting parts of parameter space. How can we make them decamp to those regions if all we have is a discrete set of old positions, some of which to top it off are given an inordinate amount of *importance*?

As for the weight imbalance point, it can be solved easily. We just need to replace the weighted set of particles by an unweighted one (with uniform weights  $w_i = 1/M$ ) according to the particle weights. For instance, we could divide the  $M$  new particle locations by assigning  $M \cdot w_i$  particles to each  $\theta_i$ . Since  $M \cdot w_i$  will not in general be an integer value, the distribution is better preserved on average by drawing  $M$  samples from a multinomial distribution with  $w_i$  probability of getting  $\theta_i$  ( $w_i \leq 1$  given that  $\sum_{i=1}^M w_i = 1$ ). This is known as a bootstrap filter.

By doing this, we convert the weighted sample of particles to one where probability is conveyed by particle density. Naturally, even though this simplifies the problem, there is still an imbalance: the higher

weight particle(s) will monopolize these new samples just as they did the weight. As a consequence, we expect to have multiple particles at some sites; what is more, these sites will be no more varied than they were before.

Clearly, there's still something missing, as illustrated in figure 24. We have a weighted sample from the last iteration (e.g. the first, after we reweight the initial uniform samples according to their first datum likelihoods). Then we resample them with a bootstrap filter. Then what?




Figure 24: Illustration of the obtention of a weighted particle set and the necessary treatment to generate an efficient unweighted sample. This corresponds a segment of 23, with a change of the sampling techniques. The bottom clipping displays the following posterior, to be targeted by the darkest particle set. The 'x' markers stand for particles, their sizes representing their weights.

We must *propagate* the particles somehow, but we can't do that if we adhere to our discrete set. No matter how well it represents the current distribution in terms of expectations, it will never be continuous like itself, and so it won't allow the particles to glide away from their original spots.

One possibility is to perturb the particles about their sites, using for example a Gaussian distribution. By allowing the particles to roam around, we are likely to place them at different points in parameter space, just as we wanted; and by doing so stochastically, we can expect that the overlaid ones will move into different positions.

This amounts to interpreting the weighted sample as a mixture of Gaussian distributions. Of course, we could just as well use other distributions to transform our particles into a smooth density. These distributions are commonly called kernels, and the method is called kernel smoothing. This is described in [73], which is the main reference for the entirety of this subsection.

With Gaussian kernels, our smoothed distribution would be of the form:

$$\mathbf{P}_k(\theta) \approx \sum_{i=1}^M w_i^{(k)} \mathcal{N}(m_i^{(k)}, s_k^2) \quad (3.17)$$

We have already taken care of the weights by bootstrapping, so finalizing the kernel shape is all that's left. The reason why the Gaussian hyperparameters have been left vague is that they're tunable degrees of freedom. One obvious choice for the means is the previous particle location. After applying the bootstrap filter, we would take each sample's location as the mean for a Gaussian, sample from that Gaussian, and fix a particle at to this sample's coordinates for the  $(k + 1)$ th iteration, instead of keeping it at its original site. We would proceed like such for all particles, after which we would have a more diverse sample.

Caution should be taken however in ensuring it doesn't get *too* diverse, in the sense that we're injecting noise. The source is the Gaussians' spread, which we're yet to determine. Since we have no reason to distinguish between particles in this case, we choose a fixed standard deviation  $s_k$  that depends only on the iteration.

Ideally, it would decrease as the iterations accumulate, as we don't want to toss the particles around too aggressively: we risk scattering them through the parameter space, thereby losing information (and possibly throwing them back into low probability regions). One natural way of informing this hyperparameter is to fall back on the latest particle cloud's standard deviation, namely using a proportion:

$$s_k = h\sigma_k \quad (3.18)$$

Here  $h$  can itself be a function of the iteration number. This is also applicable in higher dimensions, where the covariance matrix would take the place of the variance. This only furthers the convenience of this idea, as it automatically considers the relative scale between parameters and their learning rates.

Even though  $\sigma_k$  and  $s_k$  are both standard deviations, they play very different roles. The former is associated with the **SMC** estimate for the variance of  $\theta$  given the distribution  $\mathbf{P}_k$  after the reweighting,

$$\sigma_k = \sqrt{\mathbb{V}_{\mathbf{P}_k}[\theta]} \quad (3.19)$$

, which in turn estimates the expected mean squared error; the latter determines the distance between the new particles and the old ones (or its order of magnitude).

Together, they will determine the *new* **SMC** standard deviation, which comes to replace the former  $\sigma_k$  and concerns the revitalized group of particles. Under our circumstances, it is given by:

$$\sigma'_k = \sqrt{1 + h^2} \cdot \sigma_k \quad (3.20)$$

This affirms our concerns. Except for the case of  $h = 0$ , where we stay by the bootstrap filter, we are increasing the uncertainty by dispersing the particles. While this may not seem very serious for small  $h$ , this over-dispersion will be amplified at each iteration we put this in practice, partially canceling out the efforts put into gathering and processing the data.

A technique called *kernel shrinkage* has been proposed as a restorative. The idea is that to compensate for stirring the particles, we pull them slightly towards the SMC mean. The more we agitate them, the more we impel them in the direction of the center, closer to the middle of the particle cloud. More specifically, to each particle we assign the kernel mean:

$$m_i^{(k)} = a\theta_i^{(k)} + (1 - a)\mu_k \quad (3.21)$$

These means will preserve the variance for  $a = \sqrt{1 - h^2}$ . Solving for  $h$ ,

$$s_k = \sqrt{1 - a^2} \cdot \sigma_k \quad (3.22)$$

That is, in priming each particle location  $\theta_i^{(k)}$  for the next iteration, we will be sampling from:

$$\theta_i^{(k+1)} \sim \mathcal{N}(m_i^{(k)}, s_k^2) \quad (3.23)$$

Putting everything together, we have a closed form for our smoothed density of 3.17, and we can sample fresh particles from it for the following iteration. We can do so before step 2a of the  $(k + 1)$ th iteration, as mentioned earlier, or after step 2b of the  $k$ -th iteration.

The choice may vary depending on the source, seemingly with an inclination toward the latter option, which we will therefore opt for. However, this is a matter of taste, since either way the samples are based on the same information and available when needed. The difference is otherwise not very relevant; fresh samples are unlikely to be a necessity in the end of the last iteration as in the beginning of the first one.

With this option, at the end of each iteration we get new samples for the following one according to the aforementioned smoothed density:

$$\theta^{(k+1)} \sim \sum_{i=1}^M w_i^{(k)} \mathcal{N}(m_i^{(k)}, s_k^2) \quad (3.24)$$

This describes the Liu-West *resampler* or *filter* (LWF). It is named after its two authors Liu and West, who both developed the underlying ideas and presented this particular algorithm [73]. The *filter* designation comes from the fact that SMC samplers are ordinarily called *particle filters*, due to how they carry through particle selection.

The method is summarized in algorithm 3. No bounds are enforced for the particles' motion, though one could easily exclude points without posterior support; the particles are not expected to stray much within this scheme, and are unlikely to move away from the originally covered region at all.

SMC samplers making use of this resampling strategy have often been adopted for quantum characterization examples. The method is presented in detail in [51], and applied and/or briefly discussed in many others (e.g. [117, 49]). It has also been used in experimental implementations [113]. A strong point in favour of the Liu West resampler is that it doesn't require likelihood evaluations, which is rarely the case for other methods.

On the flip-side, it only preserves the first two moments of the distribution (mean and variance, in our case representing the estimate and uncertainty respectively), which is not always sufficient. In particular,

**Inputs:** weighted sample  $\{\theta_i, w_i\}_{i=1}^M$ , resampling parameter  $a$ .

**Computes:** new equally weighted positions.

**Assumes:** Ability to sample from multinomial and Gaussian distributions.

**Assures:** The new sample has the same mean and variance as the original one.

1. Compute the current mean and variance of the particle cloud.

$$\mu_k = \sum_{i=0}^M w_i \theta_i \quad (3.25)$$

$$\sigma^2 = \sum_{i=0}^M w_i \theta_i^2 - \mu^2 \quad (3.26)$$

Define

$$s = \sqrt{1 - a^2} \cdot \sigma \quad (3.27)$$

2. For  $i \in \{1, \dots, M\}$ , get a sample  $\theta'_i$ :

- a) Sample a particle  $\theta_s$  from the original set  $\{\theta_i\}_{i=1}^M$  with probabilities given by their weights  $\{w_i\}_{i=1}^M$  (i.e. a multinomial distribution), and calculate

$$m_i = a\theta_s + (1 - a)\mu_k \quad (3.21)$$

- b) Sample  $\theta'_i$  from a gaussian centered at  $m_i$

$$\theta'_i \sim \mathcal{N}(m_i, s^2) \quad (3.28)$$

Algorithm 3: Liu-West resampling algorithm. The calculations assume  $\theta$  to be one-dimensional, but the generalization is straightforward.

it is utterly unfit for multimodal posteriors, except for the uninteresting case where  $a = 0$ . This also means it may converge to wrong values in the case of redundancy, which it won't be able to recover from. Furthermore, it can't be expected to yield very accurate expectations. These problems may be relieved by using more particles and resampling less often, but this is a precarious solution; not only does it defeat the purpose of using such a lightweight method, it is also bound to fail in more complex scenarios. A more sophisticated resampling strategy will be mentioned in subsection 3.3.3.

The behaviour with respect to other moments depends on the hyperparameter  $a$ , which should be adapted to the application, namely the characteristics of the likelihood function. This quantity negotiates a trade-off between variety and preservation. For  $a = 1$ , we are back to the bootstrap filter, which draws random samples (with replacement) from the original particle locations. At the other extreme, choosing  $a = 0$  would amount to making a more drastic normality assumption: we would ignore the individual particle locations and sample from a Gaussian with the same mean and variance as the particle ensemble<sup>3</sup>.

<sup>3</sup>It's worth noting that under such an assumption more lightweight strategies may be suitable, eliminating the need to store the entire particle cloud. For instance, [115] proposes a [Gaussian rejection filtering \(GRF\)](#) approach to phase estimation, where samples are produced as for  $a = 0$  and the reweightings (Bayes' updates) are replaced by an acceptance-rejection step. In that case, the particles can be treated sequentially and discarded after their contributions to quantities of interest have been accumulated, greatly easing the memory requirements. The mean and variance can be calculated from 4 common counters



High values of  $a$  are customarily picked so as to mostly preserve the structure of the particle cloud. Most references above choose  $a = 0.98$ , with the original authors applying an even higher value (0.995), which is fitting for models with pronounced non-Gaussianity. For others, a lower value may improve convergence; 0.9 is suggested in [117] for that reason.

One more thing to observe is that the filtering interrupts the parallelism of the SMC algorithm. While the particles can be separately reweighted (up to a constant factor) and propagated, the probabilities for the multinomial sampling step, which must be determined between these two tasks, depend on the whole particle group. In contrast, algorithm 1 could be executed in parallel up to the normalization step, and algorithm 2 would be serialized only by the demands of the application; the normalizing constant's use can be deferred until an expectation needs to be computed.

Finally, it should be noted that not all sources define the scope of the term *resampling* equally. It aims to emphasize the fact that we are reworking the sample we had to begin with, but what it entails may vary. In our case, equation 3.17 suggests that we join the multinomial sampling together with the normal perturbations into a unified resampling step of which the bootstrap filter is a particular case.

Nevertheless, just as often the *resampling* and *particle propagation* or *move* steps are split up in SMC algorithms, the former referring exclusively to bootstrapping. In other words, *resampling* can refer to the multinomial sampling alone [102, 54], or it can integrate the variety-introducing element as well [51]. The first case commonly occurs when the fusion isn't as natural as it is in the case of our mixture density assumption (subsection 3.3.3), which for convenience we sample from in two steps.

### 3.2.2 Sequential importance resampling

We are now nearly all set for our sequential importance (re)sampling algorithm, having elected a resampling strategy. But before adopting this new-found strategy, it is convenient to ask whether we actually want to use it.

In truth, it depends. Even leaving aside the chosen resampler's characteristics, we may not want to put it into action at every step of the algorithm. Instead of processing one datum at a time to get a weighted sample and immediately resampling, we may want to ponder whether to do so. Depending on the situation, and in particular on the effect of the fresh data, we could instead add a sequence of a few data before altering the particle locations; for that sequence, the weights would be updated directly as in algorithm 2.

Such a wait is justified when these data have caused slow variations in the distribution, making the particles as-is fit to convey its picture. In that case, spending resources to shuffle the particles around is unnecessary - and possibly even harmful, since the overreliance on circumstantial evidence may result in their unsubstantiated relocation.

This is made evident by looking at figure 23 (along with 22). After the first curve, which includes not one but *five* data, resampling is uncalled for. The weights are still fairly even, which indicates a still instead of  $M$  particles. Results of the application of this strategy can be found in subsection 4.2.1.

representative sample. Only after the second curve - where 5 more data have been added, making up a total of 10 - have the weights become unbalanced.

Somewhere between the 5<sup>th</sup> and the 10<sup>th</sup> distributions, we may have benefited from a resampling step, which would have helped prevent the sample's being virtually reduced to a meager two particles; notwithstanding, to include one such step at every iteration may be excessive.

Motivated by this, we may want to opt for a *dynamic* approach, where we decide whether to resample depending on the current particle cloud's state. A commonly used diagnostic criterion for weight disparity is the [effective sample size \(ESS\)](#), or an approximation thereof. This is a widely used criterion, and is presented in e.g. [83].

As the name suggests, the [ESS](#) tells us how many exact samples our  $M$  imperfect samples would amount to, thus assessing the statistical efficiency of the weighted sample. In general, it penalizes correlation; in our case, this is naturally expressed by the weights, as the dominance of some samples means they *count for several* while being only one.

$$\widehat{\text{ESS}} = \frac{\left(\sum_{i=1}^M W_i\right)^2}{\sum_{i=1}^M W_i^2} = \frac{1}{\sum_{i=1}^M w_i^2} \quad (3.29)$$

With this definition,

$$1 \leq \widehat{\text{ESS}} \leq M \quad (3.30)$$

After the initialization, we have  $M$  particles with uniform weights, so  $\widehat{\text{ESS}} = M$ . The same holds immediately after a resampling step. For all other possible ways of assigning the weights, the [ESS](#) is smaller: non-uniformity introduces a penalty.

It is easy to see that if  $M_X$  of the  $M$  particles have among them uniform weights  $1/M_X$ , and the other  $(1 - M_X)$  have zero weights, the [ESS](#) is  $M_X$ , which makes sense since only the non-zero weight particles contribute to the estimates. For instance, if all but two particles have collapsed, and those two particles have each a weight of  $1/2$ , the sample size is 2, as should be. In general, the more drastic the weight inequality, the lower the [ESS](#). This measures the progressive deterioration of our initially even sample: a low sample size is symptomatic of particle degeneracy.

Generally, some threshold is chosen for the [ESS](#), below which a resampling step is triggered. This threshold is often given in terms of the original particle number,  $M_{\text{thr}} = \epsilon M$ . In [28, 51] and many others a value of  $\epsilon = 0.5$  is suggested, though sometimes more conservative values are imposed.

It should be noted that the [ESS](#) is not always controlled through adaptive resampling. There are alternative ways of keeping it around or above a target value, such as choosing the sequential distributions adaptively based on a *look-ahead* into prospective posteriors [54]. Such alternatives do however tend to be more resource-intensive, whereas this one barely adds to the overall cost of the algorithm.

Finally, we point out that alternative resampling criteria may be used, e.g. the entropy of the weights [28].

Algorithm 4 lays out the steps of the [SIR](#) algorithm. It also goes by the name of sequential importance sampling (SIS) or a more wordy sequential importance sampling and resampling (SISR).

**Inputs:** data vector  $\vec{D}$  (to be used incrementally:  $\{D_1\}, \{D_1, D_2\}, \dots, \{D_1, D_2, \dots, D_N\}$ ), number of particles  $M$ , resampling parameter  $a$ , threshold effective sample size  $M_{\text{thr}}$ .

**Computes:** successive lists of pairs  $\{\theta_i^{(k)}, w_i^{(k)}\}_{i=1}^M$  for (by order)  $k \in \{1, \dots, N\}$  to be used in 3.10.

**Assumes:** ability to sample from the prior and to evaluate the likelihoods, access to a resampling routine.

**Assures:** The  $k$ -th list is computed in the  $k$ -th iteration, using only the datum  $D_k$  and the particle history (namely the last generation of particles).

1. Sample  $\{\theta_i^{(0)}\}_{i=1}^M$  from the prior distribution and set  $w_i^{(0)} = 1/M$  for all  $i \in \{1, \dots, M\}$ . Set  $\{\theta_i^{(1)}\}_{i=1}^M = \{\theta_i^{(0)}\}_{i=1}^M$ .

2. For  $k \in \{1, \dots, N\}$ :

a) For each particle  $\theta_i^{(k)} \in \{\theta_i^{(k)}\}_{i=1}^M$ , update the weight using the latest datum  $D_k$ .

$$W_i^{(k)} = w_i^{(k-1)} \cdot \mathbf{L}(\theta_i^{(k)} | D_k) \quad (3.31)$$

Accumulate these quantities to get the normalization constant  $C^{(k)} = \sum_{i=1}^M W_i^{(k)}$ .

b) For each  $\theta_i^{(k)} \in \{\theta_i^{(k)}\}_{i=1}^M$ , compute the normalized weights  $w_i^{(k)} = W_i^{(k)} / C^{(k)}$ .

c) Calculate the approximated effective sample size

$$\widehat{\text{ESS}}_k = \left( \sum_{i=1}^M [w_i^{(k)}]^2 \right)^{-1} \quad (3.32)$$

d) If  $\widehat{\text{ESS}}_k < M_{\text{thr}}$ , resample the particles  $\{\theta_i^{(k)}, w_i^{(k)}\}_{i=1}^M$  to get a new set  $\{\theta_i^{(k)}, 1/M\}_{i=1}^M$ .

e) Set  $\{\theta_i^{(k+1)}\}_{i=1}^M = \{\theta_i^{(k)}\}_{i=1}^M$ .

Algorithm 4: Sequential importance resampling (SIR) algorithm (with the prior distribution as importance function). A  $(k)$  superscript signals dependence of the marked quantity on the first  $k$  data and no more.

In spite of the fact that the exposition has been focused on a particular instance of the SIR algorithm, the method allows far more generality. There are multiple things that could have been chosen differently, such as the sequence of distributions, the means of monitoring the sample's representativeness, the resampler itself, and others.

In particular, we took the importance distribution to be the prior (a tentative latest posterior), and highlighted the convenience of this choice. Even though this is the most widely used strategy, it should be mentioned that others exist - and more importantly, they do while still enabling recursion in the weight computations. That is, a different choice won't necessarily involve reanalysing the accumulated data at each iteration.

The mechanism behind this is fairly unsurprising: it suffices to take a slightly broader view of what we just did. Just like we used the particle cloud to leave only the last datum's likelihood in equation 3.15, we can use it to *deduct* the prior probability from the calculations, even if we don't want to appoint it directly as the importance distribution. We can formalize our particle representation of the prior by means of a

transition equation:

$$\mathbf{P}(\theta^{(k+1)} \mid \theta^{(k)}) \quad (3.33)$$

This integrates our Markovian model, dictating the evolution of the states. In our case, it encodes our particle-based representation of the current distribution: the prior probability at the  $(k + 1)$ th iteration. The point is that this estimate allows us to cancel out the contributions of previous data (within our model; outside of it, this annulment is approximate).

With this, we can jump from one iteration to the next and maintain some continuity, whatever be our importance function. This is done (as before) by attributing an individualized importance distribution to each particle - for instance, by placing a kernel at its position.

$$\theta_i^{(k+1)} \sim \mathbf{P}(\cdot \mid \theta_i^{(k)}) \quad (3.34)$$

In the following iteration, since the new particles are generated from the old ones and exist in one-to-one correspondence, the weights can be reused. That of each particle is inherited by its offspring, and the new weights can be calculated - *updated* - as:

$$w_i^{(k+1)} \propto w_i^{(k)} \frac{\mathbf{L}(\theta_i^{(k+1)} \mid D_{k+1}) \cdot \mathbf{P}(\theta_i^{(k+1)} \mid \theta_i^{(k)})}{\pi(\theta_i^{(k+1)} \mid \dots)} \quad (3.35)$$

, where the importance function can be conditioned both on the data and on the particles.

In our case, it precisely corresponded to 3.33, yielding:

$$w_i^{(k+1)} \propto w_i^{(k)} \cdot \mathbf{L}(\theta_i^{(k+1)} \mid D_{k+1}) \quad (3.36)$$

For uniform weights (e.g. after applying a bootstrap filter), this reduces back to equation 3.15. If furthermore we use gaussian kernels in the transition density, we can recover our previous resampling strategy. If on the other hand we interpret the discrete sample of the prior more literally and use Dirac-delta kernels (as per 2.48), this takes us back to the basic reweighting we started out with in algorithm 2, with no repositioning. Other approaches can be found in the literature; for instance, [24] suggests using a (leave-one-out) **kernel density estimate (KDE)** as the transition density for the weight updates, despite not using kernel smoothing when refreshing the particle locations. If the particle cloud is a poor representation of the previous distribution (due to failure to approach equilibrium), this may allow it to recover for the subsequent iterations.

In conclusion, the **SIR** algorithm systematizes approximate inference in a flexible manner, without overt application-dependent restrictions (namely, without assumptions about the distributions, excluding those imposed by the employed resampler). By bringing the state of knowledge up to speed as soon as possible, it at once enables adaptive experimental design - or other *online* assessments - and conducts the discretization so as to preserve its aptness, effectively revising it according to the turns of events. In doing so, it tackles the issues brought about by finite numerical precision, averting the unavoidable fate of any *a priori* discretization: powerlessness in the face of ever-growing information.

### 3.2.3 Tempered likelihood estimation

In the previous section, we considered targeting a sequence of increasingly narrow distributions. The benefits were evident: adaptivity played a pivotal role in helping the particle cloud keep up with this narrowness, acting as a stand-in for a much larger amount of resources.

It was this breaking up of the inference process what allowed us to bring a resampler into action when needed. As we saw, this isn't necessarily always. Sometimes, progress happens too slowly to warrant it, so we may want to perform several updates before acting on it. That was considered when conceiving the algorithm; yet, we failed to consider the opposite extreme. What if progress happens too *fast* to act on it?

This is even more dangerous, because taking a too big leap may lead to particle degeneracy and taint the results. Unfortunately, it is also harder to solve, because we don't have as much freedom at the other end of the spectrum. We considered adding data one by one, which is the slowest achievable pace in this setting. However, this definition of the minimum time-step is rather arbitrary, in the sense that we don't necessarily exert control over the effect each datum induces.

Generally, the data don't make guarantees about the mildness of their alterations to the parameter space landscape. If in two different problems there is a ten-fold difference in an individual datum's impact, in one of them the resampling schedule is ten times more coarse-grained. This may well mean that resampling in between data isn't enough.

In the precession example, the effect of each datum could be manipulated with relative ease, since we could choose the evolution time in a continuum and it affected the posterior in a predictable manner. Sadly, this is not always the case. Often the likelihood function is much too complex to anticipate how it will behave, and the freedom in choosing experimental controls isn't enough to moderate its influence on the distribution.

Moreover, other hard to solve intricacies can arise. Under the presence of multimodality, be it intrinsic or due to redundancy, there is a danger of localized particle depletion. We can think of a process that randomly switches between two unknown regimes  $A$  and  $B$ , corresponding to two different parameter specifications and thus two different asymptotic modes. Further, by hypothesis, the processes induce close to mutually exclusive observations, conceivably because the modes are far apart.

Let's say that we collect a first datum while regime  $A$  was operative, and that it presents evidence in favor of a region containing the parameter vector associated with  $A$ . The particles stationed close to that of  $B$  will get zero or near zero weights, and are likely to be driven away from it. As a result, even if we next get a datum that would uphold  $B$ , there won't be support in the distribution to register that. Clearly, our approach is ill-suited to this type of scenario, and will do very poorly.

Again, this didn't present much of an issue in the precession frequency case: not only was multimodality periodic and unlikely to ever exclude the true parameter, it also tended to arise for longer times, which we could handle with care. But naturally, we can't expect this to come to pass in other applications; and even in this one, such a heedless choice may trap the sampler in local minima.

Such problems can be solved by choosing the sequence of target distributions more carefully. A

popular solution is known as *tempered* or *annealed* likelihood estimation (TLE), presented in detail in [89] and briefly described in e.g. [87, 102]. In this approach, we consider a sequence of  $S$  target distributions (not counting the prior) that are nearly proportional to powers of the posterior:

$$\mathbf{P}_s(\theta) \propto \mathbf{L}(\theta \mid \vec{D})^{\gamma_s} \mathbf{P}(\theta) \quad , \quad 0 \leq \gamma_s \leq 1 \quad (3.37)$$

The goal being to smooth the function and the prior being smooth already, making it more level through an exponent on  $\mathbf{P}(\theta)$  would serve no purpose; as such, this treatment is reserved for the likelihood, which is the challenging item.

The exact distributions are naturally given by:

$$\mathbf{P}_s(\theta) = \frac{\mathbf{L}(\theta \mid \vec{D})^{\gamma_s} \mathbf{P}(\theta)}{\int \mathbf{L}(\theta \mid \vec{D})^{\gamma_s} \mathbf{P}(\theta) d\theta} \quad , \quad s \in [0..S] \quad (3.38)$$

The  $\gamma_s$  are the tempering or annealing coefficients, also called the (inverse) *temperatures*. They must satisfy:

$$\gamma_0 = 0 < \gamma_1 < \dots < \gamma_S = 1 \quad (3.39)$$

For  $s = 0$  we have our prior distribution as before,  $\mathbf{P}(\theta)$ . For  $s = 1$  we have our posterior distribution (also as before),  $\mathbf{P}(\theta \mid \vec{D}) \propto \mathbf{L}(\theta \mid \vec{D})^{\gamma_s} \mathbf{P}(\theta)$ .

The  $s$  indices are chosen to avoid confusion with the  $k$ s from before. In that case, a datum  $D_k$  was assigned to each iteration  $k$ . Now, *all*  $D_k$  for  $k \in \{1, \dots, N\}$  are used at *each* iteration, and the sequential distributions differ solely in the exponentiation of their common likelihood function.

Using algorithm 4 as a reference, the only thing that changes is the weight updating step. As before, to jump from the  $s$ -th iteration to the  $(s + 1)$ th one, we must multiply by  $\mathbf{P}_{s+1}(\theta)$  and divide by  $\mathbf{P}_s(\theta)$  (to compensate for the samples' being distributed according to the latter, as in equation 3.6).

$$w_i^{(s+1)} \propto w_i^{(s)} \cdot \frac{\mathbf{P}_{s+1}(\theta_i^{(s+1)})}{\mathbf{P}_s(\theta_i^{(s+1)})} \propto w_i^{(k)} \cdot \mathbf{L}(\theta_i^{(s+1)} \mid \vec{D})^{\gamma_{s+1} - \gamma_s} \quad (3.40)$$

Again, we don't need to worry about scale factors, which cancel out when normalizing. We only ever need to evaluate each of the distributions up to a constant.

With this, the distribution starts out flattened, and is gradually made less so; by slowly increasing the exponent to 1, we lift the posterior up to its final form. The tempering coefficients being real numbers, they can be chosen continuously unlike before, and they have a direct impact on the distribution's shape.

The term *annealing* takes a similar sense as in *simulated annealing*, an approximate optimization strategy which in turn borrows its name from metallurgy. In common they have the fact that they concern some *particles*, physical or otherwise, that undergo random motion. In the beginning this randomness is more yielding, and they are allowed to move around quite freely; but as time passes, they are more and more bound to the desirable sites they stumble upon. In the case of the algorithms, the point is that exploration should be comprehensive at first, then progressively scale down.

In the other two cases, this is done by decreasing a *temperature*, the reduction in kinetic energy resulting in decreased mobility. In annealed likelihood estimation, it's almost as if we have an each time stronger potential energy (associated with the likelihood), which pulls the particles towards higher probability locations - more the more advanced the iteration is.

As for the idea of *tempering*, it is perhaps even more suitable. We can think of tempering a sauce, where we want to add in an ingredient that doesn't combine easily. In our case this ingredient is the likelihood, whereas the prior is the premade sauce. If we add in the likelihood all at once (to get the posterior), we will likely observe poor mixing. Instead, we add just a little bit of the likelihood, folding it in as we go; that is, consolidating the distribution through the iterations by refining how we discretize it. Although these distributions will at first still resemble the prior more than the posterior, each tempering step makes the two parts slightly more compatible, until we can finally mix them together in full.

In short, this new scheme ensures that the distribution evolves consistently and controllably at all points, whereas the old one was merely based on an empirical idea. That being said, it was a very convenient idea: it fixed the number of likelihood evaluations for the reweightings at  $N \cdot M$  ( $N$  is the number of data, and  $M$  that of particles), which for TLE becomes  $N \cdot M \cdot S$  for some number of iterations  $S$ . Also, the latter algorithm is clearly unsuited for online estimation.

Regardless, it is quite more robust and flexible than sequential updating; for that reason, it is a fairly popular strategy for dealing with target posteriors that are difficult to simulate, namely in high-dimensional spaces and in the presence of multimodality. In addition, a more meticulous strategy may increase the accuracy of the estimates produced by the samples - namely those regarding the sought-after marginal likelihood, or equivalently the model evidence (subsection 2.3.1).

Since it tends to be applied in more demanding contexts, the TLE approach is usually paired with more sophisticated move steps than the ones mentioned so far. In particular, LWF is unfit for multimodal distributions, one of the main motivating scenarios. Though work-arounds can be used to solve this, such as an intermediate clustering step [49], they can be computationally demanding while not addressing the underlying problem. In fact, failure in the face of multimodality is purely a symptom betraying an overall context of insufficient property preservation. This same context is likely to bias estimates, including those of the model evidence - hence, a high accuracy variety introducing mechanism is decisive for the reliability of expectations. The most widely used strategy will be presented later in subsection 3.3.3.

These alternatives do however tend to be quite more computationally demanding, which is amplified by the complexity of the distributions they target in TLE. More often than not, they require likelihood evaluations of their own, to be performed independently for each particle. Expectably, this aggravates the exigency of the protocol, further widening the gulf between the two algorithms.

All of the aforementioned requirements may be to a degree relaxed by making use of subsampling strategies, which can take advantage of the increased regularity of the sequential distributions; this will be considered in subsection 3.5. Unbiased estimators coupled with variance reduction techniques may reduce the number of necessary likelihood evaluations while largely preserving accuracy, making the method at once more efficient and more parallelizable [54].

### 3.2.4 General SMC schemes and estimation of the model evidence

Estimators for the model evidence, or the marginal probability, have been a recurring theme. The reason for this is clear. They aren't simply a normalizing constant, but rather one that assesses a model's worth - including *comparative* worth, with all that entails. This brings up an important matter: at some point along the chapter, a means of calculating it was lost.

In subsection 2.2.2, we found that the marginal,  $\mathbf{P}(\vec{D})$ , could be approximated using the sum of the unnormalized weights:

$$\sum_{i=1}^M W_i(D) \hat{=} \frac{M}{V} \mathbf{P}(\vec{D}) \quad (3.41)$$

This provides instant access to one of the most important estimates in Bayesian inference, dispensing with the extra calculations - including likelihood evaluations - that appealing to 3.1 would. Algorithm 1 is compatible with such a calculation, and so too is algorithm 2. But for algorithm 4 and the likes of it, the same no longer holds.

In the first one, the normalizing constant was directly available, because it was calculated at once. In the second, we normalized at each step, so the cumulative product of multiplying factors would do. In the third, things are no longer so clear.

Algorithm 2 already presents a minor complication: at each iteration, we must sum the likelihoods used to update the weights in equation 3.11, and not the resulting weights themselves. Each datum's probability - apart from a constant factor - is calculated at the corresponding step by removing the influence of the previous iterations. This is of course unless we can delay normalizing until the end, in which case we'd be back to algorithm 1. Normalization is only necessary when evaluating expectations.

When we move on to SIR, even this is out of reach. The problem lies with the resampling stage. By transforming weight into particle density, it disrupts the scheme. While this *disruption* is intentional, in this context it means that the weights no longer stand on their own as uniform samples. After the first resampling step, we can't *remove the influence* of the previous iteration anymore, because it is ingrained on the particle cloud. The available information is fused with the particles' attributes (locations and weights); if the original idea was to use nothing but their weights, their locations now contain irreplaceable information.

Simply put, we would want to calculate the expectations over the prior, but can only average over the current distribution  $\mathbf{P}_k(\theta)$ . And this is in the case of algorithm 4, where the sequence of target distributions is the same as in 2; for other sequences of distributions, the weights don't even begin to resemble the data's probabilities.

To see if this can be salvaged, we can start with a generic view of SMC where we consider the normalizing constant explicitly. The target is always a sequence of probability distributions; let's say, of length  $T + 1$ . They can be written as:

$$\mathbf{P}_t(\theta) = \frac{\eta_t(\theta)}{Z_t}, \quad t \in \{0, \dots, T\} \quad (3.42)$$



, having defined the normalizing constant:

$$Z_t \equiv \int \eta_t(\theta) d\theta \quad (3.43)$$

This covers all possibilities, including of course the previously examined cases. In algorithm 4 the  $\eta_t(\theta)$  were given by:

$$\eta_t(\theta) \equiv \mathbf{L}(\theta \mid \{D_1, \dots, D_t\}) \mathbf{P}(\theta) \quad (3.44)$$

While for TLE:

$$\eta_t(\theta) \equiv \mathbf{L}(\theta \mid \vec{D})^t \mathbf{P}(\theta) \quad (3.45)$$

In general, the reweightings abide by:

$$\mathbf{w}_i^{(t)} \propto \mathbf{w}_i^{(t-1)} \cdot \frac{\eta_t(\theta_i^{(t)})}{\eta_{t-1}(\theta_i^{(t)})} \quad (3.46)$$

, since the  $Z_t$  constants (this constancy being of course relative to  $\theta$ ) are rendered irrelevant by normalization. More explicitly,

$$W_i^{(t)} = \mathbf{w}_i^{(t-1)} \cdot \frac{\eta_t(\theta_i^{(t)})}{\eta_{t-1}(\theta_i^{(t)})} \quad (3.47)$$

$$\mathbf{w}_i^{(t)} = \frac{W_i^{(t)}}{\sum_{j=1}^M W_j^{(t)}} \quad (3.48)$$

With this, expectations can be computed at each step as:

$$\mathbb{E}_{\mathbf{P}_t(\theta)} [f(\theta)] = \int f(\theta) \mathbf{P}_t(\theta) d\theta \approx \sum_{i=1}^M f(\theta_i^{(t)}) \cdot \mathbf{w}_i^{(t)} \quad (3.49)$$

If a resampling stage does not take place at the end of the  $t$ -th iteration, the  $(t+1)$ th particle set is unchanged:  $\{\theta_i^{(t+1)}\}_{i=1}^M = \{\theta_i^{(t)}\}_{i=1}^M$ . If on the contrary a resampling mechanism is activated,  $\{\theta_i^{(t)}\}_{i=1}^M$  is recast to get  $\{\theta_i^{(t+1)}\}_{i=1}^M$ . This preparation aims to introduce variety for the following iteration, and preserves  $\mathbf{P}_t(\theta)$  (or a selection of relevant attributes).

Thus, taking into account the fact that resamplers typically have nice convergence properties, we may as well write expectation 3.49 as:

$$\mathbb{E}_{\mathbf{P}_t(\theta)} [f(\theta)] \approx \sum_{i=1}^M f(\theta_i^{(t+1)}) \cdot \mathbf{w}_i^{(t)} \quad (3.50)$$

With this, the main ideas for a generic SMC inference algorithm are established.

Our interest lies in the normalizing constants  $Z_t$ . As long as the final distribution is the posterior  $\mathbf{P}(\theta \mid \vec{D})$ , then  $Z_T$  is the marginal.

$$Z_T = \mathbf{P}(\vec{D}) = \int \mathbf{L}(\theta \mid \vec{D}) \mathbf{P}(\theta) d\theta \quad (3.51)$$

At the other end,  $Z_0$  is just the integral of the prior, which is 1 for any valid distribution.

$$Z_0 = \int \mathbf{P}(\theta) d\theta = 1 \quad (3.52)$$

Thus we can write:

$$\mathbf{P}(\vec{D}) = \frac{Z_T}{1} = \frac{Z_T}{Z_0} \quad (3.53)$$

Which in turn can be written as:

$$\mathbf{P}(\vec{D}) = \frac{Z_T}{Z_0} = \prod_{t=1}^T \frac{Z_t}{Z_{t-1}} \quad (3.54)$$

, since all the intermediate terms cancel out when expanding.

Now we need to get the individual terms, or some estimate thereof:

$$\left( \frac{Z_t}{Z_{t-1}} \right) \quad (3.55)$$

This is written in terms of two consecutive distributions, which is promising since we were initially displeased with their closeness. Developing the expression further, we can write the term corresponding to the  $t$ -th iteration as an average over the  $(t-1)$ th distribution, which is closer to what we wanted:

$$\begin{aligned} \left( \frac{Z_t}{Z_{t-1}} \right) &= \frac{\int \eta_t(\theta) d\theta}{Z_{t-1}} = \int \frac{\eta_t(\theta)}{Z_{t-1}} d\theta \\ &= \int \frac{\eta_t(\theta)}{\eta_{t-1}(\theta)} \cdot \frac{\eta_{t-1}(\theta)}{Z_{t-1}} d\theta = \int \frac{\eta_t(\theta)}{\eta_{t-1}(\theta)} \mathbf{P}_{t-1}(\theta) d\theta \end{aligned} \quad (3.56)$$

At this point we fall back on 3.50.

$$\left( \frac{Z_t}{Z_{t-1}} \right) = \int \frac{\eta_t(\theta)}{\eta_{t-1}(\theta)} \mathbf{P}_{t-1}(\theta) d\theta \approx \sum_{i=1}^M \frac{\eta_t(\theta_i^{(t)})}{\eta_{t-1}(\theta_i^{(t)})} \cdot w_i^{(t-1)} \quad (3.57)$$

Recognizing 3.47, this becomes:

$$\left( \frac{Z_t}{Z_{t-1}} \right) \approx \sum_{i=1}^M W_i^{(t)} \quad (3.58)$$

The finalized estimate is obtained by plugging 3.58 into 3.54.

$$\mathbf{P}(\vec{D}) \approx \prod_{t=1}^T \left( \sum_{i=1}^M W_i^{(t)} \right) \quad (3.59)$$

As it turns out, the marginal probability can still be estimated using nothing but sums of unnormalized weights - only now some extra work is required to see why and how.

Also, in this case there isn't a missing  $V/M$  factor, which is made clear by the fact that our derivation is correct. It was based on approximate expectations, which correspond to perfectly proper estimates, though 3.59 may make it seem like they are equivalent for the single step, single datum case.

Looking at the final formula only, the fundamental difference in this regard is brought about by the fact that at the first step we normalized after sampling from the prior, which exerts an influence in the product of 3.59. If we had just sampled from the (flat) prior, the starting weights would be  $1/V$  and sum to  $M/V$ . Since we normalized too, by attributing uniform weights, they are  $1/M$  and sum to 1. Of course, the correctness holds for any choice of a prior distribution, as the proof is clearly inclusive.

This restores the ability to obtain the model evidence basically *for free* along with the adoption of SMC. Because the act of determining the sums in 3.59 is intrinsic to the algorithm (e.g. the  $C^{(k)}$  in algorithm 4), one need only retrofit it with a calculation of their cumulative product.

### 3.3 Markov Chain Monte Carlo and the Metropolis-Hastings algorithm

Markov Chain Monte Carlo (MCMC) comprises a family of methods that rely on Markov chains to produce samples from a probability distribution. In its simplest form, the construction of such a chain rests only upon the ability to evaluate the probability at specific points; more advanced methods may require gradient information as well. For this section, the main references are [5, 103, 15].

This ties into what has been discussed so far. It had been noted that, if we could sample from the posterior directly, then to compute its associated expectations those samples would be all we needed from it. It would suffice to evaluate the function whose expected value is to be calculated at those samples. This can be seen as an alternative to 3.1, and a special case of 3.6 - which for the case of posterior expectations becomes:

$$\mathbb{E}_{\mathbf{P}(\theta|\vec{D})} [f(\theta)] = \int f(\theta) \mathbf{P}(\theta | \vec{D}) d\theta \approx \frac{1}{M} \sum_{i=1}^M \frac{f(\theta_i) \cdot \mathbf{P}(\theta_i | \vec{D})}{\pi(\theta_i)} \quad , \quad \theta_i \sim \pi(\cdot) \quad (3.60)$$

The choice of the posterior as importance distribution materializes our idea.

$$\mathbb{E}_{\mathbf{P}(\theta|\vec{D})} [f(\theta)] \approx \frac{1}{M} \sum_{i=1}^M f(\theta_i) \quad , \quad \theta_i \sim \mathbf{P}(\theta | \vec{D}) \quad (3.61)$$

This is a somewhat familiar idea, since strategies reminiscent of this one were used all throughout section 3.2. First, we used a similar approach to cancel out the most part of the influence each distribution had in its corresponding weight updates, effectively removing the overlap of information with their predecessors. And second, we ultimately used this very approach, it being precisely the end goal of SMC; the last iteration produces samples from the posterior. Accordingly, 3.10 is just an execution of 3.61 (for  $k = N$  and uniform weights).

There, we adopted a group of particles that we collectively propagated, allowing them to *communicate* at several points while traversing through a sequence of distributions that little by little morphed into the posterior. The difference is that now we will pick a single *particle* (which isn't usually called a particle, but

rather a state) and take it on a lone trip, with the purpose of producing samples (consecutive states) from a fixed distribution. Naturally, this distribution is the posterior.

Both methods are sequential in a sense, but the sense is not the same. In **SMC**, the fundamental sequentiality comes from targeting a sequence of progressive distributions; the samples from the posterior are produced all at once in the last step, as are those from all intermediate distributions in their respective steps. In **MCMC**, a Markov chain is sequential in the sense that it produces samples one after the other, but they all respect the same distribution. There isn't a series of distributions, only one of samples.

To sample from the posterior without slowly warming up to it as we did before, we must build a personalized Markov chain. This describes a stochastic process that travels through a sequence of states, in our case points in parameter space. The transition from one state to another is probabilistic; importantly, it depends only on the original state. In other words, the chain has no memory. Its current state fully determines the probability distribution of states it may transition to next. We write the probability of moving from a current state  $\theta^{(i)}$  to some prospective  $\theta^{(i+1)}$  as:

$$T(\theta^{(i+1)} | \theta^{(i)}) \quad (3.62)$$

The shift of the  $i$  index to a superscript illustrates the contrast between **SMC** and **MCMC**.

Thus, if the chain is at state  $\theta^{(i)}$ , the following state can be obtained by sampling:

$$\theta^{(i+1)} \sim T(\cdot | \theta^{(i)}) \quad (3.63)$$

For our purposes, we need a chain whose equilibrium distribution is the posterior distribution. For that, it must of course be invariant under the posterior. That is, if we have a group of samples from the posterior and perform a Markov transition on each (using parallel Markov chains), we should get another such group. If we instead perform multiple transitions on the same sample, the history of states will also constitute a sample from the posterior, and we can in theory obtain expected values by averaging over the registered trajectory of the Markov chain.

$$\mathbb{E}_{\mathbf{P}(\theta|\vec{D})} [f(\theta)] \approx \frac{1}{M} \sum_{i=1}^M f(\theta^{(i)}) \quad , \quad \theta^{(i)} \in \mathbf{t}_{\text{MC}} \quad (3.64)$$

A sufficient condition for having some distribution  $\mathbf{P}_X(\theta)$  as stationary distribution of the Markov chain is that of *detailed balance*, or reversibility:

$$T(\theta' | \theta) \mathbf{P}_X(\theta) = T(\theta | \theta') \mathbf{P}_X(\theta') \quad (3.65)$$

Essentially, this says that the probability of observing a transition from state  $\theta$  to state  $\theta'$  or vice-versa should be the same. In our case  $\mathbf{P}_X(\theta) = \mathbf{P}(\theta | \vec{D})$ . It's interesting to note the similarity between this equation and that which gave rise to Bayes' rule back in chapter 2.

The condition 3.65 can easily be achieved using rejection sampling. We divide our transition distribution into a proposal  $q$ , and an acceptance/rejection step  $a$ :

$$T(\theta' | \theta) = q(\theta' | \theta) \cdot a(\theta', \theta) \quad (3.66)$$

We can preserve the freedom of choosing the proposal distribution by tailoring the acceptance probability to it, effectively performing a *correction*.

$$a(\theta', \theta) = \frac{q(\theta | \theta')\mathbf{P}_X(\theta')}{q(\theta' | \theta)\mathbf{P}_X(\theta)} \quad (3.67)$$

With this, it is easy to see that equation 3.65 holds.

$$T(\theta' | \theta)\mathbf{P}_X(\theta) = q(\theta' | \theta)a(\theta', \theta)\mathbf{P}_X(\theta) \quad (3.68)$$

$$= q(\theta' | \theta) \frac{q(\theta | \theta')\mathbf{P}_X(\theta')}{q(\theta' | \theta)\mathbf{P}_X(\theta)} \mathbf{P}_X(\theta) = T(\theta | \theta')\mathbf{P}_X(\theta') \quad (3.69)$$

A caveat is that  $a$ , being a probability, mustn't be higher than one, and the ratio in 3.67 guarantees no such thing. We can patch this by having an exception for the case where the ratio exceeds unity, bringing it down to 1.

$$a(\theta', \theta) = \min \left( 1, \frac{q(\theta | \theta')\mathbf{P}_X(\theta')}{q(\theta' | \theta)\mathbf{P}_X(\theta)} \right) \quad (3.70)$$

This solution is called the **Metropolis-Hastings (MH)** algorithm [84, 56]; there are other methods still within **MCMC**.

An interesting thing to note is that the theoretical framework is unaffected by constant factors, which cancel out in the ratio as well as in all the equations. In particular, the probability distribution doesn't need to be normalized. This is a crucial point, because when we let go of the multi-particle arrangement of **SMC** we lost our means of calculating the normalizing constant for the posterior. This way, we can still formulate the transitions by evaluating the likelihood and the prior only.

Similarly, we lost some ease in assessing convergence and robustness when changing methods. In **SMC**, the congruity of the particle cloud conveniently briefed us on where it stood, whereas an individual chain offers no such measure. This too can be fixed, though less directly, by starting multiple chains at different locations and running them in parallel. The ensemble can then used to extract statistics<sup>4</sup>. In particular, their proximity is used to evaluate consistency: large discrepancies should raise suspicion. Such a scheme resembles **SMC**, but it is completely parallelizable because the particles don't interact .

Having laid out these ideas, we are ready to construct a theoretically sound Markov chain. The procedure is described in algorithm 5.

As long as the stationary distribution is the correct one, the Markov chain will converge as desired<sup>5</sup>. As such, the starting point  $\theta^{(0)}$  can be chosen at random, though it will have an effect on convergence. Likewise, we are free to choose a proposal distribution  $q(\theta' | \theta)$ , but it strongly affects the performance of the algorithm.

The chain *will* eventually reach its steady state, or be close enough - at which point we get exact samples from the posterior, or accurate enough ones. Nonetheless, what this *eventually* offers is scant

<sup>4</sup>Similarly, if the intel provided by the particle cloud in **SMC** is insufficient, the particles too can be split into groups that evolve independently of each other, and their outcomes compared [24].

<sup>5</sup>This is an oversimplified idea; some transitions may cause the chain to become faulty due to complex geometrical attributes of the target distribution [15]. This may frustrate ergodicity, a characteristic which underpins **MCMC** - thus invalidating estimates.

**Inputs:** starting state  $\theta^{(0)}$ .

**Computes:** sequence of states  $\mathbf{t}_{\text{MC}} = \{\theta^{(i)}\}_{i=1}^M$  to be used in 3.64.

**Assumes:** Ability to evaluate the prior, the likelihoods and a predefined proposal density  $q(\theta' | \theta)$ , and additionally to sample from the latter and from a binomial distribution.

**Assures:** The samples are distributed according to the posterior distribution as  $M \rightarrow \infty$ .

1. For  $i \in \{1, \dots, M\}$ :

a) Sample a proposal

$$\theta^* \sim q(\cdot | \theta^{(i-1)}) \quad (3.71)$$

b) Compute the acceptance probability

$$a(\theta^*, \theta^{(i-1)}) = \min \left( 1, \frac{q(\theta^{(i-1)} | \theta^*) \mathbf{L}(\theta^* | \vec{D}) \mathbf{P}(\theta^*)}{q(\theta^* | \theta^{(i-1)}) \mathbf{L}(\theta^{(i-1)} | \vec{D}) \mathbf{P}(\theta^{(i-1)})} \right) \quad (3.72)$$

c) With probability  $a(\theta^*, \theta^{(i-1)})$ , set

$$\theta^{(i)} = \theta^* \quad (3.73)$$

Otherwise, keep  $\theta^{(i)} = \theta^{(i-1)}$ .

Algorithm 5: Metropolis-Hastings algorithm for posterior sampling in Bayesian inference.

assurance. If care is not taken, runtime until convergence may be too long for its fruits to ever be reaped. Even if this is not true, in practice two provisions may be considered for accuracy's sake:

- A number of initial samples is to be discarded, because the chain won't have reached equilibrium yet. This is called a **burn-in** period.
- We may want to keep only every  $L$ th sample, discarding the ones in between. This is called **thinning**. The point is that consecutive samples will be correlated due to the localized proposals; this correlation can be made negligible at or after lag  $L$  (for large enough  $L$ ). Correlated samples may bias calculations. This is particularly relevant if a *compact* set of samples is desired, due to e.g. memory or post-processing constraints [72].

Both of these points are greatly impacted by the proposal distribution. The time required to reach the stationary regime (or get close enough), called the mixing time, greatly depends on it, as it determines how fast the chain explores the space. For the same reason, it influences the serial correlation between samples. If it causes minuscule perturbations, the chain will mix slowly; however, countering this isn't as simple as opting for bold proposals.

Ideally, the transitions *should* produce distant moves; but since the ratio  $\mathbf{P}(\theta^*)/\mathbf{P}(\theta^{(i-1)})$  is a factor in the acceptance probability, unchecked distant proposals may suffer from high rejection rates, which is self-defeating since the chain will stall. This is of course not applicable if the proposals are well founded and can systematically find regions of high probability, but that how to generate such proposals is an open problem.

By now, it is clear that the single most important implementation detail is the proposal distribution, which takes a similar role to our resampler from before. It can affect convergence speed, computational efficiency, and even correctness. The following two sections we will go over two well-known approaches: random walk Metropolis, and Hamiltonian Monte Carlo. Both are mostly based on [15], one of the main references listed for this section.

### 3.3.1 Random walk Metropolis

A common simplification of the Metropolis-Hastings algorithm is called simply the *Metropolis* algorithm. It describes the original method, which Wilfred Hastings later generalized it to the general case. The only difference is that the proposal distribution is chosen to be symmetric, i.e.

$$q(\theta' | \theta) = q(\theta | \theta') \quad (3.74)$$

With this, the acceptance ratio from 3.70 becomes:

$$a(\theta', \theta) = \min \left( 1, \frac{\mathbf{P}_X(\theta')}{\mathbf{P}_X(\theta)} \right) \quad (3.75)$$

This has a nice interpretation: each point's chances of being picked grow in proportion with the probability the target distribution assigns to it. Naturally, if we want to sample from a distribution, we mean to favor precisely the points that it favors, and that too with the insistence that it favors them.

In its earliest version, the algorithm utilized gaussian proposals centered at the starting state, which are obviously symmetric.

$$\theta' \sim \mathcal{N}(\theta, \sigma) \quad (3.76)$$

$$q(\theta' | \theta) = \mathcal{N}(\theta' | \theta, \sigma) \quad (3.77)$$

The standard deviation is often tuned to maintain the acceptance ratio around a pre-defined value, though more sophisticated methods can be employed [43].

Along with 3.75, 3.76 defines a protocol which we will call *random walk Metropolis (RWM)*. It is based on stochastic exploration corrected for probability density. Thinking of the behaviour of such a strategy in response to an increase in dimensions, we can see how it balances the tension between probability density and parameter space volume (mentioned in section 3.1). On the one hand, the Gaussian generalizes to a multivariate normal distribution, which relays the behaviour of high dimensional spaces through the way the tails take up increasing volume. On the other, the correction of 3.75 always privileges higher probability regions. Together, they will negotiate the intended compromise, focusing on the region that produces significant contributions to the integral.

On the downside, equation 3.76 describes a random walk: the direction and speed of exploration are not customized to the target function. Due to this fact, the chain is likely to move in small steps, and further to often retrace them since there is no element of continuity between transitions (which there could still be in a Markovian process: it can be achieved via the target function's geometry rather than past states).

This tends to worsen with increased dimensions. As the region that yields consequential contributions to the expectations becomes more and more concentrated and the volume outside of it grows, uninformed proposals become each time more inadequate. Either they are bold and likely to fall outside the narrow non-negligible probability region (wherefore they are rejected), or they are conservative and slow down the chain.

For simple problems, **RWM** may still perform well within a reasonable span of time. Unfortunately, its issues can quickly get out of hand, bringing unacceptable inefficiency. In that case, the differential structure of the posterior distribution may aid the production of more educated guesses. This is a commonly exploited resource in more advanced **MCMC** methods, namely Hamiltonian Monte Carlo<sup>6</sup>.

### 3.3.2 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (**HMC**), originally called Hybrid Monte Carlo (also **HMC**) [29], comes to solve the problems that a naive random walk brings, and reinstates the efficacy of **MCMC** for high-dimensional spaces. This is approached at length in references [15], which contains an intuitive overview on **MCMC** methods with a focus on Hamiltonian Monte Carlo, and [88], which is a more detailed review of **HMC**. Finally, the reference manual and source code of the statistical inference platform *Stan* [42, 105] presents most if not all of the state-of-the-art **HMC** techniques, along with select bibliography in the first case.

We have suggested that gradient information may come in handy, but its application is not so direct: we don't simply want to steer towards high probability regions, and so it necessitates some correction. This correction should be based on differential geometry, but a physical analogy provides a more accessible explanation.

It is inspired in the Boltzmann distribution, which assigns occupation probabilities to states based on their energy. For some state  $i$  with energy  $E_i$  and at temperature  $T$ , the occupation  $p_i$  is:

$$p_i \propto e^{\frac{-E_i}{kT}} \quad (3.78)$$

,  $k$  being the Boltzmann constant.

We regard our possible  $\theta$ s as states, and our goal is that the chain finds them with probability  $\mathbf{P}_X(\theta)$ . Thus the energy our *physical system* has in association with the target distribution is given by:

$$\mathbf{P}_X(\theta) \propto e^{\frac{-U(\theta)}{kT}} \quad (3.79)$$

We take this as a potential energy, regarding  $\theta$  as a position - hence the use of the letter  $U$ . We then choose:

$$U(\theta) = -\ln(\mathbf{P}_X(\theta)) = -\ell(\theta) \quad (3.80)$$

This energy is the negative of the log-likelihood  $\ell(\theta)$ , a commonly used construction by virtue of its computational convenience - it turns products of small (positive) numbers into sums of large (negative)

<sup>6</sup>Another common approach is Langevin Monte Carlo, also called the Metropolis-adjusted Langevin algorithm, which is a particular case of Hamiltonian Monte Carlo [8].



numbers. The energy offset is irrelevant, as is a multiplying constant in the probability; the implementation is insensitive to them, as it should. This choice of reference produces an always negative energy,  $U(\theta) \in ] - \infty, 0]$  for  $\mathbf{P}_X(\theta) \in [0, 1]$ .

Additionally, we introduce an auxiliary *momentum* variable, or vector - it should have as many dimensions as there are parameters. It supplies a *kinetic energy*, which will help us explore the parameter space in a more orderly fashion. The effect is that now we work on a phase space with double the dimensions as compared to the parameter space.

$$\theta \rightarrow (\theta, p) \quad (3.81)$$

We also introduce a distribution over the momentum, which may be conditioned on the position:  $\mathbf{P}(p \mid \theta)$ . The joint position/momentum distribution must admit the original distribution as marginal distribution. That way, we can transform samples from this augmented distribution into samples from the target distribution simply by discarding the momentum component.

$$\mathbf{P}(\theta, p) = \mathbf{P}(p \mid \theta)\mathbf{P}_X(\theta) \quad (3.82)$$

Reformulating the probability of occupation, we can define a (classical) Hamiltonian that gives the total energy of the system.

$$\mathbf{P}(\theta, p) = e^{-H(\theta, p)} \quad (3.83)$$

$$H(\theta, p) = U(\theta) + K(p, q) = -\ln(\mathbf{P}_X(\theta)) - \ln(\mathbf{P}(p \mid \theta)) \quad (3.84)$$

While the position distribution is given by the problem, the momentum distribution can be picked at will for the implementation. If we forgo position dependence and borrow the kinetic energy from classical mechanics, we get<sup>7</sup>:

$$K = \frac{p^2}{2\mathcal{M}} \rightarrow \mathbf{P}(p \mid \theta) = e^{-\frac{p^2}{2\mathcal{M}}} \quad (3.85)$$

With a corrective offset in the kinetic energy<sup>8</sup>, the momentum distribution associated with 3.85 becomes an easy-to-sample-from Gaussian, with mean 0 and (co)variance given by the mass. This is called an Euclidean-Gaussian kinetic energy [15]. The mass must be tuned according to the problem, but is often chosen to be diagonal (commonly a multiple of the identity matrix). This will be discussed later in this subsection. For now, we have:

$$p \sim \mathcal{N}(\cdot \mid 0, M) \quad (3.86)$$

It is important that we can sample from the momentum distribution, because we don't want it to be a hindrance to the Markov chain's journey. Here it is trivial to sample a momentum vector from the correct distribution before even beginning to construct the chain. This way, we can start the chain with a perfect

<sup>7</sup>The one-dimensional case is considered for simplicity, but the generalization is straightforward:  $K = p^T M^{-1} p / 2$ . The mass  $M$  must be a symmetric positive-definite matrix. Other elements can be similarly generalized.

<sup>8</sup>This offset is the logarithm of the Gaussian's normalizing constant, which depends on the mass, but it doesn't affect the dynamics as long as  $\mathcal{M}$  is independent of  $\theta$ .

sample on the momentum side, and no burn-in period is required as far as that side is concerned. This is in contrast with the position distribution, where we can't sample a starting  $\theta$  directly from the target distribution; the inability to do so is the very reason why we use **MCMC**. In that case, the starting state generally falls far from the mode(s), and it is the chain's task to conduct it to higher probability regions.

We then start the process at some starting position  $\theta(t = 0)$ , which is chosen as for any Markov chain, and with some momentum  $p(t = 0)$ , which is sampled directly from the selected momentum distribution. The objective is to advance the chain while preserving the augmented distribution; or equivalently, while preserving the energy given by the Hamiltonian of equation 3.84. We can do so by having the *state*, regarded as a mechanical system, undergo Hamiltonian dynamics.

This completes the necessary structure for the construction of a **HMC** proposal, which comes to replace the Gaussian sampling of random walk Metropolis. It makes use of the improvised momentum to slide along the potential determined by equation 3.80, as the augmented system explores the extended space (a *mechanical phase space*) that is built from the original one (the parameter space). This is depicted in figure 25.




Figure 25: Illustration of the process behind a HMC proposal (for a gaussian distribution).

Simulating the evolution prompted by a Hamiltonian is conceptually simple: it is described by Hamilton's equations, and we are in possession of all of their elements. Assuming the same position-independent kinetic energy as before, they take the form:

$$\begin{cases} \frac{d\theta}{dt} = \frac{\partial H}{\partial p} \\ \frac{dp}{dt} = -\frac{\partial H}{\partial \theta} \end{cases} \leftrightarrow \begin{cases} \frac{d\theta}{dt} = \frac{\partial K}{\partial p} \\ \frac{dp}{dt} = -\frac{\partial U}{\partial \theta} \end{cases} \leftrightarrow \begin{cases} \frac{d\theta}{dt} = \frac{p}{\mathcal{M}} \\ \frac{dp}{dt} = -\frac{\partial U}{\partial \theta} \end{cases} \quad (3.87)$$

To evolve the Markov chain, we must simulate the dynamics this system induces for some evolution time  $\Delta t$ .

$$\left(\theta(0), p(0)\right) \rightarrow \left(\theta(\Delta t), p(\Delta t)\right) \quad (3.88)$$

For simplicity we write as before

$$(\theta, p) \rightarrow (\theta', p') \quad (3.89)$$

If we can do execute this perfectly, the Hamiltonian will be preserved, as will the position and momentum distributions, producing perfect samples. But as per usual, we won't, and must resort to approximate numerical methods. In that case we can regard the last state of the trajectory as an elaborate proposal  $\theta' \sim q(\cdot | \theta)$ , and perform a Metropolis-Hastings step to account for discretization errors. The acceptance rate is as before (defining  $v \equiv (\theta, p)$ ):

$$a(v', v) = \min\left(1, \frac{q(v | v')\mathbf{P}(v')}{q(v' | v)\mathbf{P}(v)}\right) \quad (3.90)$$

This is a problem. Unless there is no movement, the probability of recovering the initial state of the trajectory  $(\theta, p)$  when starting at its final state  $(\theta', p')$  is null, just like a sliding object will never revert its motion to go back uphill. Hamiltonian dynamics are deterministic, fixing the ratio of proposal probabilities at 0/1 - which sets the acceptance probability to zero. For any  $(\theta_0, p_0)$ , there is only one possible final state  $(\theta_F, p_F)$ , so any other point has null probability of being proposed.

Fortunately, the correction is straightforward: it suffices to invert the momentum before making the proposal ( $p' = -p(\Delta t)$ ). Should the trajectory have started at  $(\theta', -p)$ , it would have been the exact inverse of the one starting at  $(\theta', -p)$ . This renders the proposal symmetric (1/1), in which case we're back to the Metropolis case:

$$a(\theta', \theta) = \min\left(1, \frac{\mathbf{P}(v')}{\mathbf{P}(v)}\right) = \min\left(1, \frac{e^{-H(v')}}{e^{-H(v)}}\right) = \min\left(1, e^{H(v)-H(v')}\right) \quad (3.91)$$

This acceptance probability resembles the Boltzmann factor, depending on the difference between energies. The more the new state lowers the energy as compared to the old one, the more likely it is to be accepted. Also, we can see that the sign of the momentum doesn't matter as long as the kinetic energy's dependence on it is quadratic, so although the sign flip is required for the sake of theoretical soundness no practical adjustments are necessary other than using 3.91 and not 3.90 to compute the acceptance probability (it being that the latter would mean setting it to 0).

With this, we get exact samples from the augmented density, whose projection into the original space obeys our target distribution. What is more, if the energy is approximately preserved, which it should be for reasonable integration methods, we will have that  $a(\theta', \theta) \approx 1$ : an acceptance rate nearing 100%. This is how the methodology affords efficient exploration of the space, producing the ideal combination of distant proposals and high acceptance probabilities - of course, at the cost of the transition kernel's higher complexity.

Here the *reasonability* of a given integrator may take an unusual form as compared to other applications. In particular, it is important that the integrator be a symplectic map, or equivalently a canonical transformation. This means that it should preserve the form of Hamilton's equations, though not necessarily the energy itself. This means that even if it doesn't exactly obey our ideal Hamiltonian, it acts under the action of a *perturbed* one that approximates it. This prevents the trajectory from drifting - it instead oscillates around the correct one, and its fluctuations average out unbiasedly instead of building up for long integration times.

An example of a symplectic integrator is the elementary leapfrog integrator, which interleaves momentum and position updates to advance the system while evaluating the gradients at midpoints. The algorithm is presented in 6. Note the time reversibility: if we start at the final point and integrate backwards in time, we will retrace our steps all the way back to the initial point.

**Inputs:** initial position  $\theta(0)$  and momentum  $p(0)$ , integration stepsize  $\epsilon$ , mass  $\mathcal{M}$ , integration time  $\Delta t$ .

**Computes:** sequence of position  $\theta(t)$  and momentum  $p(t)$  states for  $0 < t \lesssim \Delta t$ , where the two instances of  $t$  are matched for the last state.

**Assumes:** Ability to evaluate the prior and the likelihoods, and additionally their gradients.

**Assures:** Preserves the form of Hamilton's equations, though not necessarily the Hamiltonian itself.

1. Define

$$U(\theta) = \mathbf{L}(\theta \mid \vec{D})\mathbf{P}(\theta) \quad (3.92)$$

$$\theta_x, p_x = \theta(x \cdot \epsilon), p(x \cdot \epsilon) \quad (3.93)$$

2. Advance the momentum by a half timestep.

$$p_{1/2} = p_0 - \frac{\epsilon}{2} \cdot \nabla U(\theta_0) \quad (3.94)$$

3. For  $k \in \{1, \dots, L\}$ , with  $L \equiv \text{round}(\Delta t/\epsilon)$ :

- a) Advance the position by a timestep.

$$\theta_k = \theta_{k-1} + \epsilon \cdot \frac{p_{k-1/2}}{\mathcal{M}} \quad (3.95)$$

- b) If  $k \neq L$ , advance the momentum by a timestep.

$$p_{k+1/2} = p_{k-1/2} - \epsilon \cdot \nabla U(\theta_k) \quad (3.96)$$

4. Advance the momentum by a half timestep.

$$p_L = p_{L-1/2} - \frac{\epsilon}{2} \cdot \nabla U(\theta_L) \quad (3.97)$$

Algorithm 6: Leapfrog integration for solving Hamilton's equations in the context of Bayesian posterior sampling.

The leapfrog algorithm can be generalized to higher order, though the increase in accuracy may not

justify the cost. Even though this method is very simple, it beats others whose smaller but coherent errors may prove disastrous for the Hamiltonian trajectory.

Another benefit is that even though they are not immune to problems that frustrate exploration (and thus accuracy), these problems are easy to recognize. This is because the same geometrical features - namely high curvature - that bias estimators will affect numerical stability, causing the trajectory to diverge. Such problems can be easily identified so compensatory strategies can be employed. Overall, their simplicity and theoretical grounds facilitate the task of assessing their effectiveness.

The integrator's performance depends on some hyperparameters, namely the stepsize  $\epsilon$  and the path length  $L$ . Together, they determine the integration time  $\epsilon \cdot L$ . For a fixed integration time, the stepsize determines how many iterations  $L$  must be performed, which defines the cost of the algorithm; in particular, the number of evaluations of the likelihood and its gradient scale inversely with  $\epsilon$ . Because the accuracy and stability of the integration depends strongly on this same  $\epsilon$ , which determines the roughness of time discretization, a good balance should be sought when choosing it.

Having discussed the integrator, we can finalize the main details of [HMC](#). The procedure is laid out in [algorithm 7](#). One may note that an important detail is missing in its construction: it doesn't explicitly cater to constraints (parameter space boundaries). They're left to be controlled by the fact that their prior probability will be zero, which will cause a minimum overflow. A way to treat restraints is to consider the effect of attributing an infinite potential energy to out-of-bounds regions, and handle it within the [integration 6](#). This leads to the position ( $\theta$ ) advancement step being followed by a *rebounding* motion (the position is mirrored relative to the boundary and the momentum is negated) if any positional boundary has been crossed [[88](#), [24](#)].

Note that all that is used from the generated phase space trajectory of [6](#) is the last state, which is proposed as a state for the actual trajectory of the Markov chain. There exist some alternative approaches that don't treat intermediary states simply as an auxiliary construct but rather as proposals by themselves. More advanced strategies may drop the Metropolis correction entirely and instate some other scheme that safeguards correctness. A short overview is left for [subsection 3.4.2](#). In order to get valid proposals for such a scheme, the leapfrog [algorithm 6](#) must be adapted to produce valid points, by synchronizing the position and momentum coordinates. This is more commonly known as the velocity Verlet method. It could be done by evaluating the momentum at integer steps as well, e.g. moving the edging half-steps inside the loop and doing 3 updates per iteration instead of 2. Alternatively, one could drop midpoint evaluations, or at least change the role of the momentum and position since the latter has more expensive updates (log-gradient evaluations).

As compared to random walk Metropolis, well-tuned [HMC](#) only presents random walk behaviour between energy levels. This is induced by the random choice of a momentum variable that lifts the position coordinates up to a double-dimensional (as compared to the original parameter space) phase space, joining the position coordinate to establish the energy<sup>9</sup>. Other than that commencing step, the exploration of

<sup>9</sup>Even this random walk aspect can be lessened by only partially refreshing the momentum between transitions [[88](#)].

**Inputs:** starting state  $\theta^{(0)}$ , mass  $\mathcal{M}$ , integration time  $\Delta t$ .

**Computes:** sequence of states  $\mathbf{t}_{\text{MC}} = \{\theta^{(i)}\}_{i=1}^M$  to be used in 3.64.

**Assumes:** Ability to evaluate the prior and the likelihoods, to sample from gaussian and binomial distributions, and to integrate Hamilton's equations.

**Assures:** The samples are distributed according to the posterior distribution as  $M \rightarrow \infty$ .

1. Define

$$H(\theta, p) = -\ln \left( \mathbf{L}(\theta \mid \vec{D}) \mathbf{P}(\theta) \right) - \frac{p^2}{2\mathcal{M}} \quad (3.98)$$

$$v^{(i)} \equiv (\theta^{(i)}, p^{(i)}) \quad (3.99)$$

2. For  $i \in \{1, \dots, M\}$ :

a) Sample the initial momentum to define the initial state.

$$p^{(i-1)} \sim \mathcal{N}(\cdot \mid \mu = 0, \sigma^2 = M) \quad (3.100)$$

b) Simulate Hamilton's equations for  $H(\theta^{(i-1)}, p^{(i-1)})$  for a time  $\Delta t$  and define the proposal according to the final state.

$$v^* \equiv (\theta^*, p^*) \equiv \left( \theta(\Delta t), -p(\Delta t) \right) \quad (3.101)$$

c) Compute the acceptance probability.

$$a(v^*, v) = \min \left( 1, \exp \left( H(v^{(i-1)}) - H(v^*) \right) \right) \quad (3.102)$$

d) With probability  $a(v^*, v^{(i-1)})$ , set

$$\theta^{(i)} = \theta^* \quad (3.103)$$

Otherwise, keep  $\theta^{(i)} = \theta^{(i-1)}$ . The momentum can be discarded.

Algorithm 7: Hamiltonian Monte Carlo algorithm for posterior sampling in Bayesian inference.

the energy level set is deterministic - fixed by the integrator - and informed by the geometry of the target distribution, which brings improved efficacy.

The energy transition distribution is induced by the sampling of the momentum, and should be adapted to the marginal energy distribution of the target distribution. This is similar to how the **RWM** proposal should be tuned to the distribution itself. On the other hand, the efficacy of the exploration for a fixed energy depends on the integration method and time. As for the integration time, there is a trade-off between completeness of exploration and cost, so the goal would be to keep it *just long enough* to characterize the energy level set before the returns diminish. On the opposite extreme, its being too small may not bring gains as compared to exploration via a random walk. This will be further discussed in subsection 3.4.2.

As for the mass, its inverse should resemble the covariance of the target distribution. Intuitively, this makes sense: the sharper the distribution, the *heavier* the system should be, so that the moves are

less adventurous. In the general case, the mass is a matrix, and it rotates and rescales the augmented parameter space, reshaping the structures that must be traversed. Naturally, this means that it may help smooth the way for the integrator.

For a Gaussian distribution, optimal results can be achieved with a Gaussian kinetic energy, but in the general case that is no longer true. As such, a position dependence may be worked in to better treat local features of the distribution's landscape [46].

All of these considerations were based on [15], which presents an accessible but thorough introduction to HMC and the main aspects concerning its implementation. It additionally contains illustrations that aid the visualization of many of the points approached here (and more).

Despite its ability to propose distant moves with high acceptance probabilities, HMC is not without its shortcomings. Namely, its performance is strongly affected by the tuning of the parameters  $\epsilon$ ,  $L$  and  $M$  (and possibly others in more advanced implementations), which is a complex task. Furthermore, it is far harder to implement than RWM, and rather computationally demanding - though this is in terms of cost per iteration, which per se doesn't matter much. Also, it is not well suited to multi-modal distributions, and can get trapped in isolated energy minima, especially if the modes are far apart and/or differ in height (since even if the chain gets there the acceptance probability will be low due to the difference in energy). However, not only is this weakness common to most other approaches to MCMC, it also can be mitigated via improvements to the basic HMC algorithm [11, 88, 69, 53, 13] or by embedding the chain in an interacting ensemble (as in SMC).

The circumstances under which the penny plain version of HMC shines the most are high-dimensionality, hard-to-resolve (elongated and/or sharp) shapes, and near-normality - all of which greatly benefit from the fictitious momentum variable. Due to this enhancement, HMC chains are capable of accomplishing geometric ergodicity for a wider range of functions than RWM or other simplistic approaches, in which case the expectation estimates abide by a CLT instead of achieving correctness only in the asymptotic limit (as for a usual Markov chain) [13]. Even if this does not hold, progress tends to occur on a much faster time scale for a geometric method as compared to random walk based exploration. Further, even in a worst case scenario HMC offers an interesting perk: powerful diagnostic tools. This is owing to the fact that when it does fail, it tends to fail quite spectacularly, facilitating fault detection and correction [13], as mentioned regarding the construction of the numerical trajectory.

These advantages of HMC will be demonstrated in chapter 4 by means of an illustrative set of examples.

### 3.3.3 Markov kernels within SMC

Despite the alluring simplicity of their underlying principle, MCMC methods have many drawbacks of their own. Coming near their promised correctness presents many hurdles, since theoretical ergodicity often doesn't transfer well to practice. Not only is the stationary regime hard to reach, but also hard to recognize.

For these reasons and others, they do worse in many aspects as compared to SMC. Namely, they:

- Are not suitable for online estimation;

- Tend to require more likelihood evaluations (especially when considering **SMC-SIR** with the prior as importance function and likelihood free resamplers like Liu West's);
- Don't deal well with complex features of the likelihood landscape, such as multimodality - especially when the modes are sharp and distant from one another (Markov chains are likely to get stuck in the one closest to the starting point for the duration of any realistic runtime);
- Do not offer as effortless a way of computing of the marginal likelihood/model evidence;
- Can be hard to tune.

As such, they can often lag behind as a competitor to **SMC**. However, they don't have to be a competitor - they can be an ally. Their invariance properties make them perfect to refresh the particle cloud while exactly preserving the distribution it represents. For that reason, **MCMC** kernels are widely used within **SMC** for particle propagation, [87, 102, 54, 97]. After performing the reweighting and multinomial sampling steps as before, the uniformly weighted particles are moved in parallel using one or more **MCMC** transitions. A **RWM** kernel is perhaps the most popular approach, but **HMC** may significantly enhance performance. When using **HMC** moves, **SMC** is often termed sequential Hamiltonian Monte Carlo (**sequential Hamiltonian Monte Carlo (SHMC)**).

As compared to **LWF**, this scheme preserves the entire distribution and not just its first two moments. This automatically solves the issues in capturing multimodality, by getting directly to the root of the problem.

The aid is reciprocal, and the **SMC** framework also complements the Markov transitions well in several ways. When they're used as variety-introducing mutations, a burn-in period is not required, because the particle cloud is already a Monte Carlo approximation of the desired distribution before they are performed. If it does not approximate it exactly, these moves should aid convergence, since the Markov chain's stationary distribution is precisely the target.

Discarding lag samples is also unnecessary, given that the distribution is characterized by weighted particle density and not successive Markov chain states. However, correlation between samples is still relevant, despite not being strictly necessary for formal correctness in the asymptotic limit (infinite number of particles). It dictates how effectively variety is introduced in the particle positions, which strongly impacts how well covered the parameter space is. Correlation between samples can be used for deciding by how many Markov moves each particle should be displaced, in order to assure an effective exploration [54].

The two simplifications mentioned above are direct benefits of shifting responsibility from sequential states to parallel particles, but there are others. For instance, the particle cloud structure can provide an insight into the tuning of the **MCMC** kernels [18], whose efficacy strongly depends on their parameters. This is particularly critical when using Markov transitions which are very sensitive to the parameters and where choosing them is a challenging problem, as with **HMC**. For instance, the mass matrix should resemble the covariance of the target distribution, which can be approximated by that of the particle cloud.

Finally, the fact that **SMC** is more supple with respect to the targeted densities brings added flexibility, relaxing the **MCMC** requirement that the target density be a constant function of the parameters.

Thus constructed **SMC** samplers enjoy, unlike before, absolute formal correction (asymptotically), in



the sense that the resampling step doesn't introduce error. They can also be interpreted as *interacting* Markov chains. This is perhaps the most common implementation of particle filters of this sort, and it is popular across various fields of knowledge. For that reason, it is known by many different names, both more specific and more general: interacting Metropolis algorithms, Feynman-Kac particle models (mostly in measure-theoretic literature), etc. [86, 85].

It can also be seen as a genetic algorithm which alternates between mutation and selection phases. For the duration of the *mutation* periods, the particles are allowed to evolve independently through Markov transitions. The *selection* stage is performed based on a probabilistic survival of the fittest criterion: the most apt particles *reproduce*, and the least apt ones *die*.

All the advantages of exact *mutations* come at the cost of extra likelihood evaluations. Given that a finite number of particles is unable to characterize in full a continuous probability density, it was expected that to thoroughly preserve it consulting the particle cloud would not be enough. The extra required resources are significant, making the propagation step the most costly part of the algorithm. For the scheme of algorithm 4, which is particularly frugal with respect to likelihood evaluations, moving a particle at some step  $k$  will demand at least twice as many resources as have been used for all the reweightings of the particle up to that point ( $2k$ ). And this is for *RWM* mutations - *HMC* requires on top of that  $L \cdot k$  gradient evaluations,  $L$  often being of the order of the hundreds or more and the gradients being rather expensive for large datasets and/or complex functions.


The use of Markov kernels also makes the algorithm more memory-intensive, because all the data must be preserved throughout the runtime (to be used at least in the move steps). While in the case of tempered likelihood estimation this was already the case, this still brings some added difficulties, namely if meaning to exploit parallelism [54]. While parallel Markov chains are in theory *embarrassingly parallel*, large shared datasets make them less so due to the overwhelming memory requirements.

All of this may be dispensable if the mean and variance offer a sufficient description of the distribution, but not for more complex targets or if the accuracy of expectations is of concern. In particular, the model evidence estimator is sensitive to the propagation strategy. Further, not only do Markov steps in themselves improve the accuracy of this estimator, they can also be structured in a way to enable recycling the proposals for the construction of a different and more robust one. This strategy is proposed in [102].

Either way, the *SMC* samplers described here have been employed to deal with complex statistical functions, including tempered inference with *HMC* move kernels [24, 54] (the most demanding combination of all alternatives reviewed here). Cost doesn't matter *per se*, and for many applications the value of these costly but dependable strategies is completely unmatched by simpler but unavailing ones.

### 3.4 Implementations of Hamiltonian Monte Carlo

The previous section contained superficial considerations on the tuning of *HMC*'s parameters (namely the mass matrix and integration step size and path length), but mostly left out structural modifications to the



(a) Binomial likelihood model (the example explored throughout sections 2.1 and 2.2). 100 trials were contemplated in the data vector.

(b) Squared-cosine likelihood model (the example from subsection 2.3.3). 50 measurements were contemplated in the data vector, using evenly spaced times on  $[0, 10[$ . The thin dotted vertical lines (in blue) mark the log-likelihood asymptotes, where the likelihood is null.

Figure 26: Posterior probability (crimson) and log-probability (blue) for two different likelihood models and flat priors. The dashed vertical lines mark the real parameter used to generate the data  $D$ .

algorithm or its target, instead focusing on the *standard* version of **HMC**. This section will present a short survey of a few such modifications, namely alternative takes on generating a trajectory, selecting a sample, and determining the integration path length. Before that, a short aside on features **HMC** struggles with, which can themselves be resolved by more sophisticated implementations.

### 3.4.1 Conditions for applicability

The most determining factor for the efficiency of **HMC** is whether it is well-suited for the target function. A few cases where it is not have been briefly mentioned, such as multi-modality. But one needn't go so far to watch **HMC** fail.

Let us take for instance the example of section 2.3.3. It is a rather simple example, where we can expect the posterior to be approximately normal after just a few data. The distribution was uni-modal, and univariate; **HMC** is known to excel in much more complex scenarios. Yet, it doesn't in this one, and upon implementing it we would find it to be no more effective than **RWM**.

The reason why can be seen in figure 26, which shows the likelihood and log-likelihood for the *coin* model (constant binomial function) and for the precession model (time-varying sinusoidal likelihood).

In the first case (subfigure 26a), the log-likelihood is a smooth curve that **HMC** can easily sail through - as much as it may not be the most resource-effective method for such a simple target. But in the second case (subfigure 26b), the log-likelihood looks grim for **HMC**.

The problem is that the likelihood function (given in 2.84) has roots. This should if anything be a *nice* attribute: it may allow for ruling out some hypotheses based on a single datum. However, the resulting discontinuous log-likelihood is a problem for HMC<sup>10</sup>.

Our HMC algorithm assumed the *potential energy* to be continuously differentiable within some boundaries, which means that no points should have null likelihoods. That they do creates potential barriers that are overridden only if the discretized dynamics happens to leap over them by chance. As a consequence, the periodic roots of the precession example then create a minefield of potential wells that trap the particle in.

If these discontinuities affected the likelihood landscape at single points, we would with unbounded probability never encounter them; unfortunately, their neighbourhoods have a disruptive effect as well due to how the energy tends to infinity. This is aggravated by limited numerical precision; after a point the trajectory will become unstable, as the energy and its gradient overflow. When they do, the particle won't even be able to slide back into representable probability territory.

Even though the particles may be able to escape by resampling induced relocation or luck, if one is to rely on such methods much cheaper alternatives exist than HMC, in particular confessedly luck-based methods (random walks). If there aren't many such asymptotes, a possible way of escaping these regions would be to screen for vanishing acceptance rates and perform a RWM step after HMC if they're found. Obviously, such a sequence is posterior-invariant if each step is, and by slightly moving the particle one may hope to introduce some chance that it crosses the potential wall and is positioned more favorably to Hamiltonian dynamics. This is attempted in section 4.2.1. However, it may lead to expensive proposals often being constructed just to be (with near unit probability) discarded and replaced by cheaper ones, which is rather unproductive.

### 3.4.2 Sampling a trajectory and state

This section considers variations in how the HMC proposals are picked. It is based on [15, 59].

In the HMC implementation of algorithm 7, the last state obtained by integrating was proposed as sample, and all others ignored. A Metropolis-Hastings acceptance/rejection step was then performed for correction, which was necessary due to the simulation of the Hamiltonian dynamics' not being exact. In that case, HMC could be seen as a particular (and remarkably creative) case of the Metropolis-Hastings algorithm.

The reason why the last state is generally adopted as proposal is that it tends to be more distant to the initial point, meaning we expect to get less correlated samples. However, this isn't a perfect rule, nor is it the only possible one. Clearly, the energy conserving properties of the symplectic integrator apply to the full trajectory. As such, any other state in the trajectory could be proposed, and accepted or rejected based on its own coordinates.

<sup>10</sup>Though some ways to deal with discrete or discontinuous likelihoods have recently been proposed [91], they are certainly not tolerated by construction.

As a matter of fact, a sample from the trajectory will asymptotically correspond to a sample from the augmented target distribution (i.e. as the path length goes to infinity). Accordingly, we could pick a point from the trajectory uniformly at random, and perform the corresponding Metropolis correction.

While this is a valid strategy, it is not the best, because each trajectory point is assigned equal probability. This isn't an issue in what respects correctness, as long as the correction is performed, but we may get unnecessarily low acceptance probabilities when there would have been better proposals in the generated trajectory. In this scenario, we would with considerable probability be forced to discard the point, and the construction of the entire trajectory would be in vain.

An alternative is to propose the points with probabilities proportional to the augmented probability density at their phase space locations. From here on, this is what is intended with the phrase "sampling a point from the trajectory" unless otherwise specified. This could be done using for example slice sampling, or simply multinomial sampling. It requires extra target evaluations, but also allows for devising more efficient sampling schemes.

The most basic approach would be to generate the whole trajectory while keeping track of all points, and then sample one of them using some sampling strategy of choice. We must in this case register not only the points' positions, but also their momentum coordinates, so as to be able to evaluate the augmented target density at each of them. As this scheme by itself safeguards the invariance of the target distribution under the Markov moves, the Metropolis correction is no longer necessary (see [15] for a formal analysis). As a matter of fact, this is no longer a Metropolis scheme.

The fact that all states must be kept in memory is highly discouraging, but better tactics can be used that improve memory management all the while performing equivalently. This can be done by choreographing the trajectory expansion in such a way that we can sample from it along its construction. The key thing to note is that if we have some initial point and want to integrate the Hamilton's equations for some number of discrete time-steps  $L$ , picking a trajectory uniformly at random from all ( $L$ ) possible ones and *then* sampling a point from that selected trajectory amounts to sampling from all reachable points under the same conditions.

Picking a trajectory at random from all that could contain the initial point given a fixed trajectory length is straightforward: it suffices to simulate the dynamics incrementally, successively choosing between forward or backwards time evolution (i.e. we would use positive or negative time-steps in the leapfrog integration). We would start at the current extremes of the trajectory (last/latest in time when integrating forward, first/earliest in time when integrating backwards) and integrate until the desired length was achieved.

Essentially, we would be progressively appending extra states to the previously obtained trajectory, producing a new and longer one. These states can be added one or a few at a time in an additive scheme, but for faster progression a multiplicative scheme can be used. A common choice is to append a new trajectory as long as the last one, effectively doubling the trajectory length at each iteration.

In that case, we could generate a trajectory  $A$  of length  $W$  and keep a representative sample from it, along with the cumulative target density associated with the ensemble of points (which we will call its weight). Then we would pick at random a direction for the ensuing simulation, and take the corresponding

extreme of trajectory  $A$  as the initial state for a second trajectory  $B$ . Finally, we would build trajectory  $B$  until its length matched that of  $A$ . After this, we could concatenate those trajectories, and keep the sample from  $A$  or  $B$  with probabilities proportional to the weights of the entire trajectories. This provides a sample from the concatenated trajectory  $A + B$ .

In the first iteration, the  $W$  above would be one. That means that each trajectory would contain a single point; the first one would be constituted by the initial state, and the second one by one of the two possible adjoining states in phase space (given the integrator parameters, namely the stepsize, the evolution is deterministic; the only random choice is the direction for the time evolution). In this case, the trajectories' weights would be simply the extended density at their only constituents, and the samples representing them would be the only possible one for each. As such, this base case corresponds exactly to sampling a point from the whole two-state trajectory. Joining the two parts together, we would get a length-two trajectory, and a sample representing it. Summing up the individual trajectories' weights, we'd also get a weight for the compound trajectory.

Then we would integrate again; backwards in time with 50% probability, and forward with another 50%. This integration would start one of the two points obtained so far, depending on their relative time positioning. If choosing a doubling scheme, we would advance by two integrator steps. After that, we would proceed similarly for the length four trajectory. With the additive one-by-one scheme, the new trajectory would always contain a single point, and its weight would always be the density at that point. In either case, this would be repeated until the intended path length was reached.

So far we have assumed sampling proportional to the trajectory weights, but a higher efficacy can be achieved by introducing bias. This is another benefit of progressive sampling, that wouldn't be possible in the first two schemes outlined above (picking a state uniformly and performing a Metropolis correction, and sampling from the whole trajectory after building it). It is possible to attribute higher probability to the newest trajectory's sample, reducing correlation - and that too while preserving correctness, as exposed in [15].

In the scheme proposed therein, should the newest trajectory outweigh the former or their weights match, the newest trajectory's representative sample is kept, resulting in a proposal which is further away from the initial point (which obviously always falls within the oldest/previous trajectory segment). Even if it doesn't, the new trajectory will be favored over the old one. The Markov chain's relevant properties remain unharmed, because this bias is cancelled out if we consider all possible initial points: should the initial point have been in the newest trajectory instead, the success probability would be inverted with respect to the trajectories.

### 3.4.3 Dynamic implementations: the no U-turn sampler

We have seen that the opportunity to introduce anti-correlation measures makes it beneficial to build the trajectory as we go, until it achieves some *desired length*. This length is the tunable  $L$  of subsection 3.3.2. It had been noted that it is a critical parameter for the performance of HMC, it being that the trajectory

should stop short of dwindling returns. However, a preliminary calibration is speculative at best, given that we don't know what the trajectory will look like. Why not then judge this in the light of its current state, if we are to be sitting at its frontier? Such is the idea presented in [15, 59], which we will again follow.

When constructing the trajectory along the way, one can opt for a *dynamic* scheme, where the number of integration time-steps is chosen adaptively based on some termination criterion. This can bring sizeable gains, since even perfect static tuning is in general sub-optimal: because the random momentum resampling at each Markov chain state lifts the sample up to a random energy, the ideal path length varies at each step, even for the same position coordinates.

In general, to different initial phase space states correspond different energies. Depending on the energy, a longer or shorter evolution may be required to properly explore the associated phase space region. If using a one-size-fits-all approach for the path length selection, we will be under-exploring some energy level sets (falling short of the returns they might otherwise have contributed with), and over-exploring others (wasting computational resources for little benefit). It is then clear it would be beneficial to introduce some adaptivity in the extent of this exploration.

This can be realized by defining a criterion with whose satisfaction the trajectory draws to a close, finalizing the state transition. A dynamic criterion can be contingent on the current standing of the trajectory, and is usually made to depend in its endpoints - earliest and latest in time, respectively  $(\theta^{(-)}, p^{(-)})$  and  $(\theta^{(+)}, p^{(+)})$ . A well advised choice is the no-U-turn termination criterion (with which HMC is called the No U-turn sampler, or NUTS) [59, 105, 15]:

Terminate if:

$$p^{(+)} \cdot (\theta^{(-)} - \theta^{(+)}) > 0 \text{ AND/OR } -p^{(-)} \cdot (\theta^{(+)} - \theta^{(-)}) > 0 \quad (3.104)$$

Each of these conditions signals that the velocity at an endpoint of the trajectory has the same direction<sup>11</sup> as the distance vector from that endpoint's position to the other's. That is, integrating further from that extreme while keeping the other one fixed is expected to bring the endpoints closer together - or in other words, a U-turn is made. This is a rough indication that most of the exploration has taken place, and continuing it will bring reduced returns. Note that when integrating backwards the effective momentum is negated, since rewinding time reverts the direction of the velocity.

Some sources suggest taking the disjunction of the conditions [59], while others propose the conjunction [15]. The former option may be too aggressive, since it stops as soon as any of the points moves towards the other (which may be moving away still faster); by contrast, the latter ensures that the endpoint are indeed coming towards each other, which may be too conservative. Which is best will depend on the application.

An important point to note is that this criterion can't be applied naively, under penalty of forfeiting correctness. To ensure reversibility, any dynamic termination scheme must be implemented carefully.

In order for the Markov chain to admit the target distribution as stationary distribution, any two states in every trajectory must generate that trajectory with equal probability. Sampling from all the phase space

<sup>11</sup>In the sense that the vectors make an acute angle, not that they're in perfect alignment.

points contained in it with probabilities proportional to the extended density then preserves the augmented target, and in particular the original target.

In static schemes, the trajectory generated by any initial state can easily be picked at random from all possibilities by randomly switching between time directions, and this principle will be respected. However, when considering dynamic schemes where adaptivity is employed in the choice of path length things get somewhat trickier. Instead of expanding our trajectory a fixed amount, we expand it until a requisite is satisfied. The problem is that we can't neglect all the ways this requisite *could have been* satisfied, had the construction gone differently. In other words, the termination criterion might have cheated other trajectory elements out of their chances of reconstructing the trajectory. A point belonging to a trajectory it had no chances of originating is a clear violation of our condition.

If for instance we have some current trajectory, append a point to an extreme, and that point fulfills the termination criterion with any point other than its opposite extreme, that would mean that had the trajectory started at it, it would have stopped earlier. The aforementioned reversibility principle would then be breached, shattering the theoretical foundations of the Markov chain and rendering the construction invalid for our purposes. More specifically, this defies the validity of any estimator obtained by simulation.

This can be corrected by scanning through newcomer points that could have matched pairwise to provoke an early termination, had the trajectory started somewhere else; they are intruders, and should be excluded. That is, one must ensure that *any* point that we welcome into our trajectory would have been able to, as a starting point, recreate it from end to end. Here an exponential expansion strategy performs much better than an additive one: it allows for saving on these correction checks, because it is more selective in the trajectories it allows [15].

To any trajectory built through a doubling scheme corresponds a balanced binary tree, so for any path that we want to concatenate with our previous trajectory we need only check its balanced binary sub-trees (of which there are as many as nodes). By contrast, a one-by-one addition scheme would require any new arrival to be compared with every single one of the older states.

Independently of the strategy, if for any new state in the to-be-appended trajectory the termination criterion would have brought the process to a halt, the construction must be at once interrupted. The entirety of the new trajectory is to be discarded, and the old one to be taken as is. This is the downside of a bold expansion scheme: if considering a single addition at a time, we would never discard more than one point, whereas when doubling we may have to discard as many as we actually use (the latest correct trajectory).

Although it may sound unreasonable that we'd be asked to dispose of a trajectory packed with potential samples, the method is still less wasteful than a one-size-fits-all approach, and is commonly used with a doubling progression scheme.

An additional point to note is that when biasing sample selection (as discussed in the previous section) within such a scheme, one may choose to introduce bias in the construction of the entire trajectory (including all the sub-trees when constructing a new stretch) or only when appending a new trajectory to the old one - that is, when doubling. The original paper [59], which coined the term **NUTS**, suggests opting



for the latter option. Details for the efficient implementation of all the ideas mentioned here (and others) can be found in that paper, as well as performance assessments.

The algorithm was implemented along with other sample picking strategies, and the results are shown for a Rosenbrock function in subsection 4.1.2. The no U-turn sampler performs better even when matching the static and dynamic strategies for total states covered, i.e. including the discarded ones.

## 3.5 Subsampling strategies

Hamiltonian Monte Carlo is a very powerful method, and one which has enjoyed a lot of success - but it is also quite computationally costly. Monte Carlo methods are in general quite computationally demanding, but this is particularly true of HMC. In inference, likelihood evaluations are the most expensive resource, and HMC requires both itself and its gradient to be evaluated at each point of the phase space trajectory brought about by the (discretized) Hamiltonian dynamics. It is then clear why subsampling strategies may be of interest when scaling up complexity (namely dimensions and/or dataset size) - the data batch size may determine whether HMC is an efficient or even viable solution in view of the available memory and processing resources. As an added benefit, it may increase efficacy if employing parallel processing, which may be too memory-intensive if the full dataset is to be available to each worker [54].

Unfortunately, naive subsampling strategies can be catastrophic for the performance of HMC, debasing the very features that make it so appealing [14]. As such, subsampling strategies must be developed either carefully, or not at all. This section will present a shallow view of a few proposed solutions in the literature. The main references are [14, 20, 96, 23, 54]. Subsection 3.5.1 presents a simple and fairly general estimator, whereas 3.5.2 and 3.5.3 outline the ideas of stochastic gradient (SG)-HMC and energy conserving subsampling (ECS)-HMC respectively. The latter is more prudent approach, relying on a pseudo-marginal method to refresh the subsample so a HMC (within Gibbs) step can be performed in a seemly manner.

### 3.5.1 Unbiased estimators and variance reduction with control variates

We start by discussion likelihood estimators, based on [96, 23, 54]. Ideally we would use the exact likelihood and log-likelihood, which depend on all data:

$$L(\theta) = \prod_{k=1}^N L_k(\theta | D_k) \equiv \prod_{k=1}^N L_k(\theta) \quad (3.105)$$

$$\ell(\theta) = \sum_{k=1}^N \ell(\theta | D_k) \equiv \sum_{k=1}^N \ell_k(\theta) \quad (3.106)$$

However, this may be infeasible or impractical for large datasets, which motivates the use of mini-batches sampled from the original data record  $\vec{D}$ . That is, we would like to construct an unbiased estimator based on a subsample of  $m < N$  observations.



If  $\widehat{\ell}_m(\theta)$  is an unbiased estimator of the log-likelihood based on subsampling  $m$  data and  $\widehat{\sigma}_m^2(\theta)$  estimates its standard deviation, we can estimate the likelihood by:

$$\widehat{\mathbf{L}}_m(\theta) = \exp\left(\widehat{\ell}_m(\theta) - \frac{1}{2}\widehat{\sigma}_m^2(\theta)\right) \quad (3.107)$$

The motivation is that, by linearity of expectation, an unbiased estimator can easily be found for the log-likelihood, and  $\widehat{\mathbf{L}}(\theta)$  will be unbiased too if  $\widehat{\ell}_m(\theta)$  is normal (which is likely for large  $N$  as per the CLT) and  $\widehat{\sigma}_m^2(\theta)$  is exact. The variance of the estimator corrects for bias; if it is approximate (a more realistic assumption), the estimator in 3.107 is approximately unbiased.

As for the log-likelihood estimator, the starting point is:

$$\widehat{\ell}_m(\theta) = \frac{N}{m} \sum_{j=1}^m \ell_{u_j}(\theta) \quad (3.108)$$

, where the  $u_j$  are subsampling indices sampled at random:  $u_j \in \{1, \dots, N\}$ . Once we fix the  $m$  elements of  $\vec{u}$ ,  $\widehat{\ell}_m(\theta)$  is a value that estimates  $\widehat{\ell}(\theta)$ . Evaluating it has cost  $\mathcal{O}(m)$ .

Clearly,  $\mathbb{E}[\widehat{\ell}_m(\theta)] = \ell_m(\theta)$  by linearity of expectation. However, subsampling tends to increase the variance in the estimate. A well-known variance reduction technique for Monte Carlo estimates is the use of control variates<sup>12</sup>. They exploit the fact that the expected value is unaltered if we add a term and cancel out its expectation:

$$\widehat{\ell}_m(\theta) = \frac{N}{m} \sum_{j=1}^m \ell_{u_j}(\theta) + C \left( \frac{N}{m} \sum_{j=1}^m q_{u_j}(\theta) - \sum_{k=1}^N q_k(\theta) \right) \quad (3.109)$$

The  $q_k(\theta)$  are auxiliary random variables called *control variates*. Even though the expectation is unchanged, the variance is not, and can be minimized through the choice of the  $q_k$  and  $C$ . If we have some way of approximating  $\ell_k(\theta)$ , one way of achieving this is by choosing  $q_k(\theta) \approx \ell_k(\theta)$  and setting  $C = -1$  to get the so-called (for obvious reasons) *difference estimator*:

$$\widehat{\ell}_m(\theta) = \sum_{k=1}^N q_k(\theta) + \frac{N}{m} \sum_{j=1}^m \left( \ell_{u_j}(\theta) - q_{u_j}(\theta) \right) \quad (3.110)$$

The  $q_k(\theta)$  can be Taylor series around some reference value  $\theta^*$  for the distribution<sup>13</sup>. For MCMC within SMC, the reference  $\theta^*$  can be the current SMC mean  $\bar{\theta}$ , and for MCMC it can be determined in a preliminary processing phase. In both cases, the coefficients for the Taylor expansion can be calculated once using the full dataset and are then used for all parallel particle updates/transitions (for SMC/MCMC respectively), signifying  $\mathcal{O}(1)$  cost (i.e. no likelihood evaluations, only distances with respect to  $\theta^*$ ). This implies that 3.110 has cost  $\mathcal{O}(m)$  just like 3.108.

<sup>12</sup>Of course, other variance reduction techniques exist, some of which have been successfully employed in HMC [70].

<sup>13</sup>This assumes unimodality (unless provisions are made, such as clustering the data to form multiple centroids  $\theta_k^*$  to serve as references).

Finally,  $\widehat{\sigma}_m^2(\theta)$  too can be obtained from the subsample as the variance of the differences in 3.110 (between parenthesis). This clears the requirements for the computation of the estimator in 3.107, which can be used in any sampling strategy, namely any MCMC and SMC method (for both the weight updates and the particle propagation steps in the latter case). In HMC, the potential energy would become:

$$\widehat{U} = -\ln\left(\widehat{\mathbf{L}}_m(\theta)\right) \quad (3.111)$$

We don't use  $\widehat{U} = -\widehat{\ell}_m(\theta)$  directly because we are interested in the unbiasedness of the target distribution<sup>14</sup>, the potential energy being simply a derived construct. The HMC updates require the log-gradient through  $-\nabla\widehat{U}$ , which can be obtained from the sample as well by differentiating 3.110 and its variance. Most of these calculations are laid out in [96], which additionally shows that the error in expected values associated with the perturbed posterior decreases as  $\mathcal{O}(1/(Nm^2))$ . In particular, this is the expected convergence rate of the calculated mean and variance for the posterior distribution.

For instance, letting  $m = \mathcal{O}(\sqrt{N})$  means that their error decreases quadratically with the total number of observations, which is very favorable. However, this is assuming that  $\theta^*$  is the posterior mode based on the full dataset. For SMC, there are convenient statistics at hand via the particle cloud, whose mean we can use directly; for MCMC, auxiliary methods such as gradient descent become necessary. In that case, *stochastic* gradient descent based on a fixed subsample may be a less purpose-defeating assumption, but it slows down convergence. If the gradient is also based on  $m = \mathcal{O}(\sqrt{N})$  observations, the rate becomes  $\mathcal{O}(N^{-1/2})$ .

Being in possession of an estimator, we can simply use it instead of the complete likelihood (or its byproducts) when performing inference. However, this doesn't necessarily guarantee a good cost-to-performance ratio; improvements in the estimator *do* bring improved results - but this is as compared to a crude approximation. When the reference is the fully proper algorithm, this is not necessarily so.

In the case of HMC, it's particularly tempting to subsample the gradients, which are responsible for the lion's share of the cost. The subsampling indices could be sampled only at the beginning of each trajectory (picking a new  $\vec{u}$  for each iteration of algorithm 7), or within the trajectory itself (picking a new  $\vec{u}$  for each iteration of algorithm 6).

If we do this but impose the usual Metropolis corrections based on the full dataset, which is in itself costly, we correct the bias that we are introducing. However, we are likely to dramatically lower the acceptance rate, since the energy is no longer stable. If we skip this correction to avoid processing all the data, the trajectory will be biased - just how much depends on the quality of the estimators, the features of the data, and the details of the implementation. Resampling the subsampling indices at each step of the trajectory performs better than not, but still not satisfactorily so [14]. As such, this approach severely compromises the scalability that HMC prides itself on. In this context, some revisions of the algorithm have been suggested.

<sup>14</sup>Though one *could* ignore this and opt for a more biased but simpler estimator based on a sub-product, which may still do well depending on the implementation details and structure of the data.

### 3.5.2 Stochastic gradient Hamiltonian Monte Carlo

Reference [20] adopts an intuitive view of the dynamics behind HMC in an attempt to remedy the downfalls of subsampling. They point out that using a subsampled gradient amounts to using a noisy gradient, which by the CLT they write as:

$$\Delta\widehat{U}(\theta) \approx \Delta U(\theta) + \mathcal{N}(0, V) \quad (3.112)$$

, where  $V$  is the (co)variance of the stochastic gradient noise and may depend on  $\theta$ . This defines a diffusion matrix that shows an additional dependence on the stepsize,  $B = \epsilon V/2$ , which in turn determines the unwanted random force that affects the dynamics.

It is shown that the dynamics associated with 3.112 fail to preserve the target distribution, leading to divergent trajectories. The potentially catastrophic impact of noise in the dynamics is demonstrated in figure 27 for a simple oscillator.




Figure 27: Effect of noise on the phase space trajectory for a harmonic oscillator model ( $U = \theta^2/2$ ). The trajectory with correct dynamics is plotted in black, that with a noisy gradient in green, and that with a noisy gradient compensated by a friction term in a gradation from blue to red (by ascending order of elapsed time). This example is based on an identical one from [20].

To mend this problem, the authors suggest introducing a friction term inspired in the Langevin equation (a stochastic differential equation initially used to describe Brownian motion [106]):

$$\frac{dp}{dt} = -\frac{\partial\widehat{U}}{\partial\theta} - \gamma\frac{p}{M} \quad (3.113)$$

This presents a modification of 3.87. For  $\gamma = B$ , equilibrium is restored. In general, the exact noise model won't be known exactly, but rather approximately; the algorithm can be adapted accordingly. The

variance  $V$  can be estimated using the inverse of the Fisher information matrix based on a data mini-batch. All details can be found in [20]. Results of implementing the algorithm are shown in subsection 4.2.3.

The authors suggest omitting the Metropolis corrections entirely to spare costs, and instead controlling the quality of the sampler (namely the bias) through the discretization of the dynamics (by keeping the integration stepsize small). However, there is then a sensitive trade-off between bias and integration cost for a fixed evolution time. Overall, even though it helps control divergence, this scheme may undermine the efficiency of HMC, especially in high-dimensional scenarios [14, 23].

### 3.5.3 Hamiltonian Monte Carlo with energy conserving subsampling

A possible pathway to a more correct subsampling HMC strategy is an extended target density. We will consider an approach presented in [96] for MCMC in general, in [23] for HMC in particular, and still more particularly in [54] for SMC with HMC move steps. The idea is to have the target distribution encompass the subsampling indices, as the random variable that they are:

$$\mathbf{P}_X(\theta, \vec{u}) = \widehat{\mathbf{L}}_m(\theta, \vec{u})\mathbf{P}(\theta)\mathbf{P}(\vec{u}) \quad (3.114)$$

When performing HMC steps, only the posterior is targeted; that is, the Hamiltonian dynamics operate for a constant  $\vec{u}$ <sup>15</sup>. Accordingly, when we further augment the distribution with a momentum variable, it respects  $\theta$  only, and matches it in dimension. Thus, within HMC:

$$\mathbf{P}_X(\theta, \vec{u}) = \exp\left(-\widehat{H}_{\vec{u}}(\theta, \vec{p})\right)\mathbf{P}(\vec{u}) \quad (3.115)$$

This doesn't mean that  $\widehat{H}_{\vec{u}}(\theta, \vec{p})$  doesn't depend on  $\vec{u}$ : it does, because now the potential energy is  $\widehat{U}(\theta) = -\ln \widehat{\mathbf{L}}_m(\theta, \vec{u}) - \ln \mathbf{P}(\theta)$ . However, this vector is fixed within the Hamiltonian, its variability being relegated to this vector's own distribution  $\mathbf{P}(\vec{u})$ .

Still,  $\vec{u}$  must be updated, or else the extension serves no purpose. That can be done with a simple Metropolis step (using a uniform proposal distribution, which is clearly symmetric too), as long as  $\theta$  is held fixed. Updating each parameter by conditioning it on all others is a strategy called *Gibbs sampling*, a special case of the Metropolis-Hastings algorithm where we sample each variable individually while sweeping through the parameter vector [5].

Here the separation is between  $\theta$  and  $\vec{u}$ . We want to sample  $\vec{u}$  according to  $\mathbf{P}_X(\theta, \vec{u})$  conditionally on  $\theta$ ; we can do so with a MH step. This strategy is called Metropolis-within-Gibbs, or more descriptively one-variable-at-a-time Metropolis. We fix  $\theta$ , sample new indices uniformly, and accept them with probability:

$$a(\vec{u}', \vec{u}) = \min\left(1, \frac{\widehat{\mathbf{L}}_m(\theta, \vec{u}')}{\widehat{\mathbf{L}}_m(\theta, \vec{u})}\right) \quad (3.116)$$

<sup>15</sup>Auxiliary variables have been attributed a momenta of their own in HMC [1], but typically they're continuous, as in latent variable models. HMC isn't made for discrete domains like that of  $\vec{u}$ , though accommodations can be made [91]

Note that  $\theta$  is fixed, corresponding in practice to the latest  $\theta$  iterate. This is an instance of *pseudo-marginal* Metropolis-Hastings (*pseudo-marginal (PM)-MH*), a version of *MH* where the likelihood isn't available analytically and a non-negative unbiased estimator takes its place. This preserves correctness (the chain's invariant distribution is retained), earning marginal methods the title of *exact approximations* [114].

A common example comes from latent variable models whose integrals can't be solved analytically: importance sampling can be used to get a Monte Carlo estimator of the likelihood (this requires only pointwise evaluations of the integrand, and is for that reason often called a *likelihood-free* method). An alternative would be to target the joint distribution of parameters and latent variables. This is less efficient than the *marginal* methods (such as the ones considered so far), which target the intended distribution directly; as such, pseudo-marginal methods attempt to bring the approach closer to them [4, 1].

In our subsampling case, we introduce not evaluation sites but subsampling indices, since we are truncating a sum and not an integral. Either way, this amounts to extending the target density to encompass some extra random variables, now the auxiliary samples used to estimate the likelihood [4, 108, 114, 96]. They can be refreshed at each iteration; this refreshment is factored in when computing the joint acceptance rate [4, 108] (as it is in 3.116). That there is variability in the compound *MH* step for even fixed  $\theta$  - equivalently, that the distribution is extended - slows down convergence, how much depending on the variance of the estimator (hence the attempted reduction through e.g. control variates) [114].

One may wonder whether the samples can't simply be picked anew at each iteration to instance the estimator and then discarded. This is sometimes done, and given the self-explanatory name of Monte Carlo within Metropolis (in the case of a Monte Carlo likelihood estimator). However, in that case, absolute validity is only achieved in the case of perfect *estimation*; otherwise the scheme is rather an *inexact approximation* whose accuracy grows with the number of samples [4].

A possible complication of 3.116 is index stagnation. Should the estimator significantly overestimate the likelihood associated with some  $\vec{u}$ , others will suffer from low acceptance, *jamming* the sampler and affecting validity. To solve this, correlation can be introduced between consecutive proposals, by updating only a segment of the previous index set in what is called a *block pseudo-marginal (bPM)* strategy [108]. This should make the acceptance rate more stable, preventing the monopoly that could happen otherwise.

After the index updating step, a *HMC* step can be performed while holding  $\vec{u}$  fixed (targeting 3.107). The *HMC* moves preserve the energy because the *MH* correction step targets precisely the energy according to which the Hamiltonian dynamics are simulated, guaranteeing high acceptance as well as correctness.

In short, *ECS* is achieved by wrapping the *HMC* transitions in a Gibbs update, in which a block Metropolis-Hastings step is used to pick the subsampling indices (conditioned on the latest particle location) before the usual *HMC* step (conditioned on the indices). This amounts to a pseudo-marginal approach where the posterior is augmented to include some supplementary sample identifiers, the joint density being left invariant by the composite - Gibbs - Markov transition comprised of two *MH* steps (one of which is *HMC*). Marginalizing over the indices  $\vec{u}$  should yield the intended  $\theta$  samples.

This approach was applied to *TLE-SMC* (with *HMC-ECS* for particle propagation); the results are shown in subsection 4.2.3. All of the implementation details can be found in [54].

# Applied sampling and inference examples

This chapter presents a few test cases for the methods presented in chapters 2 and 3. The tests on actual quantum devices are left for chapter 5.

Section 4.1 applies statistical simulation techniques to hard-to-sample-from density functions. These functions are non-negative, describing probability distributions up to a constant - but not Bayesian posterior distributions. They consist of a 6-dimensional Gaussian (subsection 4.1.1), a Rosenbrock function (4.1.2), and a smiley-like structure built from a pre-determined set of points (4.1.3). Their complexity is due to high dimensionality, sharpness and multimodality respectively. All the subsections test MCMC as described in section 3.3. Two instances of Metropolis-Hastings are considered: random walk Metropolis, and Hamiltonian Monte Carlo.

The performances of static, progressive and dynamic variations of HMC (as per subsection 3.4.2) are compared in subsection 4.1.2. Everywhere else, a standard static last-state implementation of HMC (algorithm 7) is implicit. Subsection 4.1.3 concerns a smiley-shaped trimodal distribution, to which SMC is applied in addition to the other methods. Two different Markov kernels are tested: RWM and HMC.

In section 4.2, sampling methods are instead used in inference problems, using data generated according to the intended likelihoods. Subsection 4.2.1 applied Gaussian rejection filtering and MCMC (RWM and HMC, section 3.3) to phase estimation (subsection 2.3.3). Subsection 4.2.2 targets the precession example of subsection 2.3.3 using the same MCMC methods, and additionally SMC-SIR with LWF (subsections 3.2.2 and 3.2.1) as well as with HMC move steps (this being typically abbreviated as SHMC). The example is easily made to result in multi(bi)-modal posteriors, the nemesis of the Liu-West kernel shrinkage method - effectively orchestrating the ruin which befalls it in a wide array of realistic applications.

Finally, subsection 4.2.3 generalizes the precession likelihood to higher dimensions, to test the methods in more challenging settings. Results are shown for SG-HMC (subsection 3.5.2) within SMC-SIR, and for TLE-SHMC-ECS (subsections 3.2.3, 3.3.2, 3.5.3) with control variates (section 3.5.1). Finally, some heuristics for choosing likelihood hyperparameters (*experiment controls*) are tested.

## 4.1 Sampling problems

These examples are based on [24].

### 4.1.1 6-dimensional normal distribution

The target for this subsection is a multivariate normal distribution. This example showcases the benefits of **HMC** as compared to **RWM** or other basic **MH** algorithms. The distribution was chosen to have mean  $\vec{\mu} = [20, 20, 20, -20, -20, -20]$  and identity covariance. The results are shown in figure 28.

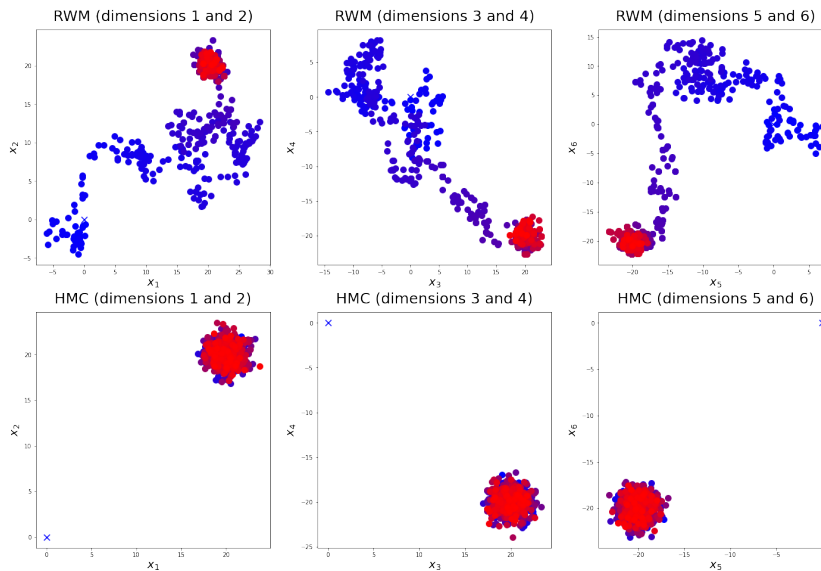


Figure 28: Trajectories generated using **RWM** (top) and **HMC** (bottom) to sample from a multivariate normal distribution with 6 dimensions/parameters. Both chains started at the origin and were evolved for  $10^3$  states. These states are plotted as points gradating from blue to red (from earlier to later iterations), whereas the starting point is marked with an 'x'. The dimensions are plotted pairwise for convenience.

The **RWM** proposal variance was such that the mean acceptance rate was 42%, whereas for **HMC** it was 100% with  $L = 30$ ,  $\epsilon = 0.05$ ,  $\mathcal{M} = \hat{\mathbf{I}}_6$ . **HMC** manages a golden combination of distant proposals and very high acceptance probabilities, whereas **RWM** struggles a fair bit more to reach the mode. It should be mentioned that they're not matched for resources; each **HMC** step is by nature more expensive, due to how it is informed by the target's geometry. However, larger and larger distances will render **RWM**'s wandering less and less productive, hence it is bound to become less cost-effective be it sooner or later.

Furthermore, not only does the **HMC** chain mix faster, there is also less correlation between consecutive samples, as can be seen from both the unimportance of the starting point and the uniform red color in figure 28: the latest iterations alone cover the mode thoroughly and evenly, which can't be said of **RWM**. This reduces the number of samples necessary to comply with accuracy requirements for the expectations.

The performance of **HMC** is especially impressive in this example: its proposals are particularly good if the target itself is Gaussian, the usual Gaussian distribution over the momentum is chosen, and the covariance matrices of these distributions are alike [15]. While in practice sampling from a Gaussian isn't a great achievement, by the **CLT** many posteriors will approach normality for large enough datasets. Thus, the virtues that shine through in this particular case can be expected to manifest in others.



### 4.1.2 Rosenbrock function

The target for this subsection is a Rosenbrock function as defined as in [24] (where it seems to be exponentiated as compared to its usual form):

$$g(x, y) = \exp \left[ \frac{1}{8} (-5(y - x^2)^2 - x^2) \right] \quad (4.1)$$

The resulting density resembles a narrow smile with sharp corners. Its shape is plotted in figure 29. The key thing to note other than its overall form is that the corners of the *mouth* reach all the way up to  $y \approx 55$ , with a probability density that decreases rapidly with increasing  $y$ . This feature is especially hard to resolve, because Markov chains struggle to explore such narrow and low probability regions.




Figure 29: 3-dimensional plot of the Rosenbrock function from equation 4.1.

The results of using **RWM** and **HMC** to sample from this function are shown in figure 30. The **RWM** proposal standard deviation was fixed at  $\text{Cov}_{\text{prop}} = \text{diagonal}(0.15)$ , yielding a 68% acceptance rate, whereas for **HMC** the acceptance rate was 98% with  $L = 2^6$ ,  $\epsilon = 0.1$ ,  $\mathcal{M} = \hat{\mathbf{I}}_2$ .

Clearly, **HMC** achieves a resolution that **RWM** does not: the momentum is vital for climbing up the elongated corners of the shape. Owing to it, **HMC** comes much closer to reaching the regions with largest values of  $y$  and non-null density (which occur at  $y \approx 55$ ), reaching roughly  $y \approx 40$  while **RWM** fails to achieve even  $y = 10$ . This result could be made still better by increasing  $L$  or the number of transitions.






Figure 30: Trajectories generated using RWM (30a) and HMC (30b) to sample from a (bivariate) Rosenbrock function. Both chains started at the origin and were evolved for  $10^4$  states. These states are plotted as points gradating from blue to red (from earlier to later iterations).

Another way of improving it without increasing any iteration number is to refine the algorithm. Figure 30b used the best known HMC variant, where the last state of the trajectory is deterministically appointed as the proposal, and a probabilistic MH accept/reject step is performed (algorithm 7). Figure 31 shows the results of using the variants discussed in subsections 3.4.2 and 3.4.3 - namely uniform sampling, biased progressive sampling and the NUTS. The best performing one is the NUTS, which exhibits a rather impressive ascent to the edges all the way up to the maximal values at  $y \approx 55$  (though they're visibly less populated than the main body, in conformity with figure 29).

In these cases, a doubling progression scheme was used for generating the trajectories. In the static biased strategy, bias was introduced both when generating and concatenating trajectories, as this was the option that worked better. For the NUTS, however, only when joining trajectories were the probabilities skewed, as in the original implementation of [59]. The difference relative to the strategy exposed therein is that multinomial sampling was used instead of slice sampling for picking the proposals.

The HMC parameters  $\epsilon$ ,  $\mathcal{M}$  were the same as in the previous case (figure 30b). Due to using a dynamic termination criterion, the NUTS doesn't allow for pre-determining the path length  $L$ . For assessing the performance, the average generation trajectory length per transition  $L_{avg}$  was computed, which determines the total cost. For a fair comparison, all other methods were chosen to have  $L > L_{avg}$ ; the NUTS still outperformed them. An important point is that this  $L_{avg}$  includes *all* the computed transitions, counting those that had to be discarded to guard reversibility. By contrast, the other methods always use all the generated states; again, the NUTS outperforms them regardless. This proves what was claimed in subsection 3.4.3.

Another subtlety mentioned in that section was the use of a conjunction or a disjunction of the 2 U-turn conditions. For figure 31c, an OR was used, but the AND was also tested. It increased  $L_{avg}$  to  $2^{7.7}$  while barely enhancing the exploration, so in this case the more restrictive condition worked best.




Figure 31: Trajectories generated using different variants of HMC to sample from a (bivariate) Rosenbrock function. They differ (both between themselves and from the standard case of 30b) in how the trajectories are generated and their samples picked. The sample selection methods are: uniform sampling (31a), biased progressive sampling (31b) and NUTS with partly biased progressive sampling (31c). Figures 31a and 31b used static implementations with  $L = 2^6$ , whereas 31c's dynamic termination produced on average  $L = 2^{5.4}$ . All chains started at the origin and were evolved for  $10^4$  states. These states are plotted as points gradating from blue to red (from earlier to later iterations).

### 4.1.3 Smiley kernel density estimate

The target for this subsection is a smiley face. A smiley-like density can be formed by joining together three functions similar to the one of the previous section, and rotating or flattening them to taste. Instead of using that density directly, we sample 768 points from it and smooth them into a continuous density. This smoothing was achieved by a gaussian kernel density estimate (KDE) (figure 32a).

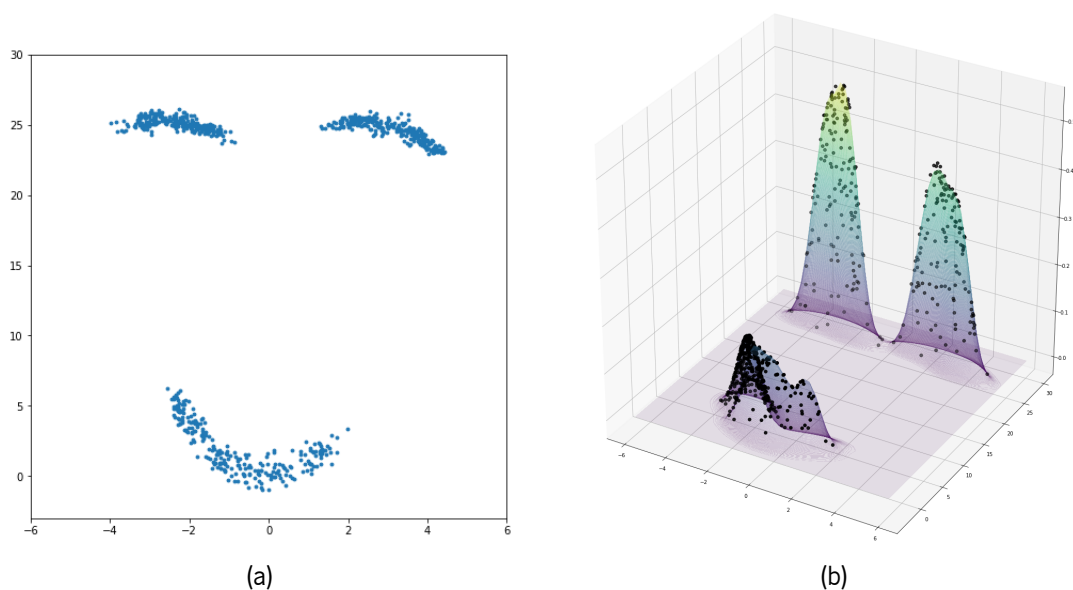


Figure 32: Target bivariate smiley function (32a), generated using 768 gaussian kernels, and 3-d scatter of the particle cloud produced by SHMC (32b) plotted over a surface plot of the target.

The goal is to make the data divisible: by only considering a fraction of the points, we can make the function shallower. This provides a way of creating a sequence of targets for **SMC** without having to consider inference. The data can be divided into a list of chunks, and that list traversed at the rate of an item per iteration. The target density at iteration  $i$  is the data contained by the first  $i$  elements (counting from 1). The 768 points were split into 8 data-chunks, 7 containing 100 points and 1 containing the leftover 68. The first iteration of **SMC** then targets the first one, the second iteration targets the first one plus the second one, and so on until the eighth iteration.

Note that in principle obtaining the problem-statement set of points is in itself a hard sampling problem. However, it can be made easier by splitting the probability density into 3 local densities and starting 3 separate Markov chains around the center of each eye and the mouth. That was the adopted approach, with 268 samples attributed to each feature; **HMC** chains were used, with some burn-in and lag samples discarded for better coverage. When testing, the information of the modes' number and locations were withheld from the samplers, whose particles were initially placed by a normal distribution centered at  $[0, 12.5]$  and with diagonal covariance  $\text{diag}(10, 20)$ .

Figure 32b shows the result of applying **SHMC** in these conditions, and with 768 points. For the **SMC** reweightings, a leave-one-out **KDE** of the particle cloud was used in the denominator instead of the previous target. This idea from [24] was mentioned in section 3.2.2, and aided coverage. The **HMC** parameters were  $L = 20$ ,  $\epsilon = 0.05$ ,  $\mathcal{M} = \hat{\mathbf{I}}_2$ , producing a 99% acceptance rate, and resampling was performed at each step. The sampler succeeded in finding all modes and covering them satisfactorily.

Other methods were tested, namely **SMC** and parallel Markov chains; figure 33 presents a selection of results. All methods used 128 particles for 8 steps, and the same prior (a multivariate gaussian). The particle reduction is significant, making the problem at once more challenging for the sampler and less challenging for the processing unit. Note that **MCMC** typically relies on longer evolutions, but this wouldn't make for an equitable comparison: for the parallel methods, all steps are performed on the final density, so even when matching them for iterations they're already more costly than **SMC**.

The conditions for **HMC** were as before, producing 96% acceptance. **RWM** used  $\sigma_{\text{prop}} = 0.05$ , producing 78% and 66% acceptance within and outside **SMC** respectively. The difference between methods is clear on two fronts: the coupling of **SMC** greatly benefits the particle positioning and mode evenness as compared to non-interacting chains, and the insight of **HMC** seems to be an as strong asset.

## 4.2 Inference problems

### 4.2.1 Phase estimation

This example is based on [115], for the idea of Bayesian (iterative) phase estimation in general as for the implementation of **GRF** and adaptive heuristics in particular. All of this was explored in subsection 2.3.3.

Figure 34 shows the results of applying **MCMC** and **GRF** to learning a phase set at  $\phi_{\text{real}} = 0.5$  using 100 measurements and a flat prior on  $[0, 2\pi]$ . The parameter space is periodic. **GRF** works under a





Figure 33: Trajectories generated using different SMC/MCMC sampling methods to sample from a (bivariate) smiley KDE. Results are shown for 4 methods: parallel RWM chains (33a), SMC with RWM moves (33b), parallel HMC chains (33c), and SHMC (33d).

normality assumption, where the distributions (the successive posteriors) are represented in the period-long interval by two quantities: the arithmetic mean and variance of the particle locations. With this, some distributions over the phase are bound to be misrepresented. For instance, if the distribution is sharp and centered at  $2\pi$ , this description will relay a broad gaussian centered at  $\pi$  due to a bisection at the borders. To account for this, a wrapped normal distribution (on the unit circle) can be considered. A clever alternative fix is suggested in [115]; the authors propose screening for this type of problem using the variance and treating it as a special case using  $\pi$  shifts. That was the option opted for here.



(a) Trajectory generated by MCMC with RWM (black dots) and HMC (red dots) transitions. The HMC proposals prompted essentially all bigger leaps, but were most often unavailing. The experiment controls were chosen semi-arbitrarily in advance, and all the data are used at each step.

(b) Evolution of the distribution means through the iterations of GRF. The experiment controls were chosen adaptively. Each iteration  $i$  uses only the  $i$ th datum and the  $(i + 1)$ th iteration's results.

Figure 34: Evolution of the sampling process through the iterations for MCMC with RWM and HMC move steps (34a) and for GRF (34b). Both consider 100 experiments total, but their design and use are different. The horizontal coordinate is the phase (angle) value. The points are positioned vertically and upward according to the iteration number, e.g. the topmost point corresponds to the final result. The likelihoods based on the full datasets are plotted in blue in a second axis for visualization.

For GRF, the experimental controls were chosen adaptively according to section 2.4.3; the data were

considered one at a time (per iteration), meaning 100 steps. For **MCMC**, all combinations of  $M \in [1..10]$  and  $\theta \in [0..9] \cdot \pi/5$  were used, and 30 transitions were realized. Note that even though this corresponds to fewer iterations, each of them considers all the data, in contrast with **GRF** where a single datum is contemplated per step (along with the previous results). This makes **GRF** more lightweight at the expense of correctness. It also enables adaptivity, bringing more informative data: after roughly 20 measurements, the results are already quite decisive, whereas with so few data **MCMC** typically fails to converge.

The problem of applying **HMC** to this type of problem was discussed in 3.4.1. To manage that, a **RWM** step was performed after the **HMC** one whenever the acceptance probability for the latter fell below 0.01. This resulted in only 3.3% **HMC** steps in spite of **HMC** acceptance's nearing 100%. In other words, all points attempted **HMC**, but for 96.7% this resulted in vanishing probability of moving; a select few had near unit probability of relocation. It is clear that the **HMC** transitions are the driving force behind the bold transitions; all big leaps that the chain carried off was produced by **HMC**. Notwithstanding, this is a hard problem to manage, especially for longer evolutions, and requires lavishing resources on the proposals for rare benefit. Even though it can be controlled by fine-graining the integration further, the effect is limited and resource-intensive.


## 4.2.2 Precession and bimodality

This example is based on [49]. It has been mentioned in subsection 2.3.3.

Figure 34 shows the results of using **SMC-SIR** with **LWF** and **MCMC** to learn a frequency  $\omega_{\text{real}} = 0.5$  using 100 measurements and a flat prior on  $\omega \in ]0, 10]$ . The times were chosen in increments of 0.08, decided with a parameter sweep. Note that what matters is not the units, but only the relative scale between times and frequencies. They're always found paired up in the form  $\omega \cdot t$ .

**MCMC** was implemented in the same conditions as section 4.2.1, resulting in roughly 100% mean **HMC** acceptance but only 6.7% **HMC** steps. Clearly, we face the same problem as before. Both used 100 data; **MCMC** used 30 steps, and **SMC** was implemented as in algorithm 4, with 100 particles and the **LWF** parameter  $a = 0.98$ . Despite using more steps and particles, it actually uses fewer likelihood evaluations than **MCMC**:  $T = 100 \cdot 100$  for the 100 reweightings of 100 particles. Had **MCMC** ran for 100 iterations, it would use as many with **RWM** alone, whereas **HMC** would add to that  $LT$  gradient evaluations. Naturally, performing both at some iterations accumulates their costs, which end up higher even for 30 steps (and would still even if **RWM** were dispensable, especially if considering that the cost of evaluating the likelihood is usually much less significant than that of computing its gradient). Also, **MCMC** by itself is less robust, especially for higher frequencies (or equivalently longer evolution times) and/or broader priors, as well as incompatible with online processing (unlike **SIR**).

As discussed in 3.3.3, Markov chains can be made more robust by coupling. We now switch from **MCMC** per se to Markov transitions within an **SMC** scheme, and compare this with **SMC-LWF** (both with **SIR**). The advantages can be made obvious by extending the domain to encompass negative frequencies. Due to the cosine function's (formula 2.84) symmetry, this gives rise to multimodality, in which case **LWF**'s

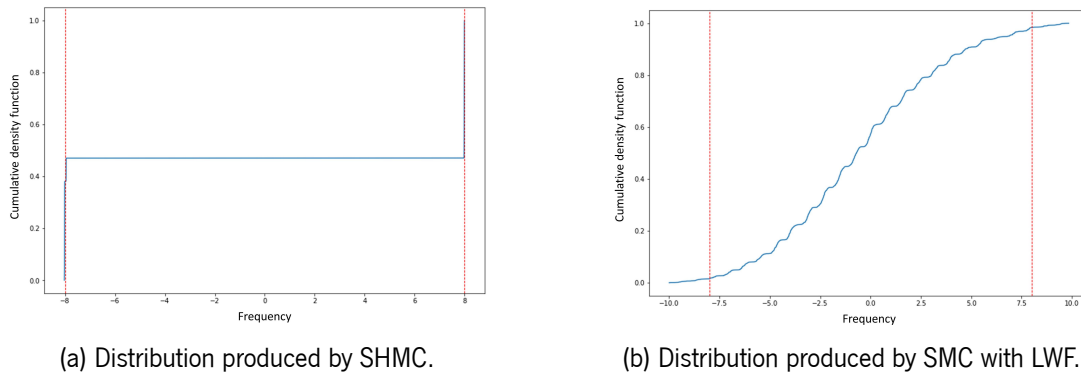


(a) Trajectory generated by MCMC with RWM (black dots) and HMC (red dots) transitions. All data are used at each step. HMC proposals prompted essentially all bigger leaps, but were most often unavailing.

(b) Evolution of the distribution means through the iterations of SIR with LWF. Each iteration  $i$  uses only the  $i$ th datum and the  $(i + 1)$ th iteration's results.

Figure 35: Evolution of the sampling process through the iterations for MCMC with RWM and HMC move steps (35a) and SMC with LWF (35b). Both consider 100 experiments total, and the same set of pre-determined controls. The points are positioned vertically and upward according to the iteration number, e.g. the topmost point is the final result; the horizontal coordinate is the phase (angle) value. The likelihood is plotted in blue on a second axis for visualization.

performance is disastrous (as was to be expected). This was given as a motivation for seeking alternative methods in the reference cited for this section, [49]. The issue can be seen clearly in figure 36.



(a) Distribution produced by SHMC.

(b) Distribution produced by SMC with LWF.

Figure 36: Cumulative density functions obtained via MCMC with RWM and HMC move steps (36a) and SMC with LWF (36b), for a prior with support over negative frequencies. The modes are marked with dashed vertical red lines.

The conditions are the same as before, except for the fact that 2000 particles were used instead of 100 to make for smoother curves and the frequency was fixed at  $\omega = 8$  (or equivalently  $\omega = -8$ ).

Finally, we present results for TLE (section 3.2.3). We test TLE-SHMC, both based on the full data and with energy conserving subsampling and control variates (TLE-SHMC-ECS, section 3.5). In the latter case, the likelihood evaluations were based on a subset of the data for both the reweightings and HMC. The subsampling indices for each iteration and particle were chosen in bPM-MH steps (for 3 blocks, inducing

roughly 67% correlation).

The applied control variates being gradient based (Taylor expansions), they didn't work well due to high target curvature. This can be controlled in several ways. One option would be to keep the evolution times short, but this is inefficient because they bring little information after a point. Another one would be to narrow down the prior. This can be done by a warm up phase on the first (lower times) observations if prior knowledge is insufficient to justify it. Chapter 5 will show the warming up alternative; for these tests, a less general prior was chosen. The posterior distributions are presented in figure 37 as KDEs in the domain of the prior. The real frequency was set at  $\omega = 0.8$ .




Figure 37: Kernel density estimates of the posterior resulting from full data SHMC and from subsampling 50/400 observations (SHMC-ECS).

The 400 measurement times were chosen randomly up to  $t_{\max} = 100$ , and when subsampling 50/400 were evaluated. Apart from that, all conditions were matched for the two methods. The tempering coefficients were 10 and evenly spaced. The resulting standard deviations were  $\sigma_{\text{FD}} = 8.2 \times 10^{-4}$  and  $\sigma_{\text{subs}} = 8.6 \times 10^{-4}$  for the full data and subsampling cases respectively.

### 4.2.3 Multivariate function (2 and 4-dimensional estimation)

This section considers a multi-parameter generalization of the two previous examples, with the purpose of testing methods in more challenging settings. This can be achieved by choosing a vector  $\vec{\omega}$  containing  $dim$  parameters  $\omega_d$  and choosing a still binomial ( $D \in \{0, 1\}$ ) likelihood:

$$L(\vec{\omega} | D) = \frac{1}{dim} \sum_{d=1}^{dim} \cos^2\left(\frac{\omega_d t}{2}\right) \quad (4.2)$$

Clearly, this introduces redundancy, resulting in multimodality. The number of modes reflects in how many ways the frequencies can be distributed by the dimensions, which are symmetric: the parameters all play identical roles. In other words, there are as many modes as there are possible orders for the

parameters. This total is given by the number of permutations (without repetitions) of the dimensions, which is the factorial  $dim!$ .

First, we can use this to illustrate the imperativeness of resampling in a more realistic setting than that of chapter 5. Figure 38 shows how a particle number that converges to the right modes with the help of resampling fails miserably without. An immense increase in particle density can help (figure 38c), but the per-particle returns are minute - and this is for 2 dimensions only. The final particles' proximity to the modes is at best dictated by how close to them the luckiest initial grid points happened to fall, and luck is undependable for high dimensional spaces.

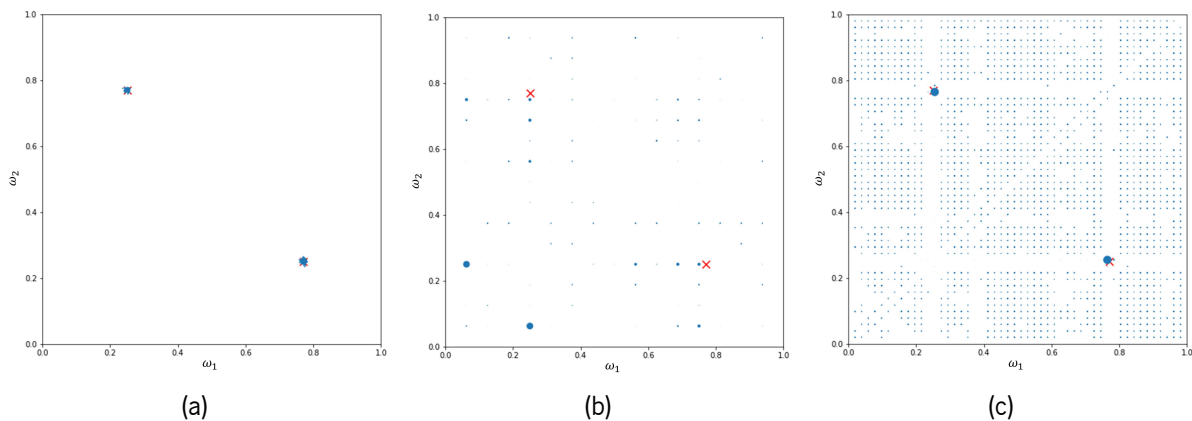


Figure 38: Results of 2-dimensional inference using SIR with RWM propagation for  $15^2$  particles (38a), a grid for  $15^2$  particles (38b), and a grid for  $50^2$  particles (38c). The modes are marked with red 'x' markers.

With this we can also subsample in a more general setting; namely SG-HMC (section 3.5.2), which hasn't been tested yet. The control variates are not applicable in this case, and so were left out. The effect of introducing friction is shown in figure 39. It helps control divergence, but the effect is limited.




Figure 39: Results of using stochastic gradients in HMC within SIR while introducing friction (39b) and not (39a). The modes are marked with red 'x' markers, but are covered by the particles.

We also test multi-dimensional generalizations of SIR-powered adaptivity. As a reference, we use random times on  $t \in [0, 100]$  for a domain normalized to 1 in all dimensions (i.e. a unit hyper-cube).



In subsection 2.4.3, an adaptive heuristic with  $t_k = 1/\sigma_{k-1}$  was suggested. Unfortunately, the standard deviation is a quite less meaningful metric under multimodality. To preserve the intuition, one possibility is to cluster the particles and consider an average of the standard deviations concerning each cluster. However, clustering is a difficult task, especially in a (realistic) setting where the number of modes is not known. A more lightweight option is to consider space occupation, another measure of uncertainty. This was the selected route; occupation was computed by partitioning the space into (adaptively many) cubic cells. Finally, a third option was considered where the times are chosen offline but increase through time. Exponentially increasing times didn't do very well in this context. In short, three methods were considered:

- Random:  $t^{(\max)} = 100$
- Adaptive:  $t_k \propto (\text{occupation\_rate}_{k-1} \cdot \widehat{\text{ESS}}_{(k-1)} / n_{\text{particles}})^{-1}$
- Increasing random:  $t_k^{(\max)} = C_1 \cdot (\text{floor}(k/C_2) + 1)$

The ratio of the effective sample size to the number of particles (its upper bound) is necessary to increase the times in the iterations where resampling doesn't take place. It could also be smoothed or escalated by some power. For the last case,  $C_1 = 100$ ,  $C_2 = 20$  were chosen (with a parameter sweep). Results for selected runs are shown in figure 40.




Figure 40: Achieved results using random (40a), adaptive (40b) and random but increasing (40c) measurement times. The modes are marked with red 'x' markers.

The inter-run results are more variable, and taking medians shows a different picture. Figure 41 shows the final particle clouds for the successful runs with median final standard deviation.

For defining *success*, assessing performance, and averaging standard deviations, the mode locations were used. Success was considered as the conjunction of accuracy, precision, correctness, and mode coverage. For evaluating these points, thresholds were considered for the average distance to a mode, the variance around each mode, the difference between real and estimated error, and the discrepancy between the number of particles attributed to each mode. Table 1 shows the quantitative results, including success rate. The considered metric was the weighted average of standard deviations associated with each mode, each particle being assigned to its closest one. This is what is meant by *standard deviation*.




Figure 41: Median successful results using random (40a), adaptive (40b) and random but increasing (40c) measurement times. The modes are marked with red 'x' markers.

|                    | Median std. deviation | Mean std. deviation  | Success rate |
|--------------------|-----------------------|----------------------|--------------|
| <b>Random</b>      | $2.0 \times 10^{-4}$  | $2.1 \times 10^{-3}$ | 65%          |
| <b>Adaptive</b>    | $7.3 \times 10^{-5}$  | $1.6 \times 10^{-3}$ | 85%          |
| <b>Inc. random</b> | $1.8 \times 10^{-4}$  | $1.7 \times 10^{-3}$ | 53%          |

Table 1: Results obtained for 2-dimensional estimation, using different choices of measurement times.

The adaptive method shows significant improvement in the median standard deviation, but this advantage shrinks if the mean is considered instead. This is a clear sign of instability in the inter-run performance. Another disadvantage is that the complexity of tuning and cost of computing occupation increase for higher dimensions, due to a more complicated interplay between volume and probability density.

Finally, figure 42 shows the results of 4-dimensional estimation (24 modes) using the the SMC of algorithm 4 and TLE-SMC, both with nearly 100% HMC moves. Both used  $12^4$  particles, 250 data, and random times on  $]0, 100]$ . For TLE, 5 coefficients were used. The added robustness of TLE can be seen clearly, as can the increased complexity coming from the extra dimensions.

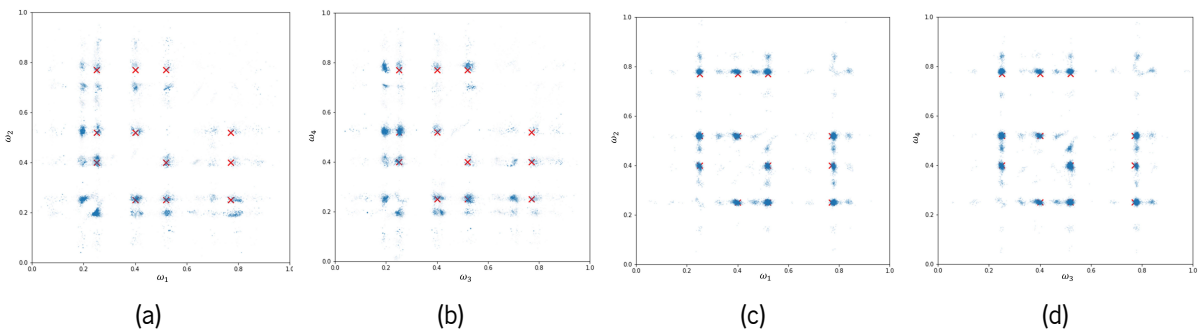


Figure 42: Results of 4-parameter estimation (with dimensions plotted pairwise) for two variations of SMC: SIR with the prior as importance function (42a and 42b), and TLE (42c and 42d). The modes are marked with red 'x' markers.

## Experiments on quantum hardware

This chapter presents a few quantum characterization experiments. Section 5.1 contains a short formal description of their theory. Section 5.2 shows results obtained by applying the parameter estimation techniques described in the previous chapters - namely Bayesian inference, adaptive experimental design, SMC, MCMC, and subsampling strategies - to the open system dynamics of quantum computers.

### 5.1 Models for open quantum systems

This section will introduce the elementary theory behind the experiments whose results are to be presented in the ensuing sections. It is mostly based on [71, 94, 95, 7, 90, 62, 63, 82] (by order of relevance). Additional sources will be cited throughout when applicable, namely for subsection 5.1.5.

Subsections 5.1.1 and 5.1.2 present methods for the visualization of mixed states and formal treatment of open quantum systems respectively. The former point is useful for the latter, as will become clear in sections 5.1.3 and 5.1.4 when analysing two paradigmatic noise models for open quantum systems.

At last, section 5.1.5 elaborates upon the precession dynamics of subsection 2.3.3 to create applied examples which are at once more realizable than the original example (given the available quantum machines) and more interesting than simple open quantum system effects. When observed in quantum devices, these phenomena too are inevitably joined by environment-induced effects (such as the ones to be explored in the two preceding subsections 5.1.3 and 5.1.4), bringing added complexity to characterization problems.

#### 5.1.1 Inside the Bloch sphere: visualizing mixed states

While *prototypical* quantum states can be written as state vectors, which consist of a linear combination of basis states, this often fails to present a realistic description. It assumes knowledge to be limited only by fundamental quantum uncertainty, and otherwise perfect. Such a representation fails to deliver when dealing with real, noisy quantum devices. In particular, these devices are ultimately *open* quantum systems, whereas this restrictive view treats but closed ones. As eloquently put in [93], quantum phenomena happen not in a Hilbert space, but rather in a lab.

This prompts an extension of the formalism to encompass *mixed* states. These are comprised of an ensemble of pure quantum states  $|\psi_i\rangle$ , each of which occurs with some probability  $p_i$ . A convenient description is brought by density operators  $\rho$ , which generalize the state vector description to all states,

pure or not:

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i| \quad (5.1)$$

, where each  $|\psi_i\rangle$  is a state vector corresponding to a pure state.

Properties can then be obtained by taking statistics over the weighted sum. These are *classical* statistics, an unwelcome feature because interference is limited. The amplitudes corresponding to different  $|\psi_i\rangle$  don't interact with one another as they do within superposition states; instead, the ensemble behaves as a mixture of isolated states. In other words, each term in the sum is closed off, and quantum effects can be observed only within its own scope. Mixed states describe uncertainty, but not in the quantum sense; they represent flawed knowledge, these flaws extending further than the fundamental limits imposed by quantum nature.

Single qubit pure states find a nice depiction on the Bloch sphere. For this representation, we write the state vector in the form:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\varphi}\sin\left(\frac{\theta}{2}\right)|1\rangle \quad (5.2)$$

, for  $\{\theta, \varphi\} \in [0, \pi] \times [0, 2\pi[$ .

Unlike (single qubit) pure states, mixed states can't be represented on the Bloch sphere; they can however be represented *in it*. This is clear by vector algebra: the vectors corresponding to a mixed state can be obtained as a convex combination of pure state vectors, a combination which will have at most unit norm (this edge case occurring precisely in the single  $i$  case, i.e. for a pure state).

A handy generalization is given by writing:

$$\rho = \frac{1}{2}(\hat{\mathbf{I}} + \vec{v} \cdot \vec{\sigma}) = \begin{bmatrix} 1 + v_z & v_x - iv_y \\ v_x + iv_y & 1 - v_z \end{bmatrix} \quad (5.3)$$


, for  $\vec{\sigma} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)$ ,  $\vec{v} = (v_x, v_y, v_z)$ , and  $v_x, v_y, v_z \in [0, 1]$ .

This switch from two angles (fixed-radius polar coordinates) to three lengths (cartesian coordinates) provides a convenient way to picture general states. Both representations are depicted in figure 43.

It follows easily from 5.3 that each element  $v_i$  of the vector  $\vec{v}$  corresponds to the expectation value of the corresponding Pauli operator  $\sigma_i$ , given by  $\langle\sigma_i\rangle = \text{Tr}[\sigma_i\rho]$ . This is the reason for the convenience. As an example, the Bloch vectors pointing to the  $x/y/z$  positive and negative poles correspond to the positive and negative eigenvalue eigenstates of  $\hat{\sigma}_x/\hat{\sigma}_y/\hat{\sigma}_z$ , with e.g.  $\vec{v} = (\pm 1, 0, 0)$  representing the  $|\pm\rangle$  state. These intuitive implications are just consequences of the fact that the projection onto each axis  $\hat{\mathbf{i}}$  represents the expected value of its associated observable  $\hat{\sigma}_i$ .

For any valid density matrix, the Bloch vector length satisfies  $\|\vec{v}\| \leq 1$ . It can be shown that the purity is given as a function of this norm:

$$\text{Tr}[\rho^2] = \frac{1}{2}(1 + \|\vec{v}\|^2) \quad (5.4)$$



(a) Bloch sphere representation of a pure state. The Bloch coordinates are the angles  $(\theta, \varphi)$ .

(b) Bloch ball representation of an arbitrary state, pure or mixed. The Bloch coordinates are the 3-dimensional cartesian coordinates  $\vec{v} = (v_x, v_y, v_z)$ , which correspond to the vector of expectations of the Pauli operators.

Figure 43: Bloch representations for pure and general states.

That is, the center of the sphere corresponds to the maximally mixed state, and unit length vectors correspond to pure states. Other mixed states will lie somewhere in between, inside the Bloch ball.

Ideally, all computations would unfold on the surface of the Bloch sphere, but environment effects and experimental error pull them inside - signifying information loss. In general, error channels result in deformations of the Bloch sphere, namely compressions. These error channels often depict unwanted interactions with the environment, which may result in energy and coherence loss. The following section introduces fundamentals concerning the symbolic treatment of open quantum systems, whereas the two subsequent ones present representative cases that illustrate artifacts commonly found in real-world practice due to imperfect isolation.

### 5.1.2 The Kraus operator sum representation

The time evolution of a quantum wavefunction is shaped by the Schrödinger equation, one of the most emblematic cornerstones of quantum mechanics. Often, we are forced to make room for incomplete information by generalizing it to density operators, which can describe mixed states. But even still, the usual form still has a propensity for being overwhelmingly broad - in that case, we must make room for incomplete *observations*.

Schrödinger's equation describes the unitary time evolution of a system; this is *any* system's time evolution; but we're not exactly talking about systems that fit inside a glass vial. More specifically, we

speak of closed quantum systems, whose time evolution is given by:

$$\rho(t) = \hat{U}\rho(0)\hat{U}^\dagger \quad (5.5)$$

, with time-dependent  $\hat{U}$ .

Unless the system under study is perfectly isolated, it will enjoy unitary evolution *jointly with its surroundings*. If we can't observe these surroundings as well, unitarity is no longer assured, and equation 5.5 is not to be applied directly.

We can think of any subject of study, such as a qubit, as a system coupled to an environment, or bath - such as the lab, or the entire outside world. Naturally, external effects vary in degree, and some may be negligible. Nonetheless, rarely does it ever occur that they are *all* negligible, which means that realistic analyses must in general contemplate them. The result of the coupling is that our subject, strictly speaking, ceases to function as a system on its own right; in particular, information can be *leaked* to the external environment.

To partial observations on a system corresponds partial information. As such, it is not very surprising that focusing on a subsystem may cause it to behave like a mixed state. To enable the discussion of some simple instances where this happens, we will now go over the basic formalism required for describing them.

We suppose we have a system of interest  $s$ , which we will just call system (despite the fact that it is just a subsystem of the bigger bipartite system that includes the environment too), living on some Hilbert space  $\mathcal{H}_s$ . The environment  $e$  has a space of its own,  $\mathcal{H}_e$ , and together they exist in  $\mathcal{H}_s \otimes \mathcal{H}_e$ .

The *partial trace* over the environment is a transformation  $(\mathcal{H}_s \otimes \mathcal{H}_e) \mapsto \mathcal{H}_s$  mapping density operators describing the composite system to reduced density operators describing only system  $s$ . It is obtained by sandwiching the composite operator between the orthonormal vectors belonging to the environment's eigenbasis:

$$\rho_s \equiv \text{Tr}_e [\rho^2] = \sum_\alpha \langle \alpha | \rho | \alpha \rangle_e \quad (5.6)$$

The  $\{ | \alpha \rangle \}$  span  $\mathcal{H}_e$  and act on it alone. There are as many as the environment's Hilbert space dimension.

We assume a separable initial state, and further that the environment is initially in a pure state  $| \mu_0 \rangle$ <sup>1</sup>. From the joint unitary evolution governed by  $U(t)$ , we can find the map acting exclusively on  $s$  as:

$$\rho_s(t) = \sum_\alpha \langle \alpha | \hat{U}(\rho_s(0) \otimes | \mu_0 \rangle \langle \mu_0 |) \hat{U}^\dagger | \alpha \rangle = \sum_\alpha \hat{K}_\alpha \rho_s(0) \hat{K}_\alpha^\dagger \quad (5.7)$$

Where we have defined:

$$\hat{K}_\alpha \equiv \langle \alpha | \hat{U} | \mu_0 \rangle \quad (5.8)$$

In the context of quantum information science, quantum maps are often called channels. Time evolution is itself a map, so a channel can represent the evolution  $\rho_s(0) \mapsto \rho_s(t)$ . In this context, equation

<sup>1</sup>This is only for the sake of simplicity, and the generalization is straightforward [71].

5.7 is often written alternatively as:

$$\mathcal{E}[\rho_s] = \sum_{\alpha} \hat{K}_{\alpha} \rho_s \hat{K}_{\alpha}^{\dagger} \quad (5.9)$$

The sum form of equations 5.7 and 5.9 is generally termed an **operator-sum representation (OSR)**, and the  $\hat{K}_{\alpha}$  are called Kraus operators. They may be time-dependent, just like  $\hat{U}$ .

These operators must still respect a normalization condition, which guarantees that they preserve any density operators' own normalization (unit trace) - that is, that they map valid states to still valid states. The appropriate condition can be shown to be:

$$\sum_{\alpha} \hat{K}_{\alpha}^{\dagger} \hat{K}_{\alpha} = I \quad (5.10)$$

If there exists a single Kraus operator, this normalization condition imposes that it be unitary, in which case we recover Schrödinger dynamics as a special case. Otherwise, the system is not closed, but thanks to the partial trace over the environment we single out its evolution. That is, tracing out the environment leaves the Kraus operators to act on it alone, describing its specific dynamics. Of course, this doesn't mean that we isolate the system, but rather that we factor in the uncertainty contributed by the fact that it is an *open* system.

Considering the unitary evolution of the composite system (subsystem of interest + environment), a transformation  $\hat{U}$  given by  $|\psi_0\rangle \mapsto |\psi_t\rangle$  can be seen as  $|\psi_t\rangle\langle\psi_0|$ , meaning that it *eats* a state  $|\psi_0\rangle$  and *produces* a state  $|\psi_t\rangle$ . This simplifies the interpretation of 5.8. Tracing out a constant pure initial state  $|\mu_0\rangle$  removes the initial environment state from the side to the left of  $\mapsto$ , leaving still its basis states to be traced out on the right-hand side. Then, the Kraus operators can be found by ignoring the fixed environment initial state and applying to the right-hand-side [95]:

$$\hat{K}_{\alpha} = {}_e\langle\alpha|\hat{U} \quad (5.11)$$

All of these ideas will become clearer with the ensuing basic examples. In the following section, referring to the system is assumed to be the default, so the  $s$  subscript will be sometimes dropped.

### 5.1.3 Energy relaxation: the amplitude damping channel

Energy loss is a problem that plagues many applications across the fields of physics, and quantum computation is no exception. It is unavoidable that a qubit which is not in its ground state will with some probability release energy to its environment, signifying amplitude damping - and thus a loss of information.

Here we will analyse the effect energy relaxation exerts in a single qubit, which can be visualized in the Bloch ball. We consider a perfect two-level system and assume the energy flow to be unidirectional; that is, the environment does not supply energy to the system. Under these conditions, the possible transitions for the qubit are from the excited state  $|1\rangle$  to the ground state  $|0\rangle$ . We suppose they occur with a time-independent probability  $p_1$  per unit time. The index is chosen in anticipation of the fact that a second channel will be considered later.

The effect of the channel is to map:

$$|1\rangle \mapsto |0\rangle, \text{ with probability } p_1 \quad (5.12)$$

The state  $|0\rangle$  is left untouched with probability 1, whereas  $|1\rangle$  is preserved only with probability  $1 - p_1$ .

Under these conditions, one of the Kraus operators will be a lowering operator:

$$\hat{K}_1 = \sqrt{p_1} |0\rangle\langle 1| = \begin{bmatrix} 0 & \sqrt{p_1} \\ 0 & 0 \end{bmatrix} \quad (5.13)$$

Another Kraus operator should map the ground state component to itself, i.e. it should represent the component  $|0\rangle\langle 0|$ . Along with the normalization condition, this produces:

$$\hat{K}_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-p_1} \end{bmatrix} = |0\rangle\langle 0| + \sqrt{1-p_1} |1\rangle\langle 1| \quad (5.14)$$

The meaning of the lower-right element may become clearer by noting the extreme cases. For  $p_1 = 0$ , this operator is just the identity, and the only Kraus operator. There is no energy loss, and unitary evolution is recovered. If  $p_1 = 1$ , then it is  $|0\rangle\langle 0|$ . The result is that  $|1\rangle$  is always mapped to  $|0\rangle$ , because there is unit probability of decay. In the intermediate cases, the density operator element  $\rho_{11}$  is reduced by the multiplying factor  $\sqrt{1-p_1}$ . The higher the emission probability, the larger this decrease: this compensates for the excited state population removed (brought to the ground state) by  $\hat{K}_1$ .

The normalization condition may seem to remove intuition from the process, but it can be recovered by a different view. Assuming the system starts out in a pure state, there are two possible branches for the unknown system evolution  $\psi(0) \mapsto \psi(dt)$ :

$$|0\rangle_s \otimes |0\rangle_e \mapsto |0\rangle_s \otimes |0\rangle_e \quad (5.15)$$

$$|1\rangle_s \otimes |0\rangle_e \mapsto \sqrt{1-p_1} |1\rangle_s \otimes |0\rangle_e + \sqrt{p_1} |0\rangle_s \otimes |1\rangle_e \quad (5.16)$$

, where we assume two possible environment states: having absorbed energy emitted by the system and not, respectively  $|1\rangle_e$  and  $|0\rangle_e$ .

The operator  $\hat{K}_1$  is obtained by tracing over the environment state  $|1\rangle_e$  - it leaves to the system a single possible evolution path, from  $|1\rangle_s$  to  $|0\rangle_s$ . The operator  $\hat{K}_0$  comes from tracing over  $|0\rangle_e$ . In this case, two paths are possible for the system: either it was always in state  $|0\rangle_s$ , or it was and still is in state  $|1\rangle_s$ . However, not having observed the environment in state  $|1\rangle_e$  reduces the likelihood that the system was originally in state  $|1\rangle_s$ , since this fact would have yielded non-null probability of observing a different outcome whereas  $|0\rangle_s$  would not.

Now we would like a more convenient expression - one that would give us the time evolution at once. The first step is getting the explicit expression for the density matrix the channel maps to. Let us write the



system's state at any time  $t$  as:

$$\rho(t) = \begin{bmatrix} \rho_{00}(t) & \rho_{01}(t) \\ \rho_{10}(t) & \rho_{11}(t) \end{bmatrix} \quad (5.17)$$

We then apply the OSR to an initial state  $\rho(0)$  to get:

$$\rho(dt) = \hat{\mathbf{K}}_0 \rho_s(0) \hat{\mathbf{K}}_0^\dagger + \hat{\mathbf{K}}_1 \rho_s(0) \hat{\mathbf{K}}_1^\dagger = \begin{bmatrix} \rho_{00}(0) + \rho_{11}(0) \cdot p_1 & \rho_{01}(0) \cdot (1 - p_1)^{1/2} \\ \rho_{10}(0) \cdot (1 - p_1)^{1/2} & \rho_{11}(0) \cdot (1 - p_1) \end{bmatrix} \quad (5.18)$$

This is an interesting result: energy loss doesn't only reduce the excited state population, it also damps the off-diagonal elements - effectively washing off quantum coherence (though more slowly than it reduces  $\rho_{11}$ ). Clear off-diagonals in a density matrix mean an absence of quantum effects, or quantum coherence.

To see how the probability of finding the system in state  $|1\rangle_s$  changes with time, we can consider a relaxation rate  $\gamma_1 = p_1/dt$  and the map to be applied  $n$  times until a time  $t = ndt$ . Then,  $p_1 = \gamma_1 dt = \gamma_1 t/n$ , and:

$$\rho_{11}(t) = \rho_{11}(0) \left(1 - \frac{\gamma_1 t}{n}\right)^n \quad (5.19)$$

Taking the continuous time limit  $n \rightarrow \infty$  and recognizing  $\lim_{n \rightarrow \infty} (1 + x/n)^n = e^x$  as the Euler number limit:

$$\rho_{11}(t) = e^{-\gamma_1 t} \rho_{11}(0) = e^{-t/T_1} \rho_{11}(0) \quad (5.20)$$

, where  $T_1 \equiv 1/\gamma_1$  is known as the characteristic energy relaxation time.

As for the other elements of the density operator, the off-diagonal terms can readily be seen to decay at a rate  $\gamma_1/2$ , or equivalently to enjoy a characteristic time twice as large. Finally, it follows easily from the trace-preservation property of the Kraus operators that the ground state population  $\rho_{00}$  increases by the summand  $\rho_{11}(t) = (1 - e^{-\gamma_1 t})\rho_{11}(0)$ .

This can be visualized as a compression of the Bloch sphere, where all points move closer to the north pole except for the north pole itself (as long as  $p_1 > 0$ ). Its shape is compressed along all directions, but twice as strongly in the  $\hat{z}$  direction, and  $|0\rangle$  suffers no displacement. The magnitudes of both the shift and the compression are proportional to  $\gamma_1$ . Details on the geometric interpretation can be found in [71].


The key thing to note is that because we don't actually measure the environment, we have only a probabilistic understanding of what transpired, which decreases the purity of the qubit's state.

The characteristic time  $T_1$  can be estimated using the circuit of figure 44 for several waiting times  $\Delta t$  since the initialization at  $|1\rangle$ . The pi pulse realizes a rotation of  $\pi$  radians in the Bloch sphere, taking the  $|0\rangle$  state to the  $|1\rangle$  state. It can be a  $\hat{\sigma}_x$  gate, in which case the rotation is around the  $x$  axis.

After the qubit is brought to the excited state  $|1\rangle$  (at the time reference  $\Delta t = 0$ ), the probability of measuring  $|1\rangle$  decays as:

$$\mathbf{P}(1 | \Delta t) = e^{-\Delta t/T_2} \quad (5.21)$$

This model can be used to perform inference, or to fit a curve by regression.

Figure 44: Pulse schedule for measuring  $T_1$ .

### 5.1.4 Dephasing: the phase damping channel

Another common source of noise is *phase damping*. This map is characterized by a phase flip which occurs with some probability  $p_\varphi$ :

$$\rho(dt) = (1 - p) \cdot \rho(0) + p \cdot \hat{\sigma}_z \rho(0) \hat{\sigma}_z \quad (5.22)$$

Clearly, a possible set of Kraus operators is now:

$$\hat{K}_0 = \sqrt{1 - p_\varphi} \cdot \hat{I} \quad (5.23)$$

$$\hat{K}_1 = \sqrt{p_\varphi} \cdot \hat{\sigma}_z \quad (5.24)$$

The second operator flips the sign of the off-diagonals, which slowly erases them since they partially cancel out their own original versions (left unaltered with probability  $1 - p$ ). It is easy to see that acting on the Bloch representation of pure states with  $\hat{\sigma}_z$  flips the sign of the Bloch coordinate  $\varphi$ , the phase, hence the term *dephasing*. For general states, it flips the  $x$  and  $y$  Bloch vector components. The evolution of the density operator is given by:

$$\rho(dt) = \hat{K}_0 \rho_s(0) \hat{K}_0^\dagger + \hat{K}_1 \rho_s(0) \hat{K}_1^\dagger = \begin{bmatrix} \rho_{00}(0) & \rho_{01}(0) \cdot (1 - 2p_\varphi) \\ \rho_{10}(0) \cdot (1 - 2p_\varphi) & \rho_{11}(0) \end{bmatrix} \quad (5.25)$$

Taking the continuous time limit as before, now for  $p_\varphi = \gamma_\varphi t/n$ , yields:

$$\rho(t) = \begin{bmatrix} \rho_{00}(0) & e^{-t/T_\varphi} \rho_{01}(0) \\ e^{-t/T_\varphi} \rho_{10}(0) & \rho_{11}(0) \end{bmatrix} \quad (5.26)$$

, where  $T_\varphi \equiv 1/(2\gamma_\varphi)$  is defined as the characteristic time for this noise channel.

The effect on the off-diagonal terms is similar to what we observed in the previous section, meaning that the  $x$  and  $y$  Bloch components shrink, but now the diagonal is left intact. In other words, this channel has no effect on the  $z$  eigenbasis.

The Bloch vector can easily be found to evolve as:

$$\vec{v}(t) = (e^{-t/T_\varphi}, e^{-t/T_\varphi}, 1) \cdot \vec{v}(0) \quad (5.27)$$

That is, it recedes in the horizontal plane towards the vertical ( $z$ ) axis, as its  $x$  and  $y$  components suffer from damping until they reach 0 (at which point the state is maximally mixed along those directions, the pairs of eigenstates being equally probable so as to yield null expectation).

If both dephasing and energy loss are present, they combine to produce an effective coherence time  $T_2^*$  given by:

$$\frac{1}{T_2^*} = \frac{1}{2T_1} + \frac{1}{T_\varphi} \quad (5.28)$$

This coherence time  $T_2^*$  is known as  $T_2$  star to distinguish it from another  $T_2$  that will be discussed shortly. It can be estimated using the circuit 45 for several waiting times  $\Delta t$  between half-pi pulses.




Figure 45: Pulse schedule for measuring  $T_2^*$ .

The purpose of the half-pi pulses is to take the Bloch vector to the equatorial plane, where the  $T_2^*$  is most relevant, and then back for the computational basis measurement. In the  $x$  basis, the function is simply:

$$\mathbf{P}(+ | \Delta t) = \frac{1}{2} + \frac{e^{-\Delta t/T_2^*}}{2} \quad (5.29)$$

A Hadamard gate directly translates this to  $\mathbf{P}(0 | \Delta t)$ , but any rotation that takes the Bloch vector to the  $x$ - $y$  plane will do, such as rotations of  $\pi/2$  radians around the  $x$  or  $y$  axis (though then if the angle is the same for both pulses the probability  $\mathbf{P}(0 | \Delta t)$  is complementary to 5.29).

The sources of this decay can be many, but some of them are actually reversible, unlike the energy loss of the previous section. In particular, a common reason for this effect is field inhomogeneities that lead to low frequency noise. This can occur for an ensemble of physical qubits due to local field variations that change their individual Larmor frequencies, or similarly if there are small temporal fluctuations in the qubit's resonance frequency.

Looking at figure 45, this will cause the  $\pi/2$  pulses to be slightly off-resonance by a varying amount, causing a series of *accidental* Ramsey experiments as will be described in the following section. We can think of a set of Bloch vectors that start out at state  $|+\rangle$ , but each of which rotates around the  $z$  axis at a slightly different frequency. If we measure such an ensemble, the final states will be inconsistent across several measurements, signifying a loss of purity.

Each such vector is pulsed into the equatorial plane state and slowly precesses within it at some individual frequency  $\delta$ , instead of staying in place as would be desirable. After a time  $\Delta t$ , it will have covered an angular distance of  $\delta \cdot \Delta t$ . But if halfway through the evolution we apply a  $\pi$  pulse, its position will be mirrored relative to the vertical plane, so that it is  $\delta \cdot \Delta t$  away from  $|-\rangle$ .

Allowing it to evolve for the other half of the free evolution time will cancel out its former motion, leaving it deterministically at  $|-\rangle$  (assuming instantaneous pulses and that  $\delta$  is constant). If we have a temporal or spatial ensemble of spins, all of their shifts will be erased by this inversion pulse, which effectively reunites them - so that they can jointly describe a pure state.

This is generally called a spin echo or Hahn echo experiment [55], the term *echo* alluding to the way the noise retraces its steps to repeat its own pattern until it cancels itself out for a measurement at  $\Delta t$ . The result is a longer characteristic time, called  $T_2$ : by refocusing the spins, it eliminates a significant source of noise, improving  $T_2^*$ . It essentially works as a high-pass filter.  $T_2$  is often called the dephasing time constant. The circuit for measuring this time is shown on figure 46.




Figure 46: Pulse schedule for measuring  $T_2$ . This is known as a *Hahn echo* or *spin echo* sequence.

To prolong the coherence time still further, higher order methods can be used. Here we will use the simple circuit from 46, the purpose being mostly to make the coherence time more stable. The crude  $T_2^*$  is typically quite short, and suffers from strong fluctuations. This interferes with the comparison of results across multiple methods, which can hardly be expected to be consistent.

With this we conclude the open system effects to be characterized in section 5.2. Despite being very simple, these two examples actually capture the essence of two important types of noise present in real quantum devices. In practice, they can stem from different and complex physical mechanisms - but these straightforward descriptions present a simple view of their net result, providing useful and widely used noise models.

### 5.1.5 Spin precession

We will now consider unitary dynamics that combine with the open system effects of subsections 5.1.3 and 5.1.4 to produce likelihood functions whose characterization is more challenging than that of simple exponential models. They take the form of sinusoids. The main references are [119, 52, 41, 112, 25].

In section 2.3.3, we considered the evolution induced by the Hamiltonian:

$$\hat{H}_0 = \frac{\Omega_0}{2} \hat{\sigma}_z = \frac{1}{2} \begin{bmatrix} \Omega_0 & 0 \\ 0 & -\Omega_0 \end{bmatrix} \quad (5.30)$$

Here we change the  $\omega$  of before into  $\Omega_0$  for reasons that will shortly become clear. This Hamiltonian describes **Larmor precession**, and it can be realized by applying a magnetic field in the  $\hat{z}$  direction [119]. We saw that the resulting wavefunction evolution was given for some initial Bloch coordinates  $(\theta_0, \varphi_0)$  by:

$$|\psi(t)\rangle = \cos(\theta_0/2) |0\rangle + e^{i\varphi_0 + \Omega_0 t} \sin(\theta_0/2) |1\rangle \quad (5.31)$$

The angular frequency  $\Omega_0$ , called the Larmor frequency, is proportional to the magnitude of the applied field,  $\Omega_0 = \gamma B$  (using Planck units so that  $\hbar = 1$ ).

Similar phenomena can be observed by building on the Hamiltonian of 5.30, some of which are useful for control or characterization of two level systems. Considering a qubit subject to 5.30, we can add a  $\hat{\sigma}_x$  contribution which rotates around the  $z$  axis:

$$\hat{H} = \hat{H}_0 + \hat{H}_d = \frac{\Omega_0}{2}\hat{\sigma}_z + \frac{\Omega_d}{2}(e^{i\omega t}\hat{\sigma}_+ + e^{-i\omega t}\hat{\sigma}_-) = \begin{bmatrix} \Omega_0 & \Omega_d e^{-i\omega t} \\ \Omega_d e^{i\omega t} & -\Omega_0 \end{bmatrix} \quad (5.32)$$

, where  $\hat{\sigma}_+ = |1\rangle\langle 0|$  and  $\hat{\sigma}_- = |0\rangle\langle 1|$  are the raising and lowering operators for the  $z$  basis. This addition corresponds to a periodic driving field that rotates around the  $z$  axis (i.e. perpendicularly to the polarization).

Applying the time-dependent Schrödinger equation to a state vector in the  $z$  basis, written in the form  $|\psi(t)\rangle = a|0\rangle + b|1\rangle$ , produces a system of two equations - one for each amplitude's derivative. The explicit time dependence on  $e^{i\omega t}$  can be removed by transforming to a rotating frame of reference (e.g. by applying the unitary transformation  $\hat{U} = e^{-i\omega t/2}|0\rangle\langle 0| + e^{i\omega t/2}|1\rangle\langle 1|$ ), which yields a simplified pair of equations. These correspond to an *effective* Hamiltonian that is time-independent:

$$\hat{H}_{eff} = \frac{1}{2} \begin{bmatrix} \delta & \Omega_d \\ \Omega_d & -\delta \end{bmatrix}, \quad \delta \equiv \Omega_0 - \omega \quad (5.33)$$

This Hamiltonian describes the same dynamics as the previous one, but now in a rotating referential. Comparing with 5.32, we can see that in the rotating frame the driving field appears at rest ( $\omega = 0$ ), and the Larmor frequency has been replaced by the difference  $\delta$  between itself and the normal field's rotation frequency, this difference being called the *detuning*.

Away from resonance,  $|\delta| \gg 0$ , the detuning  $\delta$  induces dynamics similar to those of 5.30. However, the driving field leads to coherent switching between basis states, and in resonance ( $\omega \approx \Omega_0 \leftrightarrow \delta \approx 0$ ) the most relevant phenomenon are the  $|0\rangle \leftrightarrow |1\rangle$  transitions.

The time evolution of the state vector can most easily be described by diagonalizing  $\hat{H}_{eff}$ , i.e. solving the characteristic polynomial in the rotating frame and then the eigenequation (which is just the time-independent Schrödinger equation) for each eigenvalue  $E_{\pm}$ . Working in the energy eigenbasis of  $\hat{H}_{eff}$  and assuming a  $|0\rangle$  initial state, we get:

$$|\psi(t)\rangle = \cos(\theta/2)e^{-itE_+}|E_+\rangle + \sin(\theta/2)e^{-itE_-}|E_-\rangle \quad (5.34)$$

These energy eigenstates  $|E_{\pm}\rangle$ , around which the Bloch vector precesses, belong to the  $x$ - $z$  plane, making an angle  $\alpha = \tan^{-1}(\Omega_d/\delta)$  with the  $z$  axis. This amounts to a rotation of the quantization axis. That is, the basis of the Bloch sphere is altered according to the ratio between the driving amplitude  $\Omega_d$  and the detuning  $\delta$ . The larger the first one is compared to the second, the further the axes move away from the  $z$  axis, and the closest to the  $x$  axis.

If there is no driving ( $\Omega_d = 0, \omega = 0$ ), then  $\alpha = 0$  and the eigenvectors overlap with those of the  $z$  basis. In that case, the Bloch vector precesses around the  $\hat{z}$  direction and the frequency is  $\Omega_0$ ; we recover

the original Larmor oscillations. For no detuning (in resonance,  $\delta = 0$ ), they instead coincide with those of the  $x$  basis, and the frequency is  $\Omega_d$ . All of this can be seen intuitively from the Hamiltonian 5.33.

From 5.34, we can get the probability of state  $|1\rangle$  at a time  $t$ :

$$\mathbf{P}(1 | t) = \frac{\Omega_d^2}{\Omega_d^2 + \delta^2} \sin^2 \left( \frac{\sqrt{\Omega_d^2 + \delta^2}}{2} \cdot t \right) \quad (5.35)$$

This formula describes what is called **Rabi flopping**, the cyclic evolution of the amplitudes of the two energy levels of a qubit under an oscillatory driving field [52]. For finite detuning, the oscillation frequency is increased and the amplitude is reduced.

This can be used to implement quantum gates, since equation 5.35 can be solved for the time given a desired Bloch sphere position. One can initialize the qubit at  $|0\rangle$  and apply resonant microwave pulses for some duration  $\Omega_d \cdot t = 2\theta_X$  to move the Bloch vector by an angle  $\theta_X$  (this  $\theta$  being the Bloch sphere angle with the  $z$  axis, and the same of figure 43a and equation 5.2). The direction too can be chosen according to the intended results.

Other than control, a possible application for this type of setup is characterization. Equation 5.35 can be used to determine the qubit frequency  $\Omega_0$  to some precision, by applying transverse radiation of known parameters and observing the system's behaviour. For this purpose, better results can be achieved by using another method, called Ramsey's method [112, 41]. Norman Ramsey was originally a student of the physicist responsible for the original technique, Isidor Rabi, and improved upon it.

The idea is to replace the application of a field for some evolution time  $\Delta t$  by two shorter duration pulses separated in space/time. Shortening the interaction length is positive, because it reduces the degree to which the oscillations are affected by experimental error (e.g. field inhomogeneities).

The experimental scheme is based on two  $\pi/2$  pulses of frequency  $\omega$  detuned by  $\delta$  with respect to an estimate of the qubit's resonant frequency  $\tilde{\Omega}_0$  (e.g. a difference of the order of  $MHz$  vs.  $GHz$ ). Their application is separated by free evolution for a time  $\Delta t$  during which the only action on the system is the parallel field  $\hat{H}_0$ . During this period, with no driving, the Bloch vector oscillates at a frequency  $\delta = \Omega_0 - \omega$  around the  $z$  axis,  $\Omega_0$  being the *real* resonant frequency and the quantity of interest.

The last applied pulse's frequency determines the free evolution, during which the effective Hamiltonian is the Larmor-type Hamiltonian given by equation 5.33 when  $\Omega_d = 0$ . During this period, the driving is off, so its amplitude is indeed null. In the end, we apply a second identical pulse, after which measuring in the  $z$  basis gives (again assuming initialization at  $|0\rangle$ ):

$$\mathbf{P}(1 | t) = \cos^2 \left( \frac{\delta}{2} \cdot \Delta t \right) \quad (5.36)$$

If the pulse frequency  $\tilde{\Omega}_0$  exactly matched  $\Omega_0$ , no oscillations would be observed, because under those circumstances  $\delta = 0$ . In the likely case that it doesn't, the difference between these values can be learned: it corresponds to the observed oscillation frequency  $\delta_{\text{meas}}$ , which can be estimated by performing

measurements for several  $\Delta t$  and performing inference of fitting a curve. Its value can then be used to correct the estimate.

This technique is most commonly called **Ramsey interferometry**, and its formulation has been central for high precision metrology (including applications as sensitive as atomic clocks [112]). The basic sequence to be used for the Ramsey experiments is depicted in figure 47.




Figure 47: Pulse schedule for estimating the detuning  $\delta_{\text{meas}}$  in a Ramsey experiment. This is known as a *Ramsey sequence*. Note that  $\delta$  is an experimental control determining the pulse frequency. The frequency estimate can then be corrected to higher precision as  $\Omega_0 = \tilde{\Omega}_0 + \delta_{\text{meas}}$

In practice, the oscillation pattern is unlikely to exactly recreate equation 5.36, due to experimental error and environment effects. The protocol suffers from decoherence, similar to what has been described in section 5.1.4, as well as other sources of error such as fluctuations in the detuning. The result is that the sinusoidal wave suffers from damping, which can be modeled by an exponential decay envelope:

$$P(1 | t) = e^{-t/T_2^*} \cos^2\left(\frac{\delta}{2} \cdot \Delta t\right) + \frac{1 - e^{-t/T_2^*}}{2} \quad (5.37)$$

This imposes a symmetric effect on the basis states, since decoherence is expected to affect  $|+\rangle$  and  $|-\rangle$  equally (the  $\pi/2$  pulses translate these results into the computational basis). This expression can also be derived by considering Lorentzian noise in the real parameter, and replacing it with a pair of hyperparameters [51].

Just as some of the damping effects of section 5.1.4 were remediable, Ramsey sequences can benefit from error mitigation techniques for improved robustness. We refer to [112] for discussion on how to realize these refinements. Based on the ideas exposed therein, we construct a very simple improved version of the circuit from figure 47, shown in figure 48. More sophisticated enhancements can be found in [112].

This is called by the authors a Hahn-Ramsey scheme, owing to the way it borrows from the Ramsey and Hahn echo experiments. Due to cancelling out some random errors, it improves the contrast of the Ramsey fringes, and eases their decay. Under the presence of noise, the achievable precision in determining  $\delta$  greatly benefits from this.

Within the scope of this work, the purpose of such an arrangement is to make the coherence time large relative to the duration of at least a few oscillation cycles. One may then assume it to be negligible in the scale of the measurement times, which takes us from equation 5.37 back to 5.36 - effectively halving




Figure 48: Pulse schedule for measuring  $\Omega_0 = \tilde{\Omega}_0 + \delta_{\text{meas}}$  in an echoed Ramsey experiment. Note that  $\delta$  is an experimental control determining the pulse frequency. This is a simplification of the ideas in [112], who called this type of setup a *Hahn-Ramsey sequence*.

the number of parameters. This allows for testing methods that would otherwise not be feasible due to processing constraints (namely TLE variations).

Unlike the ones proposed in subsections 5.1.3 and 5.1.4, Ramsey-type experiments necessarily require pulse-level control; to enforce the chosen detuning  $\delta$ , the frequency of the pulses should be customizable, whereas pre-defined unitaries are tuned to the resonant frequency.

The  $\pi/2$  pulses may be constructed from scratch by running a Rabi fit to determine the amplitude given the intended duration, or vice-versa. Calibrating these pulses is the type of task the Ramsey experiment should realize, but naturally a starting point is required. An alternative is to adapt standard operations, namely rotations around the  $x$  or  $y$  axis. Since gates are calibrated to be in resonance, the frequency should be altered by a small detuning  $\delta$ .

Of course, changes in frequency necessitate a correction in the pulse duration and/or amplitude (for a fixed angle, e.g.  $\pi/2$ ), which should be adapted by solving 5.35 for  $\Delta t$ . Depending on the intended precision, bringing the pulses slightly off-resonance while maintaining the amplitude  $\Omega_d$  and the pulse duration may be sufficiently accurate. For a  $\delta$  small enough as compared to the resonant frequency  $\Omega_0$ , the effect in the pulse duration should be negligible.

## 5.2 Characterization of quantum devices

This section presents the results of applying Bayesian inference to the problem of learning the dynamical parameters of quantum devices (open quantum systems), using various sampling strategies. The subjects of characterization were a few IBM quantum machines.

Subsection 5.2.1 details considerations relevant to the experiments, namely how they were executed and their performance assessed. Subsections 5.2.2 and 5.2.3 present the results of coherence time estimation ( $T_2$  from subsection 5.1.4 and  $T_1$  from 5.1.3 respectively). Subsections 5.2.4 and 5.2.5 concern Ramsey estimation, as per subsection 5.1.5. In the former, the basic pulse sequence is taken, so the induced oscillation frequency  $\delta$  is targeted along with  $T_2^*$ . In the latter, the results of applying a refocusing pulse are demonstrated, and more resource-intensive techniques are applied to the characterization of the frequency, which becomes one-dimensional (assuming  $T_2^* = \infty$  for short enough evolution times).



### 5.2.1 Experimental setup

The experiments were performed using IBM quantum services [61], Qiskit [6] and Qiskit-Pulse [2], a pulse-level programming kit.

Except where indicated differently, measurement times were chosen non-adaptively on a grid, and the posterior distribution was discretized via **SIR** with the prior as importance function (algorithm 4, subsection 3.2.2) and **MCMC** move steps (algorithm 5, (sub)sections 3.3 and 3.3.3). Additionally, the chosen Markov kernel was unless otherwise stated **RWM** (subsection 3.3.1) with a proposal variance proportional to the current **SMC** estimate of the variance. This proportionality was tuned to obtain a stable particle acceptance rate around 65%.

For each set of results, the backend estimates at the time of data collection are presented where relevant and applicable. The latest known coherence times are supplied by IBM at any given time. However, these estimates are not very reliable: variations as high as 50% were observed between consecutive calibrations. Intervals of hours between calibrations are not uncommon, sometimes lasting up to a full day. Within that time span, the time constants tend to suffer from significant fluctuations.

As such, curve fits were used as a benchmark to validate and contextualize results. They additionally allow for comparing the achieved uncertainty in view of the data requirements, yielding a more complete performance assessment. The fitting was done using two methods: the SciPy [111] function `curve_fit` (non-linear least squares fit) with default parameters, and Qiskit's native coherence fitters (from the `BaseCoherenceFitter` class).

The former was used for direct comparison, since it can work on the same data (measurement times and pulse schedules) as the inference protocol, whereas the default fitters take only gate sequences. Regardless, the methods were approximately matched for maximum evolution time, and for each of them the measurement times were evenly spaced from some near-zero value to this upper limit.

The SciPy fitter is the same tool that Qiskit coherence fitters use internally, but Qiskit's error estimates are more reliable. They are adjusted according to the variance of each measurement over all shots, through the optional argument `sigma` of `curve_fit`. It was made evident by the obtained results that the default parameters result in the errors' being grossly overestimated. For a fair comparative analysis, the quantitative results of the Qiskit fitters are the ones presented.

In the case of the inference, whose performance is more variable due to using fewer data, these quantitative results were obtained using a number of runs and taking median values for both the parameter estimates (means) and their associated uncertainties (standard deviations). In some cases, the evolution of the uncertainty through the iterations of the **SMC** algorithm is also represented. Finally, the results will where relevant be represented graphically as the curve associated with the determined parameter(s).

In the last case, the same graphs present data points corresponding to the (*measurement time*, *probability*) tuples obtained using custom pulse schedules and a curve fitted specifically to them. This is done using external (non-Qiskit) curve fitting. The data used for the points and for these curve fits are exactly the same, whereas the inference results use a lower number of shots. This serves as an

assessment of predictive power. Curve fits tend to require higher resolution if reasonable results are to be obtained.

In the cases where direct quantitative comparison is intended, the same backend was used, and maximal time proximity for the measurements across several methods was aimed for. Even still, unavoidable factors make it so that gathering datasets in succession where possible does not guarantee that each registered identical behaviour. Figure 49 shows the time evolution of the estimated values using inference (whose results can be more local in time due to relying on fewer shots). The delays are due to queuing, execution times and repetition of runs to extract statistics. As such, discrepancies may in some cases be attributable to sources other than error in the estimation.

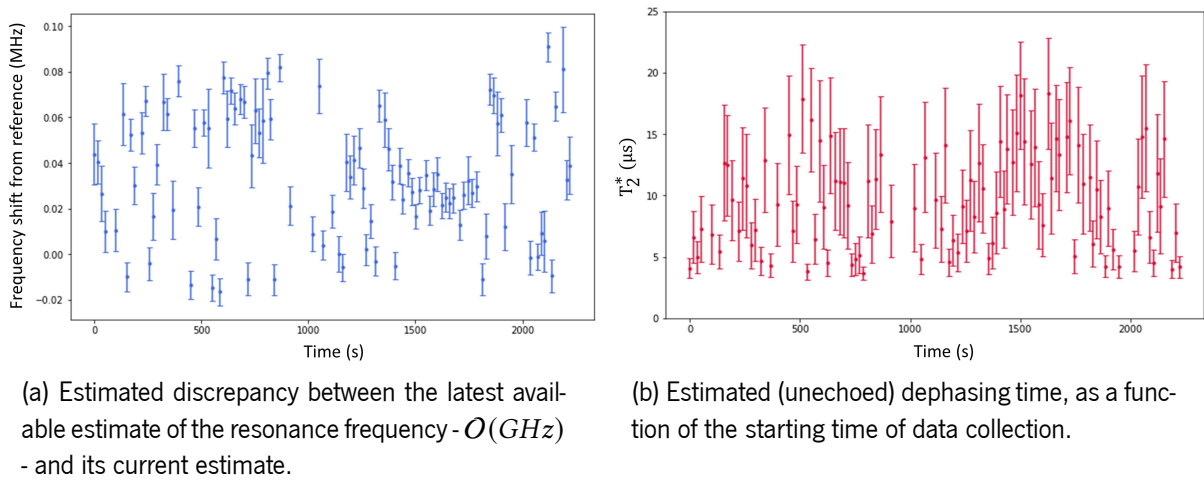


Figure 49: Estimated time fluctuations of two quantities of interest for the IBMQ backend `ibmq_armonk`. These are the parameters to be targeted in the Ramsey experiment of section 5.2.4. The  $x$  coordinates are the starting times of data collection. The error bars represent standard deviations.

Unlike curve fits, inference can work with all-different times for a single measurement each. However, Qiskit only allows up to 75 different pulse schedules per job submitted, each of which can be repeated for up to 8192 shots. To streamline the data collection process, a maximum of 75 different times is used for the data vectors, the total desired number of experiments being distributed by these times.

It should be noted that the problem of estimating time constants is fairly simple, and not necessarily a case where inference stands out; oscillations are a more interesting application. Even so, they make an interesting test case, as they elucidate the impact of intrinsic differences between likelihood models (in terms of learning rate, most favorable SMC sequences, particle number, effective sample size, order of magnitude of the evolution times, etc.). They also help understand the difficulties estimating an oscillation frequency along with a damping constant brings, due to conflicts between what is optimal for each.

For the Ramsey experiments, the probabilities of 0 and 1 are switched with respect to equations 5.36 and 5.37. This is because a the half-pi pulses were adapted from a Hadamard gate - using the default instruction `u2(0,  $\pi$ )` -, which is an involution. Because it represents a rotation of  $\pi$  around the  $x + z$  axis, when applied twice to  $|0\rangle$  it brings the state back to  $|0\rangle$  rather than to  $|1\rangle$ . This is without prejudice

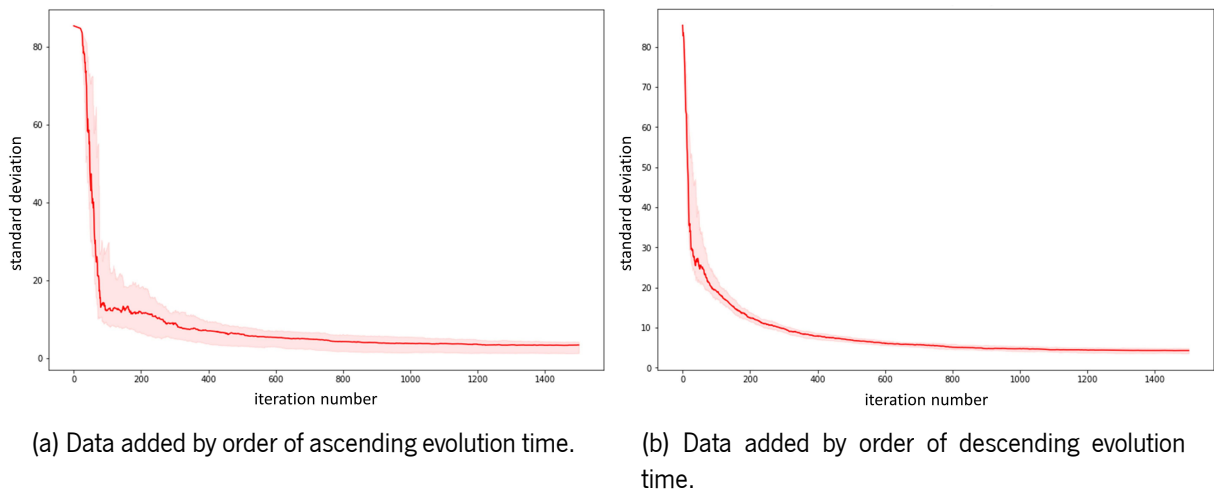


Figure 50: Evolution of the standard deviation of the dephasing time  $T_2$  during the inference process, for the IBMQ device `ibmq_rome` and two different orders of data addition. Both results were based on the same dataset containing 1500 observations. The shaded region represents the interquartile range.

to the covered angle's being as intended  $\pi/2$  on the  $x - z$  plane.

### 5.2.2 Hahn echo experiment ( $T_2$ estimation)

This subsection presents the results of estimating the dephasing time  $T_2$  (subsection 5.1.4) of a few IBM quantum devices.

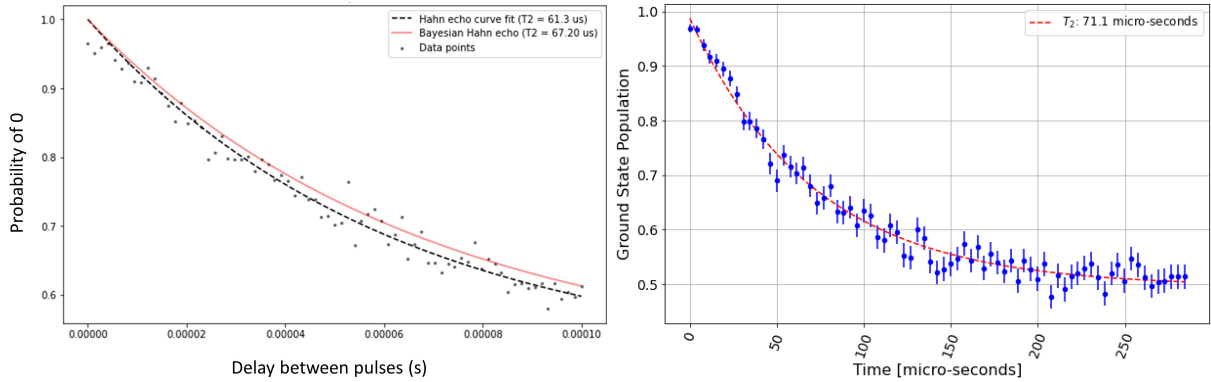
In the case of exponential decay models such as this one, the sampler benefits from long evolution times being considered sooner, because near the origin the information is more challenging. This stands in contrast with the precession frequency example taken as an example throughout chapters 2 and 3, but follows the same line of reasoning.

Figure 50 illustrates this fact. The only difference between graphs is the order by which the data are fed to the sampler. In the case of figure 50a the shortest times were considered first, whereas in that of 50b the opposite was done.

The sharp drop in the uncertainty followed by a plateau and the larger inter-run variability in the first case are a sign of a sharp drop in effective sample size and undue particle concentration (owing to particle collapse). Conversely, the second graph shows a steadier learning process. Thus, the longer measurement times were included earliest in the sequence of SMC target distributions.

Due to systematic errors, the coefficient and constant of equation 5.21 may not be exactly  $1/2$ . In that case, it may be beneficial to allow for more freedom in the coordinate at origin and horizontal asymptote, which is done by default in the Qiskit fits. This problem is illustrated in figure 51. Equation 5.21 should then be changed to  $\mathbf{P}(1 | \Delta t) = A \exp(-t/T_2) + B$ .

To remedy this, one can include a rough estimate of the hyperparameters  $A$  and  $B$  in the inference model. This can be achieved with a higher shot count for  $t = 0$  and some large  $t \rightarrow \infty$  (implying a constant measurement overhead). The fraction of 0 outcomes for each of these times can be identified



(a) Results obtained using Bayesian inference (solid red line) and a curve fit (dashed black line), superposed on the data points used for the latter (gray dots), which used 512 shots each. For the inference, only 4% (20) of these shots per point were used.

(b) Results obtained using the Qiskit coherence fitter for  $T_2$ , with 512 shots per data point. The model of equation 5.21 is relaxed to adapt the coordinates at  $t = 0$  and  $t = \infty$ .

Figure 51: Curves produced by  $T_2$  estimates for the IBMQ device `ibmq_rome`, using 75 different measurement times. In figure 51a the model of equation 5.21 was used, whereas 51b relaxes it to include two constants  $A$  and  $B$ .

with  $A$  and  $B$  respectively. The result of using 512 shots for these model calibration measurements before the inference process is shown in figure 52. Though they were chosen to match the fitter, rougher estimates can be used, e.g. 50. The curves can be seen to fit the data better as compared to figure 51a.




Figure 52: Curves produced by  $T_2$  estimates for the IBMQ device `ibmq_casablanca`, using also estimates for  $A, B \neq 1/2$ . The results were obtained using Bayesian inference (solid red line) and a curve fit (dashed black line), superposed on the data points used for the latter (gray dots). The curve fit relied on 512 shots per point, including the extremes to which it was adapted. For the inference, 20 shots per point were used, and 100 for each of the hyperparameter-adjusting ( $A$  and  $B$ ) measurements.

With these adjustments in mind, we can compare the quantitative results of inference and the Qiskit fitter. The experiments were in this case done on the `ibmq_casablanca` backend.

|                           | $T_2$ ( $\mu s$ ) | Standard deviation | Shot number                          |
|---------------------------|-------------------|--------------------|--------------------------------------|
| <b>Bayesian inference</b> | 52.51             | 4.4                | $20 \cdot 75 + 512 \cdot 2 = 2\,524$ |
| <b>Qiskit fitter</b>      | 52.39             | 3.7                | $512 \cdot 75 = 38\,400$             |
| <b>Qiskit fitter</b>      | 56.70             | 13.9               | $169 \cdot 15 = 2\,535$              |

Table 2: Dephasing time estimation results for the IBMQ device `ibmq_casablanca`. The registered backend estimate was at the time  $T_2 = 57.58\mu s$  for the two first lines, and  $48.31\mu s$  for the last one (as mentioned before, variations such as these are just an experimental artifact).

The conditions for the inference were as follows. A 50 particle **SMC** approximation was used, with a flat prior on  $]0, 250]\mu s$ . A single Markov (**RWM**) move was used per particle per step. The threshold effective sample size for resampling was set at  $\widehat{ESS} = 0.8$ . A total of 75 different measurement times was used for 20 shots each, signifying 1500 data/steps, plus  $2 \cdot 512$  additional shots to customize  $A$  and  $B$ . Of the total steps, 33 (2%) triggered a resampling stage on average. The measurement times were chosen in constant increments within  $t \in ]0, 115]\mu s$ . The results were obtained by taking medians over 100 runs split equally by 10 different datasets.

The results are shown in table 2, and compared to those of the Qiskit fitter. For comparison purposes, the latter was applied in two sets of conditions. First, using the same number of measurement times and range, but for 512 shots each instead of 20. In this case, both the obtained values and their associated uncertainties were close, in spite of the discrepancy between shot counts. Second, allowing the curve fitter roughly as many total measurements as were used for the inference; the associated graph can be seen in figure 53. The measurement times were uniformly spaced out in the same range as for the inference, but with fewer different times and more shots. The achieved error was over 3 times larger than that obtained by inference.




Figure 53: Curve produced by the Qiskit coherence fitter for  $T_2$  for the IBMQ device `ibmq_casablanca`, using 15 different measurement times for 169 shots each (a total of 2 535 shots).

The reason for using different experiments is that simply providing the fitter the exact data that was used for the inference yielded an error in the order of the thousands of microseconds. Starting from there,

the variance could be reduced by reducing the number of measurement times, originally 75, by which the shots were distributed. This was done with a granularity of 10 times. After a point, the trend was reversed; in the table, the lowest achieved standard deviation is presented. Though the measurement schemes were not exhaustively picked to be the best, still less were those chosen for the inference: they were simply arrayed on intervals for which the signal was still sure to be passably clear.

In conclusion, both comparative cases strongly suggest that that inference is more robust for extracting information from a small dataset.

### 5.2.3 Quantifying energy loss (T<sub>1</sub> estimation)

For this subsection, the parameter of interest is the energy relaxation time  $T_1$  (subsection 5.1.3). The experiment is identical to that of the previous section, except for the fact that the occupation of  $|1\rangle$  tends to 0% and not 50%. Figure 54 shows the evolution of the uncertainty as the inference proceeds.




Figure 54: Evolution of the standard deviation during the estimation of the energy relaxation time constant  $T_1$  of IBMQ device `ibmq_guadalupe`. The dataset contemplated 1500 observations. The shaded region represents the interquartile range.

The conditions were similar to section 5.2.2, apart from two points. The prior was set on  $]0, 100]\mu\text{s}$  (this doesn't have much of an impact after the first few iterations), and the measurement times on  $]0, 50]\mu\text{s}$ . A standard deviation of  $\sigma = 2.9\mu\text{s}$  for an estimate of  $T_1 = 62.45\mu\text{s}$  was achieved.

### 5.2.4 Ramsey experiment (2-dimensional estimation)

This subsection presents the results of performing Ramsey experiments (subsection 5.1.5) on IBM quantum devices. The aim was to use the circuit from figure 47 to estimate the detuning  $\delta$  and coherence time  $T_2^*$  from equation 5.37. Instead of targeting the latter directly, the factor  $\gamma_2^* = 1/T_2^*$  is considered, so that the two parameters are matched in units (though not necessarily scale) as suggested in [51].

As for the frequency estimation, relevant considerations have been laid out in subsection 2.3.3 and the ensuing sections. Unlike for the isolated characteristic time estimation of subsections 5.2.2 and 5.2.3, the data were considered by order of increasing evolution times, as this was found to work best in this

case. Note that although this tends to streamline *frequency* estimation, it has the opposite effect for the characteristic time, as demonstrated in subsection 5.2.2 for the non-oscillatory case. Of course, once these parameters are joined together in a likelihood function, there exists a mutual influence in their learning rates, creating a more delicate interplay. The tension between what is best for each parameter complicates the decision process, given that they play such different but inseparable roles.

An example the evolution of standard deviations through the iterations is shown in figure 55, where the difference between parameter learning rates is evident.


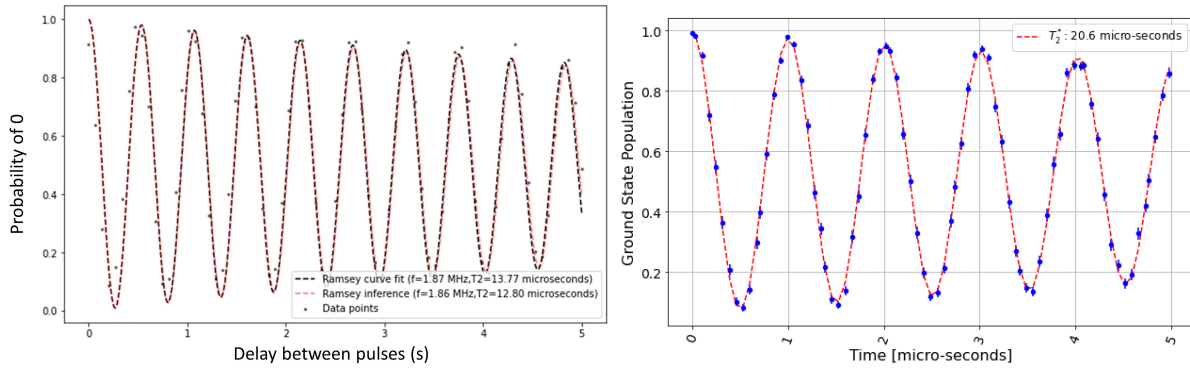


Figure 55: Evolution of the standard deviations during the estimation of the detuning frequency  $\delta$  and coherence factor  $\gamma_2^*$  of IBMQ device `ibmq_armonk`. The dataset contemplated 150 observations. The solid curves are the median results for 300 runs split by 100 datasets. The shaded regions represent interquartile ranges.

The default Qiskit fitter for  $T_2^*$  takes as an input parameter the number of oscillations to be induced during the time span, rather than a detuning frequency. It is geared towards coherence time estimation, and not frequency calibration. Thus, this is the quantity whose estimation is to be compared, though the frequency still brings added complexity. An example of the results obtained by inference and by curve fitting is presented in figure 56. For reference, the uncertainties in the estimation of the induced oscillation frequency were 0.4% for inference, nearly 0% for curve fitting with 256 times more shots, and 15% for curve fitting with the same shots.

As discussed in section 5.1.4,  $T_2^*$  is susceptible to large variations, so any comparison should be taken with a grain of salt. Qiskit does not make estimates of this constant available (despite still supplying a built-in fitter for it), likely due to its volatility.

We will now present some quantitative final results for a specific setting. For the inference, a 225 particle SMC approximation was used, with a flat prior on  $f \in ]0.0, 5.0] MHz$ ,  $T_2^* \in [3.00, 25.00] \mu s$ . Three Markov moves were used per particle per step. The threshold effective sample size for resampling was set at  $\widehat{ESS} = 0.8$ . A total of 75 different measurement times was used for 2 shots each, signifying



(a) Results obtained using Bayesian inference (solid red line) and a curve fit (dashed black line), superposed on the data points used for the latter (gray dots), which used 512 shots each. For the inference, only 0.4% (2) of the shots per point were used.

(b) Results obtained using the Qiskit coherence fitter for  $T_2^*$ , with 512 shots per data point.

Figure 56: Curves produced by  $\delta$  (labeled  $f$ ) and  $T_2^*$  estimates for the IBMQ device `ibmq_guadalupe`, using 75 different measurement times.

150 data/steps. Of these, 28 (19%) triggered a resampling stage on average. The measurement times were chosen in constant increments within  $t \in ]2.0, 5] \mu\text{s}$ . The results were obtained by taking medians over 100 runs split equally by 10 different datasets.

For a detuning of  $\delta = 1.83\text{MHz}$ , the final estimates were  $\delta = 1.863\text{MHz}$  and  $T_2^* = 9\mu\text{s}$ , and their associated uncertainties  $\sigma = 0.007\text{MHz}$  and  $3\mu\text{s}$ . Table 3 compares the  $T_2^*$  results between this method and Qiskit's. We again observe that the *real* values are quite fickle, which explains the disparities.

|                           | $T_2^*$ ( $\mu\text{s}$ ) | Standard deviation | Total shot count         |
|---------------------------|---------------------------|--------------------|--------------------------|
| <b>Bayesian inference</b> | 9                         | 3                  | $75 \cdot 2 = 150$       |
| <b>Qiskit fitter</b>      | 21                        | 1                  | $75 \cdot 512 = 38\,400$ |
| <b>Qiskit fitter</b>      | 9                         | 30                 | $15 \cdot 10 = 150$      |

Table 3: Unechoed dephasing time estimation results for the IBMQ device `ibmq_armonk`.

Once more, the last row presents the result of running a curve fit for as many total measurements as were included in the inference's dataset (the corresponding graph is shown in figure 57 for visualization), which was 10 times worse than the inference's. The registered standard deviation was the best one obtained among multiple good-faith attempts. They were chosen as for the case of table 2, only now the time range too had to be adapted (shortened) to obtain reasonable results. These adaptations were at all times in benefit of the curve fitter. Most choices of running curve fits with 150 total shots yielded infinite variance, including when matching the scheme used for inference. Again, the latter dealt better with limited information.






Figure 57: Curve produced by the Qiskit coherence fitter for  $T_2^*$  for the IBMQ device `ibmq_guadalupe`, using 10 different measurement times for 15 shots each (a total of 150 shots).

For completeness, we note that the fitter achieved uncertainties of 0 and 15% for the frequency, for the higher and equal shot counts respectively (second and third lines of table 3). The key point is that once the data allowance was matched, the fitter performed roughly 40 times worse for this parameter, an even starker contrast. The induced oscillation frequencies were matched for order of magnitude.

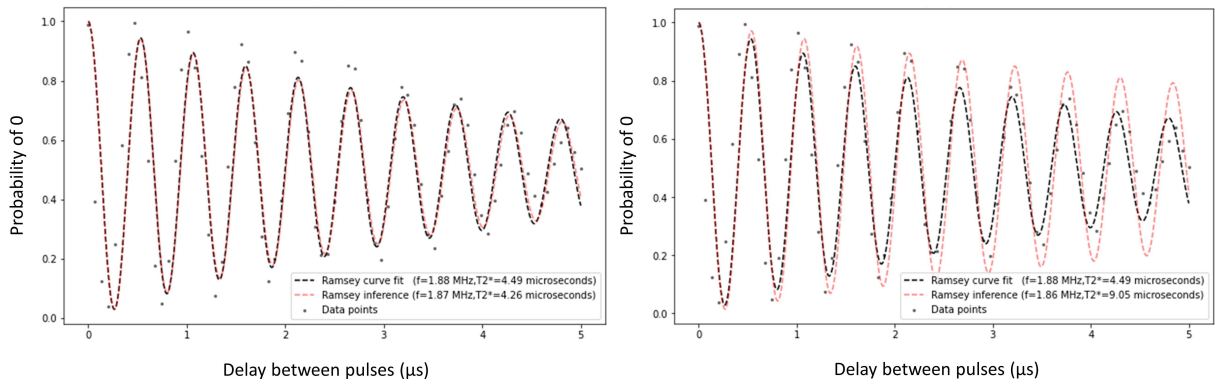
Figure 56a suggests that the inference results fit the data well. Expectably, the results are more variable when the coherence times are shorter. In particular, they may be quite local in time; due to using few shots, they capture short-time pictures of the system (see also figure 49). The results tend to be more consistent for evolution times well within the coherence time, where the signal is less noisy. This is illustrated in figure 58. In 58a, inference was performed based on a dataset extracted from the presented data points; it underestimated the earliest data points' oscillation amplitude. Using a more general dataset (figure 58b) tends to fit the earliest points of the dataset better than inference based on the dataset itself.

### 5.2.5 Echoed Ramsey experiment (1-dimensional estimation)

This subsection presents the results of performing *Hahn-Ramsey* experiments on a few IBM quantum devices. The aim was to use the circuit from figure 48 to estimate the detuning  $\delta$  from equation 5.36. Since learning a single parameter is less demanding than learning two, this allows for testing methods that require more computational resources while still targeting a more interesting example than those of sections 5.2.2 and 5.2.3.

The impact of the pi-pulse is illustrated in figure 59. When comparing figure 59b to the original case of figure 59a, we can see that the decoherence in the represented interval (up to  $5\mu s$ ) is very slight. For measurements within that interval, we can approximate the amplitude as 1 while still obtaining reasonable results. The coherence time was improved by around 400%.

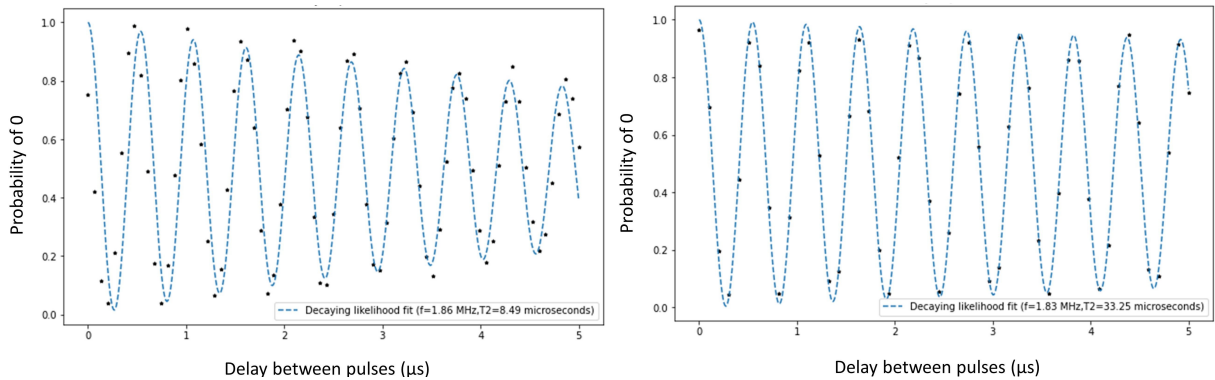
This was used to test adaptive experimental design on `ibmq_armonk` (section 2.4). The strategy was compared to the non-adaptive (offline) estimation of before. This serializes the data collection process. Due to limited usage of the IBMQ devices, only 15 data/steps were used for both methods, and 100 runs. Still, the difference was clear (figure 60).



(a) Bayesian inference results plotted along with the data they were based on (4% of the information contained by these data was used).

(b) Bayesian inference results plotted along with a preliminary set of gathered data, which preceded most of the data the inference was based on.

Figure 58: Ramsey experiment inference results obtained using 100 runs over the same dataset, or distributed by successive datasets spanning a longer period of time. The inference results (dashed red lines) are plotted against the same set of data points (gray dots), from which the data for the inference was extracted in the case of figure 58a. A curve fit (dashed black lines) tailored to these points is included for reference. Each of the data points used 512 shots, whereas for the inference only 2 were used for each time. The targeted device was `ibmq_armonk`, with 75 different measurement times.



(a) Data points generated using the Ramsey sequence of figure 47.

(b) Data points generated using the Ramsey sequence of figure 48.

Figure 59: Effect of inserting a refocusing pulse in a Ramsey sequence for `ibmq_armonk`. Each of the 75 data points (black dots) was obtained using 512 shots. To aid visualization, a curve fit (dashed blue lines) performed using SciPy for the function of equation 5.36 is plotted along.

A 100 particle SMC approximation was used, with a flat prior on  $\delta \in ]0.0, 10.0] MHz$ . One Markov move was used per particle per step. The threshold effective sample size for resampling was set at  $\widehat{ESS} = 0.5$ . A total of 15 different measurement times was used for a single shot each, signifying 15 steps. Of these, 7 (47%) and 9 (60%) triggered a resampling stage on average for the adaptive and offline methods respectively. The offline measurement times were chosen in constant increments within  $t \in ]0.2, 2] \mu s$ , and the adaptive ones were chosen from 20 guesses around  $1/\sigma_{curr}$  to have the lowest variance (as suggested in subsection 2.4.3). The results were obtained by taking medians over 100 runs




Figure 60: Evolution of the standard deviation during the inference process for adaptive (blue line) and offline (red line) estimation. The shaded regions represent interquartile ranges.

split equally by 10 different datasets. The quantitative results are presented in table 4.

|                 | $\delta$ (MHz) | Standard deviation | Total shot count | Precision ( $\sigma^2 \Delta t_{acc}$ ) |
|-----------------|----------------|--------------------|------------------|---|
| <b>Adaptive</b> | 1.8            | 0.1                | 15               | 0.38                                    |
| <b>Offline</b>  | 2.3            | 0.8                | 15               | 9.3                                     |

Table 4: Results of using adaptive and offline inference to learn a detuning frequency on the IBMQ device `ibmq_armonk`. The detuning relative to the prevailing backend estimate was  $1.83\text{MHz}$ .

An additional quantity was considered: the *precision*, defined as in [99] ( $\sigma^2 \Delta t_{acc}$ , where  $\Delta t_{acc}$  is the cumulative evolution time). This is meant to penalize long measurement times, and judge whether the adaptive method unfairly benefits from lengthier evolutions. According to this metric, adaptivity still outperforms the offline method. Note that penalties of this type can be incorporated into the utility function, if curtailing simulation runtime is critical.

Next, TLE with HMC move steps was attempted (figure 61) for the backend `ibmq_armonk`. A 100 particle SMC approximation was used, with a flat prior on  $f \in ]0.0, 10.0]\text{MHz}$ . One Markov (HMC) move was used per particle per step. The HMC hyperparameters were set to  $M = \sigma_{curr}^2$ ,  $L = 20$ ,  $\epsilon = 0.001$ , producing an acceptance rate of 68%. The threshold effective sample size for resampling was set at  $\widehat{ESS} = 0.5$ . A total of 75 different measurement times was used for a single shot each. The tempering coefficients were spaced on a grid for a total of 10, signifying 10 steps. Of these, 3 (30%) triggered a resampling stage on average. The measurement times were chosen in constant increments within  $t \in ]0.2, 10]\mu\text{s}$ . The results were obtained by taking medians over 100 runs split equally by 10 different datasets. The quantitative results are presented in table 5.

Finally, subsampling TLE with RWM move steps with control variates was tested. The methods are those of subsection 3.5. As has been mentioned in chapter 4, a narrowed down the prior is required for accuracy due to the control variates. This was achieved via a warm up on the lower times. The results are presented in figure 62.

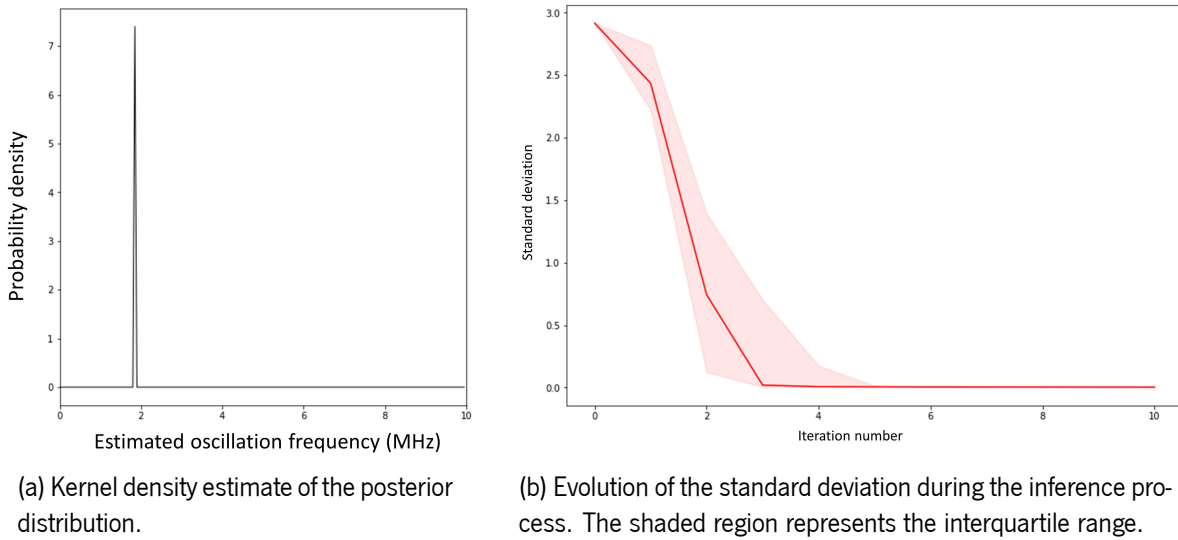


Figure 61: Results of using TLE with HMC move steps to learn a detuning frequency on the IBMQ device `ibmq_armonk`. The dataset contemplated 75 observations, and ten tempering coefficients were used.

|                           | $\delta$ (MHz) | Standard deviation | Total shot count |
|---------------------------|----------------|--------------------|------------------|
| <b>Bayesian inference</b> | 1.830          | $6 \times 10^{-3}$ | 75               |


Table 5: Hahn-Ramsey frequency estimation results for the IBMQ device `ibmq_armonk`, using TLE with HMC move steps. The detuning relative to the prevailing backend estimate was  $1.83\text{MHz}$ .

The experiment was performed in similar conditions to the previous one, except for a few aspects. The number of particles was 200, and they were resampled at each step. A total of 75 different measurement times was used for 5 shots each, totaling 375 data. The quantitative results are presented in table 6

|                        | $\delta$ (MHz) | Standard deviation   | Data batch size |
|------------------------|----------------|----------------------|-----------------|
| <b>Full data TLE</b>   | 1.831          | $1.2 \times 10^{-3}$ | 375             |
| <b>Subsampling TLE</b> | 1.832          | $1.0 \times 10^{-3}$ | 150             |
| <b>Warm-up</b>         | 1.83           | $1.3 \times 10^{-2}$ | 150             |

Table 6: Hahn-Ramsey frequency estimation results for the IBMQ device `ibmq_armonk`, using TLE with HMC move steps based on the full data or subsampling. The detuning relative to the prevailing backend estimate was  $1.83\text{MHz}$ .

The difference in the estimator between the full data and subsampling cases was negligible. The standard deviation, albeit slightly more variable, was also close between methods.



(a) Kernel density estimates of the full data posterior (375 observations) and of the warmed up distribution (150/375 observations).

(b) Kernel density estimates of the full data posterior (375 observations) and of the subsampling posterior (150/375 observations).

Figure 62: Results of using TLE with RWM move steps and data subsampling with control variates to learn a detuning frequency on the IBMQ device `ibmq_armonk`. A close up of the distribution is presented, though the prior covered  $]0, 10]MHz$ . The dataset contemplated 375 observations, and 10 tempering coefficients were used, 3 of which were employed in the warm up in the subsampling case. When warming up or subsampling, each likelihood evaluation relied on 150/375 data, in the former case the 150 ones with lower evolution times.

## Discussion and conclusions

The alliance between Bayesian inference and Monte Carlo methods is a powerful tool for quantum metrology, and one which can negotiate costs on all fronts. It offers a springboard for achieving optimal balance between accuracy and each of the key resources: classical processing and memory, number of quantum measurements, number of quantum simulation runs, and simulation runtime.

Moreover, these strategies enjoy remarkable generality, barely making assumptions about the problem at hand. This is a large part of what popularized them as versatile instruments for data science, but applies even within the realm of quantum mechanics. They supply a way of characterizing quantum systems via experimental outcomes, in the broadest possible sense. These *quantum systems* can be digital computers, sensors, analog simulators, circuits, or others; and the experiments can be any, as long as their outcomes provide an insight into the system's inner workings. This makes them a suitable solution for countless quantum problems, including archetypal characterization tasks (tomography, parameter estimation) but also routines as fundamental as phase estimation. Thus, the methodologies exposed herein are more than means for the validation of quantum systems: they are also one more example of how machine learning and other classical processing strategies can serve pragmatic quantum algorithms, and assist in the extraction of maximal utility from state-of-the-art quantum devices.

One of the scenarios where the Bayesian approach shines the most is when the available information is scarce, or its obtention costly. For small datasets, its performance was shown to be vastly superior to that of standard regression-based curve fitting. Other notable contexts for application are those where processing data as it arrives is an advantage. In particular, this capability paves the way for achieving fundamental limits of metrology in the quantum regime. Adaptive quantum measurements informed by Bayesian experimental design, another perk afforded by the paradigm's far-reaching nature, present an interesting path towards Heisenberg-limited estimation. Importantly, this approach relinquishes onerous state preparation and prohibitively long coherence times in favor of classical resources. As with other NISQ-friendly protocols, this grants current devices the opportunity to observe quantum advantage with greatly eased requirements. Even with the simple experiments and sparing use of optimization considered here, their potential was clear, as empirical strategies for roughly locally optimal designs greatly improved the achievable precision. Equivalently, adaptivity promises to bring a sharp drop in the number of measurements required for attaining some precision threshold.

Another aspect that can have a drastic impact in performance is the representation of the posterior distribution. In Bayesian inference, given a generative model for the data, this is in principle the only source of inexactitude, as the results aren't otherwise assumed asymptotic. All sensible inquiries (in a

---

statistical/measure-theoretic sense) to this distribution, and thus all results, are entrusted to this representation through expectation integrals. Such an ability to forsake optimization is crucial for dependability given complex models; and the lack thereof is the pitfall of many machine learning algorithms, whose naiveté becomes fatal in high-dimensional spaces.

The literature on Bayesian algorithms for quantum phenomena has favored uncomplicated strategies for relaying the results, but these are bound to become ineffective as more complex systems are targeted. And who better than modern day inference to foretell of their prospective future? Simple methods are fitting for simple problems, and/or when processing speed or memory requirements are the main concerns. However, they can quickly become unworkable, forfeiting the richness - and even the correctness - of Bayesian statistics. Not only does this compromise the results, it may also misinform experimental decisions. Ultimately, problems such as these will frustrate attempts at learning Hamiltonians with many parameters; and naturally, generalizing the description to open systems only adds to the difficulty.

In that case, high performance [MCMC](#) methods are the foremost solution, as anticipated by their widespread use in the statistical sphere (and for Bayesian inference in particular). These Markov chains can be run independently, or coupled within a [SMC](#) scheme. Due to benefitting from informative inter-particle interactions, the latter option can create especially robust and fast algorithms; as a bonus, the ensemble offers complimentary estimators of the the model evidence. Keeping [SIR](#) but substituting [LWF](#) with Markov kernels is particularly interesting, due to the fact that many advantages are kept, such as the compatibility with online processing. Markov moves inside [SMC](#) are a widely used strategy in statistics and across the multitude of fields that rely on it, especially when correctness and generality are a concern. For offline processing, [TLE](#) may offer still more correctness, namely for multimodal and high-dimensional distributions. It also offers a more stable understructure for subsampling, which may overcome its typically higher resource count to bring quite favorable scaling.

As for the construction of the Markov transitions themselves, the best practice is to tailor them to the problem. For continuous probability densities, [HMC](#) is doubtlessly a strong contender, owing to the solid geometric foundations that underpin its many empirical successes. In some cases, its implementation may require adjustments; in others, such an elaborate method may be unwarranted, as others can fill in for a fraction of the cost. Regardless, it is to be expected that scaling up the complexity will eventually require upskilling the walker, and [HMC](#)'s is arguably the most skillful of them all - the more so as the original framework is built upon.

From scrupulous experimental design to advanced sampling strategies, the tactics explored throughout this dissertation are interesting directions to be expanded further in future research, as attested by their aptness for learning the dynamical parameters of noisy quantum systems based on meager datasets. Resource intensive as they may seem, they are ultimately more cost-effective once the problem becomes complex enough - which this problem is bound to. For any progress that is made, the means for its validation should trail after; this is a necessary - if challenging - imposition for reaching new regimen of physics in applied practice. As science pushes the frontiers of what is realisable, efficient characterization will play a fundamental role in guarding the path to quantum advantage.

# Bibliography

- [1] J. Alenlöv, A. Doucet, and F. Lindsten. “Pseudo-Marginal Hamiltonian Monte Carlo”. In: *Journal of Machine Learning Research* 22.141 (2021), pp. 1–45. url: <http://jmlr.org/papers/v22/19-486.html>.
- [2] T. Alexander et al. “Qiskit pulse: programming quantum computers through the cloud with pulses”. In: *Quantum Science and Technology* 5.4 (Aug. 2020), p. 044006. issn: 2058-9565. doi: 10.1088/2058-9565/aba404. url: <http://dx.doi.org/10.1088/2058-9565/aba404>.
- [3] A. R. Alves. *Quantum parameter estimation via Bayesian learning (Masters project)*. <https://github.com/alexandra-frca/bayesianlearning>. 2021.
- [4] C. Andrieu and G. O. Roberts. “The pseudo-marginal approach for efficient Monte Carlo computations”. In: *The Annals of Statistics* 37.2 (Apr. 2009). doi: 10.1214/07-aos574. url: <https://doi.org/10.1214/07-aos574>.
- [5] C. Andrieu et al. “An Introduction to MCMC for Machine Learning”. In: *Machine Learning* 50.1 (Jan. 2003), pp. 5–43. issn: 1573-0565. doi: 10.1023/A:1020281327116. url: <https://doi.org/10.1023/A:1020281327116>.
- [6] M. S. ANIS et al. *Qiskit: An Open-source Framework for Quantum Computing*. 2021. doi: 10.5281/zenodo.2573505.
- [7] A. Asfaw et al. *Learn Quantum Computation Using Qiskit*. 2020. url: <http://community.qiskit.org/textbook>.
- [8] A. Barbu and S. Zhu. “Hamiltonian and Langevin Monte Carlo”. In: Springer Singapore, Feb. 2020, pp. 281–325. isbn: 978-981-13-2970-8. doi: 10.1007/978-981-13-2971-5\_9.
- [9] D. W. Berry et al. “How to perform the most accurate possible phase measurements”. In: *Physical Review A* 80.5 (Nov. 2009). doi: 10.1103/physreva.80.052114. url: <https://doi.org/10.1103/physreva.80.052114>.
- [10] A. Beskos et al. “Geometric MCMC for infinite-dimensional inverse problems”. In: *Journal of Computational Physics* 335 (Apr. 2017), pp. 327–351. doi: 10.1016/j.jcp.2016.12.041. url: <https://doi.org/10.1016/j.jcp.2016.12.041>.
- [11] M. J. Betancourt. *Adiabatic Monte Carlo*. 2015. arXiv: 1405.3489 [stat.ME].



- [12] M. J. Betancourt. *Efficient Bayesian inference with Hamiltonian Monte Carlo*. <https://www.youtube.com/watch?v=pHsuIaPbNbY&t=3095s>. Machine Learning Summer School Iceland, 2014.
- [13] M. J. Betancourt. *Scalable Bayesian Inference with Hamiltonian Monte Carlo*. <https://www.youtube.com/watch?v=jUSZboSq1zg&t=1395s>. London Machine Learning Meetup, 2018.
- [14] M. J. Betancourt. *The Fundamental Incompatibility of Hamiltonian Monte Carlo and Data Subsampling*. 2015. arXiv: 1502.01510 [stat.ME].
- [15] M. Betancourt. *A Conceptual Introduction to Hamiltonian Monte Carlo*. 2018. arXiv: 1701.02434 [stat.ME].
- [16] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. “Variational Inference: A Review for Statisticians”. In: *Journal of the American Statistical Association* 112.518 (Apr. 2017), pp. 859–877. doi: 10.1080/01621459.2017.1285773. url: <https://doi.org/10.1080%2F01621459.2017.1285773>.
- [17] D. Bontemps. “Bernstein–von Mises theorems for Gaussian regression with increasing number of regressors”. In: *The Annals of Statistics* 39.5 (Oct. 2011). doi: 10.1214/11-aos912. url: <https://doi.org/10.1214%2F11-aos912>.
- [18] A. Buchholz, N. Chopin, and P. E. Jacob. “Adaptive Tuning of Hamiltonian Monte Carlo Within Sequential Monte Carlo”. In: *Bayesian Analysis* 16.3 (Sept. 2021). doi: 10.1214/20-ba1222. url: <https://doi.org/10.1214%2F20-ba1222>.
- [19] B. A. Chase and J. M. Geremia. “Single-shot parameter estimation via continuous quantum measurement”. In: *Phys. Rev. A* 79 (2 Feb. 2009), p. 022314. doi: 10.1103/PhysRevA.79.022314. url: <https://link.aps.org/doi/10.1103/PhysRevA.79.022314>.
- [20] T. Chen, E. B. Fox, and C. Guestrin. *Stochastic Gradient Hamiltonian Monte Carlo*. 2014. arXiv: 1402.4102 [stat.ME].
- [21] S. L. Cotter et al. “MCMC Methods for Functions: Modifying Old Algorithms to Make Them Faster”. In: *Statistical Science* 28.3 (Aug. 2013). doi: 10.1214/13-sts421. url: <https://doi.org/10.1214%2F13-sts421>.
- [22] K. Craigie et al. “Resource-efficient adaptive Bayesian tracking of magnetic fields with a quantum sensor”. In: *Journal of Physics: Condensed Matter* 33.19 (Apr. 2021), p. 195801. doi: 10.1088/1361-648x/abe34f. url: <https://doi.org/10.1088%2F1361-648x%2Fabe34f>.
- [23] K.-D. Dang et al. “Hamiltonian Monte Carlo with Energy Conserving Subsampling”. In: *Journal of Machine Learning Research* 20.100 (2019), pp. 1–31. url: <http://jmlr.org/papers/v20/17-452.html>.

- [24] R. Daviet. “Inference with Hamiltonian Sequential Monte Carlo Simulators”. In: *SSRN Electronic Journal* (2016). doi: [10.2139/ssrn.2888242](https://doi.org/10.2139/ssrn.2888242). url: <https://doi.org/10.2139/2Fssrn.2888242>.
- [25] I. Deutsch. *Short Course in Quantum Information: Lecture 8*. 2005. url: [http://info.phys.unm.edu/~ideutsch/Courses/Sandia%20F05/QI\\_Lecture8.pdf](http://info.phys.unm.edu/~ideutsch/Courses/Sandia%20F05/QI_Lecture8.pdf) (visited on 09/26/2021).
- [26] M. Dobišček et al. “Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark”. In: *Physical Review A* 76.3 (Sept. 2007). doi: [10.1103/physreva.76.030306](https://doi.org/10.1103/physreva.76.030306). url: <https://doi.org/10.1103/2Fphysreva.76.030306>.
- [27] A. Doucet, N. Freitas, and N. Gordon. “An Introduction to Sequential Monte Carlo Methods”. In: *Sequential Monte Carlo Methods in Practice*. Springer New York, 2001, pp. 3–14. doi: [10.1007/978-1-4757-3437-9\\_1](https://doi.org/10.1007/978-1-4757-3437-9_1). url: [https://doi.org/10.1007/2F978-1-4757-3437-9\\_1](https://doi.org/10.1007/2F978-1-4757-3437-9_1).
- [28] A. Doucet and A. Johansen. “A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later”. In: *Handbook of Nonlinear Filtering* 12 (Jan. 2009).
- [29] S. Duane et al. “Hybrid Monte Carlo”. In: *Physics Letters B* 195.2 (1987), pp. 216–222. issn: 0370-2693. doi: [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X). url: <https://www.sciencedirect.com/science/article/pii/037026938791197X>.
- [30] A. W. F. Edwards. “Likelihood, Expanded Edition.” In: *Biometrics* 49.4 (Dec. 1993), p. 1285. doi: [10.2307/2532282](https://doi.org/10.2307/2532282). url: <https://doi.org/10.2307/2F2532282>.
- [31] W. Edwards, H. Lindman, and L. J. Savage. “Bayesian statistical inference for psychological research.” In: *Psychological Review* 70.3 (May 1963), pp. 193–242. doi: [10.1037/h0044139](https://doi.org/10.1037/h0044139). url: <https://doi.org/10.1037/2Fh0044139>.
- [32] T. J. Evans, R. Harper, and S. T. Flammia. *Scalable Bayesian Hamiltonian learning*. 2019. arXiv: [1912.07636](https://arxiv.org/abs/1912.07636) [quant-ph].
- [33] C. Ferrie, C. Granade, and D. Cory. “Adaptive Hamiltonian Estimation Using Bayesian Experimental Design”. In: *AIP Conference Proceedings* 1443 (Nov. 2011). doi: [10.1063/1.3703632](https://doi.org/10.1063/1.3703632).
- [34] C. Ferrie. “Quantum model averaging”. In: *New Journal of Physics* 16.9 (Sept. 2014), p. 093035. doi: [10.1088/1367-2630/16/9/093035](https://doi.org/10.1088/1367-2630/16/9/093035). url: <https://doi.org/10.1088/2F1367-2630/2F16/2F9/2F093035>.
- [35] C. Ferrie and C. E. Granade. “Likelihood-Free Methods for Quantum Parameter Estimation”. In: *Physical Review Letters* 112.13 (Apr. 2014). doi: [10.1103/physrevlett.112.130402](https://doi.org/10.1103/physrevlett.112.130402). url: <https://doi.org/10.1103/2Fphysrevlett.112.130402>.

- [36] C. Ferrie, C. E. Granade, and D. G. Cory. “How to best sample a periodic probability distribution, or on the accuracy of Hamiltonian finding strategies”. In: *Quantum Information Processing* 12.1 (Apr. 2012), pp. 611–623. doi: [10.1007/s11128-012-0407-6](https://doi.org/10.1007/s11128-012-0407-6). url: <https://doi.org/10.1007%2Fs11128-012-0407-6>.
- [37] R. P. Feynman. “Simulating Physics with Computers”. In: *Feynman and Computation*. CRC Press, 2002, pp. 133–153. doi: [10.1201/9780429500459-11](https://doi.org/10.1201/9780429500459-11). url: <https://doi.org/10.1201%2F9780429500459-11>.
- [38] L. J. Fiderer, J. Schuff, and D. Braun. “Neural-Network Heuristics for Adaptive Bayesian Quantum Estimation”. In: *PRX Quantum* 2.2 (Apr. 2021). doi: [10.1103/prxquantum.2.020303](https://doi.org/10.1103/prxquantum.2.020303). url: <https://doi.org/10.1103%2Fprxquantum.2.020303>.
- [39] D. A. Freedman. “On the Asymptotic Behavior of Bayes Estimates in the Discrete Case”. In: *The Annals of Mathematical Statistics* 34.4 (Dec. 1963), pp. 1386–1403. doi: [10.1214/aoms/1177703871](https://doi.org/10.1214/aoms/1177703871). url: <https://doi.org/10.1214/aoms/1177703871>.
- [40] D. A. Freedman. “On the Asymptotic Behavior of Bayes Estimates in the Discrete Case II”. In: *The Annals of Mathematical Statistics* 36.2 (Apr. 1965), pp. 454–456. doi: [10.1214/aoms/1177700155](https://doi.org/10.1214/aoms/1177700155). url: <https://doi.org/10.1214%2Faoms%2F1177700155>.
- [41] M. Frimmer and L. Novotny. “The classical Bloch equations”. In: *American Journal of Physics* 82.10 (Oct. 2014), pp. 947–954. doi: [10.1119/1.4878621](https://doi.org/10.1119/1.4878621). url: <https://doi.org/10.1119%2F1.4878621>.
- [42] A. Gelman, D. Lee, and J. Guo. “Stan”. In: *Journal of Educational and Behavioral Statistics* 40.5 (Oct. 2015), pp. 530–543. doi: [10.3102/1076998615606113](https://doi.org/10.3102/1076998615606113). url: <https://doi.org/10.3102%2F1076998615606113>.
- [43] A. Gelman and C. Pasarica. “Adaptively Scaling the Metropolis Algorithm Using Expected Squared Jumped Distance”. In: *SSRN Electronic Journal* (2007). doi: [10.2139/ssrn.1010403](https://doi.org/10.2139/ssrn.1010403). url: <https://doi.org/10.2139%2Fssrn.1010403>.
- [44] A. A. Gentile et al. “Learning models of quantum systems from experiments”. In: *Nature Physics* 17.7 (Apr. 2021), pp. 837–843. doi: [10.1038/s41567-021-01201-7](https://doi.org/10.1038/s41567-021-01201-7). url: <https://doi.org/10.1038%2Fs41567-021-01201-7>.
- [45] V. Giovannetti. “Quantum-Enhanced Measurements: Beating the Standard Quantum Limit”. In: *Science* 306.5700 (Nov. 2004), pp. 1330–1336. doi: [10.1126/science.1104149](https://doi.org/10.1126/science.1104149). url: <https://doi.org/10.1126%2Fscience.1104149>.
- [46] M. Girolami and B. Calderhead. “Riemann manifold Langevin and Hamiltonian Monte Carlo methods”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73.2 (Mar. 2011), pp. 123–214. doi: [10.1111/j.1467-9868.2010.00765.x](https://doi.org/10.1111/j.1467-9868.2010.00765.x). url: <https://doi.org/10.1111%2Fj.1467-9868.2010.00765.x>.

- [47] Google Colaboratory. 2021. url: <https://colab.research.google.com/>.
- [48] C. Granade, J. Combes, and D. G. Cory. “Practical Bayesian tomography”. In: 18.3 (Mar. 2016), p. 033024. doi: [10.1088/1367-2630/18/3/033024](https://doi.org/10.1088/1367-2630/18/3/033024). url: <https://doi.org/10.1088/1367-2630/18/3/033024>.
- [49] C. Granade and N. Wiebe. “Structured filtering”. In: *New Journal of Physics* 19.8 (Aug. 2017), p. 083014. doi: [10.1088/1367-2630/aa77cf](https://doi.org/10.1088/1367-2630/aa77cf). url: <https://doi.org/10.1088/1367-2630/aa77cf>.
- [50] C. Granade et al. *QInfer: Library for Statistical Inference in Quantum Information*. Sept. 2016. doi: [10.5281/zenodo.157007](https://doi.org/10.5281/zenodo.157007). url: <http://dx.doi.org/10.5281/zenodo.157007>.
- [51] C. E. Granade et al. “Robust online Hamiltonian learning”. In: *New Journal of Physics* 14.10 (Oct. 2012), p. 103013. doi: [10.1088/1367-2630/14/10/103013](https://doi.org/10.1088/1367-2630/14/10/103013). url: <https://doi.org/10.1088/1367-2630/14/10/103013>.
- [52] R. Gross, A. Marx, and F. Deppe. *Lecture notes on applied superconductivity, 10.3: control of quantum two-level systems*. 2013. url: [https://www.wmi.badw.de/teaching/Lecturenotes/AS/AS2013\\_Chapter10\\_2\\_Slides.pdf](https://www.wmi.badw.de/teaching/Lecturenotes/AS/AS2013_Chapter10_2_Slides.pdf) (visited on 05/08/2021).
- [53] M. Gu and S. Sun. *Variational Langevin Hamiltonian Monte Carlo for Distant Multi-modal Sampling*. 2019. arXiv: [1906.00229](https://arxiv.org/abs/1906.00229) [cs.LG].
- [54] D. Gunawan et al. “Subsampling sequential Monte Carlo for static Bayesian models”. In: *Statistics and Computing* 30.6 (Sept. 2020), pp. 1741–1758. doi: [10.1007/s11222-020-09969-z](https://doi.org/10.1007/s11222-020-09969-z). url: <https://doi.org/10.1007/s11222-020-09969-z>.
- [55] E. Hahn. “Spin Echoes”. In: *Phys. Rev.* 80 (Nov. 1950), pp. 580–594. doi: [10.1063/1.3066708](https://doi.org/10.1063/1.3066708).
- [56] W. K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1 (Apr. 1970), pp. 97–109. doi: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97). url: <https://doi.org/10.1093/biomet/57.1.97>.
- [57] B. L. Higgins et al. “Entanglement-free Heisenberg-limited phase estimation”. In: *Nature* 450.7168 (Nov. 2007), pp. 393–396. doi: [10.1038/nature06257](https://doi.org/10.1038/nature06257). url: <https://doi.org/10.1038/nature06257>.
- [58] I. Hincks et al. *Hamiltonian Learning with Online Bayesian Experiment Design in Practice*. 2018. arXiv: [1806.02427](https://arxiv.org/abs/1806.02427) [quant-ph].
- [59] M. Hoffman and A. Gelman. “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo”. In: *Journal of Machine Learning Research* 15 (Nov. 2011).
- [60] F. Huszár and N. M. T. Houlby. “Adaptive Bayesian quantum tomography”. In: *Phys. Rev. A* 85 (5 May 2012), p. 052120. doi: [10.1103/PhysRevA.85.052120](https://doi.org/10.1103/PhysRevA.85.052120). url: <https://link.aps.org/doi/10.1103/PhysRevA.85.052120>.

- [61] IBM Quantum. 2021. url: <https://quantum-computing.ibm.com/>.
- [62] IBM Quantum team. *Introduction to Quantum Computing and Quantum Hardware*. 2020.
- [63] IBM Quantum team. *Qiskit Summer School on Quantum Machine Learning*. 2021.
- [64] H. Jeffreys. *Theory of probability*. Oxford University Press, 1961. Chap. 5 Significance tests: one new parameter. isbn: 0-19-850368-7.
- [65] R. E. Kass and A. E. Raftery. “Bayes Factors”. In: *Journal of the American Statistical Association* 90.430 (June 1995), pp. 773–795. doi: [10.1080/01621459.1995.10476572](https://doi.org/10.1080/01621459.1995.10476572). url: <https://doi.org/10.1080%2F01621459.1995.10476572>.
- [66] T. Kim and B.-S. Choi. “Efficient decomposition methods for controlled-R n using a single ancillary qubit”. In: *Scientific Reports* 8.1 (Apr. 2018). doi: [10.1038/s41598-018-23764-x](https://doi.org/10.1038/s41598-018-23764-x). url: <https://doi.org/10.1038%2Fs41598-018-23764-x>.
- [67] A. Kitaev, A. Shen, and M. Vyalıy. *Classical and Quantum Computation*. American Mathematical Society, May 2002. doi: [10.1090/gsm/047](https://doi.org/10.1090/gsm/047). url: <https://doi.org/10.1090%2Fgsm%2F047>.
- [68] A. Y. Kitaev. “Quantum measurements and the abelian stabilizer problem”. In: *Electronic Colloquium on Computational Complexity* 3 (1996).
- [69] S. Lan, J. Streets, and B. Shahbaba. *Wormhole Hamiltonian Monte Carlo*. 2014. arXiv: [1306.0063](https://arxiv.org/abs/1306.0063) [stat.CO].
- [70] Z. Li et al. “Stochastic gradient Hamiltonian Monte Carlo with variance reduction for Bayesian inference”. In: *Machine Learning* 108.8-9 (July 2019), pp. 1701–1727. doi: [10.1007/s10994-019-05825-y](https://doi.org/10.1007/s10994-019-05825-y). url: <https://doi.org/10.1007%2Fs10994-019-05825-y>.
- [71] D. A. Lidar. *Lecture Notes on the Theory of Open Quantum Systems*. 2020. arXiv: [1902.00967](https://arxiv.org/abs/1902.00967) [quant-ph].
- [72] W. Link and M. Eaton. “On thinning of chains in MCMC”. In: *Methods in Ecology and Evolution* 3 (June 2011), pp. 112–115. doi: [10.1111/j.2041-210X.2011.00131.x](https://doi.org/10.1111/j.2041-210X.2011.00131.x).
- [73] J. Liu and M. West. “Combined Parameter and State Estimation in Simulation-Based Filtering”. In: *Sequential Monte Carlo Methods in Practice*. Springer New York, 2001, pp. 197–223. doi: [10.1007/978-1-4757-3437-9\\_10](https://doi.org/10.1007/978-1-4757-3437-9_10). url: [https://doi.org/10.1007%2F978-1-4757-3437-9\\_10](https://doi.org/10.1007%2F978-1-4757-3437-9_10).
- [74] S. Lloyd. “Universal Quantum Simulators”. In: *Science* 273.5278 (Aug. 1996), pp. 1073–1078. doi: [10.1126/science.273.5278.1073](https://doi.org/10.1126/science.273.5278.1073). url: <https://doi.org/10.1126%2Fscience.273.5278.1073>.
- [75] J. M. Lourenço. *The NOVAthesis L<sup>A</sup>T<sub>E</sub>X Template User’s Manual*. NOVA University Lisbon. 2021. url: <https://github.com/joaomlourenco/novathesis/raw/master/template.pdf>.

- [76] H.-H. Lu et al. “A controlled-NOT gate for frequency-bin qubits”. In: *npj Quantum Information* 5.1 (Mar. 2019). doi: [10.1038/s41534-019-0137-z](https://doi.org/10.1038/s41534-019-0137-z). url: <https://doi.org/10.1038/s41534-019-0137-z>.
- [77] J. M. Lukens et al. “A practical and efficient approach for Bayesian quantum state estimation”. In: *New Journal of Physics* 22.6 (July 2020), p. 063038. doi: [10.1088/1367-2630/ab8efa](https://doi.org/10.1088/1367-2630/ab8efa). url: <https://doi.org/10.1088/1367-2630/ab8efa>.
- [78] J. M. Lukens et al. “Computationally efficient Bayesian quantum state tomography”. In: *2020 IEEE Photonics Conference (IPC)*. IEEE, Sept. 2020. doi: [10.1109/ipc47351.2020.9252416](https://doi.org/10.1109/ipc47351.2020.9252416). url: <https://doi.org/10.1109/ipc47351.2020.9252416>.
- [79] A. Lumino et al. “Experimental Phase Estimation Enhanced by Machine Learning”. In: *Physical Review Applied* 10.4 (Oct. 2018). doi: [10.1103/physrevapplied.10.044033](https://doi.org/10.1103/physrevapplied.10.044033). url: <https://doi.org/10.1103/physrevapplied.10.044033>.
- [80] T. T. Mai. *Efficient adaptive MCMC implementation for Pseudo-Bayesian quantum tomography*. 2021. arXiv: [2106.00577](https://arxiv.org/abs/2106.00577) [stat.AP].
- [81] T. T. Mai and P. Alquier. “Pseudo-Bayesian quantum tomography with rank-adaptation”. In: *Journal of Statistical Planning and Inference* 184 (May 2017), pp. 62–76. doi: [10.1016/j.jspi.2016.11.003](https://doi.org/10.1016/j.jspi.2016.11.003). url: <https://doi.org/10.1016/j.jspi.2016.11.003>.
- [82] D. Manzano. “A short introduction to the Lindblad master equation”. In: *AIP Advances* 10.2 (Feb. 2020), p. 025106. doi: [10.1063/1.5115323](https://doi.org/10.1063/1.5115323). url: <https://doi.org/10.1063/1.5115323>.
- [83] L. Martino, V. Elvira, and F. Louzada. “Effective sample size for importance sampling based on discrepancy measures”. In: *Signal Processing* 131 (Feb. 2017), pp. 386–401. doi: [10.1016/j.sigpro.2016.08.025](https://doi.org/10.1016/j.sigpro.2016.08.025). url: <https://doi.org/10.1016/j.sigpro.2016.08.025>.
- [84] N. Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (June 1953), pp. 1087–1092. doi: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114).
- [85] P. del Moral. *An introduction to Feynman-Kac particle methods in statistical learning and rare event simulation*. 2014.
- [86] P. D. Moral. *Feynman-Kac Formulae - Genealogical and Interacting Particle Systems with Applications*. Springer New York, 2004. doi: [10.1007/978-1-4684-9393-1](https://doi.org/10.1007/978-1-4684-9393-1). url: <https://doi.org/10.1007/978-1-4684-9393-1>.
- [87] P. D. Moral, A. Doucet, and A. Jasra. “Sequential Monte Carlo samplers”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68.3 (June 2006), pp. 411–436. doi: [10.1111/j.1467-9868.2006.00553.x](https://doi.org/10.1111/j.1467-9868.2006.00553.x). url: <https://doi.org/10.1111/j.1467-9868.2006.00553.x>.



- [88] R. Neal. “MCMC Using Hamiltonian Dynamics”. In: *Chapman & Hall/CRC Handbooks of Modern Statistical Methods*. Chapman and Hall/CRC, May 2011. doi: [10.1201/b10905-6](https://doi.org/10.1201/b10905-6). url: <https://doi.org/10.1201%2Fb10905-6>.
- [89] R. M. Neal. “Annealed importance sampling”. In: *Statistics and Computing* 11.2 (Apr. 2001), pp. 125–139. issn: 1573-1375. doi: [10.1023/A:1008923215028](https://doi.org/10.1023/A:1008923215028). url: <https://doi.org/10.1023/A:1008923215028>.
- [90] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. doi: [10.1017/CB09780511976667](https://doi.org/10.1017/CB09780511976667).
- [91] A. Nishimura, D. B. Dunson, and J. Lu. “Discontinuous Hamiltonian Monte Carlo for discrete parameters and discontinuous likelihoods”. In: *Biometrika* 107.2 (Mar. 2020), pp. 365–380. doi: [10.1093/biomet/asz083](https://doi.org/10.1093/biomet/asz083). url: <https://doi.org/10.1093%2Fbiomet%2Fasz083>.
- [92] S. Paesani et al. “Experimental Bayesian Quantum Phase Estimation on a Silicon Photonic Chip”. In: *Physical Review Letters* 118.10 (Mar. 2017). doi: [10.1103/physrevlett.118.100503](https://doi.org/10.1103/physrevlett.118.100503). url: <https://doi.org/10.1103%2Fphysrevlett.118.100503>.
- [93] A. Peres, ed. *Quantum Theory: Concepts and Methods*. Springer Netherlands, 2002. doi: [10.1007/0-306-47120-5](https://doi.org/10.1007/0-306-47120-5). url: <https://doi.org/10.1007%2F0-306-47120-5>.
- [94] J. Preskill. *Lecture Notes for Ph219/CS219: Quantum Information, Chapter 2*. 2015. url: [http://theory.caltech.edu/~preskill/ph219/chap2\\_13.pdf](http://theory.caltech.edu/~preskill/ph219/chap2_13.pdf) (visited on 08/28/2021).
- [95] J. Preskill. *Lecture Notes for Ph219/CS219: Quantum Information, Chapter 3*. 2018. url: [http://theory.caltech.edu/~preskill/ph219/chap3\\_15.pdf](http://theory.caltech.edu/~preskill/ph219/chap3_15.pdf) (visited on 08/28/2021).
- [96] M. Quiroz et al. “Speeding up MCMC by Efficient Data Subsampling”. In: *Journal of the American Statistical Association* 114 (Mar. 2018), pp. 1–35. doi: [10.1080/01621459.2018.1448827](https://doi.org/10.1080/01621459.2018.1448827).
- [97] M. P. Rios and H. F. Lopes. “The Extended Liu and West Filter: Parameter Learning in Markov Switching Stochastic Volatility Models”. In: *State-Space Models*. Springer New York, 2013, pp. 23–61. doi: [10.1007/978-1-4614-7789-1\\_2](https://doi.org/10.1007/978-1-4614-7789-1_2). url: [https://doi.org/10.1007%2F978-1-4614-7789-1\\_2](https://doi.org/10.1007%2F978-1-4614-7789-1_2).
- [98] J. Salvatier, T. Wiecki, and C. Fonnesbeck. “Probabilistic programming in Python using PyMC3”. In: (Jan. 2016). doi: [10.7287/PEERJ.PREPRINTS.1686V1](https://doi.org/10.7287/PEERJ.PREPRINTS.1686V1).
- [99] R. Santagati et al. “Magnetic-Field Learning Using a Single Electronic Spin in Diamond with One-Photon Readout at Room Temperature”. In: *Phys. Rev. X* 9 (2 Apr. 2019), p. 021019. doi: [10.1103/PhysRevX.9.021019](https://link.aps.org/doi/10.1103/PhysRevX.9.021019). url: <https://link.aps.org/doi/10.1103/PhysRevX.9.021019>.
- [100] A. Sergeevich et al. “Characterization of a qubit Hamiltonian using adaptive measurements in a fixed basis”. In: *Physical Review A* 84.5 (Nov. 2011). doi: [10.1103/physreva.84.052315](https://doi.org/10.1103/physreva.84.052315). url: <https://doi.org/10.1103%2Fphysreva.84.052315>.

- [101] S. Sivaganesan. “Sensitivity of some posterior summaries when the prior is unimodal with specified quantiles”. In: *Canadian Journal of Statistics* 19.1 (Mar. 1991), pp. 57–65. doi: [10.2307/3315536](https://doi.org/10.2307/3315536). url: <https://doi.org/10.2307/3315536>.
- [102] L. F. South, A. N. Pettitt, and C. C. Drovandi. “Sequential Monte Carlo Samplers with Independent Markov Chain Monte Carlo Proposals”. In: *Bayesian Analysis* 14.3 (Sept. 2019). doi: [10.1214/18-ba1129](https://doi.org/10.1214/18-ba1129). url: <https://doi.org/10.1214/18-ba1129>.
- [103] J. S. Speagle. *A Conceptual Introduction to Markov Chain Monte Carlo Methods*. 2020. arXiv: [1909.12313](https://arxiv.org/abs/1909.12313) [stat.OT].
- [104] K. M. Svore, M. B. Hastings, and M. Freedman. “Faster phase estimation”. In: *Quantum Information and Computation* 14.3&4 (Mar. 2014), pp. 306–328. doi: [10.26421/qic14.3-4-7](https://doi.org/10.26421/qic14.3-4-7). url: <https://doi.org/10.26421/qic14.3-4-7>.
- [105] S. D. Team. *Stan Modeling Language Users Guide and Reference Manual 2.27*. 2021. url: <http://community.qiskit.org/textbook>.
- [106] D. Tong. *Lectures on Kinetic Theory*. 2012. url: <https://www.damtp.cam.ac.uk/user/tong/kintheory/three.pdf> (visited on 08/15/2021).
- [107] G. Tóth and I. Apellaniz. “Quantum metrology from a quantum information science perspective”. In: *Journal of Physics A: Mathematical and Theoretical* 47.42 (Oct. 2014), p. 424006. doi: [10.1088/1751-8113/47/42/424006](https://doi.org/10.1088/1751-8113/47/42/424006). url: <https://doi.org/10.1088/1751-8113/47/42/424006>.
- [108] M. -N. Tran et al. *The Block Pseudo-Marginal Sampler*. 2017. arXiv: [1603.02485](https://arxiv.org/abs/1603.02485) [stat.ME].
- [109] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, Oct. 1998. Chap. 10.2 Bernstein–von Mises Theorem. doi: [10.1017/cbo9780511802256](https://doi.org/10.1017/cbo9780511802256). url: <https://doi.org/10.1017/cbo9780511802256>.
- [110] J. Vandekerckhove, D. Matzke, and E.-J. Wagenmakers. *Model Comparison and the Principle of Parsimony*. Oxford University Press, Dec. 2015. doi: [10.1093/oxfordhb/9780199957996.013.14](https://doi.org/10.1093/oxfordhb/9780199957996.013.14). url: <https://doi.org/10.1093/oxfordhb/9780199957996.013.14>.
- [111] P. Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [112] N. V. Vitanov et al. “Fault-tolerant Hahn-Ramsey interferometry with pulse sequences of alternating detuning”. In: *Physical Review A* 91.3 (Mar. 2015). doi: [10.1103/physreva.91.033406](https://doi.org/10.1103/physreva.91.033406). url: <https://doi.org/10.1103/physreva.91.033406>.
- [113] J. Wang et al. “Experimental quantum Hamiltonian learning”. In: *Nature Physics* 13.6 (Mar. 2017), pp. 551–555. issn: 1745-2481. doi: [10.1038/nphys4074](https://doi.org/10.1038/nphys4074). url: <http://dx.doi.org/10.1038/nphys4074>.



- 
- [114] D. J. Warne, R. E. Baker, and M. J. Simpson. “A practical guide to pseudo-marginal methods for computational inference in systems biology”. In: *Journal of Theoretical Biology* 496 (July 2020), p. 110255. doi: [10.1016/j.jtbi.2020.110255](https://doi.org/10.1016/j.jtbi.2020.110255). url: <https://doi.org/10.1016%2Fj.jtbi.2020.110255>.
- [115] N. Wiebe and C. Granade. “Efficient Bayesian Phase Estimation”. In: *Physical Review Letters* 117.1 (June 2016). doi: [10.1103/physrevlett.117.010503](https://doi.org/10.1103/physrevlett.117.010503). url: <https://doi.org/10.1103%2Fphysrevlett.117.010503>.
- [116] N. Wiebe et al. “Hamiltonian Learning and Certification Using Quantum Resources”. In: *Physical Review Letters* 112.19 (May 2014). doi: [10.1103/physrevlett.112.190501](https://doi.org/10.1103/physrevlett.112.190501). url: <https://doi.org/10.1103%2Fphysrevlett.112.190501>.
- [117] N. Wiebe et al. “Quantum Hamiltonian learning using imperfect quantum resources”. In: *Physical Review A* 89.4 (Apr. 2014). doi: [10.1103/physreva.89.042314](https://doi.org/10.1103/physreva.89.042314). url: <https://doi.org/10.1103%2Fphysreva.89.042314>.
- [118] B. P. Williams and P. Lougovski. “Quantum state estimation when qubits are lost: a no-data-left-behind approach”. In: *New Journal of Physics* 19.4 (Apr. 2017), p. 043003. doi: [10.1088/1367-2630/aa65de](https://doi.org/10.1088/1367-2630/aa65de). url: <https://doi.org/10.1088%2F1367-2630%2Faa65de>.
- [119] B. Zwiebach. *Two state systems*. 2013. url: [https://ocw.mit.edu/courses/physics/8-05-quantum-physics-ii-fall-2013/lecture-notes/MIT8\\_05F13\\_Chap\\_07.pdf](https://ocw.mit.edu/courses/physics/8-05-quantum-physics-ii-fall-2013/lecture-notes/MIT8_05F13_Chap_07.pdf) (visited on 09/20/2021).
- [120] M. Zwiernik, C. A. Pérez-Delgado, and P. Kok. “General Optimality of the Heisenberg Limit for Quantum Metrology”. In: *Physical Review Letters* 105.18 (Oct. 2010). doi: [10.1103/physrevlett.105.180402](https://doi.org/10.1103/physrevlett.105.180402). url: <https://doi.org/10.1103%2Fphysrevlett.105.180402>.



