# Frank-Wolfe Bayesian Quadrature: Probabilistic Integration with Theoretical Guarantees 

François-Xavier Briol<br>Department of Statistics<br>University of Warwick<br>f-x.briol@warwick.ac.uk

Chris J. Oates<br>School of Mathematical and Physical Sciences<br>University of Technology, Sydney<br>christopher.oates@uts.edu.au

Mark Girolami<br>Department of Statistics<br>University of Warwick<br>\& The Alan Turing Institute for Data Science<br>m.girolami@warwick.ac.uk

Michael A. Osborne<br>Department of Engineering Science<br>University of Oxford<br>mosb@robots.ox.ac.uk


#### Abstract

There is renewed interest in formulating integration as a statistical inference problem, motivated by obtaining a full distribution over numerical error that can be propagated through subsequent computation. Current methods, such as Bayesian Quadrature, demonstrate impressive empirical performance but lack theoretical analysis. An important challenge is therefore to reconcile these probabilistic integrators with rigorous convergence guarantees. In this paper, we present the first probabilistic integrator that admits such theoretical treatment, called Frank-Wolfe Bayesian Quadrature (FWBQ). Under FWBQ, convergence to the true value of the integral is shown to be up to exponential and posterior contraction rates are proven to be up to super-exponential. In simulations, FWBQ is competitive with state-of-the-art methods and out-performs alternatives based on Frank-Wolfe optimisation. Our approach is applied to successfully quantify numerical error in the solution to a challenging Bayesian model choice problem in cellular biology.


## 1 Introduction

Computing integrals is a core challenge in machine learning and numerical methods play a central role in this area. This can be problematic when a numerical integration routine is repeatedly called, maybe millions of times, within a larger computational pipeline. In such situations, the cumulative impact of numerical errors can be unclear, especially in cases where the error has a non-trivial structural component. One solution is to model the numerical error statistically and to propagate this source of uncertainty through subsequent computations. Conversely, an understanding of how errors arise and propagate can enable the efficient focusing of computational resources upon the most challenging numerical integrals in a pipeline.
Classical numerical integration schemes do not account for prior information on the integrand and, as a consequence, can require an excessive number of function evaluations to obtain a prescribed level of accuracy [21]. Alternatives such as Quasi-Monte Carlo (QMC) can exploit knowledge on the smoothness of the integrand to obtain optimal convergence rates [7]. However these optimal rates can only hold on sub-sequences of sample sizes $n$, a consequence of the fact that all function evaluations are weighted equally in the estimator [24]. A modern approach that avoids this problem is to consider arbitrarily weighted combinations of function values; the so-called quadrature rules (also called cubature rules). Whilst quadrature rules with non-equal weights have received compar-
atively little theoretical attention, it is known that the extra flexibility given by arbitrary weights can lead to extremely accurate approximations in many settings (see applications to image de-noising [3] and mental simulation in psychology [13]).

Probabilistic numerics, introduced in the seminal paper of [6], aims at re-interpreting numerical tasks as inference tasks that are amenable to statistical analysis ${ }^{1}$ Recent developments include probabilistic solvers for linear systems [14] and differential equations [5, 28]. For the task of computing integrals, Bayesian Quadrature (BQ) [22] and more recent work by [20] provide probabilistic numerics methods that produce a full posterior distribution on the output of numerical schemes. One advantage of this approach is that we can propagate uncertainty through all subsequent computations to explicitly model the impact of numerical error [15]. Contrast this with chaining together classical error bounds; the result in such cases will typically be a weak bound that provides no insight into the error structure. At present, a significant shortcoming of these methods is the absence of theoretical results relating to rates of posterior contraction. This is unsatisfying and has likely hindered the adoption of probabilistic approaches to integration, since it is not clear that the induced posteriors represent a sensible quantification of the numerical error (by classical, frequentist standards).

This paper establishes convergence rates for a new probabilistic approach to integration. Our results thus overcome a key perceived weakness associated with probabilistic numerics in the quadrature setting. Our starting point is recent work by [2], who cast the design of quadrature rules as a problem in convex optimisation that can be solved using the Frank-Wolfe (FW) algorithm. We propose a hybrid approach of [2] with BQ, taking the form of a quadrature rule, that (i) carries a full probabilistic interpretation, (ii) is amenable to rigorous theoretical analysis, and (iii) converges orders-of-magnitude faster, empirically, compared with the original approaches in [2]. In particular, we prove that super-exponential rates hold for posterior contraction (concentration of the posterior probability mass on the true value of the integral), showing that the posterior distribution provides a sensible and effective quantification of the uncertainty arising from numerical error. The methodology is explored in simulations and also applied to a challenging model selection problem from cellular biology, where numerical error could lead to mis-allocation of expensive resources.

## 2 Background

### 2.1 Quadrature and Cubature Methods

Let $\mathcal{X} \subseteq \mathbb{R}^{d}$ be a measurable space such that $d \in \mathbb{N}_{+}$and consider a probability density $p(x)$ defined with respect to the Lebesgue measure on $\mathcal{X}$. This paper focuses on computing integrals of the form $\int f(x) p(x) \mathrm{d} x$ for a test function $f: \mathcal{X} \rightarrow \mathbb{R}$ where, for simplicity, we assume $f$ is square-integrable with respect to $p(x)$. A quadrature rule approximates such integrals as a weighted sum of function values at some design points $\left\{x_{i}\right\}_{i=1}^{n} \subset \mathcal{X}$ :

$$
\begin{equation*}
\int_{\mathcal{X}} f(x) p(x) \mathrm{d} x \approx \sum_{i=1}^{n} w_{i} f\left(x_{i}\right) \tag{1}
\end{equation*}
$$

Viewing integrals as projections, we write $p[f]$ for the left-hand side and $\hat{p}[f]$ for the right-hand side, where $\hat{p}=\sum_{i=1}^{n} w_{i} \delta\left(x_{i}\right)$ and $\delta\left(x_{i}\right)$ is a Dirac measure at $x_{i}$. Note that $\hat{p}$ may not be a probability distribution; in fact, weights $\left\{w_{i}\right\}_{i=1}^{n}$ do not have to sum to one or be non-negative. Quadrature rules can be extended to multivariate functions $f: \mathcal{X} \rightarrow \mathbb{R}^{d}$ by taking each component in turn.
There are many ways of choosing combinations $\left\{x_{i}, w_{i}\right\}_{i=1}^{n}$ in the literature. For example, taking weights to be $w_{i}=1 / n$ with points $\left\{x_{i}\right\}_{i=1}^{n}$ drawn independently from the probability distribution $p(x)$ recovers basic Monte Carlo integration. The case with weights $w_{i}=1 / n$, but with points chosen with respect to some specific (possibly deterministic) schemes includes kernel herding [4] and Quasi-Monte Carlo (QMC) [7]. In Bayesian Quadrature, the points $\left\{x_{i}\right\}_{i=1}^{n}$ are chosen to minimise a posterior variance, with weights $\left\{w_{i}\right\}_{i=1}^{n}$ arising from a posterior probability distribution.

Classical error analysis for quadrature rules is naturally couched in terms of minimising the worstcase estimation error. Let $\mathcal{H}$ be a Hilbert space of functions $f: \mathcal{X} \rightarrow \mathbb{R}$, equipped with the inner

[^0]product $\langle\cdot, \cdot\rangle_{\mathcal{H}}$ and associated norm $\|\cdot\|_{\mathcal{H}}$. We define the maximum mean discrepancy (MMD) as:
\[

$$
\begin{equation*}
\operatorname{MMD}\left(\left\{x_{i}, w_{i}\right\}_{i=1}^{n}\right):=\sup _{f \in \mathcal{H}:\|f\|_{\mathcal{H}}=1}|p[f]-\hat{p}[f]| . \tag{2}
\end{equation*}
$$

\]

The reader can refer to [29] for conditions on $\mathcal{H}$ that are needed for the existence of the MMD. The rate at which the MMD decreases with the number of samples $n$ is referred to as the 'convergence rate' of the quadrature rule. For Monte Carlo, the MMD decreases with the slow rate of $\mathcal{O}_{P}\left(n^{-1 / 2}\right)$ (where the subscript $P$ specifies that the convergence is in probability). Let $\mathcal{H}$ be a RKHS with reproducing kernel $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and denote the corresponding canonical feature map by $\Phi(x)=$ $k(\cdot, x)$, so that the mean element is given by $\mu_{p}(x)=p[\Phi(x)] \in \mathcal{H}$. Then, following [29]

$$
\begin{equation*}
\operatorname{MMD}\left(\left\{x_{i}, w_{i}\right\}_{i=1}^{n}\right)=\left\|\mu_{p}-\mu_{\hat{p}}\right\|_{\mathcal{H}} \tag{3}
\end{equation*}
$$

This shows that to obtain low integration error in the RKHS $\mathcal{H}$, one only needs to obtain a good approximation of its mean element $\mu_{p}$ (as $\forall f \in \mathcal{H}: p[f]=\left\langle f, \mu_{p}\right\rangle_{\mathcal{H}}$ ). Establishing theoretical results for such quadrature rules is an active area of research [1].

### 2.2 Bayesian Quadrature

Bayesian Quadrature (BQ) was originally introduced in [22] and later revisited by [11, 12] and [23]. The main idea is to place a functional prior on the integrand $f$, then update this prior through Bayes' theorem by conditioning on both samples $\left\{x_{i}\right\}_{i=1}^{n}$ and function evaluations at those sample points $\left\{\mathrm{f}_{i}\right\}_{i=1}^{n}$ where $\mathrm{f}_{i}=f\left(x_{i}\right)$. This induces a full posterior distribution over functions $f$ and hence over the value of the integral $p[f]$. The most common implementation assumes a Gaussian Process (GP) prior $f \sim \mathcal{G P}(0, k)$. A useful property motivating the use of GPs is that linear projection preserves normality, so that the posterior distribution for the integral $p[f]$ is also a Gaussian, characterised by its mean and covariance. A natural estimate of the integral $p[f]$ is given by the mean of this posterior distribution, which can be compactly written as

$$
\begin{equation*}
\hat{p}_{\mathrm{BQ}}[f]=\mathrm{z}^{T} K^{-1} \mathrm{f} \tag{4}
\end{equation*}
$$

where $\mathrm{z}_{i}=\mu_{p}\left(x_{i}\right)$ and $K_{i j}=k\left(x_{i}, x_{j}\right)$. Notice that this estimator takes the form of a quadrature rule with weights $\mathrm{w}^{\mathrm{BQ}}=\mathrm{z}^{T} K^{-1}$. Recently, [27] showed how specific choices of kernel and design points for BQ can recover classical quadrature rules. This begs the question of how to select design points $\left\{x_{i}\right\}_{i=1}^{n}$. A particularly natural approach aims to minimise the posterior uncertainty over the integral $p[f]$, which was shown in [16, Prop. 1] to equal:

$$
\begin{equation*}
v_{\mathrm{BQ}}\left(\left\{x_{i}\right\}_{i=1}^{n}\right)=p\left[\mu_{p}\right]-\mathrm{z}^{T} K^{-1} \mathrm{z}=\mathrm{MMD}^{2}\left(\left\{x_{i}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right) \tag{5}
\end{equation*}
$$

Thus, in the RKHS setting, minimising the posterior variance corresponds to minimising the worst case error of the quadrature rule. Below we refer to Optimal BQ (OBQ) as BQ coupled with design points $\left\{x_{i}^{\mathrm{OBQ}}\right\}_{i=1}^{n}$ chosen to globally minimise (5). We also call Sequential BQ (SBQ) the algorithm that greedily selects design points to give the greatest decrease in posterior variance at each iteration. OBQ will give improved results over SBQ, but cannot be implemented in general, whereas SBQ is comparatively straight-forward to implement. There are currently no theoretical results establishing the convergence of either BQ, OBQ or SBQ.
Remark: (5) is independent of observed function values f . As such, no active learning is possible in SBQ (i.e. surprising function values never cause a revision of a planned sampling schedule). This is not always the case: For example [12] approximately encodes non-negativity of $f$ into BQ which leads to a dependence on $f$ in the posterior variance. In this case sequential selection becomes an active strategy that outperforms batch selection in general.

### 2.3 Deriving Quadrature Rules via the Frank-Wolfe Algorithm

Despite the elegance of BQ , its convergence rates have not yet been rigorously established. In brief, this is because $\hat{p}_{\mathrm{BQ}}[f]$ is an orthogonal projection of $f$ onto the affine hull of $\left\{\Phi\left(x_{i}\right)\right\}_{i=1}^{n}$, rather than e.g. the convex hull. Standard results from the optimisation literature apply to bounded domains, but the affine hull is not bounded (i.e. the BQ weights can be arbitrarily large and possibly negative). Below we describe a solution to the problem of computing integrals recently proposed by [2], based on the FW algorithm, that restricts attention to the (bounded) convex hull of $\left\{\Phi\left(x_{i}\right)\right\}_{i=1}^{n}$.

```
Algorithm 1 The Frank-Wolfe (FW) and Frank-Wolfe with Line-Search (FWLS) Algorithms.
Require: function \(J\), initial state \(g_{1}=\bar{g}_{1} \in \mathcal{G}\) (and, for FW only: step-size sequence \(\left\{\rho_{i}\right\}_{i=1}^{n}\) ).
    for \(i=2, \ldots, n\) do
        Compute \(\bar{g}_{i}=\operatorname{argmin}_{g \in \mathcal{G}}\left\langle g,(D J)\left(g_{i-1}\right)\right\rangle_{\times}\)
        [For FWLS only, line search: \(\rho_{i}=\operatorname{argmin}_{\rho \in[0,1]} J\left((1-\rho) g_{i-1}+\rho \bar{g}_{i}\right)\) ]
        Update \(g_{i}=\left(1-\rho_{i}\right) g_{i-1}+\rho_{i} \bar{g}_{i}\)
    end for
```

The Frank-Wolfe (FW) algorithm (Alg. 1), also called the conditional gradient algorithm, is a convex optimization method introduced in [9]. It considers problems of the form $\min _{g \in \mathcal{G}} J(g)$ where the function $J: \mathcal{G} \rightarrow \mathbb{R}$ is convex and continuously differentiable. A particular case of interest in this paper will be when the domain $\mathcal{G}$ is a compact and convex space of functions, as recently investigated in [17]. These assumptions imply the existence of a solution to the optimization problem.

At each iteration $i$, the FW algorithm computes a linearisation of the objective function $J$ at the previous state $g_{i-1} \in \mathcal{G}$ along its gradient $(D J)\left(g_{i-1}\right)$ and selects an 'atom' $\bar{g}_{i} \in \mathcal{G}$ that minimises the inner product a state $g$ and $(D J)\left(g_{i-1}\right)$. The new state $g_{i} \in \mathcal{G}$ is then a convex combination of the previous state $g_{i-1}$ and of the atom $\bar{g}_{i}$. This convex combination depends on a step-size $\rho_{i}$ which is pre-determined and different versions of the algorithm may have different step-size sequences.

Our goal in quadrature is to approximate the mean element $\mu_{p}$. Recently [2] proposed to frame integration as a FW optimisation problem. Here, the domain $\mathcal{G} \subseteq \mathcal{H}$ is a space of functions and taking the objective function to be:

$$
\begin{equation*}
J(g)=\frac{1}{2}\left\|g-\mu_{p}\right\|_{\mathcal{H}}^{2} . \tag{6}
\end{equation*}
$$

This gives an approximation of the mean element and $J$ takes the form of half the posterior variance (or the $\mathrm{MMD}^{2}$ ). In this functional approximation setting, minimisation of $J$ is carried out over $\mathcal{G}=\mathcal{M}$, the marginal polytope of the RKHS $\mathcal{H}$. The marginal polytope $\mathcal{M}$ is defined as the closure of the convex hull of $\Phi(\mathcal{X})$, so that in particular $\mu_{p} \in \mathcal{M}$. Assuming as in [18] that $\Phi(x)$ is uniformly bounded in feature space (i.e. $\exists R>0: \forall x \in \mathcal{X},\|\Phi(x)\|_{\mathcal{H}} \leq R$ ), then $\mathcal{M}$ is a closed and bounded set and can be optimised over.

In order to define the algorithm rigorously in this case, we introduce the Fréchet derivative of $J$, denoted $D J$, such that for $\mathcal{H}^{*}$ being the dual space of $\mathcal{H}$, we have the unique map $D J: \mathcal{H} \rightarrow \mathcal{H}^{*}$ such that for each $g \in \mathcal{H},(D J)(g)$ is the function mapping $h \in \mathcal{H}$ to $(D J)(g)(h)=\langle g-\mu, h\rangle_{\mathcal{H}}$. We also introduce the bilinear map $\langle\cdot, \cdot\rangle_{\times}: \mathcal{H} \times \mathcal{H}^{*} \rightarrow \mathbb{R}$ which, for $F \in \mathcal{H}^{*}$ given by $F(g)=$ $\langle g, f\rangle_{\mathcal{H}}$, is the rule giving $\langle h, F\rangle_{\times}=\langle h, f\rangle_{\mathcal{H}}$.
A particular advantage of this method is that it leads to 'sparse' solutions which are linear combinations of the atoms $\left\{\bar{g}_{i}\right\}_{i=1}^{n}$ [2]. In particular this provides a weighted estimate for the mean element:

$$
\begin{equation*}
\hat{\mu}_{\mathrm{FW}}:=g_{n}=\sum_{i=1}^{n}\left(\prod_{j=i+1}^{n}\left(1-\rho_{j-1}\right) \rho_{i-1}\right) \bar{g}_{i}:=\sum_{i=1}^{n} w_{i}^{\mathrm{FW}} \bar{g}_{i}, \tag{7}
\end{equation*}
$$

where by default $\rho_{0}=1$ which leads to all $w_{i}^{\mathrm{FW}} \in[0,1]$ when $\rho_{i}=1 /(i+1)$. A typical sequence of approximations to the mean element is shown in Fig. 1 (left), demonstrating that the approximation quickly converges to the ground truth (in black). Since minimisation of a linear function can be restricted to extreme points of the domain, the atoms will be of the form $\bar{g}_{i}=\Phi\left(x_{i}^{\mathrm{FW}}\right)=k\left(\cdot, x_{i}^{\mathrm{FW}}\right)$ for some $x_{i}^{\mathrm{FW}} \in \mathcal{X}$. The minimisation in $g$ over $\mathcal{G}$ from step 2 in Algorithm 1 therefore becomes a minimisation in $x$ over $\mathcal{X}$ and this algorithm therefore provides us design points. In practice, at each iteration $i$, the FW algorithm hence selects a design point $x_{i}^{\mathrm{FW}} \in \mathcal{X}$ which induces an atom $\bar{g}_{i}$ and gives us an approximation of the mean element $\mu_{p}$. We denote by $\hat{\mu}_{\mathrm{FW}}$ this approximation after $n$ iterations. Using the reproducing property, we can show that the FW estimate is a quadrature rule:

$$
\begin{equation*}
\hat{p}_{\mathrm{FW}}[f]:=\left\langle f, \hat{\mu}_{\mathrm{FW}}\right\rangle_{\mathcal{H}}=\left\langle f, \sum_{i=1}^{n} w_{i}^{\mathrm{FW}} \bar{g}_{i}\right\rangle_{\mathcal{H}}=\sum_{i=1}^{n} w_{i}^{\mathrm{FW}}\left\langle f, k\left(\cdot, x_{i}^{\mathrm{FW}}\right)\right\rangle_{\mathcal{H}}=\sum_{i=1}^{n} w_{i}^{\mathrm{FW}} f\left(x_{i}^{\mathrm{FW}}\right) . \tag{8}
\end{equation*}
$$

The total computational cost for FW is $\mathcal{O}\left(n^{2}\right)$. An extension known as FW with Line Search (FWLS) uses a line-search method to find the optimal step size $\rho_{i}$ at each iteration (see Alg. 1).


Figure 1: Left: Approximations of the mean element $\mu_{p}$ using the FWLS algorithm, based on $n=$ $1,2,5,10,50$ design points (purple, blue, green, red and orange respectively). It is not possible to distinguish between approximation and ground truth when $n=50$. Right: Density of a mixture of 20 Gaussian distributions, displaying the first $n=25$ design points chosen by FW (red), FWLS (orange) and SBQ (green). Each method provides well-spaced design points in high-density regions. Most FW and FWLS design points overlap, partly explaining their similar performance in this case.

Once again, the approximation obtained by FWLS has a sparse expression as a convex combination of all the previously visited states and we obtain an associated quadrature rule. FWLS has theoretical convergence rates that can be stronger than standard versions of FW but has computational cost in $\mathcal{O}\left(n^{3}\right)$. The authors in [10] provide a survey of FW-based algorithms and their convergence rates under different regularity conditions on the objective function and domain of optimisation.
Remark: The FW design points $\left\{x_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}$ are generally not available in closed-form. We follow mainstream literature by selecting, at each iteration, the point that minimises the MMD over a finite collection of $M$ points, drawn i.i.d from $p(x)$. The authors in [18] proved that this approximation adds a $\mathcal{O}\left(M^{-1 / 4}\right)$ term to the MMD, so that theoretical results on FW convergence continue to apply provided that $M(n) \rightarrow \infty$ sufficiently quickly. Appendix A provides full details. In practice, one may also make use of a numerical optimisation scheme in order to select the points.

## 3 A Hybrid Approach: Frank-Wolfe Bayesian Quadrature

To combine the advantages of a probabilistic integrator with a formal convergence theory, we propose Frank-Wolfe Bayesian Quadrature (FWBQ). In FWBQ, we first select design points $\left\{x_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}$ using the FW algorithm. However, when computing the quadrature approximation, instead of using the usual FW weights $\left\{w_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}$ we use instead the weights $\left\{w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}$ provided by BQ. We denote this quadrature rule by $\hat{p}_{\text {FWBQ }}$ and also consider $\hat{p}_{\text {FWLSBQ }}$, which uses FWLS in place of FW. As we show below, these hybrid estimators (i) carry the Bayesian interpretation of Sec. 2.2 (ii) permit a rigorous theoretical analysis, and (iii) out-perform existing FW quadrature rules by orders of magnitude in simulations. FWBQ is hence ideally suited to probabilistic numerics applications.

For these theoretical results we assume that $f$ belongs to a finite-dimensional RKHS $\mathcal{H}$, in line with recent literature [2, 10, 17, 18]. We further assume that $\mathcal{X}$ is a compact subset of $\mathbb{R}^{d}$, that $p(x)>0$ $\forall x \in \mathcal{X}$ and that $k$ is continuous on $\mathcal{X} \times \mathcal{X}$. Under these hypotheses, Theorem 1 establishes consistency of the posterior mean, while Theorem 2 establishes contraction for the posterior distribution.
Theorem 1 (Consistency). The posterior mean $\hat{p}_{\text {FWBQ }}[f]$ converges to the true integral $p[f]$ at the following rates:

$$
\left|p[f]-\hat{p}_{\text {FWBQ }}[f]\right| \leq M M D\left(\left\{x_{i}, w_{i}\right\}_{i=1}^{n}\right) \leq\left\{\begin{array}{cl}
\frac{2 D^{2}}{R} n^{-1} & \text { for } F W B Q  \tag{9}\\
\sqrt{2} D \exp \left(-\frac{R^{2}}{2 D^{2}} n\right) & \text { for } F W L S B Q
\end{array}\right.
$$

where the FWBQ uses step-size $\rho_{i}=1 /(i+1), D \in(0, \infty)$ is the diameter of the marginal polytope $\mathcal{M}$ and $R \in(0, \infty)$ gives the radius of the smallest ball of center $\mu_{p}$ included in $\mathcal{M}$.

Note that all the proofs of this paper can be found in Appendix B. An immediate corollary of Theorem 1 is that FWLSBQ has an asymptotic error which is exponential in $n$ and is therefore superior to that of any QMC estimator [7]. This is not a contradiction - recall that QMC restricts attention to uniform weights, while FWLSBQ is able to propose arbitrary weightings. In addition we highlight a robustness property: Even when the assumptions of this section do not hold, one still obtains atleast a rate $\mathcal{O}_{P}\left(n^{-1 / 2}\right)$ for the posterior mean using either FWBQ or FWLSBQ [8].

Remark: The choice of kernel affects the convergence of the FWBQ method [15]. Clearly, we expect faster convergence if the function we are integrating is 'close' to the space of functions induced by our kernel. Indeed, the kernel specifies the geometry of the marginal polytope $\mathcal{M}$, that in turn directly influences the rate constant $R$ and $D$ associated with FW convex optimisation.

Consistency is only a stepping stone towards our main contribution which establishes posterior contraction rates for FWBQ. Posterior contraction is important as these results justify, for the first time, the probabilistic numerics approach to integration; that is, we show that the full posterior distribution is a sensible quantification (at least asymptotically) of numerical error in the integration routine:

Theorem 2 (Contraction). Let $S \subseteq \mathbb{R}$ be an open neighbourhood of the true integral $p[f]$ and let $\gamma=\inf _{r \in S^{C}}|r-p[f]|>0$. Then the posterior probability mass on $S^{c}=\mathbb{R} \backslash S$ vanishes at a rate:

$$
\operatorname{prob}\left(S^{c}\right) \leq\left\{\begin{array}{cl}
\frac{2 \sqrt{2} D^{2}}{\sqrt{\pi} R \gamma} n^{-1} \exp \left(-\frac{\gamma^{2} R^{2}}{8 D^{4}} n^{2}\right) & \text { for FWBQ }  \tag{10}\\
\frac{2 D}{\sqrt{\pi \gamma}} \exp \left(-\frac{R^{2}}{2 D^{2}} n-\frac{\gamma^{2}}{2 \sqrt{2} D} \exp \left(\frac{R^{2}}{2 D^{2}} n\right)\right) & \text { for FWLSBQ }
\end{array}\right.
$$

where the FWBQ uses step-size $\rho_{i}=1 /(i+1), D \in(0, \infty)$ is the diameter of the marginal polytope $\mathcal{M}$ and $R \in(0, \infty)$ gives the radius of the smallest ball of center $\mu_{p}$ included in $\mathcal{M}$.

The contraction rates are exponential for FWBQ and super-exponential for FWLBQ, and thus the two algorithms enjoy both a probabilistic interpretation and rigorous theoretical guarantees. A notable corollary is that OBQ enjoys the same rates as FWLSBQ, resolving a conjecture by Tony O'Hagan that OBQ converges exponentially [personal communication]:
Corollary. The consistency and contraction rates obtained for $F W L S B Q$ apply also to $O B Q$.

## 4 Experimental Results

### 4.1 Simulation Study

To facilitate the experiments in this paper we followed [1, 2, 11, 18] and employed an exponentiatedquadratic (EQ) kernel $k\left(x, x^{\prime}\right):=\lambda^{2} \exp \left(-1 / 2 \sigma^{2}\left\|x-x^{\prime}\right\|_{2}^{2}\right)$. This corresponds to an infinitedimensional RKHS, not covered by our theory; nevertheless, we note that all simulations are practically finite-dimensional due to rounding at machine precision. See Appendix E for a finitedimensional approximation using random Fourier features. EQ kernels are popular in the BQ literature as, when $p$ is a mixture of Gaussians, the mean element $\mu_{p}$ is analytically tractable (see Appendix C). Some other ( $p, k$ ) pairs that produce analytic mean elements are discussed in [1].

For this simulation study, we took $p(x)$ to be a 20-component mixture of 2D-Gaussian distributions. Monte Carlo (MC) is often used for such distributions but has a slow convergence rate in $\mathcal{O}_{P}\left(n^{-1 / 2}\right)$. FW and FWLS are known to converge more quickly and are in this sense preferable to MC [2]. In our simulations (Fig. 2, left), both our novel methods FWBQ and FWLSBQ decreased the MMD much faster than the FW/FWLS methods of [2]. Here, the same kernel hyper-parameters $(\lambda, \sigma)=(1,0.8)$ were employed for all methods to have a fair comparison. This suggests that the best quadrature rules correspond to elements outside the convex hull of $\left\{\Phi\left(x_{i}\right)\right\}_{i=1}^{n}$. Examples of those, including BQ, often assign negative weights to features (Fig. S1 right, Appendix D).

The principle advantage of our proposed methods is that they reconcile theoretical tractability with a fully probabilistic interpretation. For illustration, Fig. 2 (right) plots the posterior uncertainty due to numerical error for a typical integration problem based on this $p(x)$. In-depth empirical studies of such posteriors exist already in the literature and the reader is referred to [3, 13, 22] for details.

Beyond these theoretically tractable integrators, SBQ seems to give even better performance as $n$ increases. An intuitive explanation is that SBQ picks $\left\{x_{i}\right\}_{i=1}^{n}$ to minimise the MMD whereas



Figure 2: Simulation study. Left: Plot of the worst-case integration error squared ( $\mathrm{MMD}^{2}$ ). Both FWBQ and FWLSBQ are seen to outperform FW and FWLS, with SBQ performing best overall. Right: Integral estimates for FWLS and FWLSBQ for a function $f \in \mathcal{H}$. FWLS converges more slowly and provides only a point estimate for a given number of design points. In contrast, FWLSBQ converges faster and provides a full probability distribution over numerical error shown shaded in orange ( $68 \%$ and $95 \%$ credible intervals). Ground truth corresponds to the dotted black line.

FWBQ and FWLSBQ only minimise an approximation of the MMD (its linearisation along $D J$ ). In addition, the SBQ weights are optimal at each iteration, which is not true for FWBQ and FWLSBQ. We conjecture that Theorem 1 and 2 provide upper bounds on the rates of SBQ. This conjecture is partly supported by Fig. 1 (right), which shows that SBQ selects similar design points to FW/FWLS (but weights them optimally). Note also that both FWBQ and FWLSBQ give very similar result. This is not surprising as FWLS has no guarantees over FW in infinite-dimensional RKHS [17].

### 4.2 Quantifying Numerical Error in a Proteomic Model Selection Problem

A topical bioinformatics application that extends recent work by [19] is presented. The objective is to select among a set of candidate models $\left\{M_{i}\right\}_{i=1}^{m}$ for protein regulation. This choice is based on a dataset $\mathcal{D}$ of protein expression levels, in order to determine a 'most plausible' biological hypothesis for further experimental investigation. Each $M_{i}$ is specified by a vector of kinetic parameters $\theta_{i}$ (full details in Appendix D). Bayesian model selection requires that these parameters are integrated out against a prior $p\left(\theta_{i}\right)$ to obtain marginal likelihood terms $L\left(M_{i}\right)=\int p\left(\mathcal{D} \mid \theta_{i}\right) p\left(\theta_{i}\right) \mathrm{d} \theta_{i}$. Our focus here is on obtaining the maximum a posteriori (MAP) model $M_{j}$, defined as the maximiser of the posterior model probability $L\left(M_{j}\right) / \sum_{i=1}^{m} L\left(M_{i}\right)$ (where we have assumed a uniform prior over model space). Numerical error in the computation of each term $L\left(M_{i}\right)$, if unaccounted for, could cause us to return a model $M_{k}$ that is different from the true MAP estimate $M_{j}$ and lead to the mis-allocation of valuable experimental resources.

The problem is quickly exaggerated when the number $m$ of models increases, as there are more opportunities for one of the $L\left(M_{i}\right)$ terms to be 'too large' due to numerical error. In [19], the number $m$ of models was combinatorial in the number of protein kinases measured in a high-throughput assay (currently $\sim 10^{2}$ but in principle up to $\sim 10^{4}$ ). This led [19] to deploy substantial computing resources to ensure that numerical error in each estimate of $L\left(M_{i}\right)$ was individually controlled. Probabilistic numerics provides a more elegant and efficient solution: At any given stage, we have a fully probabilistic quantification of our uncertainty in each of the integrals $L\left(M_{i}\right)$, shown to be sensible both theoretically and empirically. This induces a full posterior distribution over numerical uncertainty in the location of the MAP estimate (i.e. 'Bayes all the way down'). As such we can determine, on-line, the precise point in the computational pipeline when numerical uncertainty near the MAP estimate becomes acceptably small, and cease further computation.

The FWBQ methodology was applied to one of the model selection tasks in [19]. In Fig. 3(left) we display posterior model probabilities for each of $m=352$ candidates models, where a low number ( $n=10$ ) of samples were used for each integral. (For display clarity only the first 50 models are shown.) In this low- $n$ regime, numerical error introduces a second level of uncertainty that we quantify by combining the FWBQ error models for all integrals in the computational pipeline; this is summarised by a box plot (rather than a single point) for each of the models (obtained by sampling - details in Appendix D). These box plots reveal that our estimated posterior model probabilities are


Figure 3: Quantifying numerical error in a model selection problem. FWBQ was used to model the numerical error of each integral $L\left(M_{i}\right)$ explicitly. For integration based on $n=10$ design points, FWBQ tells us that the computational estimate of the model posterior will be dominated by numerical error (left). When instead $n=100$ design points are used (right), uncertainty due to numerical error becomes much smaller (but not yet small enough to determine the MAP estimate).
completely dominated by numerical error. In contrast, when $n$ is increased through 50, 100 and 200 (Fig. 3. right and Fig. S2), the uncertainty due to numerical error becomes negligible. At $n=200$ we can conclude that model 26 is the true MAP estimate and further computations can be halted. Correctness of this result was confirmed using the more computationally intensive methods in [19].

In Appendix D we compared the relative performance of FWBQ, FWLSBQ and SBQ on this problem. Fig. S1 shows that the BQ weights reduced the MMD by orders of magnitude relative to FW and FWLS and that SBQ converged more quickly than both FWBQ and FWLSBQ.

## 5 Conclusions

This paper provides the first theoretical results for probabilistic integration, in the form of posterior contraction rates for FWBQ and FWLSBQ. This is an important step in the probabilistic numerics research programme [15] as it establishes a theoretical justification for using the posterior distribution as a model for the numerical integration error (which was previously assumed [11, 12, 20, 23, 27, e.g.]). The practical advantages conferred by a fully probabilistic error model were demonstrated on a model selection problem from proteomics, where sensitivity of an evaluation of the MAP estimate was modelled in terms of the error arising from repeated numerical integration.

The strengths and weaknesses of BQ (notably, including scalability in the dimension $d$ of $\mathcal{X}$ ) are well-known and are inherited by our FWBQ methodology. We do not review these here but refer the reader to [22] for an extended discussion. Convergence, in the classical sense, was proven here to occur exponentially quickly for FWLSBQ, which partially explains the excellent performance of BQ and related methods seen in applications [12, 23], as well as resolving an open conjecture. As a bonus, the hybrid quadrature rules that we developed turned out to converge much faster in simulations than those in [2], which originally motivated our work.
A key open problem for kernel methods in probabilistic numerics is to establish protocols for the practical elicitation of kernel hyper-parameters. This is important as hyper-parameters directly affect the scale of the posterior over numerical error that we ultimately aim to interpret. Note that this problem applies equally to BQ , as well as related quadrature methods [2, 11, 12, 20] and more generally in probabilistic numerics [28]. Previous work, such as [13], optimised hyper-parameters on a perapplication basis. Our ongoing research seeks automatic and general methods for hyper-parameter elicitation that provide good frequentist coverage properties for posterior credible intervals, but we reserve the details for a future publication.

## Acknowledgments

The authors are grateful for discussions with Simon Lacoste-Julien, Simo Särkkä, Arno Solin, Dino Sejdinovic, Tom Gunter and Mathias Cronjäger. FXB was supported by EPSRC [EP/L016710/1]. CJO was supported by EPSRC [EP/D002060/1]. MG was supported by EPSRC [EP/J016934/1], an EPSRC Established Career Fellowship, the EU grant [EU/259348] and a Royal Society Wolfson Research Merit Award.

## References

[1] F. Bach. On the Equivalence between Quadrature Rules and Random Features. arXiv:1502.06800, 2015.
[2] F. Bach, S. Lacoste-Julien, and G. Obozinski. On the Equivalence between Herding and Conditional Gradient Algorithms. In Proceedings of the 29th International Conference on Machine Learning, pages 1359-1366, 2012.
[3] Y. Chen, L. Bornn, N. de Freitas, M. Eskelin, J. Fang, and M. Welling. Herded Gibbs Sampling. Journal of Machine Learning Research, 2015. To appear.
[4] Y. Chen, M. Welling, and A. Smola. Super-Samples from Kernel Herding. In Proceedings of the Conference on Uncertainty in Artificial Intelligence, pages 109-116, 2010.
[5] P. Conrad, M. Girolami, S. Särkkä, A. Stuart, and K. Zygalakis. Probability Measures for Numerical Solutions of Differential Equations. arXiv:1506.04592, 2015.
[6] P. Diaconis. Bayesian Numerical Analysis. Statistical Decision Theory and Related Topics IV, pages 163-175, 1988.
[7] J. Dick and F. Pillichshammer. Digital Nets and Sequences - Discrepancy Theory and Quasi-Monte Carlo Integration. Cambridge University Press, 2010.
[8] J. C. Dunn. Convergence Rates for Conditional Gradient Sequences Generated by Implicit Step Length Rules. SIAM Journal on Control and Optimization, 18(5):473-487, 1980.
[9] M. Frank and P. Wolfe. An Algorithm for Quadratic Programming. Naval Research Logistics Quarterly, 3:95-110, 1956.
[10] D. Garber and E. Hazan. Faster Rates for the Frank-Wolfe Method over Strongly-Convex Sets. In Proceedings of the 32nd International Conference on Machine Learning, pages 541-549, 2015.
[11] Z. Ghahramani and C. Rasmussen. Bayesian Monte Carlo. In Advances in Neural Information Processing Systems, pages 489-496, 2003.
[12] T. Gunter, R. Garnett, M. Osborne, P. Hennig, and S. Roberts. Sampling for Inference in Probabilistic Models with Fast Bayesian Quadrature. In Advances in Neural Information Processing Systems, 2014.
[13] J.B. Hamrick and T.L. Griffiths. Mental Rotation as Bayesian Quadrature. In NIPS 2013 Workshop on Bayesian Optimization in Theory and Practice, 2013.
[14] P. Hennig. Probabilistic Interpretation of Linear Solvers. SIAM Journal on Optimization, 25:234-260, 2015.
[15] P. Hennig, M. Osborne, and M. Girolami. Probabilistic Numerics and Uncertainty in Computations. Proceedings of the Royal Society A, 471(2179), 2015.
[16] F. Huszar and D. Duvenaud. Optimally-Weighted Herding is Bayesian Quadrature. In Uncertainty in Artificial Intelligence, pages 377-385, 2012.
[17] M. Jaggi. Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization. In Proceedings of the 30th International Conference on Machine Learning, volume 28, pages 427-435, 2013.
[18] S. Lacoste-Julien, F. Lindsten, and F. Bach. Sequential Kernel Herding : Frank-Wolfe Optimization for Particle Filtering. In Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, pages 544-552, 2015.
[19] C.J. Oates, F. Dondelinger, N. Bayani, J. Korkola, J.W. Gray, and S. Mukherjee. Causal Network Inference using Biochemical Kinetics. Bioinformatics, 30(17):i468-i474, 2014.
[20] C.J. Oates, M. Girolami, and N. Chopin. Control Functionals for Monte Carlo Integration. arXiv: 1410.2392, 2015.
[21] A. O'Hagan. Monte Carlo is Fundamentally Unsound. Journal of the Royal Statistical Society, Series D, 36(2):247-249, 1984.
[22] A. O’Hagan. Bayes-Hermite Quadrature. Journal of Statistical Planning and Inference, 29:245-260, 1991.
[23] M. Osborne, R. Garnett, S. Roberts, C. Hart, S. Aigrain, and N. Gibson. Bayesian Quadrature for Ratios. In Proceedings of the 15th International Conference on Artificial Intelligence and Statistics, pages 832840, 2012.
[24] A. B. Owen. A Constraint on Extensible Quadrature Rules. Numerische Mathematik, pages 1-8, 2015.
[25] A. Rahimi and B. Recht. Random Features for Large-Scale Kernel Machines. Advances in Neural Information Processings Systems, 2007.
[26] C. Rasmussen and C.K.I. Williams. Gaussian Processes for Machine Learning. MIT Press, 2006.
[27] S. Särkkä, J. Hartikainen, L. Svensson, and F. Sandblom. On the Relation between Gaussian Process Quadratures and Sigma-Point Methods. arXiv:1504.05994, 2015.
[28] M. Schober, D. Duvenaud, and P. Hennig. Probabilistic ODE solvers with Runge-Kutta means. In Advances in Neural Information Processing Systems 27, pages 739-747, 2014.
[29] B. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. Lanckriet. Hilbert Space Embeddings and Metrics on Probability Measures. Journal of Machine Learning Research, 11:1517-1561, 2010.
[30] R.A. Weinberg. The Biology of Cancer, volume 1. Garland Science, 2006.

## Supplementary Material

## Appendix A: Details for the FWBQ and FWLSBQ Algorithms

A high-level pseudo-code description for the Frank-Wolfe Bayesian Quadrature (FWBQ) algorithm is provided below.

```
Algorithm 2 The Frank-Wolfe Bayesian Quadrature (FWBQ) Algorithm
Require: function \(f\), reproducing kernel \(k\), initial point \(x_{0} \in \mathcal{X}\).
    Compute design points \(\left\{x_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}\) using the FW algorithm (Alg. 1).
    Compute associated weights \(\left\{w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\) using BQ (Eqn. 4).
    Compute the posterior mean \(\hat{p}_{\mathrm{FWBQ}}[f]\), i.e. the quadrature rule with \(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\).
    Compute the posterior variance \(v_{\mathrm{BQ}}\left(\left\{x_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}\right)\) using BQ (Eqn. 5).
    Return the full posterior \(\mathcal{N}\left(\hat{p}_{\mathrm{FWBQ}}, v_{\mathrm{BQ}}\left(\left\{x_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}\right)\right)\) for the integral \(p[f]\).
```

Frank-Wolfe Line-Search Bayesian Quadrature (FWLSBQ) is simply obtained by substituting the Frank-Wolfe algorithm with the Frank-Wolfe Line-Search algorithm. In this appendix, we derive all of the expressions necessary to implement both the FW and FWLS algorithms (for quadrature) in practice. All of the other steps can be derived from the relevant equations as highlighted in Algorithm 2 above.

The FW/FWLS are both initialised by the user choosing a design point $x_{1}^{\mathrm{FW}}$. This can be done either at random or by choosing a location which is known to have high probability mass under $p(x)$. The first approximation to $\mu_{p}$ is therefore given by $g_{1}=k\left(\cdot, x_{1}^{\mathrm{FW}}\right)$. The algorithm then loops over the next three steps to obtain new design points $\left\{x_{i}^{\mathrm{FW}}\right\}_{i=2}^{n}$ :
Step 1) Obtaining the new Frank-Wolfe design points $x_{i+1}^{F W}$.
At iteration $i$, the step consists of choosing the point $\bar{x}_{i}^{\mathrm{FW}}$. Let $\left\{w_{l}^{(i)}\right\}_{l=1}^{i-1}$ denote the Frank-Wolfe weights assigned to each of the previous design points $\left\{x_{l}^{\mathrm{FW}}\right\}_{l=1}^{i-1}$ at this new iteration, given that we choose $x$ as our new design point. The choice of new design point is done by computing the derivative of the objective function $J\left(g_{i-1}\right)$ and finding the point $x^{*}$ which minimises the inner product:

$$
\begin{equation*}
\arg \min _{g \in \mathcal{G}}\left\langle g,(D J)\left(g_{i-1}\right)\right\rangle_{\times} \tag{11}
\end{equation*}
$$

To do so, we need to obtain an equivalent expression of the minimisation of the linearisation of $J$ (denoted $D J$ ) in terms of kernel values and evaluations of the mean element $\mu_{p}$. Since minimisation of a linear function can be restricted to extreme points of the domain, we have that

$$
\begin{equation*}
\arg \min _{g \in \mathcal{G}}\left\langle g,(D J)\left(g_{i-1}\right)\right\rangle_{\times}=\arg \min _{x \in \mathcal{X}}\left\langle\Phi(x),(D J)\left(g_{i-1}\right)\right\rangle_{\times} \tag{12}
\end{equation*}
$$

Then using the definition of $J$ we have:

$$
\begin{equation*}
\arg \min _{x \in \mathcal{X}}\left\langle\Phi(x),(D J)\left(g_{i-1}\right)\right\rangle_{\times}=\arg \min _{x \in \mathcal{X}}\left\langle\Phi(x), g_{i-1}-\mu_{p}\right\rangle_{\mathcal{H}} \tag{13}
\end{equation*}
$$

where

$$
\begin{align*}
\left\langle\Phi(x), g_{i-1}-\mu_{p}\right\rangle_{\mathcal{H}} & =\left\langle\Phi(x), \sum_{l=1}^{i-1} w_{l}^{(i-1)} \Phi\left(x_{l}\right)-\mu_{p}\right\rangle_{\mathcal{H}} \\
& =\sum_{i=1}^{i-1} w_{l}^{(i-1)}\left\langle\Phi(x), \Phi\left(x_{l}\right)\right\rangle_{\mathcal{H}}-\left\langle\Phi(x), \mu_{p}\right\rangle_{\mathcal{H}}  \tag{14}\\
& =\sum_{l=1}^{i-1} w_{l}^{(i-1)} k\left(x, x_{l}\right)-\mu_{p}(x) .
\end{align*}
$$

Our new design point $x_{i}^{\mathrm{FW}}$ is therefore the point $x^{*}$ which minimises this expression. Note that this equation may not be convex and may require us to make use of approximate methods to find the
minimum $x^{*}$. To do so, we sample $M$ points (where $M$ is large) independently from the distribution $p$ and pick the sample which minimises the expression above. From [18] this introduces an additive error term of size $\mathcal{O}\left(M^{-1 / 4}\right)$, which does not impact our convergence analysis provided that $\mathrm{M}(\mathrm{n})$ vanishes sufficiently quickly. In all experiments we took $M$ between 10,000 and 50,000 so that this error will be negligible.
It is important to note that sampling from $p(x)$ is likely to not be the best solution to optimising this expression. One may, for example, be better off using any other optimisation method which does not require convexity (for example, Bayesian Optimization). However, we have used sampling as the result from [18] discussed above allows us to have a theoretical upper bound on the error introduced.

Step 2) Computing the Step-Sizes and Weights for the Frank-Wolfe and Frank-Wolfe Line-Search Algorithms.
Computing the weights $\left\{w_{l}^{(i)}\right\}_{l=1}^{n}$ assigned by the FW/FWLS algorithms to each of the design points is obtained using the equation:

$$
\begin{equation*}
w_{l}^{(i)}=\prod_{j=l+1}^{i}\left(1-\rho_{j-1}\right) \rho_{l-1} \tag{15}
\end{equation*}
$$

Clearly, this expression depends on the choice of step-sizes $\left\{\rho_{l}\right\}_{l=1}^{i}$. In the case of the standard Frank-Wolfe algorithm, this step-size sequence is a an input from the algorithm and so computing the weights is straightforward. However, in the case of the Frank-Wolfe Line-Search algorithm, the choice of step-size is optimized at each iteration so that $g_{i}$ minimises $J$ the most.

In the case of computing integrals, this optimization step can actually be obtained analytically. This analytic expression will be given in terms of values of the kernel values and evaluations of the mean element.

First, from the definition of $J$

$$
\begin{align*}
J\left((1-\rho) g_{i-1}+\rho \Phi\left(x_{i}\right)\right)= & \frac{1}{2}\left\langle(1-\rho) g_{i-1}+\rho \Phi\left(x_{i}\right)-\mu_{p},(1-\rho) g_{i-1}+\rho \Phi\left(x_{i}\right)-\mu_{p}\right\rangle_{\mathcal{H}} \\
= & \frac{1}{2}\left[(1-\rho)^{2}\left\langle g_{i-1}, g_{i-1}\right\rangle_{\mathcal{H}}+2(1-\rho) \rho\left\langle g_{i-1}, \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}\right. \\
& +2 \rho^{2}\left\langle\Phi\left(x_{i}\right), \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}-2(1-\rho)\left\langle g_{i-1}, \mu_{p}\right\rangle_{\mathcal{H}} \\
& \left.-2 \rho\left\langle\Phi\left(x_{i}\right), \mu_{p}\right\rangle_{\mathcal{H}}+\left\langle\mu_{p}, \mu_{p}\right\rangle_{\mathcal{H}}\right] \tag{16}
\end{align*}
$$

Taking the derivative of this expression with respect to $\rho$, we get:

$$
\begin{align*}
\frac{\partial J\left((1-\rho) g_{i-1}+\rho \Phi\left(x_{i}\right)\right)}{\partial \rho}= & \frac{1}{2}\left[-2(1-\rho)\left\langle g_{i-1}, g_{i-1}\right\rangle_{\mathcal{H}}+2(1-2 \rho)\left\langle g_{i-1}, \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}\right. \\
& \left.+2 \rho\left\langle\Phi\left(x_{i}\right), \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}+2\left\langle g_{i-1}, \mu_{p}\right\rangle_{\mathcal{H}}-2\left\langle\Phi\left(x_{i}\right), \mu_{p}\right\rangle_{\mathcal{H}}\right] \\
= & \rho\left[\left\langle g_{i-1}, g_{i-1}\right\rangle_{\mathcal{H}}-2\left\langle g_{i-1}, \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}+\left\langle\Phi\left(x_{i}\right), \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}\right. \\
= & \rho\left\|g_{i-1}-\Phi\left(x_{i}\right)\right\|_{\mathcal{H}}^{2}-\left\langle g_{i-1}-\Phi\left(x_{i}\right), g_{i-1}-\mu_{p}\right\rangle_{\mathcal{H}} \tag{17}
\end{align*}
$$

Setting this derivative to zero gives us the following optimum:

$$
\begin{equation*}
\rho^{*}=\frac{\left\langle g_{i-1}-\mu_{p}, g_{i-1}-\Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}}{\left\|g_{i-1}-\Phi\left(x_{i}\right)\right\|_{\mathcal{H}}^{2}} \tag{18}
\end{equation*}
$$

Clearly, differentiating a second time with respect to $\rho$ gives $\left\|g_{i-1}-\Phi\left(x_{i}\right)\right\|_{\mathcal{H}}^{2}$, which is non-negative and so $\rho^{*}$ is a minimum. One can show using geometrical arguments about the marginal polytope $\mathcal{M}$ that $\rho^{*}$ will be in $[0,1]$ [17].

The numerator of this line-search expression is

$$
\begin{align*}
\left\langle g_{i-1}-\mu_{p}, g_{i-1}-\Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}= & \left\langle g_{i-1}, g_{i-1}\right\rangle_{\mathcal{H}}-\left\langle\mu_{p}, g_{i-1}\right\rangle_{\mathcal{H}} \\
& -\sum_{l=1}^{i-1} w_{l}^{(i-1)} k\left(x_{l}, x_{i}\right)+\mu_{p}\left(x_{i}\right) \\
= & \sum_{l=1}^{i-1} \sum_{m=1}^{i-1} w_{l}^{(i-1)} w_{m}^{(i-1)} k\left(x_{l}, x_{m}\right)  \tag{19}\\
& -\sum_{l=1}^{i-1} w_{l}^{(i-1)}\left[k\left(x_{l}, x_{i}\right)+\mu_{p}\left(x_{l}\right)\right]+\mu_{p}\left(x_{i}\right) .
\end{align*}
$$

Similarly the denominator is

$$
\begin{align*}
\left\|g_{i-1}-\Phi\left(x_{i}\right)\right\|_{\mathcal{H}}^{2} & =\left\langle g_{i-1}-\Phi\left(x_{i}\right), g_{i-1}-\Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}} \\
& =\left\langle g_{i-1}, g_{i-1}\right\rangle_{\mathcal{H}}-2\left\langle g_{i-1}, \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}}+\left\langle\Phi\left(x_{i}\right), \Phi\left(x_{i}\right)\right\rangle_{\mathcal{H}} \\
& =\sum_{l=1}^{i-1} \sum_{m=1}^{i-1} w_{l}^{(i-1)} w_{m}^{(i-1)} k\left(x_{l}, x_{m}\right)-2 \sum_{l=1}^{i-1} w_{l}^{(i-1)} k\left(x_{l}, x_{i}\right)+k\left(x_{i}, x_{i}\right) \tag{20}
\end{align*}
$$

Clearly all expressions provided here can be vectorised for efficient computational implementation. Step 3) Computing a new approximation of the mean element.

The final step consists of updating the approximation of the mean element, which can be done directly by setting:

$$
\begin{equation*}
g_{i}=\left(1-\rho_{i}\right) g_{i-1}+\rho_{i} \bar{g}_{i} \tag{21}
\end{equation*}
$$

## Appendix B: Proofs of Theorems and Corollaries

Theorem (Consistency). The posterior mean $\hat{p}_{\mathrm{FWBQ}}[f]$ converges to the true integral $p[f]$ at the following rates:

$$
\left|p[f]-\hat{p}_{\text {FWBQ }}[f]\right| \leq M M D\left(\left\{x_{i}, w_{i}\right\}_{i=1}^{n}\right) \leq\left\{\begin{array}{cl}
\frac{2 D^{2}}{R} n^{-1} & \text { for } F W B Q \\
\sqrt{2} D \exp \left(-\frac{R^{2}}{2 D^{2}} n\right) & \text { for } F W L S B Q
\end{array}\right.
$$

where the FWBQ uses step-size $\rho_{i}=1 /(i+1), D \in(0, \infty)$ is the diameter of the marginal polytope $\mathcal{M}$ and $R \in(0, \infty)$ gives the radius of the smallest ball of center $\mu_{p}$ included in $\mathcal{M}$.
Proof. The posterior mean in BQ is a Bayes estimator and so the MMD takes a minimax form [16]. In particular, the BQ weights perform no worse than the FW weights:

$$
\begin{equation*}
\operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right)=\inf _{\mathrm{w} \in \mathbb{R}^{n}} \operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}\right\}_{i=1}^{n}\right) \leq \operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}\right) \tag{22}
\end{equation*}
$$

Now, the values attained by the objective function $J$ along the path $\left\{g_{i}\right\}_{i=1}^{n}$ determined by the FW(/FWLS) algorithm can be expressed in terms of the MMD as follows:

$$
\begin{equation*}
J\left(g_{n}\right)=\frac{1}{2}\left\|\hat{\mu}_{\mathrm{FW}}-\mu_{p}\right\|_{\mathcal{H}}^{2}=\frac{1}{2} \mathrm{MMD}^{2}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{FW}}\right\}_{i=1}^{n}\right) . \tag{23}
\end{equation*}
$$

Combining (22) and (23) gives

$$
\begin{equation*}
\left|p[f]-\hat{p}_{\mathrm{FWBQ}}[f]\right| \leq \operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right)\|f\|_{\mathcal{H}} \leq 2^{1 / 2} J^{1 / 2}\left(g_{n}\right) \tag{24}
\end{equation*}
$$

since $\|f\|_{\mathcal{H}} \leq 1$. To complete the proof we leverage recent analysis of the FW algorithm with steps $\rho_{i}=1 /(n+1)$ and the FWLS algorithm. Specifically, from [2] Prop. 1] we have that:

$$
J\left(g_{n}\right) \leq\left\{\begin{array}{cl}
\frac{2 D^{4}}{R^{2}} n^{-2} & \text { for FW with step size } \rho_{i}=1 /(i+1)  \tag{25}\\
D^{2} \exp \left(-R^{2} n / D^{2}\right) & \text { for FWLS }
\end{array}\right.
$$

where $D$ is the diameter of the marginal polytope $\mathcal{M}$ and $R$ is the radius of the smallest ball centered at $\mu_{p}$ included in $\mathcal{M}$.

Theorem (Contraction). Let $S \subseteq \mathbb{R}$ be an open neighbourhood of the true integral $p[f]$ and let $\gamma=\inf _{r \in S^{C}}|r-p[f]|>0$. Then the posterior probability mass on $S^{c}=\mathbb{R} \backslash S$ vanishes at a rate:

$$
\operatorname{prob}\left(S^{c}\right) \leq\left\{\begin{array}{cl}
\frac{2 \sqrt{2} D^{2}}{\sqrt{\pi} R \gamma} n^{-1} \exp \left(-\frac{\gamma^{2} R^{2}}{8 D^{4}} n^{2}\right) & \text { for } F W B Q, \rho_{i}=1 /(i+1) \\
\frac{2 D}{\sqrt{\pi} \gamma} \exp \left(-\frac{R^{2}}{2 D^{2}} n-\frac{\gamma^{2}}{2 \sqrt{2} D} \exp \left(\frac{R^{2}}{2 D^{2}} n\right)\right) & \text { for } F W L S B Q
\end{array}\right.
$$

where $D \in(0, \infty)$ is the diameter of the marginal polytope $\mathcal{M}$ and $R \in(0, \infty)$ gives the radius of the smallest ball of center $\mu_{p}$ included in $\mathcal{M}$.
Proof. We will obtain the posterior contraction rates of interest using the bounds on the MMD provided in the proof of Theorem 1. Given an open neighbourhood $S \subseteq \mathbb{R}$ of $p[f]$, we have that the complement $S^{c}=\mathbb{R} \backslash S$ is closed in $\mathbb{R}$. We assume without loss of generality that $S^{c} \neq \emptyset$, since the posterior mass on $S^{c}$ is trivially zero when $S^{c}=\emptyset$. Since $S^{c}$ is closed, the distance $\gamma=$ $\inf _{r \in S^{c}}|r-p[f]|>0$ is strictly positive. Denote the posterior distribution by $\mathcal{N}\left(m_{n}, \sigma_{n}^{2}\right)$ where we have that $m_{n}:=\hat{p}_{\mathrm{FWBQ}}[f]$ where $\hat{p}_{\mathrm{FWBQ}}=\sum_{i=1}^{n} w_{i}^{\mathrm{BQ}} \delta\left(x_{i}^{\mathrm{FW}}\right)$ and $\sigma_{n}:=\operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right)$. Directly from the supremum definition of the MMD we have:

$$
\begin{equation*}
\left|p[f]-m_{n}\right| \leq \sigma_{n}\|f\|_{\mathcal{H}} \tag{26}
\end{equation*}
$$

Now the posterior probability mass on $S^{c}$ is given by

$$
\begin{equation*}
M_{n}=\int_{S^{c}} \phi\left(r \mid m_{n}, \sigma_{n}\right) \mathrm{d} r \tag{27}
\end{equation*}
$$

where $\phi\left(r \mid m_{n}, \sigma_{n}\right)$ is the p.d.f. of the posterior normal distribution. By the definition of $\gamma$ we get the upper bound:

$$
\begin{align*}
M_{n} & \leq \int_{-\infty}^{p[f]-\gamma} \phi\left(r \mid m_{n}, \sigma_{n}\right) \mathrm{d} r+\int_{p[f]+\gamma}^{\infty} \phi\left(r \mid m_{n}, \sigma_{n}\right) \mathrm{d} r  \tag{28}\\
& =1+\Phi(\underbrace{\frac{p[f]-m_{n}}{\sigma_{n}}}_{(*)}-\frac{\gamma}{\sigma_{n}})-\Phi(\underbrace{\frac{p[f]-m_{n}}{\sigma_{n}}}_{(*)}+\frac{\gamma}{\sigma_{n}}) \tag{29}
\end{align*}
$$

From (26) we have that the terms $(*)$ are bounded by $\|f\|_{\mathcal{H}} \leq 1<\infty$ as $\sigma_{n} \rightarrow 0$, so that asymptotically we have:

$$
\begin{align*}
M_{n} & \lesssim 1+\Phi\left(-\gamma / \sigma_{n}\right)-\Phi\left(\gamma / \sigma_{n}\right)  \tag{30}\\
& =\operatorname{erfc}\left(\gamma / \sqrt{2} \sigma_{n}\right) \sim\left(\sqrt{2} \sigma_{n} / \sqrt{\pi} \gamma\right) \exp \left(-\gamma^{2} / 2 \sigma_{n}^{2}\right) \tag{31}
\end{align*}
$$

Finally we may substitute the asymptotic results derived in the proof of Theorem 1 for the MMD $\sigma_{n}$ into (31) to complete the proof.

Corollary. The consistency and contraction rates obtained for FWLSBQ apply also to OBQ.
Proof. By definition, OBQ chooses samples that globally minimise the MMD and we can hence bound this quantity from above by the MMD of FWLSBQ:

$$
\begin{equation*}
\operatorname{MMD}\left(\left\{x_{i}^{\mathrm{OBQ}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right)=\inf _{\left\{x_{i}\right\}_{i=1}^{n} \in \mathcal{X}} \operatorname{MMD}\left(\left\{x_{i}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right) \leq \operatorname{MMD}\left(\left\{x_{i}^{\mathrm{FW}}, w_{i}^{\mathrm{BQ}}\right\}_{i=1}^{n}\right) . \tag{32}
\end{equation*}
$$

Consistency and contraction follow from inserting this inequality into the above proofs.

## Appendix C: Computing the Mean Element for the Simulation Study

We compute an expression for $\mu_{p}(x)=\int_{-\infty}^{\infty} k\left(x, x^{\prime}\right) p\left(x^{\prime}\right) \mathrm{d} x^{\prime}$ in the case where $k$ is an exponentiated-quadratic kernel with length scale hyper-parameter $\sigma$ :

$$
\begin{equation*}
k\left(x, x^{\prime}\right):=\lambda^{2} \exp \left(\frac{-\sum_{i=1}^{d}\left(x_{i}-x_{i}^{\prime}\right)^{2}}{2 \sigma^{2}}\right)=\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \phi\left(x \mid x^{\prime}, \Sigma_{\sigma}\right) \tag{33}
\end{equation*}
$$

where $\Sigma_{\sigma}$ is a d-dimensional diagonal matrix with entries $\sigma^{2}$, and where $p(x)$ is a mixture of ddimensional Gaussian distributions:

$$
\begin{equation*}
p(x)=\sum_{l=1}^{L} \rho_{l} \phi\left(x \mid \mu_{l}, \Sigma_{l}\right) \tag{34}
\end{equation*}
$$

(Note that, in this section only, $x_{i}$ denotes the $i$ th component of the vector $x$.) Using properties of Gaussian distributions (see Appendix A. 2 of [26]) we obtain

$$
\begin{align*}
\mu_{p}(x) & =\int_{-\infty}^{\infty} k\left(x, x^{\prime}\right) p\left(x^{\prime}\right) \mathrm{d} x^{\prime} \\
& =\int_{-\infty}^{\infty} \lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \phi\left(x^{\prime} \mid x, \Sigma_{\sigma}\right) \times\left(\sum_{l=1}^{L} \rho_{l} \phi\left(x^{\prime} \mid \mu_{l}, \Sigma_{l}\right)\right) \mathrm{d} x^{\prime} \\
& =\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \rho_{l} \int_{-\infty}^{\infty} \phi\left(x^{\prime} \mid x, \Sigma_{\sigma}\right) \times \phi\left(x^{\prime} \mid \mu_{l}, \Sigma_{l}\right) \mathrm{d} x^{\prime}  \tag{35}\\
& =\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \rho_{l} \int_{-\infty}^{\infty} a_{l}^{-1} \phi\left(x^{\prime} \mid c_{l}, C_{l}\right) \mathrm{d} x^{\prime} \\
& =\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \rho_{l} a_{l}^{-1}
\end{align*}
$$

where we have:

$$
\begin{equation*}
a_{l}^{-1}=(2 \pi)^{-\frac{d}{2}}\left|\Sigma_{\sigma}+\Sigma_{l}\right|^{-\frac{1}{2}} \exp \left(-\frac{1}{2}\left(x-\mu_{l}\right)^{T}\left(\Sigma_{\sigma}+\Sigma_{l}\right)^{-1}\left(x-\mu_{l}\right)\right) \tag{36}
\end{equation*}
$$

This last expression is in fact itself a Gaussian distribution with probability density function $\phi\left(x \mid \mu_{l}, \Sigma_{l}+\Sigma_{\sigma}\right)$ and we hence obtain:

$$
\begin{equation*}
\mu_{p}(x):=\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \rho_{l} \phi\left(x \mid \mu_{l}, \Sigma_{l}+\Sigma_{\sigma}\right) \tag{37}
\end{equation*}
$$

Finally, we once again use properties of Gaussians to obtain

$$
\begin{align*}
\int_{-\infty}^{\infty} \mu_{p}(x) p(x) \mathrm{d} x= & \int_{-\infty}^{\infty}\left[\lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \rho_{l} \phi\left(x \mid \mu_{l}, \Sigma_{l}+\Sigma_{\sigma}\right)\right] \\
& \times\left[\sum_{m=1}^{L} \rho_{m} \phi\left(x \mid \mu_{m}, \Sigma_{m}\right)\right] \mathrm{d} x \\
= & \lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \sum_{m=1}^{L} \rho_{l} \rho_{m} \int_{-\infty}^{\infty} \phi\left(x \mid \mu_{l}, \Sigma_{l}+\Sigma_{\sigma}\right) \phi\left(x \mid \mu_{m}, \Sigma_{m}\right) \mathrm{d} x \\
= & \lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \sum_{m=1}^{L} \rho_{l} \rho_{m} a_{l m}^{-1} \\
= & \lambda^{2}(\sqrt{2 \pi} \sigma)^{d} \sum_{l=1}^{L} \sum_{m=1}^{L} \rho_{l} \rho_{m} \phi\left(\mu_{l} \mid \mu_{m}, \Sigma_{l}+\Sigma_{m}+\Sigma_{\sigma}\right) \tag{38}
\end{align*}
$$

Other combinations of kernel $k$ and density $p$ that give rise to an analytic mean element can be found in the references of [1].

## Appendix D: Details of the Application to Proteomics Data

Description of the Model Choice Problem

The 'CheMA' methodology described in [19] contains several elements that we do not attempt to reproduce in full here; in particular we do not attempt to provide a detailed motivation for the mathematical forms presented below, as this requires elements from molecular chemistry. For our present purposes it will be sufficient to define the statistical models $\left\{M_{i}\right\}_{i=1}^{m}$ and to clearly specify the integration problems that are to be solved. We refer the reader to [19] and the accompanying supplementary materials for a full biological background.

Denote by $\mathcal{D}$ the dataset containing normalised measured expression levels $y_{S}\left(t_{j}\right)$ and $y_{S}^{*}\left(t_{j}\right)$ for, respectively, the unphosphorylated and phosphorylated forms of a protein of interest ('substrate') in a longitudinal experiment at time $t_{j}$. In addition $\mathcal{D}$ contains normalised measured expression levels $y_{E_{i}}^{*}\left(t_{j}\right)$ for a set of possible regulator kinases ('enzymes', here phosphorylated proteins) that we denote by $\left\{E_{i}\right\}$.

An important scientific goal is to identify the roles of enzymes (or 'kinases') in protein signaling; in this case the problem takes the form of variable selection and we are interested to discover which enzymes must be included in a model for regulation of the substrate $S$. Specifically, a candidate model $M_{i}$ specifies which enzymes in the set $\left\{E_{i}\right\}$ are regulators of the substrate $S$, for example $M_{3}=\left\{E_{2}, E_{4}\right\}$. Following [19] we consider models containing at most two enzymes, as well as the model containing no enzymes.
Given a dataset $\mathcal{D}$ and model $M_{i}$, we can write down a likelihood function as follows:

$$
\begin{equation*}
L\left(\theta_{i}, M_{i}\right)=\prod_{n=1}^{N} \phi\left(\frac{y_{S}^{*}\left(t_{n+1}\right)-y_{S}^{*}\left(t_{n}\right)}{t_{n+1}-t_{n}} \left\lvert\, \frac{-V_{0} y_{S}^{*}\left(t_{n}\right)}{y_{S}^{*}\left(t_{n}\right)+K_{0}}+\sum_{E_{j} \in M_{i}} \frac{V_{j} y_{E_{j}}^{*}\left(t_{n}\right) y_{S}\left(t_{n}\right)}{y_{S}\left(t_{n}\right)+K_{j}}\right., \sigma_{\mathrm{err}}^{2}\right) . \tag{39}
\end{equation*}
$$

Here the model parameters are $\theta_{i}=\left\{\mathrm{K}, \mathrm{V}, \sigma_{\text {err }}\right\}$, where $(\mathrm{K})_{j}=K_{j},(\mathrm{~V})_{j}=V_{j}, \phi$ is the normal p.d.f. and the mathematical forms arise from the Michaelis-Menten theory of enzyme kinetics. The $V_{j}$ are known as 'maximum reaction rates' and the $K_{j}$ are known as 'Michaelis-Menten parameters'. This is classical chemical notation, not to be confused with the kernel matrix from the main text. The final parameter $\sigma_{\text {err }}$ defines the error magnitude for this 'approximate gradient-matching' statistical model.

The prior specification proposed in [19] and followed here is

$$
\begin{align*}
\mathrm{K} & \sim \phi_{T}\left(K \mid 1,2^{-1} \mathrm{I}\right)  \tag{40}\\
\sigma_{\mathrm{err}} \mid \mathrm{K} & \sim p\left(\sigma_{\mathrm{err}}\right) \propto 1 / \sigma_{\mathrm{err}}  \tag{41}\\
\mathrm{~V} \mid \mathrm{K}, \sigma & \sim \phi_{T}\left(V \mid 1, N \sigma_{\mathrm{err}}^{2}\left(\mathrm{X}(\mathrm{~K})^{T} \mathrm{X}(\mathrm{~K})\right)^{-1}\right) \tag{42}
\end{align*}
$$

where $\phi_{T}$ denotes a Gaussian distribution, truncated so that its support is $[0, \infty)$ (since kinetic parameters cannot be non-negative). Here $\mathrm{X}(\mathrm{K})$ is the design matrix associated with the linear regression that is obtained by treating the K as known constants; we refer to [19] for further details.

Due to its careful design, the likelihood in Eqn. 39 is partially conjugate, so the following integral can be evaluated in closed form:

$$
\begin{equation*}
L\left(\mathrm{~K}, M_{i}\right)=\int_{0}^{\infty} \int_{0}^{\infty} L\left(\theta_{i}, M_{i}\right) p\left(\mathrm{~V}, \sigma_{\text {err }} \mid \mathrm{K}\right) \mathrm{dVd} \sigma_{\text {err }} \tag{43}
\end{equation*}
$$

The numerical challenge is then to compute the integral

$$
\begin{equation*}
L\left(M_{i}\right)=\int_{0}^{\infty} L\left(\mathrm{~K}, M_{i}\right) p(\mathrm{~K}) \mathrm{dK} \tag{44}
\end{equation*}
$$

for each candidate model $M_{i}$. Depending on the number of enzymes in model $M_{i}$, this will either be a 1-, 2- or 3-dimensional numerical integral. Whilst such integrals are not challenging to compute on a per-individual basis, the nature of the application means that the values $L\left(M_{i}\right)$ will be similar for many candidate models and, when the number of models is large, this demands either a very precise calculation per model or a careful quantification of the impact of numerical error on the subsequent inferences (i.e. determining the MAP estimate). It is this particular issue that motivates the use of probabilistic numerical methods.

## Description of the Computational Problem

We need to compute integrals of functions with domain $\mathcal{X}=[0, \infty)^{d}$ where $d \in\{1,2,3\}$ and the sampling distribution $p(x)$ takes the form $\phi_{T}\left(x \mid 1,2^{-1} \mathrm{I}\right)$. The test function $f(x)$ corresponds to $L\left(\mathrm{~K}, M_{i}\right)$ with $x=\mathrm{K}$. This is given explicitly by the $g$-prior formulae as:

$$
\begin{align*}
L\left(\mathrm{~K}, M_{i}\right) & =\frac{1}{(2 \pi)^{N / 2}} \frac{1}{(N+1)^{d / 2}} \Gamma\left(\frac{N}{2}\right) b_{N}^{-\frac{N}{2}}  \tag{45}\\
b_{N} & =\frac{1}{2}\left(\mathrm{Y}^{T} \mathrm{Y}+\frac{1}{N} 1^{T} \mathrm{X}^{T} \mathrm{X} 1-\mathrm{V}_{N}^{T} \Omega_{N} \mathrm{~V}_{N}\right)  \tag{46}\\
\mathrm{V}_{N} & =\Omega_{N}^{-1}\left(\frac{1}{N} \mathrm{X}^{T} \mathrm{X} 1+\mathrm{X}^{T} \mathrm{Y}\right)  \tag{47}\\
\Omega_{N} & =\left(1+\frac{1}{N}\right) \mathrm{X}^{T} \mathrm{X}  \tag{48}\\
(\mathrm{Y})_{n} & =\frac{y_{S}^{*}\left(t_{n+1}\right)-y_{S}^{*}\left(t_{n}\right)}{t_{n+1}-t_{n}} \tag{49}
\end{align*}
$$

where for clarity we have suppressed the dependence of X on K. For the Frank-Wolfe Bayesian Quadrature algorithm, we require that the mean element $\mu_{p}$ is analytically tractable and for this reason we employed the exponentiated-quadratic kernel with length scale $\lambda$ and width scale $\sigma$ parameters:

$$
\begin{equation*}
k\left(x, x^{\prime}\right)=\lambda^{2} \exp \left(-\frac{\sum_{i=1}^{d}\left(x_{i}-x_{i}^{\prime}\right)^{2}}{2 \sigma^{2}}\right) \tag{51}
\end{equation*}
$$

For simplicity we focussed on the single hyper-parameter pair $\lambda=\sigma=1$, which produces:

$$
\begin{align*}
\mu_{p}(x) & =\int_{0}^{\infty} k\left(x, x^{\prime}\right) p\left(x^{\prime}\right) \mathrm{d} x^{\prime}  \tag{52}\\
& =\int_{0}^{\infty} \exp \left(-\sum_{i=1}^{d}\left(x_{i}-x_{i}^{\prime}\right)^{2}\right) \phi_{T}\left(x^{\prime} \mid 1,2^{-1} \mathrm{I}\right) \mathrm{d} x^{\prime}  \tag{53}\\
& =2^{-d / 2}(1+\operatorname{erf}(1))^{-d} \prod_{i=1}^{d} \exp \left(-\frac{\left(x_{i}-1\right)^{2}}{2}\right)\left(1+\operatorname{erf}\left(\frac{x_{i}+1}{\sqrt{2}}\right)\right), \tag{54}
\end{align*}
$$

where $\phi_{T}$ is the p.d.f. of the truncated Gaussian distribution introduced above and erf is the error function. To compute the posterior variance of the numerical error we also require the quantity:

$$
\int_{0}^{\infty} \int_{0}^{\infty} k\left(x, x^{\prime}\right) p(x) p\left(x^{\prime}\right) \mathrm{d} x \mathrm{~d} x^{\prime}=\int_{0}^{\infty} \mu_{p}(x) p(x) \mathrm{d} x= \begin{cases}0.629907 \ldots & \text { for } d=1  \tag{55}\\ 0.396783 \ldots & \text { for } d=2 \\ 0.249937 \ldots & \text { for } d=3\end{cases}
$$

which we have simply evaluated numerically. We emphasise that principled approaches to hyperparameter elicitation are an important open research problem that we aim to address in a future publication (see discussion in the main text). The values used here are scientifically reasonable and serve to illustrate key aspects of our methodology.
FWBQ provides posterior distributions over the numerical uncertainty in each of our estimates for the marginal likelihoods $L\left(M_{i}\right)$. In order to propagate this uncertainty forward into a posterior distribution over posterior model probabilities (see Figs. 3 in the main text and S 2 below), we simply sampled values $\hat{L}\left(M_{i}\right)$ from each of the posterior distributions for $L\left(M_{i}\right)$ and used these samples values to construct posterior model probabilities $\hat{L}\left(M_{i}\right) / \sum_{j} \hat{L}\left(M_{j}\right)$. Repeating this procedure many times enables us to sample from the posterior distribution over the posterior model probabilities (i.e. two levels of Bayes' theorem). This provides a principled quantification of the uncertainty due to numerical error in the output of our primary Bayesian analysis.

## Description of the Data

The proteomic dataset $\mathcal{D}$ that we considered here was a subset of the larger dataset provided in [19]. Specifically, the substrate $S$ was the well-studied 4E-binding protein 1 (4EBP1) and the enzymes
$E_{j}$ consisted of a collection of key proteins that are thought to be connected with 4EBP1 regulation, or at least involved in similar regulatory processes within cellular signalling. Full details, including experimental protocols, data normalisation and the specific choice of measurement time points are provided in the supplementary materials associated with [19].
For this particular problem, biological interest arises because the data-generating system was provided by breast cancer cell lines. As such, the textbook description of 4EBP1 regulation may not be valid and indeed it is thought that 4EBP1 dis-regulation is a major contributing factor to these complex diseases (see [30]). We do not elaborate further on the scientific rationale for model-based proteomics in this work.


Figure S1: Comparison of quadrature methods on the proteomics dataset. Left: Value of the MMD $^{2}$ for FW (black), FWLS (red), FWBQ (green), FWLSBQ (orange) and SBQ (blue). Once again, we see the clear improvement of using Bayesian Quadrature weights and we see that Sequential Bayesian Quadrature improves on Frank-Wolfe Bayesian Quadrature and Frank-Wolfe LineSearch Bayesian Quadrature. Right: Empirical distribution of weights. The dotted line represent the weights of the Frank-Wolfe algorithm with line search, which has all weights $w_{i}=1 / n$. Note that the distribution of Bayesian Quadrature weights ranges from -17.39 to 13.75 whereas all versions of Frank-Wolfe have weights limited to $[0,1]$ and have to sum to 1 .


Figure S2: Quantifying numerical error in a model selection problem. Marginalisation of model parameters necessitates numerical integration and any error in this computation will introduce error into the reported posterior distribution over models. Here FWBQ is used to model this numerical error explicitly. Left: At $n=50$ design points the uncertainty due to numerical error prevents us from determining the true MAP estimate. Right: At $n=200$ design points, models 16 and 26 can be better distinguished as the uncertainty due to numerical error is reduced (model 26 can be seen to be the MAP estimate, although some uncertainty about this still remains even at this value of $n$, due to numerical error).

## Appendix E: FWBQ algorithms with Random Fourier Features

In this section, we will investigate the use of random Fourier features (introduced in [25]) for the FWLS and FWLSBQ algorithms. An advantage of using this type of approximation is that the cost of manipulating the Gram matrix, and in particular of inverting it, goes down from $\mathcal{O}\left(n^{3}\right)$ to $\mathcal{O}\left(n D^{2}\right)$ for some user-defined constant $D$ which controls the quality of approximation. This could make Bayesian Quadrature more competitive against other integration methods such as MCMC or QMC. Furthermore, the kernels obtained using this method lead to finite-dimensional RKHS, which therefore satisfy the assumptions required for the theory in this paper to hold. This will be the aspect that we will focus on. In particular, we will show empirically that exponential convergence may be possible even when the RKHS is infinite-dimensional.
We will re-use the 20 -component mixture of Gaussians example with $d=2$ from our simulation studies, but using instead a random Fourier approximation of the exponentiated-quadratic (EQ) kernel $k\left(x, x^{\prime}\right):=\lambda^{2} \exp \left(-1 / 2 \sigma^{2}\left\|x-x^{\prime}\right\|_{2}^{2}\right)$ with $(\lambda, \sigma)=(1,0.8)$ and $M=10000$.

Following Bochner's theorem, we can always express translation invariant kernels in Fourier space:

$$
\begin{equation*}
k\left(x, x^{\prime}\right)=\int_{\mathcal{W}} g(w) \exp \left(j w\left(x-x^{\prime}\right)\right) \mathrm{d} w=\mathbb{E}\left[\exp \left(j w^{T} x\right) \exp \left(j w^{T} x^{\prime}\right)\right] \tag{56}
\end{equation*}
$$

where $w \sim g(w)$ for $g(w)$ being the Fourier transform of the kernel. One can then use a Monte Carlo approximation of the kernel's Fourier expression with $D$ samples whenever $g$ is a p.d.f.. Our approximated kernel will then lead to a $D$-dimensional RKHS and will be given by:

$$
\begin{equation*}
k\left(x, x^{\prime}\right) \approx \frac{1}{D} \sum_{j=1}^{D} z_{w_{j}, b_{j}}(x) z_{w_{j}, b_{j}}\left(x^{\prime}\right)=\hat{k}_{D}\left(x, x^{\prime}\right) \tag{57}
\end{equation*}
$$

where $z_{w_{j}, b_{j}}(x)=\sqrt{2} \cos \left(w_{j}^{T} x+b_{j}\right)$ and $b_{j} \sim[0,2 \pi]$ uniformly. Random Fourier features approximations are unbiased and, in the specific case of a $d$-dimensional EQ kernel with $\lambda=1$, we have to samples from the following Fourier transform:

$$
\begin{equation*}
g(w)=\left(\frac{2 \pi}{\sigma^{2}}\right)^{-\frac{d}{2}} \exp \left(-\frac{\sigma^{2}\|w\|_{2}^{2}}{2}\right) \tag{58}
\end{equation*}
$$

which is a $d$-dimensional Gaussian distribution with zero mean and covariance matrix with all diagonal elements equal to $\left(1 / \sigma^{2}\right)$.

The impact on the MMD from the use of random Fourier features to approximate the kernel for both the FWLS and FWLSBQ algorithms is demonstrated in Figure S3. In this example, the quadrature rule uses the kernel with random features but the MMD is calculated using the original $\mathcal{H}$-norm. The reason for using this $\mathcal{H}$-norm is to have a unique measure of distance between points which can be compared.
Clearly, we once again have that the rate of convergence of the FWLSBQ is much faster than FWLS when using the exact kernel. The same phenomena is observed for the method with high number of random features $(D=5000)$. This suggests that both the choice of design points and the calculation of the BQ weights is not strongly influenced by the approximation. It is also interesting to notice that the rates of convergence is very close for the exact and $D=5000$ methods (atleast when $n$ is small), potentially suggesting that exponential convergence is possible for the exact method. This is not so surprising in itself since using a Gaussian kernel represents a prior belief that the integrand of interest is very smooth, and we can therefore expect fast convergence of the method.

However, in the case with a smaller number of random features is used ( $D=1000$ ), we actually observe a very poor performance of the method, which is mainly due to the fact that the weights are not well approximated anymore.

In summary, the experiments in this section suggest that the use of random features is a potential alternative for scaling up Bayesian Quadrature, but that one needs to be careful to use a high enough number of features. The experiments also give hope of having very similar convergence for infinitedimensional and finite-dimensional spaces.


Figure S3: Random Fourier Features (RFF) for Bayesian Quadrature. RFF are used to approximate the EQ kernel in the example of the simulation study. The $\mathrm{MMD}^{2}$ is plotted in the case where the EQ kernel is used (FWLS: blue; FWLSBQ: black), as well as when a using random features with $D=1000$ (FWLS: red; FWLSBQ: purple) and $D=5000$ (FWLS: orange; FWLSBQ: green).


[^0]:    ${ }^{1}$ A detailed discussion on probabilistic numerics and an extensive up-to-date bibliography can be found at http://www.probabilistic-numerics.org

