

Replica analysis of overfitting in generalized linear regression models

A C C Coolen^{1,2,3,5} , M Sheikh⁴ , A Mozeika³ ,
F Aguirre-Lopez⁴  and F Antenucci²

¹ Dept of Biophysics, Radboud University, 6525AJ Nijmegen, The Netherlands

² Saddle Point Science Ltd, 35A South St, London W1K2XF, United Kingdom

³ London Inst for Mathematical Sciences, 35A South St, London W1K2XF, United Kingdom

⁴ Dept of Mathematics, King's College London, London WC2R 2LS, United Kingdom

E-mail: a.coolen@science.ru.nl

Received 20 April 2020, revised 18 June 2020

Accepted for publication 25 June 2020

Published 13 August 2020



CrossMark

Abstract

Nearly all statistical inference methods were developed for the regime where the number N of data samples is much larger than the data dimension p . Inference protocols such as maximum likelihood (ML) or maximum *a posteriori* probability (MAP) are unreliable if $p = \mathcal{O}(N)$, due to overfitting. This limitation has for many disciplines with increasingly high-dimensional data become a serious bottleneck. We recently showed that in Cox regression for time-to-event data the overfitting errors are not just noise but take mostly the form of a bias, and how with the replica method from statistical physics one can model and predict this bias and the noise statistics. Here we extend our approach to arbitrary generalized linear regression models (GLM), with possibly correlated covariates. We analyse overfitting in ML/MAP inference without having to specify data types or regression models, relying only on the GLM form, and derive generic order parameter equations for the case of L_2 priors. Second, we derive the probabilistic relationship between true and inferred regression coefficients in GLMs, and show that, for the relevant hyperparameter scaling and correlated covariates, the L_2 regularization causes a predictable direction change of the coefficient vector. Our results, illustrated by application to linear, logistic, and Cox regression, enable one to correct ML and MAP inferences in GLMs systematically for overfitting bias, and thus extend their applicability into the hitherto forbidden regime $p = \mathcal{O}(N)$.

Keywords: generalized linear models, overfitting, regression, replica method

(Some figures may appear in colour only in the online journal)

⁵ Author to whom any correspondence should be addressed.

1. Introduction

Extensive quantities of data are now available in many commercial, scientific and medical settings, due to the decreasing cost of high throughput measurement devices and data storage, and rapidly increased computing power. Here we will be concerned with data where each sample is a pair (z, s) , with $z \in \mathbb{R}^p$ (the input, or covariate vector) and with an output variable s . The latter can be real-valued, discrete, or even composite. The aim is to determine from a given set $\mathcal{D} = \{(z_1, s_1), \dots, (z_N, s_N)\}$ of randomly drawn historic samples whether there is information in z about s , and to predict the value of s associated with any vector z . In parametric statistical inference one approaches this question by postulating a parametrized probabilistic model $p(s|z, \theta)$ for the dependence of s on z , followed by defining a function $\Omega(\theta|\mathcal{D})$ whose minimization gives a sensible estimate for the parameters θ . In this work we will focus on generalized linear models (GLMs, [14]), which are regression models $p(s|z, \theta)$ in which the covariates z enter strictly via a linear combination $\beta \cdot z = \sum_{\mu=1}^p \beta_{\mu} z_{\mu}$, with coefficients $\beta \in \mathbb{R}^p$. For the $\Omega(\theta|\mathcal{D})$ function, common choices are $\Omega(\theta|\mathcal{D}) = -\log p(\mathcal{D}|\theta)$ (ML regression) and $\Omega(\theta|\mathcal{D}) = -\log p(\theta|\mathcal{D})$ (MAP regression). Here $p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta)$, and $p(\mathcal{D}|\theta) = \prod_i p(s_i|z_i, \theta)$. MAP requires the specification of a prior parameter distribution $p(\theta)$. ML and MAP can be seen as approximations of a computationally often intractable Bayesian approach, where one works with the full posterior distribution $p(\theta|\mathcal{D})$. MAP replaces this posterior by a delta peak at the most probable point, and ML follows from MAP by choosing a flat prior $p(\theta)$. ML performs well when $p \ll N$, but its estimators become increasingly incorrect in the high-dimensional regime, which for GLMs involves both the number N of samples and the number p of covariates diverging, with finite ratio $\zeta = p/N$. Remedial attempts can be categorized into corrective [1–8] or preventative ones [9]. The former route seeks to construct better estimators via a power series in N^{-1} (for fixed p), with the ML estimator as zeroth term, but becomes unwieldy beyond the linear term [10]. Also various computational resampling recipes were proposed [11, 12], and a wider family of estimators beyond ML/MAP [13]. All these remedial and corrective approaches tend to work only when the data dimension p is small and fixed. The most popular remedial approaches to overfitting in the statistics and machine learning communities are regularization, i.e. MAP inference with optimized priors, and variable selection, i.e. regression with low-dimensional representations of the vectors z .

We tackle the high-dimensional statistics regime adopting a statistical physics perspective. Statistical physics provided many contributions to our understanding of information processing in this regime, that is not captured in traditional statistics. Statistical physics tends to deal with ‘typical-case’ scenarios, unlike the ‘worst-case’ analysis more prevalent in statistics and computer science (see e.g. [15]). The two approaches are complementary, with ‘typical-case’ behaviour becoming relevant if the ‘worst-case’ scenario is very rare as $N \rightarrow \infty$. Within statistical physics, the techniques from the field of spin glasses have been particularly effective, especially replica theory and the cavity method (and related message passing algorithms).

The replica method [16] gives relatively simple prescriptions for computing asymptotic joint distributions of model parameters, and allows one to predict asymptotic values of statistical estimators. It led to valuable progress in various areas of computer science [17], in particular in machine learning [18–22]. Although it is adaptable to many machine learning problems, the replica method is not always provably exact. Hence considerable effort has been dedicated to proving rigorously the replica predictions in specific settings, for problems originating from statistical physics [23], and in machine learning (e.g. low-rank matrix factorisation

[24]). Alternative methods were also proposed to derive the replica results, based on the elegant interpolation technique [25], and later extended to Bayesian inference [26].

Inference problems are intrinsically algorithmic. One ideally wants computationally efficient methods for finding the answer to any problem instance. In this aspect, the insights from statistical physics originate from the iterative procedure for computing marginals in Ising spin models with pair interactions, known as the Thouless–Anderson–Palmer (TAP) equations [27]. When implemented correctly [28], it is equivalent to the belief propagation (BP) approach in computer science [29, 30], as was realised in [31]. For continuous variables and multibody interactions, the TAP approach to inference is now commonly referred to as approximate message passing (AMP) [32]. AMP is computationally simple, usually competitive with the fastest algorithms, rigorously characterized in the large system size limit by the cavity method (or state evolution), and used to predict accurately performance metrics such as the mean-squared error (MSE) or the detection accuracy [33–37]. A generalization of AMP to arbitrary priors and component-wise output functions is found in [38], which coined the name generalized approximate message passing (GAMP) for the generalized linear model. Message passing tools were also used to study logistic regression in the high-dimensional regime, for ML and MAP [39, 40]. Implications of overfitting for likelihood ratio tests in the $p = \mathcal{O}(N)$ regime were explored in [41].

Despite its successes, the application of AMP to real-world problems has been limited by its dependence on a Bayes-optimal setting and specific structural features of the data matrix. When there is model mismatch, replica symmetry may be broken and AMP may stop converging [42, 43]. When the distributions are unknown, one could try to find a minimax estimator over a class of distributions [44], or combine GAMP with expectation–maximization (EM) estimation [45–49]. Alternatively, AMP can be modified to a replica symmetric broken (RSB) structure, but this algorithm becomes computationally more demanding with each RSB step, and requires introducing additional parameters whose values are not easily determined [42]. Nevertheless, for linear models the main limitation of AMP is often the structure of the data matrix. AMP’s original version holds only for i.i.d. sub-Gaussian random data matrices [33, 34, 38], and AMP is fragile with regard to alternative choices. For example, it diverges for even mildly ill-conditioned or non-zero-mean data matrices [50–52]. Several heuristic strategies have been proposed for inducing AMP to converge [50, 52–54] but their effectiveness is limited. Other algorithms for linear regression have been designed using approximations of belief propagation (BP) and/or free-energy minimization, such as adaptive TAP [55], expectation consistent approximation (EC) [56–58], (S-transform AMP) S-AMP [59], and (orthogonal AMP) OAMP [60]. Numerical experiments suggest that some are indeed more robust than AMP, but their convergence has not been rigorously determined. Currently, the AMP-like algorithm with a rigorous underpinning that is able to handle the broadest class of data matrices is vector approximate message passing (VAMP) [61], which converges correctly for all large random data matrices that are right-orthogonally invariant. While this class of matrices relax the need for fully independent matrix elements, it is still an excessively specific requirement for most practical applications.

In this work we consider GLMs and focus on generalising the structure of the data matrix, employing the replica method—the most versatile of our theoretical tools. We build on recent studies [62, 63] which gave an accurate quantitative analysis of overfitting in (regularized) Cox models [69, 70] for time-to-event data. We show how the calculations of [62, 63] can be generalized to ML or MAP regression with arbitrary GLMs. Here we consider only models with a single linear combination of covariates (which includes logistic and ordinal class regression, perceptrons [66], and other survival analysis models such as frailty and random effects models [67]). Generalization to models with multiple linear combinations (e.g. neural networks with

hidden layers, or survival analysis with competing risks [68]) is straightforward. We analyse overfitting in ML or MAP inference with GLMs without requiring the entries of the data matrix to be uncorrelated. We only assume that there is no model mismatch, and that the N covariate vectors $\{z_i\}$ are drawn independently from some distribution $p(z)$. This distribution may describe correlated entries, provided some weak conditions on the spectrum of the correlation matrix are met. We refer to this setting as row-independent data matrix. Using only the generalized linear form of the models, we derive generic replica order parameter equations within the replica-symmetric ansatz (RS), for the case of Gaussian priors⁶. Only at the stage of solving order parameter equations will one have to specify model details. We also calculate the probabilistic relationship between true and inferred association coefficients, and show that, when covariates are correlated and L_2 regularizers are sufficiently strong to have an effect, the latter induce a predictable direction change of the coefficient vector. For linear regression problems, similar studies of MAP estimators are present in literature. Again these assume i.i.d. elements or some form of rotational invariance for the data matrix, either within an RS [71, 72] or an RSB ansatz [73]. Within the setting of rotationally invariant data matrices, the authors of [74] define an oracle version of VAMP and show rigorously that the corresponding state evolution converges to the MSE state evaluated via replica theory in [57, 71, 75]. Similar proofs of replica results can be found also for MMSE estimators in linear regression, see e.g. [76–78] for Gaussian data matrices and [79] for rotationally invariant data matrices. They cannot immediately be extended to the row-independent data matrices we consider in this work. In [45] the replica method is used to analyse properties of the MAP estimator in compressed sensing. Our present setting differs in two key aspects: we consider generalized linear models and we remove the need for i.i.d. entries of the data matrix. There is presently no AMP-like algorithm that provably works for independently drawn covariate vectors with correlated entries. The present RS replica calculation, however, is able to deal with such more realistic data, and the mathematical physics literature provides evidence that any failures of the replica solution in practical applications are most likely to reflect model mismatch, i.e. a breaking of the replica symmetry, as opposed to fundamental features of the analytical continuation. Our limitation to ML and MAP estimators is also not crucial, and the results could be easily extended to e.g. the minimum mean square error (MMSE) estimator. We concentrate on ML and MAP because they are less computationally demanding, in the absence of an AMP-like algorithm for row-independent data matrices, and their practical evaluation is usually feasible using standard numerical methods.

This paper is organized as follows. We first generalize in section 2 the replica analysis of [62, 63] to arbitrary GLMs. In section 3 we calculate the quantitative relation between true and inferred association parameters in the overfitting regime, for potentially correlated covariates, and show how our results can be used to compute new estimators that are decontaminated for overfitting distortions (via bias removal, or MSE minimization). In section 4 we test our theoretical predictions via application to linear, logistic and Cox regression, recovering some known results as a test, and deriving several new ones. We close with a discussion of present and future work. Most of the more technical calculations are relegated to appendices, to focus the main text on the key ideas and outcomes. In contrast to most analytical studies on overfitting in literature, our theory is not limited to linear models, to uncorrelated covariates, to scalar outputs, or to models with output noise. Our present results enable one to correct ML and MAP inferences in generalized linear regression models for overfitting distortions, and thereby

⁶The present limitation to L_2 (i.e. Gaussian) priors is not critical, alternative choices such as L_1 (or lasso) priors simply lead to more complicated integrals.

extend the applicability of these popular regression methods into the hitherto forbidden regime $p = \mathcal{O}(N)$.

2. General theory of GLM regression

2.1. Definitions and notation

In generalized linear regression models, the probability (density) of observing an outcome $s \in \Omega$ depends on the values of covariates $\mathbf{z} \in \mathbb{R}^p$ via an expression of the form $\text{Prob}(s|\mathbf{z}) = p(s|\boldsymbol{\beta} \cdot \mathbf{z}/\sqrt{p}, \theta)$, with θ representing any auxiliary parameters that are not coupled to covariates. The covariates appear only in an inner product with a vector $\boldsymbol{\beta}$ of so-called association parameters. The outcome set Ω can be continuous, discrete, or a combination of both (for multi-valued outcomes), and the auxiliary parameters θ can even be a function, as in the Cox model [69]. We consider MAP inference without model mismatch, where upon observing a data set $\mathcal{D} = \{(\mathbf{z}_1, s_1), \dots, (\mathbf{z}_N, s_N)\}$ in which all samples (\mathbf{z}_i, s_i) are assumed to have been drawn randomly and independently from a distribution of the form $p(\mathbf{z}, s) = p(\mathbf{z})p(s|\boldsymbol{\beta}^* \cdot \mathbf{z}/\sqrt{p}, \theta^*)$, the inferred parameters $(\hat{\boldsymbol{\beta}}, \hat{\theta})$ are those that maximize the Bayesian posterior parameter probability $p(\boldsymbol{\beta}, \theta|\mathcal{D})$:

$$\begin{aligned} p(\boldsymbol{\beta}, \theta|\mathcal{D}) &= \frac{p(\mathcal{D}|\boldsymbol{\beta}, \theta)p(\boldsymbol{\beta}, \theta)}{\int d\boldsymbol{\beta}' d\theta' p(\mathcal{D}|\boldsymbol{\beta}', \theta')p(\boldsymbol{\beta}', \theta')} \\ &= \frac{p(\boldsymbol{\beta}, \theta) \prod_{i=1}^N p(s_i|\boldsymbol{\beta} \cdot \mathbf{z}_i/\sqrt{p}, \theta)}{\int d\boldsymbol{\beta}' d\theta' p(\boldsymbol{\beta}', \theta') \prod_{i=1}^N p(s_i|\boldsymbol{\beta}' \cdot \mathbf{z}_i/\sqrt{p}, \theta')}. \end{aligned} \quad (2.1)$$

Hence, upon taking a logarithm and discarding an irrelevant constant,

$$(\hat{\boldsymbol{\beta}}, \hat{\theta})_{\text{MAP}} = \text{argmax}_{\boldsymbol{\beta}, \theta} \left\{ \sum_{i=1}^N \log p\left(s_i \left| \frac{\boldsymbol{\beta} \cdot \mathbf{z}_i}{\sqrt{p}}, \theta \right.\right) + \log p(\boldsymbol{\beta}, \theta) \right\}. \quad (2.2)$$

Choosing a regression model implies choosing a parametrization $p(s|\xi, \theta)$ and a prior $p(\boldsymbol{\beta}, \theta)$. We recover ML from MAP regression by choosing the prior to be constant. Our convention to define GLMs in terms of $\boldsymbol{\beta} \cdot \mathbf{z}/\sqrt{p}$ as opposed to $\boldsymbol{\beta} \cdot \mathbf{z}$ ensures that even for $p \rightarrow \infty$ the components of $\boldsymbol{\beta}$ will typically scale as $\beta_\mu = \mathcal{O}(1)$. Following mainstream literature, we will for simplicity choose nontrivial priors only for the coefficients $\boldsymbol{\beta}$, where their inclusion is indeed most critical, so $p(\boldsymbol{\beta}, \theta) \propto p(\boldsymbol{\beta})$.

For instance, the simplest GLM is linear regression, where one has outcomes from $\Omega = \mathbb{R}$, two auxiliary parameters $\theta = (\beta_0, \sigma)$ with $\beta_0 \in \mathbb{R}$ and $\sigma > 0$, and

$$p(s|\xi, \beta_0, \sigma) = (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{1}{2}(s-\xi-\beta_0)^2/\sigma^2}. \quad (2.3)$$

In logistic regression, which can be seen as a stochastic generalization of the binary perceptron [66], we have $\Omega = \{-1, 1\}$, one auxiliary parameter $\theta = \beta_0 \in \mathbb{R}$, and

$$p(s|\xi, \beta_0) = \frac{1}{2} + \frac{1}{2}s \tanh(\xi + \beta_0) = \frac{e^{s(\xi + \beta_0)}}{2 \cosh(\xi + \beta_0)}. \quad (2.4)$$

In Cox regression [69] without censoring we have $\Omega = [0, \infty)$ and a functional auxiliary parameter $\theta = \{\lambda(t)\}$, the base hazard rate, with

$$p(s|\xi, \lambda) = \lambda(s) \exp \left[\xi - e^\xi \int_0^s ds' \lambda(s') \right] = -\frac{\partial}{\partial s} \exp \left[-e^\xi \int_0^s ds' \lambda(s') \right]. \quad (2.5)$$

For Cox regression with censoring, $\Omega = [0, \infty) \times \{0, 1\}$ (the outcome is a pair (t, r) of an event time t and a binary label r indicating whether the event was a primary one or censoring), with two functional auxiliary parameters $\theta = \{\lambda_0(t), \lambda_1(t)\}$ (the base rates of the primary and the censoring events), and

$$p(t, r|\xi, \lambda_0, \lambda_1) = \lambda_r(t) \exp \left[\xi \delta_{r1} - \int_0^t dt' \lambda_0(t') - e^\xi \int_0^t dt' \lambda_1(t') \right]. \quad (2.6)$$

In proportional hazards ordinal class regression with C discrete possible outcomes we have $\Omega = \{1, 2, \dots, C\}$ and $\theta = (\lambda_2, \dots, \lambda_C) \in \mathbb{R}^{C-1}$, with

$$p(c|\xi, \lambda_2, \dots, \lambda_C) = (1 - \tilde{p}(c|\xi, \lambda_c)) \prod_{c' > c}^{C+1} \tilde{p}(c'|\xi, \lambda_{c'}) \quad (2.7)$$

$$\tilde{p}(1|\xi, \lambda) = 0, \quad \tilde{p}(C+1|\xi, \lambda) = 1, \quad 1 < c \leq C: \tilde{p}(c|\xi, \lambda) = e^{-\lambda \exp(\xi)}.$$

2.2. The information-theoretic overfitting measure

We follow closely the procedure in [62, 63], which can be adapted to arbitrary GLMs with only minimal change. We start from the observation that MAP regression for any model of the type $p(s|z, \theta)$ (whether or not of the GLM form) is equivalent to minimization over the model parameters θ of the quantity

$$\Omega(\theta|\mathcal{D}) = D(\hat{p}_{\mathcal{D}}||p_{\theta}) - N^{-1} \log p(\theta). \quad (2.8)$$

Here $\hat{p}_{\mathcal{D}}$ is the empirical distribution of covariates and outcomes in the data, $\hat{p}(s, z|\mathcal{D}) = N^{-1} \sum_{i \leq N} \delta(s - s_i) \delta(z - z_i)$, $D(\hat{p}_{\mathcal{D}}||p_{\theta})$ is the Kullback–Leibler distance

$$D(\hat{p}_{\mathcal{D}}||p_{\theta}) = \int dz ds \hat{p}(s, z|\mathcal{D}) \log \left(\frac{\hat{p}(s|z, \mathcal{D})}{p(s|z, \theta)} \right), \quad (2.9)$$

and p_{θ} is the assumed parametrized regression model, with covariate-conditioned outcome probabilities $p(s|z, \theta)$. For discrete variables, delta functions are replaced by Kronecker delta symbols⁷. Assuming that our data were indeed generated from a model of the assumed form, with (unknown) parameters θ^* , a transparent overfitting measure can be defined as $E(\theta^*, \mathcal{D}) = \min_{\theta} \Omega(\theta|\mathcal{D}) - \Omega(\theta^*|\mathcal{D})$, giving

$$E(\theta^*, \mathcal{D}) = \min_{\theta} \left\{ \frac{1}{N} \sum_{i=1}^N \log \left(\frac{p(s_i|z_i, \theta^*)}{p(s_i|z_i, \theta)} \right) + \frac{1}{N} \log \left(\frac{p(\theta^*)}{p(\theta)} \right) \right\}. \quad (2.10)$$

Perfect regression would give $E(\theta^*, \mathcal{D}) = 0$, finding $E(\theta^*, \mathcal{D}) < 0$ implies overfitting, and finding $E(\theta^*, \mathcal{D}) > 0$ implies underfitting. In GLMs with the MAP regression protocol (2.2) the parameters would be $\theta = (\beta, \theta)$. Our calculations focus on evaluating the average of (2.10) over the possible realizations of the data set \mathcal{D} , whose samples are drawn randomly from $p(s, z) = p(s|z, \theta^*)p(z)$ for some $p(z)$. The average is handled using the replica identity

$$\langle \log Z \rangle = \lim_{n \rightarrow 0} n^{-1} \log \langle Z^n \rangle, \quad (2.11)$$

⁷ Note that this definition is very different from the Kullback–Leibler distance employed in e.g. [64, 65], which, in our present notation, measures the deviation between the parameter posteriors $p(\theta|\mathcal{D})$ computed before and after removal of a single data sample.

and we write the minimization as the computation of the ground state energy density of a statistical mechanical system with degrees of freedom $\theta \in \mathbb{R}^p$ and Hamiltonian

$$H(\theta|\theta^*, \mathcal{D}) = \sum_{i=1}^N \log \left(\frac{p(s_i|z_i, \theta^*)}{p(s_i|z_i, \theta)} \right) + \log \left(\frac{p(\theta^*)}{p(\theta)} \right). \quad (2.12)$$

We can thus model MAP regression as the zero noise limit of a stochastic minimization of $H(\theta|\theta^*, \mathcal{D})$ at inverse noise level γ , giving, with the help of (2.11),

$$\langle E(\theta^*, \mathcal{D}) \rangle_{\mathcal{D}} = \lim_{\gamma \rightarrow \infty} E_{\gamma}(\theta^*), \quad (2.13)$$

$$\begin{aligned} E_{\gamma}(\theta^*) &= -\frac{\partial}{\partial \gamma} \frac{1}{N} \left\langle \log \int d\theta e^{-\gamma H(\theta|\theta^*, \mathcal{D})} \right\rangle_{\mathcal{D}} \\ &= -\lim_{n \rightarrow 0} \frac{\partial}{\partial \gamma} \frac{1}{Nn} \log \int d\theta^1 \dots d\theta^n \left\langle e^{-\gamma \sum_{\alpha=1}^n H(\theta^{\alpha}|\theta^*, \mathcal{D})} \right\rangle_{\mathcal{D}} \\ &= -\lim_{n \rightarrow 0} \frac{\partial}{\partial \gamma} \frac{1}{Nn} \log \int d\theta^1 \dots d\theta^n \prod_{\alpha=1}^n \left[\frac{p(\theta^{\alpha})}{p(\theta^*)} \right]^{\gamma} \\ &\quad \times \left\{ \int dz ds p(z) p(s|z, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|z, \theta^{\alpha})}{p(s|z, \theta^*)} \right]^{\gamma} \right\}^N. \end{aligned} \quad (2.14)$$

Integrals over outcome variables become summations when these variables are discrete, and integrals over functional parameters are interpreted as path integrals. In the alternative limit $\gamma \rightarrow 1$ the quantity $E_{\gamma}(\theta^*)$ would involve the average over all data realizations of the Bayesian estimator, $E_1(\theta^*) = \langle \int d\theta p(\theta|\mathcal{D}) \Omega(\theta|\mathcal{D}) - \Omega(\theta^*|\mathcal{D}) \rangle_{\mathcal{D}}$.

Application of expression (2.14) to generalized linear regression models implies setting $\theta \rightarrow (\beta, \theta)$ and $p(s|z, \theta) \rightarrow p(s|\beta \cdot z / \sqrt{p}, \theta)$, so that we obtain

$$\langle E(\beta^*, \theta^*, \mathcal{D}) \rangle_{\mathcal{D}} = \lim_{\gamma \rightarrow \infty} E_{\gamma}(\beta^*, \theta^*), \quad (2.15)$$

with

$$\begin{aligned} E_{\gamma}(\beta^*, \theta^*) &= -\lim_{n \rightarrow 0} \frac{\partial}{\partial \gamma} \frac{1}{Nn} \log \int d\theta^1 \dots d\theta^n \int d\beta^1 \dots d\beta^n \prod_{\alpha=1}^n \left[\frac{p(\beta^{\alpha})}{p(\beta^*)} \right]^{\gamma} \\ &\quad \times \left\{ \int dz ds p(z) p(s|z, \beta^*, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|z, \beta^{\alpha}, \theta^{\alpha})}{p(s|z, \beta^*, \theta^*)} \right]^{\gamma} \right\}^N. \end{aligned} \quad (2.16)$$

2.3. Replica analysis in the regime $p = \mathcal{O}(N)$

In appendix A we compute (2.16) for $N, p \rightarrow \infty$ with fixed ratio $\zeta = p/N$, assuming $p(z)$ to be a zero-average distribution on \mathbb{R}^p , and for $L2$ priors $p(\beta) \propto \exp(-\eta\beta^2)$.⁸ We include this derivation, which follows [63], for completeness. The outcome of the regression process is

⁸ Note that this latter choice would become $p(\beta) \propto \exp(-\eta p \beta^2)$ for the alternative convention where the association coefficients are not rescaled by \sqrt{p} , i.e. for GLMs written as $p(s|\beta \cdot z, \theta)$.

characterized by the values of a finite number of order parameters⁹, in terms of which one can quantify the relation between inferred and true regression coefficients. The result of appendix A is:

$$\begin{aligned} \lim_{N \rightarrow \infty} E_\gamma(\beta^*, \theta^*) &= \int \text{D}y_0 \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \log p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) - \zeta \eta S^2 \\ &+ \eta \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma + ga} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta\gamma + ga)^2} \right\rangle \right. \\ &+ \left. \left\langle \frac{1}{2\eta\gamma + ga} \right\rangle - f \left\langle \frac{a}{(2\eta\gamma + ga)^2} \right\rangle \right\} - \int \text{D}z \text{D}y_0 \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \\ &\times \frac{\int \text{D}y p^\gamma(s|uy + wy_0 + vz, \theta) \log p(s|uy + wy_0 + vz, \theta)}{\int \text{D}y p^\gamma(s|uy + wy_0 + vz, \theta)}, \end{aligned} \quad (2.17)$$

with the shorthand $\text{D}z = (2\pi)^{-\frac{1}{2}} e^{-\frac{1}{2}z^2} \text{d}z$. Brackets denote averages over the limit $p \rightarrow \infty$ of the eigenvalue spectrum $\varrho(a)$ of the $p \times p$ covariate correlation matrix \mathbf{A} , with entries $A_{\mu\nu} = \int \text{d}z p(z) z_\mu z_\nu$. This result depends on the true association parameter vector β^* only via the asymptotic rescaled amplitude $S^2 = \lim_{p \rightarrow \infty} p^{-1} \beta^{*2}$, assuming the components of β^* to have been drawn randomly from a symmetric distribution with finite second and fourth moments. Of the covariate covariance matrix \mathbf{A} we only require that its eigenvalue spectrum obeys $\lim_{p \rightarrow \infty} \int \text{d}a \varrho(a) a \in \mathbb{R}$ and $\lim_{p \rightarrow \infty} p^{-1} \int \text{d}a \varrho(a) a^2 = 0$. The order parameters (u, v, w, f, g, θ) are determined by extremization of the following quantity, which acts as a free energy density:

$$\Psi_{\text{RS}}(u, v, w, f, g, \theta) = \Psi_{\text{RS}}^A(u, v, w, f, g) - \Psi_{\text{RS}}^B(u, v, w, \theta), \quad (2.18)$$

with

$$\begin{aligned} \Psi_{\text{RS}}^A(\dots) &= -\frac{1}{2} \zeta (g + f) u^2 - \frac{1}{2} \zeta g (v^2 + w^2) \\ &+ \frac{1}{2} \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma + ga} \right\rangle^{-1} + \langle \log(2\eta\gamma + ga) \rangle + f \left\langle \frac{a}{2\eta\gamma + ga} \right\rangle \right\}, \end{aligned} \quad (2.19)$$

$$\Psi_{\text{RS}}^B(\dots) = \int \text{D}y_0 \text{D}z \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \log \int \text{d}y p^\gamma(s|uy + wy_0 + vz, \theta). \quad (2.20)$$

Here θ^* are the true (unknown) auxiliary model parameters assumed to have been used to generate the observed data. The physical meaning of the two main order parameters, expressed in terms of the MAP-inferred parameters $\hat{\beta}$ and the true parameters β^* responsible for the data, is

$$v = \lim_{p \rightarrow \infty} \left\langle \frac{1}{\sqrt{p}} \left\{ \hat{\beta} \cdot \mathbf{A} \hat{\beta} - \frac{(\beta^* \cdot \mathbf{A} \hat{\beta})^2}{\beta^* \cdot \mathbf{A} \beta^*} \right\}^{\frac{1}{2}} \right\rangle_{\mathcal{D}} \quad (2.21)$$

⁹The order parameters have been determined within the so-called replica-symmetric (RS) ansatz, which implies the assumption that the stochastic optimization process at finite γ is ergodic.

$$w = \lim_{p \rightarrow \infty} \left\langle \frac{1}{\sqrt{p}} \frac{\beta^* \cdot A \hat{\beta}}{\sqrt{\beta^* \cdot A \beta^*}} \right\rangle_{\mathcal{D}} \quad (2.22)$$

Perfect regression, i.e. $\hat{\beta} = \beta^*$, would give $v = 0$ and $w = \lim_{p \rightarrow \infty} \left\langle \sqrt{\beta^* \cdot A \beta^* / p} \right\rangle_{\mathcal{D}}$.

In the limit $\gamma \rightarrow \infty$ the maximization of the posterior becomes deterministic, and we recover the formulae describing MAP inference. In the precursor studies [62, 63] it was found that the canonical scaling of the RS order parameters for large γ is

$$u = \tilde{u} / \sqrt{\gamma}, \quad v, w, \theta = \mathcal{O}(1), \quad g = \tilde{g}\gamma, \quad f = \tilde{f}\gamma^2. \quad (2.23)$$

Assuming this scaling to hold more generally gives

$$\lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \Psi_{\text{RS}}^A(\dots) = \frac{1}{2} \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} + \tilde{f} \left[\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle - \tilde{u}^2 \right] - \tilde{g}(v^2 + w^2) \right\}, \quad (2.24)$$

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \Psi_{\text{RS}}^B(\dots) &= \int \text{D}y_0 \text{D}z \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) \\ &\times \lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \log \int \text{d}y e^{-\frac{1}{2}y^2} p^\gamma(s|\tilde{u}y/\sqrt{\gamma} + wy_0 + vz, \theta) \\ &= \int \text{D}y_0 \text{D}z \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) \max_{y \in \mathbb{R}} \left[\log p(s|\tilde{u}y + wy_0 + vz, \theta) - \frac{1}{2}y^2 \right] \\ &= \int \text{D}y_0 \text{D}z \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) \max_{\xi \in \mathbb{R}} \left[\log p(s|\xi, \theta) - \frac{(\xi - wy_0 - vz)^2}{2\tilde{u}^2} \right]. \end{aligned} \quad (2.25)$$

We next abbreviate $\Xi_A(\tilde{f}, \tilde{g}, \tilde{u}, v, w) = \lim_{\gamma \rightarrow \infty} \gamma^{-1} \Psi_{\text{RS}}^A(\dots)$ and $\Xi_B(\tilde{u}, v, w, \theta) = \lim_{\gamma \rightarrow \infty} \gamma^{-1} \Psi_{\text{RS}}^B(\dots)$, and write the various averages as $\langle \dots \rangle = \int \text{D}y_0 \text{D}z \dots$ and $\langle f(s) \rangle_s = \int \text{d}s p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) f(s)$ (with the integral over s replaced by a sum if s is discrete). We also define

$$\xi(\mu, \sigma, s, \theta) = \operatorname{argmax}_{\xi \in \mathbb{R}} \left[\log p(s|\xi, \theta) - \frac{1}{2}(\xi - \mu)^2 / \sigma^2 \right]. \quad (2.26)$$

A sufficient condition for $\xi(\mu, \sigma, s, \theta)$ to exist is that $\operatorname{argmax}_{\xi \in \mathbb{R}} p(s|\xi, \theta)$ exists for all (s, θ) , which we have found to be true in all GLM models considered so far. Note that $\xi(\mu, \sigma, s, \theta)$ is the solution of

$$\frac{\partial}{\partial \xi} \log p(s|\xi, \theta) = (\xi - \mu) / \sigma^2. \quad (2.27)$$

Hence we may write the model-independent part of the quantity to be extremized as

$$\Xi_A(\tilde{f}, \tilde{g}, \tilde{u}, v, w) = \frac{1}{2} \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} + \tilde{f} \left[\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle - \tilde{u}^2 \right] - \tilde{g}(v^2 + w^2) \right\}, \quad (2.28)$$

and the model-dependent part as

$$\begin{aligned} \Xi_B(\tilde{u}, v, w, \theta) &= \left\langle \left\langle \left\langle \max_{\xi \in \mathbb{R}} \left[\log p(s|\xi, \theta) - \frac{1}{2}(\xi - wy_0 - vz)^2 / \tilde{u}^2 \right] \right\rangle_s \right\rangle \right\rangle \\ &= \left\langle \left\langle \left\langle \left[\log p(s|\xi, \theta) - \frac{(\xi - wy_0 - vz)^2}{2\tilde{u}^2} \right] \Big|_{\xi = \xi(wy_0 + vz, \tilde{u}, s, \theta)} \right\rangle_s \right\rangle \right\rangle. \end{aligned} \quad (2.29)$$

The RS order parameter equations can then be written as

$$\frac{\partial \Xi_A}{\partial \tilde{f}} = \frac{\partial \Xi_A}{\partial \tilde{g}} = \frac{\partial \Xi_B}{\partial \theta} = 0, \quad (2.30)$$

$$\frac{\partial \Xi_A}{\partial \tilde{u}} = \frac{\partial \Xi_B}{\partial \tilde{u}}, \quad \frac{\partial \Xi_A}{\partial v} = \frac{\partial \Xi_B}{\partial v}, \quad \frac{\partial \Xi_A}{\partial w} = \frac{\partial \Xi_B}{\partial w}. \quad (2.31)$$

In appendix C we analyse and simplify these RS order parameter equations further, and find that we can rewrite our closed MAP order parameter equation set as:

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \quad (2.32)$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2, \quad (2.33)$$

$$\left\langle \left\langle [\xi(wy_0 + vz, \tilde{u}, s, \theta) - wy_0 - vz]^2 \right\rangle_s \right\rangle = -\zeta \tilde{f} \tilde{u}^4, \quad (2.34)$$

$$\left\langle \left\langle (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) \right\rangle_s \right\rangle = 1 - \zeta \tilde{g} \tilde{u}^2, \quad (2.35)$$

$$\left\langle \left\langle \left\langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \frac{\partial \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*)}{\partial y_0} \right\rangle_s \right\rangle \right\rangle = \zeta w \tilde{u}^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1}, \quad (2.36)$$

$$\left\langle \left\langle \left\langle \frac{\partial \log p(s|\xi, \theta)}{\partial \theta} \Big|_{\xi = \xi(wy_0 + vz, \tilde{u}, s, \theta)} \right\rangle_s \right\rangle \right\rangle = 0. \quad (2.37)$$

The function $\xi(\mu, \sigma, s, \theta)$, defined as the solution of (2.27), obeys $\lim_{\sigma \rightarrow 0} \xi(\mu, \sigma, s, \theta) = \mu$. Its partial derivative $(\partial_1 \xi)(\mu, \sigma, s, \theta)$ follows upon working out the partial derivative with respect to μ of (2.27),

$$0 = \frac{1}{\sigma^2} + \frac{\partial \xi}{\partial \mu} \frac{\partial}{\partial \xi} \left[\frac{\partial}{\partial \xi} \log p(s|\xi, \theta) + \frac{\mu - \xi}{\sigma^2} \right]. \quad (2.38)$$

Hence

$$(\partial_1 \xi)(\mu, \sigma, s, \theta) = \left[1 - \sigma^2 \frac{\partial^2 \log p(s|\xi, \theta)}{\partial \xi^2} \right]_{\xi = \xi(\mu, \sigma, s, \theta)}^{-1} \quad (2.39)$$

and

$$(\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) = \left[1 - \tilde{u}^2 \frac{\partial^2 \log p(s|\xi, \theta)}{\partial \xi^2} \right]_{\xi = \xi(wy_0 + vz, \tilde{u}, s, \theta)}^{-1}. \quad (2.40)$$

2.4. The limit $\zeta \rightarrow 0$ for ML regression

For $\eta = 0$ we revert back to ML regression. Here equations (2.32) and (2.33) simplify to $\tilde{g} = 1/\tilde{u}^2$ and $\tilde{f} = -v^2/\tilde{u}^4$, equation (2.37) remains unaltered, and the three equations (2.34)–(2.36) referring to extremization over (\tilde{u}, v, w) simplify to

$$\langle\langle\langle[\xi(wy_0 + vz, \tilde{u}, s, \theta) - wy_0 - vz]^2\rangle_s\rangle\rangle = \zeta v^2, \quad (2.41)$$

$$\langle\langle\langle(\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta)\rangle_s\rangle\rangle = 1 - \zeta, \quad (2.42)$$

$$\langle\langle\langle\left\langle\xi(wy_0 + vz, \tilde{u}, s, \theta) \frac{\partial \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*)}{\partial y_0}\right\rangle_s\rangle\rangle = \zeta w. \quad (2.43)$$

As a test, let us consider the classical regime $\zeta \rightarrow 0$ where the number of samples is much larger than the number of covariates. We can show relatively easily, for any model $p(s|\xi, \theta)$, that in this limit the remaining ML equations are solved by the correct solution $(\tilde{u}, v, w, \theta) = (0, 0, S\langle a \rangle^{\frac{1}{2}}, \theta^*)$, as one should expect. To see this we use

$$\lim_{\tilde{u} \rightarrow 0} \xi(wy_0 + vz, \tilde{u}, s, \theta) = wy_0 + vz, \quad (2.44)$$

$$\lim_{\tilde{u} \rightarrow 0} (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) = 1. \quad (2.45)$$

Upon inserting $(\tilde{u}, v, w, \theta) = (0, 0, S\langle a \rangle^{\frac{1}{2}}, \theta^*)$ we find that (2.41) and (2.42) are trivially satisfied, whereas (2.37) and (2.43) reduce, respectively, to the trivial statements

$$\begin{aligned} 0 &= \left\langle\left\langle\left\langle\frac{\partial \log p(s|wy_0, \theta^*)}{\partial \theta^*}\right\rangle_s\right\rangle\right\rangle \\ &= \int \mathcal{D}y_0 \int ds \frac{\partial}{\partial \theta^*} p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) = 0 \end{aligned} \quad (2.46)$$

and

$$\begin{aligned} 0 &= \left\langle y_0 \left\langle\left\langle\frac{\partial \log p(s|wy_0, \theta^*)}{\partial y_0}\right\rangle_s\right\rangle\right\rangle \\ &= S\langle a \rangle^{\frac{1}{2}} \int \mathcal{D}y_0 y_0 \int ds \frac{\partial}{\partial y_0} p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*) = 0. \end{aligned} \quad (2.47)$$

3. Link between true and inferred association parameters

3.1. Replica evaluation of the joint distribution

We can also calculate the (probabilistic) relation¹⁰ between the true parameters β_μ^* and the MAP-inferred values $\hat{\beta}_\mu$ in regression models of the GLM form. The relevant object to be computed in the case of MAP regression with prior $p(\beta)$ and data $\mathcal{D} = \{(z_1, s_1), \dots, (z_N, s_N)\}$ is the joint distribution

$$\mathcal{P}(\beta, \beta^* | \mathcal{D}) = \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \frac{\int d\theta d\beta e^{\gamma \log p(\theta, \beta | \mathcal{D})} \delta(\beta - \beta_\mu)}{\int d\theta d\beta e^{\gamma \log p(\theta, \beta | \mathcal{D})}} \delta(\beta^* - \beta_\mu^*), \quad (3.1)$$

¹⁰Note that, although anticipated at the time, this was not yet done in the previous studies [62, 63].

with the posterior parameter likelihood

$$p(\theta, \beta | \mathcal{D}) = \frac{p(\beta, \theta) \prod_{i=1}^N p(s_i | \beta \cdot \mathbf{z}_i / \sqrt{p}, \theta)}{\int d\beta' d\theta' p(\beta', \theta') \prod_{i=1}^N p(s_i | \beta' \cdot \mathbf{z}_i / \sqrt{p}, \theta')}. \quad (3.2)$$

The limit $\gamma \rightarrow \infty$ ensures that the integrations in (3.1) are dominated by the parameter values where $p(\theta, \beta | \mathcal{D})$ is maximized. Hence the fraction in (3.1) reduces to $\delta(\beta - \hat{\beta}_\mu)$, where $\hat{\beta}_\mu$ is the MAP estimator of the parameter β_μ , given dataset \mathcal{D} . Note that for $\gamma \rightarrow 1$ expression (3.1) would have reduced to the joint distribution of true coefficients and their Bayesian estimators, viz.

$$\mathcal{P}(\beta, \beta^* | \mathcal{D}) = \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^*) \int d\theta d\beta p(\theta, \beta | \mathcal{D}) \delta(\beta - \beta_\mu). \quad (3.3)$$

As always we are interested mainly in the typical form of $\mathcal{P}(\beta, \beta^* | \mathcal{D})$, so we average over the possible realizations of the data set \mathcal{D} , assuming all samples (\mathbf{z}_i, s_i) to be drawn randomly and independently from $p(\mathbf{z})p(s | \beta^* \cdot \mathbf{z} / \sqrt{p}, \theta^*)$:

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \langle \mathcal{P}(\beta, \beta^* | \mathcal{D}) \rangle_{\mathcal{D}} \\ &= \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^*) \left\langle \frac{\int d\theta d\beta e^{\gamma \log p(\theta, \beta | \mathcal{D})} \delta(\beta - \beta_\mu)}{\int d\theta d\beta e^{\gamma \log p(\theta, \beta | \mathcal{D})}} \right\rangle_{\mathcal{D}}. \end{aligned} \quad (3.4)$$

Upon inserting (3.2) into (3.4) we obtain for this \mathcal{D} -independent joint distribution:

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^*) \\ &\quad \times \left\langle \frac{\int d\theta d\beta \left[p(\beta) p(\theta) \prod_{i=1}^N p(s_i | \beta \cdot \mathbf{z}_i / \sqrt{p}, \theta) \right]^\gamma \delta(\beta - \beta_\mu)}{\int d\theta d\beta \left[p(\beta) p(\theta) \prod_{i=1}^N p(s_i | \beta \cdot \mathbf{z}_i / \sqrt{p}, \theta) \right]^\gamma} \right\rangle_{\mathcal{D}}. \end{aligned} \quad (3.5)$$

We can evaluate (3.5) using the following alternative form of the replica identity, which can be shown to be equivalent to the previous version $\langle \log Z \rangle = \lim_{n \rightarrow 0} n^{-1} \log \langle Z^n \rangle$,

$$\begin{aligned} \left\langle \frac{\int dx w(x, y) f(x)}{\int dx w(x, y)} \right\rangle_y &= \lim_{n \rightarrow 0} \left\langle \left[\int dx w(x, y) f(x) \right] \left[\int dx w(x, y) \right]^{n-1} \right\rangle_y \\ &= \lim_{n \rightarrow 0} \int \left[\prod_{\alpha=1}^n dx^\alpha \right] f(x^1) \left\langle \prod_{\alpha=1}^n w(x^\alpha, y) \right\rangle_y. \end{aligned} \quad (3.6)$$

Application of this identity to (3.5), with the choices $x \rightarrow (\theta, \beta)$, $y \rightarrow \mathcal{D}$, $w(x, y) \rightarrow [p(\beta, \theta) \prod_{i=1}^N p(s_i | \beta \cdot \mathbf{z}_i / \sqrt{p}, \theta)]^\gamma$, and $f(x) \rightarrow \delta(\beta - \beta_\mu)$, followed by working out the definition of the average over the data sets \mathcal{D} , gives

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^*) \int \left\{ \prod_{\alpha=1}^n d\theta^\alpha d\beta^\alpha [p(\beta^\alpha) p(\theta^\alpha)]^\gamma \right\} \delta(\beta - \beta_\mu^1) \\ &\quad \times \left\{ \int dz ds p(\mathbf{z}) p(s | \beta^* \cdot \mathbf{z} / \sqrt{p}, \theta^*) \prod_{\alpha=1}^n [p(s | \beta^\alpha \cdot \mathbf{z} / \sqrt{p}, \theta^\alpha)]^\gamma \right\}^N. \end{aligned} \quad (3.7)$$

In appendix D we calculate the limit $p, N \rightarrow \infty$ of (3.7), for finite ratio $\zeta = p/N$. This builds on the replica calculation in appendix A. For the simplest case of uncorrelated and normalized covariates, i.e. $\mathbf{A} = \mathbb{I}$, we find that

$$\lim_{N \rightarrow \infty} \mathcal{P}(\beta | \beta^*) = \frac{1}{v\sqrt{2\pi}} e^{-\frac{1}{2}(\beta - w\beta^*/S)^2/v^2} \quad (3.8)$$

This confirms what was suggested by simulation data and exploited in [62]: if we plot inferred versus true association parameters in a plane, we will find for $L2$ priors and uncorrelated covariates a linear cloud with slope w/S and zero-average Gaussian noise of width v . We have now proved this analytically, for *any* generalized linear model. For the more tricky case of correlated covariates, i.e. arbitrary covariance matrices \mathbf{A} subject only to the spectral conditions $\lim_{p \rightarrow \infty} \langle a \rangle \in \mathbb{R}$ and $\lim_{p \rightarrow \infty} p^{-1} \langle a^2 \rangle \in \mathbb{R}$ of appendix B, we find

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \frac{\delta(\beta^* - \beta_\mu^0)}{\sqrt{2\pi |\tilde{f}| [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A} (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\mu}}} \quad (3.9) \\ &\times e^{-\frac{1}{2} [\beta + \tilde{d}_0 [(\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1})^{-1} \beta^0]_{\mu}]^2 / |\tilde{f}| [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A} (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\mu}}. \end{aligned}$$

Once more the inferred vector $\hat{\beta}$ depends linearly on the true vector β^* , supplemented with Gaussian noise. However, in the presence of covariate correlations, we obtain a scalar relation $\hat{\beta} = \kappa\beta^* + \text{noise}$ typically only when $\eta = 0$ (i.e. no regularizer, giving ML regression).

Expression (3.9) is consistent with the more general propositions

$$\langle \hat{\beta} \rangle = -\tilde{d}_0 [\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1}]^{-1} \beta^*, \quad (3.10)$$

$$\langle \hat{\beta}_\mu \hat{\beta}_\nu \rangle - \langle \hat{\beta}_\mu \rangle \langle \hat{\beta}_\nu \rangle = |\tilde{f}| [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A} (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\nu}. \quad (3.11)$$

Using expression (D.10) for d_0 , and $c_0 = S\langle a \rangle^{\frac{1}{2}} w$, we can write (3.10) also as

$$\langle \hat{\beta} \rangle = \frac{w\langle a \rangle^{\frac{1}{2}}}{S} \left\langle \frac{a^2}{2\eta + a\tilde{g}} \right\rangle^{-1} [\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1}]^{-1} \beta^*. \quad (3.12)$$

When covariates are correlated, MAP regression with Gaussian priors sufficiently strong to have an impact will for finite $\zeta > 0$ not just rescale the length of the inferred association vectors but also change its direction. Only for small η or weak correlations (or if by accident β^* happens to be an eigenvector of \mathbf{A}) will the relation between $\langle \hat{\beta} \rangle$ and β^* reduce to scalar multiplication. This is reminiscent of what happens for linear regression, and we will explore the connection in more detail in a subsequent section. It is interesting that the form of the above expressions is universal; GLM model dependencies enter only via the order parameters $(\tilde{d}_0, \tilde{f}, \tilde{g})$.

Since the vectors and matrices in (3.10), (3.11) have diverging dimensionality as $p \rightarrow \infty$, it is not possible to derive these propositions directly using the steepest descent route followed in deriving (3.9). Only for linear ML regression will we be able to confirm (3.10), (3.11) rigorously. However, in addition to simulation experiments with different models (described in a subsequent section), one can envisage several indirect mathematical tests of expressions (3.10) and (3.11). First, they can be used to compute the two order parameters c_0 and C , testing their consistency with our RS order parameter equations derived earlier. This gives, using appendix B,

$$c_0 = \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p \langle \hat{\beta}_\mu \rangle A_{\mu\nu} \beta_\nu^* \quad (3.13)$$

$$= \frac{w \langle a \rangle^{\frac{1}{2}}}{S} \left\langle \frac{a^2}{2\eta + a\tilde{g}} \right\rangle^{-1} \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p \beta_\mu^* [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A}^2]_{\mu\nu} \beta_\nu^* = w \langle a \rangle^{\frac{1}{2}} S,$$

$$\begin{aligned} C &= \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p \langle \hat{\beta}_\mu \hat{\beta}_\nu \rangle A_{\mu\nu} \\ &= \lim_{p \rightarrow \infty} \frac{1}{p} \text{Tr} \left[|\tilde{f}| [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A}^2 (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}] \right] \\ &\quad + \frac{w^2 \langle a \rangle}{S^2} \left\langle \frac{a^2}{2\eta + a\tilde{g}} \right\rangle^{-2} \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p [(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A}^3 (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\nu} \beta_\mu^* \beta_\nu^* \\ &= |\tilde{f}| \left\langle \frac{a^2}{(2\eta + a\tilde{g})^2} \right\rangle + w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + a\tilde{g}} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + a\tilde{g})^2} \right\rangle. \end{aligned} \quad (3.14)$$

Clearly, (3.13) is identical to the result of combining the first identity of (A.26) with the expression for \tilde{S} in (A.29). Similarly, upon combining the third identity of (A.26) with $\lim_{\gamma \rightarrow \infty} u = \lim_{\gamma \rightarrow \infty} \tilde{u}/\sqrt{\gamma} = 0$, we find that (3.14) gives in the limit $\gamma \rightarrow \infty$:

$$v^2 + w^2 = |\tilde{f}| \left\langle \frac{a^2}{(2\eta + a\tilde{g})^2} \right\rangle + w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + a\tilde{g}} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + a\tilde{g})^2} \right\rangle, \quad (3.15)$$

which, in combination with $\tilde{f} < 0$, reproduces equation (2.33). Hence one can compute the correct RS order parameter equations from (3.11), (3.12). Secondly, in the ML limit $\eta \rightarrow 0$, where we know that $|\tilde{f}| = v^2 \tilde{g}^2$, the formulae are seen to simplify as follows, confirming *en passant* an ansatz made in [63]:

$$\langle \hat{\beta} \rangle = (w/S \langle a \rangle^{\frac{1}{2}}) \beta^*, \quad \langle \hat{\beta}_\mu^2 \rangle - \langle \hat{\beta}_\mu \rangle^2 = v^2 (\mathbf{A}^{-1})_{\mu\mu}. \quad (3.16)$$

As a third test we can also verify from (3.10), (3.11) our earlier results for uncorrelated and normalized covariates. Substitution of the appropriate values $\mathbf{A} = \mathbb{I}$, $\tilde{f} = -v^2/\tilde{u}^4$, $2\eta + \tilde{g} = \tilde{u}^{-2}$, and $\tilde{d}_0 = -w/S\tilde{u}^2$ into (3.10), (3.11) gives indeed the correct expressions

$$\langle \hat{\beta} \rangle = (w/S) \beta^*, \quad \langle \hat{\beta}_\mu^2 \rangle - \langle \hat{\beta}_\mu \rangle^2 = v^2. \quad (3.17)$$

3.2. Correction of association parameters for overfitting effects

To work out the replica order parameters and all associated theoretical predictions in practice, we first need to estimate the true covariate correlation matrix \mathbf{A} from the available covariate samples $\{z_1, \dots, z_N\}$ (which is a standard statistical problem in portfolio theory), and the value of S (which controls the amplitude of the unknown vector β^*). The latter can be found for sufficiently large p by evaluation of $p^{-1} \hat{\beta} \cdot \mathbf{A} \hat{\beta}$, using the outcome $\hat{\beta}$ of MAP/ML inference on the given data, in combination with equation (3.14). One similarly uses the MAP/ML inferred auxiliary parameters $\hat{\theta}$ together with their associated order parameter equations that express the link between $\hat{\theta}$ and θ^* to eliminate the need to know θ^* . Once \mathbf{A} , S , θ^* , and the solution of our

RS equations are known, expressions (3.10), (3.11) allow us to construct alternative estimators from the MAP estimator $\hat{\beta}$ of the association parameters, decontaminated from the distorting effects of overfitting. To compactify notation we first define two $p \times p$ matrices \mathbf{G} and Ξ :

$$\mathbf{G} = |\tilde{d}_0|[\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1}]^{-1}, \quad (3.18)$$

$$\Xi = |\tilde{f}|(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}\mathbf{A}(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}, \quad (3.19)$$

with which (3.10) and (3.11) become

$$\langle \hat{\beta} \rangle = \mathbf{G}\beta^*, \quad \langle \hat{\beta}_\mu \hat{\beta}_\nu \rangle - \langle \hat{\beta}_\mu \rangle \langle \hat{\beta}_\nu \rangle = \Xi_{\mu\nu}. \quad (3.20)$$

Both \mathbf{G} and Ξ are symmetric matrices, which commute, and $\Xi = |\tilde{f}|\tilde{d}_0^{-2}\mathbf{G}\mathbf{A}^{-1}\mathbf{G}$. We will limit ourselves to linear correction protocols of the form $\hat{\beta} \rightarrow \hat{\beta}^* = \mathbf{F}\hat{\beta}$, where the correction matrix \mathbf{F} is restricted to be non-stochastic. One could *a priori* envisage several natural criteria for determining \mathbf{F} , dependent upon the desired properties of the new estimator $\hat{\beta}^*$, such as:

- (a) Removal of the inference bias, i.e. $\langle \hat{\beta}^* \rangle = \beta^*$.
- (b) Minimization of the expected MSE (mean squared error) $\sum_{\mu=1}^p \langle (\hat{\beta}_\mu^* - \beta_\mu^*)^2 \rangle$.
- (c) Minimization of the expected generalization error.

In appendix E we show that, somewhat counterintuitively, minimization of the generalization error can lead to nonsensical results (an excessive bias and a hyperconfident outcome prediction model), and should therefore not be used. We will next compute the correction matrices and corresponding new estimators for the more reliable criteria (a) and (b) in explicit form.

Criterion (a), removal of inference bias, is immediately seen to require choosing $\mathbf{F}_{\text{opt}} = \mathbf{G}^{-1}$, giving the new and unbiased estimator

$$\hat{\beta}^* = |\tilde{d}_0|^{-1}[\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1}]\hat{\beta}. \quad (3.21)$$

Its variance is

$$\langle \hat{\beta}_\mu^{*2} \rangle - \langle \hat{\beta}_\mu^* \rangle^2 = (|\tilde{f}|/\tilde{d}_0^2)(\mathbf{A}^{-1})_{\mu\mu}. \quad (3.22)$$

Next we work out criterion (b) for large p , assuming the various matrices to obey the conditions of appendix B, so that we may use expressions such as $p^{-1}\sum_{\mu\nu \leq p}\beta_\mu^*\beta_\nu^*M_{\mu\nu} = (S^2/p)\text{Tr}\mathbf{M} + o(1)$. The objective function to be minimized over \mathbf{F} then becomes, after a rescaling by p to ensure that it is $\mathcal{O}(1)$ as $p \rightarrow \infty$:

$$\begin{aligned} \Omega(\mathbf{F}) &= \frac{1}{p} \langle (\mathbf{F}\hat{\beta} - \beta^*)^2 \rangle \\ &= \frac{1}{p} \sum_{\mu} \left(\sum_{\nu\rho} F_{\mu\nu}F_{\mu\rho} \langle \hat{\beta}_\nu \hat{\beta}_\rho \rangle + (\beta_\mu^*)^2 - 2\beta_\mu^* \sum_{\nu} F_{\mu\nu} \langle \hat{\beta}_\nu \rangle \right) \\ &= \frac{1}{p} \sum_{\mu\nu\rho} F_{\mu\nu}F_{\mu\rho} [\Xi_{\nu\rho} + (\mathbf{G}\beta^*)_{\nu}(\mathbf{G}\beta^*)_{\rho}] + S^2 - \frac{2}{p} \sum_{\mu\nu\rho} \beta_\mu^* F_{\mu\nu} G_{\nu\rho} \beta_\rho^* + o(1) \\ &= \frac{1}{p} \text{Tr}(\mathbf{F}\Xi\mathbf{F}^\dagger) + \frac{S^2}{p} \sum_{\mu\nu\rho\lambda} F_{\mu\nu}F_{\mu\rho} G_{\nu\lambda} G_{\rho\lambda} + S^2 - \frac{2S^2}{p} \sum_{\mu\nu} F_{\mu\nu} G_{\nu\mu} + o(1) \\ &= \Omega_1(\mathbf{F}) + \Omega_2(\mathbf{F}) + o(1). \end{aligned} \quad (3.23)$$

with

$$\Omega_1(\mathbf{F}) = \frac{1}{p} \text{Tr}(\mathbf{F}\Xi\mathbf{F}^\dagger), \quad \Omega_2(\mathbf{F}) = \frac{S^2}{p} \text{Tr}[(\mathbf{F}\mathbf{G}-\mathbb{I})^\dagger(\mathbf{F}\mathbf{G}-\mathbb{I})]. \quad (3.24)$$

Removal of the inference bias gives $\Omega_2(\mathbf{F}) = 0$ (achieved for $\mathbf{F} = \mathbf{G}^{-1}$, following the previous criterion (a)), and removal of the inference noise gives $\Omega_1(\mathbf{F}) = 0$ (achieved for $\mathbf{F} = \mathbf{0}$, or for any \mathbf{F} if $\Xi = \mathbf{0}$, i.e. if the MAP inference is already noise-free). Hence we can interpret $\Omega_1(\mathbf{F})$ as the error contribution from the noise, and $\Omega_2(\mathbf{F})$ as the error contribution from the bias. In criterion (a) we minimized $\Omega_2(\mathbf{F})$ and this would generally increase $\Omega_1(\mathbf{F})$. Minimizing (3.23) requires balancing the two error sources. This is the bias-variance trade-off in inference [80, 81]. However, since $\Omega(\mathbf{F})$ is quadratic in \mathbf{F} we can find the location of the overall minimum in explicit form:

$$\mathbf{F}_{\text{opt}} = (\Xi/S^2 + \mathbf{G}^2)^{-1}\mathbf{G}, \quad (3.25)$$

$$\Omega_{\text{min}} = \frac{1}{p} \text{Tr}(S^{-2}\mathbb{I} + \Xi^{-1}\mathbf{G}^2)^{-1}. \quad (3.26)$$

4. Applications to specific regression models

We now apply the generic replica symmetric MAP order parameter equations (2.32)–(2.37), where $\xi(\mu, \sigma, s, \theta)$ represents the solution of (2.27), to different regression models of the GLM family. We test the predictions of the theory for MAP and ML regression against measurements of simulations with different outcome types and models, and with synthetic data. In all cases we will for simplicity choose the covariate distribution $p(\mathbf{z})$ to be Gaussian, with zero average but potentially correlated components $\{z_\mu\}$. For the $p \times p$ covariance matrix \mathbf{A} with entries $A_{\mu\nu} = \langle z_\mu z_\nu \rangle$ we will make the following choice, with $\epsilon \in [0, 1]$:

$$\begin{aligned} A_{\mu\mu} &= 1 \\ A_{\mu, \mu+1} &= A_{\mu+1, \mu} = \epsilon \\ A_{\mu\nu} &= 0 \text{ for all other entries.} \end{aligned} \quad (4.1)$$

This describes pairwise correlated covariates. The matrix (4.1) obeys the conditions in appendix B, and is trivially diagonalised to give $\varrho(a) = \frac{1}{2}\delta(a-1+\epsilon) + \frac{1}{2}\delta(a-1-\epsilon)$, enabling precise tests of the predictions of the theory.

4.1. Linear regression

Replica equations for MAP linear regression. The simplest case of a GLM corresponds to linear regression, where the outcomes of ML and MAP regression can in principle be computed in explicit form. It therefore serves as the simplest test for our general equations. In linear regression we have $\theta = (\beta_0, \Sigma) \in \mathbb{R}^2$ and

$$p(s|\xi, \theta) = (2\pi\Sigma^2)^{-\frac{1}{2}} e^{-\frac{1}{2}(s-\xi-\beta_0)^2/\Sigma^2}. \quad (4.2)$$

Here we find that

$$\xi(\mu, \sigma, s, \theta) = \frac{\mu\Sigma^2 + \sigma^2(s - \beta_0)}{\Sigma^2 + \sigma^2}. \quad (4.3)$$

Hence $(\partial_1 \xi)(\mu, \sigma, s, \theta) = \Sigma^2 / (\sigma^2 + \Sigma^2)$, and upon working out the relevant derivatives of $p(s|\xi, \theta)$, we find the following closed set of MAP order equations:

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \quad (4.4)$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2, \quad (4.5)$$

$$\int \mathrm{D}t \left\langle \left\langle (\beta_0^* - \beta_0 + (S\langle a \rangle^{\frac{1}{2}} - w)y_0 + \Sigma^* t - vz)^2 \right\rangle_s \right\rangle = -\zeta \tilde{f} (\Sigma^2 + \tilde{u}^2)^2, \quad (4.6)$$

$$\frac{1}{\tilde{u}^2 + \Sigma^2} = \zeta \tilde{g}, \quad (4.7)$$

$$\frac{S}{\Sigma^2 + \tilde{u}^2} = \frac{\zeta w \langle a \rangle^{\frac{1}{2}}}{\left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle}, \quad (4.8)$$

$$\beta_0^* - \beta_0 = 0, \quad (4.9)$$

$$\int \mathrm{D}t \left\langle \left\langle [\beta_0^* - \beta_0 + (S\langle a \rangle^{\frac{1}{2}} - w)y_0 + \Sigma^* t - vz]^2 \right\rangle_s \right\rangle = \frac{(\Sigma^2 + \tilde{u}^2)^2}{\Sigma^2}. \quad (4.10)$$

Thus one always has $\beta_0 = \beta_0^*$, and the other equations can be compactified to

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \quad 1/\Sigma^2 = -\zeta \tilde{f}, \quad \frac{1}{\tilde{u}^2 + \Sigma^2} = \zeta \tilde{g}, \quad (4.11)$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2, \quad (4.12)$$

$$\frac{S}{\Sigma^2 + \tilde{u}^2} = \zeta w \langle a \rangle^{\frac{1}{2}} \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1}, \quad (4.13)$$

$$(S\langle a \rangle^{\frac{1}{2}} - w)^2 + v^2 + \Sigma^{*2} = (\Sigma^2 + \tilde{u}^2)^2 / \Sigma^2. \quad (4.14)$$

Via substitutions one can reduce these coupled equations to a single nonlinear equation for \tilde{g} , the numerical solution of which then generates the other order parameters ($v, w, \tilde{f}, \tilde{u}, \Sigma$). With the short-hand $\alpha_{k\ell} = \langle a^k / (2\eta + \tilde{g}a)^\ell \rangle$ this equation takes the form

$$\tilde{g}^{-1} = \zeta (1 - \zeta \tilde{g} \alpha_{11}) [S^2 (\alpha_{10} - 2\tilde{g} \alpha_{21} + \tilde{g}^2 \alpha_{32}) + \Sigma^{*2}] + \zeta \tilde{g} \alpha_{22}. \quad (4.15)$$

Similarly, using the above formulae we can also simplify the predictions (3.10), (3.11) to

$$\langle \hat{\beta} \rangle = (\mathbb{I} + \frac{2\eta}{\tilde{g}} \mathbf{A}^{-1})^{-1} \beta^*, \quad (4.16)$$

$$\langle \hat{\beta}_\mu^2 \rangle - \langle \hat{\beta}_\mu \rangle^2 = \frac{1}{\zeta \Sigma^2 \tilde{g}^2} [(\mathbb{I} + \frac{2\eta}{\tilde{g}} \mathbf{A}^{-1}) \mathbf{A}^{-1} (\mathbb{I} + \frac{2\eta}{\tilde{g}} \mathbf{A}^{-1})]_{\mu\mu}. \quad (4.17)$$

For uncorrelated covariates, i.e. $\mathbf{A} = \mathbb{I}$, these results are consistent with the well-known asymptotic behaviour of linear estimators with large random measurement matrices [82–85]. Setting

$\eta = 0$ brings us from MAP regression to ML regression. Here we find that the above equations reduce after some simple manipulations to

$$w = S\langle a \rangle^{\frac{1}{2}}, \quad \Sigma = \Sigma^* \sqrt{1-\zeta}, \quad \tilde{u} = \Sigma^* \sqrt{\zeta}, \quad v = \Sigma^* \sqrt{\frac{\zeta}{1-\zeta}}, \quad (4.18)$$

$$\langle \hat{\beta} \rangle = \beta^*, \quad \langle \hat{\beta}_\mu^2 \rangle - \langle \hat{\beta}_\mu \rangle^2 = \frac{\zeta}{1-\zeta} (\Sigma^*)^2 (\mathbf{A}^{-1})_{\mu\mu}. \quad (4.19)$$

Thus also the association parameters will on average be inferred correctly in ML, but there will be increasing overfitting induced noise (diverging at the transition point $\zeta = 1$), and underestimation of the true uncertainty Σ^* in the outcome predictions.

Direct solution. For linear regression we can go beyond testing the replica predictions against numerical simulations, since the regression problem allows for exact solution. The parameter to be inferred are β and Σ , whose MAP estimators are

$$(\hat{\beta}, \hat{\Sigma}) = \arg \min_{\beta, \Sigma} \left\{ \frac{1}{2\Sigma^2} \sum_{i=1}^N (s_i - \frac{\beta \cdot \mathbf{z}_i}{\sqrt{p}})^2 + N \log \Sigma + \eta \beta^2 \right\}. \quad (4.20)$$

This minimization results in the following coupled equations, with the empirical $p \times p$ covariance matrix $\hat{\mathbf{A}}$ with entries $\hat{A}_{\mu\nu} = N^{-1} \sum_{i \leq N} z_{i\mu} z_{i\nu}$:

$$\hat{\Sigma}^2 = \frac{1}{N} \sum_{i=1}^N (s_i - \frac{\hat{\beta} \cdot \mathbf{z}_i}{\sqrt{p}})^2, \quad (4.21)$$

$$\hat{\beta}_\mu = \sum_{\nu} \left(2\eta\zeta \hat{\Sigma}^2 \mathbb{I} + \hat{\mathbf{A}} \right)_{\mu\nu}^{-1} \frac{\zeta}{\sqrt{p}} \sum_{i=1}^N z_{i\nu} s_i. \quad (4.22)$$

The direct solution is formulated in terms of the empirical covariate covariance matrix $\hat{\mathbf{A}}$, whereas the replica analysis involves the true population covariance matrix \mathbf{A} . To understand the connection between the two descriptions, we need to compute disorder-averaged quantities from the above equations. To do this, we assume, as in the replica analysis, that the data are generated by a linear model of the type (4.2), with unknown parameters (β^*, Σ^*) . Hence $s_i = \beta^* \cdot \mathbf{z}_i / \sqrt{p} + \Sigma^* \xi_i$, in which all ξ are i.i.d. random variables, drawn from $p(\xi) = (2\pi)^{-\frac{1}{2}} e^{-\frac{1}{2}\xi^2}$. We will show below that in ML regression $\hat{\Sigma}$ is self-averaging for $p \rightarrow \infty$, so that for large p we can evaluate the distribution of inferred association parameters, averaged over all possible realizations of the data, i.e. over all $\{\xi_i, \mathbf{z}_i\}$:

$$\begin{aligned} P(\hat{\beta}) &= \left\langle \left\langle \delta \left[\hat{\beta} - (2\eta\zeta \hat{\Sigma}^2 \mathbb{I} + \hat{\mathbf{A}})^{-1} \left(\hat{\mathbf{A}} \beta^* + \frac{\zeta \Sigma^*}{\sqrt{p}} \sum_{i=1}^N z_i \xi_i \right) \right] \right\rangle \right\rangle_{\{\xi, \mathbf{z}\}} \\ &= \left\langle \int \frac{d\mathbf{x}}{(2\pi)^d} e^{i\mathbf{x} \cdot (\hat{\beta} - \hat{\mathbf{G}} \beta^*) - \frac{1}{2} \zeta (\Sigma^*)^2 \mathbf{x} \cdot \hat{\mathbf{G}} \hat{\mathbf{A}}^{-1} \hat{\mathbf{G}} \mathbf{x}} \right\rangle_{\{\mathbf{z}\}} \\ &= \int d\hat{\mathbf{A}} W(\hat{\mathbf{A}}) \mathcal{N}(\hat{\beta} | \hat{\mathbf{G}} \beta^*, 2\zeta \Sigma^{*2} \hat{\mathbf{G}} \hat{\mathbf{A}}^{-1} \hat{\mathbf{G}}), \end{aligned} \quad (4.23)$$

with $\hat{\mathbf{G}} = (\mathbb{I} + 2\eta\zeta \hat{\Sigma}^2 \hat{\mathbf{A}}^{-1})^{-1}$, and

$$\mathcal{N}(\beta | \mu, \Sigma) = \frac{e^{-\frac{1}{2}(\beta - \mu) \cdot \Sigma^{-1} (\beta - \mu)}}{\sqrt{(2\pi)^d \text{Det } \Sigma}}, \quad (4.24)$$

$$P(\hat{\mathbf{A}}) = \left\langle \prod_{\mu\nu=1}^p \delta \left[\hat{A}_{\mu\nu} - \frac{1}{N} \sum_{i=1}^N z_{i\mu} z_{i\nu} \right] \right\rangle_{\{z\}}. \quad (4.25)$$

Thus $P(\hat{\beta})$ is an average of Gaussian distributions, each weighted by the measure $P(\hat{\mathbf{A}})$ of empirical covariate covariance matrices. The integral in (4.23) is still defined over all $p \times p$ matrices $\hat{\mathbf{A}}$.

In appendix F we evaluate (4.25) and (4.23) further for the choice $p(z) = [(2\pi)^{-p} \text{Det } \mathbf{A}]^{\frac{1}{2}} e^{-\frac{1}{2}z \cdot \mathbf{A}z}$, and show that here $P(\hat{\beta})$ can be written as the following integral over the space Ω_p of symmetric positive definite matrices, involving the Wishart distribution $W(\hat{\mathbf{A}})$ [87] with N degrees of freedom:

$$P(\hat{\beta}) = \int_{\Omega_p} d\hat{\mathbf{A}} W(\hat{\mathbf{A}}) \mathcal{N}(\hat{\beta} | \hat{\mathbf{G}}\beta^*, 2\zeta \Sigma^{*2} \hat{\mathbf{G}}\hat{\mathbf{A}}^{-1} \hat{\mathbf{G}}), \quad (4.26)$$

with

$$W(\hat{\mathbf{A}}) = \frac{e^{-\frac{1}{2}N \text{Tr}(\hat{\mathbf{A}}\mathbf{A}^{-1})} (\text{Det } \hat{\mathbf{A}})^{\frac{1}{2}(N-p-1)}}{\mathcal{Z}(\mathbf{A})} \quad (4.27)$$

$$\begin{aligned} \mathcal{Z}(\mathbf{A}) &= \int_{\Omega_p} d\hat{\mathbf{A}} e^{-\frac{1}{2}N \text{Tr}(\hat{\mathbf{A}}\mathbf{A}^{-1})} (\text{Det } \hat{\mathbf{A}})^{\frac{1}{2}(N-p-1)} \\ &= \left(\frac{2}{N}\right)^{Np/2} \pi^{\frac{1}{4}p(p-1)} (\text{Det } \mathbf{A})^{\frac{1}{2}N} \prod_{j=\frac{1}{2}(N-p+1)}^{\frac{1}{2}N} \Gamma(j). \end{aligned} \quad (4.28)$$

From the properties of the Wishart distribution follow average and variance of the entries of $\hat{\mathbf{A}}$, which confirm, as expected, that $\hat{A}_{\mu\nu} = A_{\mu\nu} + \mathcal{O}(N^{-\frac{1}{2}})$:

$$\langle \hat{A}_{\mu\nu} \rangle = A_{\mu\nu}, \quad \langle \hat{A}_{\mu\nu}^2 \rangle - \langle \hat{A}_{\mu\nu} \rangle^2 = N^{-1}A_{\mu\nu} + N^{-2}A_{\mu\mu}A_{\nu\nu}. \quad (4.29)$$

The integral in (4.26) is still nontrivial, so we now focus on the case of linear ML regression and take the limit $\eta \rightarrow 0$, where our previous result simplifies considerably. We define the $p \times p$ matrix with entries $C_{\mu\nu} = p^{-1}(\hat{\beta}_\mu - \beta_\mu^*)(\hat{\beta}_\nu - \beta_\nu^*)$. This leads to

$$\begin{aligned} P(\hat{\beta}) &= \int_{\Omega_p} \frac{d\hat{\mathbf{A}} e^{-\frac{1}{2}N \text{Tr}(\hat{\mathbf{A}}\mathbf{A}^{-1})} (\text{Det } \hat{\mathbf{A}})^{\frac{1}{2}(N-p-1)}}{\mathcal{Z}(\mathbf{A})} \mathcal{N}(\beta | \beta^*, \zeta(\Sigma^*)^2 \hat{\mathbf{A}}^{-1}) \\ &= \frac{[2\pi\zeta(\Sigma^*)^2]^{-p/2}}{\mathcal{Z}(\mathbf{A})} \int_{\Omega_p} d\hat{\mathbf{A}} e^{-\frac{1}{2}N \text{Tr}[\hat{\mathbf{A}}(\mathbf{A}^{-1} + C/(\Sigma^*)^2)]} (\text{Det } \hat{\mathbf{A}})^{\frac{1}{2}(N-p)}. \end{aligned} \quad (4.30)$$

This is again an integral of the Wishart form that can be evaluated analytically, now with $N + 1$ degrees of freedom. Thus we get

$$P(\hat{\beta}) = \frac{N^{Np/2 - \frac{1}{2}p(N+1)} \Gamma(\frac{1}{2}(N+1)) \sqrt{\text{Det } \mathbf{A}}}{[\pi\zeta(\Sigma^*)^2]^{p/2} \Gamma(\frac{1}{2}(N-p+1))} [\text{Det}(\mathbb{I} + \mathbf{A}C/(\Sigma^*)^2)]^{-\frac{N+1}{2}}. \quad (4.31)$$

Finally we use the identity

$$\text{Det}(\mathbb{I} + \mathbf{A}C/(\Sigma^*)^2) = 1 + \frac{(\hat{\beta} - \beta^*) \cdot \mathbf{A}(\hat{\beta} - \beta^*)}{p(\Sigma^*)^2} \quad (4.32)$$

to show that $P(\hat{\beta})$ is for any (p, N) a multivariate student's t -distribution with $N - p + 1$ degrees of freedom:

$$P(\hat{\beta}) = \pi^{-p/2} \frac{\Gamma(\frac{1}{2}(N+1))}{\Gamma(\frac{1}{2}(N-p+1))} \frac{\sqrt{\text{Det } \mathbf{A}}}{[p(\Sigma^*)^2]^{p/2}} \left[1 + \frac{(\hat{\beta} - \beta^*) \cdot \mathbf{A}(\hat{\beta} - \beta^*)}{p(\Sigma^*)^2} \right]^{-\frac{N+1}{2}}. \quad (4.33)$$

Equivalently we can write $\hat{\beta} = \beta^* + \mathbf{A}^{\frac{1}{2}} \mathbf{x}$, where

$$P(\mathbf{x}) = \frac{\Gamma(\frac{1}{2}(N+1))}{\Gamma(\frac{1}{2}(N-p+1))} \frac{\pi^{-p/2}}{[p(\Sigma^*)^2]^{p/2}} \left[1 + \frac{\mathbf{x}^2}{p(\Sigma^*)^2} \right]^{-\frac{N+1}{2}} \quad (4.34)$$

Mean and covariance matrix of (4.33) are in the limit $N, p \rightarrow \infty$, with $p/N = \zeta$ fixed, exactly as predicted by the replica theory, since

$$\langle \hat{\beta}_\mu \rangle = \beta_\mu^*, \quad (4.35)$$

$$\langle \hat{\beta}_\mu \hat{\beta}_\nu \rangle - \langle \hat{\beta}_\mu \rangle \langle \hat{\beta}_\nu \rangle = \frac{\zeta(\Sigma^*)^2}{1-\zeta} (\mathbf{A}^{-1})_{\mu\nu} + \mathcal{O}\left(\frac{1}{N}\right). \quad (4.36)$$

Along the same lines one can also compute higher order moments of $P(\hat{\beta})$, giving results such as

$$\left\langle \left[\frac{1}{p} (\hat{\beta} - \beta^*)^2 \right]^2 \right\rangle - \left\langle \frac{1}{p} (\hat{\beta} - \beta^*)^2 \right\rangle^2 = \mathcal{O}(p^{-1}) \quad (4.37)$$

Although (4.33) is itself not a Gaussian distribution, for the marginal distribution of any finite set of components of $\hat{\beta}$ it predicts Gaussian statistics in the limit $p, N \rightarrow \infty$ with fixed ratio $\zeta = p/N$. It is a general property of the multivariate student's t -distribution that all its marginals also obey multivariate student's t -distributions [89]. Let us define the set of indices corresponding to non-marginalized components of $\hat{\beta}$ as $\mathcal{S} \subset \{1, \dots, p\}$, and write these components as $\tilde{\beta} = \{\beta_\mu, \mu \in \mathcal{S}\}$. We also define an $|\mathcal{S}| \times |\mathcal{S}|$ matrix $\tilde{\mathbf{A}}$, defined by the property that $(\tilde{\mathbf{A}}^{-1})_{\mu\nu} = (\mathbf{A}^{-1})_{\mu\nu}$ for all $(\mu, \nu) \in |\mathcal{S}|^2$. Then the marginal distribution for $\tilde{\beta}$ is [89]:

$$P(\tilde{\beta}) \propto \left[1 + \frac{(\tilde{\beta} - \tilde{\beta}^*) \cdot \tilde{\mathbf{A}}(\tilde{\beta} - \tilde{\beta}^*)}{p(\Sigma^*)^2} \right]^{-\frac{1}{2}(N-p+1+|\mathcal{S}|)}. \quad (4.38)$$

For $|\mathcal{S}|$ finite and $p, N \rightarrow \infty$ with $p/N = \zeta$ fixed, we can expand this and find

$$P(\tilde{\beta}) \propto e^{-\frac{1}{2} \frac{1-\zeta}{\zeta(\Sigma^*)^2} (\tilde{\beta} - \tilde{\beta}^*) \cdot \tilde{\mathbf{A}} (\tilde{\beta} - \tilde{\beta}^*) + \mathcal{O}(p^{-1})}. \quad (4.39)$$

So in the relevant limit, exact evaluation of $P(\hat{\beta})$ gives for linear ML regression a Gaussian distribution for the marginals (if $|\mathcal{S}|$ is finite), with, upon using $(\tilde{\mathbf{A}}^{-1})_{\mu\nu} = (\mathbf{A}^{-1})_{\mu\nu}$ and in accordance with (4.35), (4.36),

$$\mu, \nu \in \mathcal{S} : \langle \hat{\beta}_\mu \rangle = \beta_\mu^*, \quad \langle \hat{\beta}_\mu \hat{\beta}_\nu \rangle - \langle \hat{\beta}_\mu \rangle \langle \hat{\beta}_\nu \rangle = \frac{\zeta(\Sigma^*)^2}{1-\zeta} (\mathbf{A}^{-1})_{\mu\nu}. \quad (4.40)$$

We conclude from the above analysis that, in those cases where exact evaluation enables direct comparison with the predictions of the replica theory (i.e. for linear ML regression), there is full agreement between the two, and that the two propositions (3.10) and (3.11) hold. Going

beyond ML to do the same test for linear MAP regression requires evaluation of the integral in (4.26), which we have so far been unable to do.

The direct calculation of the statistics of the inferred noise parameter $\hat{\Sigma}$ in ML linear regression also confirms the replica prediction. After some simple manipulations one finds that $\hat{\Sigma}^2$ can be written in terms of the data as

$$\hat{\Sigma}^2 = \frac{\Sigma^{*2}}{N} \sum_{i,j=1}^N \xi_i \left\{ \delta_{ij} - \frac{1}{N} \mathbf{z}_i \cdot \hat{\mathbf{A}}^{-1} \mathbf{z}_j \right\} \xi_j. \quad (4.41)$$

We define the $p \times N$ matrix \mathbf{Z} with entries $Z_{\mu i} = z_{i\mu} / \sqrt{N}$. The characteristic function of the distribution $P(\hat{\Sigma}^2)$ of $\hat{\Sigma}^2$ over the realizations of the outcome noise $\boldsymbol{\xi}$ is

$$\begin{aligned} \phi(k) &= \int \mathcal{D}\boldsymbol{\xi} \, e^{ik \frac{\Sigma^{*2}}{N} \boldsymbol{\xi} \cdot [\mathbb{I} - \mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z}] \boldsymbol{\xi}} \\ &= \left[\text{Det} \left(\left(1 - \frac{2ik}{N} \Sigma^{*2}\right) \mathbb{I} + \frac{2ik}{N} \Sigma^{*2} \mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z} \right) \right]^{-\frac{1}{2}}. \end{aligned} \quad (4.42)$$

We note that $\mathbf{Z}\mathbf{Z}^\dagger = \hat{\mathbf{A}}$, from which it follows in turn that $(\mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z})^2 = \mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z}$. Hence $\mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z}$ is a projection matrix, with eigenvalues 0 and 1. Moreover, since $\text{Tr}(\mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z}) = p$, we know in fact that it has precisely ζN eigenvalues 1 and $(1 - \zeta)N$ eigenvalues 0. Hence, for any realization of the covariates we have

$$\text{Det} \left[\left(1 - \frac{2ik\Sigma^{*2}}{N}\right) \mathbb{I} + \frac{2ik\Sigma^{*2}}{N} \mathbf{Z}^\dagger \hat{\mathbf{A}}^{-1} \mathbf{Z} \right] = \left(1 - \frac{2ik\Sigma^{*2}}{N}\right)^{(1-\zeta)N}, \quad (4.43)$$

so that

$$\phi(k) = \left(1 - \frac{2ik\Sigma^{*2}}{N}\right)^{-\frac{1}{2}(1-\zeta)N}. \quad (4.44)$$

We recognize that this is the characteristic function of the gamma distribution, with average $(1 - \zeta)\Sigma^{*2}$ and width $\Sigma^{*2} \sqrt{2(1-\zeta)}/N$. Hence $\hat{\Sigma}$ obeys the gamma distribution and is self-averaging with respect to the realization of the data for $N \rightarrow \infty$, and $\lim_{N \rightarrow \infty} \hat{\Sigma} = \Sigma^* \sqrt{1-\zeta}$, confirming the prediction in (4.18) of the replica theory.

Numerical simulations of MAP linear regression with correlated covariates. The result of solving numerically the RS order parameter equations in the presence of covariate correlations of the type (4.1) is shown in figure 1, where we plot the resulting values of the order parameters v and w and the inferred noise strength Σ together with regression simulation data (for synthetic Gaussian covariates), as functions of the ratio $\zeta = p/N$. In these experiments we chose $\beta_0 = \beta_0^* = 0$, for simplicity. Once more we observe excellent agreement between theory and simulation. In the top row we also plot for each parameter combination the MAP-inferred parameters $\hat{\beta}_\mu$ versus the corresponding true association strengths β_μ^* , for pooled data from 20 regressions. In the two columns on the right we see that the width of the data cloud (top picture) reflects inference noise in the case where $\epsilon = 0$ (no covariate correlations), with a larger v , whereas for $\epsilon = 0.75$ the inference noise is reduced, so that there the wider cloud reflects the correlation-induced bias.

These data also enable us to test our two protocols (3.21), (3.25) for correcting the MAP estimator of the association parameters for the distortions caused by overfitting. See figure 2. We plot the slopes of the data clouds of estimators versus true parameter values, as shown for $\zeta = 0.5$ in the top row of figure 1, including the MAP estimator (red circles), the minimum




Figure 1. Results of linear MAP regression simulations with $Np = 400\,000$ and $\beta_0^* = 0$, for different combinations of η (regularizer), ϵ (covariate correlations), and Σ^* (true noise strength). In all cases $S = 2$. Top row: inferred versus true association parameters for $\zeta = 0.5$. Second and third row: order parameters w and v plotted versus ζ . Bottom row: inferred noise strength Σ versus ζ . Each simulation data point represents average and standard deviation computed over 400 synthetic data sets and regressions. Solid curves: theoretical predictions obtained by solving the RS equations.

bias estimator (3.21) (blue squares), and the minimum MSE estimator (3.25) (green crosses). For linear regression with small regularizers (i.e. small η) the slopes of the data clouds for the MAP estimator are very close to unity, the inference bias is weak, and hence the red and blue data points coincide (debiasing makes no difference). For stronger correlations and stronger regularization (column on the right), this is no longer true. Similarly, the MSE values of the minimum MSE estimator (3.25) (green) are as predicted indeed always identical to or below those of the other two estimators.




Figure 2. Tests of the correction protocols for MAP estimators, applied to the data of the previous figure. Red circles: the MAP estimator $\hat{\beta}_\mu$. Blue squares: the corrected estimator (3.21), aimed at removing inference bias. Green crosses: the corrected estimator (3.25), aimed at minimizing the MSE. The top row show as a function of ζ the slopes of the relation between the three estimators and the true values β_μ^* of association parameters (the slopes of the data clouds in the top row of the previous figure, computed via least squares analysis); this slope would be unity for unbiased estimators. The bottom row shows the values of $\sqrt{\text{MSE}}$, where $\text{MSE} = p^{-1} \sum_{\mu \leq p} (\hat{\beta}_\mu - \beta_\mu^*)^2$.

4.2. Logistic regression

Equations for MAP logistic regression. In logistic regression we have $s \in \{-1, 1\}$ (alternatively one could define $s \in \{0, 1\}$; with our present choice the equations will be somewhat more compact), with $\theta = \beta_0 \in \mathbb{R}$, and

$$p(s|\xi, \beta_0) = \frac{e^{s(\xi + \beta_0)}}{2 \cosh(\xi + \beta_0)}. \tag{4.45}$$

Hence $\partial_\xi \log p(s|\xi, \beta_0) = s - \tanh(\xi + \beta_0)$ and $\partial_\xi^2 \log p(s|\xi, \beta_0) = \tanh^2(\xi + \beta_0) - 1$. We will now compute the various model-dependent building blocks of our general order parameter equations (2.32)–(2.37). The function $\xi(\mu, \sigma, s, \beta_0)$ is the solution of

$$s - \tanh(\xi + \beta_0) = (\xi - \mu)/\sigma^2. \tag{4.46}$$

We switch from ξ to the new variable $x = s(\xi + \beta_0)$, so $sx - \beta_0 = \xi$. Now

$$1 - [x - s(\beta_0 + \mu)]/\sigma^2 = \tanh(x). \tag{4.47}$$

We next define $\tilde{x}(\mu, \sigma)$ as the solution of the following transcendental equation, whose solution is unique since the right-hand side increases monotonically from -1 to 1 , and the left-hand side decreases monotonically from $+\infty$ to $-\infty$:

$$\tilde{x}(\mu, \sigma) : \text{solution of } \tanh(x) = 1 - (x - \mu)/\sigma^2. \tag{4.48}$$

Graphical inspection shows that $\tilde{x}(\mu, \sigma)$ increases monotonically with both μ and $\sigma \geq 0$, and that $\mu = \tilde{x}(\mu, 0) \leq \tilde{x}(\mu, \sigma) \leq \tilde{x}(\mu, \infty) = \infty$. In terms of \tilde{x} we may write

$$\xi(\mu, \sigma, s, \beta_0) = s\tilde{x}(s(\beta_0 + \mu), \sigma) - \beta_0. \tag{4.49}$$

We can now work out the relevant derivatives required in our equations:

$$(\partial_1 \xi)(\mu, \sigma, s, \beta_0) = [1 + \sigma^2 [1 - \tanh^2(\xi + \beta_0)]]^{-1}, \tag{4.50}$$

$$\frac{\partial}{\partial \beta_0} \log p(s|\xi, \beta_0) = s - \tanh(\xi + \beta_0), \tag{4.51}$$

$$\frac{\partial}{\partial y_0} \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) = S\langle a \rangle^{\frac{1}{2}} [s - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*)]. \tag{4.52}$$

Upon substituting the above model-specific expressions for logistic regression into our general RS order parameter equations (2.32)–(2.37), we obtain

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \tag{4.53}$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2, \tag{4.54}$$

$$\langle \langle [s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) - \beta_0 - wy_0 - vz]^2 \rangle_s \rangle \rangle = -\zeta \tilde{f} \tilde{u}^4, \tag{4.55}$$

$$\langle \langle \langle [1 + \tilde{u}^2 [1 - \tanh^2(s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}))]]^{-1} \rangle_s \rangle \rangle = 1 - \zeta \tilde{g} \tilde{u}^2, \tag{4.56}$$

$$\langle \langle \langle [s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) - \beta_0] [s - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*)] \rangle_s \rangle \rangle = \frac{\zeta w \tilde{u}^2 \langle a \rangle^{\frac{1}{2}}}{S \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle}, \tag{4.57}$$

$$\langle \langle \langle s - \tanh(s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u})) \rangle_s \rangle \rangle = 0. \tag{4.58}$$

Equations for ML logistic regression. For $\eta = 0$ we revert from MAP to ML regression. Here we find the usual model-independent simplifications $\tilde{g} = 1/\tilde{u}^2$ and $\tilde{f} = -v^2/\tilde{u}^4$, the covariate correlations (if present) drop out of the theory, and the remaining equations simplify to

$$\langle \langle \langle [s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) - \beta_0 - wy_0 - vz]^2 \rangle_s \rangle \rangle = \zeta v^2, \tag{4.59}$$

$$\langle \langle \langle [1 + \tilde{u}^2 [1 - \tanh^2(\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}))]]^{-1} \rangle_s \rangle \rangle = 1 - \zeta, \tag{4.60}$$

$$\langle \langle \langle [s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) - \beta_0] [s - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*)] \rangle_s \rangle \rangle = \frac{\zeta w}{S\langle a \rangle^{\frac{1}{2}}}, \tag{4.61}$$

$$\langle \langle \langle s - \tanh(s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u})) \rangle_s \rangle \rangle = 0. \tag{4.62}$$

For numerical evaluation it is helpful to write (4.62) in an alternative form, exploiting the equation that defines $\tilde{x}(\mu, \sigma)$: $\tanh(\tilde{x}(\mu, \sigma)) = 1 - [\tilde{x}(\mu, \sigma) - \mu]/\sigma^2$. We see that

$$\begin{aligned} \tanh(s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u})) &= s \tanh(\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u})) \\ &= s - \frac{1}{\tilde{u}^2} [s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) - (\beta_0 + wy_0 + vz)]. \end{aligned} \tag{4.63}$$

This enables us to write (4.62) as

$$\beta_0 = \langle\langle s\tilde{x}(s(\beta_0 + wy_0 + vz), \tilde{u}) \rangle\rangle_s. \quad (4.64)$$

Upon finally writing in explicit form all the averages, and after some simple rewriting, we obtain four equations that can be solved numerically via fixed-point iteration:

$$\zeta v^2 = \int Dy_0 Dz \left\{ \frac{1}{2} \left[1 + \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] [\tilde{x}(\beta_0 + wy_0 + vz, \tilde{u}) - (\beta_0 + wy_0 + vz)]^2 + \frac{1}{2} \left[1 - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] [\tilde{x}(-(\beta_0 + wy_0 + vz), \tilde{u}) + (\beta_0 + wy_0 + vz)]^2 \right\}, \quad (4.65)$$

$$\zeta = \int Dy_0 Dz \left\{ \frac{1}{2} \left[1 + \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] \frac{\tilde{u}^2 [1 - \tanh^2(\tilde{x}(\beta_0 + wy_0 + vz, \tilde{u}))]}{1 + \tilde{u}^2 [1 - \tanh^2(\tilde{x}(\beta_0 + wy_0 + vz, \tilde{u}))]} + \frac{1}{2} \left[1 - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] \frac{\tilde{u}^2 [1 - \tanh^2(\tilde{x}(-(\beta_0 + wy_0 + vz), \tilde{u}))]}{1 + \tilde{u}^2 [1 - \tanh^2(\tilde{x}(-(\beta_0 + wy_0 + vz), \tilde{u}))]} \right\}, \quad (4.66)$$

$$w = \frac{S\langle a \rangle^{\frac{1}{2}}}{2\zeta} \int Dy_0 Dz \left[1 - \tanh^2(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] [\tilde{x}(\beta_0 + wy_0 + vz, \tilde{u}) + \tilde{x}(-(\beta_0 + wy_0 + vz), \tilde{u})], \quad (4.67)$$

$$\beta_0 = \int Dy_0 Dz \left\{ \frac{1}{2} \left[1 + \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] \tilde{x}(\beta_0 + wy_0 + vz, \tilde{u}) - \frac{1}{2} \left[1 - \tanh(S\langle a \rangle^{\frac{1}{2}} y_0 + \beta_0^*) \right] \tilde{x}(-(\beta_0 + wy_0 + vz), \tilde{u}) \right\}. \quad (4.68)$$

ML logistic regression for data with zero offset. Further simplifications arise when we have $\beta_0^* = 0$. By symmetry of the Gaussian averages we now immediately obtain $\beta_0 = 0$ in (4.68), which leaves us with just three coupled equations to be solved numerically. Upon using wherever possible the symmetry of the Gaussian averages these final equations take the relatively simple form:

$$\zeta v^2 = \int Dy_0 Dz \left[1 + \tanh(S\langle a \rangle^{\frac{1}{2}} y_0) \right] [\tilde{x}(wy_0 + vz, \tilde{u}) - (wy_0 + vz)]^2, \quad (4.69)$$

$$\zeta = \int Dy_0 Dz \left[1 + \tanh(S\langle a \rangle^{\frac{1}{2}} y_0) \right] \frac{\tilde{u}^2 [1 - \tanh^2(\tilde{x}(wy_0 + vz, \tilde{u}))]}{1 + \tilde{u}^2 [1 - \tanh^2(\tilde{x}(wy_0 + vz, \tilde{u}))]}, \quad (4.70)$$

$$\zeta w = S\langle a \rangle^{\frac{1}{2}} \int Dy_0 Dz \left[1 - \tanh^2(S\langle a \rangle^{\frac{1}{2}} y_0) \right] \tilde{x}(wy_0 + vz, \tilde{u}). \quad (4.71)$$

In figure 3 we plot the resulting values of the order parameters v and w , whose physical meaning is given in (2.21), (2.22), as functions of $\zeta = p/N$, together with the corresponding results of regression simulations on synthetic data with Gaussian covariates. The agreement between theory and simulations is very good.

MAP logistic regression with correlated covariates. The result of solving numerically the MAP equations (4.53)–(4.58) in the presence of covariate correlations of the type (4.1) is shown in figure 4, where we plot the resulting values of the order parameters v and w together with regression simulation data (for synthetic Gaussian covariates) as functions of the ratio




Figure 3. Regression simulations (markers) versus theoretical predictions (lines) for logistic ML regression with $Np = 400\,000$, $\mathbf{A} = \mathbb{I}$, and $\beta_0^* = 0$. Left: order parameter w versus ζ , for different values of S . Right: order parameter v versus ζ , for different values of S . Each simulation data point represents average and standard deviation computed over 400 synthetic data sets and regressions.

$\zeta = p/N$. In these experiments we chose $\beta_0 = \beta_0^* = 0$; we will address the intercept parameter below. Once more we observe excellent agreement between theory and simulation. In the top row we also plot for each parameter combination the MAP-inferred parameters $\hat{\beta}_\mu$ versus the corresponding true association strengths $\beta_{\mu^*}^*$, for pooled data from 20 regressions and $\zeta = 0.5$.

Again we can also for logistic regression test our two protocols (3.21), (3.25) for correcting the MAP estimator of the association parameters for the distortions caused by overfitting. See figure 5. The slopes of the data clouds of estimators versus true parameter values, as shown for $\zeta = 0.5$ in the top row of figure 4, are indeed typically away from unity (implying inference bias), both for the MAP estimator (red circles) and the minimum MSE estimator (3.25) (green crosses). For the debiased estimator (3.21) (blue squares), in contrast, the slope is indeed unity, indicating that bias has been removed successfully. Similarly, the MSE values of the minimum MSE estimator (3.25) (green) are as predicted indeed always below those of the other two estimators.

In order to test prediction (3.9) for the distribution of inferred regression parameters we next generated 10^6 data sets, each with $p = 500$ and $N = 1000$ (so $\zeta = 0.5$), with Gaussian covariates that are pairwise correlated according to (4.1) and $\epsilon = 0.75$. The true association parameters were drawn as i.i.d. Gaussian random variables with amplitude $S = 1$. After carrying our regularized logistic regression with $\eta = 0.05$, we carried out on each of the resulting MAP estimators $\{\hat{\beta}_\mu\}$ of each dataset the specific linear transformation that according to (3.9) should transform these into zero-average and unit variance Gaussian random variables (using the order parameters computed from the theory). Upon creating for each value of μ a histogram of the rescaled estimators $\hat{\beta}'_\mu$, we obtain 500 histograms which according to theory should all collapse asymptotically to a zero average unit variance Gaussian. The result is shown in figure 6. This figure confirms that, even for the modest values of p and N used, the predicted Gaussian statistics of the estimators with the predicted values of average and width given in (3.9) are indeed correct.

Intercept parameter for imbalanced class sizes. Having training data with vastly different outcome class sizes leads to the minority outcome rarely being predicted [90] in logistic regression. As this imbalance increases, especially in the overfitting regime the intercept term β_0 in




Figure 4. Results of logistic MAP regression simulations with $Np = 400\,000$ and $\beta_0^* = 0$, for different combinations of η (regularizer strength), ϵ (covariate correlations), and S (true association strengths). Top row: inferred versus true association parameters for the ratio $\zeta = p/N = 0.5$. Middle and lower row: order parameters w and v plotted versus ζ . Each simulation data point represents average and standard deviation computed over 400 synthetic data sets and regressions. Solid curves give the theoretical predictions obtained by solving the RS equations.

parametrized models diverges [91], and all new samples are assigned the majority outcome. Medical data often exhibit large imbalances between numbers of diseased and healthy samples, with the clinically important decision relying on identifying the rare cases correctly. Similarly, in financial fraud detection there may be millions of legitimate transactions against a handful of fraudulent ones, and we seek to identify the minority class. Existing methods to mitigate the effect of class imbalance have focussed on data pre-processing [92, 93] or incorporating a cost function into the classification algorithm [90]. While these methods are useful to the practitioner, theoretical explanations are limited [91, 94]. Our present theory enables us to investigate class imbalance effects analytically.

The outcome class imbalance in logistic regression data is measured by $m = N^{-1} \sum_{i=1}^N s_i \in [-1, 1]$. Averaging over the data in (4.45) gives the expectation value

$$\langle m \rangle = \int dz p(z) \tanh(\beta_0^* + \beta^* \cdot z / \sqrt{p}). \tag{4.72}$$




Figure 5. Tests of the correction protocols for MAP estimators, applied to the data of the previous figure. Red circles: the MAP estimator $\hat{\beta}_\mu$. Blue squares: the corrected estimator (3.21), aimed at removing inference bias. Green crosses: the corrected estimator (3.25), aimed at minimizing the MSE. The top row show as a function of ζ the slopes of the relation between the three estimators and the true values β_μ^* of association parameters (the slopes of the data clouds in the top row of the previous figure, computed via least squares analysis); this slope would be unity for unbiased estimators. The bottom row shows the values of $\sqrt{\text{MSE}}$, where $\text{MSE} = p^{-1} \sum_{\mu \leq p} (\hat{\beta}_\mu - \beta_\mu^*)^2$.

Figure 7 shows our theoretical predictions for the inferred order parameter β_0 in ML and MAP logistic regression, together with the result of regression simulations. For any given level of outcome imbalance, controlled by β_0^* , the bias in the inferred intercept increases with ζ . Regularization mitigates this effect, leading to the possibility of correcting inferred class imbalances, in spite of the regularization being applied only to the coefficients $\{\beta_\mu\}_{\mu=1}^p$, not to β_0 itself. This situation is reminiscent of Cox’s survival analysis model [69], where the inferred hazard rate (given by Breslow’s estimator [95]) can be expressed in terms of the inferred regression coefficients, and thereby inherits their inference bias. In the $\eta > 0$ case, we find the intercept inflation differs from that of the association parameters, due to their rescaling with \sqrt{p} . Again the agreement between theory and experiment is very satisfactory.

4.3. Regularized Cox regression

Equations for MAP Cox regression. In Cox regression without censoring¹¹ we have $s = t \in [0, \infty)$, $\theta = \{\lambda(t)\}$ (the so-called base hazard rate, a nonnegative function on the time interval $[0, \infty)$ such that $\Lambda(t) = \int_0^t dt' \lambda(t')$ diverges for $t \rightarrow \infty$), and

$$p(t|\xi, \lambda) = \lambda(t)e^{\xi - \exp(\xi)\Lambda(t)}, \tag{4.73}$$

¹¹ Also the more complicated case of Cox regression with censoring falls within the scope of our present GLM equations, provided the censoring events are non-informative (as with end-of-trial censoring). For informative censoring, such as censoring caused by nontrivial competing risks, we first need to generalize the theory further to models in which outcome probabilities involve multiple linear combinations of covariates. This should be straightforward, but will be the subject of a future study.




Figure 6. Histograms of rescaled values of the MAP estimators, each defined as $\hat{\beta}'_\mu = \left[\hat{\beta}_\mu + \tilde{d}_0 [(\tilde{g}\mathbf{I} + 2\eta\mathbf{A}^{-1})^{-1}\beta^*]_\mu \right] / \sqrt{|\tilde{f}|[(\tilde{g}\mathbf{A} + 2\eta\mathbf{I})^{-1}\mathbf{A}(\tilde{g}\mathbf{A} + 2\eta\mathbf{I})^{-1}]_{\mu\mu}}$, computed for 10^6 independent regularized logistic regression experiments with correlated covariates ($\epsilon = 0.75$, $\eta = 0.01$, $S = 1$, $p = 500$ and $N = 1000$). The resulting 500 histograms are plotted together in the present figure, one for each value of μ . According to the theoretical prediction (3.9), these histograms should all asymptotically become zero average and unit average distributions. This is indeed seen to be the case.

$$\log p(t|\xi, \lambda) = \log \lambda(t) + \xi - e^\xi \Lambda(t). \tag{4.74}$$

Hence

$$\frac{\partial}{\partial \xi} \log p(t|\xi, \lambda) = 1 - e^\xi \Lambda(t), \quad \frac{\partial^2}{\partial \xi^2} \log p(t|\xi, \lambda) = -e^\xi \Lambda(t). \tag{4.75}$$

The analysis of overfitting in MAP Cox regression with arbitrary covariate covariance matrices \mathbf{A} was first carried out in [63]. We will now show that from the general equations (2.32)–(2.37) one indeed recovers the results of [63], and more (e.g. the explicit link between true and inferred association parameters).

We will first compute the various model-dependent building blocks of the RS equations. The relevant (functional and partial) derivatives of the logarithm of the outcome probability density are

$$\frac{\delta}{\delta \lambda(t)} \log p(s|\xi, \lambda) = \frac{\delta(s-t)}{\lambda(t)} - e^\xi \theta(s-t), \tag{4.76}$$

$$\frac{\partial}{\partial y_0} \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^*) = S\langle a \rangle^{\frac{1}{2}} \left[1 - e^{S\langle a \rangle^{\frac{1}{2}} y_0} \Lambda^*(s) \right]. \tag{4.77}$$

The function $\xi(\mu, \sigma, s, \beta_0)$ is here the solution of

$$1 - e^\xi \Lambda(s) = (\xi - \mu) / \sigma^2. \tag{4.78}$$

Upon switching from ξ to the variable $x = \mu - \xi + \sigma^2$, we can solve x in explicit form:

$$x = W(\sigma^2 e^{\mu + \sigma^2} \Lambda(s)). \tag{4.79}$$




Figure 7. Predicted values of the offset parameter β_0 are drawn as solid curves, for $\mathbf{A} = \mathbb{I}$ (uncorrelated covariates), $S = 1$, and three regularization strengths ($\eta = 0$: blue circles; $\eta = 0.025$: red squares; $\eta = 0.05$: green crosses). Full circles give the average values of $\hat{\beta}_0$ found in MAP regression, for 400 simulations with Gaussian covariates and $NP = 400\,000$. Standard deviations are not shown in order to reduce visual clutter, but range between 0.003 for small ζ and 0.1 for large ζ . The true offset used in generating the data was $\beta_0^* = 0.25$, representing an average class imbalance of 42 : 58 according to (4.72).

Here $W(x)$ denotes the Lambert W -function. i.e. the inverse of $f(x) = x \exp(x)$, with derivative $W'(x) = W(x)/x[1 + W(x)]$. It then follows that

$$\xi(\mu, \sigma, s, \lambda) = \mu + \sigma^2 - W(\sigma^2 e^{\mu + \sigma^2} \Lambda(s)), \tag{4.80}$$

$$\frac{\partial}{\partial \mu} \xi(\mu, \sigma, s, \lambda) = \frac{1}{1 + W(\sigma^2 e^{\mu + \sigma^2} \Lambda(s))}. \tag{4.81}$$

Order parameter equations. We insert the above formulae into our RS order parameter equations (2.32)–(2.37), and use identities such as $\langle \exp(S \langle a \rangle^{\frac{1}{2}} y_0) \Lambda^*(s) \rangle_s = 1$, to simplify our equations to the following set:

$$\tilde{u}^2 = \left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle, \tag{4.82}$$

$$v^2 = w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle, \tag{4.83}$$

$$-\zeta \tilde{f} \tilde{u}^4 = \left\langle \left\langle \left[\tilde{u}^2 - W(\sigma^2 e^{w y_0 + v z + \tilde{u}^2} \Lambda(s)) \right]^2 \right\rangle_s \right\rangle, \tag{4.84}$$

$$\zeta \tilde{g} \tilde{u}^2 = \left\langle \left\langle \left\langle \frac{W(\sigma^2 e^{w y_0 + v z + \tilde{u}^2} \Lambda(s))}{1 + W(\sigma^2 e^{w y_0 + v z + \tilde{u}^2} \Lambda(s))} \right\rangle_s \right\rangle \right\rangle, \tag{4.85}$$

$$\frac{\zeta w \tilde{u}^2 \langle a \rangle^{\frac{1}{2}}}{S \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle} = - \left\langle \left\langle \left\langle W(\tilde{u}^2 e^{wy_0 + vz + \tilde{u}^2} \Lambda(s)) \left[1 - e^{S\langle a \rangle^{\frac{1}{2}} y_0} \Lambda^*(s) \right] \right\rangle_s \right\rangle \right\rangle, \tag{4.86}$$

$$\frac{\left\langle \left\langle \delta(s-t) \right\rangle_s \right\rangle}{\lambda(t)} = \left\langle \left\langle \left\langle e^{wy_0 + vz + \tilde{u}^2 - W(\tilde{u}^2 e^{wy_0 + vz + \tilde{u}^2} \Lambda(s))} \theta(s-t) \right\rangle_s \right\rangle \right\rangle. \tag{4.87}$$

The first three are immediately recognised from [63]. With the identity $\exp(-W(x)) = W(x)/x$, we find also that (4.85) reduces to the corresponding equation in [63]. This leaves only the identification of equation (4.86). Let us start from the corresponding equation in [63], which reads:

$$\begin{aligned} \zeta w \left[\tilde{g} - \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} \right] &= \frac{1}{\tilde{u}^2} \left\langle \left\langle y_0 \left\langle W(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(s)) \right\rangle_s \right\rangle \right\rangle \\ &= \frac{1}{\tilde{u}^2} \left\langle \left\langle \frac{\partial}{\partial y_0} \left\langle W(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(s)) \right\rangle_s \right\rangle \right\rangle \\ &= \frac{1}{\tilde{u}^2} \left\langle \left\langle \left\langle \frac{\partial W(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(s))}{\partial y_0} \right. \right. \right. \\ &\quad \left. \left. \left. + W(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(s)) \frac{\partial \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda_0^*)}{\partial y_0} \right\rangle_s \right\rangle \right\rangle \\ &= w \zeta \tilde{g} \\ &\quad + \frac{S\langle a \rangle^{\frac{1}{2}}}{\tilde{u}^2} \left\langle \left\langle \left\langle W(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(s)) \left[1 - e^{S\langle a \rangle^{\frac{1}{2}} y_0} \Lambda^*(s) \right] \right\rangle_s \right\rangle \right\rangle. \end{aligned} \tag{4.88}$$

From this we directly recover (4.86), as required. This confirms that from our general theory for GLMs we can indeed recover also for the example choice of regularized Cox regression the complete results of [63]. Moreover, we now have the additional identities (3.10), (3.11), which were not available in that earlier study, and reveal the nontrivial impact of regularization in the more realistic scenario of correlated covariates (which cannot be extracted from overlap order parameters alone).

5. Discussion

In this paper we have extended to arbitrary generalized linear regression models (GLM) the replica analysis of overfitting in MAP and ML inference that was developed initially in [62, 63] for Cox regression [69] with time-to-event data. Parameter inference methods such as MAP and ML were derived and work well for the regime $p \ll N$, where p is the dimensionality and N is the number of samples. But they can produce large inference errors when $p = \mathcal{O}(N)$. This seriously hampers statistical inference in high dimensions, and thereby limits progress in many data-driven scientific disciplines.

In all GLMs, the ML/MAP overfitting-induced parameter inference errors consist of a combination of a reproducible bias and excess noise, both of which disappear when $p/N \rightarrow 0$

but become more prominent as the ratio p/N increases. In the regime $p, N \rightarrow \infty$ with fixed $\zeta = p/N$, the replica method enables us to predict analytically both this inference bias and the distribution of the excess noise, expressed in terms of the true parameters of the model that generated the data, and the distribution of covariates from which the samples were drawn. In contrast to some recent alternative approaches, such as [39, 40, 78], by using the replica method we are not restricted to uncorrelated covariates or to models with output noise only, and we can calculate in explicit form the relation between MAP/ML estimators and the true (but unknown) model parameters responsible for the data. Covariate correlations are in fact found to play an important role in this relation. Our results pave the way for correcting ML and MAP inferences in GLMs systematically for overfitting bias, and thereby extend the applicability of such models into the hitherto forbidden regime $p = \mathcal{O}(N)$.

We found that in our analysis, the choice of outcome data types and regression models can be left until after the derivation of replica symmetric order parameter equations (there is no evidence for replica symmetry breaking, which is reasonable since we have assumed absence of model mismatch). Our derivation relies only on the generalized linear form of GLMs and on choosing L_2 priors. Hence the replica calculation need not be repeated for every new GLM model instance; as always with the replica method, it served as a relatively painless and elegant but powerful vehicle for arriving at a closed set of order parameter equations, together with formulae expressing the relation between the ML/MAP parameter estimators and the true values of these parameters. The latter equations can serve as a natural and convenient starting point for practical applications, even for those with no interest in their derivation. We illustrate our results in this paper by applying the general theory to linear, logistic, and Cox regression, and find excellent agreement with simulations and earlier results. We have limited our experiments to $\zeta \leq 1$. In ML regression this marks the point by which a phase transition will have occurred (even earlier in logistic regression), whereas in MAP regression with $\eta > 0$ once can in fact continue regression and find agreement between theory and simulations into the $\zeta > 1$ regime (data not shown here).

This work can be extended in both practical and theoretical ways, several of which are presently being explored. Our theory was built upon the idealized scenario of knowledge of the underlying data-generating model. To put it into practice, the variance S of the true regression parameters and the population covariance matrix \mathbf{A} need to be estimated. The former is available through the inferred MAP estimators and the existing order parameters. The latter can be estimated from the empirical covariate statistics, building on methods such as [86, 96–98]. For time-to-event models, the next obvious step would be to include censored data. More general extensions of the present theory include working with alternative non-Gaussian priors, inspecting more pathological models or data where some of our mathematical assumptions no longer hold, or generalizing the theory to regression models with multiple linear predictors, such as multinomial regression [99], multiple risks and latent classes in survival analysis [70, 100], and multilayer neural networks [66, 101, 102].

Acknowledgments

The authors gratefully acknowledge valuable and stimulating discussions with Sir David Cox and Dr Heather Battey. MS is supported by the Biotechnology and Biological Sciences Research Council (award 1668568) and GSK Ltd. AM is supported by Cancer Research UK (award C45074/A26553) and the UK's Medical Research Council (award MR/R014043/1). FAL is supported through a scholarship from Conacyt (Mexico).

Appendix A. Derivation of the generic saddle point form

A.1. Preparation

We start with expression (2.16), with the $L2$ prior $p(\beta) \propto \exp(-p\eta\beta^2)$:

$$E_\gamma(\beta^*, \theta^*) = -\lim_{n \rightarrow 0} \frac{\partial}{\partial \gamma} \frac{1}{Nn} \log \int d\theta^1 \dots d\theta^n \int d\beta^1 \dots d\beta^n \prod_{\alpha=1}^n \left[\frac{p(\beta^\alpha)}{p(\beta^*)} \right]^\gamma \\ \times \left\{ \int dz ds p(z) p(s|z, \beta^*, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|z, \beta^\alpha, \theta^\alpha)}{p(s|z, \beta^*, \theta^*)} \right]^\gamma \right\}^N. \quad (\text{A.1})$$

The covariate distribution $p(z)$ is assumed to have zero mean and covariance matrix \mathbf{A} , with entries $A_{\mu\nu} = \int dz p(z) z_\mu z_\nu$. We consider the regime where $N, p \rightarrow \infty$ with fixed ratio $\zeta = p/N$. Following [62, 63] we next introduce

$$p(\mathbf{y}|\beta^0, \dots, \beta^n) = \int dz p(z) \prod_{\alpha=0}^n \delta \left[y^\alpha - \frac{\beta^\alpha \cdot z}{\sqrt{p}} \right], \quad (\text{A.2})$$

where $\mathbf{y} = \{y^0, y^1, \dots, y^n\} \in \mathbb{R}^{n+1}$ (which in survival analysis would be interpreted as risk scores) and $\beta^0 \equiv \beta^*$. Now

$$E_\gamma(\beta^*, \theta^*) = -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int d\theta^1 \dots d\theta^n \int d\beta^1 \dots d\beta^n \prod_{\alpha=1}^n \left[\frac{p(\beta^\alpha)}{p(\beta^0)} \right]^\gamma \\ \times \left\{ \int d\mathbf{y} p(\mathbf{y}|\beta^0, \dots, \beta^n) \int ds p(s|y^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|y^\alpha, \theta^\alpha)}{p(s|y^0, \theta^*)} \right]^\gamma \right\}^N. \quad (\text{A.3})$$

To proceed we assume that $p(\mathbf{y}|\beta^0, \dots, \beta^n)$ is Gaussian, via the central limit theorem. Since $\int dz p(z) z = \mathbf{0}$, the distribution $p(\mathbf{y}|\beta^0, \dots, \beta^n)$ is now given by

$$p(\mathbf{y}|\beta^0, \dots, \beta^n) = \frac{e^{-\frac{1}{2} \mathbf{y} \mathbf{C}^{-1} \{\beta\} \mathbf{y}}}{\sqrt{(2\pi)^{n+1} \det \mathbf{C} \{\beta\}}}. \quad (\text{A.4})$$

It is determined in full by the $(n+1) \times (n+1)$ covariance matrix $\mathbf{C} \{\beta\}$, with entries

$$C_{\alpha\rho} \{\beta\} = \int dz p(z) \left(\frac{\beta^\alpha \cdot z}{\sqrt{p}} \right) \left(\frac{\beta^\rho \cdot z}{\sqrt{p}} \right) = \frac{1}{p} \beta^\alpha \cdot \mathbf{A} \beta^\rho. \quad (\text{A.5})$$

For each replica pair (α, ρ) we use the integral representation of the Dirac delta function, and rescale the conjugate integration parameter by p , substituting

$$1 = \int dC_{\alpha\rho} \delta \left[C_{\alpha\rho} - \frac{1}{p} \beta^\alpha \cdot \mathbf{A} \beta^\rho \right] = \int \frac{dC_{\alpha\rho} d\hat{C}_{\alpha\rho}}{2\pi/p} e^{ip\hat{C}_{\alpha\rho} (C_{\alpha\rho} - \frac{1}{p} \beta^\alpha \cdot \mathbf{A} \beta^\rho)}, \quad (\text{A.6})$$

in order to simplify expression (A.3) to

$$\begin{aligned}
 E_\gamma(\beta^*, \theta^*) &= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \{d\theta^1 \dots d\theta^n\} \int d\mathbf{C} d\hat{\mathbf{C}} \frac{e^{ip \sum_{\alpha,\rho=0}^n \hat{C}_{\alpha\rho} C_{\alpha\rho}}}{(2\pi/p)^{(n+1)^2}} \\
 &\times \left[\int \frac{d\mathbf{y} e^{-\frac{1}{2}\mathbf{y} \mathbf{C}^{-1} \mathbf{y}}}{\sqrt{(2\pi)^{n+1} \det \mathbf{C}}} \int ds p(s|y^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|y^\alpha, \theta^\alpha)}{p(s|y^0, \theta^*)} \right]^\gamma \right]^N \\
 &\times \int d\beta^1 \dots d\beta^n e^{-\eta\gamma \sum_{\alpha=1}^n [(\beta^\alpha)^2 - (\beta^0)^2] - i \sum_{\alpha,\rho=0}^n \hat{C}_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho}. \quad (\text{A.7})
 \end{aligned}$$

A.2. Conversion into a saddle point problem

We next transform $\hat{\mathbf{C}} = -\frac{1}{2}i\mathbf{D}$, define $\tilde{\beta} \equiv \mathbf{A}^{\frac{1}{2}}\beta$ and introduce the $np \times np$ matrix Ξ and the np -dimensional vector ξ , with entries

$$\Xi_{\alpha\mu;\beta\nu} = 2\eta\gamma\delta_{\alpha\beta}(\mathbf{A}^{-1})_{\mu\nu} + \delta_{\mu\nu}D_{\alpha\beta}, \quad \xi_\mu^\alpha = -D_{0\alpha}\tilde{\beta}_\mu^0 \quad (\text{A.8})$$

The Gaussian integral in (A.7) then becomes

$$\begin{aligned}
 &\int \left(\prod_{\alpha=1}^n d\tilde{\beta}^\alpha e^{-\eta\gamma\tilde{\beta}^\alpha \cdot \mathbf{A}^{-1}\tilde{\beta}^\alpha} \right) e^{-\frac{1}{2}\sum_{\alpha,\rho=1}^n D_{\alpha\rho}\tilde{\beta}^\alpha \cdot \tilde{\beta}^\rho - \sum_{\rho=1}^n D_{0\rho}\tilde{\beta}^0 \cdot \tilde{\beta}^\rho} \\
 &= \frac{(2\pi)^{\frac{np}{2}}}{\sqrt{\det \Xi}} e^{\frac{1}{2}\xi \cdot \Xi^{-1} \xi}. \quad (\text{A.9})
 \end{aligned}$$

Let $\{a_\mu\}$ and $\{b_\alpha\}$ denote the eigenvalues of \mathbf{A} and \mathbf{D} . The two terms \mathbf{P} and \mathbf{Q} of Ξ , with components $P_{\alpha\mu,\beta\nu} = 2\eta\gamma\delta_{\alpha\beta}(\mathbf{A}^{-1})_{\mu\nu}$ and $Q_{\alpha\mu,\beta\nu} = \delta_{\mu\nu}D_{\alpha\beta}$, commute. The eigenvectors of Ξ can therefore be written as $\{\hat{\mathbf{u}}^{\mu\alpha}\}$, with components $\hat{u}_{\nu\rho}^{\mu\alpha} = u_\rho^\alpha v_\nu^\mu$, and where $\sum_{\rho \leq n} D_{\lambda\rho} u_\rho^\alpha = b_\alpha u_\lambda^\alpha$ and $\sum_{\nu \leq p} A_{\lambda\nu} v_\nu^\mu = a_\mu v_\lambda^\mu$, and where both are normalised according to $\sum_{\rho \leq n} (u_\rho^\alpha)^2 = \sum_{\nu \leq p} (v_\nu^\mu)^2 = 1$. The eigenvalues of Ξ are then $\xi_{\mu\alpha} = 2\eta\gamma/a_\mu + b_\alpha$, and

$$\det \Xi = \prod_{\mu=1}^p \prod_{\alpha=1}^n \left(\frac{2\eta\gamma}{a_\mu} + b_\alpha \right), \quad (\Xi^{-1})_{\alpha\mu,\alpha'\mu'} = \sum_{\beta=1}^n \sum_{\nu=1}^p \frac{u_\alpha^\beta v_\mu^\nu u_{\alpha'}^\beta v_{\mu'}^\nu}{2\eta\gamma/a_\nu + b_\beta}. \quad (\text{A.10})$$

Hence the integral (A.9) can be written as

$$\frac{(2\pi)^{\frac{np}{2}} e^{\frac{1}{2}\xi \cdot \Xi^{-1} \xi}}{\sqrt{\det \Xi}} = e^{\frac{1}{2}np \log(2\pi) - \frac{1}{2}np \log(2\eta\gamma/a+b) + \frac{1}{2}np \langle (\xi \cdot \hat{\mathbf{u}})^2 (2\eta\gamma/a+b)^{-1} \rangle}, \quad (\text{A.11})$$

where the averages are over the eigenvalues and orthonormal eigenvectors of Ξ , i.e. $\langle f(a, b, \hat{\mathbf{u}}) \rangle = (np)^{-1} \sum_{\mu=1}^p \sum_{\alpha=1}^n f(a_\mu, b_\alpha, \hat{\mathbf{u}}^{\mu\alpha})$. Since $p = \zeta N$ with $\zeta > 0$, the integrals over \mathbf{C} , $\hat{\mathbf{C}}$ and the base hazard rates in (A.7) can for $N \rightarrow \infty$ be evaluated by steepest descent, provided the limits $n \rightarrow 0$ and $N \rightarrow \infty$ commute. Expression (A.11) then enables us to write the result as

$$\lim_{N \rightarrow \infty} E_\gamma(\beta^*, \theta^*) = \frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{n} \text{extr} \Psi(\mathbf{C}, \mathbf{D}, \theta^1 \dots \theta^n), \quad (\text{A.12})$$

with

$$\begin{aligned} \Psi(\mathbf{C}, \mathbf{D}, \theta^1 \dots \theta^n) &= -\frac{1}{2}\zeta \left[\sum_{\alpha, \rho=0}^n D_{\alpha\rho} C_{\alpha\rho} - \frac{1}{p} D_{00} (\tilde{\beta}^0)^2 \right] + \frac{1}{2}(n+1-n\zeta) \log(2\pi) \\ &+ \frac{1}{2} \log \det \mathbf{C} - n\eta\zeta\gamma S^2 + \frac{1}{2}n\zeta \left\langle \log \left(\frac{2\eta\gamma}{a} + b \right) \right\rangle - \frac{1}{2}n\zeta \left\langle \frac{(\boldsymbol{\xi} \cdot \hat{\mathbf{u}})^2}{2\eta\gamma/a+b} \right\rangle \\ &- \log \int d\mathbf{y} e^{-\frac{1}{2}\mathbf{y} \mathbf{C}^{-1} \mathbf{y}} \int d\mathbf{s} p(\mathbf{s} | \mathbf{y}^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(\mathbf{s} | \mathbf{y}^\alpha, \theta^\alpha)}{p(\mathbf{s} | \mathbf{y}^0, \theta^*)} \right]^\gamma. \end{aligned} \quad (\text{A.13})$$

where $S^2 = \lim_{p \rightarrow \infty} p^{-1}(\beta^0)^2$. Differentiating $\Psi(\dots)$ with respect to D_{00} immediately gives $C_{00} = p^{-1}\beta^0 \cdot \mathbf{A}\beta^0 \equiv \tilde{S}^2$.

A.3. Replica symmetric saddle points

Replica symmetric (RS) saddle points are fully invariant under all permutations of the replica labels $\{1, \dots, n\}$. For the present model the RS ansatz takes the form

$$\begin{aligned} \theta^\alpha &= \theta, & C_{0\alpha} &= c_0 & C_{\alpha\rho} &= C\delta_{\alpha\rho} + c(1 - \delta_{\alpha\rho}), \\ D_{0\alpha} &= d_0, & D_{\alpha\rho} &= D\delta_{\alpha\rho} + d(1 - \delta_{\alpha\rho}). \end{aligned} \quad (\text{A.14})$$

Both \mathbf{C} and \mathbf{D} are positive definite, so $C > c$ and $D > d$. We may now write

$$\mathbf{C} = \begin{pmatrix} C_{00} & c_0 & \dots & \dots & c_0 \\ c_0 & C & c & \dots & c \\ \vdots & c & C & \dots & c \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_0 & c & c & \dots & C \end{pmatrix}, \quad \mathbf{C}^{-1} = \begin{pmatrix} B_{00} & b_0 & \dots & \dots & b_0 \\ b_0 & B & b & \dots & b \\ \vdots & b & B & \dots & b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b_0 & b & b & \dots & B \end{pmatrix}.$$

\mathbf{C} has two nondegenerate eigenvalues λ_\pm with $\lambda_+ \lambda_- = [C + (n-1)c]C_{00} - nc_0^2$, and a further $n-1$ fold degenerate eigenvalue $\lambda_0 = C - c$. Hence

$$\log \det \mathbf{C} = \log C_{00} + n \log(C-c) + \frac{n(c-c_0^2/C_{00})}{C-c} + \mathcal{O}(n^2). \quad (\text{A.15})$$

The entries of \mathbf{C}^{-1} are found to be

$$B_{00} = \frac{C + (n-1)c}{C_{00}[C + (n-1)c] - nc_0^2}, \quad b_0 = -\frac{c_0}{C_{00}[C + (n-1)c] - nc_0^2}, \quad (\text{A.16})$$

$$B = b + \frac{1}{C-c}, \quad b = \frac{c_0^2 - cC_{00}}{(C_{00}[C + (n-1)c] - nc_0^2)(C-c)}. \quad (\text{A.17})$$

Hence

$$\mathbf{y} \cdot \mathbf{C}^{-1} \mathbf{y} = B_{00}(y^0)^2 + (B-b) \sum_{\alpha=1}^n (y^\alpha)^2 + b \left(\sum_{\alpha=1}^n y^\alpha \right)^2 + 2b_0 y^0 \sum_{\alpha=1}^n y^\alpha. \quad (\text{A.18})$$

The matrix \mathbf{D} has one eigenvalue $D + (n-1)d$ with eigenvector $\mathbf{v} = (1, \dots, 1)$, and the $n-1$ fold degenerate eigenvalue $D - d$ with eigenspace $(1, \dots, 1)^\perp$. Hence

$$\left\langle \log \left(\frac{2\eta\gamma}{a} + b \right) \right\rangle = \left\langle \log \left(\frac{2\eta\gamma}{a} + D - d \right) \right\rangle + \left\langle \frac{da}{2\eta\gamma + (D-d)a} \right\rangle + \mathcal{O}(n). \quad (\text{A.19})$$

Similarly, using the RS form of $\xi_\mu^\alpha = -d_0(A^{\frac{1}{2}}\beta^0)_\mu$, we may write

$$\left\langle \frac{(\boldsymbol{\xi} \cdot \hat{\mathbf{u}})^2}{2\eta\gamma/a+b} \right\rangle = d_0^2 \left\langle \frac{a^2(\beta^0 \cdot \mathbf{v})^2}{2\eta\gamma+(D-d)a} \right\rangle + O(n). \tag{A.20}$$

Inserting the above RS expressions into (A.13), and using $C_{00} = \tilde{S}^2$, then gives us

$$\begin{aligned} \frac{1}{n}\Psi(\dots) &= -\frac{1}{2}\zeta(2d_0c_0 + DC - dc) + \frac{1}{2}(1-\zeta)\log(2\pi) - \eta\zeta\gamma S^2 + O(n) \\ &+ \frac{1}{2} \left[\log(C-c) + \frac{c-c_0^2/\tilde{S}^2}{C-c} \right] - \frac{1}{2}\zeta d_0^2 \left\langle \frac{a^2(\beta^0 \cdot \mathbf{v})^2}{2\eta\gamma+(D-d)a} \right\rangle \\ &+ \frac{1}{2}\zeta \left\langle \log \left(\frac{2\eta\gamma}{a} + D-d \right) \right\rangle + \frac{1}{2}\zeta \left\langle \frac{da}{2\eta\gamma+(D-d)a} \right\rangle + \frac{1}{2n} \log(\tilde{S}^2 B_{00}) \\ &- \frac{1}{n} \log \int DzDy_0 \int ds p(s|y_0/\sqrt{B_{00}}, \theta^*) \\ &\times \left[\int dy e^{-\frac{1}{2}(B-b)y^2 + y(iz\sqrt{b}-b_0y_0/\sqrt{B_{00}})} \frac{p^\gamma(s|y, \theta)}{p^\gamma(s|y_0/\sqrt{B_{00}}, \theta^*)} \right]^n. \end{aligned} \tag{A.21}$$

We note that

$$B_{00}^{-1} = \tilde{S}^2 - nc_0^2/(C-c) + O(n^2), \quad B - b = 1/(C-c), \tag{A.22}$$

$$b_0 = -c_0/\tilde{S}^2(C-c) + O(n), \quad b = \frac{c_0^2 - c\tilde{S}^2}{\tilde{S}^2(C-c)^2} + O(n). \tag{A.23}$$

This enable us to write the limit $\Psi_{RS}(\dots) = \lim_{n \rightarrow \infty} n^{-1}\Psi(\dots)$ in the simpler form

$$\begin{aligned} \Psi_{RS}(\dots) &= -\frac{1}{2}\zeta \left\{ 2d_0c_0 + DC - dc + \log(2\pi) + 2\eta\gamma S^2 \right. \\ &+ d_0^2 \left\langle \frac{a^2(\beta^0 \cdot \mathbf{v})^2}{2\eta\gamma+(D-d)a} \right\rangle - \left\langle \log \left(\frac{2\eta\gamma}{a} + D-d \right) \right\rangle - \left. \left\langle \frac{da}{2\eta\gamma+(D-d)a} \right\rangle \right\} \\ &- \int DzDy_0 \int ds p(s|\tilde{S}y_0, \theta^*) \log \int Dy \frac{p^\gamma(s|y\sqrt{C-c} + z(c-c_0^2/\tilde{S}^2)^{\frac{1}{2}} + y_0c_0/\tilde{S}, \theta)}{p^\gamma(s|\tilde{S}y_0, \theta^*)}. \end{aligned} \tag{A.24}$$

Here the brackets denote averages over eigenvectors and eigenvalues of the covariate correlation matrix \mathbf{A} : $\langle f(\mathbf{v}, a) \rangle = \lim_{p \rightarrow \infty} p^{-1} \sum_{\mu=1}^p f(\mathbf{v}_\mu, a_\mu)$, with $\mathbf{A}\mathbf{v}_\mu = a_\mu\mathbf{v}_\mu$ for all $\mu = 1 \dots p$.

A.4. Simplification of the theory

We extremize (A.24) over d_0 , which removes an order parameter, and we transform

$$u = \sqrt{C-c}, \quad v = \sqrt{c-(c_0/\tilde{S})^2}, \quad w = c_0/\tilde{S}, \quad f = d, \quad g = D-d, \tag{A.25}$$

with $u, v, w \in \mathbb{D}, \infty$ and with the inverse transformations

$$c_0 = \tilde{S}w, \quad c = v^2 + w^2, \quad C = u^2 + v^2 + w^2. \tag{A.26}$$

These steps result in

$$\lim_{N \rightarrow \infty} E_\gamma(\beta^*, \theta^*) = \frac{\partial}{\partial \gamma} \text{extr}_{u,v,w,f,g,\lambda} \Psi_{\text{RS}}(u, v, w, f, g, \theta), \quad (\text{A.27})$$

in which

$$\begin{aligned} \Psi_{\text{RS}}(\dots) = & -\frac{1}{2}\zeta(g+f)u^2 - \frac{1}{2}\zeta g(v^2+w^2) - \zeta\eta\gamma S^2 \\ & + \frac{1}{2}\zeta \left\{ \tilde{S}^2 w^2 \left\langle \frac{a^2(\beta^0 \cdot v)^2}{2\eta\gamma+ga} \right\rangle^{-1} + \left\langle \log \left(\frac{2\eta\gamma+ga}{a} \right) \right\rangle + f \left\langle \frac{a}{2\eta\gamma+ga} \right\rangle \right\} \\ & - \int \text{DzDy}_0 \int \text{ds} p(s|\tilde{S}y_0, \theta^0) \log \int \text{Dy} \frac{p^\gamma(s|uy+wy_0+vz, \theta)}{p^\gamma(s|\tilde{S}y_0, \theta^0)}. \end{aligned} \quad (\text{A.28})$$

We could also extremize over f , leading to a simple expression with which to remove f and either u or g . The true association parameters β^0 are seen to enter the asymptotic theory only in quadratic functions of β^0 . In appendix B we show that, if the true associations $\{\beta_\mu^0\}$ are drawn randomly and independently from a zero-average distribution, and under mild conditions on the spectrum $\varrho(a)$ of the covariate correlation matrix \mathbf{A} , both terms will be self-averaging with respect to the realization of β^0 . Consequently, with $S^2 = \lim_{p \rightarrow \infty} p^{-1}(\beta^0)^2$ we may then write

$$\tilde{S}^2 = S^2 \langle a \rangle, \quad \left\langle \frac{a^2(\beta^0 \cdot v)^2}{2\eta\gamma+ga} \right\rangle = \left\langle \frac{S^2 a^2}{2\eta\gamma+ga} \right\rangle, \quad (\text{A.29})$$

(where we used the fact that the eigenvectors v of \mathbf{A} were normalized). Hence

$$\begin{aligned} \lim_{N \rightarrow \infty} E_\gamma(\beta^0, \theta^*) = & \int \text{Dy}_0 \int \text{ds} p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \log p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) - \zeta\eta S^2 \\ & + \eta\zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma+ga} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta\gamma+ga)^2} \right\rangle + \left\langle \frac{1}{2\eta\gamma+ga} \right\rangle \right\} \\ & - f \left\langle \frac{a}{(2\eta\gamma+ga)^2} \right\rangle - \int \text{DzDy}_0 \int \text{ds} p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \\ & \times \frac{\int \text{Dy} p^\gamma(s|uy+wy_0+vz, \theta) \log p(s|uy+wy_0+vz, \theta)}{\int \text{Dy} p^\gamma(s|uy+wy_0+vz, \theta)}. \end{aligned} \quad (\text{A.30})$$

The order parameters $\{u, v, w, f, g, \theta\}$ are computed by extremization of the following function, from which we removed any constant terms:

$$\begin{aligned} \Psi_{\text{RS}}(\dots) = & -\frac{1}{2}\zeta(g+f)u^2 - \frac{1}{2}\zeta g(v^2+w^2) \\ & + \frac{1}{2}\zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma+ga} \right\rangle^{-1} + \langle \log(2\eta\gamma+ga) \rangle + f \left\langle \frac{a}{2\eta\gamma+ga} \right\rangle \right\} \\ & - \int \text{DzDy}_0 \int \text{ds} p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \log \int \text{Dy} p^\gamma(s|uy+wy_0+vz, \theta). \end{aligned} \quad (\text{A.31})$$

Appendix B. Self-averaging with respect to true associations

The results of this appendix were derived in [63], but will be briefly recapitulated, for completeness and because they are also needed in deriving (3.10), (3.11). We investigate random variables of the form $\mathcal{R} = p^{-1}\beta^0 \cdot \mathbf{P}\beta^0$, where the true association vectors $\beta^0 = \{\beta_\mu^0\}$ are drawn randomly from some distribution $p(\beta^0)$, and \mathbf{P} is a fixed symmetric positive definite $p \times p$ matrix, which is independent of β^0 . We wish to know the conditions under which \mathcal{R} will be self-averaging, i.e. $\lim_{p \rightarrow \infty} \langle \mathcal{R} \rangle > 0$ exists, and $\lim_{p \rightarrow \infty} [\langle \mathcal{R}^2 \rangle - \langle \mathcal{R} \rangle^2] = 0$ (brackets denote averaging over $p(\beta^0)$). We assume:

- (a) The $\{\beta_\mu^0\}$ are independent and identically distributed, i.e. $p(\beta^0) = \prod_{\mu=1}^p p(\beta_\mu^0)$.
- (b) $p(\beta_\mu^0)$ is symmetric in β_μ^0 , with finite second and fourth order moments.
- (c) $\lim_{p \rightarrow \infty} p^{-1} \sum_{\mu=1}^p P_{\mu\mu} \in \mathbb{R}$.
- (d) $\lim_{p \rightarrow \infty} p^{-2} \sum_{\mu\nu=1}^p P_{\mu\nu}^2 = 0$.

Given that $S^2 = \lim_{p \rightarrow \infty} p^{-1}(\beta^0)^2$, we must identify $\langle (\beta_\mu^0)^2 \rangle = S^2$. It was shown in [63] that the above conditions are sufficient for \mathcal{R} to be self-averaging. This enabled us to infer that the following identities hold (for $g > 0$), as soon as average and width of the eigenvalue distribution $\varrho(a)$ of \mathbf{A} remain finite in the limit $p \rightarrow \infty$:

$$\lim_{p \rightarrow \infty} \frac{1}{p} \beta^0 \cdot \mathbf{A} \beta^0 = S^2 \int da \varrho(a) a, \tag{B.1}$$

$$\lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\rho=1}^p \frac{a_\rho^2 (\beta^0 \cdot \mathbf{v}^\rho)^2}{2\eta\gamma + g a_\rho} = \int da \varrho(a) \frac{S^2 a^2}{2\eta\gamma + g a}. \tag{B.2}$$

Appendix C. Further evaluation of the RS order parameter equations

One can take further steps in evaluating the RS order parameter equations (2.30), (2.31), without specifying any specific GLM model, exploiting the structural features of the theory only. For instance, the order parameter equations for (\tilde{f}, \tilde{g}) are not model dependent, and give:

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \tag{C.1}$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2. \tag{C.2}$$

For uncorrelated and normalized data, where $\varrho(a) = \delta(a - 1)$, this reduces to

$$2\eta + \tilde{g} = \tilde{u}^{-2}, \quad \tilde{f} = -v^2 / \tilde{u}^4. \tag{C.3}$$

Alternatively, for $\eta \rightarrow 0$ (ML regression) equations (C.1) and (C.2) become

$$\tilde{g} = 1 / \tilde{u}^2, \quad \tilde{f} = -v^2 / \tilde{u}^4. \tag{C.4}$$

In addition to the two partial derivatives of Ξ_A with respect to \tilde{f} and \tilde{g} , we also require derivatives with respect to (\tilde{u}, v, w) . These are

$$\frac{\partial \Xi_A}{\partial \tilde{u}} = -\zeta \tilde{f} \tilde{u}, \quad \frac{\partial \Xi_A}{\partial v} = -\zeta \tilde{g} v, \quad \frac{\partial \Xi_A}{\partial w} = \zeta w \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} - \tilde{g} \right]. \tag{C.5}$$

We also need the partial derivatives of (2.29). We note that all partial derivatives of the argument of (2.29) that are channelled indirectly via the variable ξ vanish at the point $\xi = \xi(wy_0 + vz, \tilde{u}, s, \theta)$, by definition. Hence

$$\frac{\partial \Xi_B}{\partial \tilde{u}} = \frac{1}{\tilde{u}^3} \left\langle \left\langle \left[\xi(wy_0 + vz, \tilde{u}, s, \theta) - wy_0 - vz \right]^2 \right\rangle_s \right\rangle, \tag{C.6}$$

$$\frac{\partial \Xi_B}{\partial v} = \frac{1}{\tilde{u}^2} \left\{ \left\langle \left\langle z \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle - v \right\}, \tag{C.7}$$

$$\frac{\partial \Xi_B}{\partial w} = \frac{1}{\tilde{u}^2} \left\{ \left\langle \left\langle y_0 \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle - w \right\}, \tag{C.8}$$

$$\frac{\partial \Xi_B}{\partial \theta} = \left\langle \left\langle \left\langle \frac{\partial \log p(s|\xi, \theta)}{\partial \theta} \Big|_{\xi=\xi(wy_0+vz, \tilde{u}, s, \theta)} \right\rangle_s \right\rangle \right\rangle. \tag{C.9}$$

The remaining four order parameter equations, in addition to the previously derived pair (C.1) and (C.2), then become

$$\left\langle \left\langle \left[\xi(wy_0 + vz, \tilde{u}, s, \theta) - wy_0 - vz \right]^2 \right\rangle_s \right\rangle = -\zeta \tilde{f} \tilde{u}^4, \tag{C.10}$$

$$\left\langle \left\langle z \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle \right\rangle = v (1 - \zeta \tilde{g} \tilde{u}^2), \tag{C.11}$$

$$\left\langle \left\langle y_0 \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle \right\rangle = w + \zeta w \tilde{u}^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} - \tilde{g} \right], \tag{C.12}$$

$$\left\langle \left\langle \left\langle \frac{\partial \log p(s|\xi, \theta)}{\partial \theta} \Big|_{\xi=\xi(wy_0+vz, \tilde{u}, s, \theta)} \right\rangle_s \right\rangle \right\rangle = 0. \tag{C.13}$$

Equations (C.11) and (C.12) can be simplified further upon integrating by parts over z and y_0 . We need to take care that y_0 appears also in the distribution $p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*)$ used to define the measure $\langle \dots \rangle_s$. We first turn to the average in (C.11):

$$\begin{aligned} \left\langle \left\langle z \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle \right\rangle &= \left\langle \left\langle \left\langle \frac{\partial}{\partial z} \xi(wy_0 + vz, \tilde{u}, s, \theta) \right\rangle_s \right\rangle \right\rangle \\ &= v \left\langle \left\langle (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) \right\rangle_s \right\rangle. \end{aligned} \tag{C.14}$$

Next we work on the average in (C.12):

$$\begin{aligned} \left\langle \left\langle y_0 \langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \right\rangle \right\rangle &= \left\langle \left\langle \frac{\partial}{\partial y_0} \int ds p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \xi(wy_0 + vz, \tilde{u}, s, \theta) \right\rangle \right\rangle \\ &= \left\langle \left\langle \int ds \left\{ \xi(wy_0 + vz, \tilde{u}, s, \theta) \frac{\partial}{\partial y_0} p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) \right. \right. \right. \\ &\quad \left. \left. \left. + w p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*) (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) \right\} \right\rangle \right\rangle \\ &= \left\langle \left\langle \left\langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \frac{\partial \log p(s|S\langle a \rangle^{\frac{1}{2}}y_0, \theta^*)}{\partial y_0} \right\rangle_s \right\rangle \right\rangle \\ &\quad + w \left\langle \left\langle (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) \right\rangle_s \right\rangle. \end{aligned} \tag{C.15}$$

With the above results, and upon discarding the trivial solution $v = 0$ and using (C.11) to simplify (C.12), we can rewrite our closed MAP order parameter equation set as:

$$\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle = \tilde{u}^2, \quad (\text{C.16})$$

$$w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle = v^2, \quad (\text{C.17})$$

$$\langle \langle [\xi(wy_0 + vz, \tilde{u}, s, \theta) - wy_0 - vz]^2 \rangle_s \rangle = -\zeta \tilde{f} \tilde{u}^4, \quad (\text{C.18})$$

$$\langle \langle (\partial_1 \xi)(wy_0 + vz, \tilde{u}, s, \theta) \rangle_s \rangle = 1 - \zeta \tilde{g} \tilde{u}^2, \quad (\text{C.19})$$

$$\langle \langle \left\langle \xi(wy_0 + vz, \tilde{u}, s, \theta) \frac{\partial \log p(s|S\langle a \rangle^{\frac{1}{2}} y_0, \theta^*)}{\partial y_0} \right\rangle_s \rangle \rangle = \zeta w \tilde{u}^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1}, \quad (\text{C.20})$$

$$\langle \langle \left\langle \frac{\partial \log p(s|\xi, \theta)}{\partial \theta} \Big|_{\xi=\xi(wy_0 + vz, \tilde{u}, s, \theta)} \right\rangle_s \rangle \rangle = 0. \quad (\text{C.21})$$

Appendix D. Statistics of inferred association parameters

D.1. Asymptotic form

In this appendix we give the details of the evaluation of the distribution (3.7), in the limit $p, N \rightarrow \infty$ with $\zeta = p/N$, first for uncorrelated and then for correlated covariates. We write $\beta^0 = \beta^*$, assume flat priors for the non-association parameters θ , and use the definition (A.2). Due to the limit $n \rightarrow 0$, we may also insert into the above expression without consequence quantities such as $p^{-\gamma n}(\beta^0)$ and $p^{-\gamma n}(s|y^0, \theta^*)$, in order to bring it closer to the integrals found in appendix A. The result is

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \{d\theta^1 \dots d\theta^n\} \int d\beta^1 \dots d\beta^n \delta(\beta - \beta_\mu^1) \prod_{\alpha=1}^n \left[\frac{p(\beta^\alpha)}{p(\beta^0)} \right]^\gamma \\ &\times \left\{ \int dy p(y|\beta^0, \dots, \beta^n) \int ds p(s|y^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|y^\alpha, \theta^\alpha)}{p(s|y^0, \theta^*)} \right]^\gamma \right\}^N. \end{aligned} \quad (\text{D.1})$$

We can now repeat the manipulations of appendix A, with slight modifications. It will in fact be useful to work with the more general family of factorizing priors $p(\beta) \propto \prod_{\mu \leq p} p(\beta_\mu)$, of which the Gaussian one is a special case, but which also allows us to inspect e.g. L1 priors. Our expression for $\mathcal{P}(\beta, \beta^*)$ then becomes

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \{d\theta^1 \dots d\theta^n\} \int dC d\hat{C} \frac{e^{ip \sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} C_{\alpha\rho}}}{(2\pi/p)^{(n+1)^2}} \\ &\times \left[\int \frac{dy e^{-\frac{1}{2} y C^{-1} y}}{\sqrt{(2\pi)^{n+1} \det C}} \int ds p(s|y^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(s|y^\alpha, \theta^\alpha)}{p(s|y^0, \theta^*)} \right]^\gamma \right]^N \\ &\times \int d\beta^1 \dots d\beta^n \delta(\beta - \beta_\mu^1) e^{-i \sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p \left[\frac{p(\beta_\nu^\alpha)}{p(\beta_\nu^0)} \right]^\gamma. \end{aligned} \quad (\text{D.2})$$

We proceed to the limit $p, N \rightarrow \infty$ with fixed $\zeta = p/N$. In view of our previous calculations we define the following quantity:

$$\begin{aligned} \Psi(\mathbf{C}, \hat{\mathbf{C}}, \theta^1, \dots, \theta^n) &= -i\zeta \sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} C_{\alpha\rho} + \frac{1}{2}(n+1) \log(2\pi) + \frac{1}{2} \log \det \mathbf{C} \\ &\quad - \log \int d\mathbf{y} e^{-\frac{1}{2}\mathbf{y}^T \mathbf{C}^{-1} \mathbf{y}} \int d\mathbf{s} p(\mathbf{s}|\mathbf{y}^0, \theta^*) \prod_{\alpha=1}^n \left[\frac{p(\mathbf{s}|\mathbf{y}^\alpha, \theta^\alpha)}{p(\mathbf{s}|\mathbf{y}^0, \theta^*)} \right]^\gamma \\ &\quad - \frac{1}{N} \log \int d\beta^1 \dots d\beta^n e^{-i\sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p \left[\frac{p(\beta_\nu^\alpha)}{p(\beta_\nu^0)} \right]^\gamma. \end{aligned} \quad (\text{D.3})$$

This enables us to write (D.2) as

$$\begin{aligned} \mathcal{P}(\beta, \beta^*) &= \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \int \{d\theta^1 \dots d\theta^n\} \int d\mathbf{C} d\hat{\mathbf{C}} e^{-N\Psi(\mathbf{C}, \hat{\mathbf{C}}, \theta^1, \dots, \theta^n)} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \\ &\quad \times \frac{\int d\beta^1 \dots d\beta^n \delta(\beta - \beta_\mu^1) e^{-i\sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)}{\int d\beta^1 \dots d\beta^n e^{-i\sum_{\alpha, \rho=0}^n \hat{C}_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)}. \end{aligned} \quad (\text{D.4})$$

Integrating both sides over (β, β^*) shows that the first line of (D.4) on its own would equal one. Hence for $N \rightarrow \infty$ we will be left simply with the limit $p \rightarrow \infty$ of the second line, which is an $\mathcal{O}(1)$ object, evaluated at the saddle point of $\Psi(\mathbf{C}, \hat{\mathbf{C}}, \theta^1, \dots, \theta^n)$. For Gaussian priors, the saddle point is the one computed in appendix A. Hence, upon transforming as before $\hat{\mathbf{C}} = -\frac{1}{2}i\mathbf{D}$, and choosing the saddle point values,

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \\ &\quad \times \frac{\int d\beta^1 \dots d\beta^n \delta(\beta - \beta_\mu^1) e^{-\frac{1}{2}\sum_{\alpha, \rho=0}^n D_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)}{\int d\beta^1 \dots d\beta^n e^{-\frac{1}{2}\sum_{\alpha, \rho=0}^n D_{\alpha\rho} \beta^\alpha \cdot A \beta^\rho} \prod_{\alpha=1}^n \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)}. \end{aligned} \quad (\text{D.5})$$

This expression depends on the choice of $p(\mathbf{s}|\xi, \theta)$ only indirectly, via the values of the order parameters $\{D_{\alpha\rho}\}$. We will now work out (D.5) first for uncorrelated and normalized covariates, followed by evaluation for arbitrary covariate correlations.

D.2. Uncorrelated and normalized covariates

This is the simplest case, where $A_{\mu\nu} = \delta_{\mu\nu}$. The integrations in the above formula now factorize over all components of β , giving

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \left(\lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \right) \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \\ &\quad \times \frac{\int d\beta^1 \dots d\beta^n \delta(\beta - \beta^1) e^{-\frac{1}{2}\sum_{\alpha, \rho=1}^n D_{\alpha\rho} \beta^\alpha \beta^\rho - \beta^* \sum_{\alpha=1}^n D_{\alpha 0} \beta^\alpha} \prod_{\alpha=1}^n p^\gamma(\beta^\alpha)}{\int d\beta^1 \dots d\beta^n e^{-\frac{1}{2}\sum_{\alpha, \rho=1}^n D_{\alpha\rho} \beta^\alpha \beta^\rho - \beta^* \sum_{\alpha=1}^n D_{\alpha 0} \beta^\alpha} \prod_{\alpha=1}^n p^\gamma(\beta^\alpha)}. \end{aligned} \quad (\text{D.6})$$

Hence

$$\lim_{N \rightarrow \infty} \mathcal{P}(\beta|\beta^*) = \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{\int d\beta^1 \dots d\beta^n \delta(\beta - \beta^1) e^{-\frac{1}{2} \sum_{\alpha, \rho=1}^n D_{\alpha\rho} \beta^\alpha \beta^\rho - \beta^* \sum_{\alpha=1}^n D_{\alpha 0} \beta^\alpha} \prod_{\alpha=1}^n p^\gamma(\beta^\alpha)}{\int d\beta^1 \dots d\beta^n e^{-\frac{1}{2} \sum_{\alpha, \rho=1}^n D_{\alpha\rho} \beta^\alpha \beta^\rho - \beta^* \sum_{\alpha=1}^n D_{\alpha 0} \beta^\alpha} \prod_{\alpha=1}^n p^\gamma(\beta^\alpha)}. \quad (\text{D.7})$$

We next use the replica symmetric form of the matrix \mathbf{D} , i.e. $D_{\alpha\rho} = D\delta_{\alpha\rho} + d(1 - \delta_{\alpha\rho})$ and $D_{\alpha 0} = d_0$ and for $\alpha, \rho = 1 \dots n$, and carry out a Gaussian linearization:

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{P}(\beta|\beta^*) &= \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{\int \mathbf{Dz} \left[\int d\beta' e^{iz\sqrt{d}\beta' - \frac{1}{2}(D-d)(\beta')^2 - d_0\beta^*\beta'} p^\gamma(\beta') \right]^{n-1} e^{iz\sqrt{d}\beta - \frac{1}{2}(D-d)\beta^2 - d_0\beta^*\beta} p^\gamma(\beta)}{\int \mathbf{Dz} \left[\int d\beta' e^{iz\sqrt{d}\beta' - \frac{1}{2}(D-d)(\beta')^2 - d_0\beta^*\beta'} p^\gamma(\beta') \right]^n} \\ &= \lim_{\gamma \rightarrow \infty} \int \mathbf{Dz} \left[\frac{e^{-\frac{1}{2}(D-d)\beta^2 + \beta(iz\sqrt{d} - d_0\beta^*)} p^\gamma(\beta)}{\int d\beta' e^{-\frac{1}{2}(D-d)(\beta')^2 + \beta'(iz\sqrt{d} - d_0\beta^*)} p^\gamma(\beta')} \right]. \end{aligned} \quad (\text{D.8})$$

In terms of the transformed order parameters $f = d$ and $g = D - d$ this becomes

$$\lim_{N \rightarrow \infty} \mathcal{P}(\beta|\beta^*) = \lim_{\gamma \rightarrow \infty} \int \mathbf{Dz} \left[\frac{e^{-\frac{1}{2}g\beta^2 + \beta(iz\sqrt{f} - d_0\beta^*)} p^\gamma(\beta)}{\int d\beta' e^{-\frac{1}{2}g(\beta')^2 + \beta'(iz\sqrt{f} - d_0\beta^*)} p^\gamma(\beta')} \right]. \quad (\text{D.9})$$

The order parameter d_0 , which we could remove from the general theory, here needs to be computed after all. For the $L2$ (i.e. Gaussian) prior $p(\beta) \propto \exp(-\eta\beta^2)$ we can find d_0 via differentiation of (A.24) and subsequently use (B.2), giving

$$d_0 = -c_0 \left\langle \frac{a^2(\beta^0 \cdot \mathbf{v})^2}{2\eta\gamma + (D-d)a} \right\rangle^{-1} = -\frac{c_0}{S^2} \left\langle \frac{a^2}{2\eta\gamma + ga} \right\rangle^{-1}. \quad (\text{D.10})$$

For uncorrelated and normalized covariates we have $\varrho(a) = \delta(a - 1)$, so $c_0 = Sw$ and

$$d_0 = -\frac{c_0}{S^2} (2\eta\gamma + g) = -\frac{\gamma w}{S} (2\eta + \tilde{g}). \quad (\text{D.11})$$

We thus find with $f = \tilde{f}\gamma^2$,

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{P}(\beta|\beta^*) &= \lim_{\gamma \rightarrow \infty} \int \mathbf{Dz} \left[\frac{e^{\gamma \left[-\frac{1}{2}(2\eta + \tilde{g})\beta^2 + \beta(iz\sqrt{\tilde{f}} + w(2\eta + \tilde{g})\beta^*/S) \right]}}{\int d\beta' e^{\gamma \left[-\frac{1}{2}(2\eta + \tilde{g})(\beta')^2 + \beta'(iz\sqrt{\tilde{f}} + w(2\eta + \tilde{g})\beta^*/S) \right]}} \right] \\ &= \lim_{\gamma \rightarrow \infty} \frac{\sqrt{\gamma(2\eta + \tilde{g})}}{\sqrt{2\pi}} \int \mathbf{Dz} e^{-\frac{1}{2}\gamma(2\eta + \tilde{g}) \left[\beta - iz\sqrt{\tilde{f}}/(2\eta + \tilde{g}) - w\beta^*/S \right]^2}. \end{aligned} \quad (\text{D.12})$$

For $\varrho(a) = \delta(a - 1)$ we also know that $2\eta + \tilde{g} = \tilde{u}^{-2}$ and $\tilde{f} = -v^2/\tilde{u}^4$. Hence the above integral reduces to

$$\lim_{N \rightarrow \infty} \mathcal{P}(\beta|\beta^*) = \frac{1}{v\sqrt{2\pi}} e^{-\frac{1}{2}(\beta - w\beta^*/S)^2/v^2} \quad (\text{D.13})$$

This confirms what was suggested by simulation data and exploited in [62]: if we plot inferred versus true association parameters in a plane, we will find for $L2$ priors and uncorrelated covariates a linear cloud with slope w/S and zero-average Gaussian noise of width v . We have now proved this analytically, for *any* generalized linear model.

D.3. Correlated covariates

This is the more tricky case. We return to (D.5) and implement first the replica symmetry ansatz, i.e. $D_{\alpha 0} = d_0$ and $D_{\alpha\rho} = D\delta_{\alpha\rho} + d(1 - \delta_{\alpha\rho})$ for $\alpha, \rho = 1 \dots n$, so that we can proceed with our calculation:

$$\begin{aligned}
\lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \\
&\times \frac{\int d\tilde{\beta} e^{-\frac{1}{2}d\tilde{\beta} \cdot A \tilde{\beta} - d_0 \tilde{\beta} \cdot A \beta^0} \int \prod_{\alpha=1}^n [d\beta^\alpha e^{-\frac{1}{2}(D-d)\beta^\alpha \cdot A \beta^\alpha} \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)] \delta(\tilde{\beta} - \sum_{\alpha=1}^n \beta^\alpha) \delta(\beta - \beta_\mu^1)}{\int d\tilde{\beta} e^{-\frac{1}{2}d\tilde{\beta} \cdot A \tilde{\beta} - d_0 \tilde{\beta} \cdot A \beta^0} \int \prod_{\alpha=1}^n [d\beta^\alpha e^{-\frac{1}{2}(D-d)\beta^\alpha \cdot A \beta^\alpha} \prod_{\nu=1}^p p^\gamma(\beta_\nu^\alpha)] \delta(\tilde{\beta} - \sum_{\alpha=1}^n \beta^\alpha)} \\
&= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \\
&\times \frac{\int d\hat{\beta} d\tilde{\beta} W(\hat{\beta}, \tilde{\beta}) H^{n-1}(\hat{\beta}) \int d\beta^1 e^{-i\hat{\beta} \cdot \beta^1 - \frac{1}{2}(D-d)\beta^1 \cdot A \beta^1} \delta(\beta - \beta_\mu^1) \prod_{\nu=1}^p p^\gamma(\beta_\nu^1)}{\int d\hat{\beta} d\tilde{\beta} W(\hat{\beta}, \tilde{\beta}) H^n(\hat{\beta})}, \tag{D.14}
\end{aligned}$$

with

$$W(\hat{\beta}, \tilde{\beta}) = e^{i\hat{\beta} \cdot \tilde{\beta} - \frac{1}{2}d\tilde{\beta} \cdot A \tilde{\beta} - d_0 \tilde{\beta} \cdot A \beta^0}, \tag{D.15}$$

$$H(\hat{\beta}) = \int d\beta^1 e^{-i\hat{\beta} \cdot \beta^1 - \frac{1}{2}(D-d)\beta^1 \cdot A \beta^1} \prod_{\nu=1}^p p^\gamma(\beta_\nu^1). \tag{D.16}$$

Note that $\hat{\beta}, \tilde{\beta} \in \mathbb{R}^p$. For $n \rightarrow 0$ the denominator evaluates to $(2\pi)^p$. Hence expression (D.14) can be simplified to

$$\begin{aligned}
\lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \\
&\times \int \frac{d\hat{\beta} d\tilde{\beta}}{(2\pi)^p} W(\hat{\beta}, \tilde{\beta}) \left\{ \frac{\int d\beta^1 e^{-i\hat{\beta} \cdot \beta^1 - \frac{1}{2}(D-d)\beta^1 \cdot A \beta^1} \delta(\beta - \beta_\mu^1) \prod_{\nu=1}^p p^\gamma(\beta_\nu^1)}{\int d\beta^1 e^{-i\hat{\beta} \cdot \beta^1 - \frac{1}{2}(D-d)\beta^1 \cdot A \beta^1} \prod_{\nu=1}^p p^\gamma(\beta_\nu^1)} \right\}. \tag{D.17}
\end{aligned}$$

We choose the Gaussian prior $p(\beta) \propto \exp(-\eta\beta^2)$, we write $\delta(\beta - \beta_\mu^1)$ in integral form, we introduce the unit vector \hat{e}^μ with components $\hat{e}_\nu^\mu = \delta_{\mu\nu}$, we use $f = d$ and $g = D - d$, and we do the Gaussian integrals where possible. This gives

$$\begin{aligned}
\lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \frac{dk}{2\pi} e^{ik\beta} \\
&\times \int \frac{d\hat{\beta} d\tilde{\beta}}{(2\pi)^p} W(\hat{\beta}, \tilde{\beta}) \left\{ \frac{\int d\beta^1 e^{-i(\hat{\beta} + k\hat{e}^\mu) \cdot \beta^1 - \frac{1}{2}\beta^1 \cdot [(D-d)A + 2\gamma\eta\mathbb{I}]\beta^1}}{\int d\beta^1 e^{-i\hat{\beta} \cdot \beta^1 - \frac{1}{2}\beta^1 \cdot [(D-d)A + 2\gamma\eta\mathbb{I}]\beta^1}} \right\} \\
&= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \frac{dk}{2\pi} e^{ik\beta - \frac{1}{2}k^2[(D-d)A + 2\gamma\eta\mathbb{I}]_{\mu\mu}^{-1}} \\
&\times \int \frac{d\hat{\beta}}{(2\pi)^p} e^{-k\hat{e}^\mu \cdot [(D-d)A + 2\gamma\eta\mathbb{I}]^{-1}\hat{\beta}} \int d\tilde{\beta} e^{-\frac{1}{2}d\tilde{\beta} \cdot A \tilde{\beta} - \tilde{\beta} \cdot (d_0 A \beta^0 - i\hat{\beta})}
\end{aligned} \tag{D.18}$$

$$\begin{aligned}
 &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \frac{dk}{2\pi} e^{ik[\beta + d_0 \hat{e}^\mu \cdot [(D-d)\mathbf{A} + 2\gamma\eta\mathbb{I}]^{-1} \mathbf{A} \beta^0]} \\
 &\quad \times e^{-\frac{1}{2}k^2 [(D-d)\mathbf{A} + 2\gamma\eta\mathbb{I}]^{-1} - d[(D-d)\mathbf{A} + 2\gamma\eta\mathbb{I}]^{-1} \mathbf{A} [(D-d)\mathbf{A} + 2\gamma\eta\mathbb{I}]^{-1}}_{\mu\mu} \\
 &= \lim_{p \rightarrow \infty} \lim_{\gamma \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \delta(\beta^* - \beta_\mu^0) \int \frac{dk}{2\pi} e^{ik[\beta + d_0[(g\mathbf{A} + 2\gamma\eta\mathbb{I})^{-1} \mathbf{A} \beta^0]_\mu]} \\
 &\quad \times e^{-\frac{1}{2}k^2 [(g\mathbf{A} + 2\gamma\eta\mathbb{I})^{-1} - f(g\mathbf{A} + 2\gamma\eta\mathbb{I})^{-1} \mathbf{A} (g\mathbf{A} + 2\gamma\eta\mathbb{I})^{-1}]_{\mu\mu}}. \tag{D.19}
 \end{aligned}$$

Next we use the scaling with γ of the order parameters, $f = \tilde{f}\gamma^2$, $g = \tilde{g}\gamma$ and $d_0 = \gamma\tilde{d}_0$. For the integrals to converge we must have $\tilde{f} < 0$ (which follows from solving the order parameter equations). We can then take $\gamma \rightarrow \infty$ and do the integral over k , giving

$$\begin{aligned}
 \lim_{N \rightarrow \infty} \mathcal{P}(\beta, \beta^*) &= \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p \frac{\delta(\beta^* - \beta_\mu^0)}{\sqrt{2\pi|\tilde{f}|[(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A} (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\mu}}} \tag{D.20} \\
 &\quad \times e^{-\frac{1}{2}[\beta + \tilde{d}_0[(\tilde{g}\mathbb{I} + 2\eta\mathbf{A}^{-1})^{-1} \beta^0]_\mu]^2 / |\tilde{f}|[(\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1} \mathbf{A} (\tilde{g}\mathbf{A} + 2\eta\mathbb{I})^{-1}]_{\mu\mu}}.
 \end{aligned}$$

This is expression (3.9) in the main text.

Appendix E. Pathologies of generalization error minimization

Here we illustrate the dangers of using the generalization error as an objective function to be minimized, by using logistic regression as an example. The generalization error $E_g \in [0, 1]$ is the expected fraction of samples for which the true and the inferred model disagree on the outcome values, for samples drawn randomly from the population (as opposed to from the training set). In logistic regression we have $s = \pm 1$ and write $p(s|z, \beta) = \frac{1}{2} + \frac{1}{2}s \tanh(\beta \cdot z)$ (rescaling by \sqrt{p} is not relevant here). If the true and inferred parameters are β^* and β , the generalization error is

$$\begin{aligned}
 E_g &= \int dz p(z) \sum_{s,s'=\pm 1} \frac{1}{2} (1 - ss') p(s|z, \beta^*) p(s'|z, \beta) \\
 &= \frac{1}{2} - \frac{1}{2} \int dz p(z) \tanh(\beta^* \cdot z) \tanh(\beta \cdot z). \tag{E.1}
 \end{aligned}$$

If we were to use E_g to optimize the inferred parameters (assuming it could be estimated without explicit knowledge of the true parameters β^*), we would seek to minimize E_g over β . We note the lower bound

$$E_g \geq \frac{1}{2} - \frac{1}{2} \int dz p(z) |\tanh(\beta^* \cdot z)|. \tag{E.2}$$

While E_g indeed computes the fraction of samples for which the two models disagree on the outcome, it does *not* measure whether the two models also use the same outcome probabilities. To see this, imagine choosing $\beta = \kappa\beta^*$. Here one would find

$$\frac{d}{d\kappa} E_g = - \int dz p(z) \left(\frac{\beta^* \cdot z}{2\sqrt{p}} \right) \tanh\left(\frac{\beta^* \cdot z}{\sqrt{p}}\right) \left[1 - \tanh^2\left(\kappa \frac{\beta^* \cdot z}{\sqrt{p}}\right) \right] < 0. \tag{E.3}$$

Hence the of E_g minimum is found for $\kappa \rightarrow \infty$, where the (diverging) estimator $\beta = \kappa\beta^*$ satisfies the lower bound (E.2). The inferred model $p(s|z, \beta) = \frac{1}{2} + \frac{1}{2}s \operatorname{sgn}(\beta^* \cdot z)$ would indeed get the maximum achievable fraction of binary outcomes predicted correctly, but it would believe erroneously that it has 100% prediction accuracy.

Appendix F. Distribution of empirical covariance matrices

Here we evaluate expression (4.25) further, to convert the integral in (4.23) over all $p \times p$ matrices into an integral over positive definite and symmetric ones. We write symmetric and antisymmetric parts of matrices \mathbf{M} as \mathbf{M}^s and \mathbf{M}^a , and transform integrations over all $\hat{\mathbf{A}}$ into integrations over symmetric and antisymmetric parts. The (anti)symmetrization transformations involved induce identities such as $d\mathbf{M} = 2^{\frac{1}{2}p(p-1)} d\mathbf{M}^s d\mathbf{M}^a$ and $\delta(\mathbf{M}) = 2^{-\frac{1}{2}p(p-1)} \delta(\mathbf{M}^s)\delta(\mathbf{M}^a)$, where $d\mathbf{M}^s = \prod_{\mu \leq \nu} dM_{\mu\nu}^s$ and $d\mathbf{M}^a = \prod_{\mu < \nu} dM_{\mu\nu}^a$. Moreover,

$$\int d\mathbf{M}^a e^{i\operatorname{Tr}(\mathbf{A}^a \mathbf{M}^a)} = \pi^{\frac{1}{2}p(p-1)} \prod_{\mu < \nu} \delta(A_{\mu\nu}^a), \tag{F.1}$$

$$\int d\mathbf{M}^s e^{i\operatorname{Tr}(\mathbf{A}^s \mathbf{M}^s)} = 2^p \pi^{\frac{1}{2}p(p+1)} \prod_{\mu \leq \nu} \delta(A_{\mu\nu}^s). \tag{F.2}$$

We can now compute $P(\hat{\mathbf{A}})$ for the case where $p(\mathbf{z}) = [(2\pi)^{-p} \operatorname{Det} \mathbf{A}]^{\frac{1}{2}} e^{-\frac{1}{2}\mathbf{z} \cdot \mathbf{A} \mathbf{z}}$, giving

$$\begin{aligned} P(\hat{\mathbf{A}}) &= \int \frac{d\mathbf{Q}}{(2\pi)^{p^2}} e^{i\operatorname{Tr}(\mathbf{Q}^s \hat{\mathbf{A}}^s) + i\operatorname{Tr}(\mathbf{Q}^a \hat{\mathbf{A}}^a)} [\operatorname{Det}(\mathbb{I} + \frac{2i}{N} \mathbf{A} \mathbf{Q}^s)]^{-N/2} \\ &= \delta(\hat{\mathbf{A}}^a) \left(\frac{2i}{N}\right)^{-Np/2} (\operatorname{Det} \mathbf{A})^{-N/2} \int \frac{d\mathbf{Q}^s}{(2\pi)^{\frac{1}{2}p(p+1)}} e^{i\operatorname{Tr}(\mathbf{Q}^s \hat{\mathbf{A}}^s)} [\operatorname{Det}(\mathbf{Q}^s - \frac{1}{2} N i \mathbf{A}^{-1})]^{-N/2}. \end{aligned} \tag{F.3}$$

Thus $P(\hat{\mathbf{A}}) d\hat{\mathbf{A}} = 2^{-\frac{1}{2}p(p-1)} [\delta(\hat{\mathbf{A}}^a) d\hat{\mathbf{A}}^a] [P(\hat{\mathbf{A}}^s) d\hat{\mathbf{A}}^s]$, where

$$P(\hat{\mathbf{A}}^s) = 2^{\frac{1}{2}p(p-1)} \left(\frac{2}{N}\right)^{-Np/2} (\operatorname{Det} \mathbf{A})^{-N/2} \int \frac{d\mathbf{Q}^s}{(2\pi)^{\frac{1}{2}p(p+1)}} e^{i\operatorname{Tr}(\mathbf{Q}^s \hat{\mathbf{A}}^s)} [\operatorname{Det}(i\mathbf{Q}^s + \frac{1}{2} N \mathbf{A}^{-1})]^{-N/2}. \tag{F.4}$$

We can now forget about the antisymmetric parts of $\hat{\mathbf{A}}$, and average only over all symmetric matrices. The nontrivial integral in (F.4) is found in [88], giving

$$P(\hat{\mathbf{A}}^s) = \left(\frac{2}{N}\right)^{-Np/2} \frac{e^{-\frac{1}{2}N \operatorname{Tr}(\hat{\mathbf{A}}^s \mathbf{A}^{-1})} (\operatorname{Det} \hat{\mathbf{A}}^s)^{\frac{1}{2}(N-p-1)}}{\pi^{\frac{1}{4}p(p-1)} (\operatorname{Det} \mathbf{A})^{N/2} \prod_{j=\frac{1}{2}(N-p+1)}^{N/2} \Gamma(j)}. \tag{F.5}$$

Hence $P(\hat{\mathbf{A}}^s)$ is a Wishart distribution with N degrees of freedom. With Ω_p denoting the space of positive definite symmetric $p \times p$ matrices, and dropping the superscript s , we may then summarize our result for (4.23) as:

$$P(\hat{\beta}) = \int_{\Omega_p} d\hat{\mathbf{A}} \left(\frac{2}{N}\right)^{-Np/2} \frac{e^{-\frac{1}{2}N \operatorname{Tr}(\hat{\mathbf{A}}^s \mathbf{A}^{-1})} (\operatorname{Det} \hat{\mathbf{A}}^s)^{\frac{1}{2}(N-p-1)}}{\pi^{\frac{1}{4}p(p-1)} (\operatorname{Det} \mathbf{A})^{N/2} \prod_{j=\frac{1}{2}(N-p+1)}^{N/2} \Gamma(j)} \mathcal{N}(\hat{\beta} | \hat{\mathbf{G}} \beta^*, \zeta(\Sigma^*)^2 \hat{\mathbf{G}} \hat{\mathbf{A}}^{-1} \hat{\mathbf{G}}). \tag{F.6}$$

ORCID iDs

A C C Coolen  <https://orcid.org/0000-0002-6976-5875>

M Sheikh  <https://orcid.org/0000-0001-6566-5562>

A Mozeika  <https://orcid.org/0000-0003-1514-1650>

F Aguirre-Lopez  <https://orcid.org/0000-0002-6418-8802>

References

- [1] Bartlett M S 1953 *Biometrika* **40** 12–9
- [2] Bartlett M S 1953 *Biometrika* **40** 306–17
- [3] Haldane J B S and Smith S M 1956 *Biometrika* **43** 96–103
- [4] Anderson J A and Richardson S C 1979 *Technometrics* **21** 71–8
- [5] Cox D R and Snell E J 1968 *J. R. Stat. Soc. B* **30** 248–65
- [6] Shenton L R and Bowman K O 1963 *J. R. Stat. Soc. B* **25** 305–17
- [7] Shenton L R and Bowman K O 1969 *Sankhyā: The Indian Journal of Statistics Ser: B* **31** 379–96 (<https://www.jstor.org/stable/25051692>)
- [8] Cordeiro G M and McCullagh P 1991 *J. R. Stat. Soc. B* **53** 629–43
- [9] Firth D 1993 *Biometrika* **80** 27–38
- [10] Bowman K O and Shenton L R 1965 *Biases and Covariances of Maximum Likelihood Estimators* (Oak Ridge: Union Carbide Corporation)
- [11] Efron B 1982 *The Jackknife, the Bootstrap and Other Resampling Plans* (Philadelphia, PA: SIAM)
- [12] Quenouille M H 1956 *Biometrika* **43** 353–60
- [13] Kosmidis I and Lunardon N 2020 arXiv:2001.03786
- [14] McCullagh P and Nelder J A 1983 *Generalized Linear Models* (London: Chapman and Hall)
- [15] Mohri M, Afshin R and Talwakar A 2018 *Foundations of Machine Learning* (Cambridge, MA: MIT Press)
- [16] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and beyond* (Singapore: World Scientific)
- [17] Mézard M and Montanari A 2016 *Information, Physics, and Computation* (Oxford: University Press)
- [18] Seung H S, Sompolinsky H and Tishby N 1992 *Phys. Rev. A* **45** 6056–91
- [19] Watkin T L, Rau A and Biehl M 1993 *Rev. Mod. Phys.* **65** 499–556
- [20] Oppen M and Kinzel W 1996 *Models of Neural Networks II* (New York: Springer) pp 151–209
- [21] Dietrich R, Oppen M and Sompolinsky H 1999 *Phys. Rev. Lett.* **82** 2975–8
- [22] Gardner E 1988 *J. Phys. A: Math. Gen.* **21** 257–70
- [23] Talagrand M 2003 *Spin Glasses: A Challenge for Mathematicians—Cavity and Mean Field Models* (Berlin: Springer)
- [24] Dia M, Macris N, Krzakala F, Lesieur T and Zdeborová L 2016 *Advances in Neural Information Processing Systems* vol 29 pp 424–32
- [25] Guerra F 2003 *Commun. Math. Phys.* **233** 1–12
- [26] Barbier J and Macris N 2019 *Probab. Theor. Relat. Field* **174** 1133–85
- [27] Thouless D J, Anderson P W and Palmer R G 1977 *Phil. Mag.* **35** 593–601
- [28] Bolthausen E 2014 *Commun. Math. Phys.* **325** 333–66
- [29] Gallager R G 1968 *Information Theory and Reliable Communication* (New York: Wiley)
- [30] Pearl J 1982 *Proc. 2nd AAAI Conf. on Artificial Intelligence* pp 133–6
- [31] Kabashima Y and Saad D 1998 *Europhys. Lett.* **44** 668–74
- [32] Donoho D L, Maleki A and Montanari A 2009 *Proc. Natl Acad. Sci. USA* **106** 18914–9
- [33] Bayati M, Lelarge M and Montanari A 2015 *Ann. Appl. Probab.* **25** 753–822
- [34] Bayati M and Montanari A 2011 *IEEE Trans. Inf. Theory* **57** 764–85
- [35] Deshpande Y and Montanari A 2014 *Proc. 2014 IEEE Int. Symp. Information Theory* pp 2197–201
- [36] Matsushita R and Tanaka T 2013 *Advances in Neural Information Processing Systems* vol 26 pp 917–25
- [37] Lesieur T, Miolane L, Lelarge M, Krzakala F and Zdeborová L 2017 *Proc. 2017 IEEE Int. Symp. Information Theory* pp 511–5

- [38] Rangan S 2011 *Proc. 2011 IEEE Int. Symp. Information Theory* pp 2168–72
- [39] Sur P and Candès E J 2019 *Proc. Natl Acad. Sci. USA* **116** 14516–25
- [40] Salehi F, Abbasi E and Hassibi A 2019 arXiv:1906.03761
- [41] Sur P, Chen Y and Candès E J 2019 *Probab. Theor. Relat. Field* **175** 487–558
- [42] Antenucci F, Krzakala F, Urbani P and Zdeborová 2019 *J. Stat. Mech.* 023401
- [43] Takahashi T and Kabashima Y 2020 arXiv:2001.02824
- [44] Donoho D, Johnstone I, Maleki A and Montanari A 2011 *Proc. 2011 IEEE Int. Symp. Information Theory* pp 129–33
- [45] Krzakala F, Mézard M, Sausset F, Sun Y and Zdeborová L 2012 *Phys. Rev. X* **2** 021005
- [46] Krzakala F, Mézard M, Sausset F, Sun Y and Zdeborová L 2012 *J. Stat. Mech.* 08009
- [47] Kamilov U, Rangan S, Unser M and Fletcher A K 2012 *Advances in Neural Information Processing Systems* vol 25 pp 2438–46
- [48] Villa J and Schniter P 2011 *Conf. Record 45th Asilomar Conf. on Signals, Systems and Computers (IEEE)* pp 799–803
- [49] Villa J and Schniter P 2013 *IEEE Trans. Signal Process.* **61** 4658–72
- [50] Rangan S, Schniter P, Fletcher A K and Sarkar S 2019 *IEEE Trans. Inf. Theory* **65** 5339–51
- [51] Caltagirone F, Zdeborová L and Krzakala F 2014 *2014 IEEE Int. Symp. Information Theory* pp 1812–6
- [52] Via J, Schniter P, Rangan S, Krzakala F and Zdeborová L 2015 *IEEE Int. Conf. on Acoustics, Speech and Signal Processing* pp 2021–5
- [53] Manoel A, Krzakala F, Tramel E and Zdeborová L 2015 *Int. Conf. on Machine Learning* pp 1123–32
- [54] Rangan S, Fletcher A K, Schniter P and Kamilov U S 2016 *IEEE Trans. Inf. Theory* **63** 676–97
- [55] Opper M and Winther O 2001 *Phys. Rev. E* **64** 056131
- [56] Opper M and Winther O 2005 *J. Mach. Learn. Res.* **6** 2177–204
- [57] Kabashima Y and Vehkaperä M 2014 *2014 IEEE Int. Symp. on Information Theory* pp 226–30
- [58] Fletcher A, Sahraee-Ardakan M, Rangan S and Schniter P 2016 *2016 IEEE Int. Symp. Information Theory* pp 190–4
- [59] Cakmak B, Winther O and Fleury B H 2014 *2014 IEEE Inf. Theory Workshop* pp 192–6
- [60] Ma J and Ping L 2017 *IEEE Access* **5** 2020–33
- [61] Rangan S, Schniter P and Fletcher A K 2019 *IEEE Trans. Inf. Theory* **65** 6664–84
- [62] Coolen A C C, Barrett J E, Paga P and Perez-Vicente C J 2017 *J. Phys. A: Math. Theor.* **50** 375001
- [63] Sheikh M and Coolen A C C 2019 *J. Phys. A: Math. Theor.* **52** 384002
- [64] Peng F and Dey D K 1995 *Can. J. Stat.* **23** 199–213
- [65] López Quintero F O, Contreras-Reyes J E, Wiff R and Arellano-Valle R B 2017 *Fish. Bull.* **115** 13–26
- [66] Coolen A C C, Kühn R and Sollich P 2005 *Theory of Neural Information Processing Systems* (Oxford: University Press)
- [67] Klein J P and Moeschberger M L 2003 *Survival Analysis—Techniques for Censored and Truncated Data* (New York: Springer)
- [68] Hougaard P 2000 *Analysis of Multivariate Survival Data* (New York: Springer)
- [69] Cox D R 1972 *J. R. Stat. Soc. B* **34** 187–202
- [70] Cox D R and Oakes D 1984 *Analysis of Survival Data* (Cambridge: Chapman and Hall)
- [71] Rangan S, Goyal V and Fletcher A K 2009 *Advances in Neural Information Processing Systems* pp 1545–53
- [72] Vehkaperä M, Kabashima Y and Chatterjee S 2016 *IEEE Trans. Inf. Theory* **62** 2100–24
- [73] Bereyhi A, Müller R R and Schulz-Baldes H 2019 *IEEE Trans. Inf. Theory* **65** 7896–934
- [74] Gerbelot C, Abbara A and Krzakala F 2020 arXiv:2002.04372
- [75] Kabashima Y, Vehkaperä M and Chatterjee S 2012 *J. Stat. Mech.* 12003
- [76] Barbier J, Dia M, Macris N and Krzakala F 2016 *Proc. 54th Annual Allerton Conf. on Communication, Control and Computing* pp 625–32
- [77] Reeves G and Pfister H D 2016 *2016 IEEE Int. Symp. Information Theory* pp 665–9
- [78] Barbier J, Krzakala F, Macris N, Miolane L and Zdeborová 2019 *Proc. Natl Acad. Sci. USA* **116** 5451–60
- [79] Barbier J, Macris N, Maillard A and Krzakala F 2018 *2018 IEEE Int. Symp. Information Theory* pp 1390–4
- [80] Bishop C M 2006 *Pattern Recognition and Machine Learning* (New York: Springer)
- [81] Casella G and Berger R L 2002 *Statistical Inference* vol 2 (Pacific Grove: Duxbury)

- [82] Tse D N C and Hanly S V 1999 *IEEE Trans. Inf. Theory* **45** 641–57
- [83] Verdu S and Shamai S 1999 *IEEE Trans. Inf. Theory* **45** 622–40
- [84] Guo D and Verdu S 2005 *IEEE Trans. Inf. Theory* **51** 1983–2010
- [85] Rangan S, Fletcher A K and Goyal V K 2012 *IEEE Trans. Inf. Theory* **58** 1902–23
- [86] Higham N J 2002 *IMA J. Numer. Anal.* **22** 329–43
- [87] Wishart J 1928 *Biometrika* **20A** 32–52
- [88] Ingham A E 1933 *Math. Proc. Camb. Phil. Soc.* **29** 271–6
- [89] Nadarajah S and Kotz S 2005 *Acta Appl. Math.* **89** 53–84
- [90] Wallace B C, Small K, Brodley C E and Trikalinos T A 2011 *11th IEEE Int. Conf. on Data Mining* pp 754–63
- [91] Owen A B 2007 *J. Mach. Learn. Res.* **8** 761–73
- [92] Chawla N V 2002 *J. Artif. Intell. Res.* **16** 321–57
- [93] Drummond C and Holte R 2003 *Proc. ICML Workshop on Learning from Imbalanced Datasets II* pp 1–8
- [94] Sei T 2014 *J. Stat. Plan. Inference* **149** 116–24
- [95] Breslow N E 1972 *J. R. Stat. Soc. B* **34** 216–7
- [96] Hayden T L and Wells J 1988 *Linear Algebr. Appl.* **109** 115–30
- [97] Burda Z, Jurkiewicz J and Wactaw B 2005 *Phys. Rev. E* **71** 026111
- [98] El Karoui N 2008 *Ann. Stat.* **36** 2757–90
- [99] Obuchi T and Kabashima Y 2018 *J. Mach. Learn. Res.* **19** 2030–59
- [100] Rowley M *et al* 2017 *Stat. Med.* **36** 2100–19
- [101] MacKay D J C 2003 *Information Theory, Inference, and Learning Algorithms* (Cambridge: University Press)
- [102] Li B and Saad D 2018 *Phys. Rev. Lett.* **120** 248301