
Asymptotics of Bayesian Uncertainty Estimation in Random Features Regression

Youngsoo Baek
Department of Statistical Science
Duke University
Durham, NC 27705
youngsoo.baek@duke.edu

Samuel I. Berchuck
Department of Biostatistics & Bioinformatics
Duke University
Durham, NC 27705
sib2@duke.edu

Sayan Mukherjee*[†]
Center for Scalable Data Analysis and Artificial Intelligence
Universität Leipzig
Leipzig 04105
sayan.mukherjee@mis.mpg.de

Abstract

In this paper, we compare and contrast the behavior of the posterior predictive distribution to the risk of the maximum a posteriori (MAP) estimator for the random features regression model in the overparameterized regime. We will focus on the variance of the posterior predictive distribution (Bayesian model average) and compare its asymptotics to that of the risk of the MAP estimator. In the regime where the model dimensions grow faster than any constant multiple of the number of samples, asymptotic agreement between these two quantities is governed by the phase transition in the signal-to-noise ratio. They also asymptotically agree with each other when the number of samples grows faster than any constant multiple of model dimensions. Numerical simulations illustrate finer distributional properties of the two quantities for finite dimensions. We conjecture they have Gaussian fluctuations and exhibit similar properties as found by previous authors in a Gaussian sequence model, which is of independent theoretical interest.

1 Introduction

One of the most surprising empirical observations in deep learning is the generalization of overparameterized models that can perfectly interpolate the data. The “double descent” curve, referring to the test error first increasing then decreasing with model complexity, has been both empirically and theoretically validated for linear [Hastie et al., 2022] and nonlinear models [Ghorbani et al., 2021, Mei and Montanari, 2022, Hu and Lu, 2022]. Mei and Montanari [2022] showed that the generalization error of the random features (RF) model proposed by Rahimi and Recht [2007] does demonstrate double descent. Perhaps more surprisingly, they also showed that vanishingly small regularization can yield optimal generalization in a nearly noiseless learning task. These findings highlight the recent trend of explaining the success of machine learning through the prism of beneficial interpolation and overparameterization [Belkin, 2021].

There exists, however, another approach to overparameterized learning problems, which is the Bayesian posterior predictive distribution. In Bayesian practice, one can define the training objective

*Departments of Mathematics, Computer Science, Biostatistics & Bioinformatics, Duke University, NC

[†]Max Planck Institute for Mathematics in the Sciences, Leipzig

as the negative log-likelihood of a probabilistic model. The ridge regularized predictor for the RF model studied by Mei and Montanari [2022] is the maximum a posteriori (MAP) estimator of this probabilistic model. A “truly Bayesian” approach, on the other hand, is to derive the posterior predictive distribution which averages over different predictors and summarizes one’s uncertainty in prediction. The posterior predictive distribution is also referred to as the “Bayesian model average” in the literature [Ovadia et al., 2019, Fortuin et al., 2021]. A fundamental question in Bayesian statistics is whether Bayesian credible sets, defined as high probability regions of the posterior, are also valid confidence sets in the frequentist sense [Kleijn and van der Vaart, 2012]. While this is true in finite-dimensional settings and known as Bernstein von-Mises theorem, the behavior of “frequentist” and “Bayesian” confidence sets can be drastically different in high-to-infinite dimensions, as revealed by Cox [1993], Freedman [1999], Johnstone [2010].

In this work, we examine the frequentist properties of the Bayesian random features model. The focus is on whether the variance of the posterior predictive distribution (PPV) demonstrates similar asymptotics as the risk of the MAP estimator studied by Mei and Montanari [2022]. The simplicity of the probabilistic model simplifies the study, as the posterior predictive distribution is available in closed form and is centered around the MAP estimator, for which we know the asymptotic risk. Due to this simplicity, agreement between the risk of the MAP estimator versus the PPV directly implies good coverage property of a natural confidence set formed by the posterior predictive. However, in light of previous work of Johnstone [2010], it is to be expected that such agreement can be hard to reach in high-dimensional learning. Previous literature relevant to our problem is reviewed in Section 2.3.

In this work, we show:

1. Absence of double descent phenomenon in the PPV with vanishing regularization (vague prior). Numerical instability does arise at the “interpolation boundary” for any finite dimensions.
2. In low noise problems, the expected width of the Bayesian credible ball for highly overparameterized models can be much wider than the true risk of the MAP predictor.
3. The expected width of a credible ball asymptotically agrees with the risk when the sample size grows much faster than the model dimensions. Numerical results show the required growth rate of the sample size is unreasonable for many learning problems.
4. Numerical simulations demonstrate finer distributional properties that are currently beyond the reach of theoretical analysis. They suggest the results of Freedman [1999] and Johnstone [2010] in the Gaussian sequence model are applicable to the more complicated RF model.

The rest of the paper is organized as follows. Section 2 reviews the basic setup of the RF model and its probabilistic reformulation. Section 3 summarizes our main results based on asymptotic formulae and numerical simulations. Section 4 concludes with a discussion of the implications and future directions.

2 Background on Random Features Model

Let inputs $\mathbf{x}_i \in \mathbb{R}^d$, $i = 1, 2, \dots, n$ be drawn i.i.d. from a uniform measure (denoted τ) on a d -dimensional sphere with respect to the conventional Euclidean norm:

$$\mathbb{S}^{d-1}(\sqrt{d}) := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| = \sqrt{d}\}. \quad (1)$$

Let outputs y_i be generated by the following model:

$$y_i = f_d(\mathbf{x}_i) + \epsilon_i, \quad f_d(\mathbf{x}) = \beta_{d,0} + \langle \mathbf{x}, \boldsymbol{\beta}_d \rangle + f_d^{NL}(\mathbf{x}). \quad (2)$$

The model is decomposed into a linear component and a nonlinear component, f_d^{NL} . The random error ϵ_i ’s are assumed to be i.i.d. with mean zero, variance τ^2 , and finite fourth moment. We allow $\tau^2 = 0$, in which case the generating model is *noiseless*. In the analysis, these quantities will be assumed to obey further conditions, including an appropriate scaling so that all d -dependent quantities are $O_d(1)$.

2.1 Training with Ridge Regularization

We focus on learning the optimal *random features model* that best fits the training data. This is a class of functions

$$\mathcal{F} := \left\{ f : f(\mathbf{x}) = \sum_{j=1}^N a_j \sigma(\langle \mathbf{x}, \boldsymbol{\theta}_j \rangle / \sqrt{d}) \right\}, \quad (3)$$

which is dependent on N random features, $(\boldsymbol{\theta}_j)_{j=1}^N$, which are drawn i.i.d. from \mathbb{S}^{d-1} , and a nonlinear activation function, σ . The training objective solves a regularized least squares problem for the linear coefficients $\mathbf{a} \equiv (a_j)_{j=1}^N$:

$$\min_{\mathbf{a} \in \mathbb{R}^N} \left\{ \sum_{i=1}^n \left(y_i - \sum_{j=1}^N a_j \sigma(\langle \mathbf{x}_i, \boldsymbol{\theta}_j \rangle / \sqrt{d}) \right)^2 + d\psi_{1,d}\psi_{2,d}\lambda \|\mathbf{a}\|^2 \right\}, \quad (4)$$

where $\psi_{1,d} = N/d$, $\psi_{2,d} = n/d$, and $\lambda > 0$. The optimal weights, $\hat{\mathbf{a}} \equiv \hat{\mathbf{a}}(\lambda)$, determine an optimal ridge predictor denoted by $\hat{f} \equiv f(\cdot; \hat{\mathbf{a}}(\lambda))$. The dependence of the trained predictor on the dataset (\mathbf{X}, \mathbf{y}) and features $\boldsymbol{\Theta}$ are suppressed in notation.

There exist both practical and theoretical motivations for studying RF regression. On the practical side, RF regression has been suggested by Rahimi and Recht [2007] as a randomized approximation scheme for training kernel ridge regression (KRR) models for large datasets. On the theoretical side, (4) describes a “lazy training” algorithm for a 2-layer neural network with activation function σ . Previous studies have focused on the approximation power of such a function class [Jacot et al., 2018, Chizat et al., 2019] and the optimization landscape involved in training problems of type (4) [Mei et al., 2018, 2019].

Given a new input feature \mathbf{x} , the ridge regularized predictor for the unknown output y has the form

$$\hat{f}(\mathbf{x}) \equiv f(\mathbf{x}; \hat{\mathbf{a}}) = \sigma(\mathbf{x}^T \boldsymbol{\Theta} / \sqrt{d}) \hat{\mathbf{a}}, \quad (5)$$

with the optimal ridge weights $\hat{\mathbf{a}}$ and a resolvent matrix that defines the joint posterior of the weights:

$$\hat{\mathbf{a}} \equiv \hat{\mathbf{a}}(\lambda) := \hat{\boldsymbol{\Sigma}}(\lambda) \mathbf{Z}^T \mathbf{y} / \sqrt{d}, \quad (6)$$

$$\hat{\boldsymbol{\Sigma}}(\lambda) := \left(\mathbf{Z}^T \mathbf{Z} + \psi_{1,d} \psi_{2,d} \lambda \mathbf{I}_N \right)^{-1}. \quad (7)$$

Here, $\mathbf{Z} := \sigma(\mathbf{X} \boldsymbol{\Theta} / \sqrt{d}) / \sqrt{d}$ for input design matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and output vector $\mathbf{y} \equiv (y_i)_{i=1}^n$. Similarly we write $\sigma(\mathbf{x}) := \sigma(\mathbf{x}^T \boldsymbol{\Theta} / \sqrt{d})$. The $L^2 \equiv L^2(\mathbb{S}^{d-1}(\sqrt{d}); \tau)$ generalization error of predictor \hat{f} is defined by

$$R_{RF}(\mathbf{y}, \mathbf{X}, \boldsymbol{\Theta}, \lambda) := \mathbb{E}_{\mathbf{x}} \|f_d(\mathbf{x}) - \hat{f}(\mathbf{x})\|^2 \equiv \|f_d - \hat{f}\|_{L^2}^2, \quad (8)$$

We emphasize that (8) is a random quantity, as it depends on $(\mathbf{y}, \mathbf{X}, \boldsymbol{\Theta})$.

2.2 RF as Bayesian Model

The objective function (4) can be interpreted as a MAP estimation problem for an equivalent Bayesian model. Formally, we adopt a d -dependent weight prior distribution, denoted $p(\mathbf{a})$, and also a normal likelihood, denoted $p(\mathbf{y} | \mathbf{X}, \boldsymbol{\Theta}, \mathbf{a})$, centered around a function in the class (3) with variance ϕ^{-1} .

$$\begin{aligned} \mathbf{a} &\sim \text{Normal} \left(0, \phi^{-1} \frac{\psi_{1,d} \psi_{2,d} \lambda}{d} \mathbf{I}_N \right) \\ \mathbf{y} | \mathbf{X}, \boldsymbol{\Theta}, \mathbf{a} &\sim \text{Normal} \left(\sigma(\mathbf{X} \boldsymbol{\Theta} / \sqrt{d}) \mathbf{a}, \phi^{-1} \mathbf{I}_n \right) \end{aligned}$$

The normal likelihood model need not agree with the generating process (2), as is often the case for Bayesian deep learning. Technically and inferentially, it can be justified as a coherent update of belief given a squared error loss [Bissiri et al., 2016]. The joint likelihood of (\mathbf{y}, \mathbf{a}) is defined conditional on both the random features $\boldsymbol{\Theta}$ and an “inverse temperature” ϕ . The choice to condition on $\boldsymbol{\Theta}$ instead of learning them can be unnatural in certain settings and is mainly for the convenience of analysis.

However, we do note that the RF model has been used by Crawford et al. [2018] as an approximation of fully Bayesian kernel learning.

The posterior predictive distribution, or Bayesian model average over the posterior of \mathbf{a} , is defined as

$$p(y \mid \mathbf{x}, \mathbf{y}, \mathbf{X}, \Theta) := \int_{\mathbb{R}^N} p(y \mid \mathbf{x}, \Theta, \mathbf{a}) p(\mathbf{a} \mid \mathbf{y}, \mathbf{X}, \Theta) d\mathbf{a}, \quad (9)$$

where the posterior of \mathbf{a} is a probability distribution placing higher mass near settings of \mathbf{a} minimizing ridge objective (4):

$$p(\mathbf{a} \mid \mathbf{y}, \mathbf{X}, \Theta) \propto \exp \left[-\frac{\phi}{2} \left\{ \sum_{i=1}^n (y_i - f(\mathbf{x}; \mathbf{a}))^2 + \frac{\psi_{1,d}\psi_{2,d}\lambda}{d} \|\mathbf{a}\|^2 \right\} \right]. \quad (10)$$

This is a Gaussian measure centered around $\hat{\mathbf{a}}$ (6) and covariance matrix $\hat{\Sigma}(\lambda)/d$. Thus the posterior predictive at new input \mathbf{x} is also a normal distribution centered around \hat{f} , with variance

$$s^2(\mathbf{x}) \equiv s^2(\mathbf{x}; \lambda) := \phi^{-1} (1 + \sigma(\mathbf{x})^T \hat{\Sigma}(\lambda) \sigma(\mathbf{x}) / d). \quad (11)$$

We refer to (11) as the PPV. This quantity dictates the width of the uncertainty interval centered around (4) evaluated at \mathbf{x} . The expected PPV over \mathbf{x} is defined by

$$S_{RF}^2(\mathbf{y}, \mathbf{X}, \Theta, \lambda) := \mathbb{E}_{\mathbf{x}} [s^2(\mathbf{x})] \equiv \mathbb{E}_{\mathbf{x}} [\mathbb{V}[y \mid \mathbf{x}, \mathbf{y}, \mathbf{X}, \Theta]]. \quad (12)$$

The expectation over \mathbf{x} , similarly as in (8), yields the radius of the posterior credible ball for a posterior Gaussian process f centered around the optimal predictor \hat{f} . Furthermore, the Gaussian likelihood model simplifies the expected PPV into a decomposition of this radius and the inverse temperature of the posterior:

$$S_{RF}^2 = \int \|f - \hat{f}\|_{L^2}^2 dp(f \mid \mathbf{y}, \mathbf{X}, \Theta) + \phi^{-1}, \quad (13)$$

where $p(f \mid \mathbf{y}, \mathbf{X}, \Theta)$ is the law of Gaussian process f induced by the weight posterior $p(\mathbf{a} \mid \mathbf{y}, \mathbf{X}, \Theta)$. Thus, both summands in the display depend on the training features and are random. It is worth contrasting this definition with (12); in the next section, we explain in detail the motivation for comparing the two quantities.

The extra quantity introduced in the Bayesian formulation, ϕ , governs how much the resulting posterior should concentrate around the ridge predictor (6). Since we are interested in the regime where $N, d \rightarrow \infty$, it is reasonable to assume the scale of the likelihood, ϕ , must appropriately decrease with d , similar to how the prior on \mathbf{a} is rescaled by $1/d$. Practitioners may adopt some prior distribution on this parameter and perform hierarchical inference. We adopt a simpler, empirical Bayes approach and choose it to maximize the marginal likelihood:

$$\hat{\phi}^{-1} \equiv \hat{\phi}^{-1}(\lambda) := \frac{\langle \mathbf{y}, \mathbf{y} - \hat{f}(\mathbf{X}) \rangle}{n}. \quad (14)$$

This choice coincides with the training set error attained by predictor (4), so it will be decreasing as $N \rightarrow \infty$. If $N > n$ and $\lambda \rightarrow 0^+$, the training error vanishes as the model can perfectly interpolate the training set. We note the precise asymptotics of the training error has been already characterized by Mei and Montanari [2022] (Section 6).

A fundamental question here is whether R_{RF} and S_{RF}^2 have similar asymptotics, as both quantities summarize uncertainty about our prediction in different ways. R_{RF} is the ‘‘frequentist’s true risk,’’ which requires assumptions about the unknown generative process. S_{RF}^2 , on the other hand, is the ‘‘Bayesian’s risk’’ that can be actually computed without any model assumptions. Its asymptotics *does* depend on model assumptions, as both the prior and likelihood need not capture the generative process (2). In particular, it is desired that it agrees with R_{RF} in some limiting sense. Throughout the rest of this work, we probe the question: Do R_{RF} and S_{RF}^2 converge to the same value as $d, n, N \rightarrow \infty$?

2.3 Previous Works

We have introduced our problem of comparing two different quantities, R_{RF} and S_{RF}^2 , and provided their respective interpretations. Such comparison, between the frequentist risk and the variance of

the posterior, was done by Freedman [1999] in a white noise model, where one observes an infinite square-summable sequence with Gaussian noise. The key finding is that the distributions of two corresponding statistics, re-normalized, have different variances. They are in fact radically different distributions in that they are essentially orthogonal. Knapik et al. [2011] clarified the situation by showing frequentist coverage of credible ball depends heavily on the smoothness of the prior covariance operator. Johnstone [2010] have extended the results to a sequence of finite but growing length, which is a setup more similar to ours. Our study for the RF model addresses a much simpler question, as the theoretical results in Section 3 only address whether two statistics converge to the same limits. Unlike in the work by previous authors, the identity of limits cannot be taken for granted. A key feature driving the different asymptotics of the two quantities is that the Bayesian’s prior and the likelihood are “mismatched” with respect to the data generating process; i.e., f_d need not belong to the RF class (3). In Section 3.3, we demonstrate some distributional properties of R_{RF} and S_{RF}^2 empirically observed in numerical simulations.

Uncertainty quantification in Bayesian learning has recently garnered much attention in the theoretical deep learning literature. We highlight, among many works: Clarté et al. [2023a], characterizing exact asymptotic calibration of a Bayes optimal classifier and a classifier obtained from the generalized approximate message passing algorithm, and Clarté et al. [2023b], demonstrating double descent-like behavior of the Bayes calibration curve in RF models. In particular, quantities R_{RF} and S_{RF}^2 studied by our work can be related, respectively, to the “Bayes-optimal classifier” and “empirical risk classifier” of the latter work (Section 2). The recent work of Guionnet et al. [2023], on the other hand, explicitly addresses the model mismatch in Bayesian inference by deriving the exact, universal asymptotic formula for generalization of a misspecified Bayes estimator in a rank-1 signal detection problem.

3 Results

3.1 Asymptotic Characterization

We present the main results and illustrate them through numerical simulations. We operate under assumptions identical to those of Mei and Montanari [2022], stated below. The proofs are collected in the Supplementary Materials.

Assumption 1. Define $\psi_{1,d} = N/d$ and $\psi_{2,d} = n/d$. We assume $\psi_{1,d} \rightarrow \psi_1 < +\infty$ and $\psi_{2,d} \rightarrow \psi_2 < +\infty$ as $d \rightarrow \infty$.

Assumption 2. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be weakly differentiable and $|\sigma(x)|, |\sigma'(x)| < c_0 e^{c_1|x|}$ for some $c_0, c_1 < +\infty$. For $G \sim \text{Normal}(0, 1)$, the coefficients:

$$\mu_0 = \mathbb{E}[\sigma(G)], \mu_1 = \mathbb{E}[G\sigma(G)], \mu_*^2 = E[\sigma^2(G)] - (\mu_0^2 + \mu_1^2) \quad (15)$$

are assumed to be greater than 0 (ruling out a linear σ) and finite.

Assumption 3. The generating model (2) satisfies

$$\beta_{0,d}^2 \rightarrow F_0^2 < +\infty, \|\beta_d\|^2 \rightarrow F_1^2 < +\infty; \quad (16)$$

furthermore, $f_{NL,d}(\cdot)$ is a centered Gaussian process on $\mathbb{S}^{d-1}(\sqrt{d})$, whose covariance function has the form:

$$\mathbb{E}[f_{NL,d}(\mathbf{x}_1)f_{NL,d}(\mathbf{x}_2)] = \Sigma_d(\langle \mathbf{x}_1, \mathbf{x}_2 \rangle / d). \quad (17)$$

The kernel Σ_d satisfies

$$\mathbb{E}[\Sigma_d(x_{(1)}/\sqrt{d})] = 0, \mathbb{E}[x_{(1)}\Sigma_d(x_{(1)}/\sqrt{d})] = 0 \quad \text{and} \quad \Sigma_d(1) \rightarrow F_*^2 < +\infty, \quad (18)$$

where $x_{(1)}$ is the first entry of $\mathbf{x} \sim \text{Unif}(\mathbb{S}^{d-1}(\sqrt{d}))$. The signal-to-noise ratio (SNR) of the model is defined by

$$\rho = \frac{F_1^2}{F_*^2 + \tau^2}. \quad (19)$$

In this “linear proportional” asymptotic regime, we derive an asymptotic formula for the expected PPV akin to that of Mei and Montanari [2022] for the generalization error. Analysis shows Definition 1 of Mei and Montanari [2022], characterizing the asymptotics of R_{RF} , can be straightforwardly applied to also characterize the asymptotics of S_{RF}^2 .

Proposition 1. Denote by \mathbb{C}_+ the upper half complex plane: $\{a + bi : i = \sqrt{-1}, b > 0\}$. Let functions $\nu_1, \nu_2 : \mathbb{C}_+ \rightarrow \mathbb{C}_+$ be uniquely defined by the conditions: on \mathbb{C}_+ , $\nu_1(\xi), \nu_2(\xi)$ are analytic and uniquely satisfy the equations

$$\nu_1 = \psi_1 \left(-\xi - \nu_2 - \frac{\zeta^2 \nu_2}{1 - \zeta^2 \nu_1 \nu_2} \right)^{-1}, \quad (20)$$

$$\nu_2 = \psi_2 \left(-\xi - \nu_1 - \frac{\zeta^2 \nu_1}{1 - \zeta^2 \nu_1 \nu_2} \right)^{-1}, \quad (21)$$

when $|\nu_1(\xi)| \leq \psi_1/\text{Im}(\xi)$ and $|\nu_2(\xi)| \leq \psi_2/\text{Im}(\xi)$, with $\text{Im}(\xi) > C$ for sufficiently large constant C . Define

$$\chi = \nu_1(i\sqrt{\psi_1\psi_2\lambda/\mu_*})\nu_2(i\sqrt{\psi_1\psi_2\lambda/\mu_*}), \quad \zeta = \frac{\mu_1}{\mu_*}. \quad (22)$$

Under Assumptions 1-3 and for $\lambda > 0$,

$$S_{RF}^2(\mathbf{y}, \mathbf{X}, \Theta, \lambda) \xrightarrow{P} \mathcal{S}^2 = \frac{F_1^2}{1 - \chi\zeta^2} + F_*^2 + \tau^2 \text{ as } d, n, N \rightarrow \infty.$$

Note that \mathcal{S}^2 depends on $(\psi_1, \psi_2, \lambda)$ only through function χ . The following facts follow from this fact and the asymptotic characterization of function χ when $\lambda \rightarrow 0^+$.

Proposition 2. The following holds for the map $(\psi_1, \psi_2, \lambda) \mapsto \mathcal{S}^2$ defined in Proposition 1:

1. It is non-decreasing in λ .
2. $\lim_{\lambda \rightarrow 0^+} \mathcal{S}^2 < +\infty$ when $\psi_1 = \psi_2$.

Item 2, Proposition 2 deserves special attention, as it seems to suggest that there is no ‘‘double descent’’ in the asymptotics of \mathcal{S}^2 when using an improper prior. This fact does *not* possess an operational meaning at any finite $N = n$, due to the ordering of the limits $N, n \rightarrow \infty$ and $\lambda \rightarrow 0^+$. In fact, one may expect that the distribution of the least singular value of the relevant matrix \mathbf{Z} causes numerical instability for small λ when $N = n$. In Section 3.3, this hypothesis is empirically validated through simulations. Subtly, the theoretical prediction does accurately describe the numerical simulations for small choices of λ . On the other hand, the asymptotic prediction for the frequentist risk does diverge when $\lambda \rightarrow 0^+$ and $\psi_1 = \psi_2$.

3.2 Comparison with Generalization Error

The asymptotics of (8), the ‘‘frequentist’’ risk for our comparison, was characterized by Mei and Montanari [2022] through a theorem similar to Proposition 1. Our main interest lies in comparing the risk against the Bayesian variance in two limiting cases:

1. Highly overparameterized models where $\psi_1 \rightarrow \infty$. The number of parameters grows faster than any constant multiple of the number of samples.
2. Large sample models where $\psi_2 \rightarrow \infty$. The number of samples grows faster than any constant multiple of the number of parameters.

The first regime has been studied as the relevant regime in modern deep learning. The second regime is closer to the classical regime of asymptotic statistics where only the number of samples n diverges. Below, we re-state the asymptotic formulae for R_{RF} of Mei and Montanari [2022] in these special cases, which admits simplifications relative to when both ψ_1, ψ_2 are finite.

Proposition 3. (Theorem 4-5, Mei and Montanari [2022]) Under the notation of Proposition 1, define a function

$$\omega \equiv \omega(\zeta, \psi, \bar{\lambda}) = -\frac{(\psi\zeta^2 - \zeta^2 - 1) + \sqrt{\psi\zeta^2 - \zeta^2 - 1}}{2(\bar{\lambda}\psi + 1)}.$$

In the highly overparameterized regime, where $\psi_1 \rightarrow \infty$, asymptotic risk is given as

$$\mathcal{R}_{wide}(\rho, \zeta, \psi_2, \bar{\lambda}) = \lim_{\psi_1 \rightarrow \infty} \lim_{d \rightarrow \infty} \mathbb{E}R_{RF}(\mathbf{y}, \mathbf{X}, \Theta, \lambda) = \frac{(F_1^2 + F_*^2 + \tau^2)(\psi_2\rho + \omega_2^2)}{(1 + \rho)(\psi_2 - 2\omega_2\psi_2 + \omega_2^2\psi_2 - \omega_2^2)} + F_*^2$$

with $\omega_2 = \omega(\zeta, \psi_2, \lambda/\mu_*^2)$.

In the large sample regime, where $\psi_2 \rightarrow \infty$, asymptotic risk is given as

$$\mathcal{R}_{lsamp}(\zeta, \psi_1, \bar{\lambda}) = \lim_{\psi_2 \rightarrow \infty} \lim_{d \rightarrow \infty} \mathbb{E}R_{RF}(\mathbf{y}, \mathbf{X}, \Theta, \lambda) = \frac{F_1^2(\psi_1 \zeta^2 + \omega_1^2)}{\zeta^2(\psi_1 - 2\omega_1 \psi_1 + \omega_1^2 \psi_1 - \omega_1^2)} + F_*^2$$

with $\omega_1 = \omega(\zeta, \psi_1, \lambda/\mu_*^2)$.

A simple question, mentioned in Section 2.2, was whether the two quantities agree. It turns out that in the second regime, at least, the two formulae converge to the same limit, which is the main content of the next Proposition. In the first regime, whether the limits agree is determined by the signal-to-noise ratio ρ . Mei and Montanari [2022] showed that if ρ is larger than a certain critical threshold ρ_* , which depends on (ψ_2, ζ) , the optimal regularization is achieved by $\lambda \rightarrow 0^+$, whereas if ρ is smaller than ρ_* , there exists an optimal regularization λ_* bounded away from 0. This phase transition also determines the agreement of the risk and the PPV in the limit.

Proposition 4. (Proposition 5.2, Mei and Montanari [2022] $+\alpha$) *Define quantities*

$$\rho_*(\zeta, \psi_2) = \frac{\omega_{0,2}^2 - \omega_{0,2}}{(1 - \psi_2)\omega_{0,2} + \psi_2},$$

$$\omega_{0,2} = \omega(\zeta, \psi_2, 0)$$

under the notation of Proposition 3.

1. If $\rho < \rho_*$, then $\min_{\bar{\lambda} \geq 0} \mathcal{R}_{wide}(\rho, \zeta, \psi_2, \bar{\lambda})$, is attained at some $\lambda^{opt} > 0$:

$$\lambda^{opt} := \arg \min_{\bar{\lambda} \geq 0} \mathcal{R}_{wide}(\rho, \zeta, \psi_2, \bar{\lambda}) = \frac{\zeta^2 \psi_2 - \zeta^2 \omega_* \psi_2 + \zeta^2 \omega_* + \omega_* - \omega_*^2}{(\omega_*^2 - \omega_*) \psi_2},$$

$$\omega_* := \omega(\sqrt{\rho}, \psi_2, 0).$$

Furthermore, $\mathcal{R}_{wide}(\rho, \zeta, \psi_2, \lambda^{opt}) = \lim_{\psi_1 \rightarrow \infty} \mathcal{S}^2(\psi_1, \psi_2, \lambda^{opt}) - \tau^2$.

If, on the other hand, $\rho > \rho_*$, then $\arg \min_{\bar{\lambda} \geq 0} \mathcal{R}_{wide}(\rho, \zeta, \psi_2, \bar{\lambda}) = 0$. Furthermore, $\mathcal{R}_{wide}(\rho, \zeta, \psi_2, 0) < \lim_{\psi_1 \rightarrow \infty} \mathcal{S}^2(\psi_1, \psi_2, 0) - \tau^2$.

2. $\arg \min_{\bar{\lambda} \geq 0} \mathcal{R}_{lsamp}(\zeta, \psi_1, \bar{\lambda}) = 0$ and $\mathcal{R}_{lsamp}(\zeta, \psi_1, 0) = \lim_{\psi_2 \rightarrow \infty} \mathcal{S}^2(\psi_1, \psi_2, 0) - \tau^2$.

Note the subtraction of the noise level τ^2 from \mathcal{S}^2 . This is because S_{RF}^2 is computed based on the training error in Proposition 1, which includes both the approximation error and the variance of data, whereas R_{RF} is computed based only on the approximation error in the test set.

3.3 Numerical Simulations

In this section, we first compare evaluations of the asymptotic formulae for R_{RF} and S_{RF}^2 for varying $(\psi_1, \psi_2, \lambda)$. We highlight their difference for finite (ψ_1, ψ_2) at the optimal choice of λ for the MAP risk. We also present numerical simulations, which both validate the formulae (they concentrate fast) and empirically exhibit further interesting distributional properties (suggesting the need for second-order asymptotics). MATLAB codes used to produce simulation results are included in the Supplementary Materials.

Figure 1 shows the dramatic mismatch between R_{RF} and S_{RF}^2 in the low-noise regime. It turns out the conservativeness of the credible ball persists for the choice of λ that depends on (ψ_1, ψ_2) . The situation becomes more delicate in the high-noise regime because there exists a phase transition in the optimal choice of λ that depends on (19), which decreases with the noise variance τ^2 . Figures 2.2a and 2b compare the two curves, \mathcal{R} and \mathcal{S} , for the ‘‘optimally tuned’’ λ at which the best possible frequentist risk is attained, for a fixed pair of (ψ_1, ψ_2) . In a low-noise task, with $\rho = 5$, the ratio of the frequentist risk of the posterior mean predictor to the width of the posterior predictive credible ball is less than 1 for a wide range of ψ_1 . The situation is more nuanced and possibly more favorable for the Bayesian case in the high-noise task with $\rho = 1/5$.

While Figure 1 validates good concentration properties of R_{RF} and S_{RF}^2 with respect to their asymptotic formulae, we may want to investigate the rate of fluctuations for these quantities. Figure

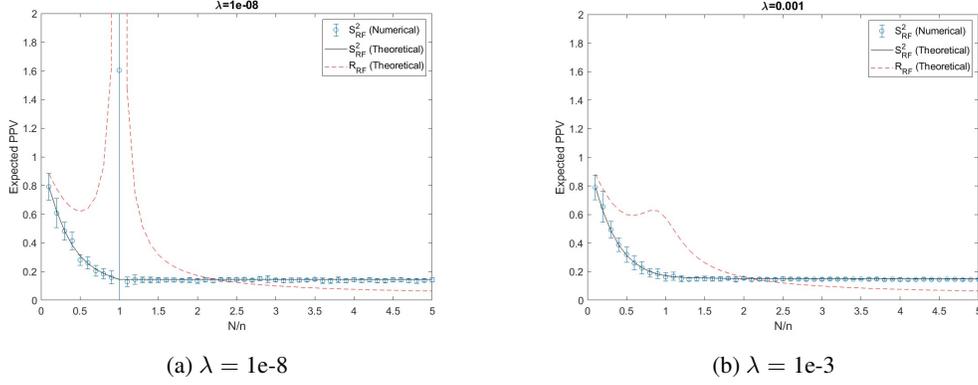


Figure 1: Comparison of asymptotic formula and 20 instances of S^2_{RF} (12). Data are generated via noiseless linear model $y = \langle x, \beta \rangle$ ($\|\beta\| = 1$, $\rho = \infty$). Activation is ReLU: $\sigma(x) = \max\{0, x\}$. d and n are fixed to 100 and 300, respectively. The asymptotic formula for R_{RF} (8) is plotted for comparison (red, dashed).

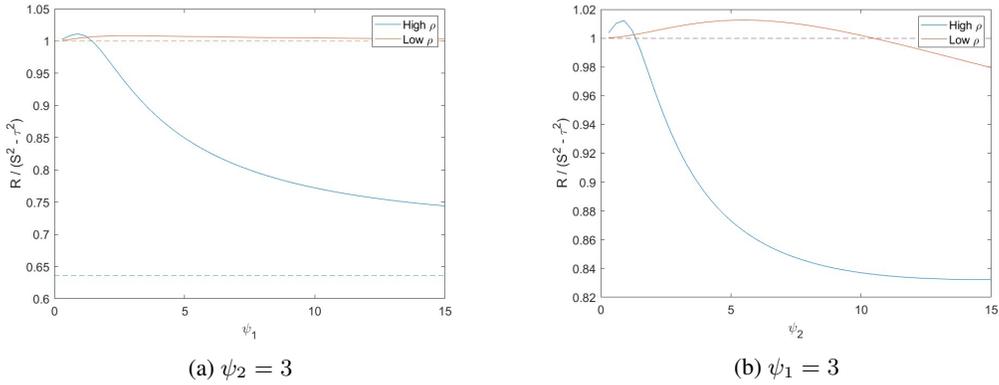


Figure 2: Ratio of $\mathcal{R}(\lambda^{opt})$ to $\mathcal{S}^2(\lambda^{opt}) - \tau^2$ as a function of ψ_1 (2a) and of ψ_2 (2b). In each plot, ψ_2 and ψ_1 are respectively fixed to 3, while $F_1 = 1$, $F_* = 0$, and $\rho = 1/\tau^2$ for noise variance $\tau^2 \in \{.2, 5\}$.

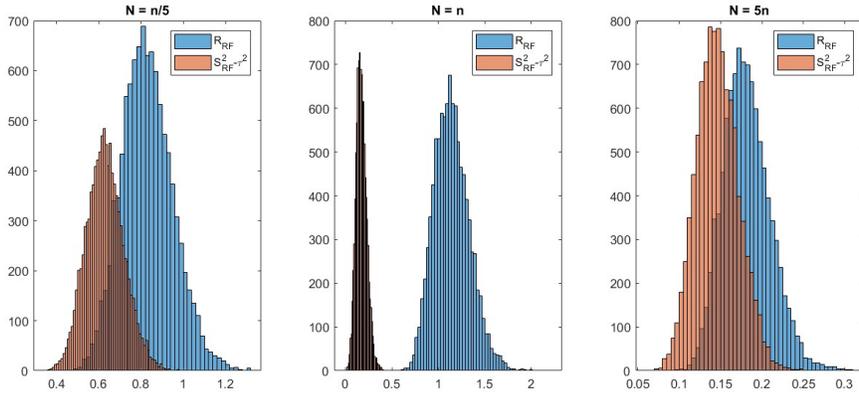


Figure 3: Histograms of $1e4$ draws of R_{RF} and $S^2_{RF} - \tau^2$ under low-noise linear model $y = \langle x, \beta \rangle + \tau^2$ with $\tau^2 = 1/5$.

3 suggests an interesting phenomenon: both quantities appear Gaussian, but are nearly orthogonal precisely near the “interpolation boundary” where R_{RF} exhibits double descent (2nd subplot). Our

empirical results should be compared with the results of Freedman [1999] and Johnstone [2010]. The Gaussianity of R_{RF} and S_{RF}^2 , if true, strengthens the agreement between the expected width S_{RF}^2 and expected risk R_{RF} immediately transfers to the agreement between the frequentist coverage of $(1 - \alpha) - \%$ Bayes credible ball around the mean and the nominal coverage. Again, such claims are generally true in finite-dimensional settings, but not true for high-dimensional settings. For instance, Freedman [1999] showed that asymptotically, the posterior variance fluctuates like a Gaussian random variable with a variance strictly smaller than that of the frequentist risk. Extracting such information is possible only if we study *second-order information* of R_{RF} and S_{RF}^2 ; in particular, suggests a Gaussian central limit for these quantities. A rigorous proof of a central limit theorem for R_{RF} and S_{RF}^2 goes beyond the scope of this work. Nevertheless, we conjecture the following for the fluctuations of R_{RF} and S_{RF}^2 :

1. Both $d(R_{RF} - \mathbb{E}R_{RF})$ and $d(S_{RF}^2 - \mathbb{E}S_{RF}^2)$ weakly converge to Gaussian distribution with appropriate variances. The faster convergence rate of d rather than \sqrt{d} is known to be common in central limit theorems for linear spectral statistics [Lytova and Pastur, 2009].
2. When $\psi_1 \leq \psi_2$, asymptotic variance of S_{RF}^2 is smaller than R_{RF} . When $\psi_1 = \psi_2$, the two distributions are nearly orthogonal. For large enough ψ_1 , the asymptotic variances are of the same order.

These conjectures are of independent interest and pose interesting theoretical challenges. We must note that second-order asymptotics has received less attention in the deep learning community. Only recently did Li et al. [2021] study the asymptotic normality of prediction risk achieved by a minimum least squares interpolator. Their result relies on a central limit theorem for linear statistics of eigenvalues of large sample covariance matrices, established by Bai and Silverstein [2004]. Many central limit theorems in random matrix theory seem insufficient to handle kernel learning or learning with random features.

4 Discussion

In sum, our calculations and empirical findings suggest that the RF model still has some interesting discoveries to offer, particularly for those interested in the implications of Bayesian deep learning. The asymptotics of the posterior predictive summaries can be very different from that of the generalization error. Numerical simulations suggest that some version of the central limit holds so that depending on the particular scaling of dimensions, the frequentist coverage of the posterior predictive distribution can be arbitrarily low or high. We now conclude with a discussion of several conjectures and the next directions suggested by these findings.

Heuristic Implications. While our technical calculations are only applicable to the RF model, we believe the general perspective of comparing frequentist versus Bayesian estimation approaches to the deep learning model can offer useful heuristics explaining empirically observed phenomena in training Bayesian deep learning models.

One suggestion is that the different regimes of asymptotics may explain away the “cold posterior effect”, first empirically observed by Wenzel et al. [2020]. The authors forcefully suggested that raising the posterior of a deep learning model to an even larger power (i.e., more concentrated) can improve generalization. That the posteriors can be too wide in highly overparameterized, near-noiseless regimes may explain this issue. Even though in this simple RF model setup, the posterior predictive mean is left unchanged from that of the posterior, it is plausible that for more complicated models this need not be so. When that posterior is much wider than the expected risk of the mean, we will mix over too many bad predicting weights, so it will be actually worse to average over the posterior than not.

Technical Improvements. The technically most interesting but quite challenging direction, suggested by our numerical simulations in Section 3.3, is to study second-order fluctuations of random matrix quantities often encountered in deep learning. While at this point a vast amount of literature exists that systematically overcomes the non-standard definitions of quantities like generalization error from the viewpoint of random matrix theory, we are aware of no study as of yet that overcomes the same issues to show a central limit-type result. On the other hand, central limit theorems for linear spectral statistics remain an active area of research in random matrix theory and are generally acknowledged to require more advanced, subtler arguments that are fundamentally different from first-order results

(i.e., convergence of empirical spectral distribution). For a start, a separate work on the central limit theorem for linear spectral statistics of kernel matrices is in preparation by the authors.

Another interesting direction is to extend the comparison between frequentist and Bayesian modeling approaches to more complex models of training dynamics. For instance, Adlam and Pennington [2020] have derived precise asymptotics of generalization error in the neural tangent kernel (NTK) regime, in which the features Θ are “learned” through a linearized approximation of the gradient flow in training. Mei et al. [2018, 2019] have suggested a different asymptotic model, in which Θ evolve through nonlinear dynamics. An interesting challenge, both technically and conceptually, is to adapt their analysis to the Bayesian setting, where the gradient flow on the space of posterior measures of Θ is approximated and specific projections are tracked.

Making the “Right” Comparison. A possible objection is that the comparison between two quantities is not appropriate: if interested in the coverage of function f_d with our credible set, then one must compare the expected width of the interval of the function space posterior, \hat{f} , instead of that of the Bayesian model average. While the objection is sensible, the fact is that the average variance of the function posterior will be an even worse uncertainty estimate for underparameterized models, whereas it will make little difference for overparameterized models. The reason is that the variance of the posterior of \hat{f} only differs from that of the posterior predictive, in (11), by the training set error term (14). The training error dominates the PPV precisely for underparameterized models ($\psi_1 < \psi_2$), so the curves of Figure 1 will now be non-*decreasing* and asymptotically convergent to a similar constant in expectation. This fact illustrates that a simple interpretation of the width of the posterior is complicated in even simplistic high-dimensional models.

Another possible objection is that the same value of λ need not be optimal for the posterior predictive. Since the asymptotic curve for S^2 is non-increasing in λ , this implies one must choose even smaller λ than λ^{opt} to obtain less conservative credible intervals, regardless of the SNR ρ . This seems an odd consequence of using overparameterized models, where a good frequentist coverage is obtained by choosing a narrower prior than that with good “contraction property” (i.e., minimal L^2 risk of posterior mean).

The final objection is that the RF model is simply not a very realistic model of practical Bayesian deep learning. We point out that quantifying the behavior of RF models as approximations of kernel learning methods remains a valid concern for Bayesian practitioners.

Acknowledgments

All authors would like to thank Matthew M. Engelhard, Boyao Li, David Page, and Alexander J Thomson for the helpful discussion around the theoretical work.

Samuel I. Berchuck would like to acknowledge support from the National Eye Institute of the National Institutes of Health under Award Number K99EY033027. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. Sayan Mukherjee would like to acknowledge partial funding from HFSP RGP005, NSF DMS 17-13012, NSF BCS 1552848, NSF DBI 1661386, NSF IIS 15-46331, NSF DMS 16-13261, as well as high-performance computing partially supported by grant 2016- IDG-1013 from the North Carolina Biotechnology Center as well as the Alexander von Humboldt Foundation, the BMBF and the Saxony State Ministry for Science.

References

- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986, 2022.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.

- Hong Hu and Yue M Lu. Sharp asymptotics of kernel ridge regression beyond the linear regime. *arXiv preprint arXiv:2205.06798*, 2022.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- Mikhail Belkin. Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numerica*, 30:203–248, 2021.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.
- Vincent Fortuin, Adrià Garriga-Alonso, Sebastian W Ober, Florian Wenzel, Gunnar Rätsch, Richard E Turner, Mark van der Wilk, and Laurence Aitchison. Bayesian neural network priors revisited. *arXiv preprint arXiv:2102.06571*, 2021.
- BJK Kleijn and AW van der Vaart. The bernstein-von-mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.
- Dennis D Cox. An analysis of bayesian inference for nonparametric regression. *The Annals of Statistics*, pages 903–923, 1993.
- David Freedman. Wald lecture: On the bernstein-von mises theorem with infinite-dimensional parameters. *The Annals of Statistics*, 27(4):1119–1141, 1999.
- Iain M Johnstone. High dimensional bernstein-von mises: simple examples. *Institute of Mathematical Statistics Collections*, 6:87, 2010.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in neural information processing systems*, 32, 2019.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR, 2019.
- Pier Giovanni Bissiri, Chris C Holmes, and Stephen G Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130, 2016.
- Lorin Crawford, Kris C. Wood, Xiang Zhou, and Sayan Mukherjee. Bayesian approximate kernel regression with variable selection. *Journal of the American Statistical Association*, 113(524):1710–1721, 2018.
- BT Knapik, AW van der Vaart, and JH van Zanten. Bayesian inverse problems with gaussian priors. *The Annals of Statistics*, 39(5):2626–2657, 2011.
- Lucas Clarté, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Theoretical characterization of uncertainty in high-dimensional linear classification. *Machine Learning: Science and Technology*, 4(2):025029, 2023a.
- Lucas Clarté, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. On double-descent in uncertainty quantification in overparametrized models. In *International Conference on Artificial Intelligence and Statistics*, pages 7089–7125. PMLR, 2023b.
- Alice Guionnet, Justin Ko, Florent Krzakala, and Lenka Zdeborová. Estimating rank-one matrices with mismatched prior and noise: universality and large deviations, 2023.

- A Lytova and L Pastur. Central limit theorem for linear eigenvalue statistics of random matrices with independent entries. *Annals of Probability*, 37(5):1778–1840, 2009.
- Zeng Li, Chuanlong Xie, and Qinwen Wang. Asymptotic normality and confidence intervals for prediction risk of the min-norm least squares estimator. In *International Conference on Machine Learning*, pages 6533–6542. PMLR, 2021.
- ZD Bai and Jack W Silverstein. Clt for linear spectral statistics of large-dimensional sample covariance matrices. *The Annals of Probability*, 32(1A):553–605, 2004.
- Florian Wenzel, Kevin Roth, Bastiaan S Veeling, Jakub Świkatkowski, Linh Tran, Stephan Mandt, Jasper Snoek, Tim Salimans, Rodolphe Jenatton, and Sebastian Nowozin. How good is the bayes posterior in deep neural networks really? *arXiv preprint arXiv:2002.02405*, 2020.
- Ben Adlam and Jeffrey Pennington. The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization. In *International Conference on Machine Learning*, pages 74–84. PMLR, 2020.

A Proofs

Proof of Proposition 1

The necessary calculations follow immediately from the proof of Mei and Montanari [2022]. Since we do not wish to replicate the long proof with only minor readjustments, we only sketch the main argument. The expected PPV consists of two random variables, denoted by

$$\widehat{\phi}^{-1}(f_d, \mathbf{X}, \Theta, \lambda) = \frac{\langle \mathbf{y}, \mathbf{y} - \widehat{f}(\mathbf{x}) \rangle}{n}, \quad r(\mathbf{X}, \Theta, \lambda) := \mathbb{E}_{\mathbf{x}}[\boldsymbol{\sigma}(\mathbf{x})^T (\mathbf{Z}^T \mathbf{Z} + \lambda \psi_1 \psi_2 \mathbf{I}_N)^{-1} \boldsymbol{\sigma}(\mathbf{x}) / d]. \quad (23)$$

It suffices to show the two quantities individually converge in probability to some limits. The first of the two is a training set error, for which an asymptotic formula was derived in Section 6, Mei and Montanari [2022]:

$$\lim_{d \rightarrow \infty} \mathbb{E} \widehat{\phi}^{-1} = -i \nu_2 (i \sqrt{\psi_1 \psi_2 \lambda} / \mu_*) \sqrt{\frac{\lambda \psi_1}{\psi_2 \mu_*^2} \left(\frac{F_1^2}{1 - \chi \zeta^2} + F_*^2 + \tau^2 \right)}, \quad (24)$$

using the definition of quantities in Proposition 1, for $\chi \equiv \chi(i \sqrt{\psi_1 \psi_2 \lambda} / \mu_*) = \nu_1 \nu_2$ ($i = \sqrt{-1}$). At this point, rearranging the fixed point equations (20) and (21) with $\xi \equiv i \sqrt{\psi_1 \psi_2 \lambda} / \mu_*$, we obtain:

$$-\psi_1 - i \sqrt{\lambda \psi_1 \psi_2} / \mu_* \nu_1 = -\psi_2 - i \sqrt{\lambda \psi_1 \psi_2} / \mu_* \nu_2 = -\chi - \frac{\zeta^2 \chi}{1 - \zeta^2 \chi}. \quad (25)$$

$$\implies -i \sqrt{\frac{\lambda \psi_1}{\psi_2 \mu_*^2}} \nu_2 = 1 + \frac{\chi}{\psi_2} \left(\frac{\zeta^2}{1 - \zeta^2 \chi} + 1 \right). \quad (26)$$

A similar formula for $\mathbb{E} r$ can also be derived, using almost exactly the same steps required for simplifying asymptotic formulae of R_{RF} and $\widehat{\phi}^{-1}$. Assumption 3 on the nonlinear component of the generating model, f_d^{NL} , implies that it admits an explicit decomposition into spherical harmonics that form a basis on \mathbb{S}^{d-1} and coefficients with Gaussian distributions (appropriately scaled). With many simplifications, also sketched out in Appendix E, Mei and Montanari [2022], we can write

$$\mathbb{E} r = \mu_1^2 \text{Tr}((\mathbf{Z}^T \mathbf{Z} + \lambda \psi_1 \psi_2 \mathbf{I}_N)^{-1} (\Theta \Theta^T / d)) + \mu_*^2 \text{Tr}((\mathbf{Z}^T \mathbf{Z} + \lambda \psi_1 \psi_2 \mathbf{I}_N)^{-1}) + o_d(1). \quad (27)$$

Next, form an $(N + n) \times (N + n)$ -dimensional block matrix \mathbf{A} , indexed by a 5-d real parameter $\mathbf{q} = (s_1, s_2, t_1, t_2, p)$:

$$\mathbf{A}(\mathbf{q}) = \begin{bmatrix} s_1 \mathbf{I}_N + s_2 \Theta \Theta^T / d & \mathbf{Z}^T + p \Theta \mathbf{X}^T / \sqrt{d} \\ \mathbf{Z} + p \mathbf{X} \Theta^T / \sqrt{d} & t_2 \mathbf{I}_n + t_2 \mathbf{X} \mathbf{X}^T / d \end{bmatrix}. \quad (28)$$

Denoting by $\xi \in \mathbb{C}_+$ the log-determinant of the resolvent matrix $\mathbf{A}(\mathbf{q}) - \xi \mathbf{I}_{N+n}$ as $G_d(\xi; \mathbf{q})$, straightforward calculus and plugging in the figures reveals

$$\mathbb{E} r = -\mu_1^2 \cdot i \frac{1}{\sqrt{\psi_1 \psi_2 \lambda}} \mathbb{E}[\partial_{s_1} G_d(i \sqrt{\lambda \psi_1 \psi_2}; 0)] - \mu_*^2 \cdot i \frac{1}{\sqrt{\psi_1 \psi_2 \lambda}} \mathbb{E}[\partial_{s_2} G_d(i \sqrt{\lambda \psi_1 \psi_2}; 0)] + o_d(1). \quad (29)$$

At this point, we pass on to the limit $d \rightarrow \infty$ and plug in appropriate asymptotic limits of the partial derivatives. Since the ξ -derivative of G_d is the trace of $\mathbf{A}(\mathbf{q}) - \xi \mathbf{I}_{N+n}$, this is possible by characterizing the asymptotics of the Stieltjes transform of the empirical spectral distribution of $\mathbf{A}(\mathbf{q})$, which is done using Proposition 8.3, Mei and Montanari [2022]. Integrating the limit over path stretching in the imaginary axis (iK with $K \rightarrow \infty$), one then rigorously passes onto the limit by Proposition 8.4, Mei and Montanari [2022]. Final simplifications using Lemma 8.1, Mei and Montanari [2022] yield the formula

$$\lim_{d \rightarrow \infty} \mathbb{E} r = -i \nu_1 (i \sqrt{\psi_1 \psi_2 \lambda} / \mu_*) \frac{\mu_*}{\sqrt{\lambda \psi_1 \psi_2}} \left(\frac{\zeta^2}{1 - \zeta^2 \chi} + 1 \right). \quad (30)$$

At this point, we may conclude, using (23), (30) and (26):

$$\widehat{\phi}^{-1}(1 + r) \xrightarrow{P} \frac{F_1^2}{1 - \zeta^2 \chi} + F_*^2 + \tau^2.$$

Proof of Proposition 2

The first point is a consequence of the fact that S_{RF}^2 is decreasing in λ and passing onto the limits. The second point is immediately seen by taking the limit of χ as $\lambda \rightarrow 0$, a formula for which is provided by Theorem 3, Mei and Montanari [2022].

Proof of Proposition 4

We want to check whether equality between asymptotic risk and the expected PPV holds. Using Proposition 3: In the highly overparameterized regime, we demand equality

$$\begin{aligned} \mathcal{R}_{wide} &= \frac{F_1^2 \psi_2 + (F_*^2 + \tau^2) \omega_2^2}{\psi_2 - 2\omega_2 \psi_2 + \omega_2^2 \psi_2 - \omega_2^2} \stackrel{*}{=} \frac{F_1^2}{1 - \omega_2}. \\ \Leftrightarrow \omega_2^3 + (\rho \psi_2 - \rho - 1) \omega_2^2 - \psi_2 \rho \omega_2 &\stackrel{*}{=} 0. \end{aligned}$$

By Proposition 13.1, Mei and Montanari [2022], this equation is met precisely by choosing $\bar{\lambda} \equiv \lambda^{opt}$, only possible if $\rho < \rho^*$ per the definition of Proposition 4. Otherwise (i.e., if $\rho > \rho^*$), the equality is unsatisfiable, and furthermore, one can use the property of the map $\omega \equiv \omega(\lambda, \psi, \zeta)$ in Lemma 13.1, Mei and Montanari [2022], to show that $\omega_2^3 + (\rho \psi_2 - \rho_1) \omega_2^2 - \psi_2 \rho \omega_2 > 0$.

In the large sample regime, we similarly want to check the equality

$$\begin{aligned} \frac{\omega_1^2 / \zeta^2 + \psi_1}{\psi_1 - 2\psi_2 \omega_1 + \psi_1 \omega_1^2 - \omega_1^2} &\stackrel{*}{=} \frac{1}{1 - \omega_1^2}. \\ \Leftrightarrow \omega_1^3 + (\zeta^2 \psi_1 - \zeta^2 - 1) \omega_1^2 - \psi_1 \zeta^2 \omega_1 &\stackrel{*}{=} 0. \end{aligned}$$

In this case, the demanded equality does not have ρ in it. Utilizing symmetry and the same argument as in Proposition 13.1, Mei and Montanari [2022], one can show that the absence of ρ implies the equality is always satisfiable by choosing $\lambda = 0$, which also corresponds to the minimizer of the asymptotic risk: $\arg \min_{\lambda \geq 0} \mathcal{R}_{lsamp} = 0$.