



Wilkinson's Bus: Weak Condition Numbers, with an Application to Singular Polynomial Eigenproblems

Martin Lotz¹ · Vanni Noferini²

Received: 28 May 2019 / Revised: 22 January 2020 / Accepted: 4 February 2020 /

Published online: 10 March 2020

© The Author(s) 2020

Abstract

We propose a new approach to the theory of conditioning for numerical analysis problems for which both classical and stochastic perturbation theories fail to predict the observed accuracy of computed solutions. To motivate our ideas, we present examples of problems that are discontinuous at a given input and even have infinite stochastic condition number, but where the solution is still computed to machine precision without relying on structured algorithms. Stimulated by the failure of classical and stochastic perturbation theory in capturing such phenomena, we define and analyse a weak worst-case and a weak stochastic condition number. This new theory is a more powerful predictor of the accuracy of computations than existing tools, especially when the worst-case and the expected sensitivity of a problem to perturbations of the input is not finite. We apply our analysis to the computation of simple eigenvalues of matrix polynomials, including the more difficult case of singular matrix polynomials. In addition, we show how the weak condition numbers can be estimated in practice.

Keywords Condition number · Stochastic perturbation theory · Weak condition number · Polynomial eigenvalue problem · Singular matrix polynomial

Mathematics Subject Classification 15A15 · 15A18 · 15B52 · 60H99 · 65F15 · 65F35

Communicated by Hans Munthe-Kaas.

✉ Vanni Noferini
vanni.noferini@aalto.fi

Martin Lotz
martin.lotz@warwick.ac.uk

¹ Mathematics Institute, The University of Warwick, Coventry, UK

² Department of Mathematics and Systems Analysis, Aalto University, Espoo, Finland

1 Introduction

The *condition number* of a computational problem measures the sensitivity of an output with respect to perturbations in the input. If the input–output relationship can be described by a differentiable function f near the input, then the condition number is the norm of the derivative of f . In the case of solving systems of linear equations, the idea of conditioning dates back at least to the work of von Neumann and Goldstine [47] and Turing [45], who coined the term. For an algorithm computing f in finite precision arithmetic, the importance of the condition number κ stems from the “rule of thumb” popularized by Higham [29, §1.6],

$$\text{forward error} \lesssim \kappa \cdot (\text{backward error}).$$

The backward error is small if the algorithm computes the exact value of f at a *nearby* input, and a small condition number would certify that this is enough to get a small overall error. Higham’s rule of thumb comes from a first-order expansion, and in practice it often holds as an approximate equality and is valuable for practitioners who wish to predict the accuracy of numerical computations. Suppose that a solution is computed with, say, a backward error equal to 10^{-16} . If $\kappa = 10^2$, then one would trust the computed value to have (at least) 14 meaningful decimal digits.

The condition number can formally still be defined when f is not differentiable, though it may not be finite. If f is not locally Lipschitz continuous at an input, then the condition number is $+\infty$; a situation clearly beyond the applicability of Higham’s rule. Inputs at which the function f is not continuous are usually referred to as *ill-posed*. Based on the worst-case sensitivity, one would usually only expect a handful of correct digits when evaluating a function at such an input, and quite possibly none.¹ On the other hand, a small condition number is not a necessary condition for a small forward–backward error ratio: it is not inconceivable that certain ill-conditioned or even ill-posed problems can be solved accurately. Consider, for example, the problem of computing an eigenvalue of the 4×4 matrix pencil (linear matrix polynomial)

$$L(x) = \begin{bmatrix} -1 & 1 & 4 & 2 \\ -2 & 3 & 12 & 6 \\ 1 & 3 & 11 & 6 \\ 2 & 2 & 7 & 4 \end{bmatrix} x + \begin{bmatrix} 2 & -1 & -5 & -1 \\ 6 & -2 & -11 & -2 \\ 5 & 0 & -2 & 0 \\ 3 & 1 & 3 & 1 \end{bmatrix}; \quad (1)$$

this is a singular matrix pencil (the determinant is identically zero) whose only finite eigenvalue is simple and equal to 1 (see Sect. 3 for the definition of an eigenvalue of a singular matrix polynomial and other relevant terminology). The input is $L(x)$ and the

¹ The number of accurate digits that can be expected when the problem is continuous but not locally Lipschitz continuous requires a careful discussion. It depends on the unit roundoff u , on the exact nature of the pathology of f , and on D . For example, computing the eigenvalues of a matrix similar to an $n \times n$ Jordan block for $n > 1$ is Hölder continuous with exponent $1/n$ but not Lipschitz continuous. Usually this translates into expecting only about $\sqrt[n]{u}$ accuracy, up to constants, when working in finite precision arithmetic. For a more complete discussion, see [28], where pathological examples of derogatory matrices are constructed, whose eigenvalues are not sensitive to finite precision computations (for fixed u), or also [33, §3.3]. For discontinuous f , however, these subtleties alone cannot justify any accurately computed decimal digits.

solution is 1. If the QZ algorithm [36], which is the standard eigensolver for pencils, is called via MATLAB’s command `eig`,² the output is:

```
>> eig(L0,-L1)

ans =

-138.1824366539536
-0.674131242894470
 1.000000000000000
 0.444114486065683
```

All but the third computed eigenvalues are complete rubbish. This is not surprising: singular pencils form a proper Zariski closed set in the space of matrix pencils of a fixed format, and it is unreasonable to expect that an unstructured algorithm would detect that the input is singular and return only one eigenvalue. Instead, being backward stable, QZ computes the eigenvalues of some nearby matrix pencil, and almost all nearby pencils have 4 eigenvalues. On the other hand, the accuracy of the approximation of the genuine eigenvalue 1 is quite remarkable. Indeed, the condition number of the problem that maps $L(x)$ to the exact eigenvalue 1 is infinite because the map from matrix pencils to their eigenvalues is discontinuous at any matrix pencil whose determinant is identically zero. To make matters worse, there exist plenty of matrix pencils arbitrarily close to $L(x)$ whose eigenvalues are *all* nowhere near 1. For example, for any $\epsilon > 0$,

$$\hat{L}(x) = L(x) + \epsilon \left(\begin{bmatrix} 0 & -1 & -4 & -1 \\ 1 & -3 & -13 & -3 \\ 0 & -2 & -8 & -2 \\ -1 & -1 & -3 & -1 \end{bmatrix} x + A \right),$$

where

$$A = \begin{bmatrix} -1 & -1 & -3 & -2 \\ -3 & -3 & -9 & -6 \\ -2 & -2 & -6 & -4 \\ -1 & -1 & -3 & -2 \end{bmatrix} \gamma_0 + \begin{bmatrix} 1 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \gamma_1 + \begin{bmatrix} 0 & -1 & -4 & -1 \\ 0 & -3 & -12 & -3 \\ 0 & -2 & -8 & -2 \\ 0 & -1 & -4 & -1 \end{bmatrix} \gamma_2 + \begin{bmatrix} 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 \end{bmatrix} \gamma_3,$$

has characteristic polynomial $\epsilon^2(\gamma_3 - x)(x^3 + \gamma_2x^2 + \gamma_1x + \gamma_0)$ and therefore, by an arbitrary choice of the parameters γ_i , can have eigenvalues literally *anywhere*. Yet, unaware of this worrying caveat, the QZ algorithm computes an excellent approximation of the exact eigenvalue: 16 correct digits! This example has not been carefully cherry picked: readers are encouraged to experiment with any singular input in order to convince themselves that QZ often computes³ accurately the (simple) eigenvalues

² MATLAB R2016a on Ubuntu 16.04.
³ Of course, if the exact solution is not known a priori, one faces the practical issue of deciding which of the computed eigenvalues is reliable. There are various ways in which this can be done in practice, such as artificially perturbing the problem; the focus of our work is on explaining why the correct solution has been shortlisted in the first place; see [32] for a more practical perspective.

of singular pencils, or singular matrix polynomials, in spite of being a discontinuous problem. See also [32] for more examples and a discussion of applications. Although the worst-case sensitivity to perturbations is indeed infinite, the *raison d'être* of the condition number, which is to predict the accuracy of computations on a computer, is not fulfilled.

Why does the QZ algorithm accurately compute the eigenvalue, when the map f describing this computational problem is not even continuous? Two natural attempts at explaining this phenomenon would be to look at *structured condition numbers* and/or *average-case (stochastic) perturbation theory*.

1. An algorithm is *structured* if it computes the exact solution to a perturbed input, where the perturbations respect some special features of the input: for example, singular, of rank 3, triangular, or with precisely one eigenvalue. The vanilla implementation of QZ used here is unstructured in the sense that it does not preserve any of the structures that would explain the strange case of the algorithm that computes an apparently uncomputable eigenvalue.⁴ It does, however, preserve the real structure. In other words, if the input is real, QZ computes the eigenvalues of a nearby *real* pencil. Yet, by taking real γ_i in the example above, it is clear that there are real pencils arbitrary close to $L(x)$ whose eigenvalues are all arbitrarily far away from 1.
2. The classical condition number is based on the worst-case perturbation of an input; as discussed in [29, §2.8], this approach tends to be overly pessimistic in practice. Numerical analysis pioneer James Wilkinson, in order to illustrate that Gaussian elimination is unstable in theory, but in practice its instability is only observed by mathematicians looking for it, is reported to have said [43]

Anyone that unlucky has already been run over by a bus.

In other words: in Wilkinson's experience, the likelihood of seeing the admittedly terrifying worst case appeared to be very small, and therefore, Wilkinson believed that being afraid of the potential catastrophic instability of Gaussian elimination is an irrational attitude. Based on this experience, Weiss et al. [48] and Stewart [41] proposed to study the effect of perturbations *on average*, as opposed to worst case; see [29, §2.8] for more references on work addressing the stochastic analysis of roundoff errors. This idea was later formalized and developed further by Armentano [4]. This approach gives some hope to explain the example above, because it is known that the set of perturbations responsible for the discontinuity of f has measure zero [18]. However, this does not imply that on average perturbations are not harmful. In fact, as we will see, the stochastic condition number for the example above (or for similar problems) is still infinite! Average-case perturbation analysis, at least in the form in which it has been used so far, is still unable to solve the puzzle.

While neither structured nor average-case perturbation theory can explain the phenomenon observed above, Wilkinson's colourful quote does contain a hint on how to proceed: shift attention from average-case analysis of perturbations to bounding

⁴ There exist algorithms able to detect and exploit the fact that a matrix pencil is singular, such as the staircase algorithm [46].

rare events. We will get back to the matrix pencil (1) in Example 5.3, where we show that our new theory does explain why this problem is solved to high accuracy using standard backward stable algorithms.

In summary, the main contributions of this paper are

1. a new species of “weak” condition numbers, which we call the *weak worst-case condition number* and the *weak stochastic condition number* that give a more accurate description of the perturbation behaviour of a computational map (Sect. 2);
2. a precise probabilistic analysis of the sensitivity of the problem of computing simple eigenvalues of singular matrix polynomials (Sects. 4 and 5);
3. an illustration of the advantages of the new concept by demonstrating that, unlike both classical and stochastic condition numbers, the weak condition numbers *are* able to explain why the apparently uncomputable eigenvalues of singular matrix polynomials, such as the eigenvalue 1 in the example above, can be computed with remarkable accuracy (Example 5.3);
4. a concrete method for bounding the weak condition numbers for the eigenvalues of singular matrix polynomials (Sect. 6).

1.1 Related Work

Rounding errors, and hence the perturbations considered, are not random [29, 1.17]. Nevertheless, the observation that the computed bounds on rounding errors are overly pessimistic has led to the study of statistical and probabilistic models for rounding errors. An early example of such a statistical analysis is Goldstine and von Neumann [27], see [29, 2.8] and the references therein for more background. Recently, Higham and Mary [31] have obtained probabilistic rounding error bounds for a wide variety of algorithms in linear algebra. In particular, they give a rigorous foundation to Wilkinson’s rule of thumb, which states that constants in rounding error bounds can be safely replaced by their square roots.

The idea of using an average, rather than a supremum, in the definition of conditioning was introduced by Weiss et al. [48] in the context of the (matrix) condition number of solving systems of linear equations, and a more comprehensive stochastic perturbation theory was developed by Stewart [41]. In [4], Armentano introduced the concept of a smooth condition number and showed that it can be related to the worst-case condition. His work uses a geometric theory of conditioning and does not extend to singular problems.

The line of work on random perturbations is not to be confused with the *probabilistic analysis of condition numbers*, where a condition number is a given function, and the distribution of this function is studied over the space of inputs (see [13] and the references therein). Nevertheless, our work is inspired by the idea of weak average-case analysis [3] that was developed in this framework. Weak average-case analysis is based on the observation, which has origins in the work of Smale [40] and Kostlan [34], that discarding a small set from the input space can dramatically improve the expected value of a condition number, shifting the focus away from the *average case* and towards bounding the probability of rare events. Our contribution is to apply this line

of thought to study random perturbations instead of random inputs. However, we stress that we do not seek to model the distribution of perturbations. The aim is to formally quantify statements such as “the set of bad perturbations is small compared to the set of good perturbations”. In other words, the (non-random) accumulation of rounding errors in a procedure would need a very good reason to give rise to a badly perturbed problem.

The conditioning of regular polynomial eigenvalue problems has been studied in detail by Tisseur [42] and by Dedieu and Tisseur in a homogeneous setting [15]. A probabilistic analysis of condition numbers (for random inputs) for such problems was given by Armentano and Beltrán [5] over the complex numbers and by Beltrán and Kozhasov [7] over the real numbers. Their work studies the distribution of the condition number on the whole space of inputs, and such an analysis only considers the condition number of regular matrix polynomials. A perturbation theory for singular polynomial eigenvalue problems was developed by de Terán and Dopico [14], and our work makes extensive use of their results. A method to solve singular generalized eigenvalue problems with plain QZ, based on applying a certain perturbation to them, is proposed in [32] (see also the references therein); note that our work goes beyond this, by showing how to estimate the weak condition number that could guarantee, often with overwhelming probability, that QZ will do fine even without any preliminary perturbation step.

1.2 Organization of the Paper

The paper is organized as follows: In Sect. 2, we review the rigorous definitions of the worst-case (von Neumann–Turing) condition number and the stochastic framework (Weiss et al., Stewart, Armentano) and comment on their advantages and limitations. We then define the weak condition numbers as quantiles and argue that, even when Wilkinson’s metaphorical bus hits von Neumann–Turing’s and Armentano–Stewart’s theories of conditioning, ours comes well endowed with powerful dodging skills. In Sect. 3, we introduce the perturbation theory of singular matrix polynomials, along with the definitions of simple eigenvalues and eigenvectors. We define the input–output map underlying our case study and introduce the directional sensitivity of such problems. In Sect. 4, which forms the core of this paper, we carry out a detailed analysis of the probability distribution of the directional sensitivity of the problems introduced in Sect. 3. In Sect. 5, we translate the probabilistic results from Sect. 4 into the language of weak condition numbers and prove the main results, Theorems 5.1 and 5.2. In Sect. 6, we sketch how our new condition numbers can be estimated in practice. Along the way we derive a simple concentration bound on the directional sensitivity of regular polynomial eigenvalue problems. Finally, in Sect. 7, we give some concluding remarks and discuss potential further applications.

2 Theories of Conditioning

For our purposes, a computational problem is a map between normed vector spaces⁵

$$f : \mathcal{V} \rightarrow \mathcal{W}, \quad D \mapsto S := f(D),$$

and we will denote the (possibly different) norms in each of these spaces by $\| \cdot \|$. Following the remark on [30, p. 56], for simplicity of exposition in this paper we focus on absolute, as opposed to relative, condition numbers. The condition numbers considered depend on the map f and an input $D \in \mathcal{V}$.

As we are only concerned with the condition of a fixed computational problem at a fixed input D , in what follows we omit reference to f and D in the notation.

Definition 2.1 (*Worst-case condition number*) The condition number of f at D is

$$\kappa = \lim_{\epsilon \rightarrow 0} \sup_{\|E\| \leq 1} \frac{\|f(D + \epsilon E) - f(D)\|}{\epsilon \|E\|}.$$

If f is Fréchet differentiable at D , then this definition is equivalent to the operator norm of the Fréchet derivative of f . However, Definition 2.1 also applies (and can even be finite) when f is not differentiable. In complexity theory [8,13], an elegant geometric definition of condition number is often used, which is essentially equivalent to Definition 2.1 under certain assumptions (which include smoothness).

The following definition is loosely derived from the work of Stewart [41] and Armentano [4], based on earlier work by Weiss et. al. [48]. In what follows, we use the terminology $X \sim \mathcal{D}$ for a random variable with distribution \mathcal{D} and $\mathbb{E}_{X \sim \mathcal{D}}[\cdot]$ for the expectation with respect to this distribution.

Definition 2.2 (*Stochastic condition number*) Let E be a \mathcal{V} -valued random variable with distribution \mathcal{D} and assume that $\mathbb{E}_{E \sim \mathcal{D}}[E] = 0$ and $\mathbb{E}_{E \sim \mathcal{D}}[\|E\|^2] = 1$. Assume that the function f is measurable. Then, the stochastic condition number is

$$\kappa_s = \lim_{\epsilon \rightarrow 0} \mathbb{E}_{E \sim \mathcal{D}} \frac{\|f(D + \epsilon E) - f(D)\|}{\epsilon \|E\|}.$$

Remark 2.3 We note in passing that Definition 2.2 depends on the choice of a measure \mathcal{D} . This measure is a parameter that the interested mathematician should choose as convenient; this is of course not particularly different than the freedom one is given in picking a norm. In fact, it is often convenient to combine these two choices, using a distribution that is invariant with respect to a given norm. Typical choices that emphasize invariance are the uniform (on a sphere) or Gaussian distributions, and the Bombieri-Weyl inner product when dealing with homogeneous multivariate polynomials [13,

⁵ One can, more generally, allow \mathcal{V} and \mathcal{W} to be anything with a notion of distance, such as general metric spaces or Riemannian manifolds. All the definitions of condition can be adapted accordingly; in this paper, we focus on the case of normed vector spaces. We will also only need such a map to be defined locally near an input of interest.

16.1]. Technically speaking, the distribution is on the *space of perturbations*, rather than the space of inputs.

If f is differentiable at D and \mathcal{V} is finite dimensional, then it was observed by Armentano [4] that the stochastic condition number can be related to the worst-case one. We illustrate this relation in a simple but instructive special case. Consider the setting⁶ where $f: \mathbb{R}^m \rightarrow \mathbb{R}^n$ ($m \geq n$) is differentiable at $D \in \mathbb{R}^m$, so that κ is the operator norm of the differential. If $\sigma_1 \geq \dots \geq \sigma_m$ denote the singular values of $df(D)$ (with $\sigma_i = 0$ for $i > n$), then $\kappa = \sigma_1$. If \mathcal{D} is the uniform distribution on the sphere, then

$$\frac{1}{m}\kappa \stackrel{(a)}{\leq} \sigma_1 \mathbb{E}_{E \sim \mathcal{D}} |E_1| \leq \mathbb{E}_{E \sim \mathcal{D}} \left[\sqrt{\sum_i \sigma_i^2 E_i^2} \right] \stackrel{(b)}{=} \mathbb{E}_{E \sim \mathcal{D}} [\|df(D)E\|_2] = \kappa_s, \quad (2)$$

where for (a) we used the fact that

$$\mathbb{E}_{E \sim \mathcal{D}} |E_1| = \frac{1}{m} \mathbb{E}_{E \sim \mathcal{D}} \|E\|_1 \geq \frac{1}{m} \mathbb{E}_{E \sim \mathcal{D}} \|E\|_2 = \frac{1}{m}$$

and for (b) we used the orthogonal invariance of the uniform distribution on the sphere. As we will see in the case of singular polynomial eigenvalue problems with complex perturbations, the bound (2) does not hold in general, as the condition number can be infinite while the stochastic condition number is bounded. However, sometimes it can happen that the stochastic condition number is also infinite, because the “directional sensitivity” (see Definition 2.4) is not an integrable function. For example, for the problem of computing the eigenvalue of the singular pencil $L(x)$ in the introduction, in spite of the fact that real perturbations are analytic for all but a proper Zariski closed set of perturbations [18], when restricting to real perturbations, we get

$$\kappa_s = \kappa = \infty.$$

Despite this, QZ computes the eigenvalue 1 with 16 digits of accuracy.

To remedy the shortcomings of the stochastic condition number as defined in 2.2, we propose a change in focus from the expected value to tail bounds and quantiles, and the key concept for that purpose is the directional sensitivity. Just as the classical worst-case condition corresponds to the norm of the derivative, the directional sensitivity corresponds to a directional derivative. And, just as a function can have some, or all, directional derivatives while still not being continuous, a computational problem can have well-defined directional sensitivities but have infinite condition number.

Definition 2.4 (*Directional sensitivity*) The *directional sensitivity* of the computational problem f at the input D with respect to the perturbation E is

$$\sigma_E = \lim_{\epsilon \rightarrow 0} \frac{\|f(D + \epsilon E) - f(D)\|}{\epsilon \|E\|}.$$

⁶ Armentano’s results apply to differentiable maps between Riemannian manifolds and cover the moments of the directional derivative as well: they are stronger and are derived with a more comprehensive approach.

The directional sensitivity takes values in $[0, \infty]$. In numerical analytic language, the directional sensitivity is the limit, for a particular direction of the backward error, of the ratio of forward and backward errors of the computational problem f ; this limit is taken letting the backward error tend to zero (again having fixed its direction), which could also be thought of as letting the unit roundoff tend to zero. See e.g. [29, §1.5] for more details on this terminology.

The directional sensitivity is, if it is finite, $\|E\|^{-1}$ times the norm of the Gâteaux derivative $df(D; E)$ of f at D in direction E . If f is Fréchet differentiable, then the Gâteaux derivative agrees with the Fréchet derivative, and we get

$$\kappa = \sup_{\|E\| \leq 1} \sigma_E.$$

If E is a \mathcal{V} -valued random variable satisfying the conditions of Definition 2.2 and if f is Gâteaux differentiable in almost all directions, then by the Fatou–Lebesgue theorem we get

$$\kappa_s = \mathbb{E}[\sigma_E].$$

When integrating, null sets can be safely ignored; however, depending on the exact nature of the divergence (or lack thereof) of the integrand when approaching those null sets, the value of the integral need not be finite. To overcome this problem and still give probabilistically meaningful statements, we propose to use instead the concept of *numerical null sets*, i.e. sets of finite but small (in a sense that can be made precise depending on, for example, the unit roundoff of the number system of choice, the confidence level required by the user, etc.) measure. This is analogous to the idea that the “numerical zero” is the unit roundoff. We next define our main characters, two classes of weak condition numbers which generalize, respectively, the classical worst-case and stochastic condition numbers.

In the following, we fix a probability space $(\Omega, \Sigma, \mathbb{P})$ and a random variable $E : \Omega \rightarrow \mathcal{V}$, where we consider \mathcal{V} endowed with the Borel σ -algebra. We further assume that

$$\mathbb{E}[E] = \int_{\Omega} E(\omega) \, d\mathbb{P}(\omega) = 0, \quad \mathbb{E}[\|E\|^2] = \int_{\Omega} \|E(\omega)\|^2 \, d\mathbb{P}(\omega) = 1.$$

The following definitions assume that σ_E is \mathbb{P} -measurable. This is the case, for example, if f is measurable and the directional (Gâteaux) derivative $df(D; E(\omega))$ exists \mathbb{P} -a.e.

Definition 2.5 (*Weak worst-case and weak stochastic condition number*) Let $0 \leq \delta < 1$ and assume that σ_E is \mathbb{P} -measurable. The δ -weak worst-case condition number and the δ -weak stochastic condition number are defined as

$$\kappa_w(\delta) := \inf\{y \in \mathbb{R} : \mathbb{P}\{\sigma_E < y\} \geq 1 - \delta\}, \quad \kappa_{ws}(\delta) := \mathbb{E}[\sigma_E \mid \sigma_E \leq \kappa_w(\delta)].$$

Remark 2.6 We note that one can give a definition of the weak worst-case and weak stochastic condition number that does not require σ_E to be a random variable, by setting

$$\kappa_w(\delta) = \inf_{\substack{\mathcal{S} \in \Sigma, \\ |\mathcal{S}| \geq 1-\delta}} \sup_{\omega \in \mathcal{S}} \sigma_{E(\omega)}, \quad \kappa_{ws}(\delta) = \inf_{\substack{\mathcal{S} \in \Sigma, \\ |\mathcal{S}| \geq 1-\delta}} \mathbb{E}[\sigma_E \mid \mathcal{S}],$$

where we used the notation $|\mathcal{S}| = \mathbb{P}(\mathcal{S})$ for the measure of a set if there is no ambiguity. This form is reminiscent of the definition of weak average-case analysis in [3], and when σ_E is a random variable, it can be shown to be equivalent to 2.5. Moreover, this slightly more general definition better illustrates the essence of the weak condition numbers: these are the (worst-case and average-case) condition numbers that ensue when one is allowed to discard a “numerically invisible” subset from the set of perturbations.

The directional sensitivity has an interpretation as (the limit of) a ratio of forward and backward errors, and hence, the new approach provides a potentially useful general framework to give probabilistic bounds on the forward accuracy of outputs of numerically stable algorithms. Moreover, as we discussed in Sect. 6, upper bounds on the weak condition numbers can be computed in practice for a natural distribution. One can therefore see δ as a parameter representing the confidence level that a user wants for the output, and any computable upper bound on κ_w becomes a practical reliability measure on the output, valid with probability $1 - \delta$. Although of course roundoff errors are not really random variables, we hope that modelling them as such can become, with this “weak theory”, a useful tool for numerical analysis problems whose traditional condition number is infinite.

3 Eigenvalues of Matrix Polynomials and Their Directional Sensitivity

Algebraically, the spectral theory of matrix polynomials is most naturally described over an algebraically closed field; however, the theory of condition is analytic in nature and it is sometimes of interest to restrict the coefficients, and their perturbations, to be real. In this section, we give a unified treatment of both real and complex matrix polynomials. For conciseness, we keep this overview very brief; interested readers can find further details in [14,18,19,26,37] and the references therein. A square matrix polynomial is a matrix $P(x) \in \mathbb{F}[x]^{n \times n}$, where $\mathbb{F} \in \{\mathbb{C}, \mathbb{R}\}$ is a field. Alternatively, we can think of it as an expression

$$P(x) = P_0 + P_1x + \dots + P_dx^d,$$

with $P_i \in \mathbb{F}^{n \times n}$. If we require $P_d \neq 0$, then the integer d in such an expression is called the *degree* of the matrix polynomial.⁷ We denote the vector space of matrix polynomials over \mathbb{F} of degree at most d by $\mathbb{F}_d^{n \times n}[x]$. A square matrix polynomial is called *singular* if $\det P(x) \equiv 0$ and otherwise regular. An element $\lambda \in \mathbb{C}$ is said to be a finite eigenvalue of $P(x)$ if

$$\text{rank}_{\mathbb{C}}(P(\lambda)) < \text{rank}_{\mathbb{F}(x)}(P(x)) =: r,$$

⁷ By convention, the zero matrix polynomial has degree $-\infty$.

where $\mathbb{F}(x)$ is the field of fractions of $\mathbb{F}[x]$, that is, the field of rational functions with coefficients in \mathbb{F} . We assume throughout rank $r \geq 1$ (which implies $n \geq 1$) and degree $d \geq 1$. The *geometric multiplicity* of the eigenvalue λ is the amount by which the rank decreases in the above definition,

$$g_\lambda = r - \text{rank}_{\mathbb{C}}(P(\lambda)).$$

There exist matrices $U, V \in \mathbb{F}[x]^{n \times n}$ with $\det(U) \in \mathbb{F} \setminus \{0\}$, $\det(V) \in \mathbb{F} \setminus \{0\}$, that transform $P(x)$ into its *Smith canonical form*,

$$U^* P(x) V = D := \text{diag}(h_1(x), \dots, h_r(x), 0, \dots, 0), \tag{3}$$

where the *invariant factors* $h_i(x) \in \mathbb{F}[x]$ are nonzero monic polynomials such that $h_i(x) | h_{i+1}(x)$ for $i \in \{1, \dots, r - 1\}$. If one has the factorizations $h_i = (x - \lambda)^{k_i} \tilde{h}_i(x)$ for some $\tilde{h}_i(x) \in \mathbb{C}[x]$, with $0 \leq k_i \leq k_{i+1}$ for $i \in \{1, \dots, r - 1\}$ and $(x - \lambda)$ not dividing any of the $\tilde{h}_i(x)$, then the k_i are called the *partial multiplicities* of the eigenvalue λ . The *algebraic multiplicity* a_λ is the sum of the partial multiplicities. Note that an immediate consequence of this definition is $a_\lambda \geq g_\lambda$. If $a_\lambda = g_\lambda$ (i.e. all nonzero k_i equal to 1), then the eigenvalue λ is said to be *semisimple*; otherwise, it is *defective*. If $a_\lambda = 1$ (i.e. $k_i = 1$ for $i = r$ and zero otherwise), then we say that λ is *simple*; otherwise, it is *multiple*.

A square matrix polynomial is *regular* if $r = n$, i.e. if $\det P(x)$ is not identically zero. A finite eigenvalue of a regular matrix polynomial is simply a root of the characteristic equation $\det P(x) = 0$, and its algebraic multiplicity is equal to the multiplicity of the corresponding root. If a matrix polynomial is not regular it is said to be *singular*. More generally, a finite eigenvalue of a matrix polynomial (resp. its algebraic multiplicity) is a root (resp. the multiplicity as a root) of the equation $\gamma_r(x) = 0$, where $\gamma_r(x)$ is the monic greatest common divisor of all the minors of $P(x)$ of order r (note that $\gamma_n(x) = \det P(x)$).

Remark 3.1 The concept of an eigenvalue, and the other definitions recalled here, is valid also in the more general setting of rectangular matrix polynomials. However, in that scenario a generic matrix polynomial has no eigenvalues [18]; as a consequence, a perturbation of a matrix polynomial with an eigenvalue would almost surely remove it. This is a fairly different setting than in the square case, and a deeper probabilistic analysis of the rectangular case is beyond the scope of the present paper.

We mention in passing that there are possible ways to extend the analysis to the rectangular case, such as embedding them in a larger square matrix polynomial or (at least in the case of pencils, or linear matrix polynomials) consider structured perturbations that do preserve eigenvalues.

3.1 Eigenvectors

To define the eigenvectors, let $\{b_1(x), \dots, b_{n-r}(x)\}$ and $\{c_1(x), \dots, c_{n-r}(x)\}$ be minimal bases [19,23,37] of $\ker P(x)$ and $\ker P(x)^*$ (as vector spaces over $\mathbb{F}(x)$), respectively. For $\lambda \in \mathbb{C}$, it is not hard to see [19,37] that $\ker_\lambda P(x) :=$

$\text{span}\{b_1(\lambda), \dots, b_{n-r}(\lambda)\}$ and $\ker_\lambda P(x)^* := \text{span}\{c_1(\lambda^*), \dots, c_{n-r}(\lambda^*)\}$ are vector spaces over \mathbb{C} of dimension $n - r$.

Note that $\ker_\lambda P(x) \subseteq \ker P(\lambda)$ and $\ker_\lambda P(x)^* \subseteq \ker P(\lambda)^*$ for $\lambda \in \mathbb{C}$, and that the difference in dimension is the geometric multiplicity, $\ker P(\lambda) - \ker_\lambda P(x) = \ker P(\lambda)^* - \ker_\lambda P(x)^* = g_\lambda$. A *right eigenvector* corresponding to an eigenvalue $\lambda \in \mathbb{C}$ is defined [19, Sec. 2.3] to be a nonzero element of the quotient space $\ker P(\lambda)/\ker_\lambda P(x)$. A *left eigenvector* is similarly defined as an element of $\ker P(\lambda)^*/\ker_\lambda P(x)^*$. In terms of the Smith canonical form (3), the last $n - r$ columns of U , evaluated at λ^* , represent a basis of $\ker_\lambda P(x)^*$, while the last $(n - r)$ columns of V , evaluated at λ , represent a basis of $\ker_\lambda P(x)$.

In the analysis, we will be concerned with a quantity of the form $|u^*P'(\lambda)v|$, where u, v are representatives of eigenvectors. It is known [19, Lemma 2.9] that $b \in \ker_\lambda P(x)$ is equivalent to the existence of a polynomial vector $b(x)$ such that $b(\lambda) = b$ and $P(x)b(x) = 0$. Then,

$$0 = \frac{d}{dx} P(x)b(x)|_{x=\lambda} = P'(\lambda)b(\lambda) + P(\lambda)b'(\lambda)$$

implies that for any representative of a left eigenvector $u \in \ker P(\lambda)^*$ we get $u^*P'(\lambda)b(\lambda) = 0$. It follows that for an eigenvalue representative v , $u^*P'(\lambda)v$ depends only the component of v orthogonal to $\ker_\lambda P(x)$, and an analogous argument also shows that this expression only depends on the component of u orthogonal to $\ker_\lambda P(x)^*$. In practice, we will therefore choose representatives u and v for the left and right eigenvalues that are orthogonal to $\ker_\lambda P(x)^*$ and $\ker_\lambda P(x)$, respectively, and have unit norm. If $P(x) \in \mathbb{F}_d^{n \times n}[x]$ is a matrix polynomial with simple eigenvalue λ , then there is a unique (up to sign) way of choosing such representatives u and v .

3.2 Perturbations of Singular Matrix Polynomials: The De Terán–Dopico Formula

Assume that $P(x) \in \mathbb{F}_d^{n \times n}[x]$, where $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, is a matrix polynomial of rank $r \leq n$, and let λ be a simple eigenvalue. Let $X = [U \ u] \in \mathbb{C}^{n \times (n-r+1)}$ be a matrix whose columns form a basis of $\ker P(\lambda)^*$, and such that the columns of $U \in \mathbb{C}^{n \times (n-r)}$ form a basis of $\ker_\lambda P(x)^*$. Likewise, let $Y = [V \ v] \in \mathbb{C}^{n \times (n-r)}$ be a matrix whose columns form a basis of $\ker P(\lambda)$, such that the columns of $V \in \mathbb{C}^{n \times (n-r)}$ form a basis of $\ker_\lambda P(x)$. In particular, v and u are representatives of, respectively, right and left eigenvectors of $P(x)$. The following explicit characterization of a simple eigenvalue is due to De Terán and Dopico [14, Theorem 2 and Eqn. (20)]. To avoid making a case distinction for the regular case $r = n$, we agree that $\det(U^*E(\lambda)V) = 1$ if U and V are empty.

Theorem 3.2 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be matrix polynomial of rank r with simple eigenvalue λ and X, Y as above. Let $E(x) \in \mathbb{F}_d^{n \times n}[x]$ be such that $X^*E(\lambda)Y$ is non-singular. Then, for small enough $\epsilon > 0$, the perturbed matrix polynomial $P(x) + \epsilon E(x)$ has exactly one eigenvalue $\lambda(\epsilon)$ of the form*

$$\lambda(\epsilon) = \lambda - \frac{\det(X^*E(\lambda)Y)}{u^*P'(\lambda)v \cdot \det(U^*E(\lambda)V)}\epsilon + O(\epsilon^2).$$

Note that in the special case $r = n$ we recover the expression for regular matrix polynomials from [42, Theorem 5] and [14, Corollary 1],

$$\lambda(\epsilon) = \lambda - \frac{u^* E(\lambda)v}{u^* P'(\lambda)v} \epsilon + O(\epsilon^2), \tag{4}$$

where u, v are left and right eigenvectors corresponding to the eigenvalue λ .

3.3 The Directional Sensitivity of a Singular Polynomial Eigenproblem

We can now describe the input–output map that underlies our analysis. By the local nature of our problem, we consider a fixed matrix polynomial $P(x) \in \mathbb{F}_d^{n \times n}[x]$ of rank r with simple eigenvalue λ and define the input–output function

$$f : \mathbb{F}_d^{n \times n}[x] \rightarrow \mathbb{C}$$

that maps $P(x)$ to λ , maps $P(x) + \epsilon E(x)$ to $\lambda(\epsilon)$ for any $E(x)$ and $\epsilon > 0$ satisfying the conditions of Theorem 3.2, and maps any other matrix polynomial to an arbitrary number other than λ .

An immediate consequence of Theorem 3.2 and our definition of the input–output map is an explicit expression for the directional sensitivity of the problem. Here we write $\|E\|$ for the Euclidean norm of the vector of coefficients of $E(x)$ as a vector in $\mathbb{F}^{n^2(d+1)}$. From now on, when talking about the “directional sensitivity of an eigenvalue in direction E ”, we implicitly refer to the input–output map f defined above.

Corollary 3.3 *Let λ be a simple eigenvalue of $P(x)$ and let $E(x) \in \mathbb{F}_d^{n \times n}[x]$ be a regular matrix polynomial. Then, the directional sensitivity of the eigenvalue λ in direction $E(x)$ is*

$$\sigma_E = \frac{1}{\|E\|} \left| \frac{\det(X^* E(\lambda) Y)}{u^* P'(\lambda)v \cdot \det(U^* E(\lambda) V)} \right|.$$

In the special case $r = n$, we have

$$\sigma_E = \frac{1}{\|E\|} \left| \frac{u^* E(\lambda)v}{u^* P'(\lambda)v} \right|.$$

For the goals in this paper, these results suffice. However, we note that it is possible to obtain equivalent formulae for the expansion that, unlike the one by De Terán and Dopico, do not involve the eigenvectors of singular polynomials.

Finally, we introduce a parameter that will enter all of our results and coincides with the inverse of the worst-case condition number in the regular case $r = n$. Choose representatives u, v of the eigenvectors that satisfy $\|u\| = \|v\| = 1$ and (if $r < n$) $U^* u = V^* v = 0$. For such a choice of eigenvectors, define

$$\gamma_P := |u^* P'(\lambda)v| \cdot \left(\sum_{j=0}^d |\lambda|^{2j} \right)^{-1/2}. \tag{5}$$

We conclude with the following variation in [42, Theorem 5]. For a proof of the following result, see [1, Lemma 2.1] or [2] for a discussion in a wider context.

Proposition 3.4 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a regular matrix polynomial and $\lambda \in \mathbb{C}$ a simple eigenvalue. Then, the worst-case condition number of the problem of computing λ is $\kappa = \gamma_P^{-1}$.*

Remark 3.5 In practice, an algorithm such as QZ applied to $P(x)$ will typically compute all the eigenvalues of a nearby matrix polynomial. Therefore, any conditioning results on the conditioning of our specific input–output map f will explain why the correct eigenvalue is found among the computed eigenvalues, but not tell us how to choose the right one in practice. For selecting the right eigenvalue, one could use heuristics, such as computing the eigenvalues of an artificially perturbed problem. For more details on these practical considerations, we refer to [32].

4 Probabilistic Analysis of the Directional Sensitivity

In this section, we study the probability distribution of the directional sensitivity of a singular polynomial eigenvalue problem. To deal with real and complex perturbations simultaneously as far as possible, we follow the convention from random matrix theory [22] and parametrize our results with a parameter β , where $\beta = 1$ if $\mathbb{F} = \mathbb{R}$ and $\beta = 2$ if $\mathbb{F} = \mathbb{C}$. We consider perturbations $E(x) = E_0 + E_1x + \dots + E_dx^d$, which we identify with the matrix $E = [E_0 \dots E_d] \in \mathbb{F}^{n \times n(d+1)}$ (each $E_i \in \mathbb{F}^{n \times n}$), and denote by $\|E\|$ the Euclidean norm of E considered as a vector in \mathbb{F}^N , where $N := n^2(d + 1)$ (equivalently, the Frobenius norm of the matrix E). When we say that E is uniformly distributed on the sphere, written $E \sim \mathcal{U}(\beta N)$ with $\beta = 1$ for real perturbations and $\beta = 2$ if E is complex, we mean that the image of E under an identification $\mathbb{F}^{n \times n(d+1)} \cong \mathbb{R}^{\beta N}$ is uniformly distributed on the corresponding unit sphere $S^{\beta N - 1}$. To avoid trivial special cases, we assume that $r \geq 1$ and $d \geq 1$, so that, in particular, $N \geq 2$.

The following theorem characterizes the distribution of the directional sensitivity under uniform perturbations.

Theorem 4.1 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a matrix polynomial of rank r and let λ be a simple eigenvalue of $P(x)$. If $E \sim \mathcal{U}(\beta N)$, where $\beta = 1$ if $\mathbb{F} = \mathbb{R}$ and $\beta = 2$ if $\mathbb{F} = \mathbb{C}$, then the directional sensitivity of λ in direction $E(x)$ satisfies*

$$\mathbb{P}\{\sigma_E \geq t\} = \begin{cases} \mathbb{P}\{Z_N/Z_{n-r+1} \geq \gamma_P^2 t^2\} & \text{if } r < n \\ \mathbb{P}\{Z_N \geq \gamma_P^2 t^2\} & \text{if } r = n \end{cases},$$

where $Z_k \sim \mathcal{B}(\beta/2, \beta(k - 1)/2)$ denotes a beta distributed random variable with parameters $\beta/2$ and $\beta(k - 1)/2$, and Z_N and Z_{n-r+1} are independent.

The proof is given later in this section, after having introduced some preliminary concepts and results. If $r = n$, then the directional sensitivity is distributed like the square root of a beta random variable, and in particular, it is bounded. Using the density

of the beta distribution, we can derive the moments and tail bounds for the distribution of the directional sensitivity explicitly.

Corollary 4.2 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a matrix polynomial of rank r and let λ be a simple eigenvalue of $P(x)$. If $E \sim \mathcal{U}(\beta N)$, where $\beta = 1$ if $\mathbb{F} = \mathbb{R}$ and $\beta = 2$ if $\mathbb{F} = \mathbb{C}$, then the expected directional sensitivity of λ in direction $E(x)$ is*

$$\mathbb{E}[\sigma_E] = \begin{cases} \frac{1}{\gamma_P} \frac{\pi}{2} \frac{\Gamma(N)\Gamma(n-r+1)}{\Gamma(N+1/2)\Gamma(n-r+1/2)} & \text{if } \mathbb{F} = \mathbb{C} \\ \infty & \text{if } \mathbb{F} = \mathbb{R} \text{ and } r < n \\ \frac{1}{\gamma_P} \frac{\Gamma(N/2)}{\sqrt{\pi}\Gamma((N+1)/2)} & \text{if } \mathbb{F} = \mathbb{R} \text{ and } r = n. \end{cases}$$

If $t \geq \gamma_P^{-1}$, then for $r < n$ we have the tail bounds

$$\mathbb{P}\{\sigma_E \geq t\} \leq \begin{cases} \frac{1}{\gamma_P^2} \frac{n-r}{N} \frac{1}{t^2} & \text{if } \mathbb{F} = \mathbb{C} \\ \frac{1}{\gamma_P} \frac{2}{\pi} \frac{\Gamma(N/2)\Gamma(n-r+1/2)}{\Gamma((N+1)/2)\Gamma((n-r)/2)} \frac{1}{t} & \text{if } \mathbb{F} = \mathbb{R}. \end{cases} \tag{6}$$

If $r = n$, then $\sigma_E \leq \gamma_P$.

Proof For the expectation, using Theorem 4.1 in the case $r < n$, we have

$$\mathbb{E}[\sigma_E] = \frac{1}{\gamma_P} \int_0^\infty \mathbb{P}\{(Z_N/Z_{n-r+1})^{1/2} \geq t\} dt = \frac{1}{\gamma_P} \mathbb{E}[(X_N/X_{n-r+1})^{1/2}],$$

where X_k denotes a $B(\beta/2, \beta(k-1)/2)$ distributed random variable. The claimed tail bounds and expected values for $r < n$ follow by applying Lemma A.1 with $k = 2$, $a = c = \beta/2$, $b = \beta(N-1)/2$, and $d = \beta(n-r)/2$. If $r = n$, the expected value follows along the lines, and the deterministic bound follows trivially from the boundedness of the beta distribution. \square

Remark 4.3 In the context of random inputs, it is common to study the logarithm of a condition number instead of the condition number itself [13,21]. Thus, even when the expected condition is not finite, the expected logarithm may still be small. Using a standard argument (see e.g. [13, Proposition 2.26]), we can deduce a bound on the expected logarithm of the directional sensitivity:

$$\mathbb{E}[\log \sigma_E] \leq \log \left(\frac{1}{\gamma_P} \frac{2}{\pi} \sqrt{\frac{n-r}{N-1}} \right) + 1.$$

The logarithm of the sensitivity is relevant as a measure for the loss of precision.

As the derivation of the bounds (6) using Lemma A.1 shows, the cumulative distribution functions in question can be expressed *exactly* in terms of integrals of hypergeometric functions. This way, the tail probabilities can be computed to high accuracy for any given t , see also Remark 4.6. However, as the derivation of the tail bounds in ‘‘Appendix A’’ also shows, the bounds given in Corollary 4.2 are sharp for




Fig. 1 Exact distribution tail of σ_E and the tail bounds for $n = 4, d = 2, r = 2,$ and $\gamma_P = 1,$ so that $N = 48$

fixed t and $n - r \rightarrow \infty,$ as well as for fixed $n - r$ and $t \rightarrow \infty.$ Figure 1 illustrates these bounds for a choice of small parameters ($n = 4, d = 2, r = 2, \gamma_P = 1$). Moreover, the bounds (6) have the added benefit of being easily interpretable. These tail bounds can be interpreted as saying that for large n and/or $d,$ it is highly unlikely that the directional sensitivity will exceed γ_P^{-1} (which by Proposition 3.4 is the worst-case condition bound in the smooth case $r = n$).

Example 4.4 Consider again the matrix pencil $L(x)$ from (1). This pencil has rank 3, and the cokernel and kernel are spanned by the vectors $p(x)$ and $q(x),$ respectively, given by

$$p(x) = \begin{bmatrix} 1 \\ 0 \\ -1 \\ 1 \end{bmatrix}, \quad q(x) = \begin{bmatrix} x \\ -2x^2 - 4x + 1 \\ x \\ 1 \end{bmatrix}.$$

The matrix polynomial has the simple eigenvalue $\lambda = 1,$ and the matrix $L(1)$ has rank 2. The cokernel $\ker L(1)^T$ and the kernel $\ker L(1)$ are spanned by the columns of the matrices X and $Y,$ given by

$$X \approx \begin{bmatrix} 0.5774 & -0.7061 \\ 0 & 0.4888 \\ -0.5774 & -0.4345 \\ 0.5774 & 0.2716 \end{bmatrix}, \quad Y \approx \begin{bmatrix} 0.1924 & -0.6873 \\ -0.9623 & -0.1322 \\ 0.1924 & 0.02644 \\ 0 & 0.7137 \end{bmatrix}.$$

Let u be the second column of X and let v be the second column of $Y.$ The vectors u and v are orthogonal to $\ker_\lambda L(x)^T = \text{span}\{p(1)\}$ and $\ker_\lambda L(x) = \text{span}\{q(1)\}$ and have unit norm. We therefore have

$$\gamma_L := \frac{|u^T L'(1)v|}{\sqrt{2}} = 0.08223$$




Fig. 2 Exact distribution tail of σ_E for the matrix pencil $L(x)$ from (1) and the theoretically computed tail bound (6)

Hence, $\gamma_L^{-1} = 12.16$. Figure 2 shows the result of comparing the distribution of σ_E , found empirically, with the bounds obtained in Theorem 4.1. The relative error in the plot is of order 10^{-5} .

The plan for the rest of this section is as follows: In Sect. 4.1, we recall some facts from probability theory and random matrix theory. In Sect. 4.2, we discuss the QR decomposition of a random matrix, and in Sect. 4.3 we use this decomposition to prove Theorem 4.1

4.1 Probabilistic Preliminaries

We write $g \sim \mathcal{N}^1(\mu, \Sigma)$ for a normal distributed (Gaussian) random vector g with mean μ and covariance matrix Σ , and $g \sim \mathcal{N}^2(\mu, \Sigma)$ for a complex Gaussian vector; this is a \mathbb{C}^n -valued random vector with expected value μ , whose real and imaginary parts are independent real Gaussian random vectors with covariance matrix $\Sigma/2$ (a special case are real and complex scalar random variables, $\mathcal{N}^\beta(\mu, \sigma^2)$). We denote the uniform distribution on a sphere S^{n-1} by $\mathcal{U}(n)$. Every Gaussian vector $g \sim \mathcal{N}^1(0, I_n)$ can be written as a product $g = rq$ with r and q independent, where $r \sim \chi(n)$ is χ -distributed with n degrees of freedom, and $q \sim \mathcal{U}(n)$.

4.1.1 Projections of Random Vectors

The squared projected lengths of Gaussian and uniform distributed random vectors can be described using the χ^2 and the beta distribution, respectively. A vector X is χ^2 -distributed with k degrees of freedom, $X \sim \chi^2(k)$, if the cumulative distribution function (cdf) is

$$\mathbb{P}\{X \leq x\} = \frac{1}{2^{k/2}\Gamma(k/2)} \int_0^x t^{k/2-1} e^{-t/2} dt.$$

The special case $\chi^2(2)$ is the exponential distribution with parameter $1/2$, written $\exp(1/2)$. The beta distribution $B(a, b)$ is defined for $a, b > -1$ and has cdf supported on $[0, 1]$,

$$\mathbb{P}\{X \leq x\} = \frac{1}{B(a, b)} \int_0^x t^{a-1} (1-t)^{b-1} dt,$$

where $B(a, b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ is the beta function. For a vector $x \in \mathbb{F}^n$, denote by $\pi_k(x)$ the projection onto the first k coordinates and by $\|\pi_k(x)\|^2 = |x_1|^2 + \dots + |x_k|^2$ its squared length. The following facts are known:

- If $g \sim \mathcal{N}^\beta(0, I_n)$, then $\beta\|\pi_k(g)\|^2 \sim \chi^2(\beta k)$;
- If $q \sim \mathcal{U}(n)$, then $\|\pi_k(q)\|^2 \sim B(k/2, (n-k)/2)$.

The first claim is a standard fact about the normal distribution and can be derived directly from it, see, for example, [9]. The statement for the uniform distribution can be derived from the Gaussian one, but also follows by a change of variables from expressions for the volume of tubular neighbourhoods of subspheres of a sphere, see for example [13, Section 20.2]. Since all the distributions considered are orthogonally (in the real case) or unitarily (in the complex case) invariant, these observations hold for the projection of a random vector onto *any* k -dimensional subspace, not just the first k coordinates.

4.1.2 Random Matrix Ensembles

If $P(x)$ is a singular matrix polynomial with a simple eigenvalue λ , then the set of perturbation directions for which the directional sensitivity is not finite is a proper Zariski closed subset, see Theorem 3.2. It is therefore natural and convenient to consider probability measures on the space of perturbations that have measure zero on proper Zariski closed subsets. This is the case, for example, if the measure is absolutely continuous with respect to the Lebesgue measure. In this paper, we will work with real and complex Gaussian and uniform distributions. For a detailed discussion of the random matrix ensembles used here, we refer to [24, Chapters 1–2].

For a random matrix, we write $G \sim \mathcal{G}_n^\beta(\mu, \sigma^2)$ if each entry of G is an independent $\mathcal{N}^\beta(\mu, \sigma^2)$ random variable, and call this a Gaussian random matrix. In the case $\beta = 2$, this is called the *Ginibre ensemble* [25]. Centred ($\mu = 0$) Gaussian random matrices are orthogonally (if $\beta = 1$) or unitarily (if $\beta = 2$) invariant ([35, Lemma 1]), and the joint density of their entries is given by

$$\frac{1}{(2\pi/\beta)^{\beta n^2/2}} e^{-\frac{\beta\|G\|^2}{2}},$$

which takes into account the fact the real and imaginary parts of the entries of a complex Gaussian have variance $1/2$. In addition, we consider the *circular real ensemble* $CRE(n)$ for real orthogonal matrices in $O(n)$ and the *circular unitary ensemble* $CUE(n)$ [20] for unitary matrices in $U(n)$, where both distributions correspond to the unique Haar probability measure on the corresponding groups.

4.2 The Probabilistic QR Decomposition

Any non-singular matrix $A \in \mathbb{F}^{n \times n}$ has a unique QR decomposition $A = QR$, where $Q \in O(n)$ (if $\mathbb{F} = \mathbb{R}$) or $U(n)$ (if $\mathbb{F} = \mathbb{C}$), and $R \in \mathbb{F}^{n \times n}$ is upper triangular with $r_{ii} > 0$ [44, Part II]. The following proposition describes the distribution of the factors Q and R in the QR decomposition of a (real or complex) Gaussian random matrix.

Proposition 4.5 *Let $G \sim \mathcal{G}_n^\beta(0, 1)$ be a Gaussian random matrix, $\beta \in \{1, 2\}$. Then G can be factored uniquely as $G = QR$, where $R = (r_{jk})_{1 \leq j \leq k \leq n}$ is upper triangular and*

- $Q \sim \text{CUE}(n)$ if $\beta = 2$ and $Q \sim \text{CRE}(n)$ if $\beta = 1$;
- $\beta r_{ii}^2 \sim \chi^2(\beta(n - i + 1))$ for $i \in \{1, \dots, n\}$;
- $r_{jk} \sim \mathcal{N}^\beta(0, 1)$ for $1 \leq j < k \leq n$.

Moreover, all these random variables are independent.

An easy and conceptual derivation of the distribution of Q can be found in [35], while the distribution of R can be deduced from the known expression for the Jacobian of the QR decomposition [22, 3.3].

4.3 Proof of Theorem 4.1

In this section, we present the proofs of Theorem 4.1 and the corollaries that follow from it. To simplify notation, we set $\ell = n - r + 1$. Recall from Corollary 3.3 the expression

$$\sigma_E = \frac{1}{\|E\|} \left| \frac{\det(X^*E(\lambda)Y)}{u^*P'(\lambda)v \cdot \det(U^*E(\lambda)V)} \right|,$$

where the columns of $X = [U \ u], Y = [V \ v] \in \mathbb{F}^{n \times \ell}$ are orthonormal bases of $\ker P(\lambda)^*$ and $\ker P(\lambda)$, the columns of U, V represent bases of $\ker_\lambda P(x)^*$ and $\ker_\lambda P(x)$, respectively, and γ_P is defined in (5).

Proof of Theorem 4.1 We first assume $r < n$. By the scale invariance of the directional sensitivity σ_E , we consider Gaussian perturbations $E \sim \mathcal{N}^\beta(0, \sigma^2 I_{\beta N})$ (recall that we interpret E as a vector in \mathbb{F}^N), where $\sigma^2 = (\sum_{j=0}^d |\lambda|^{2j})^{-1}$. This scaling ensures that the entries of $E(\lambda)$ are independent $\mathcal{N}^\beta(0, 1)$ random variables. Since the distribution of $E(\lambda)$ is orthogonally/unitarily invariant, the quotient $|\det(X^*E(\lambda)Y)|/|\det(U^*E(\lambda)V)|$ has the same distribution as the quotient $|\det(G)/\det(\overline{G})|$, where G is the upper left $\ell \times \ell$ submatrix of $E(\lambda)$ and \overline{G} the upper left $(\ell - 1) \times (\ell - 1)$ matrix. For the distribution considered, G is almost surely invertible, with inverse $H = G^{-1}$. By Cramer’s rule, $|\det(G)/\det(\overline{G})| = |h_{\ell\ell}|^{-1}$. We are thus interested in the distribution

$$\mathbb{P}\{\sigma_E \geq t\} = \mathbb{P}\left\{ \left| \frac{\det(X^*E(\lambda)Y)}{\det(U^*E(\lambda)V)} \right| \geq |u^*P'(\lambda)v|t\|E\| \right\} = \mathbb{P}\left\{ \frac{1}{|h_{\ell\ell}|} \geq |u^*P'(\lambda)v|t\|E\| \right\},$$

where $h_{\ell\ell}$ is the lower right corner of the inverse of an $\ell \times \ell$ Gaussian matrix G . To study the distribution of $|h_{\ell\ell}|^{-1}$, we resort to the probabilistic QR decomposition discussed in Sect. 4.2. If $G = QR$ is the unique QR decomposition of G with positive diagonals in R , then the inverse is given by $H = R^{-1}Q^*$, and a direct inspection reveals that the lower right element $h_{\ell\ell}$ of H is $h_{\ell\ell} = q_{\ell\ell}^*/r_{\ell\ell}$.

From Sect. 4.2, it follows that $Q \sim \text{CRE}(n)$ or $\text{CUE}(n)$, and $\beta r_{\ell\ell}^2 \sim \chi^2(\beta)$. Moreover, each column of Q is uniformly distributed on the sphere $S^{\beta\ell-1}$, so that $|q_{\ell\ell}|^2 \sim \text{B}(\beta/2, \beta(\ell - 1)/2)$ (by Sect. 4.1.1), and $\{r_{\ell\ell}^2, |q_{\ell\ell}|^2\}$ are independent. We therefore get

$$\mathbb{P} \left\{ \frac{1}{|h_{\ell\ell}|} \geq |u^* P'(\lambda)v|t \|E\| \right\} = \mathbb{P} \left\{ r_{\ell\ell}^2 \geq |q_{\ell\ell}|^2 |u^* P'(\lambda)v|^2 t^2 \|E\|^2 \right\}.$$

Setting $\gamma_P = |u^* P'(\lambda)v| \cdot (\sum_{j=0}^d |\lambda|^{2j})^{-1/2}$ (see 5), we arrive at

$$\mathbb{P}\{\sigma_E \geq t\} = \mathbb{P} \left\{ r_{\ell\ell}^2 \geq |q_{\ell\ell}|^2 \gamma_P^2 t^2 \left(\sum_{j=0}^d |\lambda|^{2j} \right) \cdot \|E\|^2 \right\}$$

Let $p_0 = (1, \lambda, \dots, \lambda^d)^T / (\sum_{i=0}^d |\lambda|^{2i})^{1/2}$. Then we can rearrange the coefficients of $E(x)$ to a matrix $F \in \mathbb{F}^{n^2 \times (d+1)}$ so that

$$\|Fp_0\|^2 = \|E(\lambda)\|^2 \cdot \left(\sum_{i=0}^d |\lambda|^{2i} \right)^{-1}.$$

Moreover, if $Q = [p_0 \ p_1 \ \dots \ p_d]$ is an orthogonal/unitary matrix with p_0 as the first column, then

$$\|E\|^2 = \|F\|^2 = \|FQ\|^2 = \left(\sum_{i=0}^d |\lambda|^{2i} \right)^{-1} \left(\|E(\lambda)\|^2 + \sum_{j=1}^d \|F \tilde{p}_j\|^2 \right),$$

where $\tilde{p}_j = p_j \cdot (\sum_{i=0}^d |\lambda|^{2i})^{1/2}$. If we denote by G^c the vector consisting of those entries of $E(\lambda)$ that are not in G , then

$$\|E(\lambda)\|^2 = \|G\|^2 + \|G^c\|^2 = \|R\|^2 + \|G^c\|^2.$$

It follows that

$$\left(\sum_{j=0}^d |\lambda|^{2j} \right) \cdot \|E\|^2 = \|R\|^2 + \|G^c\|^2 + \sum_{j=1}^d \|F \tilde{p}_j\|^2.$$

Therefore, the factor $r_{\ell\ell}^2$, itself a square of a (real or complex) Gaussian, is a summand in a sum of squares of $N = n^2(d + 1)$ Gaussians, and the quotient

$$\frac{r_{\ell\ell}^2}{\left(\sum_{j=0}^d |\lambda|^{2j}\right) \|E\|^2}$$

is equal to the squared length of the projection of a uniform random vector in $S^{\beta N-1}$ onto the first β coordinates. By Sect. 4.1.1, this is $B(\beta/2, \beta(N - 1)/2)$ distributed. Denoting this random variable by Z_N and $|q_{\ell\ell}|^2$ by Z_ℓ , we obtain

$$\mathbb{P}\{\sigma_E \geq t\} = \mathbb{P}\{Z_N \geq \gamma_P^2 t^2 Z_\ell\}.$$

This establishes the claim in the case $r < n$. If $r = n$, we use the expression (see 4),

$$\sigma_E = \left| \frac{u^* E(\lambda)v}{\|E\| \|u^* P'(\lambda)v\|} \right|,$$

where u and v are eigenvectors. By orthogonal/unitary invariance, σ_E^2 has the same distribution as the squared norm of a Gaussian. By the same argument as above, we can bound $\|E\|$ in terms of $\|E(\lambda)\|$, and the quotient with $\|E(\lambda)\|^2$ is then the squared projected length of the first β coordinates of a uniform distributed vector in $S^{\beta N-1}$, which is $B(\beta/2, \beta(N - 1)/2)$ distributed. \square

Remark 4.6 If $N+$ is large, then for a (real or complex) Gaussian perturbation with entry-wise variance $1/N$, by Gaussian concentration (see [11, Theorem 5.6]), $\|E\|$ is close to 1 with high probability:

$$\mathbb{P}\{|\|E\| - 1| \geq t\} \leq 2e^{-Nt^2/2}.$$

This means that the distribution of $\|E\|\sigma_E$ for a Gaussian perturbation will be close to that of σ_E for a uniform perturbation. Even for moderate sizes of d and n , the result can be numerically almost indistinguishable.

In fact, when G is Gaussian, then the distribution can be expressed explicitly as

$$\mathbb{P}\{|h_{\ell\ell}|^{-1} \geq t\} = {}_1F_1(1, \ell; -Nt^2),$$

where ${}_1F_1(a, b; z)$ denotes the confluent hypergeometric function (this follows by mimicking the proof of Theorem 4.1, expressing the distribution in terms of a quotient of a χ^2 and a beta random variable, and writing out the resulting integrals). Similarly, using the same computations as in the proof of Lemma A.1, we get the exact expression

$$\mathbb{P}\{|h_{\ell\ell}|^{-1} \geq t\} = \begin{cases} 1 - {}_2F_1(1 - N, 1, \ell; t^2) & \text{if } t \leq 1 \\ {}_2F_1(1 - \ell, 1, N; t^{-2}) & \text{if } t \geq 1, \end{cases}$$

where ${}_2F_1(a, b, c; z)$ is the hypergeometric function. The case distinction corresponds to different branches of the solution of the hypergeometric differential equation. See [38,39] for more on computing with hypergeometric functions.

5 Weak Condition Numbers of Simple Eigenvalues of Singular Matrix Polynomials

The tail bounds on the directional sensitivity can easily be translated into statements about condition numbers and discuss some consequences and interpretations.

Theorem 5.1 *Let $P(x) \in \mathbb{C}_d^{n \times n}[x]$ be a matrix polynomial of rank r , and let λ be a simple eigenvalue of $P(x)$. Then*

– *the worst-case condition number is*

$$\kappa = \begin{cases} \infty & \text{if } r < n, \\ \frac{1}{\gamma_P} & \text{if } r = n \end{cases} \tag{7}$$

– *the stochastic condition number, with respect to uniformly distributed perturbations, is*

$$\kappa_s = \frac{1}{\gamma_P} \frac{\pi}{2} \frac{\Gamma(N)\Gamma(n-r+1)}{\Gamma(N+1/2)\Gamma(n-r+1/2)} \tag{8}$$

– *if $r < n$ and $\delta \in (0, 1)$, then the δ -weak worst-case condition number, with respect to uniformly distributed perturbations, is bounded by*

$$\kappa_w(\delta) \leq \frac{1}{\gamma_P} \max \left\{ 1, \sqrt{\frac{n-r}{\delta N}} \right\} \tag{9}$$

The expression for the stochastic condition number involves the quotient of gamma functions, which can be simplified using the well-known bounds

$$\sqrt{x} \leq \frac{\Gamma(x+1)}{\Gamma(x+1/2)} \leq \sqrt{x+1/2}, \tag{10}$$

which hold for $x > 0$ [49]. Using these bounds on the numerator and denominator of (8), we get the more interpretable

$$\kappa_s \leq \frac{1}{\gamma_P} \frac{\pi}{2} \sqrt{\frac{n-r+1/2}{N-1/2}} \leq \frac{1}{\gamma_P} \frac{\pi}{2} \sqrt{\frac{n-r+1}{N}}.$$

The bound on the weak condition number (9) shows that $\kappa_w(1/2)$, which is the median of the same random variable of which κ_s is the expected value, is bounded by $1/\gamma_P$, which is the expression of the worst-case condition number in the regular case $r = n$.

The situation changes dramatically when considering real matrix polynomials with real perturbations, as in this case even the stochastic condition becomes infinite if

the matrix polynomial is singular. In the statement, we denote the resulting condition number with respect to real perturbations by using the superscript \mathbb{R} .

Theorem 5.2 *Let $P(x) \in \mathbb{R}^{n \times n}[x]$ be a real matrix polynomial of rank r , and let $\lambda \in \mathbb{C}$ be a simple eigenvalue of $P(x)$. Then,*

- the worst-case condition number is

$$\kappa^{\mathbb{R}} = \begin{cases} \infty & \text{if } r < n, \\ \frac{1}{\gamma_P} & \text{if } r = n \end{cases} \tag{11}$$

- the stochastic condition number, with respect to uniformly distributed real perturbations, is

$$\kappa_s^{\mathbb{R}} = \begin{cases} \infty & \text{if } r < n, \\ \frac{1}{\gamma_P} \frac{\Gamma(N/2)}{\sqrt{\pi} \Gamma((N+1)/2)} & \text{if } r = n \end{cases} \tag{12}$$

- if $r < n$ and $\delta \in (0, 1)$, then the δ -weak worst-case condition, with respect to uniformly distributed real perturbations, is

$$\kappa_w^{\mathbb{R}}(\delta) \leq \frac{1}{\gamma_P} \max \left\{ 1, \sqrt{\frac{n-r}{N}} \frac{1}{\delta} \right\}$$

- if $r < n$ and $\delta < \sqrt{(n-r)/N}$, then the δ -weak stochastic condition number satisfies

$$\kappa_{ws}^{\mathbb{R}}(\delta) \leq \frac{1}{\gamma_P} \left(\frac{1}{1-\delta} \right) \left(1 + \sqrt{\frac{n-r}{N}} \log \left(\sqrt{\frac{n-r}{N}} \delta^{-1} \right) \right) \tag{13}$$

It is instructive to compare the weak condition numbers in the singular case to the worst-case and stochastic condition number in the regular case. In the regular case ($n = r$), when replacing the worst-case with the stochastic condition we get an improvement by a factor of $\approx N^{-1/2}$, which is consistent with previous work [4] (see also Sect. 2) relating the worst-case to the stochastic condition. We will see in Sect. 6.1 that the expected value in the case $n = r$ captures the typical perturbation behaviour of the problem more accurately than the worst-case bound. Among many possible interpretations of the weak worst-case condition, we highlight the following:

- Since the bounds are monotonically decreasing as the rank r increases, we can get bounds independent of r . Specifically, we can replace the quotient $\sqrt{(n-r)/N}$ with $1/\sqrt{n(d+1)}$. This is useful since, in applications, the rank is not always known.
- While the stochastic condition number (12), which measures the *expected* sensitivity of the problem of computing a singular eigenvalue, is infinite, for $4(n-r) < N$ the *median* sensitivity is bounded by

$$\kappa_w^{\mathbb{R}}(1/2) \leq \frac{1}{\gamma_P}.$$

The median is a more robust and arguably better summary parameter than the expectation.

- Choosing $\delta = e^{-N}$ in (13), we get a weak stochastic condition bound of

$$\kappa_{ws}^{\mathbb{R}}(e^{-N}) \leq \frac{1}{\gamma_P} \left(\frac{1 + \sqrt{N(n-r)}}{1 - e^{-N}} \right).$$

That is, the condition number improves from being unbounded to sublinear in N , by just removing a set of inputs of exponentially small measure.

Example 5.3 Consider the matrix pencil $L(x)$ from (1). This matrix pencil has rank 3, with only one simple eigenvalue $\lambda = 1$. As we will see in Example 4.4, the constant γ_L appearing in the bounds is

$$\gamma_L^{-1} = 12.16.$$

In this example, $n = 4, d = 1$ and $r = 3$, so that $n - r = 1, N = n^2(d + 1) = 32$, and

$$\sqrt{\frac{n-r}{N}} = 0.1767.$$

For small enough δ , we get the (not optimized) bound $\kappa_w^{\mathbb{R}}(\delta) < 2.15 \cdot \delta^{-1}$.

It is easy to translate Corollary 4.2 into the main results, Theorem 5.1 and Theorem 5.2. For the weak stochastic condition, we need the following observation, which is a variation in [3, Lemma 2.2].

Lemma 5.4 *Let Z be a random variable such that $\mathbb{P}\{Z \geq t\} \leq C \frac{a}{t}$ for $t > a$. Then for any $t_0 > a$,*

$$\mathbb{E}[Z \mid Z \leq t_0] \leq \frac{a}{1 - \frac{Ca}{t_0}} \left(1 - C \log \left(\frac{a}{t_0} \right) \right).$$

Proof of Theorems 5.1 and 5.2 The statements about the worst-case, (7) and (11), and about the stochastic condition number, (8) and (12), follow immediately from Theorem 4.1 and Corollary 4.2.

For the weak condition number in the complex case, if $\delta \leq (n - r)/N$, then setting

$$t := \frac{1}{\gamma_P} \sqrt{\frac{n-r}{\delta N}}$$

we get $\gamma_P t \geq 1$, and therefore, using the complex tail bound from Corollary 4.2,

$$\mathbb{P}\{\sigma_E \geq t\} \leq \frac{1}{\gamma_P^2 t^2} \frac{n-r}{N} = \delta.$$

This yields $\kappa_w(\delta) \leq t$. If $\delta > (n - r)/N$, then we use the fact that the weak condition number is monotonically decreasing with δ (intuitively, the larger the set we are allowed to exclude, the smaller the condition number will be), to conclude that $\kappa_w(\delta) \leq \kappa_w(\delta_0) \leq 1/\gamma_P$, where $\delta_0 := (n - r)/N$.

For the real case, if $r < n$ we use the bound

$$\frac{\Gamma(N/2)\Gamma((n-r+1)/2)}{\Gamma((N+1)/2)\Gamma((n-r)/2)} \leq \sqrt{\frac{n-r}{N-1}},$$

which follows from (10). If $\delta < \sqrt{(n-r)/N}$, set

$$t := \frac{1}{\gamma_P} \sqrt{\frac{n-r}{N}} \frac{1}{\delta}.$$

Then

$$\mathbb{P}\{\sigma_E \geq t\} \leq \frac{1}{\gamma_P} \frac{2}{\pi} \sqrt{\frac{n-r}{N}} \frac{1}{t} = \frac{2}{\pi} \sqrt{\frac{N}{N-1}} \cdot \delta \leq \delta,$$

where for the last inequality we used the fact that $N \geq 2$. We conclude that $\kappa_w(\delta) \leq t$. If $\delta > \sqrt{(n-r)/N}$, then we use the monotonicity of the weak condition just as in the complex case. Finally, for the weak stochastic condition number in the real case, we use Lemma 5.4 with $a = \gamma_P^{-1}$, $C = \sqrt{(n-r)/N}$ and $t_0 = C(\delta\gamma_P)^{-1}$ in the conditional expectation. We just saw that $\kappa_w^{\mathbb{R}}(\delta) \leq t_0$, so that

$$\kappa_{ws}^{\mathbb{R}}(\delta) \leq \mathbb{E}[\sigma_E \mid \sigma_E \leq t_0] \leq \frac{1}{\gamma_P} \left(\frac{1}{1-\delta} \right) \left(1 + \sqrt{\frac{n-r}{N}} \log \left(\sqrt{\frac{n-r}{N}} \delta^{-1} \right) \right),$$

where we used Lemma 5.4 in the second inequality. □

6 Bounding the Weak Stochastic Condition Number

In this section, we illustrate how the weak condition number of the problem of computing a simple eigenvalue of a singular matrix polynomial can be estimated in practice. More precisely, we show that the weak condition number of a singular problem can be estimated in terms of the stochastic condition number of nearby regular problems. Before deriving the relevant estimates, given in Theorem 6.3, we discuss the stochastic condition number of regular matrix polynomials.

6.1 Measure Concentration for the Directional Sensitivity of Regular Matrix Polynomials

For the directional sensitivity in the regular case, $r = n$, the worst-case condition number is γ_P^{-1} , as was shown in Proposition 3.4. In addition, the expression for the stochastic condition number involves a ratio of gamma functions (see Corollary 4.2 or the case $r = n$ in Theorems 5.1 and 5.2). From (10), we get the approximation $\Gamma(k+1/2)/\Gamma(k) \approx \sqrt{k}$, so that the stochastic condition number for regular polynomial eigenvalue problems satisfies

$$\kappa_s \approx \frac{1}{\sqrt{N}} \kappa.$$

This is compatible with previously known results about the stochastic condition number in the smooth setting (see discussion in Sect. 2). A natural question is whether the directional sensitivity is likely to be closer to this expected value, or closer to the upper bound κ .

Theorem 4.1 describes the distribution of σ_E as that of the (scaled) square root of a beta random variable. Using the interpretation of beta random variables as squared lengths of projections of uniformly distributed vectors on the sphere (see Sect. 4.1.1), tail bounds for the distribution of σ_E therefore translate into the problem of bounding the relative volume of certain subsets of the unit sphere. A standard argument from the realm of measure concentration on spheres, Lemma 6.1, then implies that with high probability, σ_E will stay close to its mean.

Lemma 6.1 *Let $x \sim \mathcal{U}(\beta N)$ be a uniformly distributed vector on the (real or complex) unit sphere, where $\beta = 1$ if $\mathbb{F} = \mathbb{R}$ and $\beta = 2$ if $\mathbb{F} = \mathbb{C}$. Then*

$$\mathbb{P}\{|x_1| \geq t\} \leq e^{-\beta(N-1)t^2/2}.$$

Proof For complex perturbations, we get the straightforward bound

$$\mathbb{P}\{|x_1| \geq t\} \leq (1 - t^2\gamma_p^2)^{N-1} \leq e^{-(N-1)t^2}.$$

In the real case, a classic result (see [6, Lemma 2.2] for a short and elegant proof) states that the probability in question is bounded by

$$\mathbb{P}\{|x_1| \geq t\} \leq e^{-Nt^2/2}. \quad (14)$$

The claimed bound follows by replacing N with $N - 1$ for the sake of a uniform presentation. \square

The next corollary follows from the description of the distribution of σ_E in Theorem 4.1 and the characterization of beta random variables as squared projected lengths of uniform vectors from Sect. 4.1.1.

Corollary 6.2 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a regular matrix polynomial and let λ be a simple eigenvalue of $P(x)$. If $E \sim \mathcal{U}(\beta N)$, where $\beta = 1$ if $\mathbb{F} = \mathbb{R}$ and $\beta = 2$ if $\mathbb{F} = \mathbb{C}$, then for $t \leq \gamma_p^{-1}$ we have*

$$\mathbb{P}\{\sigma_E \geq t\} \leq e^{-\beta(N-1)\gamma_p^2 t^2/2}.$$

6.2 The Weak Condition Number in Terms of Nearby Stochastic Condition Numbers

It is common wisdom that computing the condition number is as hard as solving the problem at hand, so at the very least we would like to avoid making the computation of the condition estimate more expensive than the computation of the eigenvalue itself. We will therefore aim to estimate the condition number of the problem in terms of

the output of a backward stable algorithm for computing the eigenvalue and a pair of associated eigenvectors.

Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a matrix polynomial of rank $r < n$ with a simple eigenvalue $\lambda \in \mathbb{C}$, and let $E(x) \in \mathbb{F}_d^{n \times n}[x]$ be a regular perturbation. Denote by $\lambda(\epsilon)$ the eigenvalue of $P(x) + \epsilon E(x)$ that converges to λ (see Theorem 3.2), and let $u(\epsilon)$ and $v(\epsilon)$ be the corresponding left and right eigenvectors of the perturbed problem. As shown in [14, Theorem 4] (see Theorem 6.4), for all $E(x)$ outside a proper Zariski closed set, the limits

$$\bar{u} = \lim_{\epsilon \rightarrow 0} u(\epsilon), \quad \bar{v} = \lim_{\epsilon \rightarrow 0} v(\epsilon)$$

converge to representatives of left and right eigenvectors of $P(x)$ associated with λ . Whenever these limits exist and represent eigenvectors of $P(x)$, define

$$\begin{aligned} \bar{\gamma}_P &:= \bar{u}^* P'(\lambda) \bar{v} \cdot \left(\sum_{j=0}^d |\lambda|^{2j} \right)^{-1/2}, \quad \bar{\kappa} = \bar{\gamma}_P^{-1}, \quad \text{and} \\ \bar{\kappa}_s &= \bar{\gamma}_P^{-1} \begin{cases} \frac{\sqrt{\pi}}{2} \frac{\Gamma(N)}{\Gamma(N+1/2)} & \text{if } \mathbb{F} = \mathbb{C}, \\ \frac{1}{\sqrt{\pi}} \frac{\Gamma(N/2)}{\Gamma((N+1)/2)} & \text{if } \mathbb{F} = \mathbb{R}. \end{cases} \end{aligned} \tag{15}$$

Note that these parameters depend implicitly on a perturbation direction $E(x)$, even though the notation does not reflect this. The parameters $\bar{\kappa}$ and $\bar{\kappa}_s$ are the limits of the worst-case and stochastic condition numbers, $\kappa(P(x) + \epsilon E(x))$ and $\kappa_s(P(x) + \epsilon E(x))$, as $\epsilon \rightarrow 0$. Since almost sure convergence implies convergence in probability, we get

$$\mathbb{E}[\bar{\kappa}_s] = \lim_{\epsilon \rightarrow 0} \mathbb{E}[\kappa_s(P(x) + \epsilon E(x))]$$

whenever the left-hand side of this expression is finite.

A backward stable algorithm, such as vanilla QZ, computes an eigenvalue $\tilde{\lambda}$ and associated unit norm eigenvectors \tilde{u} and \tilde{v} of a nearby problem $P(x) + \epsilon E(x)$. If ϵ is small, then $\tilde{\lambda} \approx \lambda$, $\tilde{u} \approx \bar{u}$ and $\tilde{v} \approx \bar{v}$, so that we can approximate the values (15) using the output of such an algorithm. Unfortunately, this does not yet give us a good estimate of γ_P , as the definition of γ_P makes use of very special representatives of eigenvectors (recall from Sect. 3.1 that for a singular matrix polynomials, eigenvectors are only defined as equivalence classes). The following theorem shows that we can still get bounds on the weak condition numbers in terms of $\bar{\kappa}_s$.

Theorem 6.3 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a singular matrix polynomial of rank $r < n$ with simple eigenvalue $\lambda \in \mathbb{C}$. Then*

$$\kappa_w(\delta) \leq \bar{\kappa} \cdot \max \left\{ \delta^{-1/\beta} \sqrt{\frac{n-r}{N}}, 1 \right\}.$$

If $\delta \leq (n-r)/N$, then for any $\eta > 0$ we have the tail bounds

$$\mathbb{P} \left\{ \delta^{-1/\beta} \bar{\kappa}_s \geq \eta \cdot \kappa_w(\delta) \right\} \geq 1 - e^{-\beta/\eta^2}.$$

For the proof of Theorem 6.3 we recall the setting of Sect. 3. Let $X = [U \ u]$ and $Y = [V \ v]$ be matrices whose columns are orthonormal bases of $\ker P(\lambda)^*$ and $\ker P(\lambda)$, respectively, such that U and V are bases of $\ker_\lambda P(x)^T$ and $\ker_\lambda P(x)$, respectively. If $\bar{u} = u$ and $\bar{v} = v$ in (15), then $\bar{\gamma}_P = \gamma_P$. In general, however, we only get a bound. To see this, recall from Sect. 3.1 that $\bar{u}^* P'(\lambda) \bar{v}$ depends only on the component of \bar{u} that is orthogonal to $\ker_\lambda P(\lambda)^*$, and the component of \bar{v} that is orthogonal to $\ker_\lambda P(\lambda)$. In particular, $X^* P'(\lambda) Y$ has rank one, and we have (recall $\ell = n - r + 1$)

$$X^* P'(\lambda) Y = u^* P'(\lambda) v \cdot e_\ell e_\ell^*. \tag{16}$$

The key to Proposition 6.3 lies in a result analogous to Theorem 3.2 for the eigenvectors by de Terán and Dopico [14, Theorem 4].

Theorem 6.4 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be matrix polynomial of rank r with simple eigenvalue λ and X, Y as above. Let $E(x) \in \mathbb{F}_d^{n \times n}[x]$ be such that $X^* E(\lambda) Y$ is non-singular. Let ζ be the eigenvalue of the non-singular matrix pencil*

$$X^* E(\lambda) Y + \zeta \cdot X^* P'(\lambda) Y, \tag{17}$$

and let a and b be the corresponding left and right eigenvectors. Then, for small enough $\epsilon > 0$, the perturbed matrix polynomial $P(x) + \epsilon E(x)$ has exactly one eigenvalue $\lambda(\epsilon)$ as described in Theorem 3.2, and the corresponding left and right eigenvectors satisfy

$$u(\epsilon) = Xa + O(\epsilon), \quad v(\epsilon) = Yb + O(\epsilon).$$

Given a matrix polynomial $P(x)$ and a perturbation direction $E(x)$, we can therefore assume that the eigenvectors of a sufficiently small perturbation in direction $E(x)$ are approximated by $\bar{u} = Xa$ and $\bar{v} = Yb$, where a, b are the eigenvectors of the matrix pencil (17). We would next like to characterize these eigenvectors for random perturbations $E(x)$. As with the rest of this paper, the following result is parametrized by a parameter $\beta \in \{1, 2\}$ which specifies whether we work with real or complex perturbations.

Proposition 6.5 *Let $P(x) \in \mathbb{F}_d^{n \times n}[x]$ be a matrix polynomial of rank $r < n$ with simple eigenvalue $\lambda \in \mathbb{C}$, and let $E(x) \sim \mathcal{U}(\beta N)$ be a random perturbation. Let a, b be left and right eigenvectors of the linear pencil (17), let $\bar{u} = Xa$ and $\bar{v} = Yb$, and define $\bar{\gamma}_P$ as in (15). Then*

$$\mathbb{E}[\bar{\gamma}_P] \leq (\ell - 1)^{-1/2} \gamma_P, \quad \text{and} \quad \mathbb{P}\{\bar{\gamma}_P \geq \gamma_P \cdot t\} \leq e^{-\beta(\ell-1)t^2/2}.$$

Proof By scale invariance of (17), we may take $E(x)$ to be Gaussian, $E(x) \sim \mathcal{N}^\beta(0, \sigma^2 I_{\beta N})$ with $\sigma^2 = (\sum_{j=0}^d |\lambda|^{2j})^{-1}$ (so that $E(\lambda) \sim \mathcal{G}_n^\beta(0, 1)$). Set $G := X^* E(\lambda) Y$, so that $G \sim \mathcal{G}_\ell^\beta(0, 1)$. Using (16), the eigenvectors associated with (17) are then characterized as solutions of

$$a^*(G + \zeta \cdot \gamma_P e_\ell e_\ell^*) = 0, \quad (G + \zeta \cdot \gamma_P e_\ell e_\ell^*)b = 0.$$

It follows that G^*a and Gb are proportional to e_ℓ , and hence,

$$a = \frac{G^{-*}e_\ell}{\|G^{-*}e_\ell\|}, \quad b = \frac{G^{-1}e_\ell}{\|G^{-1}e_\ell\|}.$$

Clearly, each of the vectors a and b individually is uniformly distributed. They are, however, not independent. To simplify notation, set $H = G^{-1}$. For the condition estimate we get, using (16),

$$|a^* X^* P'(\lambda) Y b| = |u^* P'(\lambda) v| \cdot \frac{|e_\ell^* H e_\ell| \cdot |e_\ell^* H e_\ell|}{\|H^* e_\ell\| \cdot \|H e_\ell\|} \leq |u^* P'(\lambda) v| \cdot \frac{|e_\ell^* H e_\ell|}{\|H e_\ell\|}.$$

By orthogonal/unitary invariance of the Gaussian distribution, the random vector $q := H e_\ell / \|H e_\ell\|$ is uniformly distributed on $S^{\beta\ell-1}$. It follows that $|e^T H e_\ell| / \|H e_\ell\|$ is distributed like the absolute value of the projection of a uniform vector onto the first coordinate. For the expected value, the bound follows by observing that the expected value of such a projection is bounded by $(\ell - 1)^{-1/2}$. For the tail bound, using (14) (with N replaced by ℓ) we get

$$\mathbb{P}\{\bar{\gamma}_P \geq \gamma_P \cdot t\} = \mathbb{P}\{|a^* X^* P'(\lambda) Y b| \geq |u^* P'(\lambda) v| \cdot t\} \leq e^{-\beta(\ell-1)t^2/2}.$$

This was to be shown. □

Proof of Theorem 6.3 If $\bar{u} = Xa$ and $\bar{v} = Yb$, then

$$|\bar{u}^* P'(\lambda) \bar{v}| = |a^T X^* P'(\lambda) Y b| = |a_\ell b_\ell u^* P'(\lambda) v| \leq |u^* P'(\lambda) v|,$$

and we get the upper bound

$$\gamma_P^{-1} \leq \bar{\gamma}_P^{-1} = \bar{\kappa}.$$

For the weak condition numbers, using Theorems 5.1 and 5.2, we get the bounds

$$\kappa_w(\delta) \leq \bar{\kappa} \cdot \max \left\{ \sqrt{\frac{n-r}{N}} \frac{1}{\delta^{1/\beta}}, 1 \right\}.$$

For the tails bounds in the complex case, note that in the complex case we have

$$\begin{aligned} \mathbb{P} \left\{ (N\delta)^{-1/2} \bar{\kappa} \leq \eta \cdot \kappa_w(\delta) \right\} &= \mathbb{P} \left\{ \bar{\gamma}_P^{-1} \leq \eta \cdot \sqrt{n-r} \gamma_P^{-1} \right\} \\ &= \mathbb{P} \left\{ \bar{\gamma}_P \geq \eta^{-1} \cdot (n-r)^{-1/2} \gamma_P \right\} \leq e^{-1/\eta^2}, \end{aligned}$$

where we used Proposition 6.5 for the inequality. The real case follows in the same way. □

7 Conclusions and Outlook

The classical theory of conditioning in numerical analysis aims to quantify the susceptibility of a computational problem to perturbations in the input. While the theory serves its purpose well in distinguishing well-posed problems from problems that approach ill-posedness, it fails to explain why certain problems with high condition number can still be solved satisfactory to high precision by algorithms that are oblivious to the special structure of an input. By introducing the notions of weak and weak stochastic conditioning, we developed a tool to better quantify the perturbation behaviour of numerical computation problems for which the classical condition number fails to do so.

Our methods are based on an analysis of directional perturbations and probabilistic tools. The use of probability theory in our context is auxiliary: the purpose is to quantify the observation that the set of adversarial perturbations is small. In practice, any reasonable numerical algorithm will find the eigenvalues of a nearby regular matrix polynomial, and the perturbation will be deterministic and not random. However, as the algorithm knows nothing about the particular input matrix polynomial, it is reasonable to assume that if the set of adversarial perturbations is sufficiently small, then the actual perturbation will not be in there. Put more directly, to say that the probability that a perturbed problem has large directional sensitivity is very small is to say that a perturbation, although non-random, would need a good reason to cause damage.

The results presented continue the line of work of [3], where it is argued that, just as sufficiently small numbers are considered numerically indistinguishable from zero, sets of sufficiently small measure should be considered numerically indistinguishable from null sets. One interesting direction in which the results presented can be strengthened is to use wider classes of probability distributions, including such that are discrete, and derive equivalent (possibly slightly weaker) results. One important side effect of our analysis is a focus away from the expected value, and more towards robust measures such as the median⁸ and other quantiles.

Our results hence have a couple of important implications, or “take-home messages”, that we would like to highlight:

1. The results presented call for a critical re-evaluation of the notion of ill-posedness. It has become common practice to simply *identify* ill-posedness with having infinite condition, to the extent that condition numbers are often *defined* in terms of the inverse distance to a set of ill-posed inputs, an approach that has been popularized by J. Demmel [16,17].⁹ The question of whether the elements of such a set are actually badly behaved a practical sense is often left unquestioned. Our theory suggests that the set of inputs that are actually ill-behaved from a practical point of view can be smaller than previously thought.

⁸ The use of the median instead of the expected value in the probabilistic analysis of quantities was suggested by F. Bornemann [10].

⁹ For the complexity analysis of iterative algorithms, and in particular for problems related to convex optimization, the “distance to ill-posedness” approach may often be the most natural setting. For convex feasibility problems, for example, the ill-posed inputs form a wall separating primal from dual feasible problem instances, and closeness to this wall directly affects the speed of convergence of iterative algorithms; see [13] for more on this story.

2. Average-case analysis (and its refinement, smoothed analysis [12]) is, while well intentioned, still susceptible to the caprices of specific probability distributions. More meaningful results are obtained when, instead of analysing the behaviour of perturbations on average, one shifts the focus towards showing that the set of *adversarial* perturbations is small; ideally so small, that hitting a misbehaving perturbation would suggest the existence of a specific explanation for this rather than just bad luck. In terms of summary parameters, our approach suggests using, in line with common practice in statistics, more robust parameters such as the median instead of the mean.

A natural question that arises from the first point is: if some problems that were previously thought of as ill-posed are not (in the sense that the set of discontinuous perturbation directions is negligible), then which problems are genuinely ill-posed? In the case of polynomial eigenvalue problems, we conjecture that problems with semisimple eigenvalues are not ill-conditioned in our framework; in fact, it appears that much of the analysis performed in this section can be extended to this setting. It is not completely obvious which problems should be considered ill-posed based on this new theory. That some inputs still should can be seen, for example, by considering Jordan blocks with zeros on the diagonal; the computed eigenvalues of perturbations of the order of machine precision will not recover the correct eigenvalue in this situation. Our analysis in the semisimple case is based on the fact that the directional derivative of the function to be computed exists in sufficiently many directions.

Another consequence is that much of the probabilistic analyses of condition numbers based on the distance to ill-posedness, while still correct, can possibly be refined when using a smaller set of ill-posed inputs. In particular, it is likely that condition bounds resulting from average-case and smoothed analysis can be refined. Finally, an interesting direction would be to examine problems with high or infinite condition number that are not ill-posed in a practical sense in different contexts, such as polynomial system solving or problems arising from the discretization of continuous inverse problems.

Acknowledgements Open access funding provided by Aalto University. The spark that led to this paper was ignited at the workshop “Algebra meets numerics: condition and complexity” on November 6–7, 2017 in Berlin; we are grateful to the organizers Peter Bürgisser and Felipe Cucker for inviting us and for pointing out the work of Armentano and Stewart. In addition, the authors would like to thank Carlos Beltrán and Daniel Kressner for valuable feedback, and the anonymous referees for useful comments. We are greatly indebted to Dennis Amelunxen, whose vision of weak-average-case analysis inspired this work. We would like to acknowledge financial support from the Manchester Institute for Mathematical Sciences (MIMS) during the early stages of this project and the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the programme Approximation, Sampling and Compression in Data Science while this work was completed.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Moments and Tails for Ratios of Beta Random Variables

In appendix, we compute the expected value and tail bounds for moments of quotients of beta random variables.

Lemma A.1 *Let $a > 0, b > 0, c > 0, d > 0$ and $X \sim B(a, b), Y \sim B(c, d)$ random variables. Then for k such that $ck > 1$,*

$$\mathbb{E} \left[X^{1/k} \right] = \frac{B(a + 1/k, b)}{B(a, b)} \quad \text{and} \quad \mathbb{E} \left[(X/Y)^{1/k} \right] = \frac{B(a + 1/k, b)B(c - 1/k, d)}{B(a, b)B(c, d)},$$

If $ck = 1$, then $\mathbb{E}[(X/Y)^k] = \infty$. Moreover, the probability tails are bounded by

$$\mathbb{P}\{(X/Y)^{1/k} \geq t\} \leq \begin{cases} 1 - t^{ak} \frac{B(a+c,d)}{B(c,d)} & \text{if } t \leq 1 \\ \frac{1}{t^{ck}} \frac{B(a+c,b)}{cB(a,b)B(c,d)} & \text{if } t \geq 1 \end{cases}$$

Proof We focus on the case of the quotient $(X/Y)^{1/k}$; the statement for $X^{1/k}$ follows by simply setting $Y = 1$ in the calculations below. Set $C = 1/(B(a, b)B(c, d))$. For $ck > 1$,

$$\begin{aligned} \mathbb{E} \left[(X/Y)^{1/k} \right] &= \int_{t=0}^{\infty} \mathbb{P}\{X/Y \geq t^k\} dt \\ &= C \cdot \int_{y=0}^1 \int_{t=0}^{y^{-1/k}} \int_{x=t^k y}^1 x^{a-1} (1-x)^{b-1} y^{c-1} (1-y)^{d-1} dx dt dy \\ &\stackrel{s=y^{1/k}t}{=} C \cdot \int_{y=0}^1 \int_{s=0}^1 \int_{x=s^k}^1 x^{a-1} (1-x)^{b-1} y^{c-1/k-1} (1-y)^{d-1} dx ds dy \end{aligned} \tag{A}$$

$$\begin{aligned} &= C \cdot B(c - 1/k, d) \int_{s=0}^1 \int_{x=s^k}^1 x^{a-1} (1-x)^{b-1} dx ds \\ &= C \cdot B(c - 1/k, d) \int_{x=0}^1 \int_{s=0}^{x^{1/k}} x^{a-1} (1-x)^{b-1} ds dx \\ &= C \cdot B(a + 1/k, b)B(c - 1/k, d). \end{aligned} \tag{B}$$

If $ck = 1$, then the step from (A) to (B) breaks down, since the integral $\int_0^1 y^{-1} (1-y)^{d-1} dy$ diverges. For the tail bound, we proceed similarly. If $t \leq 1$, then

$$\begin{aligned} \mathbb{P}\{(X/Y)^{1/k} \geq t\} &= 1 - C \cdot \int_{y=0}^1 \int_{x=0}^{t^k y} x^{a-1} (1-x)^{b-1} y^{c-1} (1-y)^{d-1} dx dy \\ &\stackrel{x=(t^k y)z}{=} 1 - C \cdot t^{ak} \int_{y=0}^1 \int_{z=0}^1 z^{a-1} (1-t^k yz)^{b-1} y^{a+c-1} (1-y)^{d-1} dy dz \end{aligned}$$

$$\begin{aligned} &\stackrel{(1)}{\leq} 1 - C \cdot t^{ak} \int_{y=0}^1 \int_{z=0}^1 z^{a-1} (1-z)^{b-1} y^{a+c-1} (1-y)^{d-1} dy dz \\ &= 1 - t^{ak} \frac{B(a+c, d)}{B(c, d)}, \end{aligned}$$

where for the inequality (1) we used that $t^k yz \leq z$ when $t \leq 1$. If $t \geq 1$, then

$$\begin{aligned} \mathbb{P}\{(X/Y)^{1/k} \geq t\} &= C \cdot \int_{x=0}^1 \int_{y=0}^{t^{-k}x} x^{a-1} (1-x)^{b-1} y^{c-1} (1-y)^{d-1} dy dx \\ &\stackrel{y=(t^{-k}x)z}{=} C \cdot \frac{1}{t^{ck}} \int_{x=0}^1 \int_{z=0}^1 x^{a+c-1} (1-x)^{b-1} z^{c-1} (1-t^{-k}xz)^{d-1} dz dx \\ &\stackrel{(2)}{\leq} C \cdot \frac{1}{t^{ck}} \int_{x=0}^1 \int_{z=0}^1 x^{a+c-1} (1-x)^{b-1} z^{c-1} dz dx \\ &= \frac{1}{t^{ck}} \frac{B(a+c, b)}{cB(a, b)B(c, d)}, \end{aligned}$$

where for (2) we used that $t^{-k}xz \geq 0$. □

References

1. B. Adhikari, R. Alam, and D. Kressner. Structured eigenvalue condition numbers and linearizations for matrix polynomials. *Linear algebra and its applications*, 435(9):2193–2221, 2011.
2. R. Alam and S. Safique Ahmad. Sensitivity analysis of nonlinear eigenproblems. *SIAM Journal on Matrix Analysis and Applications*, 40(2):672–695, 2019.
3. D. Amelunxen and M. Lotz. Average-case complexity without the black swans. *Journal of Complexity*, 41:82–101, 2017.
4. D. Armentano. Stochastic perturbations and smooth condition numbers. *Journal of Complexity*, 26(2):161–171, 2010.
5. D. Armentano and C. Beltrán. The polynomial eigenvalue problem is well conditioned for random inputs. *SIAM Journal on Matrix Analysis and Applications*, 40(1):175–193, 2019.
6. K. Ball. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31:1–58, 1997.
7. K. Beltrán, C. and Kozhasov. The real polynomial eigenvalue problem is well conditioned on the average. *Foundations of Computational Mathematics*, May 2019.
8. L. Blum, F. Cucker, M. Schub, and S. Smale. *Complexity and Real Computation*. Springer-Verlag, New York, NY, USA, 1998.
9. L. Boltzmann. Referat über die Abhandlung von J.C. Maxwell: “Über Boltzmann’s Theorem betreffend die mittlere Verteilung der lebendige Kraft in einem System materieller Punkte”. *Wied. Ann. Beiblätter*, 5:403–417, 1881.
10. F. Bornemann. Private communication, 2018.
11. S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
12. P. Bürgisser. Smoothed analysis of condition numbers. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures*, pages 2609–2633. World Scientific, 2010.
13. P. Bürgisser and F. Cucker. *Condition. The Geometry of Numerical Algorithms*. Springer, Berlin-Heidelberg, Germany, 2013.

14. F. De Terán and F. M. Dopico. First order spectral perturbation theory of square singular matrix polynomials. *Linear Algebra and its Applications*, 432(4):892–910, 2010.
15. J.-P. Dedieu and F. Tisseur. Perturbation theory for homogeneous polynomial eigenvalue problems. *Linear algebra and its applications*, 358(1-3):71–94, 2003.
16. J. Demmel. On condition numbers and the distance to the nearest ill-posed problem. *Numer. Math.*, 51:251–289, 1987.
17. J. Demmel. The probability that a numerical analysis problem is difficult. *Math. Comp.*, 50:449–480, 1988.
18. A. Dmytryshyn and F. M. Dopico. Generic complete eigenstructures for sets of matrix polynomials with bounded rank and degree. *Linear Algebra and its Applications*, 535:213–230, 2017.
19. F. Dopico and V. Noferini. Root polynomials and their role in the theory of matrix polynomials. *Submitted to Linear Algebra Appl.*, 2018.
20. F. J. Dyson. The threefold way. algebraic structure of symmetry groups and ensembles in quantum mechanics. *Journal of Mathematical Physics*, 3(6):1199–1215, 1962.
21. A. Edelman. Eigenvalues and condition numbers of random matrices. *SIAM J. of Matrix Anal. and Applic.*, 9:543–556, 1988.
22. A. Edelman and N. R. Rao. Random matrix theory. *Acta Numerica*, 14:233–297, 2005.
23. G. Forney Jr. Minimal bases of rational vector spaces, with applications to multivariable linear systems. *SIAM J. Control*, 13:493–520, 1975.
24. P. J. Forrester. *Log-gases and random matrices (LMS-34)*. Princeton University Press, 2010.
25. J. Ginibre. Statistical ensembles of complex, quaternion, and real matrices. *Journal of Mathematical Physics*, 6(3):440–449, 1965.
26. I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. SIAM, Philadelphia, USA, 2009.
27. H. H. Goldstine and J. Von Neumann. Numerical inverting of matrices of high order. II. *Proceedings of the American Mathematical Society*, 2(2):188–202, 1951.
28. G. H. Golub and J. H. Wilkinson. Ill-conditioned eigensystems and the computation of the jordan canonical form. *SIAM Rev.*, 18(4):578–619, 1976.
29. N. J. Higham. *Accuracy and stability of numerical algorithms*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA, 1996.
30. N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2008.
31. N. J. Higham and T. Mary. A new approach to probabilistic rounding error analysis. 2018.
32. M. Hochstenbach, C. Mehl, and B. Plestenjak. Solving singular generalized eigenvalue problems by a rank-completing perturbation. Preprint 03-2018, TU Berlin, Berlin, 2018.
33. M. Konstantinov, D. W. Gu, V. Mehrmann, and P. Petkov. *Perturbation theory for matrix equations*, volume 9. Gulf Professional Publishing, 2003.
34. E. Kostlan. Complexity theory of numerical linear algebra. *Journal of Computational and Applied Mathematics*, 22(2-3):219–230, 1988.
35. F. Mezzadri. How to generate random matrices from the classical compact groups. *Notices of the American Mathematical Society*, 54(5):592–604, 5 2007.
36. C. B. Moler and G. W. Stewart. An algorithm for generalized matrix eigenvalue problems. *SIAM J. Numer. Anal.*, 10(2):241–256, 1973.
37. V. Noferini. The behaviour of the complete eigenstructure of a polynomial matrix under a generic rational transformation. *Electr. J. Lin. Algebra*, 22:607–624, 2012.
38. J. W. Pearson. Computation of hypergeometric functions. Master’s thesis, Oxford University, 2009.
39. J. W. Pearson, S. Olver, and M. A. Porter. Numerical methods for the computation of the confluent and gauss hypergeometric functions. *Numerical Algorithms*, 74(3):821–866, 2017.
40. S. Smale. The fundamental theorem of algebra and complexity theory. *Bulletin of the American Mathematical Society*, 4(1):1–36, 1981.
41. G. W. Stewart. Stochastic perturbation theory. *SIAM review*, 32(4):579–610, 1990.
42. F. Tisseur. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra and its Applications*, 309(1-3):339–361, 2000.
43. L. N. Trefethen. The smart money’s on numerical analysts. *SIAM News*, 45(9):1–5, 2012.
44. L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997.
45. A. Turing. Rounding-off errors in matrix processes. *Quart. J. Mech. Appl. Math.*, 1:287–308, 1948.
46. P. Van Dooren. The computation of Kronecker’s canonical form of a singular pencil. *Linear Algebra and its Applications*, 27:103–140, 1979.

47. J. von Neumann and H. Goldstine. Numerical inverting matrices of high order. *Bulleting of the AMS*, 53:1021–1099, 1947.
48. N. Weiss, G. Wasilkowski, H. Woźniakowski, and M. Shub. Average condition number for solving linear equations. *Linear Algebra and Its Applications*, 83:79–102, 1986.
49. J. G. Wendel. Note on the gamma function. *The American Mathematical Monthly*, 55(9):563–564, 1948.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.