
Toward Machine Wald

Houman Owhadi and Clint Scovel

Contents

1	Introduction	2
2	The UQ Problem Without Sample Data	4
2.1	Čebyšev, Markov, and Kreĭn	4
2.2	Optimal Uncertainty Quantification	4
2.3	Worst-Case Analysis	7
2.4	Stochastic and Robust Optimization	7
2.5	Čebyšev Inequalities and Optimization Theory	8
3	The UQ Problem with Sample Data	9
3.1	From Game Theory to Decision Theory	9
3.2	The Optimization Approach to Statistics	10
3.3	Abraham Wald	11
3.4	Generalization to Unknown Pairs of Functions and Measures and to Arbitrary Sample Data	13
3.5	Model Error and Optimal Models	14
3.6	Mean Squared Error, Variance, and Bias	15
3.7	Optimal Interval of Confidence	15
3.8	Ordering the Space of Experiments	16
3.9	Mixing Models	16
4	The Complete Class Theorem and Bayesian Inference	17
4.1	The Bayesian Approach	17
4.2	Relation Between Adversarial Model Error and Bayesian Error	18
4.3	Complete Class Theorem	19
5	Incorporating Complexity and Computation	21
5.1	Machine Wald	22
5.2	Reduction Calculus	22
5.3	Stopping Conditions	23
5.4	On the Borel-Kolmogorov Paradox	23
5.5	On Bayesian Robustness/Brittleness	24

H. Owhadi (✉) • C. Scovel

Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA
e-mail: owhadi@caltech.edu; clintscovel@gmail.com

5.6 Information-Based Complexity	25
6 Conclusion	26
Appendix	26
Construction of $\pi \odot \mathbb{D}$	26
Proof of Theorem 2	27
Conditional Expectation as an Orthogonal Projection	28
References	29

Abstract

The past century has seen a steady increase in the need of estimating and predicting complex systems and making (possibly critical) decisions with limited information. Although computers have made possible the numerical evaluation of sophisticated statistical models, these models are still designed *by humans* because there is currently no known recipe or algorithm for dividing the design of a statistical model into a sequence of arithmetic operations. Indeed enabling computers to *think as humans*, especially when faced with uncertainty, is challenging in several major ways: (1) Finding optimal statistical models remains to be formulated as a well-posed problem when information on the system of interest is incomplete and comes in the form of a complex combination of sample data, partial knowledge of constitutive relations and a limited description of the distribution of input random variables. (2) The space of admissible scenarios along with the space of relevant information, assumptions, and/or beliefs, tends to be infinite dimensional, whereas calculus on a computer is necessarily discrete and finite. With this purpose, this paper explores the foundations of a rigorous framework for the scientific computation of optimal statistical estimators/models and reviews their connections with decision theory, machine learning, Bayesian inference, stochastic optimization, robust optimization, optimal uncertainty quantification, and information-based complexity.

Keywords

Abraham Wald • Decision theory • Machine learning • Uncertainty quantification • Game theory

1 Introduction

During the past century, the need to solve large complex problems in applications such as fluid dynamics, neutron transport, or ballistic prediction drove the parallel development of computers and numerical methods for solving ODEs and PDEs. It is now clear that this development lead to a paradigm shift. Before: each new PDE required the development of new theoretical methods and the employment of large teams of mathematicians and physicists; in most cases, information on solutions was only qualitative and based on general analytical bounds on fundamental solutions. After: mathematical analysis and computer science worked in synergy to give birth to robust numerical methods (such as finite element methods) capable of solving a

large spectrum of PDEs without requiring the level of expertise of an A. L. Cauchy or level of insight of a R. P. Feynman. This transformation can be traced back to sophisticated calculations performed by arrays of *human computers* organized as parallel clusters such as in the pioneering work of Lewis Fry Richardson [1, 90], who in 1922 had a room full of clerks attempt to solve finite-difference equations for the purposes of weather forecasting, and the 1947 paper by John Von Neumann and Herman Goldstine on Numerical Inverting of Matrices of High Order [154]. Although Richardson's predictions failed due to the use of unfiltered data/initial conditions/equations and large time-steps not satisfying the CFL stability condition [90], his vision was shared by Von Neumann [90] in his proposal of the Meteorology Research Project to the US Navy in 1946, qualified by Platzman [120] as "perhaps the most visionary prospectus for numerical weather prediction since the publication of Richardson's book a quarter-century earlier."

The past century has also seen a steady increase in the need of estimating and predicting complex systems and making (possibly critical) decisions with limited information. Although computers have made possible the numerical evaluation of sophisticated statistical models, these models are still designed *by humans* through the employment of multidisciplinary teams of physicists, computer scientists, and statisticians. Contrary to the original *human computers* (such as the ones pioneered by L. F. Richardson or overseen by R. P. Feynman at Los Alamos), these *human teams* do not follow a specific algorithm (such as the one envisioned in Richardson's Forecast Factory where 64,000 human computers would have been working in parallel and at high speed to compute world weather charts [90]) because there is currently no known recipe or algorithm for dividing the design of a statistical model into a sequence of arithmetic operations. Furthermore, while *human computers* were given a specific PDE or ODE to solve, these *human teams* are not given a well-posed problem with a well-defined notion of solution. As a consequence, different *human teams* come up with different *solutions* to the design of the statistical model along with different estimates on uncertainties.

Indeed enabling computers to *think as humans*, especially when faced with uncertainty, is challenging in several major ways: (1) There is currently no known recipe or algorithm for dividing the design of a statistical model into a sequence of arithmetic operations. (2) Formulating the search for an optimal statistical estimator/model as a well-posed problem is not obvious when information on the *system of interest* is incomplete and comes in the form of a complex combination of sample data, partial knowledge of constitutive relations, and a limited description of the distribution of input random variables. (3) The space of admissible scenarios along with the space of relevant information, assumptions, and/or beliefs tends to be infinite dimensional, whereas calculus on a computer is necessarily discrete and finite.

The purpose of this paper is to explore the foundations of a rigorous/rational framework for the scientific computation of optimal statistical estimators/models for complex systems and review their connections with decision theory, machine learning, Bayesian inference, stochastic optimization, robust optimization, optimal uncertainty quantification, and information-based complexity, the most fundamental of these being the simultaneous emphasis on *computation* and *performance* as in machine learning initiated by Valiant [149].

2 The UQ Problem Without Sample Data

2.1 Čebyšev, Markov, and Kreĭn

Let us start with a simple warm-up problem.

Problem 1. Let \mathcal{A} be the set of measures of probability on $[0, 1]$ having mean less than $m \in (0, 1)$. Let μ^\dagger be an unknown element of \mathcal{A} and let $a \in (m, 1)$. What is $\mu^\dagger[X \geq a]$?

Observe that given the limited information on μ^\dagger , $\mu^\dagger[X \geq a]$ could a priori be any number in the interval $[\mathcal{L}(\mathcal{A}), \mathcal{U}(\mathcal{A})]$ obtained by computing the sup (inf) of $\mu[X \geq a]$ with respect to all possible candidates for μ^\dagger , i.e.,

$$\mathcal{U}(\mathcal{A}) := \sup_{\mu \in \mathcal{A}} \mu[X \geq a] \quad (1)$$

and

$$\mathcal{L}(\mathcal{A}) := \inf_{\mu \in \mathcal{A}} \mu[X \geq a]$$

where

$$\mathcal{A} := \{\mu \in \mathcal{M}([0, 1]) \mid \mathbb{E}_\mu[X] \leq m\}$$

and $\mathcal{M}([0, 1])$ is the set of Borel probability measures on $[0, 1]$. It is easy to observe that the extremum of (1) can be achieved only when μ is the weighted sum of a Dirac mass at 0 and a Dirac mass at a . It follows that, although (1) is an infinite dimensional optimization problem, it can be reduced to the simple one-dimensional optimization problem obtained by letting $p \in [0, 1]$ denote the weight of the Dirac mass at 1 and $1 - p$ the weight of the Dirac mass at 0: *Maximize p subject to $ap = m$* , producing the Markov bound $\frac{m}{a}$ as solution.

Problems such as (1) can be traced back to Čebyšev [77, Pg. 4] “Given: length, weight, position of the centroid and moment of inertia of a material rod with a density varying from point to point. It is required to find the most accurate limits for the weight of a certain segment of this rod.” According to Kreĭn [77], although Čebyšev did solve this problem, it was his student Markov who supplied the proof in his thesis. See Kreĭn [77] for an account of the history of this subject along with substantial contributions by Kreĭn.

2.2 Optimal Uncertainty Quantification

The generalization of the process described in Sect. 2.1 to complex systems involving imperfectly known functions and measures is the point of view of optimal uncertainty quantification (OUQ) [3, 69, 72, 96, 114, 142]. Instead of developing

sophisticated mathematical solutions, the OUQ approach is to develop optimization problems and reductions, so that their solution may be implemented on a computer, as in Bertsimas and Popescu's [15] convex optimization approach to Čebyšev inequalities, and the Decision Analysis framework of Smith [133].

To present this generalization, for a topological space \mathcal{X} , let $\mathcal{F}(\mathcal{X})$ be the space of real-valued measurable functions and $\mathcal{M}(\mathcal{X})$ be the set of Borel probability measures on \mathcal{X} . Let \mathcal{A} be an arbitrary subset of $\mathcal{F}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$, and let $\Phi: \mathcal{A} \rightarrow \mathbb{R}$ be a function producing a quantity of interest.

Problem 2. Let (f^\dagger, μ^\dagger) be an unknown element of \mathcal{A} . What is $\Phi(f^\dagger, \mu^\dagger)$?

Therefore, in the absence of sample data, in the context of this generalization, one is interested in estimating $\Phi(f^\dagger, \mu^\dagger)$, where $(f^\dagger, \mu^\dagger) \in \mathcal{F}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$ corresponds to an *unknown reality*: the function f^\dagger represents a *response function* of interest, and μ^\dagger represents the probability distribution of the inputs of f^\dagger . If \mathcal{A} represents all that is known about (f^\dagger, μ^\dagger) (in the sense that $(f^\dagger, \mu^\dagger) \in \mathcal{A}$ and that any $(f, \mu) \in \mathcal{A}$ could, a priori, be (f^\dagger, μ^\dagger) given the available information), then [114] shows that the quantities

$$\mathcal{U}(\mathcal{A}) := \sup_{(f, \mu) \in \mathcal{A}} \Phi(f, \mu) \quad (2)$$

$$\mathcal{L}(\mathcal{A}) := \inf_{(f, \mu) \in \mathcal{A}} \Phi(f, \mu) \quad (3)$$

determine the inequality

$$\mathcal{L}(\mathcal{A}) \leq \Phi(f^\dagger, \mu^\dagger) \leq \mathcal{U}(\mathcal{A}), \quad (4)$$

to be optimal given the available information $(f^\dagger, \mu^\dagger) \in \mathcal{A}$ as follows: It is simple to see that the inequality (4) follows from the assumption that $(f^\dagger, \mu^\dagger) \in \mathcal{A}$. Moreover, for any $\varepsilon > 0$, there exists a $(f, \mu) \in \mathcal{A}$ such that

$$\mathcal{U}(\mathcal{A}) - \varepsilon < \Phi(f, \mu) \leq \mathcal{U}(\mathcal{A}).$$

Consequently since all that is known about (f^\dagger, μ^\dagger) is that $(f^\dagger, \mu^\dagger) \in \mathcal{A}$, it follows that the upper bound $\Phi(f^\dagger, \mu^\dagger) \leq \mathcal{U}(\mathcal{A})$ is the best obtainable given that information, and the lower bound is optimal in the same sense.

Although the OUQ optimization problems (2) and (3) are extremely large and although some are computationally intractable, an important subclass enjoys significant and practical finite-dimensional reduction properties [114]. First, by [114, Cor. 4.4], although the optimization variables (f, μ) lie in a product space of functions and probability measures, for OUQ problems governed by linear inequality constraints on generalized moments, the search can be reduced to one over probability measures that are products of finite convex combinations of Dirac masses with explicit upper bounds on the number of Dirac masses.

Furthermore, in the special case that all constraints are generalized moments of functions of f , the dependency on the coordinate positions of the Dirac masses is eliminated by observing that the search over admissible functions reduces to a search over functions on an m -fold product of finite discrete spaces, and the search over m -fold products of finite convex combinations of Dirac masses reduces to a search over the products of probability measures on this m -fold product of finite discrete spaces [114, Thm. 4.7]. Finally, by [114, Thm. 4.9], using the lattice structure of the space of functions, the search over these functions can be reduced to a search over a finite set.

Fundamental to this development is Winkler's [169] generalization of the characterization of the extreme points of compact (in the weak topology) sets of probability measures constrained by a finite number of generalized moment inequalities defined by continuous functions to *non-compact sets of tight measures*, in particular probability measures on Borel subsets of Polish metric spaces, defined by Borel measurable moment functions, along with his [168] development of Choquet theory for weakly closed convex *non-compact* sets of tight measures. These results are based on Kendall's [71] equivalence between a linearly compact Choquet simplex and a vector lattice and results of Dubins [31] concerning the extreme points of affinely constrained convex sets in terms of the extreme points of the unconstrained set. It is interesting to note that Winkler [169] uses Kendall's result to derive a strong sharpening of Dubins result [31]. Winkler's results allow the extension of existing optimization results over measures on compact metric spaces constrained by continuous generalized moment functions to optimization over measures on Borel subsets of Polish spaces constrained by Borel measurable moment functions. For systems with symmetry, the Choquet theorem of Varadarajan [151] can be used to show that the Dirac masses can be replaced by the ergodic measures in these results. The inclusion of sets of functions along with sets of measures in the optimization problems facilitates the application to systems with imprecisely known response functions. In particular, a result of Ressel [121], providing conditions under which the map $(f, \mu) \rightarrow f_*\mu$ from function/measure pairs to the induced law is Borel measurable, facilitates the extension of these techniques from sets of *measures* to sets of *random variables*. In general, the inclusion of functions in the domain of optimization requires the development of generalized programming techniques such as generalized Benders decompositions described in Geoffrion [46]. Moreover, as has been so successful in machine learning, it will be convenient to approximate the space of measurable functions $\mathcal{F}(\mathcal{X})$ by some reproducing kernel Hilbert space $\mathcal{H}(\mathcal{X}) \subset \mathcal{F}(\mathcal{X})$ producing an approximation $\mathcal{H}(\mathcal{X}) \times \mathcal{M}(\mathcal{X}) \subset \mathcal{F}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$ to the full base space. Under the mild assumption that \mathcal{X} is an analytic subset of a Polish space and $\mathcal{H}(\mathcal{X})$ possesses a measurable feature map, it has recently been shown in [111] that $\mathcal{H}(\mathcal{X})$ is separable. Consequently, since all separable Hilbert spaces are isomorphic with ℓ^2 , it follows that the space $\ell^2 \times \mathcal{M}(\mathcal{X})$ is a universal representation space for the approximation of $\mathcal{F}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$. Moreover, in that case, since \mathcal{X} is separable and metric, so is $\mathcal{M}(\mathcal{X})$ in the weak topology, and since $\mathcal{H}(\mathcal{X})$ is Polish, it follows that the approximation space $\mathcal{H}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$ is the product of a Polish

space and a separable metric space. When furthermore \mathcal{X} is Polish, it follows that the approximation space is the product of Polish spaces and therefore Polish.

Example 1. A classic example is $\Phi(f, \mu) := \mu[f \geq a]$ where a is a safety margin. In the certification context, one is interested in showing that $\mu^\dagger[f^\dagger \geq a] \leq \varepsilon$, where ε is a safety certification threshold (i.e., the maximum acceptable μ^\dagger -probability of the system f^\dagger exceeding the safety margin a). If $\mathcal{U}(\mathcal{A}) \leq \varepsilon$, then the system associated with (f^\dagger, μ^\dagger) is safe even in the worst-case scenario (given the information represented by \mathcal{A}). If $\mathcal{L}(\mathcal{A}) > \varepsilon$, then the system associated with (f^\dagger, μ^\dagger) is unsafe even in the best-case scenario (given the information represented by \mathcal{A}). If $\mathcal{L}(\mathcal{A}) \leq \varepsilon < \mathcal{U}(\mathcal{A})$, then the safety of the system cannot be decided (although one could declare the system to be unsafe due to lack of information).

2.3 Worst-Case Analysis

The proposed solutions to Problems 1 and 2 are particular instances of worst-case analysis that, as noted by [135] and [127, p.5], is an old concept that could be summarized by the popular adage *When in doubt, assume the worst!* or:

The gods to-day stand friendly, that we may,
Lovers of peace, lead on our days to age
But, since the affairs of men rests still uncertain,
Lets reason with the worst that may befall.

Julius Caesar, Act 5, Scene 1
William Shakespeare (1564–1616)

As noted in [114], an example of worst-case analysis in seismic engineering is that of Drenick's *critical excitation* [30] which seeks to quantify the safety of a structure to the worst earthquake given a constraint on its magnitude. The combination of structural optimization (in various fields of engineering) to produce an optimal design given the (deterministic) worst-case scenario has been referred to as *optimization and anti-optimization* [35]. The main difference between OUQ and anti-optimization lies in the fact that the former is based on an optimization over (admissible) functions and measures (f, μ) , while the latter only involves an optimization over f . Because of its robustness, many engineers have adopted the (deterministic) worst-case scenario approach to UQ [35, Chap.10] when a high reliability is required.

2.4 Stochastic and Robust Optimization

Robust control [176] and robust optimization [7, 14] have been founded upon the worst-case approach to uncertainty. Recall that robust optimization describes optimization involving uncertain parameters. While these uncertain parameters are modeled as random variables (of known distribution) in stochastic programming [26],

robust optimization only assumes that they are contained in known (ambiguity) sets. Although, as noted in [35], *probabilistic methods do not find appreciation among theoreticians and practitioners alike* because “probabilistic reliability studies involve assumptions on the probability densities, whose knowledge regarding relevant input quantities is central,” the deterministic worst-case approach (limited to optimization problems over f) is sometimes “too pessimistic to be practical” [30, 35] because “it does not take into account the *improbability* that (possibly independent or weakly correlated) random variables conspire to produce a failure event” [114] (which constitutes one motivation for considering ambiguity sets involving both measures and functions). Therefore OUQ and *distributionally robust optimization* (DRO) [7, 14, 49, 53, 166, 174, 177] could be seen as middle ground between the deterministic worst-case approach of robust optimization [7, 14] and approaches of *stochastic programming* and *chance-constrained optimization* [19, 24] by defining optimization objectives and constraints in terms of expected values and probabilities with respect to imperfectly known distributions.

Although stochastic optimization has different objectives than OUQ and DRO, many of its optimization results, such as those found in Birge and Wets [16], and Ermoliev [36] and Gaivoronski, [44], are useful. In particular, the well-developed subject of Edmundson and Madansky bounds such as Edmundson [34]; Madansky [91, 92]; Gassman and Ziemba [45]; Huang, Ziemba, and Ben-Tal [57]; Frauendorfer [41]; Ben-Tal and Hochman [8]; Huang, Vertinsky, and Ziemba [56]; and Kall [67] provide powerful results. Recently Hanasusanto, Roitch, Kuhn, and Wiesemann [53] derive explicit conic reformulations for tractable problem classes and suggest efficiently computable conservative approximations for intractable ones. In some cases, e.g., Bertsimas and Popescu’s [15] and Han et al. [52], DRO/OUQ optimization problems can be reduced to convex optimization.

2.5 Čebyšev Inequalities and Optimization Theory

As noted in [114], inequalities (4) can be seen as a generalization of Čebyšev inequalities. The history of classical inequalities can be found in [70], and some generalizations in [15] and [150]; in the latter works, the connection between Čebyšev inequalities and optimization theory is developed based on the work of Mulholland and Rogers [98], Godwin [48], Isii [60–62], Olkin and Pratt [106], Marshall and Olkin [94], and the classical Markov–Krein theorem [70, pages 82 & 157], among others. We also refer to the field of majorization, as discussed in Marshall and Olkin [95], the inequalities of Anderson [5], Hoeffding [54], Joe [64], Bentkus et al. [12], Bentkus [10, 11], Pinelis [118, 119], and Boucheron, Lugosi, and Massart [20]. Moreover, the solution of the resulting nonconvex optimization problems benefit from duality theories for nonconvex optimization problems such as Rockafellar [123] and the development of convex envelopes for them, as can be found, for example, in Rikun [122] and Sherali [131].

3 The UQ Problem with Sample Data

3.1 From Game Theory to Decision Theory

To motivate the general formulation in the presence of sample data, consider another simple warm-up problem.

Problem 3. Let \mathcal{A} be the set of measures of probability on $[0, 1]$ having mean less than $m \in (0, 1)$. Let μ^\dagger be an unknown element of \mathcal{A} and let $a \in (m, 1)$. You observe $d := (d_1, \dots, d_n)$, n i.i.d. samples from μ^\dagger . What is the *sharpest* estimate of $\mu^\dagger[X \geq a]$?

The only difference between Problems 3 and 1 lies in the availability of data sampled from the underlying unknown distribution. Observe that, in presence of this sample data, the notions of *sharpest estimate* or *smallest interval of confidence* are far from being transparent and call for clear and precise definitions. Note also that if the constraint $\mathbb{E}_{\mu^\dagger}[X] \leq m$ is ignored, and the number n of sample data is large, then one could use the central limit theorem or a concentration inequality (such as Hoeffding's inequality) to derive an interval of confidence for $\mu^\dagger[X \geq a]$. A nontrivial question of practical importance is what to do when n is not large.

Writing $\Phi(\mu^\dagger) := \mu^\dagger[X \geq a]$ as the quantity of interest, observe that an estimation of $\Phi(\mu^\dagger)$ is a function (which will be written θ) of the data d . Ideally one would like to choose θ so that the estimation error $\theta(d) - \Phi(\mu^\dagger)$ is as close as possible to zero. Since d is random, a more robust notion of error is that of a statistical error $\mathcal{E}(\theta, \mu^\dagger)$ defined by weighting the error with respect to a real measurable positive loss function $V: \mathbb{R} \rightarrow \mathbb{R}$ and the distribution of the data, i.e.,

$$\mathcal{E}(\theta, \mu^\dagger) := \mathbb{E}_{d \sim (\mu^\dagger)^n} [V[\theta(d) - \Phi(\mu^\dagger)]] \quad (5)$$

Note that for $V(x) = x^2$, the statistical error $\mathcal{E}(\theta, \mu^\dagger)$ defined in (5) is the mean squared error with respect to the distribution of the data d of the estimation error. For $V(x) = \mathbb{1}_{[\gamma, \infty)}(|x|)$ defined for some $\gamma > 0$, $\mathcal{E}(\theta, \mu^\dagger)$ is the probability with respect to the distribution of d that the estimation error is larger than γ .

Now since μ^\dagger is unknown, the statistical error $\mathcal{E}(\theta, \mu^\dagger)$ of any θ is also unknown. However one can still identify the least upper bound on that statistical error through a worst-case scenario with respect to all possible candidates for μ^\dagger , i.e.,

$$\sup_{\mu \in \mathcal{A}} \mathcal{E}(\theta, \mu). \quad (6)$$

The sharpest estimator (possibly within a given class) is then naturally obtained as the minimizer of (6) over all functions θ of the data d within that class, i.e., as the minimizer of

$$\inf_{\theta} \sup_{\mu \in \mathcal{A}} \mathcal{E}(\theta, \mu). \quad (7)$$

Observe that the optimal estimator is identified independently from the observation/realization of the data and if the minimum of (7) is not achieved then one can still use a near-optimal θ . Then, when the data is observed, the estimate of the quantity of interest $\Phi(\mu^\dagger)$ is then derived by evaluating the near-optimal θ on the data d . The notion of optimality described here is that of Wald’s statistical decision theory [156–158, 160, 161], evidently influenced by Von Neumann’s game theory [153, 155]. In Wald’s formulation [157], which cites both Von Neumann [153] and Von Neumann and Morgenstern [155], the statistician finds himself in an adversarial game played against the Universe in which he tries to minimize a risk function $\mathcal{E}(\theta, \mu)$ with respect to θ in a worst-case scenario with respect to what the Universe’s choice of μ could be.

3.2 The Optimization Approach to Statistics

The optimization approach to statistics is not new and this section will now give a short, albeit incomplete, description of its development, primarily using Lehmann’s account [87]. Accordingly, it began with Gauss and Laplace with the nonparametric result referred to as the Gauss-Markov theorem, asserting that the least squares estimates are the linear unbiased estimates with minimum variance. Then, in Fisher’s fundamental paper [39], for parametric models, he proposes the maximum likelihood estimator and claims (but does not prove) that such estimators are consistent and asymptotically efficient. According to Lehmann, “the situation is complex, but under suitable restrictions Fisher’s conjecture is essentially correct . . .” The Fisher’s maximum likelihood principle was first proposed on intuitive grounds and then its optimality properties developed. However, according to Lehmann [86, Pg. 1011], Pearson then asked Neyman “Why these tests rather than any of the many other that could be proposed? This question resulted in Neyman and Pearson’s 1928 papers [104] on the likelihood ratio method, which gives the same answer as Fisher’s tests under normality assumptions. However, Neyman was not satisfied. He agreed that the likelihood ratio principle was appealing but felt that it was lacking a logically convincing justification. This then led to the publication of Neyman and Pearson [105], containing their now famous Neyman–Pearson lemma, which according to Lehmann [87], “In a certain sense this is the true start of optimality theory.” In a major extension of the Neyman–Pearson work, Huber [58] proves a *robust* version of the Neyman–Pearson lemma, in particular, providing an optimality criteria defining the robust estimator, giving rise to a rigorous theory of robust statistics based on optimality; see Huber’s Wald lecture [59]. Robustness is particularly suited to the Wald framework since robustness considerations are easily formulated with the proper choices of admissible functions and measures in the Wald framework. Another example is Kiefer’s introduction of optimality into experimental design, resulting in Kiefer’s 40 papers on Optimum Experimental Designs [74].

Not everyone was happy with “optimality” as a guiding principle. For example, Lehmann [87] states that at a 1958 meeting of the Royal Statistical Society at which Kiefer presented a survey talk [73] on Optimum Experimental Designs, Barnard quotes Kiefer as saying of procedures proposed in a paper by Box and Wilson that they “often [are] not even well-defined rules of operation.” Barnard’s reply:

in the field of practical human activity, rules of operation which are not well-defined may be preferable to rules which are.

Wynn [173], in his introduction to a reprint of Kiefer’s paper, calls this “a clash of statistical cultures.” Indeed, it is interesting to read the generally negative responses to Kiefer’s article [73] and the remarkable rebuttal by Kiefer therein. Tukey had other criticisms regarding “The tyranny of the best” in [147] and “The dangers of optimization” in [148]. In the latter he writes:

Some [statisticians] seem to equate [optimization] to statistics an attitude which, if widely adopted, is guaranteed to produce a dried-up, encysted field with little chance of real growth.

For an account of how the Exploratory Data Analysis approach of Tukey fits within the Fisher/Neyman–Pearson debate, see Lehmann [88].

Let us also remark on the influence that Student – William Sealy Gosset – had on both Fisher and Pearson. As presented in Lehmann’s [85] “‘Student’ and small-sample theory,” quoting F. N. David [79]: “I think he [Gosset] was really the big influence in statistics. . . He asked the questions and Pearson or Fisher put them into statistical language and then Neyman came to work with the mathematics. But I think most of it stems from Gosset.” The aim of Lehmann’s paper [85] is to consider to what extent David’s conclusion is justified. Indeed, the claim is surprising since Gosset is mainly known for only one contribution, that is, Student [141], with the introduction of Student’s *t*-test and its analysis under the normal distribution. According to Lehmann, “Today the pathbreaking nature of this paper is generally recognized and has been widely commented upon, . . .” Gosset’s primary concern in communicating with both Fisher and Pearson was the robustness of the test to non-normality. Lehmann concludes that “the main ideas leading to Pearson’s research were indeed provided by Student.” See Lehmann [85] for the full account, including Gosset’s relationship to the Fisher–Pearson debate, Pearson [116] for a statistical biography of Gosset, and Fisher [40] for a eulogy. Consequently, modern statistics appears to owe a lot to Gosset. Moreover, the reason for the pseudonym was a policy by Gosset’s employer, the brewery Arthur Guinness, Sons, and Co., against work done for the firm being made public. Allowing Gosset to publish under a pseudonym was a concession that resulted in the birth of the statistician Student. Consequently, the authors would like to take this opportunity to thank the Guinness Brewery for its influence on statistics today, and for such a fine beer.

3.3 Abraham Wald

Following Neyman and Pearson’s breakthrough, a different approach to optimality was introduced in Wald [156] and then developed in a sequence of papers

culminating in Wald's [161] book *Statistical Decision Functions*. Evidence of the influence of Neyman on Wald can be found in the citation of Neyman [102] in the introduction of Wald [156]. Brown [22] quotes the students of Neyman in 1966 from [103]:

The concepts of confidence intervals and of the Neyman-Pearson theory have proved immensely fruitful. A natural but far reaching extension of their scope can be found in Abraham Wald's theory of statistical decision functions. The elaboration and application of the statistical tools related to these ideas has already occupied a generation of statisticians. It continues to be the main lifestream of theoretical statistics.

Brown's purpose was to address if the last sentence in the quote was still true in 2000.

Wolfowitz [170] describes the primary accomplishments of Wald's statistical decision theory as follows:

Wald's greatest achievement was the theory of statistical decision functions, which includes almost all problems which are the *raison d'être* of statistics.

Leonard [89, Chp. 12] portrays Von Neumann's return to game theory as "partly an early reaction to upheaval and war." However he adds that eventually Von Neumann became personally involved in the war effort and "with that involvement came a significant, unforeseeable moment in the history of game theory: this new mathematics made its wartime entrance into the world, not as the abstract theory of social order central to the book, but as a problem solving technique." Moreover, on pages 278–280, Leonard discusses the statistical research groups at Berkeley, Columbia, and Princeton, in particular Wald at Columbia, and how the effort to develop inspection and testing procedures leads Wald to the development of sequential methods "apparently yielding significant economies in inspection in the Navy," leading to the publication of Wald and Wolfowitz' [162] proof of the optimality of the sequential probability ratio test and Wald's book [159] *Sequential Analysis*. Leonard's claim essentially is that the war stimulated these fine theoretical minds to pursue activities with real application value. In this regard, it is relevant to note Mangel and Samaniego's [93] stimulating description of Wald's work on aircraft survivability, along with the contemporary, albeit somewhat vague, description of "How a Story from World War II shapes Facebook today" by Wilson [167]. Indeed, in the problem of how to allocate armoring to the allied bombers based on their condition upon return from their missions, it was discovered that armoring where the previous planes had been hit was not improving their rate of return. Wald's ingenious insight was that these were the returning bombers not the ones which had been shot down. So the places where the returning bombers were hit are more likely to be the places where one *does not* need to add armoring. Evidently, his rigorous and unconventional innovations to transform this intuition into a real methodology saved many lives. Wolfowitz [170] states:

Wald not only posed his statistical problems clearly and precisely, but he posed them to fit the practical problem and to accord with the decisions the statistician was called on to make. This, in my opinion, was the key to his success—a high level of mathematical talent of the most abstract sort, and a true feeling for, and insight into, practical problems. The combination of the two in his person at such high levels was what gave him his outstanding character.

The section on Von Neumann and Remak (along with the Notes that follows it) in Kurz and Salvadori [78] describes Wald and Von Neumann's relations. Brown [21] credits Wald as the creator of the minmax idea in statistics, evidently given axiomatic justification by Gilboa and Schmeidler [47]. This certainly had something to do with his friendship with Morgenstern and his relationship with Von Neumann, who together authored the famous book [155], but this influence can be explicitly seen in Wald's [157] citation of Von Neumann [153] and Von Neumann and Morgenstern [155] in his introduction [157] of the minmax idea in statistical decision theory. Indeed, Wolfowitz states that:

... he was also spurred on by the connection between the newly announced results of [Von Neumann and Morgenstern] [155] and his own theory, and by the general interest among economists and others aroused by the theory of games.

Wolfowitz asserts that Wald's work [156] Contributions to the Theory of Statistical Estimation and Testing Hypotheses is "probably his most important paper" but that it "went almost completely unnoticed," possibly because "The use of Bayes solutions was deterrent" and "Wald did not really emphasize that he was using Bayes solutions only as a tool." Moreover, although Wolfowitz considered Wald's Statistical Decision Functions [161] his greatest achievement, he also says:

The statistician who wants to apply the results of [161] to specific problems is likely to be disappointed. Except for special problems, the complete classes are difficult to characterize in a simple manner and have not yet been characterized. Satisfactory general methods are not yet known for obtaining minimax solutions. If one is not always going to use a minimax solution (to which serious objections have been raised) or a solution satisfying some given criterion, then the statistician should have the opportunity to choose from among "representative" decision functions on the basis of their risk functions. These are not available except for the simplest cases. It is clear that much remains to be done before the use of decision functions becomes common. The theory provides a rational basis for attacking almost any statistical problem, and, when some computational help is available and one makes some reasonable compromises in the interest of computational feasibility, one can obtain a practical answer to many problems which the classical theory is unable to answer or answers in an unsatisfactory manner.

Wolfowitz [170], Morgenstern [97], and Hotelling [55] provide a description of Wald's impact at the time of his passing. The influence of Wald's minimax paradigm can also be observed on (1) decision making under severe uncertainty [134–136], (2) stochastic programming [130] (minimax analysis of stochastic problems), (3) minimax solutions of stochastic linear programming problems [175], (4) robust convex optimization [9] (where one must find the best decision in view of the worst-case parameter values within deterministic uncertainty sets), (4) econometrics [143], and (5) Savage's minimax regret model [128].

3.4 Generalization to Unknown Pairs of Functions and Measures and to Arbitrary Sample Data

In practice, complex systems of interest may involve, both an imperfectly known response function f^\dagger and an imperfectly known probability measure μ^\dagger as illustrated in the following problem.

Problem 4. Let \mathcal{A} be a set of real functions and measures of probability on $[0, 1]$ such that $(f, \mu) \in \mathcal{A}$ if and only if $\mathbb{E}_\mu[X] \leq m$ and $\sup_{x \in [0, 1]} |g(x) - f(x)| \leq 0.1$ where g is some given real function on $[0, 1]$. Let (f^\dagger, μ^\dagger) be an unknown element of \mathcal{A} and let $a \in \mathbb{R}$. Let (X_1, \dots, X_n) be n i.i.d. samples from μ^\dagger , you observe (d_1, \dots, d_n) with $d_i = (X_i, f^\dagger(X_i))$. What is the “sharpest” estimate of $\mu^\dagger[f(X) \geq a]$?

Problem 4 is an illustration of a situation in which the response function f^\dagger and the probability measure μ^\dagger are not directly observed and the sample data arrives in the form of realizations of random variables, the distribution of which is related to (f^\dagger, μ^\dagger) . To simplify the current presentation, assume that this relation is, in general, determined by a function of (f^\dagger, μ^\dagger) and use the following notation: \mathcal{D} will denote the observable space (i.e., the space in which the sample data d take values, assumed to be a metrizable Suslin space) and d will denote the \mathcal{D} -valued random variable corresponding to the observed sample data. To represent the dependence of the distribution of d on the unknown state $(f^\dagger, \mu^\dagger) \in \mathcal{A}$, introduce a measurable function

$$\mathbb{D}: \mathcal{A} \rightarrow \mathcal{M}(\mathcal{D}), \quad (8)$$

where $\mathcal{M}(\mathcal{D})$ is given the Borel structure corresponding to the weak topology, to define this relation. The idea is that $\mathbb{D}(f, \mu)$ is the probability distribution of the observed sample data d if $(f^\dagger, \mu^\dagger) = (f, \mu)$, and for this reason it may be called the *data map* (or, even more loosely, the *observation operator*). Now consider the following problem.

Problem 5. Let \mathcal{A} be a known subset of $\mathcal{F}(\mathcal{X}) \times \mathcal{M}(\mathcal{X})$ as in Problem 2 and let \mathbb{D} be a known data map as in (8). Let Φ be a known measurable semi-bounded function mapping \mathcal{A} onto \mathbb{R} . Let (f^\dagger, μ^\dagger) be an unknown element of \mathcal{A} . You observe $d \in \mathcal{D}$ sampled from the distribution $\mathbb{D}(f^\dagger, \mu^\dagger)$. What is the sharpest estimation of $\Phi(f^\dagger, \mu^\dagger)$?

3.5 Model Error and Optimal Models

As in Wald’s statistical decision theory [157], a natural notion of optimality can be obtained by formulating Problem 5 as an adversarial game in which player A chooses $(f^\dagger, \mu^\dagger) \in \mathcal{A}$ and player B (knowing \mathcal{A} and \mathbb{D}) chooses a function θ of the observed data d . As in (5) this notion of optimality requires the introduction of a risk function:

$$\mathcal{E}(\theta, (f, \mu)) := \mathbb{E}_{d \sim \mathbb{D}(f, \mu)} \left[V[\theta(d) - \Phi(f, \mu)] \right] \quad (9)$$

where $V: \mathbb{R} \rightarrow \mathbb{R}$ is a real positive measurable loss function. As in (6) the least upper bound on that statistical error $\mathcal{E}(\theta, (f, \mu))$ is obtained as through a worst-case scenario with respect to all possible candidates for (f, μ) (player's A choice), i.e.,

$$\sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}(\theta, (f, \mu)) \quad (10)$$

and an optimal estimator/model (possibly within a given class) is then naturally obtained as the minimizer of (10) over all functions θ of the data d in that class (player's B choice), i.e., as the minimizer of

$$\inf_{\theta} \sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}(\theta, (f, \mu)). \quad (11)$$

Since in real applications true optimality will never be achieved, it is natural to generalize to considering near-minimizers of (11) as near-optimal models/estimators.

Remark 1. In situations where the data map is imperfectly known (e.g., when the data d is corrupted by some noise of imperfectly known distribution), one has to include a supremum over all possible candidates $\mathbb{D} \in \mathfrak{D}$ in the calculation of the least upper bound on the statistical error.

3.6 Mean Squared Error, Variance, and Bias

For $(f, \mu) \in \mathcal{A}$ write $\text{Var}_{d \sim \mathbb{D}(f, \mu)} [\theta(d)]$ the variance of the random variable $\theta(d)$ when d is distributed according to $\mathbb{D}(f, \mu)$, i.e.,

$$\text{Var}_{d \sim \mathbb{D}(f, \mu)} [\theta(d)] := \mathbb{E}_{d \sim \mathbb{D}(f, \mu)} [(\theta(d))^2] - \left[\mathbb{E}_{d \sim \mathbb{D}(f, \mu)} [\theta(d)] \right]^2$$

The following equation, whose proof is straightforward, shows that for $V(x) = x^2$, the least upper bound on the mean squared error of θ is equal to the least upper bound on the sum of the variance of θ and the square of its bias:

$$\sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}(\theta, (f, \mu)) = \sup_{(f, \mu) \in \mathcal{A}} \left[\text{Var}_{d \sim \mathbb{D}(f, \mu)} [\theta(d)] + \left(\mathbb{E}_{d \sim \mathbb{D}(f, \mu)} [\theta(d)] - \Phi(f, \mu) \right)^2 \right]$$

Therefore, for $V(x) = x^2$, the bias/variance tradeoff is made explicit.

3.7 Optimal Interval of Confidence

Although \mathcal{E} can a priori be defined to be any risk function, taking $V(x) = \mathbb{1}_{[\gamma, \infty]}(|x|)$ (for some $\gamma > 0$) in (5) allows for a transparent and objective identification of optimal intervals of confidence. Indeed, writing,

$$\mathcal{E}_\gamma(\theta, (f, \mu)) := \mathbb{P}_{d \sim \mathbb{D}(f, \mu)} \left[|\theta(d) - \Phi(f, \mu)| \geq \gamma \right]$$

note that $\sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}_\gamma(\theta, (f, \mu))$ is the least upper bound on the probability (with respect to the distribution of d) that the difference between the true value of the quantity of interest $\Phi(f^\dagger, \mu^\dagger)$ and its estimated value $\theta(d)$ is larger than γ . Let $\epsilon \in [0, 1]$. Define

$$\gamma_\epsilon := \inf \left\{ \gamma > 0 \mid \inf_\theta \sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}_\gamma(\theta, (f, \mu)) \leq \epsilon \right\},$$

and observe that if θ_ϵ is a minimizer of $\inf_\theta \sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}_{\gamma_\epsilon}(\theta, (f, \mu))$ then $[\theta_\epsilon(d) - \gamma_\epsilon, \theta_\epsilon(d) + \gamma_\epsilon]$ is the smallest interval of confidence (random interval obtained as a function of the data) containing $\Phi(f^\dagger, \mu^\dagger)$ with probability at least $1 - \epsilon$. Observe also that this formulation is a natural extension of the OUQ formulation as described in [114]. Indeed, in the absence of sample data, it is easy to show that θ_1 is the midpoint of the optimal interval $[\mathcal{L}(\mathcal{A}), \mathcal{U}(\mathcal{A})]$.

Remark 2. We refer to [37, 38, 137] and in particular to Stein's notorious paradox [138] for the importance of a careful choice for loss function.

3.8 Ordering the Space of Experiments

A natural objective of UQ and statistics is the design of experiments, their comparisons, and the identification of optimal ones. Introduced in Blackwell [17] and Kiefer [73], with a more modern perspective in Le Cam [83] and Strasser [139], here observe that (11), as a function of \mathbb{D} , induces an order (transitive, total, but not antisymmetric) on the space of data maps that has a natural experimental design interpretation. More precisely if the data maps \mathbb{D}_1 and \mathbb{D}_2 are interpreted as the distribution of the outcome of two possible experiments, and if the value of (11) is smaller for \mathbb{D}_2 than \mathbb{D}_1 , then \mathbb{D}_2 is a preferable experiment.

3.9 Mixing Models

Given estimators $\theta_1, \dots, \theta_n$ can one obtain a better estimator by mixing those estimators? If V is convex (or quasi-convex), then the problem of finding an $\alpha \in [0, 1]^n$ minimizing the statistical error of $\sum_{i=1}^n \alpha_i \theta_i$ under the constraint $\sum_{i=1}^n \alpha_i = 1$ is a finite-dimensional convex optimization problem in α . If estimators are seen as models of reality, then this observation supports the idea that one can obtain improved models by mixing them (with the goal of achieving minimal statistical errors).

4 The Complete Class Theorem and Bayesian Inference

4.1 The Bayesian Approach

The Bayesian answer to Problem 5 is to assume that (f^\dagger, μ^\dagger) is a sample from some (prior) measure π on \mathcal{A} and then condition the expectation of $\Phi(f, \mu)$ with respect to the observation of the data, i.e., use

$$\mathbb{E}_{(f,\mu) \sim \pi, d \sim \mathbb{D}(f,\mu)} [\Phi(f, \mu) | d] \quad (12)$$

as the estimator $\theta(d)$. This requires giving \mathcal{A} the structure of a measurable space such that important quantities of interest such as $(f, \mu) \rightarrow \mu[f(X) \geq a]$ and $(f, \mu) \rightarrow \mathbb{E}_\mu[f]$ are measurable. This can be achieved using results of Ressel [121] providing conditions under which the map $(f, \mu) \rightarrow f_*\mu$ from function/measure pairs to the induced law is Borel measurable. We will henceforth assume \mathcal{A} to be a Suslin space and proceed to construct the measure of probability $\pi \odot \mathbb{D}$ of $((f, \mu), d)$ on $\mathcal{A} \times \mathcal{D}$ via a natural generalization of the Campbell measure and Palm distribution associated with a random measure as described in [68]; see also [25, Ch. 13] for a more current treatment. We refer to Sect. 6 of the appendix for the details of the construction of the distribution $\pi \odot \mathbb{D}$ of $((f, \mu), d)$ when $(f, \mu) \sim \pi$ and $d \sim \mathbb{D}(f, \mu)$, and of the marginal distribution $\pi \cdot \mathbb{D}$ of $\pi \odot \mathbb{D}$ on the data space \mathcal{D} , and the resulting regular conditional expectation (12). Consequently, the nested expectation $\mathbb{E}_{(f,\mu) \sim \pi, d \sim \mathbb{D}(f,\mu)}$ appearing in (12) will from now on be rigorously written as the expectation $\mathbb{E}_{((f,\mu),d) \sim \pi \odot \mathbb{D}}$.

Statistical error when (f^\dagger, μ^\dagger) is random. When (f^\dagger, μ^\dagger) is a random realization of π^\dagger , one may consider averaging the statistical error (9) with respect to π^\dagger and introduce

$$\mathcal{E}(\theta, \pi^\dagger) := \mathbb{E}_{((f,\mu),d) \sim \pi^\dagger \odot \mathbb{D}} [V[\theta(d) - \Phi(f, \mu)]] \quad (13)$$

When π^\dagger is an unknown element of a subset Π of $\mathcal{M}(\mathcal{A})$, the least upper bound on the statistical error (13) given the available information is obtained by taking the sup of (13) with respect to all possible candidates for π^\dagger , i.e.,

$$\sup_{\pi \in \Pi} \mathcal{E}(\theta, \pi) \quad (14)$$

When \mathcal{A} is Suslin and when (f^\dagger, μ^\dagger) is not a random sample from π^\dagger but simply an unknown element of \mathcal{A} , then a straightforward application of the reduction theorems of [114] implies that when $\Pi = \mathcal{M}(\mathcal{A})$, then (14) is equal to (11), i.e.,

$$\sup_{(f,\mu) \in \mathcal{A}} \mathcal{E}(\theta, (f, \mu)) = \sup_{\pi \in \mathcal{M}(\mathcal{A})} \mathcal{E}(\theta, \pi) \quad (15)$$

4.2 Relation Between Adversarial Model Error and Bayesian Error

When Φ has a second moment with respect to π , one can utilize the classical variational description of conditional expectation as follows: Letting $L^2_{\pi \cdot \mathbb{D}}(\mathcal{D})$ denote the space of ($\pi \cdot \mathbb{D}$ a.e. equivalence classes of) real-valued measurable functions on \mathcal{D} that are square-integrable with respect to the measure $\pi \cdot \mathbb{D}$, one has (see Sect. 6)

$$\mathbb{E}_{\pi \odot \mathbb{D}}[\Phi|d] := \arg \min_{h \in L^2_{\pi \cdot \mathbb{D}}(\mathcal{D})} \mathbb{E}_{(f, \mu, d) \sim \pi \odot \mathbb{D}} \left[(\Phi(f, \mu) - h(d))^2 \right].$$

In other words, if (f, μ) is sampled from the measure π , $\mathbb{E}_{\pi \odot \mathbb{D}}[\Phi(f, \mu)|d]$ is the best mean-square approximation of $\Phi(f, \mu)$ in the space of square-integrable functions of d . As with the regular conditional probabilities, the real-valued function on \mathcal{D}

$$\theta_\pi(d) = \mathbb{E}_{(f, \mu, D) \sim \pi \odot \mathbb{D}}[\Phi(f, \mu)|D = d], \quad d \in \mathcal{D} \quad (16)$$

is uniquely defined up to subsets of \mathcal{D} of $(\pi \cdot \mathbb{D})$ -measure zero.

Using the orthogonal projection property of the conditional expectation, one obtains that if $V(x) = x^2$, then for arbitrary θ ,

$$\mathcal{E}(\theta, \pi) = \mathcal{E}(\theta_\pi, \pi) + \mathbb{E}_{d \sim \pi \cdot \mathbb{D}} \left[\theta(d) - \theta_\pi(d) \right]^2 \quad (17)$$

Therefore, if $\Pi \subset \mathcal{M}(\mathcal{A})$ is an admissible set of priors, then (17) implies that

$$\inf_{\theta} \sup_{\pi \in \Pi} \mathcal{E}(\theta, \pi) \geq \sup_{\pi \in \Pi} \mathcal{E}(\theta_\pi, \pi).$$

In particular, when $\Pi = \mathcal{M}(\mathcal{A})$ (15) implies that

$$\inf_{\theta} \sup_{(f, \mu) \in \mathcal{A}} \mathcal{E}(\theta, (f, \mu)) \geq \sup_{\pi \in \mathcal{M}(\mathcal{A})} \mathcal{E}(\theta_\pi, \pi). \quad (18)$$

Therefore, the mean squared error of the best estimator assuming $(f^\dagger, \mu^\dagger) \in \mathcal{A}$ to be unknown is bounded below by the largest mean squared error of the Bayesian estimator obtained by assuming that (f^\dagger, μ^\dagger) is distributed according to some $\pi \in \mathcal{M}(\mathcal{A})$. In the next section, it will be shown that a complete class theorem can be used to obtain that (18) is actually an equality. In that case, (18) can be used to quantify the approximate optimality of an estimator by comparing the least upper bound $\sup_{(f, \mu) \in \mathcal{M}(\mathcal{A})} \mathcal{E}(\theta, (f, \mu))$ on the error of that estimator with $\mathcal{E}(\theta_\pi, \pi)$ for a carefully chosen π .

4.3 Complete Class Theorem

A fundamental question is whether (18) is an equality: is the adversarial error of the best estimator equal to the non-adversarial error of the worst Bayesian estimator? Is the best estimator Bayesian or an approximation thereof? A remarkable result of decision theory [156–158, 160, 161] is the complete class theorem which states (in the formulation of this paper) that if (1) the admissible measures μ are absolutely continuous with respect the Lebesgue measure, (2) the loss function V in the definition of $\mathcal{E}(\theta, (f, \mu))$ is convex in θ and (3) the decision space is compact, then optimal estimators live in the Bayesian class and non-Bayesian estimators cannot be optimal. The idea of the proof of this result is to use the compactness of the decision space and the continuity of the loss function to approximate the decision theory game by a finite game and recall that optimal strategies of adversarial finite zero-sum games are mixed strategies [99, 100].

Le Cam [81], see also [83], has substantially extended Wald’s theory in the sense that requirements of boundedness, or even finiteness, of the loss function are replaced by a requirement of lower semicontinuity, and the requirements of the compactness of the decision space and the absolute continuity of the admissible measures with respect the Lebesgue measure are removed. These vast generalizations come at some price of abstraction yet reveal the relevance and utility of an appropriate complete Banach lattice of measures. In particular, this framework of Le Cam appears to facilitate efficient concrete approximation.

As an illustration, let us describe a complete class theorem on a space of admissible measures, without the inclusion of functions, where the observation map consists of taking n -i.i.d. samples, as in Eq. (5). Let $\mathcal{A} \subset \mathcal{M}(\mathcal{X})$ be a subset of the Borel probability measures on a topological space \mathcal{X} and consider a quantity of interest $\Phi : \mathcal{A} \rightarrow \mathbb{R}$. For $\mu \in \mathcal{A}$, the data d is generated by i.i.d. sampling with respect to μ^n . That is $d \sim \mu^n$. For $\mu^\dagger \in \mathcal{A}$, the statistical error $\mathcal{E}(\theta, \mu^\dagger)$ of an estimator $\theta : \mathcal{X}^n \rightarrow \mathbb{R}$ of $\Phi(\mu^\dagger)$ is defined in terms of a loss function $V : \mathbb{R} \rightarrow \mathbb{R}$ as in (5). Define the least upper bound on that statistical error and the sharpest estimator as in (6) and (7).

Let $\Theta := \{\theta : \mathcal{X}^n \rightarrow \mathbb{R}, \theta \text{ measurable}\}$ denote the space of estimators. Since, in general, the game $\mathcal{E}(\theta, \mu), \theta \in \Theta, \mu \in \mathcal{A}$ will not have a value, that is, one will have a strict inequality:

$$\sup_{\mu \in \mathcal{A}} \inf_{\theta \in \Theta} \mathcal{E}(\theta, \mu) < \inf_{\theta \in \Theta} \sup_{\mu \in \mathcal{A}} \mathcal{E}(\theta, \mu),$$

classical arguments in game theory suggest that one extend to random estimators and random selection in \mathcal{A} . To that end, let the set of randomized estimators $\mathcal{R} := \{\hat{\theta} : \mathcal{X}^n \times \mathcal{B}(\mathbb{R}) \rightarrow [0, 1], \hat{\theta} \text{ Markov}\}$ be the set of Markov kernels. To define a topology for \mathcal{R} , define a linear space of measures as follows. Let $\mathcal{A}^n := \{\mu^n \in \mathcal{M}(\mathcal{X}^n) : \mu \in \mathcal{A}\}$ denote the corresponding set of measures generating sample data. Say that \mathcal{A}^n is dominated if there exists an $\omega \in \mathcal{M}(\mathcal{X}^n)$ such that every $\mu^n \in \mathcal{A}^n$ is absolutely continuous with respect to ω . According to

the Halmos–Savage lemma [50], see also Strasser [139, Lem. 20.3], the set \mathcal{A}^n is dominated if and only if there exists a countable mixture $\mu^* := \sum_{i=1}^{\infty} \alpha_i \mu_i^n$, with $\alpha_i \geq 0, \mu_i \in \mathcal{A}, i = 1, \dots, \infty$, and $\sum_{i=1}^{\infty} \alpha_i = 1$, such that $\mu^n \sim \mu^*, \mu \in \mathcal{A}$. A construct at the heart of Le Cam’s approach is a natural linear space notion of a mixture space of \mathcal{A} , called the L -space $L(\mathcal{A}^n) := L^1(\mu^*)$. It follows easily, see [139, Lem. 41.1], that $L(\mathcal{A}^n)$ is the set of signed measures which are absolutely continuous with respect to μ^* . When \mathcal{A} is not dominated, a natural generalization of this construction [139, Def. 41.3] due to Le Cam [81] is used. A crucial property of the L -space $L(\mathcal{A}^n)$ is that not only is it a Banach lattice (see Strasser [139, Cor. 41.4]), but by [139, Lem. 41.5] it is a complete lattice. The utility of the concept of a complete lattice to the complete class theorems can clearly be seen in the proof of the lemma in Section 2 of Wald and Wolfowitz’ [163] proof of the complete class theorem when the number of decisions and the number of distributions is finite. Then, the natural action of a randomized estimator on the bounded continuous function/mixture pairs $C_b(\mathbb{R}) \times L(\mathcal{A}^n)$ is

$$f \hat{\theta} v := \int \int f(r) \hat{\theta}(x^n, dr) v(dx^n), \quad f \in C_b(\mathbb{R}), v \in L(\mathcal{A}^n).$$

Let \mathcal{R} be endowed with the topology of pointwise convergence with respect to this action, that is, the weak topology with respect to integration against $C_b(\mathbb{R}) \times L(\mathcal{A}^n)$. Moreover, this weak topology also facilitates a definition of the space $\overline{\mathcal{R}}$ of *generalized* random estimators as bilinear real-valued maps $\vartheta : C_b(\mathbb{R}) \times L(\mathcal{A}^n) \rightarrow \mathbb{R}$ satisfying $|\vartheta(f, \mu)| \leq \|f\|_{\infty} \|\mu\|$, $\vartheta(f, \mu) \geq 0$ for $f \geq 0, \mu \geq 0$, and $\vartheta(1, \mu) = \mu(\mathcal{X})$. By [139, Thm. 42.3], the set of generalized random estimators $\overline{\mathcal{R}}$ is compact and convex, and by [139, Thm. 42.5] of Le Cam [82], \mathcal{R} is dense in $\overline{\mathcal{R}}$ in the weak topology. Moreover, when \mathcal{A}^n is dominated and one can restrict to a compact subset $C \in \mathbb{R}$ of the decision space, then Strasser [139, Cor. 42.8] asserts that $\overline{\mathcal{R}} = \mathcal{R}$.

Returning to our illustration, if one let $W_{\mu}, \mu \in \mathcal{A}$ be defined by $W_{\mu}(r) := V(r - \Phi(\mu)), r \in \mathbb{R}, \mu \in \mathcal{A}$ denote the associated family of loss functions, one can now define a generalization of the statistical error function $\mathcal{E}(\theta, \mu)$ of (5) to randomized estimators $\hat{\theta}$ by

$$\mathcal{E}(\hat{\theta}, \mu) := \int \int W_{\mu}(r) \hat{\theta}(x^n, dr) \mu^n(dx^n), \quad \hat{\theta} \in \mathcal{R}, \mu \in \mathcal{A}.$$

This definition reduces to the previous one (5) when the random estimator $\overline{\hat{\theta}}$ corresponds to a point estimator θ and extends naturally to $\overline{\mathcal{R}}$. Finally, one says that an estimator $\vartheta^* \in \overline{\mathcal{R}}$ is Bayesian if there exists a probability measure m with finite support on \mathcal{A} such that

$$\int \mathcal{E}(\vartheta^*, \mu) m(d\mu) \leq \int \mathcal{E}(\vartheta, \mu) m(d\mu), \quad \vartheta \in \overline{\mathcal{R}}.$$

The following complete class theorem follows from Strasser [139, Thm.47.9, Cor.42.8] since one can naturally compactify the decision space \mathbb{R} when the quantity of interest Φ is bounded and the loss function V is sublevel compact, that is has compact sublevel sets.

Theorem 1. *Suppose that the loss function V is sublevel compact and the quantity of interest $\Phi : \mathcal{A} \rightarrow \mathbb{R}$ is bounded. Then, for each generalized randomized estimator $\vartheta \in \overline{\mathcal{R}}$, there exists a weak limit $\vartheta^* \in \overline{\mathcal{R}}$ of Bayesian estimators such that*

$$\mathcal{E}(\vartheta^*, \mu) \leq \mathcal{E}(\vartheta, \mu), \quad \mu \in \mathcal{A}.$$

If, in addition, \mathcal{A} is dominated, then there exists such a $\vartheta^ \in \mathcal{R}$.*

A comprehensive connection of these results, where Bayesian estimators are defined only in terms of measures of finite support on \mathcal{A} , with the framework of Sect. 4 where Bayesian estimators are defined in terms of Borel measures on \mathcal{A} , is not available yet. Nevertheless it appears that much can be done in this regard. In particular, one can suspect that when \mathcal{A} is a closed convex set of probability measures equipped with the weak topology and \mathcal{X} is a Borel subset of a Polish space, that if the loss function V is convex and Φ is affine and measurable, the Choquet theory of Winkler [168, 169] can be used to facilitate this connection. Indeed, as mentioned above, complete class theorems are available for much more general loss functions than continuous or convex, more general decision spaces than \mathbb{R} , and without absolute continuity assumptions. Moreover, it is interesting to note that, although randomization was introduced to obtain minmax results, when the loss function V is strictly convex, Bayesian estimators can be shown to be *non-random*. This can be explicitly observed in the definition (16) of Bayesian estimators when $V(x) := x^2$ and is understood much more generally in Dvoretzky, Wald, and Wolfowitz [33]. We conjecture that further simplifications can be obtained by allowing *approximate* versions of complete class theorems, Bayesian estimators, optimality, and saddle points, as in Scovel, Hush, and Steinwart's [129] extension of classical Lagrangian duality theory to include approximations.

5 Incorporating Complexity and Computation

Although decision theory provides well-posed notions of optimality and performance in statistical estimation, it does not address the complexity of the actual computation of optimal or nearly optimal estimators and their evaluation against the data. Indeed, although the abstract identification of an optimal estimator as the solution of an optimization problem provides a clear objective, practical applications require the actual implementation of the estimator on a machine and its numerical evaluation against the data.

5.1 Machine Wald

The simultaneous emphasis on *performance* and *computation* can be traced back to PAC (probably approximately correct) learning initiated by Valiant [149] which has laid down the foundations of machine learning (ML). Indeed, as asserted by Wasserman in his 2013 lecture, “The Rise of the Machines” [164, Sec. 1.5]:

There is another interesting difference that is worth pondering. Consider the problem of estimating a mixture of Gaussians. In Statistics we think of this as a solved problem. You use, for example, maximum likelihood which is implemented by the EM algorithm. But the EM algorithm does not solve the problem. There is no guarantee that the EM algorithm will actually find the MLE; its a shot in the dark. The same comment applies to MCMC methods. In ML, when you say youve solved the problem, you mean that there is a polynomial time algorithm with provable guarantees.

That is, on even par with the rigorous performance analysis, machine learning also requires that solutions be efficiently implementable on a computer, and often such efficiency is established by proving bounds on the amount of computation required to produce such a solution with a given algorithm. Although Wald’s theory of optimal statistical decisions has resulted in many important statistical discoveries, looking through the three Lehmann symposia of Rojo and Pérez–Abreu [126] in 2004 and Rojo [124, 125] in 2006 and 2009, it is clear that the incorporation of the analysis of the computational algorithm, both in terms of its computational efficiency and its statistical optimality, has not begun. Therefore a natural answer to fundamental challenges in UQ appears to be the full incorporation of *computation* into a natural generalization of Wald’s statistical decision function framework, producing a framework one might call *Machine Wald*.

5.2 Reduction Calculus

The resolution of minimax problems (11) require, at an abstract level, searching in the space of all possible functions of the data. By restricting models to the Bayesian class, the complete class theorem allows to limit this search to prior distributions on \mathcal{A} , i.e., to measure over spaces of measures and functions. To enable the computation of these models, it is therefore necessary to identify conditions under which Minimax problems over measures over spaces of measures and functions can be reduced to the manipulation of finite-dimensional objects and develop the associated reduction calculus. For min or max problems over measures over spaces of measures (and possibly functions), this calculus can take the form of a reduction to a nesting of optimization problems over measures (and possibly functions for the inner part) [109, 112, 113], which, in turn, can be reduced to searches over extreme points [51, 110, 114, 142].

5.3 Stopping Conditions

Many of these optimization problems will not be tractable. However even in the tractable case, which has rigorous guarantees on the amount of computation required to obtain an approximate optima, it will be useful to have stopping criteria for the specific algorithm and the specific problem instance under consideration, which can be used to guarantee when an approximate optima has been achieved. Although in the intractable case no such guarantee will exist in general, intelligent choices of algorithms may result in the attainment of approximate optima and such tests guarantee that fact. Ermoliev, Gaivoronski, and Nedeva [36] successfully develop such stopping criteria using Lagrangian duality and generalized Bender's decompositions by Geoffrion [46] for certain stochastic optimization problems which are also relevant here. In addition, the approximation of intractable problems by tractable ones will be important. Recently, Hanasusanto, Roitch, Kuhn, and Wiesemann [53] derive explicit conic reformulations for tractable problem classes and suggest efficiently computable conservative approximations for intractable ones.

5.4 On the Borel-Kolmogorov Paradox

An oftentimes overlooked difficulty with Bayesian estimators lies in the fact that for a prior $\pi \in \mathcal{M}(\mathcal{A})$, the posterior (12) is not a measurable function of d but a convex set $\Theta(\pi)$ of measurable functions θ of d that are almost surely equal to each other under the measure $\pi \cdot \mathbb{D}$ on \mathcal{D} .

A notorious pathological consequence is the Borel–Kolmogorov paradox (see Chapter 5 of [76] and Section 15.7 of [63]). Recall that in the formulation of this paradox, one considers the uniform distribution on the two-dimensional sphere and one is interested in obtaining the conditional distribution associated with a great circle of that sphere. If the problem is parameterized in spherical coordinates, then the resulting conditional distribution is uniform for the equator but nonuniform for the longitude corresponding to the prime meridian. The following theorem suggests that this paradox is generic and dissipates the idea that it could be limited to fabricated toy examples. See also Singpurwalla and Swift [132] for implications of this paradox in modeling and inference.

Recall that for $\pi \in \mathcal{M}(\mathcal{A})$, that $\Theta(\pi)$ is defined as the convex set of measurable functions which are equal to $\pi \cdot \mathbb{D}$ -everywhere to the regular conditional expectation (12). Despite this indeterminateness, it is comforting to know that

$$\mathcal{E}(\theta_2, \pi) = \mathcal{E}(\theta_1, \pi), \quad \theta_1, \theta_2 \in \Theta(\pi).$$

Moreover, it is also easy to see that if π^\dagger is absolutely continuous with respect to π , then $\theta_1(d) = \theta_2(d)$ with $\pi^\dagger \cdot \mathbb{D}$ probability one for all $\theta_1, \theta_2 \in \Theta(\pi)$, and consequently

$$\mathcal{E}(\theta_2, \pi^\dagger) = \mathcal{E}(\theta_1, \pi^\dagger), \quad \theta_1, \theta_2 \in \Theta(\pi), \quad \pi^\dagger \prec \pi,$$

where the notation $\pi^\dagger \prec \pi$ means that π^\dagger is absolutely continuous with respect to π . The following theorem shows that this requirement of absolute continuity is necessary for all versions of conditional expectations $\theta \in \Theta(\pi)$ to share the same risk. See Sect. 6 for its proof.

Theorem 2. *Assume that $V(x) = x^2$ and that the image $\Phi(\mathcal{A})$ is a nontrivial interval. If π^\dagger is not absolutely continuous with respect to π then*

$$\frac{1}{4} \leq \frac{\sup_{\theta_1, \theta_2 \in \Theta(\pi)} (\mathcal{E}(\theta_2, \pi^\dagger) - \mathcal{E}(\theta_1, \pi^\dagger))}{(\mathcal{U}(\mathcal{A}) - \mathcal{L}(\mathcal{A}))^2 \sup_{B \in \mathcal{B}(\mathcal{D}): (\pi \cdot \mathbb{D})[B]=0} (\pi^\dagger \cdot \mathbb{D})[B]} \leq 1 \quad (19)$$

where $\mathcal{U}(\mathcal{A})$ and $\mathcal{L}(\mathcal{A})$ are defined by (2) and (3).

Remark 3. If moreover $\pi^\dagger \cdot \mathbb{D}$ is orthogonal to $\pi \cdot \mathbb{D}$, that is, there exists a set $B \in \mathcal{B}(\mathcal{D})$ such that $(\pi \cdot \mathbb{D})[B] = 0$ and $(\pi^\dagger \cdot \mathbb{D})[B] = 1$, then Theorem 2 implies that $\sup_{\theta_1, \theta_2 \in \Theta(\pi)} (\mathcal{E}(\theta_2, \pi^\dagger) - \mathcal{E}(\theta_1, \pi^\dagger))$ is larger than the statistical error of the midpoint estimator

$$\theta := \frac{\mathcal{L}(\mathcal{A}) + \mathcal{U}(\mathcal{A})}{2}.$$

As a remedy, one can try (see [144, 145] and [117]) constructing conditional expectations as disintegration or derivation limits defined as

$$\mathbb{E}_{\pi \circ \mathbb{D}}[\Phi(f, \mu) | D = d] = \lim_{B \downarrow \{d\}} \mathbb{E}_{\pi \circ \mathbb{D}}[\Phi(f, \mu) | D \in B] \quad (20)$$

where the limit $B \downarrow \{d\}$ is taken over a net of open neighborhoods of d . But as shown in [66], the limit generally depends on the net $B \downarrow \{d\}$ and the resulting conditional expectations can be distinctly different for different nets. Furthermore the limit (20) may exist/not exist on subsets of \mathcal{D} of $(\pi \cdot \mathbb{D})$ -measure zero (which, as shown above, can lead to the inconsistency of the estimator). A related important issue is that conditional probabilities can in general not be computed [2]. Observe that if the limit (20) does not exist, then Bayesian estimation of $\Phi(f, \mu)$ may have significant oscillations as the precise measurement of d becomes sharper.

5.5 On Bayesian Robustness/Brittleness

As much as classical numerical analysis shows that there are stable and unstable ways to discretize a partial differential equation, positive [13, 23, 28, 75, 80, 140, 152] and negative results [6, 27, 42, 43, 65, 84, 108, 109, 112, 113] are forming an emerging understanding of stable and unstable ways to apply Bayes' rule in practice. One aspect of stability concerns the sensitivity of posterior conclusions with respect to the underlying models and prior beliefs.

Most statisticians would acknowledge that an analysis is not complete unless the sensitivity of the conclusions to the assumptions is investigated. Yet, in practice, such sensitivity analyses are rarely used. This is because sensitivity analyses involve difficult computations that must often be tailored to the specific problem. This is especially true in Bayesian inference where the computations are already quite difficult. [165]

Another aspect concerns situations where Bayes' rule is applied iteratively and posterior values become prior values for the next iteration. Observe in particular that when posterior distributions (which are later on used as prior distributions) are only approximated (e.g., via MCMC methods), stability requires the convergence of the MCMC method in the same metric used to quantify the sensitivity of posterior with respect to the prior distributions.

In the context of the framework being developed here, recent results [108, 109, 112, 113] on the extreme sensitivity (brittleness) of Bayesian inference in the TV and Prokhorov metrics appear to suggest that robust inference, in a continuous world under finite-information, should perhaps be done with reduced/coarse models rather than highly sophisticated/complex models (with a level of coarseness/reduction depending on the available finite information) [113].

5.6 Information-Based Complexity

From the point of view of practical applications, it is clear that the set of possible models entering in the minimax problem 11 must be restricted by introducing constraints on computational complexity. For example, finding optimal models of materials in extreme environments is not the correct objective when these models require full quantum mechanics calculations. A more productive approach is to search for computationally tractable optimal models in a given complexity class. Here one may wonder if Bayesian models remain a complete class for the resulting complexity constrained minimax problems. It is also clear that computationally tractable optimal models may not use all the available information, for instance, a material model of bounded complexity should not use the state of every atom. The idea that fast computation requires computation with partial information forms the core of information-based complexity, the branch of computational complexity that studies the complexity of approximating continuous mathematical operations with discrete and finite ones up to a specified level of accuracy [101, 115, 146, 171, 172], where it is also augmented by concepts of contaminated and priced information associated with, for example, truncation errors and the cost of numerical operations. Recent results [107] suggest that decision theory concepts could be used, not only to identify reduced models but also algorithms of near-optimal complexity by reformulating the process of computing with partial information and limited resources as that of playing underlying hierarchies of adversarial information games.

6 Conclusion

Although uncertainty quantification is still in its formative stage, much like the state of probability theory before its rigorous formulation by Kolmogorov in the 1930s, it has the potential to have an impact on the process of scientific discovery that is similar to the advent of scientific computing. Its emergence remains sustained by the persistent need to make critical decisions with partial information and limited resources. There are many paths to its development, but one such path appears to be the incorporation of notions of computation and complexity in a generalization of Wald's decision framework built on Von Neumann's theory of adversarial games.

Appendix

Construction of $\pi \odot \mathbb{D}$

The below construction works when $\mathcal{A} \subseteq \mathcal{G} \times \mathcal{M}(\mathcal{X})$ for some Polish subset $\mathcal{G} \subset \mathcal{F}(\mathcal{X})$ and \mathcal{X} is Polish. Observe that since \mathcal{D} is metrizable, it follows from [4, Thm. 15.13], that, for any $B \in \mathcal{B}(\mathcal{D})$, the evaluation $\nu \mapsto \nu(B)$, $\nu \in \mathcal{M}(\mathcal{D})$, is measurable. Consequently, the measurability of \mathbb{D} implies that the mapping

$$\widehat{\mathbb{D}}: \mathcal{A} \times \mathcal{B}(\mathcal{D}) \rightarrow \mathbb{R}$$

defined by

$$\widehat{\mathbb{D}}((f, \mu), B) := \mathbb{D}(f, \mu)[B], \quad \text{for } (f, \mu) \in \mathcal{A}, B \in \mathcal{B}(\mathcal{D})$$

is a transition function in the sense that, for fixed $(f, \mu) \in \mathcal{A}$, $\widehat{\mathbb{D}}((f, \mu), \cdot)$ is a probability measure, and, for fixed $B \in \mathcal{B}(\mathcal{D})$, $\widehat{\mathbb{D}}(\cdot, B)$ is Borel measurable. Therefore, by [18, Thm. 10.7.2], any $\pi \in \mathcal{M}(\mathcal{A})$ defines a probability measure

$$\pi \odot \mathbb{D} \in \mathcal{M}(\mathcal{B}(\mathcal{A}) \times \mathcal{B}(\mathcal{D}))$$

through

$$\pi \odot \mathbb{D}[A \times B] := \mathbb{E}_{(f, \mu) \sim \pi} [\mathbb{1}_A(f, \mu) \mathbb{D}(f, \mu)[B]], \quad \text{for } A \in \mathcal{B}(\mathcal{A}), B \in \mathcal{B}(\mathcal{D}), \quad (21)$$

where $\mathbb{1}_A$ is the indicator function of the set A :

$$\mathbb{1}_A(f, \mu) := \begin{cases} 1, & \text{if } (f, \mu) \in A, \\ 0, & \text{if } (f, \mu) \notin A. \end{cases}$$

It is easy to see that π is the \mathcal{A} -marginal of $\pi \odot \mathbb{D}$. Moreover, when \mathcal{X} is Polish, [4, Thm. 15.15] implies that $\mathcal{M}(\mathcal{X})$ is Polish, and when \mathcal{G} is Polish, it follows that $\mathcal{A} \subseteq \mathcal{G} \times \mathcal{M}(\mathcal{X})$ is second countable. Consequently, since \mathcal{D} is Suslin and hence second countable, it follows from [32, Prop. 4.1.7] that

$$\mathcal{B}(\mathcal{A} \times \mathcal{D}) = \mathcal{B}(\mathcal{A}) \times \mathcal{B}(\mathcal{D})$$

and hence $\pi \odot \mathbb{D}$ is a probability measure on $\mathcal{A} \times \mathcal{D}$. That is,

$$\pi \odot \mathbb{D} \in \mathcal{M}(\mathcal{A} \times \mathcal{D}).$$

Henceforth denote $\pi \cdot \mathbb{D}$ the corresponding Bayes' sampling distribution defined by the \mathcal{D} -marginal of $\pi \odot \mathbb{D}$, and note that by (21), one has

$$\pi \cdot \mathbb{D}[B] := \mathbb{E}_{(f,\mu) \sim \pi} [\mathbb{D}(f, \mu)[B]], \quad \text{for } B \in \mathcal{B}(\mathcal{D}).$$

Since both \mathcal{D} and \mathcal{A} are Suslin, it follows that the product $\mathcal{A} \times \mathcal{D}$ is Suslin. Consequently, [18, Cor. 10.4.6] asserts that regular conditional probabilities exist for any sub- σ -algebra of $\mathcal{B}(\mathcal{A} \times \mathcal{D})$. In particular, the product theorem of [18, Thm. 10.4.11] asserts that product regular conditional probabilities

$$(\pi \odot \mathbb{D})|_d \in \mathcal{M}(\mathcal{A}), \quad \text{for } d \in \mathcal{D}$$

exist and that they are $\pi \cdot \mathbb{D}$ -a.e. unique.

Proof of Theorem 2

If $\pi^\dagger \cdot \mathbb{D}$ is not absolutely continuous with respect to $\pi \cdot \mathbb{D}$, then there exists $B \in \mathcal{B}(\mathcal{D})$ such that $(\pi \cdot \mathbb{D})[B] = 0$ and $(\pi^\dagger \cdot \mathbb{D})[B] > 0$. Let $\theta \in \Theta(\pi)$. Define

$$\theta_y(d) := \theta(d)1_{B^c}(d) + y1_B(d) \tag{22}$$

Then it is easy to see that if y is in the range of Φ , then $\theta_y \in \Theta(\pi)$. Now observe that for $y, z \in \text{Image}(\Phi)$,

$$\mathcal{E}(\theta_y, \pi^\dagger) - \mathcal{E}(\theta_z, \pi^\dagger) = \mathbb{E}_{(f,\mu,d) \sim \pi^\dagger \odot \mathbb{D}} \left[1_B(d) \left(V(y - \Phi(f, \mu)) - V(z - \Phi(f, \mu)) \right) \right]$$

Hence, for $V(x) = x^2$, it holds true that

$$\mathcal{E}(\theta_y, \pi^\dagger) - \mathcal{E}(\theta_z, \pi^\dagger) = [(y - \gamma)^2 - (z - \gamma)^2](\pi^\dagger \cdot \mathbb{D})[B]$$

with

$$\gamma := \mathbb{E}_{\pi^\dagger \circ \mathbb{D}}[\Phi | D \in B]$$

which proves

$$\begin{aligned} \sup_{\theta_2 \in \Theta(\pi)} \mathcal{E}(\theta_2, \pi^\dagger) - \inf_{\theta_1 \in \Theta(\pi)} \mathcal{E}(\theta_1, \pi^\dagger) &\geq \sup_{B \in \mathcal{B}(\mathcal{D}) : (\pi \cdot \mathbb{D})[B]=0, y, z \in \text{Image}(\Phi)} \\ &\left[(y - \mathbb{E}_{\pi^\dagger \circ \mathbb{D}}[\Phi | D \in B])^2 - (z - \mathbb{E}_{\pi^\dagger \circ \mathbb{D}}[\Phi | D \in B])^2 \right] (\pi^\dagger \cdot \mathbb{D})[B], \end{aligned}$$

and,

$$\sup_{\theta_2 \in \Theta(\pi)} \mathcal{E}(\theta_2, \pi^\dagger) - \inf_{\theta_1 \in \Theta(\pi)} \mathcal{E}(\theta_1, \pi^\dagger) \leq (\mathcal{U}(\mathcal{A}) - \mathcal{L}(\mathcal{A}))^2 \sup_{B \in \mathcal{B}(\mathcal{D}) : (\pi \cdot \mathbb{D})[B]=0} (\pi^\dagger \cdot \mathbb{D})[B].$$

To obtain the right hand side of (19) observe that (see for instance [29, Sec. 5]) there exists $B^* \in \mathcal{B}(\mathcal{D})$ such that

$$(\pi^\dagger \cdot \mathbb{D})[B^*] = \sup_{B \in \mathcal{B}(\mathcal{D}) : (\pi \cdot \mathbb{D})[B]=0} (\pi^\dagger \cdot \mathbb{D})[B]$$

and (since $\theta_2 = \theta_1$ on the complement of B^*)

$$\begin{aligned} &\sup_{\theta_1, \theta_2 \in \Theta(\pi)} (\mathcal{E}(\theta_2, \pi^\dagger) - \mathcal{E}(\theta_1, \pi^\dagger)) \\ &= \sup_{\theta_1, \theta_2 \in \Theta(\pi)} \mathbb{E}_{(f, \mu, d) \sim \pi^\dagger \circ \mathbb{D}} \left[\mathbf{1}_{B^*}(d) \left(V(\theta_2 - \Phi(f, \mu)) - V(\theta_1 - \Phi(f, \mu)) \right) \right]. \end{aligned}$$

We conclude by observing that for $V(x) = x^2$,

$$\sup_{\theta_1, \theta_2 \in \Theta(\pi)} \left(V(\theta_2 - \Phi(f, \mu)) - V(\theta_1 - \Phi(f, \mu)) \right) \leq (\mathcal{U}(\mathcal{A}) - \mathcal{L}(\mathcal{A}))^2.$$

Conditional Expectation as an Orthogonal Projection

It easily follows from Tonelli's Theorem that

$$\mathbb{E}_{\pi \cdot \mathbb{D}}[h^2] = \mathbb{E}_{\pi \circ \mathbb{D}}[h^2] = \mathbb{E}_{(f, \mu) \sim \pi} \mathbb{E}_{\mathbb{D}(f, \mu)}[h^2].$$

By considering the sub σ -algebra $\mathcal{A} \times \mathcal{B}(\mathcal{D}) \subset \mathcal{B}(\mathcal{A} \times \mathcal{D}) = \mathcal{B}(\mathcal{A}) \times \mathcal{B}(\mathcal{D})$, it follows from, e.g., Theorem 10.2.9 of [32], that $L^2_{\pi \cdot \mathbb{D}}(\mathcal{D})$ is a closed Hilbert subspace of the Hilbert space $L^2_{\pi \circ \mathbb{D}}(\mathcal{A} \times \mathcal{D})$ and the conditional expectation of Φ given the random variable D is the orthogonal projection from $L^2_{\pi \circ \mathbb{D}}(\mathcal{A} \times \mathcal{D})$ to $L^2_{\pi \cdot \mathbb{D}}(\mathcal{D})$.

References

1. Richardson, L.F.: *Weather Prediction by Numerical Process*. Cambridge Mathematical Library. Cambridge University Press, Cambridge (1922)
2. Ackerman, N.L., Freer, C.E., Roy, D.M.: On the computability of conditional probability. arXiv:1005.3014 (2010)
3. Adams, M., Lashgari, A., Li, B., McKerns, M., Mihaly, J.M., Ortiz, M., Owhadi, H., Rosakis, A.J., Stalzer, M., Sullivan, T.J.: Rigorous model-based uncertainty quantification with application to terminal ballistics. Part II: systems with uncontrollable inputs and large scatter. *J. Mech. Phys. Solids* **60**(5), 1002–1019 (2012)
4. Aliprantis, C.D., Border, K.C.: *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, 3rd edn. Springer, Berlin (2006)
5. Anderson, T.W.: The integral of a symmetric unimodal function over a symmetric convex set and some probability inequalities. *Proc. Am. Math. Soc.* **6**(2), 170–176 (1955)
6. Belot, G.: Bayesian orgulity. *Philos. Sci.* **80**(4), 483–503 (2013)
7. Ben-Tal, A., El Ghaoui, L., Nemirovski, A.: *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton (2009)
8. Ben-Tal, A., Hochman, E.: More bounds on the expectation of a convex function of a random variable. *J. Appl. Probab.* **9**, 803–812 (1972)
9. Ben-Tal, A., Nemirovski, A.: Robust convex optimization. *Math. Oper. Res.* **23**(4), 769–805 (1998)
10. Bentkus, V.: A remark on the inequalities of Bernstein, Prokhorov, Bennett, Hoeffding, and Talagrand. *Liet. Mat. Rink.* **42**(3), 332–342 (2002)
11. Bentkus, V.: On Hoeffding’s inequalities. *Ann. Probab.* **32**(2), 1650–1673 (2004)
12. Bentkus, V., Geuze, G.D.C., van Zuijlen, M.C.A.: Optimal Hoeffding-like inequalities under a symmetry assumption. *Statistics* **40**(2), 159–164 (2006)
13. Bernstein, S.N.: *Collected Works*. Izdat. “Nauka”, Moscow (1964)
14. Bertsimas, D., Brown, D.B., Caramanis, C.: Theory and applications of robust optimization. *SIAM Rev.* **53**(3), 464–501 (2011)
15. Bertsimas, D., Popescu, I.: Optimal inequalities in probability theory: a convex optimization approach. *SIAM J. Optim.* **15**(3), 780–804 (electronic) (2005)
16. Birge, J.R., Wets, R.J.-B.: Designing approximation schemes for stochastic optimization problems, in particular for stochastic programs with recourse. *Math. Prog. Stud.* **27**, 54–102 (1986)
17. Blackwell, D.: Equivalent comparisons of experiments. *Ann. Math. Stat.* **24**(2), 265–272 (1953)
18. Bogachev, V.I.: *Measure Theory*, vol. II. Springer, Berlin (2007)
19. Boj, R.I., Lorenz, N., Wanka, G.: Duality for linear chance-constrained optimization problems. *J. Korean Math. Soc.* **47**(1), 17–28 (2010)
20. Boucheron, S., Lugosi, G., Massart, P.: A sharp concentration inequality with applications. *Random Struct. Algorithms* **16**(3), 277–292 (2000)
21. Brown, L.D.: Minimacity, more or less. In: Gupta, S.S., Berger, J.O. (eds.) *Statistical Decision Theory and Related Topics V*, pp. 1–18. Springer, New York (1994)
22. Brown, L.D.: An essay on statistical decision theory. *J. Am. Stat. Assoc.* **95**(452), 1277–1281 (2000)
23. Castillo, I., Nickl, R.: Nonparametric Bernstein–von Mises theorems in Gaussian white noise. *Ann. Stat.* **41**(4), 1999–2028 (2013)
24. Chen, W., Sim, M., Sun, J., Teo, C.-P.: From CVaR to uncertainty set: implications in joint chance-constrained optimization. *Oper. Res.* **58**(2), 470–485 (2010)
25. Daley, D.J., Vere-Jones, D.: *An Introduction to the Theory of Point Processes.: General Theory and Structure. Probability and its Applications* (New York), vol. II, 2nd edn. Springer, New York (2008)
26. Dantzig, G.B.: Linear programming under uncertainty. *Manag. Sci.* **1**, 197–206 (1955)

27. Diaconis, P., Freedman, D.A.: On the consistency of Bayes estimates. *Ann. Stat.* **14**(1), 1–67 (1986). With a discussion and a rejoinder by the authors
28. Doob, J.L.: Application of the theory of martingales. In: *Le Calcul des Probabilités et ses Applications, Colloques Internationaux du Centre National de la Recherche Scientifique*, vol. 13, pp. 23–27. Centre National de la Recherche Scientifique, Paris (1949)
29. Doob, J.L.: *Measure Theory*. Graduate Texts in Mathematics, vol. 143. Springer, New York (1994)
30. Drenick, R.F.: Aseismic design by way of critical excitation. *J. Eng. Mech. Div. Am. Soc. Civ. Eng.* **99**(4), 649–667 (1973)
31. Dubins, L.E.: On extreme points of convex sets. *J. Math. Anal. Appl.* **5**(2), 237–244 (1962)
32. Dudley, R.M.: *Real Analysis and Probability*. Cambridge Studies in Advanced Mathematics, vol. 74. Cambridge University Press, Cambridge (2002). Revised reprint of the 1989 original
33. Dvoretzky, A., Wald, A., Wolfowitz, J.: Elimination of randomization in certain statistical decision procedures and zero-sum two-person games. *Ann. Math. Stat.* **22**(1), 1–21 (1951)
34. Edmundson, H.P.: Bounds on the expectation of a convex function of a random variable. Technical report, DTIC Document (1957)
35. Elishakoff, I., Ohsaki, M.: *Optimization and Anti-optimization of Structures Under Uncertainty*. World Scientific, London (2010)
36. Ermoliev, Y., Gaivoronski, A., Nedeva, C.: Stochastic optimization problems with incomplete information on distribution functions. *SIAM J. Control Optim.* **23**(5), 697–716 (1985)
37. Fisher, R.: *The Design of Experiments*. Oliver and Boyd, Edinburgh (1935)
38. Fisher, R.: Statistical methods and scientific induction. *J. R. Stat. Soc. Ser. B.* **17**, 69–78 (1955)
39. Fisher, R.A.: On the mathematical foundations of theoretical statistics. *Philos. Trans. R. Soc. Lond. Ser. A* **222**, 309–368 (1922)
40. Fisher, R.A.: “Student”. *Ann. Eugen.* **9**(1), 1–9 (1939)
41. Frauendorfer, K.: Solving SLP recourse problems with arbitrary multivariate distributions—the dependent case. *Math. Oper. Res.* **13**(3), 377–394 (1988)
42. Freedman, D.A.: On the asymptotic behavior of Bayes’ estimates in the discrete case. *Ann. Math. Stat.* **34**, 1386–1403 (1963)
43. Freedman, D.A.: On the Bernstein-von Mises theorem with infinite-dimensional parameters. *Ann. Stat.* **27**(4), 1119–1140 (1999)
44. Gaivoronski, A.A.: A numerical method for solving stochastic programming problems with moment constraints on a distribution function. *Ann. Oper. Res.* **31**(1), 347–369 (1991)
45. Gassmann, H., Ziemba, W.T.: A tight upper bound for the expectation of a convex function of a multivariate random variable. In: *Stochastic Programming 84 Part I. Mathematical Programming Study*, vol. 27, pp. 39–53. Springer, Berlin (1986)
46. Geoffrion, A.M.: Generalized Benders decomposition. *JOTA* **10**(4), 237–260 (1972)
47. Gilboa, I., Schmeidler, D.: Maxmin expected utility with non-unique prior. *J. Math. Econ.* **18**(2), 141–153 (1989)
48. Godwin, H.J.: On generalizations of Tchebycheff’s inequality. *J. Am. Stat. Assoc.* **50**(271), 923–945 (1955)
49. Goh, J., Sim, M.: Distributionally robust optimization and its tractable approximations. *Oper. Res.* **58**(4, part 1), 902–917 (2010)
50. Halmos, P.R., Savage, L.J.: Application of the Radon-Nikodym theorem to the theory of sufficient statistics. *Ann. Math. Stat.* **20**(2), 225–241 (1949)
51. Han, S., Tao, M., Topcu, U., Owhadi, H., Murray, R.M.: Convex optimal uncertainty quantification. *SIAM J. Optim.* **25**(23), 1368–1387 (2015). arXiv:1311.7130
52. Han, S., Topcu, U., Tao, M., Owhadi, H., Murray, R.: Convex optimal uncertainty quantification: algorithms and a case study in energy storage placement for power grids. In: *American Control Conference (ACC)*, 2013, Washington, DC, pp. 1130–1137. IEEE (2013)
53. Hanasusanto, G.A., Roitch, V., Kuhn, D., Wiesemann, W.: A distributionally robust perspective on uncertainty quantification and chance constrained programming. *Math. Program.* **151**(1), 35–62 (2015)

54. Hoeffding, W.: On the distribution of the number of successes in independent trials. *Ann. Math. Stat.* **27**(3), 713–721 (1956)
55. Hotelling, H.: Abraham Wald. *Am. Stat.* **5**(1), 18–19 (1951)
56. Huang, C.C., Vertinsky, I., Ziemba, W.T.: Sharp bounds on the value of perfect information. *Oper. Res.* **25**(1), 128–139 (1977)
57. Huang, C.C., Ziemba, W.T., Ben-Tal, A.: Bounds on the expectation of a convex function of a random variable: with applications to stochastic programming. *Oper. Res.* **25**(2), 315–325 (1977)
58. Huber, P.J.: Robust estimation of a location parameter. *Ann. Math. Stat.* **35**, 73–101 (1964)
59. Huber, P.J.: The 1972 Wald lecture- Robust statistics: a review. *Ann. Math. Stat.* 1041–1067 (1972)
60. Isii, K.: On a method for generalizations of Tchebycheff's inequality. *Ann. Inst. Stat. Math. Tokyo* **10**(2), 65–88 (1959)
61. Isii, K.: The extrema of probability determined by generalized moments. I. Bounded random variables. *Ann. Inst. Stat. Math.* **12**(2), 119–134; errata, 280 (1960)
62. Isii, K.: On sharpness of Tchebycheff-type inequalities. *Ann. Inst. Stat. Math.* **14**(1):185–197, 1962/1963.
63. Jaynes, E.T.: *Probability Theory*. Cambridge University Press, Cambridge (2003)
64. Joe, H.: Majorization, randomness and dependence for multivariate distributions. *Ann. Probab.* **15**(3), 1217–1225 (1987)
65. Johnstone, I.M.: High dimensional Bernstein–von Mises: simple examples. In *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown*, volume 6 of *Inst. Math. Stat. Collect.*, pages 87–98. *Inst. Math. Statist.*, Beachwood, OH (2010)
66. Kac, M., Slepian, D.: Large excursions of Gaussian processes. *Ann. Math. Stat.* **30**, 1215–1228 (1959)
67. Kall, P.: Stochastic programming with recourse: upper bounds and moment problems: a review. *Math. Res.* **45**, 86–103 (1988)
68. Kallenberg, O.: *Random Measures*. Akademie-Verlag, Berlin (1975) Schriftenreihe des Zentralinstituts für Mathematik und Mechanik bei der Akademie der Wissenschaften der DDR, Heft 23.
69. Kanga, P.-H.T., Li, B., McKerns, M., Nguyen, L.H., Ortiz, M., Owhadi, H., Sullivan, T.J.: Optimal uncertainty quantification with model uncertainty and legacy data. *J. Mech. Phys. Solids* **72**, 1–19 (2014)
70. Karlin, S., Studden, W.J.: *Tchebycheff Systems: With Applications in Analysis and Statistics*. Pure and Applied Mathematics, vol. XV. Interscience Publishers/Wiley, New York/London/Sydney (1966)
71. Kendall, D.G.: Simplexes and vector lattices. *J. Lond. Math. Soc.* **37**(1), 365–371 (1962)
72. Kidane, A.A., Lashgari, A., Li, B., McKerns, M., Ortiz, M., Owhadi, H., Ravichandran, G., Stalzer, M., Sullivan, T.J.: Rigorous model-based uncertainty quantification with application to terminal ballistics. Part I: Systems with controllable inputs and small scatter. *J. Mech. Phys. Solids* **60**(5), 983–1001 (2012)
73. Kiefer, J.: Optimum experimental designs. *J. R. Stat. Soc. Ser. B* **21**, 272–319 (1959)
74. Kiefer, J.: *Collected Works*, vol. III. Springer, New York (1985)
75. Kleijn, B.J.K., van der Vaart, A.W.: The Bernstein-Von-Mises theorem under misspecification. *Electron. J. Stat.* **6**, 354–381 (2012)
76. Kolmogorov, A.N.: *Foundations of the Theory of Probability*. Chelsea Publishing Co., New York (1956). Translation edited by Nathan Morrison, with an added bibliography by A. T. Bharucha-Reid
77. Kreĭn, M.G.: The ideas of P. L. Čebyšev and A. A. Markov in the theory of limiting values of integrals and their further development. In: Dynkin, E.B. (ed.) *Eleven Papers on Analysis, Probability, and Topology*, American Mathematical Society Translations, Series 2, vol. 12, pp. 1–122. American Mathematical Society, New York (1959)
78. Kurz, H.D., Salvadori, N.: *Understanding 'Classical' Economics: Studies in Long Period Theory*. Routledge, London/New York (2002)

79. Laird, N.M.: A conversation with F. N. David. *Stat. Sci.* **4**, 235–246 (1989)
80. Le Cam, L.: On some asymptotic properties of maximum likelihood estimates and related Bayes' estimates. *Univ. Calif. Publ. Stat.* **1**, 277–329 (1953)
81. Le Cam, L.: An extension of Wald's theory of statistical decision functions. *Ann. Math. Stat.* **26**, 69–81 (1955)
82. Le Cam, L.: Sufficiency and approximate sufficiency. *Ann. Math. Stat.* **35**, 1419–1455 (1964)
83. Le Cam, L.: *Asymptotic Methods in Statistical Decision Theory*. Springer, New York (1986)
84. Leahu, H.: On the Bernstein–von Mises phenomenon in the Gaussian white noise model. *Electron. J. Stat.* **5**, 373–404 (2011)
85. Lehmann, E.L.: “Student” and small-sample theory. *Stat. Sci.* **14**(4), 418–426 (1999)
86. Lehmann, E.L.: Optimality and symposia: some history. *Lect. Notes Monogr. Ser.* **44**, 1–10 (2004)
87. Lehmann, E.L.: Some history of optimality. *Lect. Notes Monogr. Ser.* **57**, 11–17 (2009)
88. Lenhard, J.: Models and statistical inference: the controversy between Fisher and Neyman–Pearson. *Br. J. Philos. Sci.* **57**(1), 69–91 (2006)
89. Leonard, R.: Von Neumann, Morgenstern, and the Creation of Game Theory: From Chess to Social Science, 1900–1960. Cambridge University Press, New York (2010)
90. Lynch, P.: The origins of computer weather prediction and climate modeling. *J. Comput. Phys.* **227**(7), 3431–3444 (2008)
91. Madansky, A.: Bounds on the expectation of a convex function of a multivariate random variable. *Ann. Math. Stat.* 743–746 (1959)
92. Madansky, A.: Inequalities for stochastic linear programming problems. *Manag. Sci.* **6**(2), 197–204 (1960)
93. Mangel, M., Samaniego, F.J.: Abraham Wald's work on aircraft survivability. *J. A. S. A.* **79**(386), 259–267 (1984)
94. Marshall, A.W., Olkin, I.: Multivariate Chebyshev inequalities. *Ann. Math. Stat.* **31**(4), 1001–1014 (1960)
95. Marshall, A.W., Olkin, I.: *Inequalities: Theory of Majorization and Its Applications*. Mathematics in Science and Engineering, vol. 143. Academic [Harcourt Brace Jovanovich Publishers], New York (1979)
96. McKerns, M.M., Strand, L., Sullivan, T.J., Fang, A., Aivazis, M.A.G.: Building a framework for predictive science. In: *Proceedings of the 10th Python in Science Conference (SciPy 2011)* (2011)
97. Morgenstern, O.: Abraham Wald, 1902–1950. *Econometrica: J. Econom. Soci.* 361–367 (1951)
98. Mulholland, H.P., Rogers, C.A.: Representation theorems for distribution functions. *Proc. Lond. Math. Soc.* (3) **8**(2), 177–223 (1958)
99. Nash, J.: Non-cooperative games. *Ann. Math.* (2) **54**, 286–295 (1951)
100. Nash, J.F. Jr.: Equilibrium points in n -person games. *Proc. Natl. Acad. Sci. U. S. A.* **36**, 48–49 (1950)
101. Nemirovsky, A.S.: Information-based complexity of linear operator equations. *J. Complex.* **8**(2), 153–175 (1992)
102. Neyman, J.: Outline of a theory of statistical estimation based on the classical theory of probability. *Philos. Trans. R. Soc. Lond. Ser. A* **236**(767), 333–380 (1937)
103. Neyman, J.: *A Selection of Early Statistical Papers of J. Neyman*. University of California Press, Berkeley (1967)
104. Neyman, J., Pearson, E.S.: On the use and interpretation of certain test criteria for purposes of statistical inference. *Biometrika* **20A**, 175–240, 263–294 (1928)
105. Neyman, J., Pearson, E.S.: On the problem of the most efficient tests of statistical hypotheses. *Philos. Trans. R. Soc. Lond. Ser. A* **231**, 289–337 (1933)
106. Olkin, I., Pratt, J.W.: A multivariate Tchebycheff inequality. *Ann. Math. Stat.* **29**(1), 226–234 (1958)
107. Owjadi, H.: Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games. *SIAM Rev. (Research spotlights)* (2016, to appear). arXiv:1503.03467

108. Owhadi, H., Scovel, C.: Qualitative robustness in Bayesian inference. arXiv:1411.3984 (2014)
109. Owhadi, H., Scovel, C.: Brittleness of Bayesian inference and new Selberg formulas. *Commun. Math. Sci.* **14**(1), 83–145 (2016)
110. Owhadi, H., Scovel, C.: Extreme points of a ball about a measure with finite support. *Commun. Math. Sci.* (2015, to appear). arXiv:1504.06745
111. Owhadi, H., Scovel, C.: Separability of reproducing kernel Hilbert spaces. *Proc. Am. Math. Soc.* (2015, to appear). arXiv:1506.04288
112. Owhadi, H., Scovel, C., Sullivan, T.J.: Brittleness of Bayesian inference under finite information in a continuous world. *Electron. J. Stat.* **9**, 1–79 (2015)
113. Owhadi, H., Scovel, C., Sullivan, T.J.: On the Brittleness of Bayesian Inference. *SIAM Rev. (Research Spotlights)* (2015)
114. Owhadi, H., Scovel, C., Sullivan, T.J., McKerns, M., Ortiz, M.: Optimal Uncertainty Quantification. *SIAM Rev.* **55**(2), 271–345 (2013)
115. Packel, E.W.: The algorithm designer versus nature: a game-theoretic approach to information-based complexity. *J. Complex.* **3**(3), 244–257 (1987)
116. Pearson, E.S.: ‘Student’ A Statistical Biography of William Sealy Gosset. Clarendon Press, Oxford (1990)
117. Pfanzagl, J.: Conditional distributions as derivatives. *Ann. Probab.* **7**(6), 1046–1050 (1979)
118. Pinelis, I.: Exact inequalities for sums of asymmetric random variables, with applications. *Probab. Theory Relat. Fields* 139(3-4):605–635 (2007)
119. Pinelis, I.: On inequalities for sums of bounded random variables. *J. Math. Inequal.* **2**(1), 1–7 (2008)
120. Platzman, G.W.: The ENIAC computations of 1950-gateway to numerical weather prediction. *Bull. Am. Meteorol. Soc.* **60**, 302–312 (1979)
121. Ressel, P.: Some continuity and measurability results on spaces of measures. *Mathematica Scandinavica* **40**, 69–78 (1977)
122. Rikun, A.D.: A convex envelope formula for multilinear functions. *J. Global Optim.* **10**(4), 425–437 (1997)
123. Rockafellar, R.T.: Augmented Lagrange multiplier functions and duality in nonconvex programming. *SIAM J. Control* **12**(2), 268–285 (1974)
124. Rojo, J.: Optimality: The Second Erich L. Lehmann Symposium. IMS, Beachwood (2006)
125. Rojo, J.: Optimality: The Third Erich L. Lehmann Symposium. IMS, Beachwood (2009)
126. Rojo, J., Pérez-Abreu, V.: The First Erich L. Lehmann Symposium: Optimality. IMS, Beachwood (2004)
127. Rustem, B., Howe, M.: Algorithms for Worst-Case Design and Applications to Risk Management. Princeton University Press, Princeton (2002)
128. Savage, L.J.: The theory of statistical decision. *J. Am. Stat. Assoc.* **46**, 55–67 (1951)
129. Scovel, C., Hush, D., Steinwart, I.: Approximate duality. *J. Optim. Theory Appl.* **135**(3), 429–443 (2007)
130. Shapiro, A., Kleywegt, A.: Minimax analysis of stochastic problems. *Optim. Methods Softw.* **17**(3), 523–542 (2002)
131. Sherali, H.D.: Convex envelopes of multilinear functions over a unit hypercube and over special discrete sets. *Acta Math. Vietnam.* **22**(1), 245–270 (1997)
132. Singpurwalla, N.D., Swift, A.: Network reliability and Borel’s paradox. *Am. Stat.* **55**(3), 213–218 (2001)
133. Smith, J.E.: Generalized Chebychev inequalities: theory and applications in decision analysis. *Oper. Res.* **43**(5), 807–825 (1995)
134. Sniedovich, M.: The art and science of modeling decision-making under severe uncertainty. *Decis. Mak. Manuf. Serv.* **1**(1–2), 111–136 (2007)
135. Sniedovich, M.: A classical decision theoretic perspective on worst-case analysis. *Appl. Math.* **56**(5), 499–509 (2011)
136. Sniedovich, M.: Black Swans, new Nostradamuses, Voodoo decision theories, and the science of decision making in the face of severe uncertainty. *Int. Trans. Oper. Res.* **19**(1–2), 253–281 (2012)

137. Spanos, A.: Why the Decision-Theoretic Perspective Misrepresents Frequentist Inference (2014). <https://secure.hosting.vt.edu/www.econ.vt.edu/directory/spanos/spanos10.pdf>
138. Stein, C.: Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In: Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, vol. I, pp. 197–206. University of California Press, Berkeley/Los Angeles (1956)
139. Strasser, H.: Mathematical Theory of Statistics: Statistical Experiments and Asymptotic Decision Theory, vol. 7. Walter de Gruyter, Berlin/New York (1985)
140. Stuart, A.M.: Inverse problems: a Bayesian perspective. *Acta Numer.* **19**, 451–559 (2010)
141. Student: The probable error of a mean. *Biometrika* 1–25 (1908)
142. Sullivan, T.J., McKerns, M., Meyer, D., Theil, F., Owjadi, H., Ortiz, M.: Optimal uncertainty quantification for legacy data observations of Lipschitz functions. *ESAIM Math. Model. Numer. Anal.* **47**(6), 1657–1689 (2013)
143. Tintner, G.: Abraham Wald’s contributions to econometrics. *Ann. Math. Stat.* **23**, 21–28 (1952)
144. Tjur, T.: Conditional Probability Distributions, Lecture Notes, No. 2. Institute of Mathematical Statistics, University of Copenhagen, Copenhagen (1974)
145. Tjur, T.: Probability Based on Radon Measures. Wiley Series in Probability and Mathematical Statistics. Wiley, Chichester (1980)
146. Traub, J.F., Wasilkowski, G.W., Woźniakowski, H.: Information-Based Complexity. Computer Science and Scientific Computing. Academic, Boston (1988). With contributions by A. G. Werschulz and T. Boulton
147. Tukey, J.W.: Statistical and Quantitative Methodology. Trends in Social Science, pp. 84–136. Philosophical Library, New York (1961)
148. Tukey, J.W.: The future of data analysis. *Ann. Math. Stat.* **33**, 1–67 (1962)
149. Valiant, L.G.: A theory of the learnable. *Commun. ACM* **27**(11), 1134–1142 (1984)
150. Vandenberghe, L., Boyd, S., Comanor, K.: Generalized Chebyshev bounds via semidefinite programming. *SIAM Rev.* **49**(1), 52–64 (electronic) (2007)
151. Varadarajan, V.S.: Groups of automorphisms of Borel spaces. *Trans. Am. Math. Soc.* **109**(2), 191–220 (1963)
152. von Mises, R.: Mathematical Theory of Probability and Statistics. Edited and Complemented by Hilda Geiringer. Academic, New York (1964)
153. Von Neumann, J.: Zur Theorie der Gesellschaftsspiele. *Math. Ann.* **100**(1), 295–320 (1928)
154. Von Neumann, J., Goldstine, H.H.: Numerical inverting of matrices of high order. *Bull. Am. Math. Soc.* **53**, 1021–1099 (1947)
155. Von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press, Princeton (1944)
156. Wald, A.: Contributions to the theory of statistical estimation and testing hypotheses. *Ann. Math. Stat.* **10**(4), 299–326 (1939)
157. Wald, A.: Statistical decision functions which minimize the maximum risk. *Ann. Math. (2)* **46**, 265–280 (1945)
158. Wald, A.: An essentially complete class of admissible decision functions. *Ann. Math. Stat.* **18**, 549–555 (1947)
159. Wald, A.: Sequential Analysis. 1947.
160. Wald, A.: Statistical decision functions. *Ann. Math. Stat.* **20**, 165–205 (1949)
161. Wald, A.: Statistical Decision Functions. Wiley, New York (1950)
162. Wald, A., Wolfowitz, J.: Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* **19**(3), 326–339 (1948)
163. Wald, A., Wolfowitz, J.: Characterization of the minimal complete class of decision functions when the number of distributions and decisions is finite. In: Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, pp. 149–157. University of California Press, Berkeley (1951)
164. Wasserman, L.: Rise of the Machines. Past, Present and Future of Statistical Science. CRC Press, Boca Raton (2013)

165. Wasserman, L., Lavine, M., Wolpert, R.L.: Linearization of Bayesian robustness problems. *J. Stat. Plann. Inference* **37**(3), 307–316 (1993)
166. Wiesemann, W., Kuhn, D., Sim, M.: Distributionally robust convex optimization. *Oper. Res.* **62**(6), 1358–1376 (2014)
167. Wilson, M.: How a story from World War II shapes Facebook today. IBM Watson (2012). <http://www.fastcodesign.com/1671172/how-a-story-from-world-war-ii-shapes-facebook-today>.
168. Winkler, G.: On the integral representation in convex noncompact sets of tight measures. *Mathematische Zeitschrift* **158**(1), 71–77 (1978)
169. Winkler, G.: Extreme points of moment sets. *Math. Oper. Res.* **13**(4), 581–587 (1988)
170. Wolfowitz, J.: Abraham Wald, 1902–1950. *Ann. Math. Stat.* **23**, 1–13 (1952)
171. Woźniakowski, H.: Probabilistic setting of information-based complexity. *J. Complex.* **2**(3), 255–269 (1986)
172. Woźniakowski, H.: What is information-based complexity? In *Essays on the Complexity of Continuous Problems*, pp. 89–95. European Mathematical Society, Zürich (2009)
173. Wynn, H.P.: Introduction to Kiefer (1959) Optimum Experimental Designs. In *Breakthroughs in Statistics*, pp. 395–399. Springer, New York (1992)
174. Xu, L., Yu, B., Liu, W.: The distributionally robust optimization reformulation for stochastic complementarity problems. *Abstr. Appl. Anal.* **2014**, Art. ID 469587, (2014)
175. Žáčková, J.: On minimax solutions of stochastic linear programming problems. *Časopis Pěst. Mat.* **91**, 423–430 (1966)
176. Zhou, K., Doyle, J.C., Glover, K.: *Robust and Optimal Control*. Prentice Hall, Upper Saddle River (1996)
177. Zymler, S., Kuhn, D., Rustem, B.: Distributionally robust joint chance constraints with second-order moment information. *Math. Program.* **137**(1–2, Ser. A), 167–198 (2013)