



Article Aspects of a Phase Transition in High-Dimensional Random Geometry

Axel Prüser ^{1,*}, Imre Kondor ^{2,3,4} and Andreas Engel ¹

- ¹ Institute of Physics, Carl von Ossietzky University of Oldenburg, D-26111 Oldenburg, Germany; andreas.engel@uol.de
- ² Parmenides Foundation, 82049 Pullach, Germany; kondor.imre@gmail.com
- ³ London Mathematical Laboratory, London W6 8RH, UK
- ⁴ Complexity Science Hub, 1080 Vienna, Austria
- * Correspondence: axel.prueser@uol.de

Abstract: A phase transition in high-dimensional random geometry is analyzed as it arises in a variety of problems. A prominent example is the feasibility of a minimax problem that represents the extremal case of a class of financial risk measures, among them the current regulatory market risk measure Expected Shortfall. Others include portfolio optimization with a ban on short-selling, the storage capacity of the perceptron, the solvability of a set of linear equations with random coefficients, and competition for resources in an ecological system. These examples shed light on various aspects of the underlying geometric phase transition, create links between problems belonging to seemingly distant fields, and offer the possibility for further ramifications.

Keywords: random geometry; portfolio optimization; risk measurement; disordered systems; replica theory

PACS: 05.20.-y; 05.40.-a; 05.70.Fh; 87.23.Ge

1. Introduction

A large class of problems in random geometry is concerned with the collocation of points in high-dimensional space. Applications range from optimization of financial portfolios [1], binary classifications of data strings [2] and optimal stategies in game theory [3] to the existence of non-negative solutions to systems of linear equations [4,5], the emergence of cooperation in competitive ecosystems [6,7], and linear programming with random parameters [8]. It is frequently relevant to consider the case where both the number of points *T* and the dimension of space *N* tend to infinity. This limit is often characterized by abrupt qualitative changes reminiscent of phase transitions when an external parameter or the ratio T/N vary and cross a critical value. At the same time, this high-dimensional case is amenable to methods from the statistical mechanics of disordered systems offering additional insight.

Some results obtained in different disciplines are closely related to each other without the connection always being appreciated. In the present paper, we discuss some particular cases. We will show that the boundedness of the expected maximal loss, as well as the possibility of zero variance of a random financial portfolio is closely related to the existence of a linear separable binary coloring of random points called a dichotomy. Moreover, we point out the connection with the existence of non-negative solutions to systems of linear equations and with mixed strategies in zero-sum games. On a more technical level and for the above-mentioned limit of large instances in high-dimensional spaces, we also make contact between replica calculations performed for different problems in different fields.

In addition to uncovering the common random geometrical background of seemingly very different problems, our comparative analysis sheds light on each of them from various angles and points to ramifications in their respective fields.



Citation: Prüser, A.; Kondor, I.; Engel, A. Aspects of a Phase Transition in High-Dimensional Random Geometry. *Entropy* **2021**, *23*, 805. https://doi.org/10.3390/ e23070805

Academic Editors: Ryszard Kutner, H. Eugene Stanley and Christophe Schinckus

Received: 10 May 2021 Accepted: 17 June 2021 Published: 24 June 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).

2. Dichotomies of Random Points

Consider an *N*-dimensional Euclidean space with a fixed coordinate system. Choose *T* points in this space and color them either black or white. The coloring is called a dichotomy if a hyperplane through the origin of the coordinate system exists that separates black points from white ones, see Figure 1.

To avoid special arrangements like all points falling on one line, the points are required to be in what is called a general position: the position vectors of any subset of N points should be linearly independent. Under this rather mild prerequisite, the number C(T, N) of dichotomies of T points in N dimensions only depends on T and N and not on the particular location of the points. This remarkable result was proven in several works, among them a classical paper by Cover [2]. Establishing a recursion relation for C(T, N), the explicit result was derived:





Figure 1. Two colorings of three points in two dimensions. In the **left** one, black and white points can be separated by a line through the origin; this coloring therefore represents a dichotomy. For the **right** one, no such separating line exists.

If the coordinates of the points are chosen at random from a continuous distribution, the points are in a general position with the probability one. Since there are in total 2^T different binary colorings of these points and only C(T, N) of them are dichotomies, we find for the probability that *T* random points in *N* dimensions with random coloring form a dichotomy with the cumulative binomial distribution:

$$P_{\rm d}(T,N) = \frac{C(T,N)}{2^T} = \frac{1}{2^{T-1}} \sum_{i=0}^{N-1} \binom{T-1}{i}.$$
(2)

Hence, $P_d(T, N) = 1$ for $T \le N$, $P_d(T, N) = 1/2$ for T = 2N and $P_d(T, N) \to 0$ for $T \to \infty$. The transition from $P \simeq 1$ at T = N to $P \simeq 0$ at large T becomes sharper with increasing N. This is clearly seen when considering the case of constant ratio

$$:=\frac{T}{N}$$
(3)

between the number of points and the dimension of space for different values of *N*, which shows an abrupt transition at $\alpha_c = 2$ for $N \rightarrow \infty$, cf. Figure 2.

α

For later convenience, it is useful to reformulate the condition for a certain coloring to be a dichotomy in different ways. Let us denote the position vector of point t, t = 1, ..., T, by $\boldsymbol{\xi}^t \in \mathbb{R}^N$ and its coloring by the binary variable $\zeta^t = \pm 1$. If a separating hyperplane exists, it has a normal vector $\mathbf{w} \in \mathbb{R}^N$ that fulfills

$$\zeta^t = \operatorname{sign}(\mathbf{w} \cdot \boldsymbol{\xi}^t), \qquad t = 1, \dots, T, \tag{4}$$

where we define sign(x) = 1 for $x \ge 0$ and sign(x) = -1 otherwise. With the abbreviation

$$\mathbf{r}^t := \zeta^t \boldsymbol{\xi}^t, \tag{5}$$

Equation (4) translates into $\mathbf{w} \cdot \mathbf{r}^t \ge 0$ for all t = 1, ..., T which for points in a general position, is equivalent to the somewhat stronger condition

$$\mathbf{w} \cdot \mathbf{r}^t > 0, \qquad t = 1, \dots, T. \tag{6}$$

A certain coloring ζ^t of points ξ^t is hence a dichotomy if a vector **w** exists such that (6) is fulfilled, that is, if its scalar product with all vectors \mathbf{r}^t is positive. This is quite intuitive, since by going from the vectors ξ^t to \mathbf{r}^t according to the (5), we replace all points colored black by their white-colored mirror images (or vice versa). If we started out with a dichotomy, after the transformation, all points will lie on the same side of the separating hyperplane. The meaning of Equation (6) is clear: For *T* random points in *N* dimensions with coordinates chosen independently from a symmetric distribution, there exists with probability $P_d(T, N)$ a hyperplane such that all these points lie on the same side of the hyperplane. This formulation will be crucial in Section 3 to relate dichotomies to bounded cones characterizing financial portfolios.



Figure 2. Probability $P_d(T, N)$ that *T* randomly colored points in a general position in *N*-dimensional space form a dichotomy as a function of the ratio α between *T* and *N* for different values of *N*. The transition between the limiting values P = 1 at $\alpha = 1$ and P = 0 at large α becomes increasingly sharp when *N* grows.

Singling out one particular point s = 1, ..., T, this in turn implies that there is, for any choice of *s*, a vector **w** with

$$\mathbf{w} \cdot \mathbf{r}^t > 0, \quad t = 1, \dots, T, t \neq s \quad \text{and} \quad \mathbf{w} \cdot (-\mathbf{r}^s) < 0.$$
 (7)

Consider now all vectors $\bar{\mathbf{r}}$ of the form

$$\bar{\mathbf{r}} = \sum_{t \neq s} c^t \mathbf{r}^t, \quad \text{with} \quad c^t \ge 0, \ t = 1, \dots, T, t \neq s,$$
(8)

that is, all vectors that may be written as a linear combination of the \mathbf{r}^t with $t \neq s$ and all expansion parameters c^t being non-negative. The set of these vectors $\mathbf{\bar{r}}$ is called the *non-negative cone* of the \mathbf{r}^t , $t \neq s$. Equation (7) then means that $-\mathbf{r}^s$ cannot be an element of this non-negative cone. This is clear since the hyperplane perpendicular to \mathbf{w} separates $-\mathbf{r}^s$ from this very cone, an observation that is known as Farkas' lemma [9]. Therefore, if a set of vectors \mathbf{r}^t forms a dichotomy no mirror image $-\mathbf{r}^s$ of any of them may be written as a linear combination of the remaining ones with non-negative expansion coefficients

$$\sum_{t \neq s} c^t \mathbf{r}^t \neq -\mathbf{r}^s, \qquad \forall c^t \ge 0.$$
⁽⁹⁾

Finally, adding \mathbf{r}^{s} to both sides of (9), we find

$$\sum_{t} c^{t} \mathbf{r}^{t} \neq \mathbf{o}, \quad \text{with} \quad c^{t} \ge 0, \ t = 1, \dots, T, \quad \text{and} \quad \sum_{t} c^{t} > 0, \quad (10)$$

where **o** denotes the null vector in *N* dimensions. Given *T* points \mathbf{r}^t in *N* dimensions forming a dichotomy, it is therefore impossible to find a nontrivial linear combination of these vectors with non-negative coefficients that equals the null vector.

Additionally, this corollary to the Cover result is easily intuitively understood. Assume there were some coefficients $c^t \ge 0$ that were not all zero at the same time, and that realize

$$\sum_{t} c^{t} \mathbf{r}^{t} = \mathbf{o}.$$
 (11)

If the points \mathbf{r}^t form a dichotomy, then according to (6), there is a vector \mathbf{w} that makes a positive scalar product with all of them. Multiplying (11) with this vector, we immediately arrive at a contradiction, since the l.h.s. of this equation is positive and the r.h.s. is zero.

Note that the inverse of (10) is also true: if the points do not form a dichotomy, a decomposition of the null vector of the type (11) can always be found. This is related to the fact that the non-negative cone of the corresponding position vectors is the complete \mathbb{R}^N . For if there were a vector $\mathbf{b} \in \mathbb{R}^N$ that lies not in this cone by Farkas' lemma, there would be a hyperplane separating the cone from **b**. However, the very existence of this hyperplane would qualify the points \mathbf{r}^t to be a dichotomy in contradiction to what was assumed.

In the limit $N \to \infty$, $T \to \infty$ with $\alpha = T/N$, keeping the problem of random dichotomies constant can be investigated within statistical mechanics. To make this connection explicit, we first note that no inequality in (6) is altered if **w** is multiplied by a positive constant. To decide whether an appropriate vector **w** fulfilling (6) may be found or not, it is hence sufficient to study vectors of a given length. It is convenient to choose this length as \sqrt{N} , requiring

$$\sum_{i=1}^{N} w_i^2 = N.$$
(12)

Next, we introduce for each realization of the random vectors \mathbf{r}^t an energy function

$$E(\mathbf{w}) := \sum_{t=1}^{T} \Theta\left(-\sum_{i} w_{i} r_{i}^{t}\right), \qquad (13)$$

where $\Theta(x) = 1$ if x > 0, and $\Theta(x) = 0$; otherwise it is the Heaviside step function. This energy is nothing but the number of points violating (6) for a given vector **w**. Our central quantity of interest is the entropy of the groundstate of the system, that is, the logarithm of the fraction of points on the sphere defined by (12) that realize zero energy:

$$S(\kappa, \alpha) := \lim_{N \to \infty} \frac{1}{N} \ln \frac{\int \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i^2 - N) \,\prod_{t=1}^{\alpha N} \Theta(\sum_i w_i r_t^t - \kappa)}{\int \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i^2 - N)}.$$
(14)

Here, $\delta(x)$ denotes the Dirac δ -function, and we have introduced the positive stability parameter κ to additionally sharpen the inequalities (6).

The main problem in the explicit determination of $S(\kappa, \alpha)$ is its dependence on the many random parameters r_i^t . Luckily, for large values of N deviations of S from its typical value, S_{typ} becomes extremely rare and, moreover, this typical value is given by the average over the realizations of the r_i^t :

$$S_{\text{typ}}(\kappa, \alpha) = \langle \langle S(\kappa, \alpha) \rangle \rangle.$$
(15)

The calculation of this average was performed by a classical calculation [10] which gave rise to the result:

$$S_{\text{typ}}(\kappa,\alpha) = \exp_{q} \left[\frac{1}{2} \ln(1-q) + \frac{q}{2(1-q)} + \alpha \int Dt \ln H\left(\frac{\kappa - \sqrt{q}t}{\sqrt{1-q}}\right) \right], \quad (16)$$

where the extremum is over the auxiliary quantity q, and we have used the shorthand notations

$$Dt := \frac{dt}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} \quad \text{and} \quad H(x) := \int_x^\infty Dt.$$
(17)

More details of the calculation may be found in the original reference, and in chapter 6 of [11]. Appendix A contains some intermediate steps for a closely related analysis.

Studying the limit $q \rightarrow 1$ of (16) reveals

$$S_{\text{typ}}(\kappa, \alpha) \begin{cases} > -\infty & \text{if} \quad \alpha < \alpha_c(\kappa) \\ \to -\infty & \text{if} \quad \alpha > \alpha_c(\kappa), \end{cases}$$
(18)

corresponding to a sharp transition from solvability to non-solvability at a critical value $\alpha_c(\kappa)$. This is because $\kappa = 0$ finds $\alpha_c = 2$ in agreement with (2), cf. Figure 2.

Note that Cover's result (2) holds for all values of *T* and *N*, whereas the statistical mechanics analysis is restricted to the thermodynamic limit $N \rightarrow \infty$. On the other hand, the latter can deal with all values of the stability parameter κ , whereas no generalization of Cover's approach to the case $\kappa \neq 0$ is known.

3. Phase Transitions in Portfolio Optimization under the Variance and the Maximal Loss Risk Measure

3.1. Risk Measures

The purpose of this subsection is to indicate the financial context, in which the geometric problem discussed in this paper appears. A portfolio is the weighted sum of financial assets. The weights represent the parts of the total wealth invested in the various assets. Some of the weights are allowed to be negative (short positions), but the weights sum to 1; this is called the budget constraint. Investment carries risk, and higher returns usually carry higher risk. Portfolio optimization seeks a trade-off between risk and return by the appropriate choice of the portfolio weights. Markowitz was the first to formulate the portfolio choice as a risk-reward problem [12]. Reward is normally regarded as the expected return on the portfolio. Assuming return fluctuations to be Gaussian-distributed random variables, portfolio variance offered itself as the natural risk measure. This setup made the optimization of portfolios a quadratic programming problem, which, especially in the case of large institutional portfolios, posed a serious numerical difficulty in its time. Another critical point concerning variance as a risk measure was that variance is symmetric in gains and losses, whereas investors are believed not to be afraid of big gains, only big losses. This consideration led to the introduction of downside risk measures, starting already with the semivariance [13]. Later it was recognized that the Gaussian assumption was not realistic, and alternative risk measures were sought to grasp the risk of rare but large events, and also to allow risk to be aggregated across the ever-increasing and increasingly heterogeneous institutional portfolios. Around the end of the 1980s, Value at Risk (VaR) was introduced by JP Morgan [14], and subsequently it was widely spread over the industry by their RiskMetrics methodology [15]. VaR is a high quantile, a downside risk measure (note that in the literature, the profit and loss axis is often reflected, so that losses are assigned a positive sign. It is under this convention that VaR is a high quantile, rather than a low one). It soon came under academic criticism for its insensitivity to the details of the distribution beyond the quantile, and for its lack of sub-additivity. Expected Shortfall (ES), the average loss above the VaR quantile, appeared around the turn of the century [16]. An axiomatic approach to risk measures was proposed by Artzner et al. [17] who introduced a set of postulates which any coherent risk measure was required to satisfy. ES turned out to be coherent [18,19] and was strongly advocated by academics. After a long debate, international regulation embraced it as the official risk measure in 2016 [20].

The various risk measures discussed all involved averages. Since the distributions of financial data are not known, the relative price movements of assets are observed at a number T of time points, and the true averages are replaced by empirical averages from these data. This works well if T is sufficiently large; however, in addition to all the aforementioned problems, a general difficulty of portfolio optimization lies in the fact that

the dimension *N* of institutional portfolios (the number of different assets) is large, but the number *T* of observed data per asset is never large enough, due to lack of stationarity of the time series and the natural limits (transaction costs, technical difficulties of rebalancing) on the sampling frequency. Therefore, portfolio optimization in large dimensions suffers from a high degree of estimation error, which renders the exercise more or less illusory (see e.g., [21]). Estimation of returns is even more error-prone than the risk part, so several authors disregard the return completely, and seek the minimum risk portfolio (e.g., [22–24]). We follow the same approach here.

In the two subsections that follow, we also assume that the returns are independent, symmetrically distributed random variables. This is, of course, not meant to be a realistic market model, but it allows us to make an explicit connection between the optimization of the portfolio variance under a constraint excluding short positions and the geometric problem of dichotomies discussed in Section 2. This is all the more noteworthy because analytic results are notoriously scarce for portfolio optimization with no short positions. We note that similar simplifying assumptions (Gaussian fluctuations, independence) were built into the original JP Morgan methodology, which was industry standard in its time, and influences the thinking of practitioners even today.

3.2. Vanishing of the Estimated Variance

We consider a portfolio of *N* assets with weights w_i , i = 1, ..., N. The observations r_i^t of the corresponding returns at various times t = 1, ..., T are assumed to be independent, symmetrically distributed random variables. Correspondingly, the average value of the portfolio is zero. Its variance is given by

$$\sigma_p^2 = \frac{1}{T} \sum_t \left(\sum_i w_i r_i^t \right)^2 = \sum_{i,j} w_i w_j \frac{1}{T} \sum_t r_i^t r_j^t =: \sum_{i,j} w_i w_j C_{ij}, \tag{19}$$

where C_{ij} denotes the covariance matrix of the observations. Note that the variance of a portfolio optimized in a given sample depends on the sample, so it is itself a random variable.

The variance of a portfolio obviously vanishes if the returns are fixed quantities that do not fluctuate. This subsection is not about such a trivial case. We shall see, however, that the variance optimized *under a no-short constraint* can vanish with a certain probability if the dimension N is larger than the number of observations T.

The rank of the covariance matrix is the smaller of N and T, and for $N \leq T$ the estimated variance is positive with the probability one. Thus, the optimization of variance can always be carried out as long as the number of observations T is larger than the dimension N, albeit with an increasingly larger error as T/N decreases. For large N and T and fixed $\alpha = T/N$, the estimation error increases as $\alpha/(\alpha - 1)$ with decreasing α and diverges at $\alpha \downarrow 1$ [25,26]. The divergence of the estimation error can be regarded as a phase transition. Below the critical value $\alpha_d := 1$, the optimization of variance becomes impossible. Of course, in practice, one never has such an optimization task without some additional constraints. Note that because of the possibility of short-selling (negative portfolio weights), the budget constraint (a hyperplane) in itself is not sufficient to forbid the appearance of large positive and negative positions, which then destabilize the optimization. In contrast, any constraint that makes the allowed weights finite can act as a regularizer. The usual regularizers are constraints on the norm of the portfolio vector. It was shown in [27,28] how liquidity considerations naturally lead to regularization. Ridge regression (a constraint on the ℓ_2 norm of the portfolio vector) prevents the covariance matrix from developing zero eigenvalues, and, especially in its nonlinear form [29], results in very satisfactory out-of-sample performance.

An alternative is the ℓ_1 regularizer, of which the exclusion of short positions is a special case. Together with the budget constraint, it prevents large sample fluctuations of the weights. Let us then impose the no-short ban, as it is indeed imposed in practice on a number of special portfolios (e.g., on pension funds), or, in episodes of crisis, on the

whole industry. The ban on short-selling extends the region where the variance can be optimized, but below $\alpha = 1$ the optimization acquires a probabilistic character in that the regularized variance vanishes with a certain probability, and the optimization can only be carried out when it is positive. (Otherwise, there is a continuum of solutions, namely any combination of the eigenvectors belonging to zero eigenvalues, which makes the optimized variance zero).

Interestingly, the probability of the variance vanishing is related to the problem of random dichotomies in the following way. For the portfolio variance (19) to become zero, we need to have

$$\sum_{i} w_i r_i^t = 0 \tag{20}$$

for all *t*. If we interchange *t* and *i*, we see that according to (11), this is possible as long as the *N* points in \mathbb{R}^T with position vectors $\vec{r}_i := \{r_i^t\}$ do not form a dichotomy. Hence, the probability for zero variance is from (2)

$$P_{\rm zv}(T,N) = 1 - P_{\rm d}(N,T) = 1 - \frac{1}{2^{N-1}} \sum_{i=0}^{T-1} \binom{N-1}{i} = \frac{1}{2^{N-1}} \sum_{i=T}^{N-1} \binom{N-1}{i}.$$
 (21)

Therefore, the probability of the variance vanishing is almost 1 for small α , decreases to the value 1/2 at $\alpha = 1/2$, decreases further to 0 as α increases to 1, and remains identically zero for $\alpha > 1$ [30,31]. This is similar but also somewhat complementary to the curve shown in Figure 2. Equation (21) for the vanishing of the variance was first written up in [30,31] on the basis of analogy with the minimax problem to be considered below, and it was also verified by extended numerical simulations. The above link to the Cover problem is a new result, and it is rewarding to see how a geometric proof establishes a bridge between the two problems.

In [30,31], an intriguing analogy with, for example, the condensed phase of an ideal Bose gas was pointed out. The analogous features are the vanishing of the chemical potential in the Bose gas, resp. the vanishing of the Lagrange multiplier enforcing the budget constraint in the portfolio problem; the onset of Bose condensation, resp. the appearance of zero weights ("condensation" of the solutions on the coordinate planes) due to the no-short constraint; the divergence of the transverse susceptibility, and the emergence of zero modes in both models.

3.3. The Maximal Loss

The introduction of the Maximal Loss (ML) or minimax risk measure by Young [32] in 1998 was motivated by numerical expediency. In contrast to the variance whose optimization demands a quadratic program, ML is constructed such that it can be optimized by linear programming, which could be performed very efficiently even on large datasets already at the end of the last century. Maximal Loss combines the worst outcomes of each asset and seeks the best combination of them. This may seem to be an over-pessimistic risk measure, but there are occasions when considering the worst outcomes is justifiable (think of an insurance portfolio in the time of climate change), and, as will be seen, the present regulatory market risk measure is not very far from ML.

Omitting the portfolio's return again and focusing on the risk part, the maximal loss of a portfolio is given by

$$ML := \min_{\mathbf{w}} \max_{1 \le t \le T} \left(-\sum_{i} w_{i} r_{i}^{t} \right)$$
(22)

with the constraint

$$\sum_{i} w_i = N. \tag{23}$$

We are interested in the probability $P_{ML}(T, N)$ that this minimax problem is feasible, that is, ML does not diverge to $-\infty$. To this end, we first eliminate the constraint (23) by putting

$$w_N = N - \sum_{i=1}^{N-1} w_i.$$
 (24)

This results in

$$ML := \min_{\tilde{\mathbf{w}}} \max_{1 \le t \le T} \left(-\sum_{i=1}^{N-1} w_i (r_i^t - r_N^t) - Nr_N^t \right) =: \min_{\tilde{\mathbf{w}}} \max_{1 \le t \le T} \left(-\sum_{i=1}^{N-1} w_i \tilde{r}_i^t - Nr_N^t \right)$$
(25)

with $\mathbf{\tilde{w}} := \{w_1, \dots, w_{N-1}\} \in \mathbb{R}^{N-1}$ and $\mathbf{\tilde{r}}^t := \{r_1^t - r_N^t, \dots, r_{N-1}^t - r_N^t\} \in \mathbb{R}^{N-1}$. For ML to stay finite for all choices of $\mathbf{\tilde{w}}$, the *T* random hyperplanes with normal vectors $\mathbf{\tilde{r}}^t$ have to form a bounded cone. If the points $\mathbf{\tilde{r}}^t$ form a dichotomy, then according to (6), there is a vector $\mathbf{W} \in \mathbb{R}^{N-1}$ with $\mathbf{W} \cdot \mathbf{\tilde{r}}^t > 0$ for all *t*. Since there is no constraint on the norm of $\mathbf{\tilde{w}}$, the maximal loss (25) can become arbitrarily small for $\mathbf{\tilde{w}} = \lambda \mathbf{W}$ and $\lambda \to \infty$. The cone then is not bounded. We therefore find

$$P_{\rm ML}(T,N) = P_{\rm d}(T,N-1) = \frac{1}{2^{T-1}} \sum_{i=0}^{N-2} \binom{T-1}{i}$$
(26)

for the probability that ML cannot be optimized.

In the limit $N, T \to \infty$ with $\alpha = T/N$ kept finite, (25) displays the same abrupt change as in the problem of dichotomies, a phase transition at $\alpha_c = 2$. Note that this is larger than the critical point $\alpha_d = 1$ of the unregularized variance, which is quite natural, since the ML uses only the extremal values in the data set. The probability for the feasibility of ML was first written up without proof in [1], where a comparative study of the noise sensitivity of four risk measures, including ML, was performed. There are two important remarks we can make at this point. First, the geometric consideration above does not require any assumption about the data generating process; as long as the the returns are independent, they can be drawn from any symmetric distribution without changing the value of the critical point. This is a special case of the universality of critical points discovered by Donoho and Tanner [33].

The second remark is that the problem of bounded cones is closely related to that of bounded polytopes [34]. The difference is just the additional dimension of the ML itself. If the random hyperplanes perpendicular to the vectors \tilde{r}^t form a bounded cone for ML according to (25), then they will trace out a bounded polytope on hyperplanes perpendicular to the ML axis at sufficiently high values of ML. In fact, after the replacement $N-1 \rightarrow N$ Equation (26) coincides with the result in Theorem 4 of [34] for the probability of *T* random hyperplanes forming a bounded polytope in *N* dimensions (there is a typo in Theorem 4 in [34]; the summation has to start at *i* = 0). The close relationship between the ML problem and the bounded polytope problem, on the one hand, and the Cover problem on the other hand, was apparently not clarified before.

If we spell out the financial meaning of the above result, we are led to interesting ramifications. To gain an intuition, let us consider just two assets, N = 2. If asset 1 produces a return sometimes above, sometimes below that of asset 2, then the minimax problem will have a finite solution. If, however, asset 1 dominates asset 2 (i.e., yields a return which is at least as large, and, at least at one time point, larger, than the return on asset 2 in a given sample), then, with unlimited short positions allowed, the investor will be induced to take an arbitrarily large long position in asset 1 and go correspondingly short in asset 2. This means that the solution of the minimax problem will run away to infinity, and the risk of ML will be equal to minus infinity [1]. The generalization to N assets is immediate: if among the assets there is one that dominates the rest, or there is a combination of assets that dominates some of the rest, the solution will run away to infinity, and ML will take the value of $-\infty$. This scenario corresponds to an arbitrage, and the investor gains an arbitrarily large profit without risk [35]. Of course, if such a dominance is realized in one

given sample, it may disappear in the next time interval, or the dominance relations can rearrange to display another mirage of an arbitrage.

Clearly, the ML risk measure is unstable against these fluctuations. In practice, such a brutal instability can never be observed, because there are always some constraints on the short positions, or groups of assets corresponding to branches of industries, geographic regions, and so forth. These constraints will prevent instabilities from taking place, and the solution cannot run away to infinity, but will go as far as allowed by the constraints and then stick to the boundary of the allowed region. Note, however, that in such a case, the solution will be determined more by the constraints (and ultimately by the risk manager imposing the constraints) rather than by the structure of the market. In addition, in the next period, a different configuration can be realized, so the solution will jump around on the boundary defined by the constraints.

We may illustrate the role of short positions for the instability of ML further by investigating the case of portfolio weights w_i that have to be larger than a threshold $\gamma \leq 0$. For $\gamma \to -\infty$, there are no restrictions on short positions, whereas $\gamma = 0$ corresponds to a complete ban on them. For $N, T \to \infty$ with fixed $\alpha = T/N$, the problem may be solved within the framework of statistical mechanics. The minimax problem for ML is equivalent to the following problem in linear programming: minimize the threshold variable κ under the constraints (23), $w_i \geq \gamma$, and

$$-\sum_{i} w_{i} r_{i}^{t} \leq \kappa \quad \forall t = 1, \dots, T.$$
⁽²⁷⁾

Similarly to (14), the central quantity of interest is

$$\Omega(\kappa,\gamma,\alpha) = \frac{\int_{\gamma}^{\infty} \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i - N) \prod_{t=1}^{\alpha N} \Theta(\sum_i w_i r_i^t + \kappa)}{\int_{\gamma}^{\infty} \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i - N)},\tag{28}$$

giving the fractional volume of points on the simplex defined by (23) that fulfill all constraints (27). For given α and γ , we decrease κ down to the point κ_c , where the typical value of this fractional volume vanishes. The ML is then given by $\kappa_c(\alpha, \gamma)$.

Some details of the corresponding calculations are given in the Appendix A. In Figure 3, we show some results. As discussed above, the divergence of ML for $\alpha < 2$ is indeed formally eliminated for all $\gamma > -\infty$, and the functions ML(α ; γ) smoothly interpolate between the cases $\gamma = 0$ and $\gamma \rightarrow -\infty$. However, the situation is now even more dangerous, since the unreliability of ML as a risk measure for small α remains without being deducible from its divergence.



Figure 3. *Left*: The Maximal Loss $ML = \kappa_c$ as a function of α . The analytical results (solid line) are compared to simulation results (circles) with N = 200 averaged over 100 samples. The symbol size corresponds to the statistical error. *Right*: Same as left with largely extended axis of ML.

The recognition of the instability of ML as a dominance problem has proved very fruitful and led to a series of generalizations. First, it was realized [1] that the instability of the expected shortfall, of which ML is an extreme special case, has a very similar geometric origin. (The current regulatory ES is the expected loss above a 97.5% quantile, whereas ML corresponds to 100%.) Both ES and ML are so-called coherent risk measures [17], and it was proved [35] that the root of this instability lies in the coherence axioms themselves, so every coherent risk measure suffers from a similar instability. Furthermore, it was proved [35] that the existence of a dominant/dominated pair of assets in the portfolio was a necessary and sufficient condition for the instability of ML, whereas it was only sufficient for other coherent risk measures. It follows that in terms of the variable α used in this paper (which is the reciprocal of the aspect ratio N/T used in some earlier works, such as [35–37]), the critical point of ML is a lower bound for the critical points of other coherent measures. Indeed, the critical line of ES was found to lie above the ML critical value of $\alpha_c = 2$ [36]. Value at Risk is not a coherent measure and can violate convexity, so it is not amenable to a similar study of its critical point. However, parametric VaR (that is, the quantile where the underlying distribution is given, only its expectation value and variance is determined from empirical data) is convex, and it was shown to possess a critical line that runs above that of ES [37]. The investigation of the semi-variance yielded similar results [37]. It seems, then, that the geometrical analysis of ML provides important information for a variety of risk measures, including some of the most widely used measures in the industry (VaR and ES), and also other downside risk measures.

4. Related Problems

In this section, we list a few problems from different fields of mathematics and physics that are linked to the random coloring of points in high-dimensional space and point out their connection with the questions discussed above.

4.1. Binary Classifications with a Perceptron

Feed-forward networks of formal neurons perform binary classifications of input data [38]. The simplest conceivable network of this type—the perceptron—consists of just an input layer of N units ξ_i and a single output bit $\zeta = \pm 1$ [39]. Each input ξ_i is directly connected to the output by a real valued coupling w_i . The output is computed as the sign of the weighted inputs

$$\zeta = \operatorname{sign}\left(\sum_{i=1}^{N} w_i \,\xi_i\right). \tag{29}$$

Consider now a family of random inputs $\{\xi_i^t\}$, t = 1, ..., T and ask for the probability $P_p(T, N)$ that the perceptron is able to implement a randomly chosen binary classification $\{\zeta^t\}$ of these inputs. Interpreting the vectors $\boldsymbol{\xi}^t := \{\xi_i^t\}$ as position vectors of T points in N dimensions and the required classifications ζ^t as a black/white coloring, we hence need to know the probability that this particular coloring is a dichotomy. Indeed, if a hyperplane exists that separates black points from white ones, it has a normal vector \boldsymbol{w} that gives a suitable choice for the perceptron weights to get all classifications right. Therefore, we have

$$P_{\rm p}(T,N) = P_{\rm d}(T,N) = \frac{1}{2^{T-1}} \sum_{i=0}^{N-1} \binom{T-1}{i}.$$
(30)

In the thermodynamic limit $N, T \rightarrow \infty$, this problem, together with a variety of modifications, can be analyzed using methods from the statistical mechanics of disordered systems along the lines of Equations (14)–(16), see [11].

4.2. Zero-Sum Games with Random Pay-Off Matrices

In game theory, two or more players choose among different strategies at their disposal and receive a pay-off (that may be negative) depending on the choices of all participating players. A particularly simple situation is given by a zero-sum game between two players, where one player's profit is the other player's loss. If the first player may choose among N strategies and the second among T, the setup is defined by an $N \times T$ pay-off matrix r_i^t , giving the reward for the first player if he plays strategy i and his opponent strategy t. Barring rare situations in which it is advantageous for one or both players to always choose one and the same strategy, it is known from the classical work of Morgenstern and von Neumann [40] that the best the players can do is to choose at random with different probabilities among their available strategies. The set of these probabilities p_i and q^t , respectively, is called a mixed strategy.

For large numbers of available strategies, it is sensible to investigate typical properties of such mixed strategies for random pay-off matrices. This can be done in a rather similar way to the calculation of ML presented in the Appendix A of the present paper [3]. One interesting result is that an extensive part of the probabilities p_i and q^t forming the optimal respective mixed strategies have to be identically zero: for both players, there are strategies they should never touch.

4.3. Non-Negative Solutions to Large Systems of Linear Equations

Consider a random $N \times T$ matrix r_i^t and a random vector $\mathbf{b} \in \mathbb{R}^N$. When will the system of linear equations

$$\sum_{t} r_i^t x^t = b_i, \qquad i = 1, \dots, N \tag{31}$$

typically have a solution with all x^t being non-negative? This question is related to the optimization of financial portfolios under a ban of short-selling as discussed above, and also occurs when investigating the stability of chemical or ecological problems [6,41]. Here, the x^t denotes concentrations of chemical or biological species, and hence has to be non-negative. Similar to optimal mixed strategies considered in the previous subsection, the solution typically has a number of entries x^t that are strictly zero (species that died out), the remaining ones being positive (surviving species). Again for $T = \alpha N$ and $N \to \infty$, a sharp transition at a critical value α_c separates situations with typically no non-negative solution from those in which typically such a solution can be found [4].

To make contact with the cases discussed before, it is useful to map the problem to a dual one by again using Farkas' lemma. Let us denote by

$$\bar{\mathbf{r}} = \sum_{t} c^{t} \mathbf{r}^{t}, \qquad c^{t} \ge 0, \ t = 1, \dots, T$$
(32)

the vectors in the non-negative cone of the column vectors \mathbf{r}^t of matrix r_i^t . It is clear that (31) has a non-negative solution \mathbf{x} if \mathbf{b} belongs to this cone, and that no such solution exists if \mathbf{b} lies outside the cone. In the latter case, however, there must be a hyperplane separating \mathbf{b} from the cone. Denoting the normal of this hyperplane by \mathbf{w} , we hence have the following duality: either the system (31) has a non-negative solution \mathbf{x} , or there exists a vector \mathbf{w} with

$$\mathbf{w} \cdot \mathbf{r}^t \ge 0$$
 $t = 1, \dots, T$ and $\mathbf{w} \cdot \mathbf{b} < 0.$ (33)

If the r_i^t is drawn independently from a distribution with finite first and second cumulant R and σ_r^2 , respectively, and the components b_i are independent random numbers with average B and variance σ_b^2/N , the dual problem (33) may be analyzed along the lines of (14)–(16). The result for the typical entropy of solution vectors **w** reads [4]

$$S_{\text{typ}}(\gamma,\alpha) = \exp_{q,\kappa} \left[\frac{1}{2} \ln(1-q) + \frac{q}{2(1-q)} - \frac{\kappa^2 \gamma}{2(1-q)} + \alpha \int Dt \ln H\left(\frac{\kappa - \sqrt{q}t}{\sqrt{1-q}}\right) \right], \quad (34)$$

where the parameter

$$\gamma := \left(\frac{B\,\sigma_r}{R\,\sigma_b}\right)^2\tag{35}$$

characterizes the distributions of r_i^t and b_i . The main difference to (16) is the additional extremum over κ regularized by the penalty term proportional to κ^2 . Considering the limit $q \to 1$ in (34), it is possible to determine the critical value $\alpha_c(\gamma)$ bounding the region where typically no solution **w** may be found. For nonrandom **b**, that is, $\sigma_b \to 0$ implying $\gamma \to \infty$, we find back the Cover result $\alpha_c = 2$.

The problem is closely related to a phase transition found recently in MacArthur's resource competition model [4,6,7], in which a community of purely competing species builds up a collective cooperative phase above a critical threshold of the biodiversity.

5. Discussion

In this paper, we have reviewed various problems from different disciplines, including high-dimensional random geometry, finance, binary classification with a perceptron, game theory, and random linear algebra, which all have at their root the problem of dichotomies, that is, the linear separability of points carrying a binary label and scattered randomly over a high-dimensional space. No doubt there are several further problems belonging to this class; those that spring to mind are theoretical ecology alluded to at the end of the previous Section, or linear programming with random parameters [8]. Some of these conceptual links are obvious, and have been known for decades (for example, the link between dichotomies and the perceptron), and others are far less clear at first sight, such as the relationship with the two finance problems discussed in Section 3. We regard as one of the merits of this paper the establishment of this network of conceptual connections between seemingly faraway areas of study. Apart from the occasional use of the heavy machinery of the replica theory, in most of the paper we offered transparent geometric arguments, where our only tool was basically the Farkas' lemma.

The phase transitions we encountered in all of the problems discussed here are similar in spirit to the geometric transitions discovered by Donoho and Tanner [33] and interpreted at a very high level of abstraction by [42]. One of the central features of these transitions is the universality of the critical point. This universality is different from the one observed in the vicinity of continuous phase transitions in physics, where the value of the critical point can vary widely, even between transitions belonging to the same universality class. The universality in physical phase transitions is a property of the critical indices and other critical parameters. Critical indices also appear in our abstract geometric problems, and they are universal, but we omitted their discussion which might have led far from the main theme.

At the bottom of our geometric problems, there is the optimization of a convex objective function (which is, by the way, the key to the replica symmetric solutions we found). The recent evolution of neural networks, machine learning, and artificial intelligence is mainly concerned with a radical lack of convexity, which points to the direction in which we may try to extend our studies. Another simplifying feature we exploited was the independence of the random variables. The moment that correlations appear, these problems become hugely more complicated. We left this direction for future exploration. However, it is evident that progress in any of these problems will induce progress in the other fields, and we feel that revealing their fundamental unity may help the transfer of methods and ideas between these fields. This may be the most important achievement of this analysis.

Author Contributions: Conceptualization, I.K. and A.E.; formal analysis, A.P., I.K. and A.E.; software, A.P.; writing—original draft, A.P., I.K. and A.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: A.P. and A.E. are grateful to Stefan Landmann for many interesting discussions.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Replica Calculation of Maximal Loss

In this appendix, we provide some details for the determination of the maximal loss of a random portfolio using the replica trick. The calculation is a generalization of the one presented in [3] for random zero-sum games. A presentation at full length can be found in [43]. As we pointed out in the main text, maximal loss is a special limit of the Expected Shortfall risk measure, corresponding to the so-called confidence level going to 100%. In [44] a detailed study of the behavior of ES was carried out, including the limiting case of maximal loss. That treatment is completely different from the one in here, so the present calculation can be regarded as complementary to that in [44].

The central quantity of interest is the fractional volume

$$\Omega(\kappa,\gamma,\alpha) = \frac{\int_{\gamma}^{\infty} \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i - N) \prod_{t=1}^{\alpha N} \Theta(\sum_i w_i r_i^t + \kappa)}{\int_{\gamma}^{\infty} \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i - N)}$$
(A1)

defined in (28). Although not explicitly indicated, $\Omega(\kappa, \gamma, \alpha)$ depends on all the random parameters r_i^t and is therefore by itself a random quantity. The calculation of its complete probability density $P(\Omega)$ is hopeless but for large N this distribution gets concentrated around the typical value $\Omega_{typ}(\kappa, \gamma, \alpha)$. Because Ω involves a product of many independent random factors this typical value is given by

$$\Omega_{\rm typ}(\kappa,\gamma,\alpha) = e^{\langle \ln \Omega(\kappa,\gamma,\alpha) \rangle \rangle} \tag{A2}$$

rather than by $\langle \langle \Omega(\kappa, \gamma, \alpha) \rangle \rangle$. Here $\langle \langle \dots \rangle \rangle$ denotes the average over the r_i^t . A direct calculation of $\langle \langle \ln \Omega \rangle \rangle$ is hardly possible. It may be circumvented by exploiting the identity

$$\langle \langle \ln(\Omega(\kappa,\gamma,\alpha)) \rangle \rangle = \lim_{n \to 0} \frac{1}{n} [\langle \langle \Omega^n(\kappa,\gamma,\alpha) \rangle \rangle - 1]$$
 (A3)

For natural *n* the determination of $\langle \langle \Omega^n \rangle \rangle$ is feasible. The main problem then is to continue the result to real *n* in order to perform the limit $n \to 0$.

The explicit calculation starts with

$$\left\langle \left\langle \Omega(\kappa,\gamma,\alpha)^n \right\rangle \right\rangle = \left\langle \left\langle \frac{\int_{\gamma}^{\infty} \prod_{i=1}^{N} \prod_{a=1}^{n} dw_i^a \prod_{a=1}^{n} \delta(\sum_i w_i^a - N) \prod_{t=1}^{\alpha N} \prod_{a=1}^{n} \Theta(\sum_i w_i^a r_i^t + \kappa)}{\int_{\gamma}^{\infty} \prod_{i=1}^{N} \prod_{a=1}^{n} dw_i^a \prod_{a=1}^{n} \delta(\sum_i w_i^a - N)} \right\rangle \right\rangle.$$
(A4)

Using

$$\int_{\gamma}^{\infty} \prod_{i=1}^{N} dw_i \,\delta(\sum_i w_i - N) \sim \exp\{N[1 + \ln(1 - \gamma)]\} \tag{A5}$$

for large *N* and representing the δ -functions and Θ -functions by integrals over auxiliary variables E_a , λ_t^a , and y_t^a we arrive at

$$\langle \langle \Omega(\kappa,\gamma,\alpha)^n \rangle \rangle = \exp\{-nN[1+\ln(1-\gamma)]\}$$

$$\times \int_{\gamma}^{\infty} \prod_{i,a} dw_i^a \int \prod_a \frac{dE_a}{2\pi} \exp\left[iN\sum_a E_a\left(\frac{1}{N}\sum_i w_i^a - 1\right)\right]$$

$$\times \int_{-\kappa}^{\infty} \prod_{t,a} d\lambda_t^a \int \prod_{t,a} \frac{dy_t^a}{2\pi} \exp\left(i\sum_{t,a} y_t^a \lambda_t^a\right) \Big\langle \Big\langle \exp(-i\sum_{i,t,a} y_t^a w_i^a r_i^t) \Big\rangle \Big\rangle.$$
(A6)

The average over the r_i^t may now be performed for independent Gaussian r_i^t with average zero and variance $\sigma^2 = 1/N$. The result is valid also for more general distributions. First, multiplying the variance by a constant just rescales the maximal loss but does not influence the optimal **w**. Second, for $N \rightarrow \infty$ only the first two cumulants of the distribution matter due to the central limit theorem. Crucial is, however, the assumption of the r_i^t being independent.

Performing the average we find

$$\left\langle \left\langle \exp\left(-i\sum_{i,t,a} y_t^a w_i^a r_t^i\right) \right\rangle \right\rangle = \prod_{i,t} \left[\int \frac{dr_i^t}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(r_i^t)^2}{2\sigma^2} - ir_i^t \sum_a y_t^a w_i^a\right) \right]$$
$$= \exp\left(-\frac{1}{2N} \sum_{i,t} \sum_{a,b} w_i^a w_i^b y_t^a y_t^b\right).$$
(A7)

To disentangle in (A6) the *w*-integrals from those over λ and *y* we introduce the order parameters

$$q_{ab} = \frac{1}{N} \sum_{i} w_i^a w_i^b, \quad a \ge b \tag{A8}$$

together with the conjugate ones \hat{q}_{ab} . Using standard techniques [11] we end up with

$$\langle \langle \Omega(\kappa,\gamma,\alpha)^n \rangle \rangle = \int \prod_{a \ge b} \frac{dq_{ab} d\hat{q}_{ab}}{2\pi/N} \int \prod_a \frac{dE_a}{2\pi} \times \exp\left\{ -iN \sum_{a \ge b} q_{ab} \hat{q}_{ab} - iN \sum_a E_a - nN[1 + \ln(1-\gamma)] + NG_s + \alpha NG_E \right\},$$
(A9)

where

$$G_{S} = \ln\left[\int_{\gamma}^{\infty} \prod_{a} dw^{a} \exp\left(i\sum_{a\geq b} \hat{q}_{ab}w^{a}w^{b} + i\sum_{a} E_{a}w^{a}\right)\right]$$
(A10)

and

$$G_E = \ln\left[\int_{-\kappa}^{\infty} \prod_a d\lambda^a \int \prod_a \frac{dy^a}{2\pi} \exp\left(-\frac{1}{2}\sum_{a,b} q_{ab}y^a y^b + i\sum_a y^a \lambda^a\right)\right].$$
 (A11)

For $N \rightarrow \infty$ the integrals over the order parameters in (A9) may be calculated using the saddle-point method. The essence of the so-called replica-symmetric ansatz is the assumption that the values of the order parameters at the saddle-point are invariant under permutation of the replica indices *a* and *b*. In [43] arguments are given why the replica-symmetric saddle-point should yield correct results in the present context. We therefore assume for the saddle-point values of the order parameters

$$q_{aa} = q_1 \qquad i\hat{q}_{aa} = -\frac{1}{2}\hat{q}_1 \qquad iE_a = E \qquad \forall a$$

$$q_{ab} = q_0 \qquad i\hat{q}_{ab} = \hat{q}_0 \qquad \qquad \forall a > b.$$
(A12)

which implies various simplifications in (A9)–(A11). Employing standard manipulations [11] we arrive at

$$\langle\langle \Omega(\kappa,\gamma,\alpha)^n\rangle\rangle\sim \exp\bigg\{N\underset{q_0,\hat{q}_0,q_1,\hat{q}_1,E}{\operatorname{extr}}\bigg[-\frac{n(n-1)}{2}q_0\hat{q}_0+\frac{n}{2}q_1\hat{q}_1-nE-n(1+\ln(1-\gamma))+G_S+\alpha G_E\bigg]\bigg\}.$$
(A13)

Using the shorthand notations (17) the functions G_S and G_E are now given by

$$G_{S} = \ln \int Dl \left[\exp\left(\frac{(\sqrt{\hat{q}_{0}}l + E)^{2}}{2(\hat{q}_{0} + \hat{q}_{1})}\right) \sqrt{\frac{2\pi}{\hat{q}_{0} + \hat{q}_{1}}} H \left(-\frac{\sqrt{\hat{q}_{0}}l + E - \gamma(\hat{q}_{0} + \hat{q}_{1})}{\sqrt{\hat{q}_{0} + \hat{q}_{1}}}\right) \right]^{n}$$
(A14)

and

$$G_E = \ln \int Dm \, H \left(\frac{\sqrt{q_0}m - \kappa}{\sqrt{q_1 - q_0}}\right)^n. \tag{A15}$$

We may now treat *n* as a real number and perform the limit $n \rightarrow 0$. In this way we find for the averaged entropy

$$S(\kappa,\gamma,\alpha) := \lim_{N \to \infty} \frac{1}{N} \langle \langle \ln[\Omega(\kappa,\gamma,\alpha)] \rangle \rangle = \lim_{N \to \infty} \frac{1}{N} \lim_{n \to 0} \frac{1}{n} [\langle \langle \Omega(\kappa,\gamma,\alpha)^n \rangle \rangle - 1]$$
(A16)

the expression

$$S(\kappa,\gamma,\alpha) = \operatorname{extr}_{q_0,\hat{q}_0,q_1,\hat{q}_1,E} \left[\frac{q_0\hat{q}_0}{2} + \frac{q_1\hat{q}_1}{2} - E - 1 - \ln(1-\gamma) + \frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln(\hat{q}_0 + \hat{q}_1) \right. \\ \left. + \frac{\hat{q}_0 + E^2}{2(\hat{q}_0 + \hat{q}_1)} + \int Dl\ln H\left(-\frac{\sqrt{\hat{q}_0}l + E - \gamma(\hat{q}_0 + \hat{q}_1)}{\sqrt{\hat{q}_0 + \hat{q}_1}} \right) \right.$$

$$\left. + \alpha \int Dm\ln H\left(\frac{\sqrt{q_0}m - \kappa}{\sqrt{q_1 - q_0}} \right) \right].$$
(A17)

The remaining extremization has to be done numerically. Before embarking on this task it is useful to remember that Ω and *S* are only instrumental in determining the maximal loss which in turn is given by the value κ_c of κ for which Ω tends to zero. At the same time the typical overlap q_0 between two different vectors in Ω has to tend to the self-overlap q_1 . To investigate this limit we replace the order parameter q_1 by

$$v := q_1 - q_0 \tag{A18}$$

and study the saddle-point equations for $v \rightarrow 0$. In this limit it turns out that the remaining order parameters may either also tend to zero or diverge. It is therefore convenient to make the replacements

$$\hat{q}_0 \to \frac{\hat{q}_0}{v^2}, \quad \hat{q}_1 \to \hat{w} := \frac{\hat{q}_1 + \hat{q}_0}{v}, \quad E \to \frac{E}{v}.$$
 (A19)

Rescaled in this way the saddle-point values of the order parameters remain O(1) for $v \to 0$. After some tedious calculations the saddle-point equations acquire the form

$$0 = \hat{w} - \alpha H\left(\frac{\kappa_c}{\sqrt{q_0}}\right)$$

$$0 = -\hat{q_0} + \hat{w}(q_0 + \kappa_c^2) - \alpha \sqrt{q_0} \kappa_c G\left(\frac{\kappa_c}{\sqrt{q_0}}\right)$$

$$0 = E(1 - \gamma) - \hat{w}(q_0 - \gamma) + \hat{q_0}$$

$$0 = \hat{w} - H\left(-\frac{E - \gamma \hat{w}}{\sqrt{\hat{q_0}}}\right)$$

$$0 = \hat{w}(E - 1) + \sqrt{\hat{q_0}} G\left(\frac{E - \gamma \hat{w}}{\sqrt{\hat{q_0}}}\right) + \gamma \hat{w}(1 - \hat{w})$$
(A20)

where

$$G(x) := \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}.$$
 (A21)

From the numerical solution of the system (A20) we determine $\kappa_c(\alpha, \gamma)$ as shown in Figure 3.

References

- 1. Kondor, I.; Pafka, S.; Nagy, G. Noise sensitivity of portfolio selection under various risk measures. *J. Bank. Financ.* 2007, 31, 1545–1573. [CrossRef]
- 2. Cover, T.M. Geometrical and Statistical Properties of Systems of Linear Inequalities with Applications in Pattern Recognition. *IEEE Trans. Electron. Comput.* **1965**, *EC-14*, 326–334. [CrossRef]
- 3. Berg, J.; Engel, A. Matrix Games, Mixed Strategies, and Statistical Mechanics. Phys. Rev. Lett. 1998, 81, 4999–5002. [CrossRef]
- 4. Landmann, S.; Engel, A. On non-negative solutions to large systems of random linear equations. *Physica* 2020, *A552*, 122544. [CrossRef]
- Garnier-Brun, J.; Benzaquen, M.; Ciliberti, S.; Bouchaud, J.P. A New Spin on Optimal Portfolios and Ecological Equilibria. 2021. Available online: https://arxiv.org/abs/2104.00668 (accessed on 17 June 2021).
- 6. MacArthur, R. Species packing and competitive equilibrium for many species. Theor. Popul. Biol. 1970, 1, 1–11. [CrossRef]
- Tikhonov, M.; Monasson, R. Collective Phase in Resource Competition in a Highly Diverse Ecosystem. *Phys. Rev. Lett.* 2017, 118, 048103. [CrossRef] [PubMed]
- 8. Todd, M. Probabilistic models for linear programming. Math. Oper. Res. 1991, 16, 671–693. [CrossRef]
- 9. Farkas, J. Theorie der einfachen Ungleichungen. J. Reine Angew. Math. (Crelles J.) 1902, 1902, 1–27.
- 10. Gardner, E. The space of interactions in neural network models. J. Phys. A Math. Gen. 1988, 21, 257–270. [CrossRef]
- 11. Engel, A.; Van den Broeck, C. Statistical Mechanics of Learning; Cambridge University Press: Cambridge, UK, 2001.
- 12. Markowitz, H. Portfolio selection. J. Financ. 1952, 7, 77–91.
- 13. Markowitz, H. Portfolio Selection: Efficient Diversification of Investments; J. Wiley and Sons: New York, NY, USA, 1959.
- 14. JP Morgan. Riskmetrics Technical Manual; JP Morgan: New York, NY, USA, 1995.
- 15. JP Morgan and Reuters. Riskmetrics. In Technical Document; JP Morgan: New York, NY, USA, 1996.
- 16. Acerbi, C.; Nordio, C.; Sirtori, C. Expected Shortfall as a Tool for Financial Risk Management. 2001. Available online: https://arxiv.org/abs/cond-mat/0102304 (accessed on 17 June 2021).
- 17. Artzner, P.; Delbaen, F.; Eber, J.M.; Heath, D. Coherent Measures of Risk. Math. Financ. 1999, 9, 203–228. [CrossRef]
- Acerbi, C.; Tasche, D. Expected Shortfall: A Natural Coherent Alternative to Value at Risk. *Econ. Notes* 2002, *31*, 379–388.
 [CrossRef]
- 19. Pflug, G.C. Some remarks on the value-at-risk and the conditional value-at-risk. In *Probabilistic Constrained Optimization*; Uryasev, S., Ed.; Springer: Berlin/Heidelberg, Germany, 2000; pp. 272–281.
- 20. Basel Committee on Banking Supervision. *Minimum Capital Requirements for Market Risk;* Basel Committee on Banking Supervision: Basel, Switzerland, 2016.
- 21. Michaud, R.O. The Markowitz optimization enigma: Is 'optimized' optimal? Financ. Anal. J. 1989, 45, 31–42. [CrossRef]
- 22. Kempf, A.; Memmel, C. Estimating the global minimum variance portfolio. Schmalenbach Bus. Rev. 2006, 58, 332–348. [CrossRef]
- Basak, G.K.; Jagannathan, R.; Ma, T. A jackknife estimator for tracking error variance of optimal portfolios constructed using estimated inputs. *Manag. Sci.* 2009, 55, 990–1002. [CrossRef]
- 24. Frahm, G.; Memmel, C. Dominating estimators for minimum-variance portfolios. J. Econom. 2010, 159, 289–302. [CrossRef]

- 25. Pafka, S.; Kondor, I. Noisy Covariance Matrices and Portfolio Optimization II. Physica 2003, A 319, 487–494. [CrossRef]
- 26. Burda, Z.; Jurkiewicz, J.; Nowak, M.A. Is Econophysics a Solid Science? Acta Phys. Pol. 2003, B 34, 87–132.
- 27. Caccioli, F.; Still, S.; Marsili, M.; Kondor, I. Optimal liquidation strategies regularize portfolio selection. *Eur. J. Financ.* 2013, 19, 554–571. [CrossRef]
- Caccioli, F.; Kondor, I.; Marsili, M.; Still, S. Liquidity Risk And Instabilities In Portfolio Optimization. Int. J. Theor. Appl. Financ. 2016, 19, 1650035. [CrossRef]
- 29. Ledoit, O.; Wolf, M. Nonlinear shrinkage estimation of large-dimensional covariance matrices. *Ann. Stat.* **2012**, *40*, 1024–1060. [CrossRef]
- Kondor, I.; Papp, G.; Caccioli, F. Analytic solution to variance optimization with no short positions. J. Statitical Mech. Theory Exp. 2017, 2017, 123402. [CrossRef]
- Kondor, I.; Papp, G.; Caccioli, F. Analytic approach to variance optimization under an l₁ constraint. *Eur. Phys. J. B* 2019, 92, 8.
 [CrossRef]
- 32. Young, M.R. A minimax portfolio selection rule with linear programming solution. Manag. Sci. 1998, 44, 673–683. [CrossRef]
- 33. Donoho, D.; Tanner, J. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* **2009**, *367*, 4273–4293. [CrossRef]
- 34. Schmidt, B.K.; Mattheiss, T. The probability that a random polytope is bounded. Math. Oper. Res. 1977, 2, 292–296. [CrossRef]
- 35. Kondor, I.; Varga-Haszonits, I. Instability of portfolio optimization under coherent risk measures. *Adv. Complex Syst.* 2010, 13, 425–437. [CrossRef]
- 36. Ciliberti, S.; Kondor, I.; Mézard, M. On the Feasibility of Portfolio Optimization under Expected Shortfall. *Quant. Financ.* 2007, 7, 389–396. [CrossRef]
- 37. Varga-Haszonits, I.; Kondor, I. The instability of downside risk measures. J. Stat. Mech. Theory Exp. 2008, 2008, P12007. [CrossRef]
- 38. Hertz, J.; Krogh, A.; Palmer, R.G. Introduction to the Theory of Neural Computation; Addison-Wesley: Redwood City, CA, USA, 1991.
- 39. Rosenblatt, F. Principles of Neurodynamics: Perceptions and the Theory of Brain Mechanisms; Spartan: Washington, DC, USA, 1962.
- 40. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NJ, USA, 1953.
- 41. May, R. Will a large complex system be stable? *Nature* **1972**, 238, 413–414. [CrossRef]
- 42. Amelunxen, D.; Lotz, M.; McCoy, M.B.; Tropp, J.A. Living on the edge: A geometric theory of phase transitions in convex optimization. *Inform. Inference* **2013**, *3*, 224–294. [CrossRef]
- 43. Prüser, A. Phasenübergänge in Zufälligen Geometrischen Problemen. Master's Thesis, University of Oldenburg, Oldenburg, Germany, 2020.
- 44. Caccioli, F.; Kondor, I.; Papp, G. Portfolio optimization under expected shortfall: Contour maps of estimation error. *Quant. Financ.* **2018**, *18*, 1295–1313. [CrossRef]