



HAL
open science

Quantification des incertitudes en gestion d'actifs : méthodes à noyaux et fluctuations statistiques

Linda Chamakh

► **To cite this version:**

Linda Chamakh. Quantification des incertitudes en gestion d'actifs : méthodes à noyaux et fluctuations statistiques. Statistics [math.ST]. Institut Polytechnique de Paris, 2021. English. NNT : 2021IPPAX045 . tel-03373286

HAL Id: tel-03373286

<https://theses.hal.science/tel-03373286v1>

Submitted on 11 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2021IPPAX045

Thèse de doctorat



Quantifying uncertainty in asset management: Kernel methods and statistical fluctuations

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École Polytechnique

École doctorale n°574 École doctorale de mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques appliquées

Thèse présentée et soutenue à Palaiseau, le 31 août 2021, par

LINDA CHAMAKH

Composition du Jury :

Romuald Elie Professeur, Université Gustave Eiffel (LAMA)	Rapporteur
Lorenzo Rosasco Professeur associé, Université de Gênes et MIT (MaLGa)	Rapporteur
Emmanuel Gobet Professeur, Ecole Polytechnique (CMAP)	Directeur de thèse
Zoltán Szabó Chercheur senior, Ecole Polytechnique (CMAP)	Co-directeur de thèse
Jean-Philippe Lemor Cadre scientifique, BNP Paribas (SSH)	Co-directeur de thèse
Grégoire Loeper Cadre scientifique, BNP Paribas - Université de Monash (QFIS)	Examineur
Agnès Sulem Directeur de recherche, INRIA Paris (MATHRISK)	Examinatrice
Caroline Hillairet Professeur, ENSAE (CREST)	Présidente du jury



THÈSE
présentée pour obtenir
LE GRADE DE DOCTEUR EN SCIENCES DE
L'ÉCOLE POLYTECHNIQUE
Spécialité : Mathématiques
par
Linda Chamakh

Quantifying uncertainty in asset management : Kernel
methods and statistical fluctuations

Soutenue le 31 août 2021 devant un jury composé de :

Romuald Elie (rapporteur)
Lorenzo Rosasco (rapporteur)
Emmanuel Gobet (directeur de thèse)
Zoltán Szabó (co-directeur de thèse)
Jean-Philippe Lemor (co-encadrant de thèse)
Grégoire Loeper (examineur)
Agnès Sulem (examineur)
Caroline Hillairet (examineur).

REMERCIEMENTS

Mes premiers remerciements vont à mes encadrants de thèse, Emmanuel Gobet, Zoltán Szabó et Jean-Philippe Lemor. Je vous remercie pour votre accompagnement constant durant ces trois dernières années. Merci à Jean-Philippe, pour avoir accepté que je réalise cette thèse au sein de son équipe, et pour les réunions de travail régulières qui ont su recadrer mes idées. Je remercie Emmanuel, pour m'avoir toujours accordé des disponibilités et pour m'avoir transmis de nombreux outils de calculs et de rigueur en rédaction, et Zoltán, pour m'avoir transmis sa passion pour les méthodes à noyau.

Je remercie Romuald Élie et Lorenzo Rosasco pour l'intérêt porté à mes travaux en acceptant d'être les rapporteurs de ma thèse. Être relue par mon ancien professeur de master, Romuald Élie, me permet de clore joliment mes études, et je suis très honorée de pouvoir compter sur Lorenzo Rosasco, spécialiste des méthodes d'apprentissage. Je suis également très honorée de la présence de Agnès Sulem, Caroline Hillairet et Grégoire Loeper dans mon jury de thèse.

Je remercie la BNPP, et plus particulièrement l'équipe SSH, pour m'avoir donné l'opportunité de réaliser cette thèse CIFRE. Un merci particulier à Ahmed Bel Hadj Ayed, pour m'avoir encadrée pendant mon stage de fin d'études dans l'équipe, et m'avoir initiée à l'optimisation de portefeuille et à ses problématiques pratiques. Aux stagiaires de la première heure (Thomas, Yang, Gauthier, Valentin, Randa...) jusqu'aux collègues encore présents (Mariem, Arnaud, Samed, Alexandre), j'envoie un salut affectueux. Cette thèse ne se serait pas faite sans les calculs sur bouts de papiers à la cafétéria de la Maison Dorée.

Je remercie également l'École Polytechnique et l'école doctorale. Je salue chaleureusement la team EG, Michael Allouche et Cyril Benezet, et ma sous-équipe au CMAP, en particulier Stefano de Marco, Florian Bourgey, Clément Rey et Wanqing Wang. Je remercie Polytechnique pour m'avoir confié des missions de monitorat : donner les TD d'algèbre aux élèves étrangers fut une des expériences les plus épanouissantes de ma vie.

Plus généralement, je remercie tous mes professeurs de mathématiques et de sciences croisés au cours de ma vie (je n'en serai sous doute pas là aujourd'hui si je n'avais pas fait de concours Kangourou au collège).

Merci à Artur pour m'avoir supportée et soutenue pendant ces deux dernières années. Nos voyages à travers l'Europe auront été des bouffées d'air frais nécessaires à la bonne réalisation de cette thèse. Je suis contente d'être avec quelqu'un d'aussi génial que toi, et j'espère que tu trouveras dans cette thèse la preuve que moi non plus, je ne suis pas si nulle en maths que ça.

Le plus important pour la fin : je remercie mes parents pour leur soutien indéfectible. Merci de m'avoir soutenue, accompagnée et encouragée dans tous mes choix. Vous nous avez placés dans les meilleures conditions pour réussir, Nassim, Inès, et moi. Cette thèse, comme un moindre gage de reconnaissance, vous est dédiée.

ACKNOWLEDGEMENTS

My first thanks go to my thesis supervisors, Emmanuel Gobet, Zoltán Szabó and Jean-Philippe Lemor. I thank you for your constant support during the last three years. Thanks to Jean-Philippe, for having accepted that I carry out this thesis within his team, and for the regular work meetings which allowed me to reframe my ideas. I thank Emmanuel, for having always granted me availability and for having passed on to me many tools of calculation and rigour in writing, and Zolt'an, for having passed on to me his passion for kernel methods.

I thank Romuald Élie and Lorenzo Rosasco for their interest in my work by accepting to be the rapporteurs of my thesis. Being reviewed by my former master's professor, Romuald Élie, allows me to close my studies nicely, and I am very honoured to be able to count on Lorenzo Rosasco, specialist in learning methods. I am also very honoured by the presence of Agnès Sulem, Caroline Hillairet and Grégoire Loeper in my thesis jury.

I would like to thank the BNPP, and more particularly the SSH team, for giving me the opportunity to carry out this CIFRE thesis. A special thanks to Ahmed Bel Hadj Ayed, for having supervised me during my internship in the team, and for having introduced me to portfolio optimization and its practical problems. To the trainees of the first hour (Thomas, Yang, Gauthier, Valentin, Randa...) to the colleagues still present (Mariem, Arnaud, Samed, Alexandre), I send an affectionate greeting. This thesis would not have been possible without the calculations on pieces of paper in the cafeteria of the Maison Dorée.

I also thank the École Polytechnique and the doctoral school. I warmly greet the EG team, Michael Allouche and Cyril Benezet, and my sub-team at CMAP, in particular Stefano de Marco, Florian Bourgey, Clément Rey and Wanqing Wang. I would like to thank Polytechnique for entrusting me with some monitoring missions: giving algebra tutorials to foreign students was one of the most fulfilling experiences of my life.

More generally, I would like to thank all my maths and science teachers I have met in the course of my life (I would probably not be where I am today if I had not done the Kangaroo competition in middle school).

Thanks to Artur for supporting me during these last two years. Our travels through Europe have been a breath of fresh air necessary to the successful completion of this thesis. I am happy to be with someone as great as you, and I hope that you will find in this thesis the proof that I am not that bad in maths either.

The most important thing at the end: I thank my parents for their unwavering support. Thank you for supporting me, accompanying me and encouraging me in all my choices. You have put us in the best conditions to succeed, Nassim, Inès, and me. This thesis, as a small token of appreciation, is dedicated to you.

Préambule

Le traitement des incertitudes est un problème fondamental dans un contexte financier, et plus précisément d'optimisation de portefeuille. Dans ce contexte, les variables étudiées sont souvent dépendantes du temps, avec des queues épaisses. Dans cette thèse, on s'intéresse à des outils permettant de prendre en compte les incertitudes sous ses formes principales : incertitudes statistiques, incertitudes paramétriques et erreur de modèle, tout en gardant en tête qu'on souhaite les appliquer à ce contexte.

La première partie est consacrée à l'établissement d'inégalités de concentration, c'est-à-dire de borne sur la probabilité d'écart à sa moyenne d'un estimateur, dans le cadre de variables à queues épaisses. L'objectif de ces inégalités est de quantifier quelle confiance on peut donner à un estimateur basé sur une taille finie d'observations, et leur établissement peut se faire en respectant le caractère éventuellement non-borné et non-Gaussien des distributions sous-jacentes, via notamment l'utilisation de normes Orlicz. Dans cette thèse, nous considérons le cas d'estimateurs à distributions à queues épaisses, et nous établissons une nouvelle norme Orlicz et de nouvelles inégalités de concentration qui couvrent le cas de la distribution log-normale.

Dans la seconde partie, on traite de l'impact de l'erreur de modèle pour l'estimation de la matrice de covariance sur des rendements boursiers. Pour respecter une observation quasi-systématiquement vérifiée sur les marchés qui est que la covariance entre les rendements dépend du temps, on suppose qu'il existe un processus de covariance instantanée entre les rendements dont la valeur présente dépend de ses valeurs passées. On peut alors construire explicitement la meilleure estimée de la matrice de covariance pour un instant et un horizon d'investissement donnés. En se plaçant dans le contexte du portefeuille minimum variance, on montre qu'avec grande probabilité, la meilleure estimée de la matrice de covariance (qui est aproximable par Monte Carlo dans le modèle GARCH) constitue le choix de matrice de covariance qui donne les meilleures performances.

Dans la troisième partie, on propose de modéliser l'impact des erreurs paramétriques. Lorsqu'un modèle paramétrique est utilisé, les paramètres du modèle ne sont pas toujours parfaitement connus et bien spécifiés : l'erreur paramétrique est par exemple inhérente à une calibration de paramètres sur un jeu fini de données. Nous proposons une approche pour exprimer le ratio de Sharpe et l'allocation de portefeuille en fonction de paramètres d'entrée jugés incertains. Notre approche passe par l'adaptation d'une technique d'approximation stochastique pour le calcul de la décomposition en polynômes du chaos de la quantité d'intérêt.

Enfin, dans la dernière partie de cette thèse, on s'intéresse à l'optimisation de portefeuille avec distribution cible. Cette technique peut être formalisée sans avoir recours à aucune calibration ou hypothèse de modèle sur les rendements. Nous proposons de trouver ces portefeuilles en minimisant des mesures de divergence basées sur les fonctions noyau et la théorie du transport optimal. Ces mesures de divergence peuvent être non-bornées et ont encore peu été étudiées dans le cas fonction noyau non-bornée. On établit des nouvelles garanties de convergence basées sur les inégalités de concentration dans ce cas.

Preamble

The treatment of uncertainties is a fundamental problem in a financial context, and more precisely in portfolio optimisation. In this context, the variables studied are often time dependent, with heavy tails. In this thesis, we are interested in tools allowing to take into account uncertainties in its main forms: statistical uncertainties, parametric uncertainties and model error, keeping in mind that we wish to apply them to this context.

The first part is devoted to the establishment of concentration inequalities, i.e. bounds on the probability of deviation from the mean of an estimator, in the context of variables with heavy tails. The objective of these inequalities is to quantify the confidence that can be given to an estimator based on a finite size of observations, and their establishment can be done by respecting the possibly unbounded and non-Gaussian character of the underlying distributions, via the use of Orlicz norms. In this thesis, we consider the case of estimators with heavy-tailed distributions, and we establish a new Orlicz norm and new concentration inequalities that cover the case of the lognormal distribution.

In the second part, we discuss the impact of the model error for the estimation of the covariance matrix on stock returns. In order to respect a quasi-systematically verified observation on the markets, which is that the covariance between returns depends on time, we assume that there is an instantaneous covariance process between the returns whose present value depends on its past values. One can then explicitly construct the best estimate of the covariance matrix for a given time and investment horizon. In the context of the minimum variance portfolio, we show that with high probability, the best estimate of the covariance matrix (which can be approximated by Monte Carlo method in the GARCH model) is the choice of covariance matrix that gives the best performance.

In the third part, we propose to model the impact of parametric errors. When a parametric model is used, the parameters of the model are not always perfectly known and well specified: the parametric error is for instance inherent to a calibration of parameters on a finite data set. We propose an approach to express the Sharpe ratio and the portfolio allocation as a function of input parameters that are considered uncertain. Our approach involves the adaptation of a stochastic approximation technique for the computation of the polynomial decomposition of the quantity of interest.

Finally, in the last part of this thesis, we focus on portfolio optimization with target distribution. This technique can be formalised without the need for any calibration or model assumptions on returns. We propose to find these portfolios by minimizing divergence measures based on kernels or optimal transport. Since these divergence measures can be unbounded and have not been studied much yet in the unbounded kernel case, we establish new convergence guarantees based on concentration inequalities.

List of publications

Here is a list of articles (accepted or submitted) and working papers that were written during this thesis:

- [CGS20] L. Chamakh, E. Gobet, and Z. Szabó. Orlicz random Fourier features. *Journal of Machine Learning Research (JMLR)*, 21:1–37, 2020.
- [CGL21a] L. Chamakh, E. Gobet, and J.-P. Lemor. Asymptotic analysis of different covariance matrices estimation for minimum variance portfolio. *Submitted*.
- [CGL21b] L. Chamakh, E. Gobet, and W. Liu. Orlicz norms and concentration inequalities for β -heavy tailed random variables *Submitted*.
- L. Chamakh, and Z. Szabó. Kernel and optimal transport portfolio optimization with target distribution. *Working paper*.

Keywords

Heavy tails; concentration inequality; kernel methods; covariance matrix estimation; portfolio theory; uncertainty quantification.

CONTENTS

Introduction	10
0.1 Context of the PhD	10
0.2 Concentration inequalities for heavy-tailed random variables	14
0.3 Covariance matrix estimation for minimum variance portfolio	17
0.4 Uncertainty quantification for portfolio	19
0.5 Kernel and optimal transport portfolio optimization with target distribution	21
Introduction (en français)	24
0.6 Contexte de la thèse	24
0.7 Inégalités de concentration pour variables à queues épaisses	29
0.8 Estimation de matrice de covariance pour portefeuille à variance minimum	31
0.9 Traitement des incertitudes pour le portefeuille	33
0.10 Optimisation de portefeuille avec distribution cible par divergence kernel et transport optimal	35
I Concentration inequalities for heavy-tailed random variables	37
1 α-exponential random Fourier features (RFFs)	39
1.1 Introduction	39
1.2 Problem formulation	44
1.3 Main result	46
1.4 Properties of the Orlicz norm	52
1.5 Proofs	52
1.5.1 Proof of Remark 1(v)	52
1.5.2 Proof that polynomial growth preserves the exponential Orlicz property	53
1.5.3 Proof of Theorem 1.1	54
1.5.4 Proof of $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$ and finite maximal moments	58
1.5.5 Control if $c \geq c_{HJ}$	58
1.5.6 Bounding the driving terms of Theorem 1.1 for RFF	60
1.5.7 Proofs of Corollary 1.3.1 and 1.3.2	62
Appendices	64
1.A Additional proofs	64
1.A.1 $\Psi_\gamma^{(l)}$, the convexification of Ψ_γ	64
1.A.2 Proof of Remark 2(ii)	65

1.A.3	Proof of the properties in Section 1.4 about the Orlicz norm	65
1.B	Special functions	67
1.C	External statements	67
2	β-heavy tailed concentration	68
2.1	Introduction	68
2.2	Motivating examples and main results	70
2.2.1	Orlicz norm properties	70
2.2.2	Motivating examples of heavy-tailed distributions and adapted Orlicz norm	71
2.2.3	Ψ_β^{HT} -Orlicz norm: properties and inequalities	72
2.3	Proofs	74
2.3.1	Proof of Theorem 2.2.2	74
2.3.2	Proof of Theorem 2.2.3	79
2.3.3	Proof of Theorem 2.2.5	80
2.3.4	Proof of Proposition 2.2.1	83
2.4	Conclusion	83
II	Covariance estimation for minimum variance portfolio	84
3	Covariance estimation for minimum variance portfolio	86
3.1	Introduction	86
3.1.1	Literature background	87
3.1.2	Contribution and outline of this chapter	88
3.2	Formulation of the problem	88
3.2.1	Problem setup	89
3.2.2	Model assumptions	90
3.2.3	Main results	91
3.3	Proofs and auxiliary results	92
3.3.1	Proof of Theorem 3.2.2	92
3.3.2	Proof of Proposition 3.3.1	95
3.3.3	Proof of Propositions 3.2.1 and 3.3.2	102
3.3.4	Proof of Lemma 3.3.3	103
3.3.5	Auxiliary result	105
3.4	Model specification	105
3.4.1	Motivation: GARCH-CCC model	105
3.4.2	Stationarity, ergodicity and application of results	106
3.4.3	Additional models	108
3.5	Numerical experiments	109
3.5.1	Realistic GARCH values, fitting procedure	110
3.5.2	Simulation procedure	111
3.5.3	Empirical probability: impact of p_{\max} and d	114
3.5.4	Tail exponent Hill estimation	117
3.6	Conclusion, perspectives	118
Appendices		120
3.A	GARCH properties	120
3.A.1	Ergodicity and stationarity	120
3.A.2	GARCH-CCC ergodicity	120
3.A.3	Irreducibility and auxiliary results on the GARCH-CCC	122
3.A.4	GARCH-CCC density	123
3.B	Hill estimator	124

3.C	Additional properties	125
III	Uncertain quantification for portfolio	128
4	UQSA for portfolio Sharpe ratio and allocation	130
4.1	Introduction	130
4.1.1	UQ techniques	131
4.1.2	UQSA setting	131
4.1.3	Explored financial applications	132
4.2	UQSA algorithm: principle and main assumptions	132
4.2.1	Principle and notations	133
4.2.2	UQSA algorithm	133
4.2.3	Main assumptions and convergence results	134
4.2.4	Applications: convergence metrics	135
4.3	UQSA for the Sharpe ratio	136
4.3.1	Definition and motivation	136
4.3.2	SA scheme for the Sharpe ratio	138
4.3.3	UQSA assumptions verification	139
4.3.4	Applications: log-normal returns and specific weights	141
4.3.5	Experiments	143
4.3.6	Conclusion to UQ for the Sharpe ratio section	153
4.4	UQSA for the portfolio allocation	153
4.4.1	Portfolio optimization as SA problem	153
4.4.2	Allocation model	154
4.4.3	Application	159
4.4.4	Conclusion to UQ for portfolio analysis section	167
	Appendices	168
4.A	Proof of Lemma 4.4.3	168
4.B	UQSA algorithm: a projection approach	169
IV	Portfolio optimization with kernel and optimal transport (K.O.T.) divergence	171
5	K.O.T. portfolio optimization with target distribution	173
5.1	Introduction	173
5.2	Problem formulation and optimization	175
5.3	Results	180
5.3.1	Analytical formulas for mean embedding	180
5.3.2	Concentration of semi-explicit MMD	183
5.4	Experiments on simulated data	186
5.4.1	Advantage of semi-explicit MMD	187
5.4.2	Misspecified setting	191
5.5	Experiments on financial time series	192
5.5.1	Datasets, rebalancing, performance measures, target distributions	192
5.5.2	Numerical studies	196
5.6	Conclusion	202

Appendices	204
5.A Proofs	204
5.A.1 Proofs of auxiliary statements	204
5.A.2 Proof of concentration results (Theorem 5.2, Theorem 5.3, Theorem 5.4)	208
5.B Implementation tools	215
5.B.1 Truncated evaluation of E_i^λ	216
5.B.2 Score functions and kernel derivatives	218
5.B.3 Moments of various distributions	219
5.C External statements	219

INTRODUCTION

0.1 Context of the PhD

This thesis was carried out within the framework of a Cifre contract in the Global Markets division of BNP Paribas in the Systematic Strategies and Hybrids (SSH) research team. On the academic side, the thesis was supervised at the Centre of Applied Mathematics of the École Polytechnique by Professors Emmanuel Gobet and Zoltán Szabó. Jean-Philippe Lemor, Senior Quant and head of the SSH team, supervised the thesis for BNPP.

One of the objectives of the SSH team is to ensure the robustness of the systematic strategies proposed by the BNPP structuring team. A systematic strategy can be defined as an investment strategy where the allocation rules are contractually fixed at the time of the strategy's implementation and executed automatically according to the values taken by a set of state variables representing the state of the financial market. For example, the mean-variance portfolio strategy, initiated by Markowitz [Mar52], consists of maximising the expectation of the portfolio's gain under the constraint of a target volatility: in this case, the state variables are the variance-covariance matrix of the d underlyings making up the investment universe, and their expected returns. By denoting C the covariance matrix on the returns of the assets considered (where we call return the relative variation of the asset price), $\boldsymbol{\mu}$ their mean, \mathcal{W} the set of constraints on the allocation vector and $\sigma_T > 0$ the target volatility of the portfolio, the mean-variance portfolio is the solution of the following problem:

$$\begin{aligned} \max_{\mathbf{w} \in \mathcal{W}} \mathbf{w}^\top \boldsymbol{\mu} \\ \text{s.t. } \mathbf{w}^\top C \mathbf{w} = \sigma_T^2. \end{aligned} \tag{0.1.1}$$

In practice, this portfolio will be reallocated at a predetermined frequency (for example, every day or every month) by updating the value of these state variables (C and $\boldsymbol{\mu}$ in the Markowitz case). Typically, C and $\boldsymbol{\mu}$ are estimated empirically, on a fixed number of past returns (called the *window size* of the estimate), and at each new decision on the portfolio, new estimates of C and $\boldsymbol{\mu}$ are calculated (either on the same window size as before - sliding window approach -, or on all the available historical data).

One of the most popular indicators to evaluate the performance of a strategy is the Sharpe ratio. By definition, this ratio is equal to the expectation of the returns of the portfolio in excess of the risk-free rate, renormalized by their volatility.

The question then arises as to whether a good value of the Sharpe ratio¹ on the backtest part has much chance of lasting over time, and whether the Sharpe ratio is robust to small perturbations of its initial parameters, i.e. the parameters of the strategy (window size, recalibration

¹Typically, a good Sharpe ratio is greater than 1.

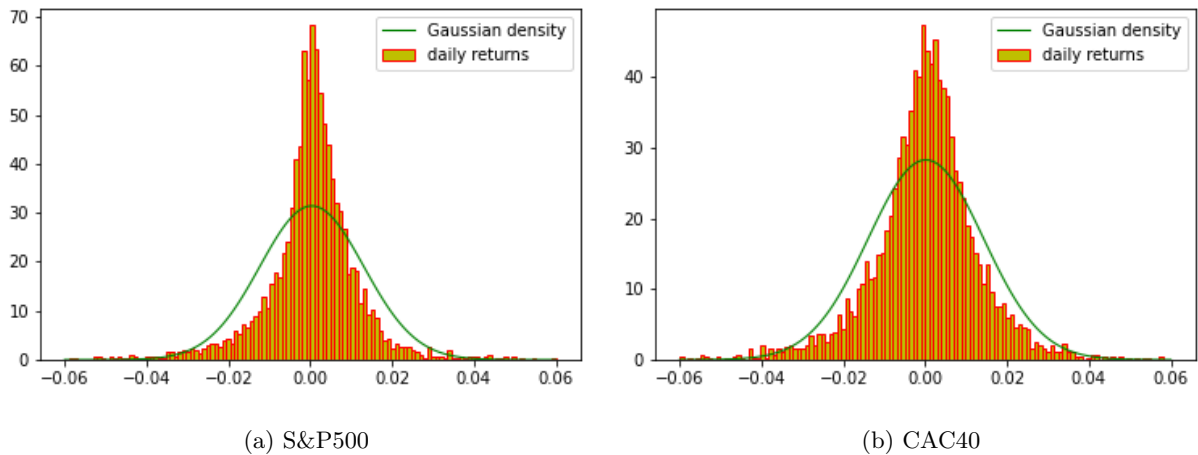


Figure 1 – Plot of daily returns of the S&500 and the CAC40 and the Gaussian distribution of the same mean and variance.

frequency, etc.), and the historical data, or backtest, of the strategy. To do this, studies called “robustness tests” are conducted by the SSH team.

It is in this industrial context that this thesis was written. The initial objective of this thesis was to adapt and apply concentration inequalities to portfolio value variables. Concentration inequalities allow to have confidence intervals on the difference between an empirical estimator, measured on a finite number of observations, and its true value. These confidence intervals are established by respecting the finite nature of the number of observations and depend on the law of the estimator for a given number of observations. On the other hand, one objective was to adapt uncertainty quantification techniques for the Sharpe ratio of a strategy, i.e. to find a way to express the Sharpe ratio as a function of its uncertain parameters.

We first looked at the statistical properties of the financial data to define the assumptions and framework under which our research should be conducted.

Distribution of returns and time dependence. The theory of the efficient market initiated at the beginning of the 20th century [Bac00, Fam71], according to which stock market variations are only white noise, has been widely questioned over the decades. This theory is based on the idea that in a market of free competition where all the players have the same information, prices automatically balance themselves independently of the past. In continuous time, this corresponds to modelling price variations as following a Brownian motion:

$$S_{t+1} - S_t = \sigma(W_t - W_{t-1}), \quad W_t \text{ Brownian motion.}$$

Osborne in 1959 proposed to model the log of price ratios as following a Gaussian distribution [Os59], a theory incorporated in the famous Black and Scholes model in 1973 [BS73].

However, and from the inception of these theories, it was clear that the Gaussian assumption of temporal independence, whether for daily price variations or log returns, was not appropriate. Indeed, as illustrated in Figures 1 and 2, the distribution of stock returns has heavier tails than the Gaussian distribution. The histograms of returns (Figure 1) show distributions with a high peak and tails slightly above the Gaussian distribution. The quantile-quantile plots (Figure 2) clearly show that the extreme quantiles of the stock returns do not coincide with those of the Gaussian distribution: the relationship should have been purely linear otherwise (blue points aligned with red line), but here the empirical quantiles deviate, indicating that the returns take more extreme values than those that would have been taken by Gaussian returns at an




Figure 2 – Quantile-quantile plots: quantile of the daily returns of the S&500 and the CAC40 as a function of the quantiles of the Gaussian distribution with reduced centre.

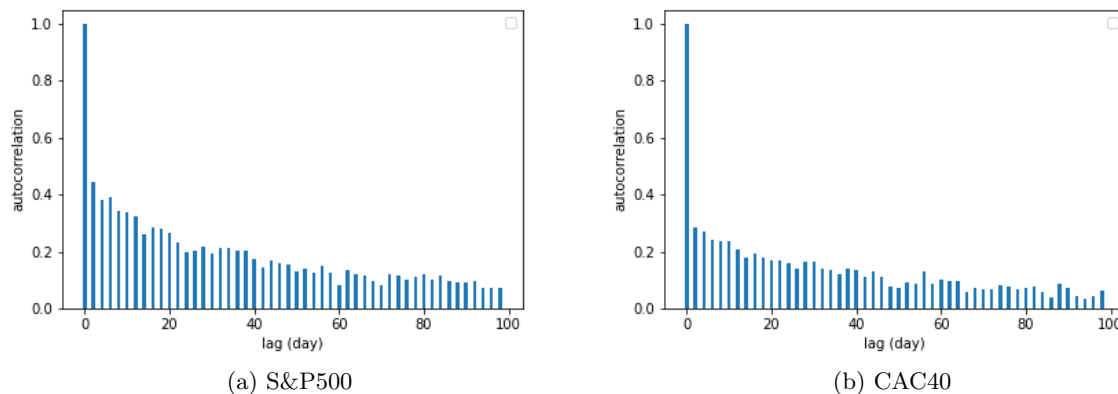


Figure 3 – Autocorrelation plot on daily returns of the S&P500 and the CAC40 in absolute value.

equivalent level of probability. Moreover, the assumption of independence over time of the volatility of returns is called into question by the high values taken by the autocorrelation of absolute returns Figure 3. Absolute returns can be considered as a proxy for the volatility of returns. Autocorrelation is defined as the correlation between absolute returns at time t and their future value at time $t + lag$: autocorrelation is here non negligible, compared to the autocorrelation of an independent white noise (see [BD16, Figure 1.13] for the autocorrelation of a white noise for example).

Mandelbrot was the first to speak about stylised facts, i.e. statistical properties verified quasi-systematically on stock market returns, in [Man63]. More recently, Cont [Con01] lists these properties. Among these stylized facts are the non-Gaussianity of daily returns and their leptokurtic nature, i.e. their tails are heavier than the Gaussian distribution, the time dependence of the variance of returns, but also the negative skewness of returns, and the volatility cluster effects.

Outline of the thesis. In what follows we detail the path followed by the thesis, with the leitmotiv of taking into account uncertainties in the context of heavy-tailed distribution for financial portfolios.

Our first objective was to focus on concentration inequalities. These inequalities allow us to




Figure 4 – Plot of volatility measured by sliding window on the daily returns of the S&P500 and the CAC40.




Figure 5 – Higher moment p_{\max} on the variance of returns calibrated with the GARCH model.

quantify the deviation between an estimator and its expectation, depending on the number of observations on which the estimator depends. The presence of heavy tails in financial returns motivated the work on new concentration inequalities adapted to this context. In the first part of this thesis, we establish concentration inequalities adapted to α -exponential (Chapter 1), or lognormal (Chapter 2) variables, which are two examples of heavy-tailed distributions.

In the second step, we were interested in the impact of taking into account the time dependence of the variance in financial returns. The initial question was whether it was better to consider an adaptive model for the estimation of the covariance matrix C in problems such as (0.6.1), or whether the simple historical estimation of C was efficient. As illustrated on the S&500 and the CAC40 (Figure 4), the variance on returns moves over time, which makes it intuitively preferable to use an adaptive model, such as the GARCH models, rather than a historical covariance on a sample size that is too large to account for short-term movements in volatility.

Moreover, the GARCH model implies a heavy-tailed distribution for the returns, and impose the existence of at most polynomial moments. Figure 5 shows the estimate of the number of finite




Figure 6 – Quantile quantile plot of correlation calibrated by sliding window between two assets.

moments from the GARCH(1,1) model for 12 financial instruments: the highest finite moment on these instruments does not exceed 6. The detailed study of these models for the performance of the minimum variance portfolio constitutes the second part (Chapter 3) of this thesis.

In the third step (Chapter 4), we sought to apply uncertainty quantification techniques to the financial portfolio setting. In uncertainty quantification, one assumes that some parameters are uncertain and seeks to “easily” express the solution of a problem in terms of these uncertain parameters, i.e. without having to solve the initial problem for each value of the uncertain parameters. Typically, this can be done by evaluating the Polynomial Chaos Expansion (PCE) of the solution of the problem, i.e. by computing a projection of the function on a polynomial basis of the uncertain parameter. We have focused on the application of these techniques to the Sharpe ratio and to portfolio allocations.

In our illustrations, we focused on the impact of uncertainty on the correlation between assets. In practice, by estimating the correlation using moving averages, over a time interval large enough to capture the level of “stationary” correlation between two assets, the distribution of calibrated correlations is relatively wide and non-uniform, as shown in Figure 6 (quantile quantile plot against the uniform distribution of the calibrated correlation on 2 financial instruments using moving averages). It was therefore interesting to quantify its impact.

Finally, the last part of this thesis (Chapter 5) is devoted to portfolio optimisation with target distribution, via kernel-based and optimal transport theory based divergence measures. From a financial point of view, the goal of this approach is to build portfolios with thinner distribution tails and smaller skewness than with traditional portfolio optimization approaches.

The parts of this thesis can be read independently. To guide the reader, and to help him in his reading choices, we give in the following a summary of each part of the thesis.

0.2 Concentration inequalities for variables with heavy-tails: an Orlicz norm and Talagrand inequality approach

Context. With the rise in the use of statistical estimators, the study of their behaviour and their deviation from the mean has become essential. Concentration inequalities allow us to quantify this deviation as a function of the size of the sample used by the estimator. Given $M \in \mathbb{N}^*$, Y_1, \dots, Y_M a sequence of random variables, generally independent, and \mathcal{F} a class of measurable functions, the concentration inequalities are in the form

$$\forall \varepsilon > 0, \quad \mathbb{P} \left(\left| \frac{1}{M} \sum_{m=1}^M Y_m - \mathbb{E}Y \right| \geq \varepsilon \right) \leq b(\varepsilon, M), \quad (\text{concentration of the sum})$$

$$\forall \varepsilon > 0, \quad \mathbb{P} \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m=1}^M f(Y_m) - \mathbb{E}f(Y) \right| \geq \varepsilon \right) \leq B(\varepsilon, M), \quad (\text{uniform concentration})$$

where b and B decrease towards 0 when M is large. In practice, b and B may depend on the distribution of $\{Y_m\}_{m \in [M]}$ and the complexity of the \mathcal{F} class. These inequalities can be interpreted as the extent to which at a given M the estimator is concentrated around its mean, converges quickly to it with M , in the case of worst-case deviation on the \mathcal{F} class for B .

Statistical learning theory has contributed to the popularity of these inequalities (see [Vap00, Chapter 3] and [Bou02]). This theory generally deals with concentration on the (bounded) error between the prediction of an algorithm and its target value independently of the distribution of the observations. This is not our framework, since in the financial context, and in many applications (cf. Chapter 1: application to the concentration of kernel derivatives approximation via random Fourier features), f is not bounded, and admits a heavy-tailed distribution. The results dealing with the case of unbounded distributions focus on the Gaussian case or under strong assumptions on the moment of the type Bernstein condition [BLM13, Section 2.3 and 2.8] not verified for a heavy-tailed distribution. We therefore seek to fill this gap.

By heavy-tailed distributions we mean distributions with tails heavier than the Gaussian or exponential distribution but less heavy than the fat-tailed Pareto distributions. We can separate these distributions into two classes:

1. the α -exponential-tailed distributions: $\mathbb{E} [e^{s|Y|^\alpha}] < \infty$ for an $\alpha \in (0, 1)$ and $s > 0$,
2. the distributions not admitting an α -exponential moment, but all the polynomial moments (typically the lognormal variables).

Our objective is therefore to establish concentration inequalities in these two cases.

Orlicz-Talagrand: motivation and sketch of proof. Given $\Psi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ continuous, increasing, such that $\Psi(0) = 0$ (called ‘‘Orlicz function’’), we define the Ψ -Orlicz norm of a random variable X taking its values in a Banach space B as:

$$\|X\|_\Psi := \inf \left\{ c > 0 : \mathbb{E} \left[\Psi \left(\frac{\|X\|_B}{c} \right) \right] \leq 1 \right\}. \quad (0.2.1)$$

Orlicz norms are useful for establishing concentration inequalities because:

1. It is possible to construct a custom Orlicz function with respect to the distribution.
2. The proof scheme for establishing concentration inequalities via Orlicz norms is (relatively) independent of the chosen Orlicz norm (with inequality being all the more robust as the Orlicz function is adapted to the distribution).

1. Choice of the function Ψ

To get tight bounds, we look for the Orlicz function with the strongest possible growth verifying $\mathbb{E} [\Psi(sX)] < \infty$ for at least one $s > 0$. Examples:

- in the α -exponential case, the limit Orlicz function, studied by [Tal89][Ada08], is

$$\Psi_\alpha(x) = e^{x^\alpha} - 1,$$

- To deal with the intermediate case, we introduce the new Orlicz function:

$$\Psi_\beta^{\text{HT}}(x) := \exp \left((\ln(x+1))^\beta \right) - 1$$

for a parameter $\beta > 1$ ($\exp(x)$ and e^x are used interchangeably throughout the manuscript). The log-normal case is covered with $\beta < 2$.

2. General proof scheme

Here for simplicity $X \in \mathbb{R}$. We can establish a robust concentration inequality on X such that $\|X\|_\Psi < \infty$ by Markov inequality by exploiting the definition of the Orlicz norm. For example, in the α -exponential case, by composing by the increasing function $\Psi_\alpha(\cdot/\|X\|_{\Psi_\alpha}) + 1$, then applying the Markov inequality, we have, for $\varepsilon > 0$

$$\begin{aligned} \mathbb{P}(|X| > \varepsilon) &= \mathbb{P}\left(\Psi_\alpha\left(\frac{|X|}{\|X\|_{\Psi_\alpha}}\right) + 1 > \Psi_\alpha\left(\frac{\varepsilon}{\|X\|_{\Psi_\alpha}}\right) + 1\right) \\ &\stackrel{\text{Markov}}{\leq} \underbrace{\left(\mathbb{E}\left[\Psi_\alpha\left(\frac{|X|}{\|X\|_{\Psi_\alpha}}\right)\right] + 1\right)}_{\leq 1} / \underbrace{\left(\Psi_\alpha\left(\frac{\varepsilon}{\|X\|_{\Psi_\alpha}}\right) + 1\right)}_{=e^{\frac{\varepsilon^\alpha}{\|X\|_{\Psi_\alpha}^\alpha}}} \leq 2 \exp\left(-\frac{\varepsilon^\alpha}{\|X\|_{\Psi_\alpha}^\alpha}\right). \end{aligned}$$

This proof scheme is used in [Ada08, vdGL13]. In practice, $X = \frac{1}{M} \sum_{m=1}^M Y_m$ or $X = \sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{m=1}^M f(Y_m)$. To guarantee $b(\varepsilon, M) = o(M)$ and $B(\varepsilon, M) = o(M)$, two inequalities will be fundamental: the Talagrand inequality [Tal89, Theorem 3: $\Psi = \Psi_\alpha$] and the maximal inequality [vW96, Lemma 2.2.2: Ψ convex+ mild growth condition]

$$\left\| \sum_{m \in [M]} Y_m \right\|_\Psi \leq C_\Psi \left(\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \right), \quad (0.2.2)$$

$$\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \leq K \Psi^{-1}(M) \max_{m \in [M]} \|Y_m\|_\Psi. \quad (0.2.3)$$

When these inequalities are verified, we easily have $b(\varepsilon, M) = o(M)$.

In the uniform case, one might be tempted to apply Talagrand's inequality directly. But we would then end up with a term in $\left\| \sum_{m \in [M]} \sup_{f \in \mathcal{F}} f(Y_m) \right\|_{L_1(B)}$ which is difficult to control.

One trick is to use a truncation technique (used by Adamczak [Ada08, Theorem 4] to establish a large deviation inequality which we have adapted and corrected to establish concentration inequalities). Let us introduce for a truncation level $c > 0$ the truncated function $\mathcal{T}_c f := -c \vee f \wedge c$ and the remainder by the truncated function $\mathcal{R}_c f := f - \mathcal{T}_c f$. We can then treat separately the concentration of the truncated part (bounded function: classical problem) and the remainder. An application of Talagrand's inequality gives the term: $\left\| \sum_{m \in [M]} \sup_{f \in \mathcal{F}} |\mathcal{R}_c f(X_m)| \right\|_{L_1(B)}$. Intuitively, by choosing c large enough, this "remainder" term becomes small. Thanks to the Hoffman-Jorgensen inequality, we can calibrate c in such a way as to guarantee that this term is negligible in front of M .

Problems encountered and contributions: In addition to setting up a fairly generic proof scheme, we have contributed to overcome the following problems:

- In the α -exponential case, the Talagrand inequality (0.2.2) is already available in the literature [Tal89, Theorem 3], but not the maximal inequality. The general maximal inequality recalled in (2.2.3) is valid for Orlicz function convex and satisfying additional mild conditions. We contributed in establishing the maximal inequality for the non-convex Ψ_α function.
- In the β -heavy case: $\Psi_\beta^{\text{HT}}(x) := e^{\ln^\beta(x+1)} - 1$, since Ψ_β^{HT} is convex, the maximal inequality is verified, but Talagrand's inequality is not available in the literature. We contributed in establishing the Talagrand inequality for the non-convex Ψ_β^{HT} .

We answer these problems in our work, but will not detail in the introduction the solutions brought.

0.3 Asymptotic analysis of the covariance matrix estimation for the minimum variance portfolio

In this section, we deal with a much more practical problem, which consists in analysing the impact of taking into account the best estimate of the covariance matrix in the minimum variance portfolio problem.

The theory of the mean-variance portfolio introduced by Markowitz consists in finding the portfolio allocation \mathbf{w} which minimises $\mathbf{w}^\top C \mathbf{w} - \gamma \mathbf{w}^\top \boldsymbol{\mu}$, where C corresponds to the covariance matrix on the returns of the assets considered, $\boldsymbol{\mu}$ their mean, and $\gamma > 0$ the investor's risk appetite. In practice, C and $\boldsymbol{\mu}$ must be estimated, and by the nature of the problem (explicit solution depending on the inverse of C) lead to a significant instability of the estimated portfolios. Many articles deal with the problem of estimating $\boldsymbol{\mu}$ [BG91, GHZ13]; we focus here on that of C with $\gamma = 0$.

We place ourselves in the following framework: we consider a universe of d assets, whose returns follow the conditional law $\mathbf{r}_{t|\mathcal{F}_{t-1}} \sim \mathcal{N}(0, V_t)$ with V_t instantaneous covariance matrix, of maximum polynomial moment p_{\max} . As explained above, a dynamic model of covariance is consistent with the stylised facts observed on the market. We consider an investment horizon H ; in this case, the *best estimate of the realized covariance matrix* is

$$cRC_{H,t} := \mathbb{E} \left[\frac{1}{H} \sum_{k=1}^H \mathbf{r}_{t+k} \mathbf{r}_{t+k}^\top | \mathcal{F}_t \right] = \frac{1}{H} \sum_{k=1}^H \mathbb{E}[V_{t+k} | \mathcal{F}_t]. \quad (0.3.1)$$

Generally, the practitioner uses a historical covariance matrix $\frac{1}{T} \sum_{k=1}^T \mathbf{r}_{t-k} \mathbf{r}_{t-k}^\top$ based on past returns. When the size of the history T tends towards infinity, this estimator tends towards the stationary covariance matrix of the process. However, when the investment horizon is short and the instantaneous covariance of the process takes different values from the stationary covariance matrix (moments of high or low short-term volatility), it is better to use $cRC_{H,t}$ to take into account the latency effects of the variance.

We consider the problem of the minimum variance portfolio: given a covariance matrix estimate C , it consists in solving

$$\text{mv}(C) := \arg \min_{\mathbf{w} \in \mathcal{W}} \mathbf{w}^\top C \mathbf{w}, \quad (0.3.2)$$

where $\mathcal{W} \subset \mathbb{R}^d$ takes into account the constraints on \mathbf{w} (typically, $\left\{ \mathbf{w} \in \mathbb{R}^d : \sum_{i=1}^d w_i = 1 \right\} \subset \mathcal{W}$). In this case, the performance metric is the *future realised variance* and the investor seeks to make it as small as possible. By denoting the matrix of *realized covariance*

$$RC_{H,t_n} = \frac{1}{H} \sum_{k=1}^H \mathbf{r}_{t_n+k} \mathbf{r}_{t_n+k}^\top, \quad (0.3.3)$$

the realised variances associated to the minimum variance portfolios (0.3.3), allocated on one hand with the conditional realized covariance (0.3.1), and on the other hand with the fixed benchmark covariance V_{ref} are written as

$$\begin{cases} R_{N,H} & := \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top RC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\ R_{N,H}^{\text{ref}} & := \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top RC_{H,t_n} \text{mv}(V_{\text{ref}}). \end{cases}$$

Our objective is to show that with high probability, $R_{N,H} < R_{N,H}^{\text{ref}}$, all the more surely as N grows. Intuitively/asymptotically, we have $R_{N,H}^{\text{ref}} - R_{N,H} \geq 0$.




Figure 7 – $\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^\infty)$, for $H = 1$ (red), $H = 5$ (blue) and $H = 21$ (green) as a function of the number of reshoots N , in log-log scale, for GARCH parameters such as $d = 50$ and $p_{\max} = 5.05$

- Let's introduce the conditional variance terms:

$$\begin{cases} cRV_{N,H} & := \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top cRC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\ cRV_{N,H}^{\text{ref}} & := \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top cRC_{H,t_n} \text{mv}(V_{\text{ref}}). \end{cases} \quad (0.3.4)$$

By definition of the minimum variance portfolio, we always have $cRV_{N,H} \leq cRV_{N,H}^{\text{ref}}$.

- Moreover, as $\mathbb{E}[RC_{H,t_n} | \mathcal{F}_{t_n}] = cRC_{H,t_n}$, the difference $R_{N,H} - cRV_{N,H}$ and $R_{N,H}^{\text{ref}} - cRV_{N,H}^{\text{ref}}$ are centered martingales. Therefore:

$$\frac{1}{N}(R_{N,H}^{\text{ref}} - R_{N,H}) = \underbrace{\frac{1}{N}(R_{N,H}^{\text{ref}} - cRV_{N,H}^{\text{ref}}) - \frac{1}{N}(R_{N,H} - cRV_{N,H})}_{\frac{1}{N} \times \text{centered martingales} \approx 0} + \underbrace{\frac{1}{N}(cRV_{N,H}^{\text{ref}} - cRV_{N,H})}_{\geq 0}.$$

Question: Can we prove via a concentration inequality that the probability of $R_{N,H} < R_{N,H}^{\text{ref}}$ tends rapidly to 0 with the number of rebalancing? As we place ourselves in the framework where V_t thus r_t has few moments (polynomial only, up to the order p_{\max}), it is difficult to obtain bounds decreasing very quickly.

Based on control of moments of order p of sum of martingale increments (Burkholder's inequality [HH80, Theorem 2.10]) and of sum of ergodic process functions (Fort-Moulines inequality [FM03, Proposition 2]), under small technical assumptions on V_t , we establish the following result: for any $H \in \mathbb{N}^*$ investment horizon, by denoting $\ell_H := \mathbb{E}\left[\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N}\right]$ the *performance gap*, there exists a constant $C > 0$ depending on H , p_{\max} , d and V_t such that, for any $N \in \mathbb{N}^*$, number of portfolio reallocations, we have:

$$\mathbb{P}\left(R_{N,H} > R_{N,H}^{\text{ref}}\right) \leq \frac{C}{\ell_H^{p_{\max}}} \frac{H^{\frac{p_{\max}}{2}}}{N^{\frac{p_{\max}}{2}}}.$$

Application with a GARCH-CCC model. Assuming a parametric model on V_t of the GARCH-Constant Conditional Correlation type and assuming the parameters of the model are known, we can estimate the optimal covariance matrix by the Monte Carlo technique. Numerically, the evolution of the probability of success tends towards 0 with the number of portfolio reallocations N (Figure 7).

Limit investment horizon. When the investment horizon H tends towards infinity, the best estimate matrix coincides with the stationary covariance matrix. In practice, we observe that when the investment horizon is of the order of magnitude of the multiple of the half-life of the process, the performance of the portfolio with estimated covariance degrades until a horizon where the choice between the covariance matrices produces equitable performances: the investor is then indifferent between the stationary or the best estimated matrix.

A problem not addressed here is the consideration of parametric uncertainties in the model. In the following section, we give a method for dealing with parametric uncertainties.

0.4 Treatment of parametric uncertainties for portfolio optimization: a chaos decomposition approach via UQSA algorithm

Understanding how parametric uncertainties perturb the outcome of an optimization problem is a fundamental problem. The field of uncertainty quantification is dealing with the set of techniques that allow to quantify this impact.

Let us suppose a parametric model on the inputs of the problem, with an uncertain parameter θ for which we suppose that we can model the uncertainty via a π law: the inputs V_θ follow a law depending on an uncertain parameter $\theta \sim \pi$. We are looking for a representation of the solution of the optimisation problem as a function of θ . Assuming that the square function integrable according to π , for any orthonormal basis of polynomials $\{B_i\}_{i \in \mathbb{N}}$, we can express the solution of the optimisation problem $z^*(\theta)$ via its projection onto this basis of polynomials:

$$z^*(\theta) = \sum_{i \geq 0} u_i B_i(\theta), \quad \forall \theta \in \Theta \subset \mathbb{R}^p, \quad \text{and} \quad \int_{\Theta} B_i(\theta) B_j(\theta) \pi(d\theta) = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{else.} \end{cases} \quad (0.4.1)$$

This is called the chaos polynomial decomposition, and it is one of the most widely used methods in uncertainty quantification [LMK10]. In practice, a truncated form of (0.4.1) is evaluated, as the truncation error generally tends quickly towards 0 (generally, $\sum_{i > m} |u_i|^2 = \mathcal{O}(m^{-\delta})$ with δ order of differentiability of the z^* function).

The UQSA algorithm [CFGS20] allows a progressive estimation of the coefficients $\{u_i\}_{i \leq m_k}$, m_k increasing, when the problem is in the form:

$$z^*(\theta) \text{ solution of } \mathbb{E}[H(z^*(\theta), V_\theta, \theta)] = 0, \quad \pi - \text{a.e.}, \quad (0.4.2)$$

In the financial context, we are interested in two cases of application of these techniques: the uncertainty management for the Sharpe ratio and for portfolio allocation when the financial growth rates follow an uncertain parametric law.

Sharpe ratio. The Sharpe ratio gives an account of the profitability of a portfolio renormalised by its risk. By denoting R the return in excess of the risk-free rate of the portfolio, it is written as

$$\text{SR} = \frac{\mathbb{E}[R]}{\sqrt{\text{Var}[R]}}.$$

One difficulty in applying the UQSA approach was to find a way to express the Sharpe ratio as a solution to a problem of the form (4.1.1). This difficulty was overcome by the following trick. By taking two independent copies of same law R and \tilde{R} , we can express the square of the Sharpe ratio as the zero of a function:

$$\text{SR}^2 \text{Var}[R] - \mathbb{E}[R]^2 = \mathbb{E} \left[\text{SR}^2 \frac{(R - \tilde{R})^2}{2} - R\tilde{R} \right] = 0.$$




Figure 8 – Evolution of the total error on the Sharpe ratio approximation in the case $m_k = \lfloor k + 1 \rfloor^b$ increasing (blue line) versus $m_k = m \in \{10, 30\}$ (dotted lines) as a function of the number of iterations k .

For our applications, we are interested in the evolution of the Sharpe ratio ratio when the correlation between assets is uncertain (with a parametric distribution modelled on the empirical distribution represented via its quantile plot in Figure 6). We illustrate the efficiency of using UQSA with an increasing number of estimated coefficients according to the iterations in Figure 8: the total error is lower with $m_k = \lfloor k + 1 \rfloor^{0.3}$ increasing rather than at m_k fixed.

Portfolio allocation. We are interested in the quantification of the uncertainties for the portfolio allocation given by

$$\arg \min_{\mathbf{w} \in \mathcal{W}} U(\mathbf{w})$$


where U is convex. Without constraints ($\mathcal{W} = \mathbb{R}^d$), the problem can generally easily be put into the form (4.1.1), by taking the gradient of U and express it as an expectation. We can then apply the UQSA algorithm to this problem. We would like to deal with the case where \mathcal{W} corresponds to the set of portfolios with positive components (purchase only) and summing to 1:

$$\mathcal{W}^{\geq 0} = \{\mathbf{w} \in \mathbb{R}^d : w_i \geq 0, \sum_{i=1}^d w_i = 1\}.$$

With constraints, we run into two difficulties:

1. The solution of the minimisation problem is not directly the zero of the gradient of U (one must introduce Lagrange multipliers and look at the KKT conditions of the problem),
2. Naively solving the problem via a stochastic approximation algorithm does not allow to stay in $\mathcal{W}^{\geq 0}$.

We propose to apply a change of variable. By introducing $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ such that $\ell(\mathbf{y}) = \mathbf{w}$, we then propose to solve the problem by $\mathbf{y} \in \mathbb{R}^{d-1}$. Numerous techniques known as compositional data analysis allow to pass from \mathbb{R}^{d-1} into $\mathcal{W}^{> 0} = \{\mathbf{w} \in \mathbb{R}^d : w_i > 0, \sum_{i=1}^d w_i = 1\}$. One of them is the log-ratio transformation, which consists in taking $w_i = \frac{e^{y_i}}{1 + \sum_{j=1}^{d-1} e^{y_j}}$, for $i = 1, \dots, d-1$, and $w_d = \frac{1}{1 + \sum_{j=1}^{d-1} e^{y_j}}$. In this thesis, we have devised a new transformation to cover the case $w_i = 0$ which we called the integral transformation and which is based on the following idea: given $f : \mathbb{R} \rightarrow \mathbb{R}_+$ such that $\int_{\mathbb{R}} f(t) dt = 1$, then we can interpret the w_i as portions of the integral of f . By introducing the positive function $p : \mathbb{R} \rightarrow \mathbb{R}_+$ such that $p(0) = 0$, there therefore exists $\mathbf{y} \in \mathbb{R}^{d-1}$ such that:



We study the cases where f corresponds to the derivative of the sigmoid function $x \mapsto \frac{1}{1+e^x}$, and a case with compact support, with f taken as the distribution function of a beta(2, 2) law.

One bottleneck of this approach is that the problem is well posed only in the *unsaturated* case, i.e. when \mathbf{w} takes its values in the interior of the simplex $\mathcal{W}^{>0}$. When \mathbf{w} admits zero component, the SA formulation can admit multiple solutions, suggesting that this approach should be improved to tackle these limit cases.

0.5 Portfolio optimization with target distribution: an approach via divergence measure based on kernel and optimal transport theories

Paradigm shift. So far, we have focused our applications on the mean-variance paradigm, and we have proposed ways of quantifying or managing the uncertainties, either by a better estimation of the first moments (best estimate approach of the covariance matrix), or by the quantification of uncertainty (to see the distribution of the portfolio or the performance knowing the parametric uncertainties). This is indeed the most widely used paradigm in portfolio construction, although it is known to create unstable portfolios that are very sensitive to parameter estimation errors.

In practice, however, the asset managers are increasingly moving away from mean-variance utility approaches towards objective functions or portfolio optimisation methods focused on portfolio risk management, for example by including fourth-order moments in the objective function to reduce the excess kurtosis of the portfolio [MZ10] or by using portfolio optimisation approaches with risk management by asset class [MRT10a, MRT10b].

The minimum divergence approach. It is in this perspective of better taking into account the high order moments of the portfolio that the approach of the “minimum divergence” portfolio with target distribution was introduced [CW12, Las19, GLLN20]: given a measure of divergence D (or semi-metric between distributions), the minimum divergence portfolio is given by

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} D(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T), \quad (0.5.1)$$

where $P = \mathbf{w}^\top R \sim \mathbb{P}_{\mathbf{w}}$ denotes the portfolio distribution, and \mathbb{P}_T the investor’s target distribution. In practice, we assume that we have access to past returns $\{\mathbf{r}_n\}_{n \in [N]}$, and we can estimate the objective function (5.1.1) empirically, or parametrically. Contrary to the authors quoted, we

take this approach with a will not to impose a parametric model on the returns, and we would like as much as possible to be able to explain the objective function according to \mathbb{P}_T :

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} D_{\mathbb{P}_T} \left(\{\mathbf{w}^\top \mathbf{r}_n\}_{n \in [N]} \right).$$

Questions. In this context, several questions arise: which divergence measure (i.e. distance between the portfolio distribution and the target distribution) to choose, which target distribution to consider, and which optimization method would be adequate for this generally non-convex problem?

Choice of distribution and divergence. We focused on a choice of parametric distributions of the generalized Gaussian type (distribution on which one can specify the level of kurtosis, or moment of order 4). This allows us to investigate the tail heaviness of the target distribution and to aim for a lower kurtosis than the Gaussian distribution. We also consider the asymmetric Gaussian, on which we can adjust the level of asymmetry, and target a more positive asymmetry than empirically observed on the returns. We consider divergence measures based on kernels and on optimal transport. Indeed, we show in this thesis that these divergence measures can be put into explicit form as a function of the target distribution, when the latter has an explicit and parametric distribution function, and have already been used in similar problems, as a replacement for maximum likelihood type methods [BBD⁺19, BBDG19].

In particular, we consider the Maximum Mean Discrepancy (MMD), which for a kernel function and two given distributions measures the “distance” between the representations of the mean of these distributions (or *mean embedding*) in the Hilbert space associated with the kernel. We show that the estimation of the MMD with unbounded kernel function (used in practice in financial applications [BF19]) has the same rate of convergence as the MMD estimators with bounded kernels.

Two other divergence measures are considered in this work: the Kernel Stein Discrepancy (KSD) and the Wasserstein distance. They are written explicitly in terms of the log derivative of the target distribution function for the KSD, and in terms of the inverse of the target distribution function for the Wasserstein distance.

Intuition and empirical behavior of these portfolios. The chosen divergence measures satisfy $D(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T) = 0$ if $\mathbb{P}_{\mathbf{w}} \stackrel{(d)}{=} \mathbb{P}_T$ and are strictly positive otherwise. In practice, there is generally no \mathbf{w} such that the distributions coincide exactly. The distribution of initial returns R is likely to have heavier tails than the target distribution. Empirically, we looked at the case where some components have heavier tails than others. Intuitively, one would expect the minimum divergence approach to systematically reduce the portfolio weight on this component. However, if the variance of this component is very high, this is not necessarily the case.

Optimisation on financial data. We have tested our approach on financial data taken from the Kenneth French library [Fre21]. To minimize our divergence measures, we used the cross-entropy method [RK04]. This method consists in randomly generating \mathbf{w} allocations in the \mathcal{W} set and focusing on the regions of \mathcal{W} that minimize the objective function as iterations proceed, by simply evaluating the function to be minimized at order 0 (no derivative calculations). This method is particularly suitable for non-convex problems. Our empirical results indicate that a target distribution with lighter and more positively skewed tails has an impact on the higher order moments of the portfolio, and that the MMD performs well on our test data.




Figure 9 – Distribution of portfolio returns based on traditional approaches (minimum variance, maximum Sharpe ratio) versus minimum MMD approach with skew Gaussian target. The distribution of the minimum divergence portfolio is slightly less skewed and leptokurtic than traditional portfolios.

In practice, the approach gives better results under normal economic conditions. Regarding this work extensions, one approach that could be considered is to take the divergence function between the empirical and target portfolios as a regularizer added to the global portfolio optimization objective.

INTRODUCTION (EN FRANÇAIS)

0.6 Contexte de la thèse

Cette thèse a été réalisée dans le cadre d'un contrat Cifre au sein du pôle Global Markets de BNP Paribas dans l'équipe de recherche en Stratégies Systématiques et Hybrides (SSH). Côté universitaire, la thèse a été encadrée au Centre de Mathématiques Appliquées de l'École Polytechniques par les Professeurs Emmanuel Gobet et Zoltán Szabó. Jean-Philippe Lemor, Quant Senior et chef de l'équipe SSH, a assuré le suivi de la thèse côté BNPP.

Un des objectifs de l'équipe SSH est de s'assurer de la robustesse des stratégies systématiques proposées par l'équipe de structuration de la BNPP. On entend par stratégie systématique une stratégie d'investissement où les règles d'allocation sont fixées contractuellement dès la mise en place de la stratégie et exécutées automatiquement suivant les valeurs prises par un ensemble de variables d'état représentant l'état du marché financier. Par exemple, la stratégie du portefeuille moyenne-variance, initiée par Markowitz [Mar52], consiste à maximiser l'espérance de gain du portefeuille sous contrainte de volatilité cible: les variables d'état sont dans ce cas la matrice de variance covariance de d sous-jacents constituant l'univers d'investissement, et leur espérance de rendements. En notant C la matrice de covariance sur les rendements des actifs considérés (où on appelle rendement la variation relative du prix de l'actif), μ leur moyenne, \mathcal{W} l'ensemble des contraintes sur le vecteur d'allocation \mathbf{w} et $\sigma_T > 0$ la volatilité cible du portefeuille, le portefeuille mean-variance est solution du problème suivant:

$$\begin{aligned} \max_{\mathbf{w} \in \mathcal{W}} \mathbf{w}^\top \boldsymbol{\mu} \\ \text{s.c } \mathbf{w}^\top C \mathbf{w} = \sigma_T^2. \end{aligned} \tag{0.6.1}$$

En pratique, on va réallouer ce portefeuille à une fréquence prédéterminée (tous les jours ou tous les mois par exemple) en actualisant la valeur de ces variables d'états (C et μ dans le cas Markowitz). Typiquement, C et μ sont estimés de façon empirique, sur un nombre fixe de rendements passés (appelée *taille de fenêtre* de l'estimation), et à chaque nouvelle prise de décision sur le portefeuille, de nouvelles estimations de C et μ sont calculées (ou bien sur la fenêtre de même taille que précédemment -approche fenêtre glissante-, ou bien sur l'ensemble des données historiques disponibles).

Un des indicateurs les plus populaires pour évaluer la performance d'une stratégie est le ratio de Sharpe. Par définition, ce ratio est égal à l'espérance des rendements du portefeuille en excès par rapport au taux sans risque, renormalisée par sa volatilité. On se pose alors la question de savoir si une bonne valeur du ratio de Sharpe² sur la partie backtest a beaucoup de chance de

²Typiquement, un bon ratio de Sharpe doit être supérieur à 1.

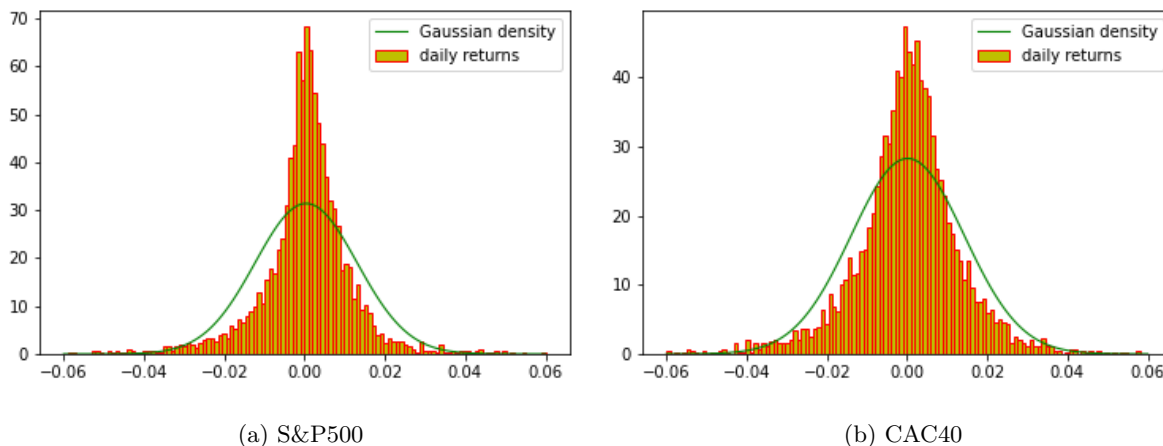


Figure 10 – Histogrammes des rendements journaliers du S&P500 et du CAC40 et de la distribution Gaussienne de même moyenne et variance.

perdurer dans le temps et si le ratio de Sharpe est robuste par rapport à de petites perturbations de ses paramètres initiaux, à savoir les paramètres de la stratégie (taille de fenêtre, fréquence de recalibration...), et les données historiques, ou backtest de la stratégie. Pour ce faire, des études appelées “tests de robustesse” sont menées par l’équipe SSH.

C’est dans ce contexte industriel que s’inscrit cette thèse. L’objectif initial de cette thèse était d’une part d’adapter et d’appliquer des inégalités de concentration à des variables type valeurs de portefeuille. Les inégalités de concentration permettent d’avoir des intervalles de confiance sur l’écart entre un estimateur empirique, mesuré sur une réalisation finie d’observations, et sa vraie valeur. Ces intervalles de confiance sont établis en respectant le caractère fini du nombre d’observations et dépendent de la loi de l’estimateur à nombre d’observations donné. D’autre part, un objectif était d’adapter les techniques de quantification d’incertitudes pour le ratio de Sharpe d’une stratégie, c’est-à-dire trouver une manière d’exprimer le ratio de Sharpe comme fonction de ses paramètres incertains.

Nous nous sommes d’abord intéressés aux propriétés statistiques des données financières pour cibler sous quelles hypothèses et cadre devaient s’effectuer nos recherches.

Distribution des rendements et dépendance en temps. La théorie du marché efficient initiée au début du 20^{ème} siècle [Bac00, Fam71] selon laquelle les variations boursières ne seraient que bruit blanc a été largement remise en cause au fil des décennies. Cette théorie repose sur l’idée que dans un marché de libre concurrence où tous les acteurs disposent de la même information, les prix s’équilibrent automatiquement indépendamment du passé. En temps continu, le modèle de Bachelier correspond à modéliser les variations de prix comme suivant un mouvement brownien:

$$S_{t+1} - S_t = \sigma(W_t - W_{t-1}), \quad W_t \text{ mouvement brownien.}$$

Osborne en 1959 propose de modéliser le log des ratios de prix comme suivant une loi Gaussienne [Osb59], théorie incorporée dans le fameux modèle de Black et Scholes en 1973 [BS73].

Cependant, et dès la création de ces théories, il était clair que l’hypothèse Gaussienne et d’indépendance temporelle, que ce soit pour les variations journalières de prix ou les log rendements, n’était pas adaptée. En effet, comme illustré Figures 10 et 11, la distribution des rendements boursiers présente des queues plus épaisses que la distribution Gaussienne. Les histogrammes des rendements (Figure 10) montre des distributions piquées avec des queues




Figure 11 – Diagrammes quantile-quantile: quantile des rendements journaliers du S&P500 et du CAC40 en fonction des quantiles de la distribution Gaussienne centrée réduite.

légèrement au-dessus de la distribution Gaussienne. Les diagrammes quantile-quantiles (Figure 11) montrent clairement que les quantiles extrêmes des rendements boursiers ne coïncident pas avec ceux de la distribution Gaussienne: la relation aurait dû être purement linéaire sinon (points bleus alignés avec ligne rouge), or ici les quantiles empiriques s'écartent, signe que les rendements prennent des valeurs plus extrêmes que celles qu'auraient prises des rendements Gaussiens à un niveau de probabilité équivalent. Par ailleurs, l'hypothèse d'indépendance dans le temps de la volatilité des rendements est remise en cause par les valeurs fortes prises par l'autocorrélation des rendements en valeur absolue Figure 12. Les rendements en valeur absolue peuvent être considérés comme un proxy de la volatilité des rendements. L'autocorrélation est définie comme la corrélation entre les rendements en valeur absolue à un instant t et leur valeur future à l'instant $t + lag$: l'autocorrélation est ici non négligeable, comparée à l'autocorrélation d'un bruit blanc indépendant (voir [BD16, Figure 1.13] pour l'autocorrélation d'un bruit blanc par exemple).

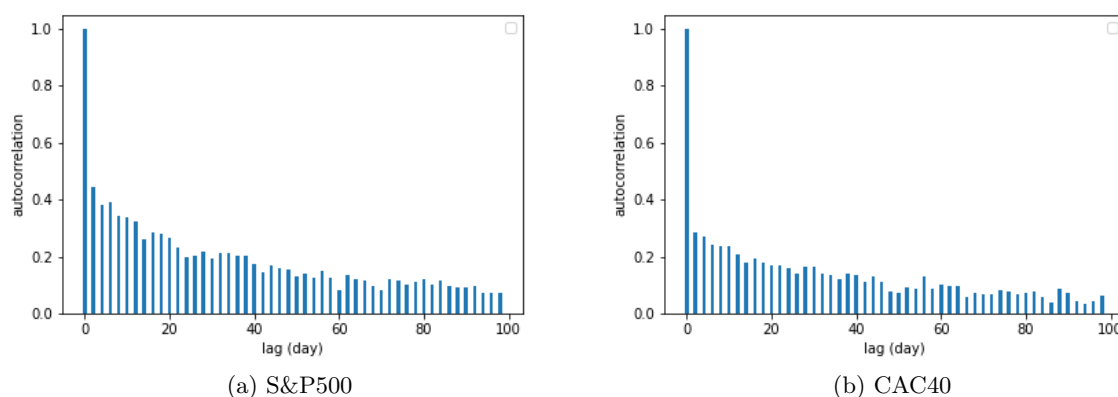


Figure 12 – Autocorrélogramme des rendements journaliers en valeur absolue du S&P500 et du CAC40.

Mandelbrot est le premier à parler de faits stylisés, c'est-à-dire de propriétés statistiques vérifiées quasi-systématiquement sur les rendements boursiers, dans [Man63]. Plus récemment, Cont [Con01] dresse une liste de ces propriétés. Parmi ces faits stylisés figurent la non-Gaussianité des rendements journaliers et leur caractère leptokurtique, c'est-à-dire leurs queues plus épaisses que la distribution Gaussienne, la dépendance en temps de la variance des rendements, mais

aussi l'asymétrie négative des rendements, et les effets de cluster de volatilité.

Plan de la thèse. Dans ce qui suit nous détaillons le cheminement suivi par la thèse, avec comme leitmotiv la prise en compte des incertitudes dans le contexte distribution à queues épaisses pour les portefeuilles financiers.

Notre objectif premier était de s'intéresser aux inégalités de concentration. Ces inégalités permettent de quantifier l'écart entre un estimateur et son espérance, en fonction du nombre d'observations dont dépend l'estimateur. La présence de queues épaisses dans les rendements financiers a motivé le travail sur des nouvelles inégalités de concentration adaptées à ce contexte. Dans la première partie de cette thèse, nous établissons des inégalités de concentration adaptées aux variables type α -exponentielle (Chapitre 1), ou log-normale (Chapitre 2), qui sont deux exemples de distributions à queues épaisses.

Dans un second temps, nous nous sommes intéressés à la dépendance temporelle de la variance dans les rendements financiers et à l'impact de leur prise en compte. La question de départ était de savoir s'il valait mieux considérer un modèle adaptatif pour l'estimation de la matrice de covariance C dans des problèmes type (0.8.2), ou si la simple estimation historique de C était efficace. Comme illustré sur le S&P500 et le CAC40 (Figure 5.9), la variance sur les rendements bouge avec le temps, ce qui rend intuitivement préférable d'utiliser un modèle adaptatif, type GARCH, plutôt qu'une covariance historique sur une taille d'échantillon trop élevée pour rendre compte des mouvements courts termes de la volatilité.




Figure 13 – Tracé de la volatilité mesurée par fenêtre glissante sur les rendements journaliers du S&P500 et du CAC40.

De plus, les hypothèses de modèle GARCH coïncident avec la présence de queues de distributions épaisses, et impose l'existence de moments au plus polynomiaux. On montre Figure 14 l'estimation du nombre de moments finis d'après le modèle GARCH(1,1) pour 12 instruments financiers: les plus hauts moments finis sur ces instruments ne dépassent pas l'ordre 6. L'étude approfondie de ces modèles pour la performance du portefeuille variance minimum constitue la seconde partie (Chapitre 3) de cette thèse.

Dans un troisième temps (Chapitre 4), nous avons cherché à appliquer des techniques de quantification d'incertitudes pour le portefeuille financier. En quantification d'incertitudes, on suppose que certains paramètres sont incertains et on cherche à exprimer "facilement" la solution d'un problème en fonction de ce paramètre incertain, c'est-à-dire sans avoir à résoudre le problème initial pour chaque valeur du paramètre incertain. Typiquement, cela peut se faire en évaluant une (approximation de l')Expansion en Polynômes du Chaos (EPC) de la solution du problème, c'est-à-dire en calculant une projection de la fonction dans une base de polynômes du paramètre incertain. Nous nous sommes concentrés sur l'application de ces techniques sur le




Figure 14 – Plus haut moment ρ_{\max} sur la variance de rendements calibrés avec le modèle GARCH.




Figure 15 – Diagramme quantile quantile: quantiles de la corrélation calibrée par fenêtre glissante entre deux actifs en fonction des quantiles de la loi uniforme.

ratio de Sharpe et l’allocation de portefeuille.

Dans nos illustrations, nous nous sommes concentrés sur l’impact de l’incertitude sur la corrélation entre des actifs. En pratique, en estimant la corrélation par moyenne mobile, sur un intervalle de temps assez large pour capter le niveau de corrélation “stationnaire” entre deux actifs, on s’aperçoit que la distribution des corrélations calibrées est relativement étendue et non-uniforme, comme montré Figure 15 (diagramme quantile quantile contre la loi uniforme de la distribution d’une corrélation calibrée sur 2 instruments financiers). Il était donc intéressant d’en quantifier l’impact.

Enfin, la dernière partie de cette thèse (Chapitre 5) est consacrée à l’optimisation de portefeuille avec distribution cible, via des mesures de divergence basées sur les kernel et la théorie du transport optimal. D’un point de vue financier, le but de cette approche est de construire des portefeuilles avec des queues de distribution moins épaisses et une asymétrie moindre qu’avec les approches traditionnelles d’optimisation de portefeuille.

Les parties de cette thèses peuvent être lues de façon indépendante. Pour guider le lecteur, et lui permettre de l’aider dans ses choix de lecture, nous donnons dans ce qui suit un résumé de chacune des parties de la thèse.

0.7 Inégalités de concentration pour les variables à queues épaisses: une approche par norme Orlicz et inégalité de Talagrand

Contexte. De par l'essor des techniques d'estimation statistique, l'étude du comportement des estimateurs statistiques et de leur écart à la moyenne devient primordiale. Les inégalités de concentration permettent de quantifier cet écart en fonction de la taille de l'échantillon utilisé par l'estimateur. Étant donné $M \in \mathbb{N}^*$, Y_1, \dots, Y_M suite de variables aléatoires, généralement indépendantes, et \mathcal{F} une classe de fonctions mesurables, les inégalités de concentration se trouvent sous la forme:

$$\forall \varepsilon > 0, \quad \mathbb{P} \left(\left| \frac{1}{M} \sum_{m=1}^M Y_m - \mathbb{E}Y \right| \geq \varepsilon \right) \leq b(\varepsilon, M), \quad (\text{concentration de la somme})$$

$$\forall \varepsilon > 0, \quad \mathbb{P} \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m=1}^M f(Y_m) - \mathbb{E}f(Y) \right| \geq \varepsilon \right) \leq B(\varepsilon, M), \quad (\text{concentration uniforme})$$

où b et B décroissent vers 0 quand M est grand. En pratique, b et B peuvent dépendre de la distribution des $\{Y_m\}_{m \in [M]}$ et de la complexité de la classe \mathcal{F} . Ces inégalités s'interprètent comme à quel point à M donné l'estimateur est concentré autour de sa moyenne, converge vite vers celle-ci avec M , dans le cas pire écart sur la classe \mathcal{F} pour B .

La théorie de l'apprentissage statistique a contribué à la popularité de ces inégalités (voir [Vap00, Chapitre 3] et [Bou02]). Cette théorie traite généralement de concentration sur l'erreur (bornée) entre la prédiction d'un algorithme et sa valeur cible indépendamment de la distribution des observations. Ce n'est pas notre cadre de travail, puisque dans un contexte financier, et dans de nombreuses applications (cf. Chapitre 1: application aux dérivées de kernels), f n'est pas bornée, et admet une distribution à queues épaisses. Les résultats abordant le cas de distributions non bornées se focalisent sur le cas Gaussien ou sous de fortes hypothèses sur les moment type condition de Bernstein [BLM13, Section 2.3 et 2.8] non vérifiée pour une distribution à queues épaisses. Nous cherchons donc à combler ce vide.

On entend par distribution à queues épaisses les distributions possédant des queues plus épaisses que la loi Gaussienne ou exponentielle mais moins épaisses que les distributions à queues lourdes, type Pareto. On peut séparer ces distributions en deux classes:

1. les distributions à queue α -exponentielle: $\mathbb{E} [e^{s|Y|^\alpha}] < \infty$ pour un $\alpha \in (0, 1)$ et un $s > 0$,
2. les distributions n'admettant pas de moment α -exponentiel, mais tous les moments polynomiaux (typiquement les variables log-normales).

Notre objectif est donc d'établir des inégalités de concentration dans ces deux cas.

Orlicz-Talagrand: motivation et techniques de preuve. Étant donnée $\Psi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ continue, croissante, telle que $\Psi(0) = 0$ (dite "fonction Orlicz"), on définit la norme Ψ -Orlicz d'une variable aléatoire X prenant ses valeurs dans un espace de Banach B comme:

$$\|X\|_\Psi := \inf \left\{ c > 0 : \mathbb{E} \left[\Psi \left(\frac{\|X\|_B}{c} \right) \right] \leq 1 \right\}. \quad (0.7.1)$$

Les normes Orlicz sont utiles pour établir des inégalités de concentration car:

1. Il est possible de construire une fonction Orlicz sur-mesure par rapport à la distribution.
2. Le schéma de preuve pour établir des inégalités de concentration par norme Orlicz est (relativement) indépendant de la norme Orlicz choisie (avec inégalité d'autant plus robuste que la fonction Orlicz est adaptée à la distribution).

1. Choix de la fonction Ψ

On va chercher la fonction Orlicz à la croissance la plus forte possible vérifiant $\mathbb{E}[\Psi(sX)] < \infty$ pour au moins un $s > 0$. Exemples:

- dans le cas α -exponentiel, la fonction Orlicz limite, étudiée par [Tal89][Ada08], est donnée par

$$\Psi_\alpha(x) = \exp(x^\alpha) - 1,$$

- pour traiter le cas intermédiaire, on introduit la nouvelle fonction Orlicz:

$$\Psi_\beta^{\text{HT}}(x) := \exp\left((\ln(x+1))^\beta\right) - 1$$

pour un paramètre $\beta > 1$. Le cas log-normal est couvert avec $\beta < 2$.

2. Schéma général de preuve

Ici par simplicité on suppose que $X \in \mathbb{R}$. On peut établir une inégalité de concentration robuste sur X telle que $\|X\|_\Psi < \infty$ par simple inégalité de Markov en exploitant la définition de la norme Orlicz. Par exemple, dans le cas α -exponentiel, en composant par la fonction croissante $\Psi_\alpha(\cdot/\|X\|_{\Psi_\alpha}) + 1$, puis en appliquant l'inégalité de Markov, on a, pour $\varepsilon > 0$:

$$\begin{aligned} \mathbb{P}(|X| > \varepsilon) &= \mathbb{P}\left(\Psi_\alpha\left(\frac{|X|}{\|X\|_{\Psi_\alpha}}\right) + 1 > \Psi_\alpha\left(\frac{\varepsilon}{\|X\|_{\Psi_\alpha}}\right) + 1\right) \\ &\stackrel{\text{Markov}}{\leq} \underbrace{\left(\mathbb{E}\left[\Psi_\alpha\left(\frac{|X|}{\|X\|_{\Psi_\alpha}}\right)\right] + 1\right)}_{\leq 1} / \underbrace{\left(\Psi_\alpha\left(\frac{\varepsilon}{\|X\|_{\Psi_\alpha}}\right) + 1\right)}_{=e^{\frac{\varepsilon^\alpha}{\|X\|_{\Psi_\alpha}^\alpha}} - 1} \leq 2 \exp\left(-\frac{\varepsilon^\alpha}{\|X\|_{\Psi_\alpha}^\alpha}\right). \end{aligned}$$

Ce schéma de preuve est utilisé dans [Ada08, vdGL13]. En pratique, $X = \frac{1}{M} \sum_{m=1}^M Y_m$ ou $X = \sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{m=1}^M f(Y_m)$. Pour garantir $b(\varepsilon, M) = \mathcal{O}(M)$ et $B(\varepsilon, M) = \mathcal{O}(M)$, deux inégalités vont être fondamentales: l'inégalité de Talagrand [Tal89, Theorem 3: $\Psi = \Psi_\alpha$] et l'inégalité maximale [vW96, Lemma 2.2.2: pour Ψ convexe + condition de croissance]

$$\left\| \sum_{m \in [M]} Y_m \right\|_\Psi \leq C_\Psi \left(\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \right), \quad (0.7.2)$$

$$\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \leq K \Psi^{-1}(M) \max_{m \in [M]} \|Y_m\|_\Psi. \quad (0.7.3)$$

Lorsque ces inégalités sont vérifiées, on a facilement $b(\varepsilon, M) = \mathcal{O}(M)$.

Dans le cas uniforme, on peut être tenté d'appliquer directement l'inégalité de Talagrand. Mais on se retrouverait alors avec un terme en $\left\| \sum_{m \in [M]} \sup_{f \in \mathcal{F}} f(Y_m) \right\|_{L_1(B)}$ difficilement contrôlable.

Une astuce est d'avoir recours à une technique de troncation (utilisée par Adamczak [Ada08, Theorem 4] pour établir une inégalité de grande déviation que nous avons adaptée et corrigée pour établir des inégalités de concentration): introduisons pour un niveau de troncation $c > 0$ la fonction tronquée $\mathcal{T}_c f := -c \vee f \wedge c$ et le reste par la fonction tronquée $\mathcal{R}_c f := f - \mathcal{T}_c f$. On peut alors traiter séparément la concentration de la partie tronquée (fonction bornée: problème classique) et du reste. Une application de l'inégalité de Talagrand fait apparaître le terme: $\left\| \sum_{m \in [M]} \sup_{f \in \mathcal{F}} |\mathcal{R}_c f(X_m)| \right\|_{L_1(B)}$. Intuitivement, en choisissant c assez grand, ce terme de "reste" devient petit. On peut grâce à l'inégalité dite de Hoffman-Jorgensen [LT13, Proposition 6.8] calibrer c de façon à garantir que ce terme soit négligeable devant M .

Problèmes rencontrés et contributions. Outre la mise en place d'un schéma assez générique de preuve, nous avons contribué à combler les problèmes suivants:

- Dans le cas α -exponentiel, l'inégalité de Talagrand (0.7.2) est déjà disponible dans la littérature [Tal89, Theorem 3], mais la fonction Ψ_α , non convexe, ne satisfait pas les hypothèses de l'inégalité maximale (0.7.3). Nous avons contribué à établir cette inégalité dans le cas α -exponentiel.
- Dans le cas β -heavy, l'inégalité maximale est vérifiée pour la fonction convexe $\Psi_\beta^{\text{HT}}(x) := \exp((\ln(x+1))^\beta) - 1$, mais l'inégalité de Talagrand n'est pas disponible dans la littérature: en se basant sur la preuve de [Tal89, Theorem 3], on a contribué à étendre l'inégalité (0.7.2) au cas $\Psi = \Psi_\beta^{\text{HT}}$.

Nous répondons à ces problèmes dans notre travail, mais n'allons pas détailler en introduction les solutions apportées.

0.8 Analyse asymptotique de l'estimation de matrice de covariance pour le portefeuille minimum variance

Dans cette partie, nous traitons d'un problème beaucoup plus pratique, qui consiste à analyser l'impact de la prise en compte de la meilleure estimée de la matrice de covariance dans le problème du portefeuille minimum variance.

La théorie du portefeuille moyenne-variance introduite par Markowitz [Mar52] consiste à trouver l'allocation de portefeuille \mathbf{w} qui minimise $\mathbf{w}^\top C \mathbf{w} - \gamma \mathbf{w}^\top \mu$, où C correspond à la matrice de covariance sur les rendements des actifs considérés, μ leur moyenne, et $\gamma > 0$ l'appétence au risque de l'investisseur. En pratique, C et μ doivent être estimés, et de par la nature du problème (solution explicite dépendant de l'inverse de C), mènent à une instabilité significative des portefeuilles estimés. De nombreux articles traitent du problème d'estimation de μ [BG91, GHZ13]; nous nous focalisons ici sur celle de C avec $\gamma = 0$.

On se place dans le cadre suivant: on considère un univers de d actifs, dont les rendements suivent la loi conditionnelle $\mathbf{r}_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, V_t)$ avec V_t matrice de covariance instantanée, de moment polynomial maximal p_{\max} . Comme expliqué précédemment, un modèle dynamique de la covariance est cohérent avec les faits stylisés observés sur le marché. On considère un horizon d'investissement H ; dans ce cas, la *meilleure estimée de la matrice de covariance réalisée* est

$$cRC_{H,t} := \mathbb{E} \left[\frac{1}{H} \sum_{k=1}^H \mathbf{r}_{t+k} \mathbf{r}_{t+k}^\top | \mathcal{F}_t \right] = \frac{1}{H} \sum_{k=1}^H \mathbb{E} [V_{t+k} | \mathcal{F}_t]. \quad (0.8.1)$$

Généralement, le praticien utilise une matrice de covariance historique $\frac{1}{T} \sum_{k=1}^T \mathbf{r}_{t-k} \mathbf{r}_{t-k}^\top$ basée sur les rendements passés. Quand la taille de l'historique T tend vers l'infini, cet estimateur tend vers la matrice de covariance stationnaire du processus. Or, lorsque l'horizon d'investissement est court et que la covariance instantanée du processus prend des valeurs différentes de la matrice de covariance stationnaire (moments de haute ou basse volatilité court terme), alors il vaut mieux pour prendre en compte des effets de latence de la variance utiliser $cRC_{H,t}$.

On considère le problème du portefeuille de variance minimum: étant donnée une matrice de covariance C , ce problème consiste à trouver l'allocation solution de

$$\text{mv}(C) := \arg \min_{\mathbf{w} \in \mathcal{W}} \mathbf{w}^\top C \mathbf{w}, \quad (0.8.2)$$

où $\mathcal{W} \subset \mathbb{R}^d$ tient compte des contraintes sur \mathbf{w} (typiquement, $\left\{ \mathbf{w} \in \mathbb{R}^d : \sum_{i=1}^d w_i = 1 \right\} \subset \mathcal{W}$). Dans ce cas, la métrique de performance est la *variance réalisée future* et l'investisseur cherche

à ce qu'elle soit la plus petite possible. En posant $RC_{H,t_n} = \frac{1}{H} \sum_{k=1}^H \mathbf{r}_{t_n+k} \mathbf{r}_{t_n+k}^\top$ (matrice de *covariance réalisée*), les variances réalisées associées au portefeuille minimum variance (0.8.2) réalloué d'une part avec la meilleure estimée de matrice de covariance (0.8.1), et d'autre part avec une matrice de covariance benchmark V_{ref} , s'écrivent comme

$$\begin{cases} R_{N,H} & := \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top RC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\ R_{N,H}^{\text{ref}} & := \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top RC_{H,t_n} \text{mv}(V_{\text{ref}}). \end{cases}$$

Notre objectif est de montrer qu'avec grande probabilité, $R_{N,H} < R_{N,H}^{\text{ref}}$, d'autant plus sûrement que N croît. Intuitivement/asymptotiquement, on a bien $R_{N,H}^{\text{ref}} - R_{N,H} \geq 0$:

- Introduisons les termes de variance conditionnelle:

$$\begin{cases} cRV_{N,H} & := \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top cRC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\ cRV_{N,H}^{\text{ref}} & := \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top cRC_{H,t_n} \text{mv}(V_{\text{ref}}). \end{cases} \quad (0.8.3)$$

Par définition du portefeuille de variance minimum, on a toujours $cRV_{N,H} \leq cRV_{N,H}^{\text{ref}}$.

- De plus, comme $\mathbb{E}[RC_{H,t_n} | \mathcal{F}_{t_n}] = cRC_{H,t_n}$, les différences $R_{N,H} - cRV_{N,H}$ et $R_{N,H}^{\text{ref}} - cRV_{N,H}^{\text{ref}}$ sont des martingales centrées. Donc:

$$\frac{1}{N}(R_{N,H}^{\text{ref}} - R_{N,H}) = \underbrace{\frac{1}{N}(R_{N,H}^{\text{ref}} - cRV_{N,H}^{\text{ref}}) - \frac{1}{N}(R_{N,H} - cRV_{N,H})}_{\frac{1}{N} \times \text{martingales centrées} \approx 0} + \underbrace{\frac{1}{N}(cRV_{N,H}^{\text{ref}} - cRV_{N,H})}_{\geq 0}.$$

Question: Peut-on prouver via une inégalité de concentration que la probabilité de $R_{N,H} < R_{N,H}^{\text{ref}}$ tend rapidement vers 0 avec le nombre de rebalancement? Comme on se place dans le cadre où V_t donc r_t a peu de moments (polynomial uniquement, jusqu'à l'ordre p_{\max}), il est difficile d'obtenir des bornes décroissant très vite.

En se basant sur des contrôles de moments d'ordre p de somme d'incrément martingale (inégalité de Burkholder [HH80, Theorem 2.10]) et de somme de fonctions de processus ergodique (inégalité de Fort-Moulines [FM03, Proposition 2]), sous de petites hypothèses techniques sur V_t , on établit le résultat suivant: pour tout $H \in \mathbb{N}^*$ horizon d'investissement, en notant $\ell_H := \mathbb{E} \left[\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N} \right]$ l'écart de performance, il existe une constante $C > 0$ dépendant de H , p_{\max} , d et de V_t telle que, pour tout $N \in \mathbb{N}^*$, nombre rebalancements du portefeuille, on a:

$$\mathbb{P} \left(R_{N,H} > R_{N,H}^{\text{ref}} \right) \leq \frac{C}{\ell_H^{p_{\max}}} \frac{H^{\frac{p_{\max}}{2}}}{N^{\frac{p_{\max}}{2}}}.$$

Application avec un modèle GARCH-CCC. En supposant un modèle paramétrique sur V_t type GARCH-Corrélation Conditionnelle Constante et en supposant les paramètres du modèle connus, on peut estimer la matrice de covariance optimale par technique de Monte Carlo. Numériquement, l'évolution de la probabilité de succès tend bien vers 0 avec le nombre de rebalancements des portefeuilles N (cf Figure 16).




Figure 16 – $\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^\infty)$, pour $H = 1$ (rouge), $H = 5$ (bleu) and $H = 21$ (vert) en fonction du nombre de rebalancement N , en échelle log-log, pour des paramètres GARCH tels que $d = 50$ et $p_{\max} = 5.05$.

Horizon d’investissement limite. Quand l’horizon d’investissement H tend vers l’infini, la matrice meilleure estimée coïncide avec la matrice de covariance stationnaire. En pratique, on observe que lorsque l’horizon d’investissement est de l’ordre de grandeur du multiple du temps de demi-vie du processus, les performances du portefeuille avec covariance estimée se dégradent jusqu’à un horizon où le choix entre les matrices de covariance produit des performances équitables: l’investisseur est alors indifférent entre la matrice stationnaire ou meilleure estimée.

Un problème non traité ici est la prise en compte des incertitudes paramétriques du modèle. Nous donnons dans la partie suivante une méthode pour traiter les incertitudes paramétriques.

0.9 Traitement des incertitudes paramétriques pour l’optimisation de portefeuille: une approche par décomposition du chaos via l’algorithme UQSA

Comprendre comment les incertitudes paramétriques perturbent le résultat d’un problème d’optimisation est un problème fondamental. Le domaine de la quantification des incertitudes est l’ensemble des techniques qui permettent de quantifier cet impact.

Supposons un modèle paramétrique sur les entrées du problème, avec un paramètre incertain θ dont on suppose qu’on peut modéliser l’incertitude via une loi π : les entrées V suivent la loi $\mu(\theta, dv)$ et $\theta \sim \pi$. On cherche une représentation de la solution du problème d’optimisation en fonction de θ . En supposant la fonction de carré intégrable selon π , pour toute base orthonormale de polynômes $\{B_i\}_{i \in \mathbb{N}}$, on peut exprimer la solution du problème d’optimisation $z^*(\theta)$ via sa projection sur cette base de polynômes:

$$z^*(\theta) = \sum_{i \geq 0} u_i B_i(\theta), \quad \forall \theta \in \Theta \subset \mathbb{R}^p, \quad \text{et} \quad \int_{\Theta} B_i(\theta) B_j(\theta) \pi(d\theta) = \begin{cases} 0 & \text{si } i \neq j, \\ 1 & \text{sinon.} \end{cases} \quad (0.9.1)$$

C’est ce qu’on appelle la décomposition en polynôme du chaos, et c’est une des méthodes les plus utilisées en quantification d’incertitudes [LMK10]. En pratique, on évalue une forme tronquée de (0.9.1), car l’erreur de troncature tend généralement vite vers 0 ($\sum_{i > m} |u_i|^2 = \mathcal{O}(m^{-\delta})$ avec δ ordre de différentiabilité de la fonction z^*).

L’algorithme UQSA [CFGS20] permet une estimation progressive des coefficients $\{u_i\}_{i \leq m_k}$,

m_k croissant, lorsque le problème se met sous la forme

$$z^*(\theta) \text{ solution de } \mathbb{E}[H(z^*(\theta), V_\theta, \theta)] = \int_{\mathcal{V}} H(z, v, \theta) \mu(\theta, dv) = 0, \quad \pi - \text{p.p.}, \quad (0.9.2)$$

Dans le contexte financier, on s'est intéressé à deux cas d'application de ces techniques: le traitement des incertitudes pour le ratio de Sharpe et pour l'allocation de portefeuille lorsque les taux d'accroissement financiers suivent une loi paramétrique incertaine.

Ratio de Sharpe. Le ratio de Sharpe [Sha66] rend compte de la rentabilité d'un portefeuille renormalisée par son risque. En notant R le rendement en excès du taux sans risque du portefeuille, il s'écrit:

$$\text{SR} = \frac{\mathbb{E}[R]}{\sqrt{\text{Var}[R]}}.$$

Une difficulté pour appliquer l'approche UQSA a été de trouver un moyen d'exprimer le ratio de Sharpe comme solution d'un problème de la forme (0.9.2). Cette difficulté a été surmontée grâce à l'astuce suivante. En prenant deux copies indépendantes et de même loi R, \tilde{R} , on peut exprimer le carré du ratio de Sharpe comme le zéro d'une fonction:

$$\text{SR}^2 \text{Var}[R] - \mathbb{E}[R]^2 = \mathbb{E} \left[\text{SR}^2 \frac{(R - \tilde{R})^2}{2} - R\tilde{R} \right] = 0.$$

Pour nos applications, on s'est intéressé à l'évolution du ratio de Sharpe lorsque la corrélation entre les actifs est incertaine (avec une distribution paramétrique calquée sur la distribution empirique représentée via son quantile plot Figure 15). On illustre l'intérêt d'utiliser UQSA avec un nombre croissant de coefficients estimés en fonction des itérations en Figure 17: l'erreur totale est plus faible avec $m_k = \lfloor k + 1 \rfloor^{0.3}$ croissant plutôt qu'à m_k fixe.




Figure 17 – Évolution de l'erreur totale sur l'approximation du ratio de Sharpe dans le cas $m_k = \lfloor k + 1 \rfloor^b$ croissant (ligne bleue) versus $m_k = m \in \{10, 30\}$ (lignes pointillées) en fonction du nombre d'itérations k .

Allocation de portefeuille. On s'intéresse à la quantification des incertitudes pour l'allocation de portefeuille donné par $\arg \min_{\mathbf{w} \in \mathcal{W}} U(\mathbf{w})$ où U est convexe. Sans contraintes ($\mathcal{W} = \mathbb{R}^d$), le problème peut généralement facilement se mettre sous la forme (0.9.2), en prenant le gradient de U et en essayant de l'exprimer comme une espérance. On peut alors appliquer l'algorithme UQSA sur ce problème. On aimerait traiter du cas où \mathcal{W} correspond à l'ensemble des portefeuilles à composantes positives (achat uniquement) et de somme 1:


$$\mathcal{W}^{\geq 0} = \left\{ \mathbf{w} \in \mathbb{R}^d : w_i \geq 0, \sum_{i=1}^d w_i = 1 \right\}.$$

Avec contraintes, on se heurte à deux difficultés:

1. La solution du problème de minimisation n'est pas directement le zéro du gradient de U (il faut introduire des multiplicateurs de Lagrange et regarder les conditions de KKT du problème),
2. Résoudre naïvement le problème via un algorithme d'approximation stochastique ne permet pas de rester dans $\mathcal{W}^{\geq 0}$.

On propose de passer par un changement de variable: on introduit $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ tel que $\ell(\mathbf{y}) = \mathbf{w}$ et de résoudre le problème en $\mathbf{y} \in \mathbb{R}^{d-1}$. De nombreuses techniques dites d'analyse de données compositionnelles [PGB11] permettent de passer de \mathbb{R}^{d-1} dans $\mathcal{W}^{>0} = \{\mathbf{w} \in \mathbb{R}^d : w_i > 0, \sum_{i=1}^d w_i = 1\}$. Typiquement, la transformation log-ratio: en prenant $w_i = \frac{e^{y_i}}{1 + \sum_{j=1}^{d-1} e^{y_j}}$, pour $i = 1, \dots, d-1$, $w_d = \frac{1}{1 + \sum_{j=1}^{d-1} e^{y_j}}$. Dans cette thèse, on a confectionner une nouvelle transformation permettant de couvrir le cas $w_i = 0$ qu'on appelle la transformation intégrale et qui repose sur l'idée suivante: étant donnée $f : \mathbb{R} \rightarrow \mathbb{R}_+$ telle que $\int_{\mathbb{R}} f = 1$, alors on peut interpréter les w_i comme des portions de l'intégrale de f . En introduisant la fonction positive $p : \mathbb{R} \rightarrow \mathbb{R}_+$ telle que $p(0) = 0$, il existe donc $\mathbf{y} \in \mathbb{R}^{d-1}$ tel que:

$$\begin{aligned} w_1 &= \int_{-\infty}^{y_1} f(t) dt, \\ w_i &= \int_{y_1 + p(y_2) + \dots + p(y_{i-1})}^{y_1 + p(y_2) + \dots + p(y_i)} f(t) dt, \\ w_d &= \int_{y_1 + p(y_2) + \dots + p(y_d)}^{\infty} f(t) dt. \end{aligned}$$



Sont étudiés les cas où f correspond à la dérivée de la fonction sigmoïd, et un cas à support compact, avec f prise comme la fonction de distribution d'une loi beta(2, 2).

Une limite de cette approche est que le problème est bien posé seulement dans le cas *non saturé*, c'est-à-dire lorsque \mathbf{w} prend ses valeurs dans l'intérieur du simplexe $\mathcal{W}^{>0}$. Dans le cas où \mathbf{w} admet des composantes nulles, la formulation en problème (0.9.2) peut admettre plusieurs solutions, suggérant que cette approche ne peut être utilisée telle qu'elle dans ces cas.

0.10 Optimisation de portefeuille avec distribution cible: une approche via mesure de divergence basée sur les théories du kernel et du transport optimal

Changement de paradigme. Jusqu'ici, on s'est concentré dans nos applications sur le paradigme moyenne-variance, et on a proposé des manières de quantifier ou de gérer les incertitudes, que ce soit par une meilleure estimation des premiers moments (approche meilleure estimée de la matrice de covariance), ou par la quantification d'incertitude (pour voir la distribution du portefeuille ou de la performance sachant les incertitudes paramétriques). C'est en effet le paradigme le plus utilisé dans la construction de portefeuille, bien qu'il soit reconnu pour créer des portefeuilles instables et très sensibles aux erreurs d'estimation des paramètres.

Or en pratique, les gestionnaires d’actifs ont de plus en plus tendance à délaisser les approches type utilité moyenne-variance pour se tourner vers des fonctions objectives ou des méthodes d’optimisation de portefeuille centrées sur la gestion du risque du portefeuille, par exemple en incluant les moments d’ordre 4 dans la fonction objective pour diminuer la kurtosis en excès du portefeuille [MZ10], ou en passant par des approches type optimisation de portefeuille avec gestion du risque par classe d’actifs [MRT10a, MRT10b].

L’approche divergence minimum. C’est dans cette optique d’une meilleure prise en compte des moments d’ordres élevés du portefeuille que s’inscrit l’approche du portefeuille” minimum divergence “avec distribution cible [CW12, Las19, GLLN20]: étant donné D mesure de divergence, le portefeuille minimum divergence est donné par

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} D(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T) \quad (0.10.1)$$

où $P = \mathbf{w}^\top R \sim \mathbb{P}_{\mathbf{w}}$ désigne la distribution du portefeuille, et \mathbb{P}_T la distribution cible de l’investisseur. En pratique, on suppose avoir accès à des rendements passés $\{\mathbf{r}_n\}_{n \in [N]}$, et on peut estimer de façon empirique, ou paramétrique, la fonction objective (0.10.1). Contrairement aux auteurs cités, nous reprenons cette approche avec une volonté de ne pas imposer de modèle paramétrique sur les rendements, et on aimerait autant que possible pouvoir expliciter la fonction objective en fonction de \mathbb{P}_T :

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} D_{\mathbb{P}_T} \left(\{\mathbf{w}^\top \mathbf{r}_n\}_{n \in [N]} \right).$$

Questions. Dans ce contexte, plusieurs questions se posent: quelle mesure de divergence (c’est-à-dire distance entre la distribution du portefeuille et la distribution cible) choisir, quelle distribution cible considérer, et quelle méthode d’optimisation serait adéquate pour ce problème généralement non convexe?

Choix de divergence et de distribution. Nous nous sommes concentrés sur un choix de distributions paramétriques type Gaussienne généralisée (distribution sur laquelle on peut spécifier le niveau de kurtosis, ou moment d’ordre 4, ce qui permet de quantifier l’épaisseur de queue de la distribution cible et viser une kurtosis moins élevée que la loi Gaussienne) et Gaussienne asymétrique (sur laquelle on peut bouger le niveau d’asymétrie, et cibler une asymétrie plus positive qu’empiriquement observée sur les rendements), et sur des mesures de divergence basées sur les noyaux et sur le transport optimal (N.T.O.). En effet, on donne dans cette thèse une formulation explicite de ces mesures de divergence en fonction de la distribution cible, quand cette dernière possède une fonction de distribution explicite et paramétrique, ce qui avait déjà été fait en partie dans des problèmes de minimisation de divergence comme méthode alternative aux méthodes type maximum de vraisemblance pour estimer cette fois des paramètres d’une distribution cible donnée [BBD⁺19, BBDG19].

En particulier, nous avons considéré comme mesure de divergence la *discrépance maximale de la moyenne* (ou Maximum Mean Discrepancy (MMD)), qui pour une fonction noyau et deux distributions données mesure la distance entre les représentations de la moyenne de ces distributions (ou *mean embedding*) dans l’espace de Hilbert associé au noyau. Nous démontrons que l’estimation de la MMD avec fonction noyau non bornée (utilisé en pratique dans des applications financières [BF19]) a le même taux de convergence que les estimateurs de MMD avec noyaux bornés.

Deux autres mesures de divergence sont considérées dans ce travail: la *discrépance de Stein avec noyau* (Kernel Stein Discrepancy (KSD)) et la distance de Wasserstein. Elles s’écrivent de façon explicite en fonction de la dérivée du log de la fonction de distribution cible pour la KSD, et en fonction de l’inverse de la fonction de répartition cible pour la distance de Wasserstein.

Intuition et comportement empirique de ces portefeuilles. Les mesures de divergences choisies satisfont $D(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T) = 0$ si $\mathbb{P}_{\mathbf{w}} \stackrel{(d)}{=} \mathbb{P}_T$ et sont strictement positives sinon. En pratique, il n'existe généralement pas de \mathbf{w} tel que les distributions coïncident exactement. La distribution des rendements initiaux R présentant des queues plus épaisses que la distribution cible par exemple. On a regardé empiriquement le cas où certaines composantes ont des queues plus épaisses que les autres. Intuitivement, on s'attendrait à ce que l'approche minimum divergence réduise systématiquement le poids du portefeuille sur cette composante. Or si la variance de cette composante est très élevée, ça n'est pas forcément le cas.

Optimisation sur données financières. Nous avons testé notre approche sur des données financières (Librairie Kenneth French [Fre21]). Pour minimiser nos mesures de divergence, nous avons eu recours à la méthode d'entropie croisée [RK04]. Cette méthode consiste à générer aléatoirement des allocations \mathbf{w} dans l'ensemble \mathcal{W} et à se concentrer sur les régions de \mathcal{W} qui minimisent la fonction objective au fur et à mesure des itérations, par simple évaluation de la fonction à minimiser à l'ordre 0 (pas de calculs de dérivés). Cette méthode est notamment adaptée aux problèmes non convexes. Nos résultats empiriques indiquent qu'une distribution cible à queues moins épaisses et plus positivement asymétriques a un impact sur les moments de plus grand ordre du portefeuille, et que la MMD offre de bonnes performances sur nos données de test.




Figure 18 – Distribution des rendements de portefeuilles basés sur des approches traditionnelles (variance minimum, Sharpe ratio maximum) versus approche MMD minimum avec cible skew Gaussienne. La distribution du portefeuille divergence minimum est légèrement moins asymétrique et leptokurtique que les portefeuilles traditionnels.

En pratique, l'approche donne de meilleurs résultats en période de conditions économiques normales. Une approche qui peut être envisagée est de considérer la fonction divergence entre portefeuille empirique et cible pourrais comme un terme de régularisation ou sous la forme d'une contrainte d'approches plus traditionnelles.

Part I

Concentration inequalities for heavy-tailed random variables

CHAPTER 1

CONCENTRATION INEQUALITIES FOR α -EXPONENTIAL RANDOM FOURIER FEATURES (RFFS) ON DERIVATIVES

Note. This chapter corresponds to an article written with E. Gobet and Z. Szabó. It has been published in *Journal of Machine Learning Research (JMLR)*, 21:1–37, 2020, under the title “Orlicz Random Fourier Features”.

Abstract. Kernel techniques are among the most widely-applied and influential tools in machine learning with applications at virtually all areas of the field. To combine this expressive power with computational efficiency numerous randomized schemes have been proposed in the literature, among which probably random Fourier features (RFF) are the simplest and most popular. While RFFs were originally designed for the approximation of kernel values, recently they have been adapted to *kernel derivatives*, and hence to the solution of large-scale tasks involving function derivatives. Unfortunately, the understanding of the RFF scheme for the approximation of higher-order kernel derivatives is quite limited due to the challenging polynomial growing nature of the underlying function class in the empirical process. To tackle this difficulty, we establish a finite-sample deviation bound for a general class of polynomial-growth functions under α -exponential Orlicz condition on the distribution of the sample. Instantiating this result for RFFs, our finite-sample uniform guarantee implies a.s. convergence with tight rate for arbitrary kernel with α -exponential Orlicz spectrum and any order of derivative.

1.1 Introduction

Kernel machines [TC04, SC08, PR16] form one of the most fundamental tools in machine learning and statistics with a wide range of successful applications. The impressive modelling power and flexibility of kernel techniques in capturing complex nonlinear relations originates from the richness of the underlying \mathcal{H}_k function class called reproducing kernel Hilbert space (RKHS, [Aro50]) associated to a $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ kernel. Kernels extend the classical notion of inner product on $\mathcal{X} = \mathbb{R}^d$ by assuming the existence of a $\phi : \mathcal{X} \rightarrow \mathcal{H}$ feature map to a Hilbert space \mathcal{H} such that $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}}$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. This simple equality (also called the kernel trick) forms the basis of kernel techniques and enables one to compute inner products implicitly without direct access to the feature of the points.

In applications one is often given $\{\mathbf{x}_n\}_{n=1}^N$ samples and is facing with an optimization problem expressed in terms of function values and derivatives¹

$$\min_{f \in \mathcal{H}_k} l \left(\left\{ \partial^{\mathbf{p}} f(\mathbf{x}_n) \right\}_{n \in [N], \mathbf{p} \in D_n}, \|f\|_{\mathcal{H}_k}^2 \right), \quad (1.1.1)$$

where $[N] = \{1, \dots, N\}$, $D_n \subset \mathbb{N}^d$, $\mathbb{N} := \{0, 1, \dots\}$, $l : \mathbb{R}^{1 + \sum_{n \in [N]} \#D_n} \rightarrow \mathbb{R}$ is a loss function, $\#D_n$ is the cardinality of the set D_n , $\partial^{\mathbf{p}} f(\mathbf{x}_n) := \frac{\partial^{p_1 + \dots + p_d} f(\mathbf{x}_n)}{\partial x_1^{p_1} \dots \partial x_d^{p_d}}$ and the RKHS \mathcal{H}_k is characterized by $f(\mathbf{x}) = \langle f, k(\cdot, \mathbf{x}) \rangle_{\mathcal{H}_k}$ ($\forall \mathbf{x} \in \mathcal{X}, \forall f \in \mathcal{H}_k$) and $k(\cdot, \mathbf{x}) \in \mathcal{H}_k$ ($\forall \mathbf{x} \in \mathcal{X}$).² The first property of RKHSs is called the reproducing property, the second one describes basic elements of \mathcal{H}_k ; combining the two properties makes the canonical feature map and feature space explicit: $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}_k}$ where $\phi(\mathbf{x}) = k(\cdot, \mathbf{x}) \in \mathcal{H}_k$.

For example by taking the quadratic loss, Tikhonov regularization, only function values ($D_n = \{\mathbf{0}\}, \forall n \in [N]$) and $\lambda > 0$, (1.1.1) reduces to kernel ridge regression

$$\min_{f \in \mathcal{H}_k} \frac{1}{N} \sum_{n \in [N]} [f(\mathbf{x}_n) - y_n]^2 + \lambda \|f\|_{\mathcal{H}_k}^2.$$

Alternatively, one can get back Hermite learning with gradient data [Zho08, SGZ10] by additionally including first-order derivatives

$$\min_{f \in \mathcal{H}_k} \frac{1}{N} \sum_{n \in [N]} \left([f(\mathbf{x}_n) - y_n]^2 + \|f'(\mathbf{x}_n) - \mathbf{y}'_n\|_2^2 \right) + \lambda \|f\|_{\mathcal{H}_k}^2, \quad \lambda > 0$$

where $f'(\mathbf{x}) = [\partial^{\mathbf{e}_1} f(\mathbf{x}); \dots; \partial^{\mathbf{e}_d} f(\mathbf{x})] \in \mathbb{R}^d$ is the derivative of f , $\mathbf{e}_j \in \mathbb{R}^d$ is the j^{th} canonical basis vector, $\|\cdot\|_2$ is the Euclidean norm and $D_n = \left\{ \mathbf{0}, \{\mathbf{e}_j\}_{j=1}^d \right\}$ ($n \in [N]$). Further examples with function derivatives are semi-supervised learning with gradient information [Zho08], nonlinear variable selection [RSM⁺10, RVM⁺13], learning of piecewise-smooth functions [LLB12], multi-task gradient learning [YWC12], structure optimization in parameter-varying ARX (autoregressive with exogenous input) processes [DTPL14], or density estimation with infinite-dimensional exponential families [SFG⁺17].

An appealing property of RKHSs is that their geometry makes the optimization problem (1.1.1) defined over function spaces computationally tractable. Indeed, assuming that l is increasing in its last argument, the $\partial^{\mathbf{p}} f(\mathbf{x}) = \langle f, \partial^{\mathbf{p}, \mathbf{0}} k(\cdot, \mathbf{x}) \rangle_{\mathcal{H}_k}$ derivative-reproducing property of kernels and the representer theorem [Zho08] guarantee that the solution of (1.1.1) has a finite-dimensional parameterization $f(\cdot) = \sum_{n \in [N]} \sum_{\mathbf{p} \in D_n} a_{n, \mathbf{p}} \partial^{\mathbf{p}, \mathbf{0}} k(\cdot, \mathbf{x}_n)$ ($a_{n, \mathbf{p}} \in \mathbb{R}$) and it is sufficient to solve

$$\min_{\mathbf{a}} l \left(\left\{ \sum_{m \in [N]} \sum_{\mathbf{q} \in D_m} a_{m, \mathbf{q}} \partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}_n, \mathbf{x}_m) \right\}_{n \in [N], \mathbf{p} \in D_n}, \sum_{\substack{n, m \in [N] \\ \mathbf{p} \in D_n, \mathbf{q} \in D_m}} a_{n, \mathbf{p}} a_{m, \mathbf{q}} \partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}_n, \mathbf{x}_m) \right) \quad (1.1.2)$$

determined by the $\partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}, \mathbf{y}) := \frac{\partial^{\sum_{i=1}^d (p_i + q_i)} k(\mathbf{x}, \mathbf{y})}{\partial x_1^{p_1} \dots \partial x_d^{p_d} \partial y_1^{q_1} \dots \partial y_d^{q_d}}$ kernel derivatives; $\mathbf{a} = (a_{n, \mathbf{p}})_{n \in [N], \mathbf{p} \in D_n} \in \mathbb{R}^{\sum_{n \in [N]} \#D_n}$.

Though kernel methods show impressive modelling power at numerous areas, due to the implicit computation of feature similarities, this flexibility comes with a computational price. Several techniques have been developed in the literature to mitigate this computational challenge

¹To have derivatives, in the sequel we assume that $\mathcal{X} = \mathbb{R}^d$.

²We use the $k(\cdot, \mathbf{x})$ shorthand to denote the function $\mathbf{y} \in \mathcal{X} \mapsto k(\mathbf{y}, \mathbf{x}) \in \mathbb{R}$ while keeping $\mathbf{x} \in \mathcal{X}$ fixed.

such as incomplete Cholesky factorization [BJ02], sub-sampling schemes [WS01, DM05, RCR17], sketching [AM15, YPW17], random Fourier features (RFF, [RR07, RR08]), their quasi-Monte Carlo [YSAM14], memory-efficient [LSS13, DXH⁺14, ZMDR19], orthogonal [YSC⁺16] or structured [BCC⁺17] variants.

In this paper we study the RFF technique which is probably the conceptually simplest and most influential approach.³ By the Bochner theorem [Rud90] a continuous, bounded, shift-invariant kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ can be written as the Fourier transform of a (finite) measure Λ , called the spectral measure

$$k(\mathbf{x}, \mathbf{y}) = \int_{\mathbb{R}^d} \cos(\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{y})) \, d\Lambda(\boldsymbol{\omega}). \quad (1.1.3)$$

The RFF method uses this representation of k to provide an explicit low-dimensional feature map approximation for the kernel values and f

$$k(\mathbf{x}, \mathbf{x}') \approx \langle \lambda(\mathbf{x}), \lambda(\mathbf{x}') \rangle_{\mathbb{R}^{2M}}, \quad \hat{f}_{\mathbf{w}}(\mathbf{x}) = \langle \mathbf{w}, \lambda(\mathbf{x}) \rangle_{\mathbb{R}^{2M}}, \quad (1.1.4)$$

where the integral representation (1.1.3) with respect to the measure Λ is replaced by an average over random points; hence the *random Fourier feature* naming. As a result, one can estimate \mathbf{w} by leveraging fast linear primal solvers. The idea has been successfully used in various contexts including differential privacy preserving [CMS11], fast function-to-function regression [ONP⁺15], learning message operators in expectation propagation [JGH⁺15], causal discovery [LPMST15, SZV19], independence testing [ZFGS17], prediction and filtering in dynamical systems [DHB⁺17], bandit optimization [LJD⁺18], or estimation of Gaussian mixture models [KBGP18].

Similarly to (1.1.4), one can consider RFF-based approximation of kernel derivatives when solving optimization tasks involving function derivatives [see (1.1.1) and (1.1.2)]. This is the strategy followed for example by [SSL⁺15] to fit distributions belonging to the infinite-dimensional exponential family, which boils down to an optimization problem with third-order kernel derivatives [SFG⁺17, Theorem 5].

The focus of this chapter is to study the approximation quality of the RFF-based kernel-derivative approximation

$$\left\| \widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k \right\|_S := \sup_{\mathbf{x}, \mathbf{y} \in S} \left| \partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}, \mathbf{y}) - \widehat{\partial^{\mathbf{p}, \mathbf{q}} k}(\mathbf{x}, \mathbf{y}) \right|,$$

where $S \subset \mathbb{R}^d$ is a compact set. Despite the large number of successful RFF applications, quite little is understood theoretically on its approximation quality. Below we provide a brief summary with particular focus on optimal guarantees and results related to kernel derivatives.

- **Kernel values** ($\mathbf{p} = \mathbf{q} = \mathbf{0}$): The uniform finite-sample bounds [RR07, SS15b] have recently been improved [SS15a] exponentially in terms of the diameter of the compact set S_M ($|S_M|$) arriving to⁴ $\|k - \hat{k}\|_{S_M} = \mathcal{O}_{a.s.} \left(\frac{\sqrt{\log |S_M|} \vee \sqrt{\log M}}{\sqrt{M}} \right)$ from $\|k - \hat{k}\|_{S_M} = \mathcal{O}_p \left(\frac{|S_M| \sqrt{\log M}}{\sqrt{M}} \right)$, where \vee denotes the maximum. The result shows that the diameter of the set S_M can grow at a $|S_M| = e^{o(M)}$ rate while still getting a consistent estimate; this rate is optimal as shown in the characteristic function literature [CT83].
- **Kernel ridge regression**: RFFs have been settled in kernel ridge regression by [RR17] via showing that using $M = o(N) = \mathcal{O}(\sqrt{N} \log N)$ random Fourier features is sufficient to get $\mathcal{O}(1/\sqrt{N})$ generalization error. Under additional γ -capacity ($\gamma \in [0, 1]$) and r -range space

³[RR07] won the 10-year test-of-time award at NIPS-2017 as a recognition of the influence of RFFs.

⁴The classical $\mathcal{O}(\cdot)$ notation up to logarithmic factors is denoted by $\tilde{\mathcal{O}}(\cdot)$; the extension of $\mathcal{O}(\cdot)$ in almost sure and convergence in probability sense are $\mathcal{O}_p(\cdot)$ and $\mathcal{O}_{a.s.}(\cdot)$.

conditions ($r \geq \frac{1}{2}$), the same authors showed that even faster, minimax optimal $\mathcal{O}\left(N^{-\frac{2r}{2r+\gamma}}\right)$ rates are achievable with $M = o(N) = \mathcal{O}\left(N^{\frac{1+\gamma(2r-1)}{2r+\gamma}} \log N\right)$ RFFs. The result improves the originally proved guarantee [RR08] holding under the pessimistic $M = \mathcal{O}(N)$ setting. The sufficiency of similar sub-linear number of RFFs with minimax guarantees was also established in the kernel least squares setting when applying mini-batched stochastic gradient descent [CRR18]. Recently the analysis of [RR17] has been further sharpened (in terms of the number of required RFFs, [LTOS19]) by leveraging the notion of effective degrees of freedom.

- **Classification with 0-1 loss:** In the classification setting with the 0-1 loss and RKHSs, [GTS18] proved that $M = o(N) = \tilde{\mathcal{O}}\left(N^{\frac{2}{2+c}}\right)$ optimized RFF features—optimized in the sense of [Bac17]—are sufficient to achieve a learning rate of $\tilde{\mathcal{O}}\left(N^{-\frac{c}{2+c}}\right)$ provided that the spectrum of the integral operator associated to the kernel decay polynomially at the rate of $\lambda_i = \mathcal{O}(i^{-c})$ with $c > 1$.⁴ The same authors showed that the learning rate can be improved to $\tilde{\mathcal{O}}(N^{-1})$ with $M = \tilde{\mathcal{O}}(\ln^d(N))$ RFF-s in case of sub-exponential spectrum, where d denotes the dimension of the inputs in the classification.
- **Kernel PCA:** [SS18] have proved that the statistical performance of kernel principal component analysis (KPCA) can be matched by $M = \mathcal{O}(N^{2/3})$ (polynomial decay) or $M = \mathcal{O}(\sqrt{N})$ (exponential decay) RFFs, depending on the eigenvalue decay of the covariance operator associated to the kernel. [UMMA18] derived a similar bound for a streaming KPCA algorithm under exponential spectrum decay condition.
- **Kernel derivatives:** If the support of the spectral measure associated to k is either bounded or it satisfies the Bernstein condition

$$\int_{\mathbb{R}^d} \frac{\left|\prod_{j=1}^d \omega_j^{p_j+q_j}\right|^n}{(\sigma_{\mathbf{p},\mathbf{q}})^n} d\Lambda(\boldsymbol{\omega}) \leq \frac{n!}{2} K^{n-2}, \quad n = 2, 3, \dots \quad (1.1.5)$$

with some constant $K \geq 1$ and $\sigma_{\mathbf{p},\mathbf{q}} = \sqrt{\int_{\mathbb{R}^d} \left|\prod_{j=1}^d \omega_j^{p_j+q_j}\right|^2 d\Lambda(\boldsymbol{\omega})}$, then

$$\left\| \widehat{\partial^{\mathbf{p},\mathbf{q}} k} - \partial^{\mathbf{p},\mathbf{q}} k \right\|_{S_M} = \mathcal{O}_{a.s.} \left(\frac{\sqrt{\log |S_M|} \vee \sqrt{\log M}}{\sqrt{M}} \right)$$

rate is achievable as shown by [SS15a] and [SS19], respectively. As a practical example, one can consider for instance the previously mentioned infinite-dimensional exponential (IE) family fitting problem where the estimation boils to solving a linear equation system with coefficients made of third-order kernel derivatives ($|\mathbf{p} + \mathbf{q}| = 3$). The IE family is defined by a kernel for which a common choice is the Gaussian; this implies a Gaussian spectrum Λ . Unfortunately, in this case the bounded support condition is not met. Similarly, the Bernstein condition (1.1.5) only holds for $|\mathbf{p} + \mathbf{q}| \leq 2$ with the Gaussian kernel, as it can be verified [SS19, Remark 'Higher-order derivatives' in Section 3] by making use of the analytical expressions for the absolute moments of the Gaussian spectrum. These limitations (summarized in Table 1.1) of the popular random Fourier features technique motivate our work and the study of widely-applied kernels with unbounded spectral support for the RFF approximation of high-order kernel derivatives. A consequence of our new estimates in Theorem 1.1 is that the *a.s.* rates previously obtained under stringent conditions (on \mathbf{p}, \mathbf{q} or Λ) are now available for any \mathbf{p}, \mathbf{q} and any spectral measure Λ with α -exponential moments (as defined in (1.1.6), $\alpha > 0$). Because Bernstein condition implies exponential moments, our result includes the one given by [SS19].

Particularly, assuming additional smoothness on the bounded shift-invariant kernel, its

Assumption on the spectral measure	Conditions on \mathbf{p}, \mathbf{q}	Convergence rate for $\left\ \widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k \right\ _{S_M}$
2nd moment exists Ref: [SS15a, Th. 1]	$\mathbf{p} = \mathbf{q} = \mathbf{0}$	$\mathcal{O}_{a.s.} \left(\frac{\sqrt{\log S_M } \sqrt{\log M}}{\sqrt{M}} \right)$
bounded support Ref: [SS15a, Th. 4]	any \mathbf{p}, \mathbf{q}	$\mathcal{O}_{a.s.} \left(\frac{\sqrt{\log S_M } \sqrt{\log M}}{\sqrt{M}} \right)$
Bernstein condition Ref: [SS19]	small \mathbf{p}, \mathbf{q}	$\mathcal{O}_{a.s.} \left(\frac{\sqrt{\log S_M } \sqrt{\log M}}{\sqrt{M}} \right)$
Orlicz condition Ref: now	any \mathbf{p}, \mathbf{q}	$\mathcal{O}_{a.s.} \left(\frac{\sqrt{\log S_M } \sqrt{\log M}}{\sqrt{M}} \right)$

Table 1.1 – Summary of RFF guarantees on kernel values and derivatives. Last line: it includes any measure Λ with a finite α -exponential moment (for some $\alpha, c > 0$, $\mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} (e^{c \|\boldsymbol{\omega}\|_2^\alpha}) < +\infty$), like the Gaussian, the inverse multiquadratic, or the sech kernel, see Corollary 1.3.2. For further examples see Table 1.2.

derivative satisfies a representation similar to (1.1.3):

$$\partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}, \mathbf{y}) = \int_{\mathbb{R}^d} \underbrace{\left[\prod_{j=1}^d \omega_j^{p_j} (-\omega_j)^{q_j} \right] c_{(\sum_{i=1}^d |p_i + q_i|)}(\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{y}))}_{=: f_{\mathbf{x}-\mathbf{y}}(\boldsymbol{\omega})} d\Lambda(\boldsymbol{\omega}),$$

where c_n is the n^{th} derivative of the $\cos(\cdot)$ function. The primary difficulty is to handle the polynomial growing nature of the

$$\mathcal{F} = \{\boldsymbol{\omega} \mapsto f_{\mathbf{x}-\mathbf{y}}(\boldsymbol{\omega}) : \mathbf{x}, \mathbf{y} \in S\}$$

function class which controls the error $\|\widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k\|_S$. We tackle this challenge by imposing the finiteness of the α -exponential Orlicz norm of the spectral measure (Λ) associated to the kernel, in other words

$$\exists \alpha > 0, c > 0 \quad \text{such that} \quad \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} \left(e^{c \|\boldsymbol{\omega}\|_2^\alpha} \right) < +\infty. \quad (1.1.6)$$

Kernels with α -exponential Orlicz spectrum include the popular Gaussian, the inverse multiquadratic, or the sech kernel; for further examples see Table 1.2 and Remark 2(ii). We establish the consistency and prove finite-sample uniform guarantees of the resulting Orlicz RFF scheme for the approximation of kernel derivatives at any order, as it is briefly illustrated in the last line of Table 1.1.

To allow this level of generality, we prove a new finite-sample deviation bound for the empirical process related to a general class of functions f with polynomial growth of the sample \mathbf{X}_m . The distribution of the latter is assumed to have finite α -exponential Orlicz norm and consequently, the random variables $f(\mathbf{X}_m)$ belong to a γ -exponential Orlicz space with index γ smaller than 1. For deriving such deviation bounds, we have been inspired by the work of [Ada08] which elegantly combines the [KR05] inequality for truncated variables, the Hoffman-Jorgensen inequality to deal with sum of residual of truncated variables, and a [Tal89] inequality

in γ -exponential Orlicz norms for sum of centered random variables. However, our work significantly differs from that of [Ada08]. First, our aims are different: [Ada08] focuses on getting large deviation bounds while we are looking for all-scale deviation bounds, which leads to a different analysis (in the application of Klein-Rio inequalities for instance). Second, we are concerned by getting upper bounds with quite explicit control. In particular, this requires a careful treatment of Orlicz-type estimates since the function $\Psi_\gamma(x) = e^{x^\gamma} - 1$ defining the Orlicz space is not convex for $\gamma < 1$ (see Figure 1.1), as opposed to the usual case; see the results in Section 1.4. We also derive sharp estimates from the Dudley entropy integral bound (Theorem 1.4), which enables us to get a tight dependency w.r.t. the diameter of the parameter space. Furthermore, we clarify the use of the Talagrand inequality (Theorem 1.2); in [Ada08, Theorem 5] it is seemingly invoked for supremum over functions while it is related to sum over centered random variables. With this novel finite-sample deviation bound, the analysis of Orlicz RFFs readily follows, using optimized inequalities.

The paper is structured as follows. Our problem is formulated in Section 1.2. The main result on the approximation quality of kernel derivatives with random Fourier features is presented in Section 1.3. Properties of the Orlicz norm are summarized in Section 1.4. Proofs are provided in Section 1.5. The appendix contains additional technical details (Section 1.A), the definition of special functions (Section 1.B) and external statements used in the proofs (Section 1.C).


1.2 Problem formulation

In this section we formally define our problem after introducing a few notations.

Notations: Let the set of natural, real and complex numbers, positive integers, positive reals, non-negative reals and non-positive integers be denoted by $\mathbb{N} = \{0, 1, \dots\}$, \mathbb{R} , \mathbb{C} , $\mathbb{Z}^+ = \{1, 2, \dots\}$, $\mathbb{R}^+ = (0, \infty)$, $\mathbb{R}^{\geq 0} = [0, \infty)$ and $\mathbb{Z}^{\leq 0} = \{0, -1, -2, \dots\}$, respectively. The positive value of $x \in \mathbb{R}$ is denoted by $(x)_+ = x \vee 0$. The Gamma function is $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} dx$ for $t \in \mathbb{C} \setminus \mathbb{Z}^{\leq 0}$. For $x \in \mathbb{R}$ the secant function is $\operatorname{sech}(x) = \frac{1}{\cosh(x)}$. Let $aS + b = \{as + b : s \in S\}$ where $S \subset \mathbb{R}$ and $a, b \in \mathbb{R}$. For an $N \in \mathbb{Z}^+$, $[N] = \{1, \dots, N\}$. Let the n^{th} derivative of the $\cos(\cdot)$ function ($n \in \mathbb{N}$) be $c_n = \cos^{(n)}$. For multi-indices $\mathbf{p}, \mathbf{q} \in \mathbb{N}^d$ and $\boldsymbol{\omega} \in \mathbb{R}^d$ let $|\mathbf{p}| = \sum_{j=1}^d p_j$, $\boldsymbol{\omega}^{\mathbf{p}} = \prod_{j=1}^d \omega_j^{p_j}$, $\partial^{\mathbf{p}} f(\mathbf{x}) = \frac{\partial^{|\mathbf{p}|} f(\mathbf{x})}{\partial x_1^{p_1} \dots \partial x_d^{p_d}}$, $\partial^{\mathbf{p}, \mathbf{q}} g(\mathbf{x}, \mathbf{y}) = \frac{\partial^{|\mathbf{p}|+|\mathbf{q}|} g(\mathbf{x}, \mathbf{y})}{\partial x_1^{p_1} \dots \partial x_d^{p_d} \partial y_1^{q_1} \dots \partial y_d^{q_d}}$. The diameter of a compact set $T \subset \mathbb{R}^d$ is denoted by $|T| = \sup_{\mathbf{x}, \mathbf{y} \in T} \|\mathbf{x} - \mathbf{y}\|_2 < \infty$. Let $S \subset \mathbb{R}^d$ be a Borel set. Let $S_\Delta = S - S = \{a - b : a \in S, b \in S\}$. The set of Borel probability measures on S is written as $\mathcal{M}_1^+(S)$. Let the Banach space of real-valued, r -power μ -integrable functions on S ($1 \leq r < \infty$) be $L^r(S, \mu)$, with $\|f\|_{L^r(S, \mu)} = [\int_S |f(x)|^r d\mu(x)]^{\frac{1}{r}}$. We use the shorthand $\mu f = \int_S f(x) d\mu(x)$ where $\mu \in \mathcal{M}_1^+(S)$ and $f \in L^1(S, \mu)$. The product measure of $\mu_1, \dots, \mu_M \in \mathcal{M}_1^+(S)$ is $\otimes_{m=1}^M \mu_m$; specifically when all the components coincide we use the shorthand $\mu^M = \otimes_{m \in [M]} \mu$. The empirical measure is $\mathbb{P}_M = \frac{1}{M} \sum_{m=1}^M \delta_{X_m}$ with δ_X being the Dirac measure concentrated on X and $X_1, \dots, X_M \sim \otimes_{m=1}^M \mu_m$. Let $(r_n)_{n \in \mathbb{N}}$ be a positive sequence. The boundedness of $\frac{X_n}{r_n}$ almost surely is denoted by $X_n = \mathcal{O}_{a.s.}(r_n)$. Let $n \in \mathbb{R}^+$. We say that an $f : \mathbb{R}^d \rightarrow \mathbb{R}$ function is of polynomial growth of order n (shortly $f \in \mathcal{F}_{\mathcal{P}(n)}$) if $\sup_{\mathbf{x} \in \mathbb{R}^d} \frac{|f(\mathbf{x})|}{1 + \|\mathbf{x}\|_2^n} < \infty$; $\mathcal{F}_{\mathcal{P}} = \cup_{n \in \mathbb{R}^+} \mathcal{F}_{\mathcal{P}(n)}$. Let us assume that $\Psi : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$ is a continuous, strictly increasing mapping, $\Psi(0) = 0$ and $\lim_{x \rightarrow \infty} \Psi(x) = \infty$. The set of \mathbb{R}^d -valued random variables having finite Ψ -Orlicz norm is defined as $L_\Psi = \left\{ \mathbf{X} : \|\mathbf{X}\|_\Psi := \inf \left\{ c > 0 : \mathbb{E} \Psi \left(\frac{\|\mathbf{X}\|_2}{c} \right) \leq 1 \right\} < +\infty \right\}$. Throughout the paper we will be particularly interested in (see Figure 1.1)

$$\Psi_\alpha : x \in \mathbb{R}^{\geq 0} \mapsto e^{x^\alpha} - 1 \in \mathbb{R}^{\geq 0} \quad (\alpha > 0),$$

in other words in random variables having finite α -exponential Orlicz norm. The fact $\mathbf{X} \in L_{\Psi_\alpha}$ is equivalent to the existence of an $s > 0$ constant such that $\mathbb{E} \left[e^{s \|\mathbf{X}\|_2^\alpha} \right] < \infty$. Random variables

Figure 1.1 – Ψ_α for different α values.

$\mathbf{X} \in L_{\Psi_2}$ and $\mathbf{X} \in L_{\Psi_1}$ are called sub-Gaussian and sub-exponential, respectively. For $f \in \mathcal{F}_{\mathcal{P}}$ and random variable \mathbf{X} having α -exponential moment ($\mathbf{X} \in L_{\Psi_\alpha}$) $\mathbb{E}f(\mathbf{X}) < \infty$. Special functions are defined in Table 1.4.

We proceed by formally defining our task. Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous, bounded and shift-invariant kernel. Then, by the Bochner theorem [Rud90] one can assume w.l.o.g. the existence of a $\Lambda \in \mathcal{M}_1^+(\mathbb{R}^d)$ spectral measure such that

$$\begin{aligned} k(\mathbf{x}, \mathbf{y}) &= \int_{\mathbb{R}^d} \cos(\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{y})) \, d\Lambda(\boldsymbol{\omega}) \\ &= \int_{\mathbb{R}^d} \cos(\boldsymbol{\omega}^\top \mathbf{x}) \cos(\boldsymbol{\omega}^\top \mathbf{y}) + \sin(\boldsymbol{\omega}^\top \mathbf{x}) \sin(\boldsymbol{\omega}^\top \mathbf{y}) \, d\Lambda(\boldsymbol{\omega}). \end{aligned}$$

Let $\mathbf{p}, \mathbf{q} \in \mathbb{N}^d$ and assume that $\int_{\mathbb{R}^d} |\boldsymbol{\omega}^{\mathbf{p}+\mathbf{q}}| \, d\Lambda(\boldsymbol{\omega}) < \infty$. In this case $\partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}, \mathbf{y})$ exists, and by the dominated convergence theorem one arrives at

$$\partial^{\mathbf{p}, \mathbf{q}} k(\mathbf{x}, \mathbf{y}) = \int_{\mathbb{R}^d} \partial^{\mathbf{p}} \cos(\boldsymbol{\omega}^\top \mathbf{x}) \partial^{\mathbf{q}} \cos(\boldsymbol{\omega}^\top \mathbf{y}) + \partial^{\mathbf{p}} \sin(\boldsymbol{\omega}^\top \mathbf{x}) \partial^{\mathbf{q}} \sin(\boldsymbol{\omega}^\top \mathbf{y}) \, d\Lambda(\boldsymbol{\omega}).$$

The integral can be estimated by Monte-Carlo technique replacing Λ with $\Lambda_M = \frac{1}{M} \sum_{m=1}^M \delta_{\boldsymbol{\omega}_m}$, $(\boldsymbol{\omega}_m)_{m \in [M]} \stackrel{\text{i.i.d.}}{\sim} \Lambda$:

$$\begin{aligned} \widehat{\partial^{\mathbf{p}, \mathbf{q}} k}(\mathbf{x}, \mathbf{y}) &= \frac{1}{M} \sum_{m=1}^M \left[\partial^{\mathbf{p}} \cos(\boldsymbol{\omega}_m^\top \mathbf{x}) \partial^{\mathbf{q}} \cos(\boldsymbol{\omega}_m^\top \mathbf{y}) + \partial^{\mathbf{p}} \sin(\boldsymbol{\omega}_m^\top \mathbf{x}) \partial^{\mathbf{q}} \sin(\boldsymbol{\omega}_m^\top \mathbf{y}) \right] \\ &= \langle \lambda_{\mathbf{p}}(\mathbf{x}), \lambda_{\mathbf{q}}(\mathbf{y}) \rangle_{\mathbb{R}^{2M}}, \end{aligned} \tag{1.2.1}$$

where $\lambda_{\mathbf{p}}(\mathbf{x}) = \frac{1}{\sqrt{M}} \left[(\partial^{\mathbf{p}} \cos(\boldsymbol{\omega}_m^\top \mathbf{x}))_{m \in [M]}; (\partial^{\mathbf{p}} \sin(\boldsymbol{\omega}_m^\top \mathbf{x}))_{m \in [M]} \right] \in \mathbb{R}^{2M}$; this is the RFF feature approximation $\lambda_{\mathbf{p}}$ in (1.1.4). For $\mathbf{p} = \mathbf{q} = \mathbf{0}$, the construction reduces to the traditional RFF technique [RR07].

This form implies that our target quantity can be written as

$$\left\| \widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k \right\|_S = \sup_{\mathbf{z} \in S_\Delta} |(\Lambda_M - \Lambda)(f_{\mathbf{z}})|, \quad f_{\mathbf{z}}(\boldsymbol{\omega}) = \boldsymbol{\omega}^{\mathbf{p}} (-\boldsymbol{\omega})^{\mathbf{q}} c_{|\mathbf{p}+\mathbf{q}|}(\boldsymbol{\omega}^\top \mathbf{z}), \tag{1.2.2}$$

$$= \sup_{f \in \mathcal{F}} |(\Lambda_M - \Lambda)(f)|, \quad \mathcal{F} = \{f_{\mathbf{z}} : \mathbf{z} \in S_\Delta\}, \tag{1.2.3}$$

thus the problem boils down to the study of supremum of empirical processes with $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$ where $n = |\mathbf{p} + \mathbf{q}| + 1$. In the next section we detail our main result about the fluctuation of such processes.

1.3 Main result

In this section we present our main result on the supremum of empirical processes of polynomial growth, and specialize it to the approximation quality of RFFs for kernel derivatives. The proofs are given in Section 1.5.

We investigate the concentration of the $\sup_{f \in \mathcal{F}} |\frac{1}{M} \sum_{m=1}^M f(\mathbf{X}_m)|$ quantity under the following assumptions:

1. **Compact parameterization:** $\mathcal{F} = \{f_{\mathbf{t}} : \mathbf{t} \in T\}$ where $f_{\mathbf{t}} : \mathbb{R}^d \rightarrow \mathbb{R}$ is parameterized by a compact set $T \subset \mathbb{R}^{d'}$.
2. **Lipschitz condition:** There exists $n \in \mathbb{R}^+$ and function $L : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq 0}$, $L \in \mathcal{F}_{\mathcal{P}(n)}$ such that
 - 2.1. $|f_{\mathbf{t}_0}(\mathbf{x})| \leq L(\mathbf{x})$ for some $\mathbf{t}_0 \in T$,
 - 2.2. $|f_{\mathbf{t}_1}(\mathbf{x}) - f_{\mathbf{t}_2}(\mathbf{x})| \leq L(\mathbf{x})\rho(\|\mathbf{t}_1 - \mathbf{t}_2\|_2)$ for all $\mathbf{x} \in \mathbb{R}^d, \mathbf{t}_1 \in T, \mathbf{t}_2 \in T$,
 - 2.3. with $\rho : [0, |T|] \rightarrow \mathbb{R}^{\geq 0}$ continuous strictly increasing mapping with $\rho(0) = 0$ such that

$$I_\rho(|T|) := \rho(|T|) \int_0^1 \sqrt{\log\left(1 + \frac{2|T|}{\rho^{-1}(u\rho(|T|))}\right)} du < \infty.$$
3. **Independence, finite α -exponential Orlicz norm:**
 - 3.1. $(\mathbf{X}_m)_{m \in [M]}$ are independent \mathbb{R}^d -valued random variables; shortly, $(\mathbf{X}_m)_{m \in [M]} \sim \otimes_{m \in [M]} \mu_m$ with $\mu_m \in \mathcal{M}_1^+(\mathbb{R}^d)$.
 - 3.2. $\exists \alpha \in \mathbb{R}^+$ such that $\|\mathbf{X}_m\|_{\Psi_\alpha} < \infty$ for all $m \in [M]$.
4. **Centering:** $\mathbb{E}[f(\mathbf{X}_m)] = 0$ for all $f \in \mathcal{F}$ and $m \in [M]$.

Under these conditions, our main result is as follows.

Theorem 1.1 (Concentration of processes with polynomial growth). *Assume that \mathcal{F} and $(\mathbf{X}_m)_{m \in [M]}$ satisfy Assumptions 1-4 and $\gamma := \frac{\alpha}{n} \leq 1$. Let \log stand for the natural logarithm, $\beta_\gamma := \Gamma\left(1 + \frac{1}{\gamma}\right)^{-\gamma}$, $\mathbb{P} := \otimes_{m \in [M]} \mu_m$, and $\|L\|_{L^2(W_{1:M})} := \sqrt{\frac{1}{M} \sum_{m \in [M]} L^2(\mathbf{X}_m)}$. Let $\Psi_\gamma^{(l)}$ be the convexification⁵ of Ψ_γ , $A_\gamma := \frac{(\Psi_\gamma^{(l)})^{-1}(1)}{\Psi_\gamma^{-1}(1)}$, $B_\gamma := (\Psi_\gamma^{(l)})^{-1}(1)$, C_γ and C_D be the constants defined in (1.C.1) and (1.C.2), and $K_\gamma := 2^{\left(\frac{1}{\gamma}-1\right)} \left(C_\gamma \left[16B_\gamma + 2^{\left(\frac{1}{\gamma}-1\right)} \left(1 + A_\gamma \right) \right] + 8A_\gamma \right)$. Then for any $\varepsilon > 0$ satisfying*

$$\varepsilon \geq 6B, \quad B := 2C_D \sqrt{d'} \frac{\mathbb{E} \left[\|L\|_{L^2(W_{1:M})} \right]}{\sqrt{M}} I_\rho(|T|), \quad (1.3.1)$$

we have

$$\mathbb{P} \left(\sup_{t \in T} \frac{1}{M} \sum_{m \in [M]} f_t(\mathbf{X}_m) \geq \varepsilon \right) \leq 2e^{-\left(\frac{M\varepsilon}{3K_\gamma \|\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|\|_{\Psi_\gamma}} \right)^\gamma} + e^{-\frac{M\varepsilon^2}{72\sigma^2 + 84c\varepsilon}} \quad (1.3.2)$$

where

$$\begin{aligned} \sigma^2 &:= \sup_{t \in T} \frac{1}{M} \sum_{m=1}^M \mathbb{E} [f_t^2(\mathbf{X}_m)], \\ c &:= \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \left[\frac{1}{\beta_\gamma} \log \left(\frac{6\Gamma\left(1 + \frac{1}{\gamma}\right) \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}{\gamma\varepsilon} \right) \right]^{\frac{1}{\gamma}} \\ &\vee 8\mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right] \in [0, +\infty). \end{aligned}$$

⁵The function Ψ_γ is not convex for $\gamma < 1$. We convexify Ψ_γ and use the Section 1.4(v) based integral control property holding for convex Ψ -s; for details on $\Psi_\gamma^{(l)}$ see Section 1.A.1.

Remark 1.

- (i) **Two-sided bound:** For $\mathbb{P}\left(\inf_{t \in T} \frac{1}{M} \sum_{m \in [M]} f_t(\mathbf{X}_m) \leq -\varepsilon\right)$ the same one-sided deviation bound can be obtained by replacing f_t with $-f_t$.
As a result one can estimate $\mathbb{P}\left(\sup_{t \in T} \left| \frac{1}{M} \sum_{m \in [M]} f_t(\mathbf{X}_m) \right| \geq \varepsilon\right)$ by twice the bound above.
- (ii) **Assumption (3):** Assumption (3a) with Assumption (4) is weaker than being i.i.d.: for example $\mathbb{E}[f_t(X_m)] = 0$ holds for $X_m = \mathcal{N}(0, \sigma_m^2)$ and $f_t(x) = c_t x^3$, but X_m -s can differ in their variance.
- (iii) **Assumption $\alpha/n \leq 1$:** This condition holds without loss of generality. Indeed, in case of $\alpha/n > 1$, one can get a modified (α', n') pair satisfying $\alpha'/n' \leq 1$ by either increasing n to the value $n' = \alpha$ using that $\mathcal{F}_{\mathcal{P}(n)} \subset \mathcal{F}_{\mathcal{P}(n')}$, or by decreasing α to the value $\alpha' = n$ using that $\|\mathbf{X}_m\|_{\Psi_\alpha} < \infty$ implies $\|\mathbf{X}_m\|_{\Psi_{\alpha'}} < \infty$ for any $\alpha' \in (0, \alpha)$.

(iv) **Proof-related remarks:**

1. **Compactness of T :** This compactness with the Lipschitz property enables one to control the covering number of \mathcal{F} .
2. **Truncated functions:** The Lipschitz property of \mathcal{F} implies that of the truncated functions: for $\forall \mathbf{x} \in \mathbb{R}^d$, \mathbf{s} and $\mathbf{t} \in T$

$$|\mathcal{T}_c f_t(\mathbf{x}) - \mathcal{T}_c f_s(\mathbf{x})| \leq |f_t(\mathbf{x}) - f_s(\mathbf{x})| \leq L(\mathbf{x})\rho(\|\mathbf{t} - \mathbf{s}\|_2), \quad (1.3.3)$$

where $\mathcal{T}_c f(\mathbf{x}) := f(\mathbf{x})\mathbb{1}_{|f(\mathbf{x})| \leq c} + c\mathbb{1}_{f(\mathbf{x}) > c} - c\mathbb{1}_{f(\mathbf{x}) < -c}$ is f soft-thresholded at level c .

3. $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$: This property is inherited (Section 1.5.4) from $L \in \mathcal{F}_{\mathcal{P}(n)}$ by the Lipschitz conditions (2a)-(2b).
4. **Finiteness of the terms in Theorem 1.1:** $\|\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|\|_{\Psi_{\frac{\alpha}{n}}}$ and $\mathbb{E}[\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|]$ are finite (see Section 1.5.4) in Theorem 1.1 by the Lipschitz assumption (2a)-(2b), $\|\mathbf{X}_m\|_{\Psi_\alpha} < \infty$ (Assumption (3b)) and $L \in \mathcal{F}_{\mathcal{P}(n)}$ (Assumption (2)).

(v) **RFF specialization:** Assuming that the α -exponential Orlicz condition holds for the spectral measure Λ associated to k ($\exists \alpha \in \mathbb{R}^+$ such that $\|\omega\|_{\Psi_\alpha} < \infty$, $\omega \sim \Lambda$),⁶ one can see (Section 1.5.1) that RFFs are covered by choosing

$$\begin{aligned} d' &= d, & f_{\mathbf{t}}(\mathbf{x}) &\leftarrow f_{\mathbf{z}}(\omega) - \Lambda f_{\mathbf{z}}, & \mathbf{t} &\leftarrow \mathbf{z}, & T &\leftarrow S_\Delta, & \mathbf{X}_m &\leftarrow \omega_m, \\ \rho(u) &= u^\beta, & \beta &= \frac{1}{1 + (\log|S_\Delta|)_+} \in (0, 1], & n &\leftarrow |\mathbf{p} + \mathbf{q}| + \beta. \end{aligned}$$

While any value of $\beta \in (0, 1]$ would meet the assumptions, allowing β to depend on the diameter of S_Δ enables us to get optimal convergence rates w.r.t. the diameter (see Corollary 1.3.2).

The terms driving the guarantee for RFF can be bounded (Section 1.5.6) as follows: there is a constant $C_{\text{RFF}} \in \mathbb{R}^+$, depending only on Λ , $|\mathbf{p} + \mathbf{q}|$, but not on $|S_\Delta|$ and M , such that

$$\begin{aligned} B &\leq \frac{C_{\text{RFF}} \sqrt{1 + (\log|S_\Delta|)_+}}{\sqrt{M}}, \\ \sigma^2 &\leq C_{\text{RFF}}, \\ \max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} \|g_{\mathbf{z}}(\omega_m)\|_{\Psi_\gamma} &\leq C_{\text{RFF}}, \\ \left\| \max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\omega_m)| \right\|_{\Psi_\gamma} &\leq C_{\text{RFF}} [\log(1 + M)]^{n/\alpha}, \\ \mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\omega_m)| \right] &\leq C_{\text{RFF}} [\log(1 + M)]^{n/\alpha}. \end{aligned} \quad (1.3.4)$$

⁶This requirement implies that $\int_{\mathbb{R}^d} |\omega^{\mathbf{p} + \mathbf{q}}| d\Lambda(\omega) < \infty$ and thus the existence of $\partial^{\mathbf{p}, \mathbf{q}} k$ for any $\mathbf{p}, \mathbf{q} \in \mathbb{N}^d$.

Using these bounds, our finite-sample uniform guarantee on Orlicz RFFs is as follows.

Corollary 1.3.1 (Orlicz RFFs for kernel derivative approximation). *Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous, bounded, shift-invariant kernel with spectral measure Λ . Suppose that Λ satisfies the α -exponential Orlicz assumption ($\exists \alpha \in \mathbb{R}^+$ such that $\|\omega\|_{\Psi_\alpha} < \infty$, $\omega \sim \Lambda$) and let $S \subset \mathbb{R}^d$ be a compact set. Let $\beta = \frac{1}{1+(\log|S_\Delta|)_+} \in (0, 1]$, let $\mathbf{p}, \mathbf{q} \in \mathbb{N}^d$, $n := |\mathbf{p} + \mathbf{q}| + \beta$, and assume that $\gamma := \frac{\alpha}{n} \leq 1$. Let $\widehat{\partial^{\mathbf{p}, \mathbf{q}} k}$ be the RFF estimate of $\partial^{\mathbf{p}, \mathbf{q}} k$ using $(\omega_m)_{m \in [M]} \stackrel{i.i.d.}{\sim} \Lambda$ samples as given in Eq. (1.2.1). Then, there exists a constant $\tilde{C} \in \mathbb{R}^+$ (depending only on Λ , $|\mathbf{p} + \mathbf{q}|$, but not on S and M) such that for any $\epsilon \geq \frac{\tilde{C} \sqrt{1+(\log|S_\Delta|)_+}}{\sqrt{M}}$,*

$$\Lambda^M \left(\left\| \widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k \right\|_S \geq \epsilon \right) \leq 2e^{-\frac{(M\epsilon)^\gamma}{\tilde{C} \log(1+M)}} + e^{-\frac{M\epsilon^2}{\tilde{C} \left(1 + \epsilon \left[\log(\tilde{C}/\epsilon) \vee \log(1+M)\right]^{1/\gamma}\right)}}. \quad (1.3.5)$$

Corollary 1.3.2 (Almost sure convergence for kernel derivative approximation). *Let $\mathbf{p}, \mathbf{q} \in \mathbb{N}^d$ and $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous, bounded, shift-invariant kernel with spectral measure Λ which satisfies the α -exponential Orlicz assumption for some $\alpha > 0$. Then, for any sequence of compact sets $(S_M)_{M=2}^\infty$ such that $(\log|S_M|)_+ = o(M)$, we have*

$$\left\| \widehat{\partial^{\mathbf{p}, \mathbf{q}} k} - \partial^{\mathbf{p}, \mathbf{q}} k \right\|_{S_M} = \mathcal{O}_{a.s.} \left(\frac{\sqrt{(\log|S_M|)_+ \vee \log M}}{\sqrt{M}} \right). \quad (1.3.6)$$

Remark 2.

- (i) **Spectral measure (Λ) examples:** Our result assumes the α -exponential Orlicz property of the spectral measure Λ associated to k . In Table 1.2 we provide various examples for Λ (with the relevant case of unbounded support) satisfying this requirement; their relation is summarized in Figure 1.2. While for the RFF approximation it is not necessary, in many of these examples the corresponding kernel value can also be computed, see Table 1.3.
- (ii) **α -exponential Orlicz assumption for tensor product kernels:** Using the α -exponential Orlicz spectral measures of Table 1.2 on \mathbb{R} , one can immediately construct Orlicz spectral measures on \mathbb{R}^d . Indeed, assume that (i) k is a product kernel, i.e. $k(\mathbf{x}, \mathbf{y}) = \prod_{i \in [d]} k_i(x_i, y_i)$, $\Lambda = \otimes_{i \in [d]} \Lambda_i$, and (ii) Λ_i , the spectral measure associated to k_i , satisfies the α_i -exponential Orlicz assumption ($\alpha_i \in \mathbb{R}^+$). Then $\omega \sim \Lambda$ is α -exponential Orlicz with $\alpha = \min_{i \in [d]} \alpha_i$; see Section 1.A.2.
- (iii) **α -exponential Orlicz vs. Bernstein assumption:** Our result complements [SS19]'s work, where the authors showed that for $d = 1$ and spectral densities $f_\lambda(\omega) \propto e^{-\omega^{2\ell}}$ the Bernstein condition (and hence fast rates) holds for $|\mathbf{p} + \mathbf{q}| \leq 2\ell = \alpha$. Indeed, we proved under the more general α -exponential Orlicz assumption the same *a.s.* convergence rates for any arbitrary order (see Corollary 1.3.2) kernel derivatives.

Spectrum	Spectral density: $f_\Lambda(\omega)$	Parameters	α
Gaussian	$\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{\omega^2}{2\sigma^2}}$	$\sigma > 0$	2
Laplace	$\frac{\sigma}{2} e^{-\sigma \omega }$	$\sigma > 0$	1
generalized Gaussian	$\frac{\alpha}{2\beta\Gamma(\frac{1}{\alpha})} e^{-\frac{ \omega }{\beta}^\alpha}$	$\alpha > 0, \beta > 0$	α
variance Gamma	$\frac{\sigma^{2b} \omega ^{b-\frac{1}{2}} K_{b-\frac{1}{2}}(\sigma \omega)}{\sqrt{\pi}\Gamma(b)(2\sigma)^{b-\frac{1}{2}}}$	$\sigma > 0, b > \frac{1}{2}$	1

Spectrum	Spectral density: $f_\Lambda(\omega)$	Parameters	α
Weibull (S)	$\frac{s}{2\lambda} \left(\frac{ \omega }{\lambda}\right)^{s-1} e^{-\left(\frac{ \omega }{\lambda}\right)^s}$	$s > 0, \lambda > 0$	s
exponentiated exponential (S)	$\frac{\alpha}{2\lambda} \left(1 - e^{-\frac{ \omega }{\lambda}}\right)^{\alpha-1} e^{-\frac{ \omega }{\lambda}}$	$\lambda > 0, \alpha > 0$	1
exponentiated Weibull (S)	$\frac{\alpha s}{2\lambda} \left(\frac{ \omega }{\lambda}\right)^{s-1} \left[1 - e^{-\left(\frac{ \omega }{\lambda}\right)^s}\right]^{\alpha-1} \times e^{-\left(\frac{ \omega }{\lambda}\right)^s}$	$s > 0, \lambda > 0, \alpha > 0$	s
Nakagami (S)	$\frac{m^m}{\Gamma(m)\Omega^m} \omega ^{2m-1} e^{-\frac{m\omega^2}{\Omega}}$	$m \geq \frac{1}{2}, \Omega > 0$	2
chi-squared (S)	$\frac{1}{2^{\frac{s}{2}+1}\Gamma(\frac{s}{2})} \omega ^{\frac{s}{2}-1} e^{-\frac{ \omega }{2}}$	$s \in \mathbb{Z}^+$	1
Erlang (S)	$\frac{\lambda^s \omega ^{s-1} e^{-\lambda \omega }}{2(s-1)!}$	$s \in \mathbb{Z}^+, \lambda > 0$	1
Gamma (S)	$\frac{1}{2\Gamma(s)\theta^s} \omega ^{s-1} e^{-\frac{ \omega }{\theta}}$	$s > 0, \theta > 0$	1
generalized Gamma (S)	$\frac{p/a^D}{2\Gamma(\frac{D}{p})} \omega ^{D-1} e^{-\left(\frac{ \omega }{a}\right)^p}$	$a > 0, D > 0, p > 0$	p
Rayleigh (S)	$\frac{ \omega }{2\sigma^2} e^{-\frac{\omega^2}{2\sigma^2}}$	$\sigma > 0$	2
Maxwell-Boltzmann (S)	$\frac{1}{\sqrt{2\pi}} \frac{\omega^2 e^{-\frac{\omega^2}{2a^2}}}{a^3}$	$a > 0$	2
chi (S)	$\frac{1}{2^{\frac{s}{2}}\Gamma(\frac{s}{2})} \omega ^{s-1} e^{-\frac{\omega^2}{2}}$	$s > 0$	2
exponential-logarithmic (S)	$-\frac{1}{2\log(p)} \frac{\beta(1-p)e^{-\beta \omega }}{1-(1-p)e^{-\beta \omega }}$	$p \in (0, 1), \beta > 0$	1
Weibull-logarithmic (S)	$-\frac{1}{2\log(p)} \frac{\alpha\beta(1-p) \omega ^{\alpha-1} e^{-\beta \omega ^\alpha}}{1-(1-p)e^{-\beta \omega ^\alpha}}$	$p \in (0, 1), \beta > 0, \alpha > 0$	α
Gamma/Gompertz (S)	$\frac{bse^{b \omega }\beta^s}{2(\beta-1+e^{b \omega })^{s+1}}$	$b > 0, \beta > 0, s > 0$	bs
hyperbolic secant	$\frac{1}{2} \operatorname{sech}\left(\frac{\pi}{2}\omega\right)$		1
logistic	$\frac{e^{-\frac{\omega}{s}}}{s\left[1+e^{-\frac{\omega}{s}}\right]^2}$	$s > 0$	1
normal-inverse Gaussian	$\frac{\alpha\delta K_1(\alpha\sqrt{\delta^2+\omega^2})}{\pi\sqrt{\delta^2+\omega^2}} e^{\delta\alpha}$	$\alpha > 0, \delta \in \mathbb{R}$	1
hyperbolic	$\frac{1}{2\delta K_1(\delta\alpha)} e^{-\alpha\sqrt{\delta^2+\omega^2}}$	$\alpha > 0, \delta \in \mathbb{R}$	1
generalized hyperbolic	$\frac{(\alpha/\delta)^\lambda}{\sqrt{2\pi}K_\lambda(\delta\gamma)} \frac{K_{\lambda-\frac{1}{2}}(\alpha\sqrt{\delta^2+\omega^2})}{\left(\frac{\sqrt{\delta^2+\omega^2}}{\alpha}\right)^{\frac{1}{2}-\lambda}}$	$\alpha > 0, \lambda \in \mathbb{R}, \delta \in \mathbb{R}$	1

Table 1.2 – Kernel spectrum examples in one dimension ($d = 1$) obeying the α -exponential Orlicz assumption. '(S)' stands for symmetrized. The symmetrization guarantees that the kernel associated to Λ is real-valued. Last column: Orlicz exponent. For the variance Gamma distribution the Orlicz exponent follows from the known $K_u(z) \sim \sqrt{\pi/(2z)}e^{-z}$ asymptotics [BNMR01, page 297] where K_u is the modified Bessel function of the second kind, as defined in

Table 1.4. Notice that the 'normal-inverse Gaussian $\xrightarrow{\delta=\sigma^2\alpha, \alpha\rightarrow\infty}$ Gaussian' limit (see Figure 1.2) changed the Orlicz exponent from 1 to 2.

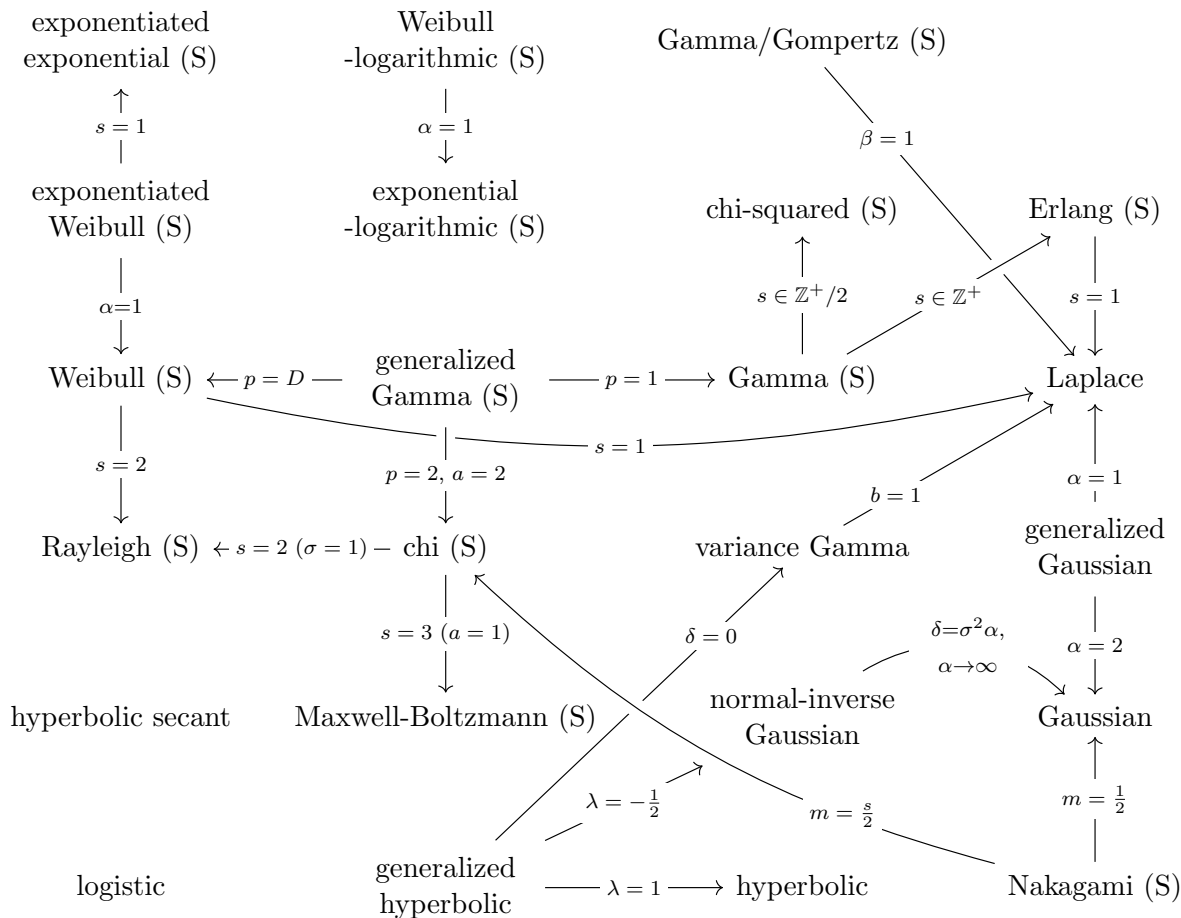


Figure 1.2 – Relation of the spectral density examples of Table 1.2. '(S)' stands for symmetrized.

Kernel name	Kernel value: $k(x, y)$	Spectrum
Gaussian	$e^{-\frac{\sigma^2(x-y)^2}{2}}$	Gaussian
inverse quadric	$\frac{\sigma^2}{\sigma^2+(x-y)^2}$	Laplace
–	$\frac{\sqrt{\pi}}{\Gamma(1/\alpha)} {}_1\Psi_1\left(\left(\frac{1}{\alpha}, \frac{2}{\alpha}\right); \left(\frac{1}{2}; 1\right), \frac{-[\beta(x-y)]^2}{4}\right)$	generalized Gaussian ^a
inverse multiquadric	$\left[\frac{\sigma^2}{\sigma^2+(x-y)^2}\right]^b$	variance Gamma
–	$\sum_{n \in 2\mathbb{N}} \frac{(-1)^{\frac{n}{2}} (x-y)^n \lambda^n}{n!} \Gamma\left(1 + \frac{n}{s}\right)$	Weibull (S) ^b
–	$\frac{[1-2i(x-y)]^{-\frac{s}{2}} + [1+2i(x-y)]^{-\frac{s}{2}}}{2}$	chi-squared (S) ^b
–	$\frac{\left[1 - \frac{i(x-y)}{\lambda}\right]^{-s} + \left[1 + \frac{i(x-y)}{\lambda}\right]^{-s}}{2}$	Erlang (S) ^b
–	$\frac{[1-\theta i(x-y)]^{-s} + [1+\theta i(x-y)]^{-s}}{2}$	Gamma (S) ^b
–	$1 - \sigma(x-y)e^{-\frac{\sigma^2(x-y)^2}{2}} \sqrt{\frac{\pi}{2}} \operatorname{erfi}\left(\frac{\sigma(x-y)}{\sqrt{2}}\right)$	Rayleigh (S) ^b
–	${}_1F_1\left(\frac{s}{2}; \frac{1}{2}; \frac{-(x-y)^2}{2}\right)$	chi (S) ^b
–	$\sum_{n \in 2\mathbb{N}} \frac{(-1)^{\frac{n}{2}} (x-y)^n \Gamma\left(\frac{n}{\alpha} + 1\right)}{-\log(p)n! \beta^{\frac{\alpha}{n}}} \operatorname{Li}_{\frac{n}{\alpha}+1}(1-p)$	Weibull-logarithmic (S) ^{b,c}
–	$\frac{1}{2} [c_\Lambda(x-y) + c_\Lambda(y-x)],$ with $c_\Lambda(t) =$ $= \beta^s \frac{sb}{sb-ti} {}_2F_1\left(s+1; -\frac{ti}{b} + s; -\frac{ti}{b} + s + 1; 1-\beta\right)$	Gamma/Gompertz (S) ^b
sech	$\operatorname{sech}(x-y)$	hyperbolic secant
–	$\frac{\pi s(x-y)}{\sinh(\pi s(x-y))}$	logistic
–	$e^{\delta[\alpha - \sqrt{\alpha^2 + (x-y)^2}]}$	normal-inverse Gaussian
–	$\frac{\alpha K_1(\delta\sqrt{\alpha^2 + (x-y)^2})}{\sqrt{\alpha^2 + (x-y)^2} K_1(\delta\alpha)}$	hyperbolic
–	$\left[\frac{\alpha}{\sqrt{\alpha^2 + (x-y)^2}}\right]^\lambda \frac{K_\lambda(\delta\sqrt{\alpha^2 + (x-y)^2})}{K_\lambda(\delta\alpha)}$	generalized hyperbolic

^aThe analytical computation of the characteristic function (and hence the kernel value) was carried out for $\alpha > 1$ [PN10].

^bIn case of symmetrization (S): $k(x, y) = \frac{1}{2} [c_\Lambda(x-y) + c_\Lambda(y-x)]$ where $c_\Lambda(t) = \mathbb{E}_{\omega \sim \Lambda}[e^{it\omega}]$ is the characteristic function of the spectral measure (on $\mathbb{R}^{\geq 0}$) before symmetrization; $i = \sqrt{-1}$.

^cThe characteristic function was obtained by [CP09].

Table 1.3 – Kernel examples for the spectral densities given in Table 1.2. The special functions ${}_1\Psi_1$, erfi , ${}_1F_1$, Li , ${}_2F_1$ and K_λ are defined in Table 1.4.

1.4 Properties of the Orlicz norm

In this section, for self-containedness we summarize the properties of $\|\cdot\|_\Psi$ which hold independently of the convexity/non-convexity of Ψ (unless explicitly required).

Let $X, X' \in \mathbb{R}^d$ be random variables, and assume that $\Psi : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$ (and similarly Φ below) is continuous, strictly increasing, $\Psi(0) = 0$ and $\lim_{x \rightarrow \infty} \Psi(x) = \infty$.

- (i) Normalization: If $X \in L_\Psi$ then $\mathbb{E} \left[\Psi \left(\frac{\|X\|_2}{\|X\|_\Psi} \right) \right] \leq 1$.
- (ii) Constant: For a $\lambda \in \mathbb{R}$ constant $\|\lambda\|_\Psi = |\lambda|/\Psi^{-1}(1)$.
- (iii) Monotonicity in Ψ : $\Psi \leq \Phi$ implies $\|X\|_\Psi \leq \|X\|_\Phi$.
- (iv) Monotonicity in the argument: If $d = 1$ and $0 \leq X \leq X'$ a.s., then $\|X\|_\Psi \leq \|X'\|_\Psi$.
- (v) Finite $\|\cdot\|_\Psi$ implies integrability: If Ψ is convex and $X \in L_\Psi$, then $\mathbb{E}[\|X\|_2] \leq \|X\|_\Psi \Psi^{-1}(1)$.
- (vi) Generalized triangle inequality: Let $X, X' \in L_{\Psi_\alpha}$ and $\alpha \in \mathbb{R}^+$. Then $X + X' \in L_{\Psi_\alpha}$ and

$$\|X + X'\|_{\Psi_\alpha} \leq 2^{\left(\frac{1}{\alpha}-1\right)_+} \left(\|X\|_{\Psi_\alpha} + \|X'\|_{\Psi_\alpha} \right).$$

- (vii) Deviation inequality from $\|\cdot\|_\Psi$: If $X \in L_\Psi$ then $\mathbb{P}(\|X\|_2 \geq c) \leq \frac{2}{\Psi(c/\|X\|_\Psi)+1}$ for any $c \geq 0$.
- (viii) Maximal inequality for $\|\cdot\|_{\Psi_\alpha}$ and $\alpha \in \mathbb{R}^+$: for any sequence $(X_m)_{m=1}^M$ of random variables in L_{Ψ_α} , we have

$$\left\| \max_{m \in [M]} \|X_m\|_2 \right\|_{\Psi_\alpha} \leq \max_{m \in [M]} \|X_m\|_{\Psi_\alpha} \left[\frac{\log(1+M)}{\log(3/2)} \right]^{1/\alpha}.$$

The proofs of these properties are available in Section 1.A.3.

1.5 Proofs

After introducing a few additional notations, we provide the proofs of our results and remarks presented in Sections 1.3 and 1.4. External statements used in the proofs are summarized in Section 1.C.

Notations: For $\gamma \in (0, 1]$ and $x \in \mathbb{R}^{\geq 0}$, let $I_\gamma(x) = \int_0^x e^{-t^\gamma} dt$ be the incomplete Gamma function. Let (Z, m) be a semi-metric space and $\epsilon \in \mathbb{R}^+$. The set $S \subseteq Z$ is said to be an ϵ -net of Z if for any $z \in Z$ there exists $s \in S$ such that $m(s, z) \leq \epsilon$. The ϵ -covering number of Z is defined as the size of the smallest ϵ -net, i.e., $N(\epsilon, m, Z) = \inf \left\{ \ell \geq 1 : \exists s_1, \dots, s_\ell \in Z \text{ such that } Z \subseteq \cup_{j=1}^\ell B_m(s_j, \epsilon) \right\}$, where $B_m(s, \epsilon) = \{z \in Z : m(z, s) \leq \epsilon\}$ is the closed ball with center $s \in Z$ and radius ϵ .

1.5.1 Proof of Remark 1(v)

In view of (1.2.2)-(1.2.3), we need to check Assumptions 1-4 with the parameterized function class

$$g_{\mathbf{z}}(\boldsymbol{\omega}) := f_{\mathbf{z}}(\boldsymbol{\omega}) - \Lambda f_{\mathbf{z}} = \boldsymbol{\omega}^{\mathbf{p}} (-\boldsymbol{\omega})^{\mathbf{q}} c_{|\mathbf{p}+\mathbf{q}|} \left(\boldsymbol{\omega}^\top \mathbf{z} \right) - \Lambda f_{\mathbf{z}}, \quad (\mathbf{z} \in S_\Delta).$$

Thanks to the α -exponential Orlicz condition on Λ and the i.i.d. property of $(\boldsymbol{\omega}_m)_{m \in [M]}$ in (1.2.1), Assumption 3 is trivially fulfilled. Assumption 4 holds by the definition of $g_{\mathbf{z}}(\cdot)$ and because the distribution of $\boldsymbol{\omega}_m$ is Λ . Assumption 1 is satisfied since S_Δ is a compact set of \mathbb{R}^d . Therefore, it remains to prove Assumption 2, with the existence of $n \in \mathbb{R}^+$ and $L \in \mathcal{F}_{\mathcal{P}(n)}$. First, notice that

$$|f_{\mathbf{z}}(\boldsymbol{\omega})| \leq \prod_{i \in [d]} |\omega_i|^{p_i+q_i} \leq \|\boldsymbol{\omega}\|_2^{|\mathbf{p}+\mathbf{q}|}. \quad (1.5.1)$$

- **Order:** (1.5.1) implies that

$$|g_{\mathbf{z}}(\boldsymbol{\omega})| \leq |f_{\mathbf{z}}(\boldsymbol{\omega})| + \Lambda |f_{\mathbf{z}}| \leq \|\boldsymbol{\omega}\|_2^{|\mathbf{p}+\mathbf{q}|} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|} \right] =: L_1(\boldsymbol{\omega}). \quad (1.5.2)$$

- **Lipschitz condition:** Let $[\mathbf{z}_1, \mathbf{z}_2] = \{a\mathbf{z}_1 + (1-a)\mathbf{z}_2 : a \in [0, 1]\}$ denote the segment connecting $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^d$. By using the mean value theorem

$$|g_{\mathbf{z}_1}(\boldsymbol{\omega}) - g_{\mathbf{z}_2}(\boldsymbol{\omega})| \leq \max_{\mathbf{z} \in [\mathbf{z}_1, \mathbf{z}_2]} \left\| \frac{\partial g_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}} \right\|_2 \|\mathbf{z}_1 - \mathbf{z}_2\|_2, \quad (1.5.3)$$

$\frac{\partial g_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}} = \frac{\partial f_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}} - \Lambda \frac{\partial f_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}}$ with $\frac{\partial f_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}} = \boldsymbol{\omega}^{\mathbf{p}}(-\boldsymbol{\omega})^{\mathbf{q}} c_{|\mathbf{p}+\mathbf{q}|+1}(\boldsymbol{\omega}^\top \mathbf{z}) \boldsymbol{\omega}$, and by using similar computations as before, one gets

$$\left\| \frac{\partial g_{\mathbf{z}}(\boldsymbol{\omega})}{\partial \mathbf{z}} \right\|_2 \leq \|\boldsymbol{\omega}\|_2^{|\mathbf{p}+\mathbf{q}|+1} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|+1} \right] =: L_2(\boldsymbol{\omega}). \quad (1.5.4)$$

As a result, to fulfill Assumption 2, we can take $L(\boldsymbol{\omega}) = \max(L_1(\boldsymbol{\omega}), L_2(\boldsymbol{\omega}))$ and $\rho(u) = u$. For such L , we have $n = |\mathbf{p} + \mathbf{q}| + 1$.

Refined L and ρ : We now derive refined L and ρ , by interpolating different bounds. From (1.5.2), we can obtain the crude estimate $|g_{\mathbf{z}_1}(\boldsymbol{\omega}) - g_{\mathbf{z}_2}(\boldsymbol{\omega})| \leq 2L_1(\boldsymbol{\omega})$, which combined with (1.5.3)-(1.5.4) gives

$$|g_{\mathbf{z}_1}(\boldsymbol{\omega}) - g_{\mathbf{z}_2}(\boldsymbol{\omega})| \leq (2L_1(\boldsymbol{\omega}))^{1-\beta} \left[\|\mathbf{z}_1 - \mathbf{z}_2\|_2 L_2(\boldsymbol{\omega}) \right]^\beta$$

for any $\beta \in (0, 1]$. Here we have used that if $0 \leq x \leq \min(x_1, x_2)$ then $x \leq x_1^{1-\beta} x_2^\beta$. It follows that one can take

$$\rho(u) = u^\beta, \quad n = |\mathbf{p} + \mathbf{q}| + \beta, \quad L(\boldsymbol{\omega}) = \max \left(L_1(\boldsymbol{\omega}), (2L_1(\boldsymbol{\omega}))^{1-\beta} L_2^\beta(\boldsymbol{\omega}) \right) \in \mathcal{F}_{\mathcal{P}(n)}.$$

For $\beta = 1$, we retrieve the former choice of L and ρ . Furthermore, we have

$$I_\rho(|T|) = |T|^\beta \int_0^1 \sqrt{\log \left(1 + \frac{2|T|}{(u|T|^\beta)^{1/\beta}} \right)} du = |T|^\beta \int_0^1 \sqrt{\log \left(1 + \frac{2}{u^{1/\beta}} \right)} du < +\infty. \quad (1.5.5)$$

Notice that the advantage of having the additional degree-of-freedom β is two-fold, and it is striking when $\beta \rightarrow 0$ (compared to $\beta = 1$). Firstly, it gives a smaller n , which has a (slight) positive impact on the control of statistical fluctuations; secondly, the dependence of $I_\rho(|T|)$ in the diameter $|T|$ is smaller through the growth exponent.

To conclude, we have proved that Orlicz RFFs fulfill the assumptions of Theorem 1.1. Later in Section 1.5.6, we will establish that $I_\rho(|T|)$ satisfies a (tight) bound w.r.t. $\sqrt{1 + (\log|T|)_+}$.

1.5.2 Proof that polynomial growth preserves the exponential Orlicz property

We show that $\|f(\mathbf{X})\|_{\Psi_\gamma} < \infty$ for $\|\mathbf{X}\|_{\Psi_\alpha} < \infty$, $f \in \mathcal{F}_{\mathcal{P}(n)}$, $n \in \mathbb{R}^+$, $\gamma = \frac{\alpha}{n}$. Indeed, by the definition of $f \in \mathcal{F}_{\mathcal{P}(n)}$, there exists $C \in \mathbb{R}^+$ such that $|f(\mathbf{x})| \leq C(1 + \|\mathbf{x}\|_2^n)$ for all $\mathbf{x} \in \mathbb{R}^d$. Hence for any $\gamma > 0$

$$|f(\mathbf{x})|^\gamma \leq C^\gamma (1 + \|\mathbf{x}\|_2^n)^\gamma \stackrel{(*)}{\leq} 2^{(\gamma-1)+} C^\gamma (1 + \|\mathbf{x}\|_2^{n\gamma}), \quad (1.5.6)$$

where in (*) we used that

$$(a + b)^\gamma \leq 2^{(\gamma-1)+} (a^\gamma + b^\gamma), \quad a, b \geq 0, \gamma > 0. \quad (1.5.7)$$

Since $\mathbf{X} \in L_{\Psi_\alpha}$ there is some $s \in \mathbb{R}^+$ for which $\mathbb{E} \left[e^{s\|\mathbf{X}\|_2^\alpha} \right] < \infty$. Combining this property with (1.5.6) and recalling that $n\gamma = \alpha$ yields

$$e^{s'|f(\mathbf{x})|^\gamma} \leq e^{s'2^{(\gamma-1)+} C^\gamma (1 + \|\mathbf{x}\|_2^\alpha)} \quad \Rightarrow \quad \mathbb{E} \left[e^{s'|f(\mathbf{X})|^\gamma} \right] \leq e^{s'2^{(\gamma-1)+} C^\gamma} \mathbb{E} \left[e^{s\|\mathbf{X}\|_2^\alpha} \right] < \infty$$

with $s' = \frac{s}{2^{(\gamma-1)+} C^\gamma}$; this shows that $f(\mathbf{X}) \in L_{\Psi_\gamma}$.

1.5.3 Proof of Theorem 1.1

By introducing the $\mathcal{R}_c f(\mathbf{x}) := f(\mathbf{x}) - \mathcal{T}_c f(\mathbf{x})$ notation of residuals obtained at level $c \in \mathbb{R}^+$ (the value of c will be specified later), we bound the target quantity by using the sub-additivity of supremum

$$\begin{aligned}
& \sup_{t \in T} \frac{1}{M} \sum_{m=1}^M \underbrace{f_t(\mathbf{X}_m)}_{\mathcal{T}_c f_t(\mathbf{X}_m) + \mathcal{R}_c f_t(\mathbf{X}_m)} \\
&= \sup_{t \in T} \frac{1}{M} \sum_{m=1}^M (\mathcal{T}_c f_t(\mathbf{X}_m) - \mathbb{E}[\mathcal{T}_c f_t(\mathbf{X}_m)] + \mathbb{E}[\mathcal{T}_c f_t(\mathbf{X}_m)] + \mathcal{R}_c f_t(\mathbf{X}_m)) \\
&\leq \underbrace{\sup_{t \in T} \frac{1}{M} \sum_{m=1}^M (\mathcal{T}_c f_t(\mathbf{X}_m) - \mathbb{E}[\mathcal{T}_c f_t(\mathbf{X}_m)])}_{\bar{Z}^{\mathcal{T}_c}} + \underbrace{\sup_{t \in T} \mathbb{E} \left[\frac{1}{M} \sum_{m=1}^M \mathcal{T}_c f_t(\mathbf{X}_m) \right]}_{\mathcal{E}^{\mathcal{T}_c}} \\
&\quad + \underbrace{\sup_{t \in T} \frac{1}{M} \sum_{m=1}^M \mathcal{R}_c f_t(\mathbf{X}_m)}_{Z^{\mathcal{R}_c}}.
\end{aligned}$$

This means that using c for which $\mathcal{E}^{\mathcal{T}_c} \leq \frac{\varepsilon}{3}$,

$$\mathbb{P} \left(\sup_{t \in T} \frac{1}{M} \sum_{m=1}^M f_t(\mathbf{X}_m) \geq \varepsilon \right) \leq \mathbb{P}(Z^{\mathcal{R}_c} \geq \varepsilon/3) + \mathbb{P}(\bar{Z}^{\mathcal{T}_c} \geq \varepsilon/3). \quad (1.5.8)$$

The structure of our proof is as follows.

1. Unbounded part ($Z^{\mathcal{R}_c}$): Based on the Talagrand and the Hoffman-Jorgensen inequalities, for large enough c (referred to as c_{HJ}) we will derive an exponential control over $\mathbb{P}(Z^{\mathcal{R}_c} \geq \varepsilon/3)$ expressed with $\|\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|\|_{\Psi_\gamma}$ which is finite by Section 1.5.4.
2. Bounded part ($\bar{Z}^{\mathcal{T}_c}$): We handle this term using the Klein-Rio inequality and the Dudley entropy integral bound. In addition, this part will give rise to the constraint (1.3.1) on ε .
3. Truncation ($\mathcal{E}^{\mathcal{T}_c}$): As $\mathbb{E}[f_t(\mathbf{X}_m)] = 0$, $\mathcal{T}_c f_t \approx f_t$ and $\mathbb{E}[\mathcal{T}_c f_t(\mathbf{X}_m)] \approx 0$ for large c (called c_{min}). The $\mathcal{E}^{\mathcal{T}_c} \leq \frac{\varepsilon}{3}$ requirement can be controlled via the integral form of the expectation of non-negative random variables and the incomplete Gamma function.

The bounding of the $Z^{\mathcal{R}_c}$, $\bar{Z}^{\mathcal{T}_c}$ and $\mathcal{E}^{\mathcal{T}_c}$ quantities is detailed in the following sections. Plugging the (1.5.9) and (1.5.11) results of the computations into (1.5.8) gives the final bound (1.3.2). The ε constraint comes from (1.5.16), provided that $c \geq c_{min} \vee c_{HJ}$. The constants c_{min} and c_{HJ} are defined in (1.5.18) and (1.5.21), respectively.

1.5.3.1 Bounding $Z^{\mathcal{R}_c}$

$\mathbb{P}(Z^{\mathcal{R}_c} \geq \varepsilon/3)$ is bounded as

$$\begin{aligned}
\mathbb{P}(Z^{\mathcal{R}_c} \geq \varepsilon/3) &\leq \mathbb{P} \left(\sup_{t \in T} \sum_{m=1}^M |\mathcal{R}_c f_t(\mathbf{X}_m)| \geq M\varepsilon/3 \right) \\
&\stackrel{(a)}{\leq} \mathbb{P} \left(\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \geq M\varepsilon/3 \right) \stackrel{(b)}{\leq} 2e^{-\left(\frac{M\varepsilon/3}{\|\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)|\|_{\Psi_\gamma}} \right)^\gamma} \\
&\stackrel{(c)}{\leq} 2e^{-\left(\frac{M\varepsilon}{3K_\gamma \|\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|\|_{\Psi_\gamma}} \right)^\gamma},
\end{aligned}$$

where in (a) we used the sub-additivity of the supremum, in (b) the deviation inequality Section (1.4)(vii) was applied, (c) holds by Section 1.5.5 for $c \geq c_{HJ}$ (the value of c_{HJ} is defined in Section 1.5.5).

1.5.3.2 Bounding $\overline{Z}^{\mathcal{T}_c}$

Below we will invoke the Klein-Rio inequality and control the expectation $\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right]$.

- **Klein-Rio inequality:** Let $g_{m,\mathbf{t}} : \mathbf{x} \in \mathbb{R}^d \mapsto \mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}) - \mathbb{E}[\mathcal{T}_c f_{\mathbf{t}}(\mathbf{X}_m)]$ and let us define the function classes

$$\mathcal{T}_c \mathcal{F}^{[M]} := \{\mathbf{g}_{\mathbf{t}} := (g_{1,\mathbf{t}}, \dots, g_{M,\mathbf{t}}) : \mathbf{t} \in T\}, \quad \mathcal{T}_c \mathcal{F} := \{\mathcal{T}_c f_{\mathbf{t}} : \mathbf{t} \in T\}.$$

- $g_{m,\mathbf{t}} \in [-2c, 2c]$ are measurable and bounded functions.
- Centering: by construction $\mathbb{E}[g_{m,\mathbf{t}}(\mathbf{X}_m)] = 0$ ($\forall m \in [M]$).
- Countability: Since $\mathbf{t} \mapsto f_{\mathbf{t}}$ is continuous, the $\sup_{\mathbf{t} \in T}$ can be restricted to rational numbers ($T \cap \mathbb{Q}^d$), one can take $T \leftarrow T \cap \mathbb{Q}^d$, and assume that $\mathcal{T}_c \mathcal{F}^{[M]}$ is countable.

If

$$\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right] \leq \varepsilon/6, \tag{1.5.10}$$

then the Klein-Rio inequality (Theorem 1.3 where the $\sup_{\mathbf{t} \in T}$ and $\sup_{\mathbf{f} \in \mathcal{T}_c \mathcal{F}^{[M]}}$ coincide) implies that

$$\begin{aligned} \mathbb{P} \left(\overline{Z}^{\mathcal{T}_c} \geq \varepsilon/3 \right) &\stackrel{(1.5.10)}{\leq} \mathbb{P} \left(\overline{Z}^{\mathcal{T}_c} - \mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right] \geq \varepsilon/6 \right) \leq e^{-\frac{M(\varepsilon/6)^2}{2(\overline{\sigma}^2 + 4c\mathbb{E}[\overline{Z}^{\mathcal{T}_c}] + 6c\varepsilon/6)}} \\ &\stackrel{(1.5.10)}{\leq} e^{-\frac{M(\varepsilon/6)^2}{2\overline{\sigma}^2 + 14c\varepsilon/6}} = e^{-\frac{M\varepsilon^2}{72\overline{\sigma}^2 + 84c\varepsilon}} \leq e^{-\frac{M\varepsilon^2}{72\sigma^2 + 84c\varepsilon}}, \end{aligned} \tag{1.5.11}$$

where the weak variance $\overline{\sigma}^2$ is defined and bounded by

$$\begin{aligned} \overline{\sigma}^2 &:= \sup_{\mathbf{t} \in T} \frac{1}{M} \sum_{m=1}^M \mathbb{E} \left[(\mathcal{T}_c f_{\mathbf{t}}(\mathbf{X}_m) - \mathbb{E}[\mathcal{T}_c f_{\mathbf{t}}(\mathbf{X}_m)])^2 \right] \leq \sup_{\mathbf{t} \in T} \frac{1}{M} \sum_{m=1}^M \mathbb{E} \left[(\mathcal{T}_c f_{\mathbf{t}}(\mathbf{X}_m))^2 \right] \\ &\leq \sup_{\mathbf{t} \in T} \frac{1}{M} \sum_{m=1}^M \mathbb{E} [f_{\mathbf{t}}^2(\mathbf{X}_m)] =: \sigma^2. \end{aligned}$$

- **Bounding $\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right]$:** We control $\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right]$ in (1.5.10) by the Dudley entropy integral bound. In this bound the covering number of $\mathcal{T}_c \mathcal{F}$ is estimated by that of the compact set $T \subset \mathbb{R}^d$ with propagation relying on Assumption (2b).
 - **Dudley entropy integral bound:** Slight modification (without absolute value) of [vW96, Lemma 2.3.1] gives

$$\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right] \leq 2\mathbb{E} [\mathcal{R}(\mathbf{X}_{1:M}, \mathcal{T}_c \mathcal{F})], \tag{1.5.12}$$

where $\mathcal{R}(\mathbf{x}_{1:M}, \mathcal{T}_c \mathcal{F}) := \mathbb{E}_{\boldsymbol{\varepsilon}} \left[\sup_{\mathbf{t} \in T} \frac{1}{M} \sum_{m=1}^M \varepsilon_m \mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}_m) \right]$ is the Rademacher average of $\mathcal{T}_c \mathcal{F}$, $\mathbf{X}_{1:M} := (\mathbf{X}_m)_{m \in [M]}$, $\mathbf{x}_{1:M} := (\mathbf{x}_m)_{m \in [M]}$, $\boldsymbol{\varepsilon} := (\varepsilon_m)_{m \in [M]}$ contains independent Rademacher variables (i.e. $\mathbb{P}(\varepsilon_m = \pm 1) = \frac{1}{2}$), and $\boldsymbol{\varepsilon}$ is independent of $\mathbf{X}_{1:M}$.

Let $Z_{\mathbf{t}}(\mathbf{x}_{1:M}) := \frac{1}{M} \sum_{m=1}^M \varepsilon_m \mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}_m)$, so $\mathcal{R}(\mathbf{x}_{1:M}, \mathcal{T}_c \mathcal{F}) = \mathbb{E}_{\boldsymbol{\varepsilon}} \left[\sup_{\mathbf{t} \in T} Z_{\mathbf{t}}(\mathbf{x}_{1:M}) \right]$, and define the pseudo-metric on T as $d(\mathbf{t}, \mathbf{s}) := \left(\frac{1}{M} \sum_{m=1}^M [\mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}_m) - \mathcal{T}_c f_{\mathbf{s}}(\mathbf{x}_m)]^2 \right)^{1/2}$.

The $\{Z_{\mathbf{t}} : \mathbf{t} \in T\}$ process is

* separable since it is continuous, and $T \subset \mathbb{R}^d$ is separable,

- * centered thanks to the Rademacher variables,
- * sub-Gaussian with respect to $M^{-1/2}d$: indeed, for any $\lambda \in \mathbb{R}^+$ and $\mathbf{t}, \mathbf{s} \in T$

$$\begin{aligned} \mathbb{E}_\varepsilon \left[e^{\lambda(Z_{\mathbf{t}} - Z_{\mathbf{s}})} \right] &\stackrel{(a)}{=} \prod_{m=1}^M \mathbb{E}_{\varepsilon_m} \left[e^{\varepsilon_m \frac{\lambda}{M} [\mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}_m) - \mathcal{T}_c f_{\mathbf{s}}(\mathbf{x}_m)]} \right] \stackrel{(b)}{\leq} \prod_{m=1}^M e^{\frac{\lambda^2}{2M^2} [\mathcal{T}_c f_{\mathbf{t}}(\mathbf{x}_m) - \mathcal{T}_c f_{\mathbf{s}}(\mathbf{x}_m)]^2} \\ &= e^{\frac{\lambda^2 (M^{-1/2}d(\mathbf{t}, \mathbf{s}))^2}{2}}. \end{aligned}$$

In (a) we used the independence of ε_m -s, (b) follows from $\mathbb{E}_{\varepsilon_m} [e^{a\varepsilon_m}] = \cosh(a) \stackrel{(*)}{\leq} e^{\frac{a^2}{2}}$ ($\forall a \in \mathbb{R}$), where $(*)$ can be obtained from the power series expansion of the $\cosh(\cdot)$ and the exponential function.

Hence Theorem 1.4 can be applied:

$$\begin{aligned} \mathcal{R}(\mathbf{x}_{1:M}, \mathcal{T}_c \mathcal{F}) &\leq C_D \int_0^\infty \sqrt{\log(N(\varepsilon, M^{-1/2}d, T))} d\varepsilon \\ &= \frac{C_D}{\sqrt{M}} \int_0^\infty \sqrt{\log(N(\varepsilon, d, T))} d\varepsilon, \end{aligned} \quad (1.5.13)$$

where we used the $N(\varepsilon, M^{-1/2}d, T) = N(M^{1/2}\varepsilon, d, T)$ identity, and applied an $\tilde{\varepsilon} = M^{1/2}\varepsilon$ substitution. We note that the above infinite integral can be truncated at $|T|_d := \sup_{\mathbf{t}, \mathbf{s} \in T} d(\mathbf{t}, \mathbf{s})$, the d -diameter of T , since $N(\varepsilon, d, T) = 1$ for $\varepsilon \geq |T|_d$.

- **Covering number:** By (1.3.3) one can relate $d(\mathbf{t}, \mathbf{s})$ and $\|\mathbf{t} - \mathbf{s}\|_2$ as

$$d(\mathbf{t}, \mathbf{s}) \leq \left[\frac{1}{M} \sum_{m=1}^M L^2(\mathbf{x}_m) \right]^{1/2} \rho(\|\mathbf{t} - \mathbf{s}\|_2) := \|L\|_{L^2(w_{1:M})} \rho(\|\mathbf{t} - \mathbf{s}\|_2),$$

which implies

$$N(\varepsilon, d, T) \leq N\left(\rho^{-1}\left(\frac{\varepsilon}{\|L\|_{L^2(w_{1:M})}}\right), \|\cdot\|_2, T\right), \quad (1.5.14)$$

$$|T|_d \leq \|L\|_{L^2(w_{1:M})} \sup_{\mathbf{t}, \mathbf{s} \in T} \rho(\|\mathbf{t} - \mathbf{s}\|_2) \leq \|L\|_{L^2(w_{1:M})} \rho(|T|). \quad (1.5.15)$$

In the last inequality the increasing property of ρ was exploited. Combining (1.5.14)-(1.5.15) with the well-known bound [vdG00, Lemma 2.5] on the covering number⁷ of a compact set $T \subset \mathbb{R}^{d'}$

$$N(\varepsilon', \|\cdot\|_2, T) \leq \left(\frac{2|T|}{\varepsilon'} + 1\right)^{d'}, \quad \forall \varepsilon' > 0,$$

(1.5.13) can be estimated further as

$$\begin{aligned} \mathcal{R}(\mathbf{x}_{1:M}, \mathcal{T}_c \mathcal{F}) &\leq \frac{C_D \sqrt{d'}}{\sqrt{M}} \int_0^{\|L\|_{L^2(w_{1:M})} \rho(|T|)} \sqrt{\log\left(\frac{2|T|}{\rho^{-1}\left(\frac{\varepsilon}{\|L\|_{L^2(w_{1:M})}}\right)} + 1\right)} d\varepsilon \\ &= C_D \sqrt{d'} \frac{\|L\|_{L^2(w_{1:M})} \rho(|T|)}{\sqrt{M}} \int_0^1 \sqrt{\log\left(1 + \frac{2|T|}{\rho^{-1}(u\rho(|T|))}\right)} du, \end{aligned}$$

⁷In our definition of the covering number, in its bound on compact sets in \mathbb{R}^d [vdG00, Lemma 2.5] and in the final Dudley entropy bound [Bar13, Lecture 11, 14] the elements of the ε -net are assumed to belong to the set covered.

where we introduced the new variable $u = \frac{\varepsilon}{\|L\|_{L^2(w_{1:M})} \rho(|T|)}$. Substituting this bound into (1.5.12) we arrive at

$$\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right] \leq 2C_D \sqrt{d'} \frac{\mathbb{E} \left[\|L\|_{L^2(W_{1:M})} \right]}{\sqrt{M}} I_\rho(|T|) =: B. \quad (1.5.16)$$

To guarantee $\mathbb{E} \left[\overline{Z}^{\mathcal{T}_c} \right] \leq \varepsilon/6$, we solve $B \leq \frac{\varepsilon}{6}$; this gives the (1.3.1) bound on ε .

1.5.3.3 Bounding $\mathcal{E}^{\mathcal{T}_c}$

- Bounding $\mathcal{E}^{\mathcal{T}_c}$ by the incomplete Gamma function (I_γ):

$$\begin{aligned} \mathbb{E} [\mathcal{T}_c f_t(\mathbf{X}_m)] &\stackrel{(a)}{=} -\mathbb{E} [\mathcal{R}_c f_t(\mathbf{X}_m)] \stackrel{(b)}{\leq} \mathbb{E} [(-f_t(\mathbf{X}_m) - c) \mathbf{1}_{f_t(\mathbf{X}_m) \leq -c}] \\ &\stackrel{(c)}{=} \int_c^\infty \mathbb{P}(-f_t(\mathbf{X}_m) \geq y) dy. \end{aligned} \quad (1.5.17)$$

In (a) we used that $\mathcal{T}_c f_t(\mathbf{x}) = f_t(\mathbf{x}) - \mathcal{R}_c f_t(\mathbf{x})$ and $\mathbb{E}[f_t(\mathbf{X}_m)] = 0$, (b) follows from

$$\mathcal{R}_c f_t(\mathbf{x}) = [f_t(\mathbf{x}) + c] \mathbf{1}_{f_t(\mathbf{x}) \leq -c} + [f_t(\mathbf{x}) - c] \mathbf{1}_{f_t(\mathbf{x}) \geq c} \geq [f_t(\mathbf{x}) + c] \mathbf{1}_{f_t(\mathbf{x}) \leq -c}.$$

(c) holds by using that for a $Z \geq 0$ random variable, $\mathbb{E}[Z] = \int_0^\infty \mathbb{P}(Z \geq z) dz$; we choose $Z = \max(0, -f_t(\mathbf{X}_m) - c)$. Therefore

$$\begin{aligned} \mathcal{E}^{\mathcal{T}_c} &= \sup_{t \in T} \mathbb{E} \left[\frac{1}{M} \sum_{m=1}^M \mathcal{T}_c f_t(\mathbf{X}_m) \right] \stackrel{(a)}{\leq} \max_{1 \leq m \leq M} \sup_{t \in T} \int_c^\infty \mathbb{P}(-f_t(\mathbf{X}_m) \geq y) dy \\ &\stackrel{(b)}{\leq} 2 \max_{1 \leq m \leq M} \sup_{t \in T} \int_c^\infty e^{-\left(\frac{y}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}\right)^\gamma} dy \\ &\stackrel{(c)}{=} 2 \max_{1 \leq m \leq M} \sup_{t \in T} \left(\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \int_{\frac{c}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}}^\infty e^{-u^\gamma} du \right) \\ &\stackrel{(d)}{=} 2 \max_{1 \leq m \leq M} \sup_{t \in T} \left(\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \left[\int_0^\infty e^{-u^\gamma} du - \int_0^{\frac{c}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}} e^{-u^\gamma} du \right] \right) \\ &\stackrel{(e)}{=} 2 \max_{1 \leq m \leq M} \sup_{t \in T} \left(\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \left[\Gamma\left(1 + \frac{1}{\gamma}\right) - I_\gamma\left(\frac{c}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}\right) \right] \right) \\ &\stackrel{(f)}{\leq} 2 \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \left[\Gamma\left(1 + \frac{1}{\gamma}\right) - I_\gamma\left(\frac{c}{\max_{m' \in [M]} \sup_{t' \in T} \|f_{t'}(\mathbf{X}_{m'})\|_{\Psi_\gamma}}\right) \right] \\ &\stackrel{(g)}{\leq} 2\Gamma\left(1 + \frac{1}{\gamma}\right) \left(1 - \left[1 - e^{-\beta_\gamma \left(\frac{c}{\max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}\right)^\gamma} \right]^{\frac{1}{\gamma}} \right) \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \\ &=: \tilde{B}, \end{aligned}$$

where (a) holds by taking maximum over $m \in [M]$ and using (1.5.17), (b) follows from the $\mathbb{P}(-f_t(\mathbf{X}_m) \geq y) \leq 2e^{-\left(\frac{y}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}\right)^\gamma}$ deviation inequality implied by Section 1.4(vii). (c) was obtained from a $u = \frac{y}{\|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}$ substitution, in (d) we decomposed the integral to have the incomplete Gamma function appear. (e) is a consequence of the definition of I_γ and the limit

$$I_\gamma(x) = \int_0^x e^{-t^\gamma} dt = \frac{1}{\gamma} \int_0^{x^\frac{1}{\gamma}} u^{\frac{1}{\gamma}-1} e^{-u} du \xrightarrow{x \rightarrow \infty} \frac{1}{\gamma} \Gamma\left(\frac{1}{\gamma}\right) = \Gamma\left(1 + \frac{1}{\gamma}\right),$$

where we applied an $u = t^\gamma$ substitution and the $\Gamma(z+1) = z\Gamma(z)$ recursion. (f) comes from the monotonicity of I_γ ($I_\gamma(x) \leq I_\gamma(y)$ if $x \leq y$). (g) follows from applying the lower bound on I_γ from Theorem 1.5 with $x = \frac{c}{\max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}$.

- **Additional truncation level bound on c :** Guaranteeing $\tilde{B} \leq \frac{\epsilon}{3}$ (and thus $\mathcal{E}^{\mathcal{T}^c} \leq \frac{\epsilon}{3}$) is equivalent to choosing c large enough such that

$$1 - \left[1 - e^{-\beta_\gamma \left(\frac{c}{\max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}} \right)^\gamma} \right]^{\frac{1}{\gamma}} \leq \frac{\epsilon}{6\Gamma\left(1 + \frac{1}{\gamma}\right) \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}.$$

Because $\gamma \leq 1$, the function $h : x \mapsto 1 - (1-x)^{\frac{1}{\gamma}}$ is concave on $[0, 1]$, and thus it is below its tangent line computed at $(0, h(0))$, i.e. $1 - (1-x)^{\frac{1}{\gamma}} \leq \frac{1}{\gamma}x$. Therefore choosing $x = e^{-\beta_\gamma \left(\frac{c}{\max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}} \right)^\gamma}$ it is enough to use c such that

$$\frac{1}{\gamma} e^{-\beta_\gamma \left(\frac{c}{\max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}} \right)^\gamma} \leq \frac{\epsilon}{6\Gamma\left(1 + \frac{1}{\gamma}\right) \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}.$$

Solving this inequality for c means that

$$c \geq c_{min} := \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma} \left[\frac{1}{\beta_\gamma} \log \left(\frac{6\Gamma\left(1 + \frac{1}{\gamma}\right) \max_{1 \leq m \leq M} \sup_{t \in T} \|f_t(\mathbf{X}_m)\|_{\Psi_\gamma}}{\gamma \epsilon} \right) \right]^{\frac{1}{\gamma}}. \quad (1.5.18)$$

1.5.4 Proof of $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$ and finite maximal moments

By Assumption (2a)-(2b), the triangle inequality and the monotonicity of ρ , one gets

$$|f_t(\mathbf{x})| \leq |f_t(\mathbf{x}) - f_{\mathbf{t}_0}(\mathbf{x})| + |f_{\mathbf{t}_0}(\mathbf{x})| \leq L(\mathbf{x}) [\rho(\|\mathbf{t} - \mathbf{t}_0\|_2) + 1] \leq L(\mathbf{x}) [\rho(|T|) + 1], \quad (1.5.19)$$

for any $\mathbf{t} \in T, \mathbf{x} \in \mathbb{R}^d$. The individual statements now can be proved as follows.

- $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$: By $L \in \mathcal{F}_{\mathcal{P}(n)}$ and (1.5.19), $f_t \in \mathcal{F}_{\mathcal{P}(n)}$ for all $\mathbf{t} \in T$, in other words $\mathcal{F} \subset \mathcal{F}_{\mathcal{P}(n)}$.
- **Finiteness of $\|\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|\|_{\Psi_{\frac{\alpha}{n}}}$:** Using (1.5.19), we get

$$\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \leq [\rho(|T|) + 1] \sum_{m=1}^M L(\mathbf{X}_m). \quad (1.5.20)$$

Thanks to Section 1.5.2, each $L(\mathbf{X}_m)$ belongs to $L_{\Psi_{\frac{\alpha}{n}}}$. Combining this with the generalized triangular inequality Section 1.4(vi) gives the claim.

- **Finiteness of $\mathbb{E}[\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)|]$:** Each $L(\mathbf{X}_m)$ is integrable (because L has a polynomial growth and the distribution of \mathbf{X}_m satisfies the α -Orlicz exponential assumption). Thus, the statement follows from (1.5.20).

1.5.5 Control if $c \geq c_{HJ}$

We show that under the assumptions of Theorem 1.1 with

$$c \geq c_{HJ} := 8\mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right] \quad (1.5.21)$$

one has

$$\left\| \sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} \leq K_\gamma \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma}. \quad (1.5.22)$$

Notice that c_{HJ} is finite by Section 1.5.4. We bound the l.h.s. of (1.5.22):

$$\begin{aligned} & \left\| \sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} = \\ & = \left\| \sum_{m=1}^M \left(\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] + \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \right\|_{\Psi_\gamma} \\ & \stackrel{(a)}{\leq} 2^{\frac{1}{\gamma}-1} \left(\left\| \sum_{m=1}^M \left(\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \right\|_{\Psi_\gamma} + \right. \\ & \quad \left. + \left\| \mathbb{E} \left[\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right\|_{\Psi_\gamma} \right) \\ & \stackrel{(b)}{\leq} 2^{\frac{1}{\gamma}-1} \left(C_\gamma \left(\mathbb{E} \left[\sum_{m=1}^M \left(\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \right] \right) + \right. \\ & \quad \left. + \left\| \max_{1 \leq m \leq M} \left| \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right| \right\|_{\Psi_\gamma} \right) + \\ & \quad \left. + \frac{1}{\Psi_\gamma^{-1}(1)} \mathbb{E} \left[\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \\ & =: 2^{\frac{1}{\gamma}-1} \left[C_\gamma (E_1 + E_2) + \frac{1}{\Psi_\gamma^{-1}(1)} E_3 \right]. \end{aligned}$$

In (a) we applied the generalized triangle inequality Section 1.4(iv) and $\left(\frac{1}{\gamma} - 1\right)_+ = \frac{1}{\gamma} - 1$ as $\gamma = \frac{\alpha}{n} \in (0, 1]$. In (b) the Talagrand inequality (1.C.1) was invoked with the $Y_m := \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} [\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)|]$ centered variables and $B := \mathbb{R}$, followed by taking the γ -Orlicz norm of the constant $\lambda := \mathbb{E} \left[\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right]$ according to Section 1.4(ii).

We continue the derivation with bounding the E_1 , E_2 and E_3 terms in (1.5.23).

• **Bounding E_1 :**

$$\begin{aligned} E_1 &= \mathbb{E} \left[\left\| \sum_{m=1}^M \left(\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \right\|_{\Psi_\gamma} \right] \\ & \stackrel{(a)}{\leq} 2 \mathbb{E} \left[\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \stackrel{(b)}{\leq} 16 \mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \\ & \stackrel{(c)}{\leq} 16 \mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right] \stackrel{(d)}{\leq} 16 \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma^l} \left(\Psi_\gamma^{(l)} \right)^{-1} (1) \\ & \stackrel{(e)}{\leq} 16 \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} \left(\Psi_\gamma^{(l)} \right)^{-1} (1), \end{aligned}$$

where in (a) we used the triangle inequality, in (b) we applied the Hoffman-Jorgensen inequality

ity (Theorem 1.1; $t_0 = 0$, $p = 1$, $B = \mathbb{R}$, $Y_m = \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)|$) with

$$\begin{aligned} \mathbb{P} \left(\sum_{m \in [M]} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| > 0 \right) &\stackrel{(f)}{=} \mathbb{P} \left(\max_{j \in [M]} \sum_{m \in [j]} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| > 0 \right) \\ &= \mathbb{P} \left(\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| > c \right) \\ &\stackrel{(g)}{\leq} \mathbb{P} \left(\max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \geq c_{HJ} \right) \stackrel{(h)}{\leq} \frac{1}{8} = \frac{1}{2 \times 4^p} \text{ with } p = 1. \end{aligned}$$

In (f) the non-negativity of Y_m was exploited; in (g) $c \geq c_{HJ}$ was used. We applied the Markov inequality and the definition of c_{HJ} in (h). (c) holds by $|\mathcal{R}_c f_t(\mathbf{X}_m)| \leq |f_t(\mathbf{X}_m)|$. In (d) and (e) we applied Section 1.4(v) with the convex $\Psi_\gamma^{(l)}$ defined in Section 1.A.1 and the monotonicity property Section 1.4(iii) with $\Psi_\gamma^{(l)} \leq \Psi_\gamma$, respectively.

• **Bounding E_2 :**

$$\begin{aligned} E_2 &= \left\| \max_{1 \leq m \leq M} \left| \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| - \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right\|_{\Psi_\gamma} \\ &\stackrel{(a)}{\leq} \left\| \max_{1 \leq m \leq M} \left(\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| + \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \right\|_{\Psi_\gamma} \\ &\stackrel{(b)}{\leq} \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| + \max_{1 \leq m \leq M} \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right\|_{\Psi_\gamma} \\ &\stackrel{(c)}{\leq} 2^{\frac{1}{\gamma}-1} \left(\left\| \max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} + \left\| \max_{1 \leq m \leq M} \mathbb{E} \left[\sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right\|_{\Psi_\gamma} \right) \\ &\stackrel{(d)}{\leq} 2^{\frac{1}{\gamma}-1} \left(\left\| \max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} + \frac{1}{\Psi_\gamma^{-1}(1)} \mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \right) \\ &\stackrel{(e)}{\leq} 2^{\frac{1}{\gamma}-1} \left(\left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} + \frac{(\Psi_\gamma^{(l)})^{-1}(1)}{\Psi_\gamma^{-1}(1)} \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma^{(l)}} \right) \\ &\stackrel{(f)}{\leq} 2^{\frac{1}{\gamma}-1} \left(1 + \frac{(\Psi_\gamma^{(l)})^{-1}(1)}{\Psi_\gamma^{-1}(1)} \right) \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma}. \end{aligned}$$

In (a) we used the triangle inequality with the monotonicity Section 1.44, (b) holds by the sub-additivity of the maximum and again the monotonicity Section 1.44, in (c) we applied the generalized triangle inequality Section 1.4(iv) and that $\left(\frac{1}{\gamma} - 1\right)_+ = \frac{1}{\gamma} - 1$ as $\gamma \in (0, 1]$, (d) holds by Section 1.4(ii) with the constant $\lambda = \mathbb{E}[\max_{1 \leq m \leq M} \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)|]$, (e) is by the monotonicity Section 1.44 as $|\mathcal{R}_c f_t(\mathbf{X}_m)| \leq |f_t(\mathbf{X}_m)|$, and by Section 1.4(v), (f) follows from $\Psi_\gamma^{(l)} \leq \Psi_\gamma$ combined with the monotonicity Section 1.4(iii).

• **Bounding E_3 :** By (b)-(e) of the E_1 derivation we have that

$$E_3 = \mathbb{E} \left[\sum_{m=1}^M \sup_{t \in T} |\mathcal{R}_c f_t(\mathbf{X}_m)| \right] \leq 8 \left\| \max_{1 \leq m \leq M} \sup_{t \in T} |f_t(\mathbf{X}_m)| \right\|_{\Psi_\gamma} (\Psi_\gamma^{(l)})^{-1}(1).$$

By adding the obtained E_1 , E_2 and E_3 bounds, we get (1.5.22) with K_γ defined in Theorem 1.1.

1.5.6 Bounding the driving terms of Theorem 1.1 for RFF

We bound the constants of Theorem 1.1 in the RFF case described in Remark 1(v).

- **The term B :** It is defined in (1.3.1). Recalling the expression (1.5.5) for $I_\rho(|T|)$ and using the Cauchy-Schwarz inequality for bounding $\mathbb{E}[\|L\|_{L^2(\omega_{1:M})}]$ by $\sqrt{\mathbb{E}_{\omega \sim \Lambda}[L^2(\omega)]}$ gives

$$B \leq 2C_D \sqrt{d} \frac{\sqrt{\mathbb{E}_{\omega \sim \Lambda}[L^2(\omega)]}}{\sqrt{M}} |S_\Delta|^\beta \int_0^1 \sqrt{\log(1 + 2u^{-1/\beta})} du.$$

We now aim at showing a tight bound for $|S_\Delta|^\beta \int_0^1 \sqrt{\log(1 + 2u^{-1/\beta})} du$ w.r.t. $|S_\Delta|$ with an appropriate choice of $\beta = \beta(|S_\Delta|)$. Indeed, let $\beta = \frac{1}{1 + (\log|S_\Delta|)_+} \in (0, 1]$. We start by proving the bound

$$I_\beta := \int_0^1 \sqrt{\log(1 + 2u^{-1/\beta})} du \leq \frac{4}{\sqrt{\beta}}, \quad \forall \beta \in (0, 1]. \quad (1.5.24)$$

By the change of variable $t = \beta \log(1 + 2u^{-1/\beta})$ (i.e. $u = \left(\frac{e^{t/\beta} - 1}{2}\right)^{-\beta}$), we get

$$I_\beta = \frac{2^\beta}{\sqrt{\beta}} \int_{\beta \log(3)}^\infty \frac{\sqrt{t} e^{t/\beta}}{(e^{t/\beta} - 1)^{\beta+1}} dt = \frac{2^\beta}{\sqrt{\beta}} \int_{\beta \log(3)}^\infty \frac{\sqrt{t}}{e^t (1 - e^{-t/\beta})^{\beta+1}} dt.$$

Using the fact that $1 - e^{-t/\beta} \geq \frac{2}{3}$ on $[\beta \log(3), +\infty)$, we arrive at

$$I_\beta \leq \frac{3^{\beta+1}}{2\sqrt{\beta}} \int_{\beta \log(3)}^\infty \sqrt{t} e^{-t} dt \stackrel{(*)}{\leq} \frac{9}{2\sqrt{\beta}} \Gamma\left(\frac{3}{2}\right) \leq \frac{4}{\sqrt{\beta}},$$

where the inequality $(*)$ is obtained by taking $\beta = 1$ in $3^{\beta+1}$ and $\beta = 0$ in the integral; hence (1.5.24) is proved. Now, using (1.5.24) with $\beta = \frac{1}{1 + (\log|S_\Delta|)_+}$ and its $|S_\Delta|^\beta = e^{\frac{\log|S_\Delta|}{1 + (\log|S_\Delta|)_+}}$ implication, we get

$$|S_\Delta|^\beta \int_0^1 \sqrt{\log(1 + 2u^{-1/\beta})} du \leq 4e^{\frac{\log|S_\Delta|}{1 + (\log|S_\Delta|)_+}} \sqrt{1 + (\log|S_\Delta|)_+} \leq 4e\sqrt{1 + (\log|S_\Delta|)_+},$$

and therefore

$$B \leq \frac{8eC_D \sqrt{d} \sqrt{\mathbb{E}_{\omega \sim \Lambda}[L^2(\omega)]} \sqrt{1 + (\log|S_\Delta|)_+}}{\sqrt{M}}.$$

- **The term σ^2 :** It is defined in Theorem 1.1. Since the variance is bounded by the second moment, $\mathbb{E}[g_{\mathbf{z}}^2(\omega_m)] \leq \mathbb{E}[f_{\mathbf{z}}^2(\omega_m)]$. Furthermore, since $(\omega_m)_{m=1}^M$ are i.i.d., the previous expectation can be bounded by $\mathbb{E}[\|\omega\|_2^{2|\mathbf{p}+\mathbf{q}|}]$ using (1.5.1). As a result, we get

$$\sigma^2 \leq \mathbb{E}_{\omega \sim \Lambda} \left[\|\omega\|_2^{2|\mathbf{p}+\mathbf{q}|} \right].$$

- **The term $\max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} \|g_{\mathbf{z}}(\omega_m)\|_{\Psi_\gamma}$ with $\gamma = \alpha/n \leq 1$ and $n = |\mathbf{p} + \mathbf{q}| + \beta$:** It appears in the definition of c (in Theorem 1.1). In view of the bound (1.5.2) which is uniform in \mathbf{z} and using property (iv) of Section 1.4, we get

$$\max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} \|g_{\mathbf{z}}(\omega_m)\|_{\Psi_\gamma} = \max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} \| |g_{\mathbf{z}}(\omega_m)| \|_{\Psi_\gamma} \leq \left\| \|\omega\|_2^{|\mathbf{p}+\mathbf{q}|} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|} \right] \right\|_{\Psi_{\alpha/n}}.$$

- **The term $\left\| \max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\omega_m)| \right\|_{\Psi_\gamma}$:** It shows up in the exponential bound (1.3.2). We invoke the maximal inequality for the γ -Orlicz norm (item (iii) of Section 1.4) with the previous estimate to obtain

$$\left\| \max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\omega_m)| \right\|_{\Psi_\gamma} \leq \left[\frac{\log(1 + M)}{\log(3/2)} \right]^{n/\alpha} \left\| \|\omega\|_2^{|\mathbf{p}+\mathbf{q}|} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|} \right] \right\|_{\Psi_{\alpha/n}}.$$

- **The term** $\mathbb{E} [\max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\boldsymbol{\omega}_m)|]$: It appears in the definition of c . Using properties (iii) and (v) of Section 1.4 and the convexification of Ψ_γ , we directly get

$$\mathbb{E} \left[\max_{1 \leq m \leq M} \sup_{\mathbf{z} \in S_\Delta} |g_{\mathbf{z}}(\boldsymbol{\omega}_m)| \right] \leq \left(\Psi_{\alpha/n}^{(l)} \right)^{-1} (1) \left[\frac{\log(1+M)}{\log(3/2)} \right]^{n/\alpha} \left\| \|\boldsymbol{\omega}\|_2^{|\mathbf{p}+\mathbf{q}|} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|} \right] \right\|_{\Psi_{\alpha/n}}.$$

Collecting the different bounds we obtain (1.3.4) by setting

$$C_{\text{RFF}}(n) := \max \left(8eC_D \sqrt{d} \sqrt{\mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} [L^2(\boldsymbol{\omega})]}, \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} \left[\|\boldsymbol{\omega}\|_2^{2|\mathbf{p}+\mathbf{q}|} \right], \right. \\ \left. 1 \vee \left(\left(\Psi_{\alpha/n}^{(l)} \right)^{-1} (1) [\log(3/2)]^{-n/\alpha} \right) \left\| \|\boldsymbol{\omega}\|_2^{|\mathbf{p}+\mathbf{q}|} + \Lambda \left[\|\cdot\|_2^{|\mathbf{p}+\mathbf{q}|} \right] \right\|_{\Psi_{\alpha/n}} \right).$$

Long (but standard) computations show that $C_{\text{RFF}}(n)$ is uniformly bounded for $n \in [|\mathbf{p}+\mathbf{q}|, |\mathbf{p}+\mathbf{q}|+1]$, and thus we can set $C_{\text{RFF}} := \sup_{n \in [|\mathbf{p}+\mathbf{q}|, |\mathbf{p}+\mathbf{q}|+1]} C_{\text{RFF}}(n)$.

1.5.7 Proofs of Corollary 1.3.1 and 1.3.2

Corollary 1.3.1 with the existence of $\tilde{C} \in \mathbb{R}^+$ is a direct consequence of Theorem 1.1 combined with Remark 1(v), in particular because C_{RFF} does not depend on S_Δ and M , and K_γ can be bounded uniformly in $n \in [|\mathbf{p}+\mathbf{q}|, |\mathbf{p}+\mathbf{q}|+1]$. The Talagrand constant C_γ is uniformly bounded w.r.t. γ provided that γ is bounded away from 0, see the proof of [Tal89, Theorem 3].

Now let us prove Corollary 1.3.2; set $\varepsilon_M = \frac{(\sqrt{6\tilde{C}\vee\tilde{C}}) \sqrt{(1+[\log|(S_M)_\Delta]_+ \vee \log(1+M))}}{\sqrt{M}}$. Observe that

- (i) ε_M satisfies the lower bound requirement on ε in Corollary 1.3.1;
- (ii) by assumption $\varepsilon_M \rightarrow 0$ as $M \rightarrow 0$ by using that $|S_\Delta| \leq 2|S|$;
- (iii) therefore

$$1 + \varepsilon_M \left[\log \left(\frac{\tilde{C}}{\varepsilon_M} \right) \vee \log(1+M) \right]^{1/\gamma} \leq 1 + \varepsilon_M \left(\left[\log \left(\frac{\tilde{C}}{\varepsilon_M} \right) \right]^{1/\gamma} + [\log(1+M)]^{1/\gamma} \right) \\ \leq 2 + \varepsilon_M [\log(1+M)]^{1/\gamma}$$

for M large enough;

$$(iv) \varepsilon_M \geq \frac{(\sqrt{6\tilde{C}\vee\tilde{C}}) \sqrt{\log(1+M)}}{\sqrt{M}}.$$

As a consequence of (iii) and (iv), setting $\delta_M := \frac{1+([\log|(S_M)_\Delta]_+ \vee \log(1+M))}{\log(1+M)} \vee 1$, we get (for M large enough)

$$\frac{M \varepsilon_M^2}{\tilde{C} \left(1 + \varepsilon_M \left[\log \left(\frac{\tilde{C}}{\varepsilon_M} \right) \vee \log(1+M) \right]^{1/\gamma} \right)} \geq \frac{6\tilde{C} \log(1+M) \delta_M}{\tilde{C} \left[2 + \frac{(\sqrt{6\tilde{C}\vee\tilde{C}}) \sqrt{\delta_M} [\log(1+M)]^{1/2+1/\gamma}}{\sqrt{M}} \right]} \\ = 6 \log(1+M) \frac{\delta_M}{2 + z_M \sqrt{\delta_M}},$$

where $z_M = \frac{(\sqrt{6\tilde{C}\vee\tilde{C}}) [\log(1+M)]^{1/2+1/\gamma}}{\sqrt{M}} \xrightarrow{M \rightarrow \infty} 0$. Since the function $\delta \in \mathbb{R}^+ \mapsto \frac{\delta}{2+z_M \sqrt{\delta}}$ is increasing and $\delta_M \geq 1$, we get (for M large enough)

$$\frac{M \varepsilon_M^2}{\tilde{C} \left(1 + \varepsilon_M \left[\log \left(\frac{\tilde{C}}{\varepsilon_M} \right) \vee \log(1+M) \right]^{1/\gamma} \right)} \geq 6 \log(1+M) \frac{1}{3}.$$

On the other hand, using (iv), we easily get $\frac{(M\epsilon_M)^\gamma}{\tilde{C} \log(1+M)} \geq \frac{\left[\left(\sqrt{6\tilde{C}\vee\tilde{C}} \sqrt{\log(1+M)\sqrt{M}} \right)^\gamma \right]}{\tilde{C} \log(1+M)} \geq 2 \log(1+M)$ for M large enough.

To sum up, in view of (1.3.5), we have proved (still for large enough M)

$$\Lambda^M \left(\left\| \widehat{\partial^{\mathbf{p}, \mathbf{q}_k}} - \partial^{\mathbf{p}, \mathbf{q}_k} \right\|_{S_M} \geq \epsilon_M \right) \leq \frac{2}{(1+M)^2} + \frac{1}{(1+M)^2}$$

and by the Borell-Cantelli lemma, we conclude to the a.s. convergence (1.3.6).

APPENDICES

In this section we give additional technical details (Section 1.A), the definition of special functions (Section 1.B), and external statements used in the proofs (Section 1.C).

1.A Additional proofs

1.A.1 $\Psi_\gamma^{(l)}$, the convexification of Ψ_γ

In the proof of Theorem 1.1 an integral control with *convex* Ψ (see Section 1.4(v)) is beneficial/applied. However, Ψ_γ is not convex for $\gamma \in (0, 1)$. To handle this issue, we convexify $\Psi_\gamma(x) = e^{x^\gamma} - 1$ in case of $\gamma \in (0, 1)$ for 'small' values of the argument.⁸

- By computing the derivatives of Ψ_γ we get that it is convex iff $x \geq x_\gamma := \left(\frac{1-\gamma}{\gamma}\right)^{\frac{1}{\gamma}}$. Indeed,

$$\begin{aligned}\Psi_\gamma'(x) &= \gamma x^{\gamma-1} e^{x^\gamma}, \\ \Psi_\gamma''(x) &= \gamma e^{x^\gamma} [(\gamma-1)x^{\gamma-2} + x^{\gamma-1}\gamma x^{\gamma-1}] \Rightarrow \\ \Psi_\gamma''(x) &= 0 \Leftrightarrow x = x_\gamma, \quad \Psi_\gamma''(x) > 0 \Leftrightarrow x > x_\gamma, \quad \Psi_\gamma''(x) < 0 \Leftrightarrow x < x_\gamma.\end{aligned}$$

- We also have to make sure that $\Psi_\gamma^{(l)}$, constructed as the line connecting $(0, 0)$ with $(x, \Psi_\gamma(x))$ glued to $\Psi_\gamma|_{[x, \infty)}$, gives a convex function, for a suitable choice of x . A geometric argument shows that it is enough to choose $x \geq x_\gamma (> 0)$ such that

$$\frac{\Psi_\gamma(x)}{x} \leq \Psi_\gamma'(x) \Leftrightarrow e^{x^\gamma} - 1 \leq \gamma x^\gamma e^{x^\gamma}.$$

Since the r.h.s. is higher order than the l.h.s., the requirement can be satisfied for large enough x ; we can choose

$$\tilde{x}_\gamma := \inf \{x \geq x_\gamma : e^{x^\gamma} - 1 \leq \gamma x^\gamma e^{x^\gamma}\},$$

and define

$$\Psi_\gamma^{(l)}(x) := \begin{cases} \frac{\Psi_\gamma(\tilde{x}_\gamma)}{\tilde{x}_\gamma} x & \text{if } x \in [0, \tilde{x}_\gamma), \\ \Psi_\gamma(x) & \text{if } x \in [\tilde{x}_\gamma, \infty). \end{cases}$$

Notice that by construction $\Psi_\gamma^{(l)} \leq \Psi_\gamma$.

⁸For $\gamma = 1$, $\Psi_\gamma^{(l)} = \Psi_\gamma$.

1.A.2 Proof of Remark 2(ii)

$\omega_i \in L_{\Psi_{\alpha_i}}$ means that $\mathbb{E}_{\omega_i \sim \Lambda_i} [e^{s_i |\omega_i|^{\alpha_i}}] < \infty$ for some $s_i \in \mathbb{R}^+$ ($\forall i \in [d]$). Let $\alpha = \min_{i \in [d]} \alpha_i$ and $\|\omega\|_\alpha := \left(\sum_{i \in [d]} |\omega_i|^\alpha \right)^{\frac{1}{\alpha}}$. Then $\|\omega\|_2 \leq \sqrt{d} \sup_{i \in [d]} |\omega_i| \leq \sqrt{d} \|\omega\|_\alpha$ ($\forall \omega \in \mathbb{R}^d$). Notice that $|\omega_i|^\alpha \leq |\omega_i|^{\alpha_i}$ if $|\omega_i| \geq 1$ and $|\omega_i|^\alpha \leq 1$ otherwise, i.e. we have $|\omega_i|^\alpha \leq |\omega_i|^{\alpha_i} + 1$ for any $\omega_i \in \mathbb{R}$. This means that taking $s = \min_{i \in [d]} s_i$ and $\tilde{s} := \frac{s}{d^{\alpha/2}} > 0$ gives

$$\begin{aligned} \|\omega\|_2^\alpha &\leq d^{\alpha/2} \|\omega\|_\alpha^\alpha = d^{\alpha/2} \sum_{i \in [d]} |\omega_i|^\alpha \leq d^{\alpha/2} \sum_{i \in [d]} (|\omega_i|^{\alpha_i} + 1), \\ \mathbb{E}_\omega \left[e^{\tilde{s} \|\omega\|_2^\alpha} \right] &\leq \mathbb{E}_\omega \left[e^{s \sum_{i \in [d]} (|\omega_i|^{\alpha_i} + 1)} \right] \\ &\stackrel{(*)}{=} \prod_{i \in [d]} \left(\mathbb{E}_{\omega_i} \left[e^{s |\omega_i|^{\alpha_i}} \right] e^s \right) \leq \prod_{i \in [d]} \left(\mathbb{E}_{\omega_i} \left[e^{s_i |\omega_i|^{\alpha_i}} \right] e^s \right) < \infty, \end{aligned}$$

where we used the independence of ω_i -s in (*). We got that $\mathbb{E}_{\omega \sim \Lambda} \left[e^{\tilde{s} \|\omega\|_2^\alpha} \right] < \infty$ which implies that $\omega \in L_{\Psi_\alpha}$.

1.A.3 Proof of the properties in Section 1.4 about the Orlicz norm

- **Properties (i)-(iv)**: These properties are well-known and directly follow from the definition of the Orlicz norm.
- **Property (v)**: The case $\|X\|_\Psi = 0$ gives a trivial inequality and can be discarded. Since Ψ is bounded from below by an increasing affine function, $X \in L_\Psi$ implies that X is integrable. Combining (i) with Jensen's inequality gives $\Psi \left(\frac{\mathbb{E}[\|X\|_2]}{\|X\|_\Psi} \right) \leq \mathbb{E} \left[\Psi \left(\frac{\|X\|_2}{\|X\|_\Psi} \right) \right] \leq 1$, and the result follows.
- **Property (vi)**: It is well-known that the usual triangle inequality holds for $\alpha \geq 1$. We now focus on the case $\alpha \in (0, 1]$. Set $c := (\|X\|_{\Psi_\alpha}^\alpha + \|X'\|_{\Psi_\alpha}^\alpha)^{1/\alpha}$, $p := \frac{c^\alpha}{\|X\|_{\Psi_\alpha}^\alpha}$ and $q := \frac{c^\alpha}{\|X'\|_{\Psi_\alpha}^\alpha}$, and notice that $\frac{1}{p} + \frac{1}{q} = 1$. Then, combining (1.5.7) with $\gamma = \alpha \in (0, 1)$ and the Hölder inequality with the conjugate exponents (p, q) yields

$$\begin{aligned} \limsup_{m \rightarrow +\infty} \mathbb{E} \left[e^{\left(\frac{m \wedge \|X + X'\|_2}{c} \right)^\alpha} \right] &\leq \limsup_{m \rightarrow +\infty} \mathbb{E} \left[e^{\left(\frac{m \wedge \|X\|_2}{c} \right)^\alpha} e^{\left(\frac{m \wedge \|X'\|_2}{c} \right)^\alpha} \right] \\ &\leq \left(\limsup_{m \rightarrow +\infty} \mathbb{E} \left[e^{m \wedge \frac{p \|X\|_2^\alpha}{c^\alpha}} \right] \right)^{1/p} \left(\limsup_{m \rightarrow +\infty} \mathbb{E} \left[e^{m \wedge \frac{q \|X'\|_2^\alpha}{c^\alpha}} \right] \right)^{1/q} \\ &\stackrel{\text{item (i)}}{\leq} 2^{1/p} 2^{1/q} = 2. \end{aligned}$$

Therefore, $X + X' \in L_{\Psi_\alpha}$ and $\|X + X'\|_{\Psi_\alpha} \leq c$ by the definition of the α -Orlicz norm. Applying (1.5.7) with $\gamma = 1/\alpha$ we get $c \leq 2^{\left(\frac{1}{\alpha}-1\right)_+} (\|X\|_{\Psi_\alpha} + \|X'\|_{\Psi_\alpha})$ and hence the claimed result is proved.

- **Property (vii)**: This is a direct consequence of the Markov inequality.
- **Property (viii)**: A similar statement was proved by [vW96, Lemma 2.2.2], but under the assumption that Ψ_α is convex (which holds only if $\alpha \geq 1$) and without explicit constant. Our statement is valid for any $\alpha > 0$ with explicit control.
 - **A first inequality**: Let $\alpha \in \mathbb{R}^+$. We claim that for any $x_0 > 0$ and any $x, y \geq 1$, we have

$$\Psi_\alpha \left(x_0^{1/\alpha} x \right) \Psi_\alpha \left(x_0^{1/\alpha} y \right) \leq \Psi_\alpha \left(x_0^{1/\alpha} \right) \Psi_\alpha \left(x_0^{1/\alpha} xy \right). \quad (1.A.1)$$

Because $\Psi_\alpha(x) = \Psi_1(x^\alpha)$ where $\Psi_1(x) =: \Psi(x) = e^x - 1$, the inequality for $\alpha = 1$ clearly implies those for all $\alpha > 0$. To prove the inequality for $\alpha = 1$, let x_0 and x be fixed, and

set $H(y) = \Psi(x_0)\Psi(x_0xy) - \Psi(x_0x)\Psi(x_0y)$. One has

$$\begin{aligned} H'(y) &= x_0x\Psi(x_0)e^{x_0xy} - x_0\Psi(x_0x)e^{x_0y} \\ &= x_0^2xe^{x_0}e^{x_0x} \left[\frac{\Psi(x_0)}{x_0e^{x_0}}e^{x_0x(y-1)} - \frac{\Psi(x_0x)}{x_0xe^{x_0x}}e^{x_0(y-1)} \right], \\ \frac{\Psi(x_0)}{x_0e^{x_0}} &= \frac{1 - e^{-x_0}}{x_0} = \int_0^1 e^{-ux_0} du \geq \int_0^1 e^{-ux_0x} du = \frac{\Psi(x_0x)}{x_0xe^{x_0x}}, \\ e^{x_0x(y-1)} &\geq e^{x_0(y-1)}, \end{aligned}$$

where we used $x_0 > 0$, $x, y \geq 1$ at the two last inequalities. This shows that $H'(y) \geq 0$, and since $H(1) = 0$ we have $H(y) \geq 0$ for any $y \geq 1$. Consequently, (1.A.1) is proved.

- **Final maximal inequality:** We follow the arguments of [vW96, Lemma 2.2.2] with slight modifications. The inequality (1.A.1) can be rewritten as

$$\Psi_\alpha(x) \leq \Psi_\alpha\left(x_0^{1/\alpha}\right) \Psi_\alpha\left(xy/x_0^{1/\alpha}\right) / \Psi_\alpha(y), \quad \forall x, y \geq x_0^{1/\alpha}. \quad (1.A.2)$$

Set $c = \max_{m \in [M]} \|X_m\|_{\Psi_\alpha} / x_0^{1/\alpha}$ and let $y \geq x_0^{1/\alpha}$.

- * If $\frac{\max_{m \in [M]} \|X_m\|_2}{cy} \leq x_0^{1/\alpha}$, then we have the crude bound $\Psi_\alpha\left(\frac{\max_{m \in [M]} \|X_m\|_2}{cy}\right) \leq \Psi_\alpha\left(x_0^{1/\alpha}\right) = \Psi(x_0)$.
- * If $\frac{\max_{m \in [M]} \|X_m\|_2}{cy} \geq x_0^{1/\alpha}$, then (1.A.2) yields

$$\begin{aligned} \Psi_\alpha\left(\frac{\max_{m \in [M]} \|X_m\|_2}{cy}\right) &\leq \Psi(x_0) \Psi_\alpha\left(\frac{\max_{m \in [M]} \|X_m\|_2}{\max_{m \in [M]} \|X_m\|_{\Psi_\alpha}}\right) / \Psi_\alpha(y) \\ &\leq \sum_{m \in [M]} \Psi(x_0) \Psi_\alpha(\|X_m\|_2 / \|X_m\|_{\Psi_\alpha}) / \Psi_\alpha(y). \end{aligned}$$

Consequently, in both cases we have

$$\Psi_\alpha\left(\frac{\max_{m \in [M]} \|X_m\|_2}{cy}\right) \leq \Psi(x_0) \left[1 + \sum_{m \in [M]} \Psi_\alpha(\|X_m\|_2 / \|X_m\|_{\Psi_\alpha}) / \Psi_\alpha(y) \right].$$

Taking expectation and using Property (i), we arrive at

$$\mathbb{E} \left[\Psi_\alpha\left(\frac{\max_{m \in [M]} \|X_m\|_2}{cy}\right) \right] \leq \Psi(x_0) \left[1 + \frac{M}{\Psi_\alpha(y)} \right].$$

Let us choose x_0 such that $\Psi(x_0) < 1$. In this case the choice $y = x_0^{1/\alpha} \vee \Psi_\alpha^{-1}\left(\frac{M}{1/\Psi(x_0)-1}\right)$ ensures that the above bound is valid and smaller than 1. Consequently, by the definition of α -Orlicz norm, we get $\|\max_{m \in [M]} \|X_m\|_2\|_{\Psi_\alpha} \leq cy$. The choice $x_0 = \log(3/2)$ satisfies the previous requirement: $\Psi(x_0) = \frac{1}{2} < 1$. In this case $y = [\log(3/2) \vee \log(1 + M)]^{1/\alpha} = [\log(1 + M)]^{1/\alpha}$ since $M \geq 1$. We have obtained the claimed Property (viii).

1.B Special functions

Name	Definition
Modified Bessel function of the first kind	$J_a(x) = \sum_{n \in \mathbb{N}} \frac{1}{n! \Gamma(n+a+1)} \left(\frac{x}{2}\right)^{2n+a}$
Modified Bessel function of the second kind	$K_a(x) = \frac{\pi}{2} \frac{J_{-a}(x) - J_a(x)}{\sin(a\pi)}$
Fox-Wright generalized hyperbolic function	${}_1\Psi_1((a, A); (b, B); x) = \sum_{n \in \mathbb{N}} \frac{\Gamma(a+An)}{\Gamma(b+Bn)} \frac{x^n}{n!}$
(Imaginary) error function	$\operatorname{erfi}(x) = \sum_{n \in \mathbb{N}} \frac{2}{\sqrt{\pi}} \frac{x^{2n+1}}{n!(2n+1)}$
Kummer's confluent hypergeometric function	${}_1F_1(a; b; x) = \sum_{n \in \mathbb{N}} \frac{a^{(n)} x^n}{b^{(n)} n!}$
Polylogarithm function	$\operatorname{Li}_a(x) = \sum_{n \in \mathbb{Z}^+} \frac{x^n}{n^a}$
Ordinary hyperbolic function	${}_2F_1(a, b; c; z) = \sum_{n \in \mathbb{N}} \frac{a^{(n)} b^{(n)} z^n}{c^{(n)} n!}$

Table 1.4 – Definition of special functions. $J_a(x), K_a(x)$: $x \in \mathbb{R}$ and a is non-integer; when a is an integer the limit is taken. ${}_1\Psi_1((a, A); (b, B); x)$: $a \in \mathbb{R}^+, b \in \mathbb{R}^+, x \in \mathbb{R}, A \in \mathbb{R}^+, B \in \mathbb{R}^+$ and $1+B > A$; ${}_1\Psi_1((a, 1); (b, 1); x) = \frac{\Gamma(a)}{\Gamma(b)} {}_1F_1(a; b; x)$. $\operatorname{erfi}(x)$: $x \in \mathbb{R}$. ${}_1F_1(a; b; x)$: $a \in \mathbb{R}^+, b \in \mathbb{R}^+, x \in \mathbb{R}$. $\operatorname{Li}_a(x)$: $a \in \mathbb{R}, x \in \mathbb{R}, |x| < 1$. ${}_2F_1(a, b; c; z)$: $a \in \mathbb{C}, b \in \mathbb{C}, c \in \mathbb{C} \setminus \mathbb{Z}^{\leq 0}, z \in \mathbb{C}$ and $|z| < 1$; for $|z| \geq 1$ its analytical continuation is taken. For $n \in \mathbb{N}$, $a^{(n)}$ is the rising factorial of a defined as $a^{(n)} = \frac{\Gamma(a+n)}{\Gamma(a)}$ where $a \in \mathbb{C} \setminus \mathbb{Z}^{\leq 0}$ and $a+n \in \mathbb{C} \setminus \mathbb{Z}^{\leq 0}$.

1.C External statements

In this subsection we state external statements which were used to derive our results. Below B stands for a separable Banach space, $L_p(B)$ is the space of B -valued p -integrable functions. The norm $\|\cdot\|_{\Psi_\alpha}$ is defined analogously to \mathbb{R}^d by changing $\|\cdot\|_2$ to $\|\cdot\|_B$.

Theorem 1.1. (*Hoffman-Jorgensen inequality, [LT13, Proposition 6.8]*) Let $p > 0$, $M \in \mathbb{Z}^+$, $(Y_m)_{m \in [M]}$ be independent random variables in $L_p(B)$, $S_m := \sum_{j=1}^m Y_j$ for $m \in [M]$, $t_0 = \inf \{t > 0 : \mathbb{P}(\max_{1 \leq m \leq M} \|S_m\|_B > t)\} \leq (2 \times 4^p)^{-1}$. Then

$$\mathbb{E} \left[\max_{m \in [M]} \|S_m\|_B^p \right] \leq 2 \times 4^p \mathbb{E} \left[\max_{m \in [M]} \|Y_m\|_B^p \right] + 2(4t_0)^p.$$

Theorem 1.2. [*Tal89, Theorem 3*] Let $\gamma \in (0, 1]$. Then, there is a constant C_γ such that for all finite sequence $(Y_m)_{m \in [M]}$ of independent, mean zero, integrable random variables in $L_{\Psi_\gamma}(B)$, we have

$$\left\| \sum_{m \in [M]} Y_m \right\|_{\Psi_\gamma} \leq C_\gamma \left(\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\gamma} \right). \quad (1.C.1)$$

Theorem 1.3. (*Klein-Rio inequality for supremum of empirical process, [KR05, Theorems 1.1-1.2]*) Let $M \in \mathbb{Z}^+$, $c \in \mathbb{R}^+$, $(X_m)_{m \in [M]}$ be independent B -valued random variables, and \mathcal{F} a countable set of $\mathbf{f} := (f_1, \dots, f_M)$ measurable functions from B into $[-c, c]^M$ such that $\mathbb{E}[f_m(X_m)] = 0$ for all $m \in [M]$. Define $Z := \sup_{\mathbf{f} \in \mathcal{F}} \frac{1}{M} \sum_{m \in [M]} f_m(X_m)$, $\sigma^2 := \frac{1}{M} \sup_{\mathbf{f} \in \mathcal{F}} \mathbb{E} \left[\sum_{m \in [M]} f_m^2(X_m) \right]$. Then, for any $t \geq 0$ the following right and left-hand sided deviation inequalities hold

$$\mathbb{P}(Z - \mathbb{E}[Z] \geq t) \leq e^{-\frac{Mt^2}{2(\sigma^2 + 2c \mathbb{E}[Z]) + 3ct}}, \quad \mathbb{P}(Z - \mathbb{E}[Z] \leq -t) \leq e^{-\frac{Mt^2}{2(\sigma^2 + 2c \mathbb{E}[Z]) + 2ct}}.$$

Theorem 1.4. (Dudley entropy integral bound)⁹ Let $\{Z_t : t \in T\}$ be a zero-mean separable stochastic process that is sub-Gaussian w.r.t. a pseudo-metric d on the indexing set T , in other words for every $\lambda \in \mathbb{R}$ $\mathbb{E} [e^{\lambda(Z_t - Z_s)}] \leq e^{\frac{\lambda^2 d(s,t)^2}{2}}$ ($\forall s, t \in T$). Then there exists a universal constant C_D such that

$$\mathbb{E} \left[\sup_{t \in T} Z_t \right] \leq C_D \int_0^\infty \sqrt{\log N(\varepsilon, d, T)} d\varepsilon, \quad (1.C.2)$$

where $N(\varepsilon, d, T)$ denotes the covering number.

Theorem 1.5. ([Alz97, Theorem 1])¹⁰ Let $\gamma \in (0, 1]$, $\beta_\gamma := \Gamma\left(1 + \frac{1}{\gamma}\right)^{-\gamma}$, $x \in \mathbb{R}^{\geq 0}$, $I_\gamma(x) := \int_0^x e^{-t^\gamma} dt$. Then $(1 - e^{-\beta_\gamma x^\gamma})^{\frac{1}{\gamma}} \leq \frac{I_\gamma(x)}{\Gamma(1+1/\gamma)} \leq (1 - e^{-x^\gamma})^{\frac{1}{\gamma}}$.

⁹See [vW96, Corollary 2.2.8] for a general statement. Regarding the numerical value of C_D , [van16, Corollary 5.25] proves that one can take $C_D = 12$ whereas [Bar13, Lecture 14] suggests a slightly smaller constant $C_D = 8\sqrt{2}$.

¹⁰The statement here follows by taking the limit of the cited result at $\gamma = 1$ and $x = 0$.

CHAPTER 2

CONCENTRATION INEQUALITIES FOR β -HEAVY TAILED RANDOM VARIABLES

Note. This chapter corresponds to an article written with E. Gobet and W. Liu. It had been submitted.

Abstract. We establish a new concentration-of-measure inequality for the sum of independent random variables with β -heavy tail. This includes exponential of Gaussian distributions (a.k.a. log-normal distributions), or exponential of Weibull distributions, among others. These distributions have finite polynomial moments at any order but may not have finite α -exponential moments. We exhibit a Orlicz norm adapted to this setting of β -heavy tails, we prove a new Talagrand inequality for the sum and a new maximal inequality. As consequence, a bound on the deviation probability of the sum from its mean is obtained, as well as a bound on uniform deviation probability.

2.1 Introduction

Understanding how sample statistical fluctuations impact prediction errors is crucial in learning algorithms. Typically, we are interested in bounding the probability that a sum of random variables exceeds a certain threshold, essentially in quantifying the deviation of the sum from its expectation. In other words, we aim at analyzing how fast the sum concentrates around its expectation. Take the notation $[M]$ for all integers from 1 to M included. For independent and centered random variables $(Y_m)_{m \in [M]}$ taking value in a Banach space $(B, \|\cdot\|_B)$, the quantity of interest takes the form $\mathbb{P}\left(\left\|\sum_{m \in [M]} Y_m\right\|_B > \epsilon\right) \leq f(\epsilon, M)$ for the most explicit and tightest possible function f . The bounded, sub-Gaussian or the sub-exponential random variables have been largely covered by the literature (for example, via Bennett and Bernstein inequalities - see [BLM13] for an extensive review of main concentration inequalities techniques), as well as the case of alpha-exponential tails [CGS20] (Chapter 1: random variables Y s.t. there exists $\alpha > 0, c > 0$, such that $\mathbb{E}\left[\exp\left(\frac{\|Y\|_B^\alpha}{c}\right)\right] < \infty$). The fat-tailed case, for which the moment generating function does not exist but some polynomial moments exist, can be tackled for example via Burkholder or Fuk-Nagaev type of inequalities [Rio17, Mar17]. These inequalities are based on the existence and on the bounding of polynomial moments of the random variables. In this chapter, we focus on the heavy-tailed random variables case, in the limit case when no α -exponential moment is finite but every polynomial moment exist.

Orlicz norm. Orlicz norm [KR61] provides a nice tool to study the statistical fluctuations of an estimator for a given family of distributions. Consider an Orlicz function $\Psi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, that is a continuous non-decreasing function, vanishing in zero and with $\lim_{x \rightarrow +\infty} \Psi(x) = +\infty$, and define the Ψ -Orlicz norm of the B -valued random variable Y by

$$\|Y\|_\Psi := \inf \left\{ c > 0 : \mathbb{E} \left[\Psi \left(\frac{\|Y\|_B}{c} \right) \right] \leq 1 \right\}. \quad (2.1.1)$$

With the additional property that Ψ is convex, Orlicz functions are commonly referred to as "Young functions" (or "N-functions" as in [KR61]). Van de Geer and Lederer [vdGL13] exhibit in their work a "Bernstein-Orlicz" norm (the "(L)-Bernstein-Orlicz" norm) adapted to sub-Gaussian and sub-exponential tails and provide deviation inequalities for suprema of functions of random variables [vdGL13, Theorem 8]. The (L)-Bernstein-Orlicz norm is the Ψ_L -Orlicz norm with

$$\Psi_L(z) = \exp \left[\left(\frac{\sqrt{1 + 2Lz} - 1}{L} \right)^2 \right] - 1.$$

Clearly $\|Y\|_{\Psi_L} < \infty$ implies the existence of exponential moment. As shown in Wellner [Wel17], it is possible to generalize these results to any Orlicz function $\Psi(x) = e^{h(x)} - 1$ with h convex. It requires again the existence of exponential moment which is not our framework. We would like to go beyond and do not assume any α -exponential moment.

As a new Orlicz function able to handle heavy-tail situations, we will consider:

$$\Psi_\beta^{\text{HT}}(x) := \exp\left((\ln(x+1))^\beta\right) - 1, \quad x \geq 0, \quad (2.1.2)$$

for a parameter $\beta > 1$. We say that Y is β -heavy tailed if there exists a $c > 0$ s.t.

$$\mathbb{E} \left[\Psi_\beta^{\text{HT}} \left(\frac{\|Y\|_B}{c} \right) \right] < \infty.$$

Typically, we aim at encompassing situations like $Y = \exp\left(|G|^{\frac{2}{\beta}}\right)$ where G is a Gaussian random variable; the case $\beta = 2$ corresponds to log-normal tails. See Section 2.2.2 for various examples.

Observe that when (2.1.1) is finite with $\Psi = \Psi_\beta^{\text{HT}}$, Y has finite polynomial moment of order p for any $p > 0$, but may not have α -exponential moments. Besides, our β -heavy tailed setting is closely related to *long-tail* modelling¹, which is used for instance in queuing applications [Asm03, Chapter 10].

Deviation inequalities for sum via Talagrand and Markov inequalities. What we call Talagrand inequality is an inequality of type:

$$\left\| \sum_{m \in [M]} Y_m \right\|_\Psi \leq C_\Psi \left(\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \right). \quad (2.1.3)$$

Talagrand [Tal89, Theorem 3] showed that this inequality is satisfied with $\Psi_\alpha(x) := e^{x^\alpha} - 1$. For the sake of presentation, let us consider i.i.d. $(Y_m)_{m \in [M]}$. The first term is then $\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} \leq \mathcal{O}(\sqrt{M})$ by Bukholder inequality when $B \subseteq \mathbb{R}$ or the more general inequality of [Pis16, Proposition 4.35] when B is a Hilbert space or a Banach space of type 2.

When the maximal inequality is satisfied, that is under the form of [vW96, Lemma 2.2.2.], the second term is bounded by

$$\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \leq K_\Psi \Psi^{-1}(M) \max_{m \in [M]} \|Y_m\|_\Psi.$$

¹typically $S(x) := \mathbb{P}(\|Y\|_B > x) = \exp(-(\ln(1+x))^\beta)$ for which $\lim_{x \rightarrow +\infty} S(x+t)/S(x) = 1$ for any $t > 0$.

Hence, for any $\epsilon > 0$, denoting $X := \frac{1}{M} \left\| \sum_{m \in [M]} Y_m \right\|_B$, thanks to the Markov inequality, the Talagrand inequality and the two previous norm controls, we get

$$\begin{aligned} \mathbb{P}(X \geq \epsilon) &\stackrel{\text{Sect. 2.2.1-(iii)}}{\leq} \frac{2}{1 + \Psi(\epsilon/\|X\|_\Psi)} = 2 \left(1 + \Psi \left(\frac{\epsilon M}{\left\| \sum_{m \in [M]} Y_m \right\|_\Psi} \right) \right)^{-1} \\ &\leq 2 \left(1 + \Psi \left(\frac{\epsilon M}{C'_\Psi(\Psi^{-1}(M) + \sqrt{M})} \right) \right)^{-1}. \end{aligned} \quad (2.1.4)$$

In particular, for $\Psi = \Psi_\alpha$, the above inequality simplifies to:

$$\mathbb{P}(X \geq \epsilon) \leq 2 \exp \left(-C'_\alpha \left(\epsilon \sqrt{M} \right)^\alpha \right).$$

It is then possible to extend this type of inequality to suprema of functions as done in Chapter 1, in the spirit of [Ada08]. In any case, a key element to derive these concentration inequalities is the Talagrand inequality (2.1.3).

Contribution. The purpose of this chapter is mainly to establish the Talagrand inequality for $\Psi = \Psi_\beta^{\text{HT}}$, to tackle β -heavy tailed random variables as a difference with previous contributions available in the literature, and to derive some ready-to-use consequences. Note that this particular Orlicz function (2.1.2) is not at all part of the general result established by Talagrand [Tal89, Proposition 12], which states that the inequality (2.1.3) holds for Orlicz function of the form $\Psi(x) := e^{x\zeta(x)}$ with ζ *non-decreasing* for x large enough and satisfying $\limsup_{u \rightarrow +\infty} \frac{\zeta(e^u)}{\zeta(u)} < +\infty$; indeed, in our setting, one easily checks that $\zeta(x) = x^{-1} \ln \left(\Psi_\beta^{\text{HT}}(x) \right) = x^{-1} (\ln(\exp((\ln(x+1))^\beta) - 1))$ is *decreasing* for x large.

Outline. In Section 2.2, we recall the motivating example and define the adapted Orlicz function. Then we state our main results: Talagrand inequality (Theorem 2.2.2), maximal inequality (Theorem 2.2.3), pointwise and uniform deviation estimates (Corollary 2.2.4 and Theorem 2.2.5). Section 2.3 is devoted to the proofs. In all these results, some universal constants appear: we do not investigate the question of having the best possible constants.

2.2 Motivating examples and main results

In this section we recall the properties of $\|\cdot\|_\Psi$ which hold independently of the convexity of Ψ . We then give examples of distributions satisfying a β -heavy-tailness property and we present our main results on sum and supremum of sum of empirical processes with β -heavy tails.

2.2.1 Orlicz norm properties

Although $\|\cdot\|_\Psi$ defined in (2.1.1) may not satisfy in general the triangle inequality, we keep calling it Orlicz norm for the sake of simplicity. For a given Banach space $(B, \|\cdot\|_B)$ over the field \mathbb{R} , we denote $L_\Psi(B) := \{Y : \Omega \rightarrow B \text{ s.t. } \|Y\|_\Psi < +\infty\}$ the set of B -valued random variables with finite Ψ -Orlicz norm. For self-containedness we summarize a few well-known properties of the $\|\cdot\|_\Psi$ norm, for a given Orlicz function Ψ , which hold independently of the convexity of Ψ (unless explicitly required). See [KR61], or in Chapter 1 Section 1.4.

- (i) Normalization: If $Y \in L_\Psi(B)$ then $\mathbb{E} \left[\Psi \left(\frac{\|Y\|_B}{\|Y\|_\Psi} \right) \right] \leq 1$.
- (ii) Homogeneity: If $Y \in L_\Psi(B)$ and $c \in \mathbb{R}$ then $cY \in L_\Psi(B)$ and $\|cY\|_\Psi = |c| \|Y\|_\Psi$.

(iii) Deviation inequality: If $Y \in L_\Psi(B)$ then $\mathbb{P}(\|Y\|_B \geq c) \leq \frac{2}{\Psi(c/\|Y\|_\Psi)+1}$ for any $c \geq 0$.

(iv) If Ψ is convex, $\|\cdot\|_\Psi$ satisfies to the triangle inequality.

2.2.2 Motivating examples of heavy-tailed distributions and adapted Orlicz norm

In this section we recall the motivating example of this study and we give various examples satisfying the β -heavy-tailness property.

2.2.2.1 Log-normal distribution

Let Y be a scalar random variable with log-normal distribution, i.e.

$$\ln(Y) \stackrel{d}{=} \mathcal{N}(\mu, \sigma^2),$$

with $\sigma > 0$. The distribution of Y admits the density

$$f_Y(y; \mu, \sigma) := \frac{1}{\sigma\sqrt{2\pi y}} e^{-\frac{(\ln y - \mu)^2}{2\sigma^2}} \mathbf{1}_{y>0}.$$

Let us investigate what kind of Orlicz function Ψ can be used to have $\|Y\|_\Psi < \infty$. In particular, we search for $\Psi(x) = \exp(\xi(x)) - 1$ such that ξ is non-decreasing, $\xi(0) = 0$ and $\lim_{x \rightarrow +\infty} \xi(x) = +\infty$ in order to ensure that $\Psi(0) = 0$ and $\lim_{x \rightarrow +\infty} \Psi(x) = +\infty$. Let $c > 0$, observe that

$$\mathbb{E} \left[\exp \left\{ \left(\left| \xi \left(\frac{|Y|}{c} \right) \right| \right) \right\} \right] < \infty \implies \liminf_{x \rightarrow \infty} \xi \left(\frac{x}{c} \right) - \frac{(\ln x)^2}{2\sigma^2} = -\infty. \quad (2.2.1)$$

Consider the following functions for $\beta > 0$:

1. $\xi_\beta(x) = (\ln(x+1))^\beta$, $x \geq 0$. Note that the case $\beta \leq 1$ is not much interesting in our setting since it quantifies tails with finite expectation at most (fat tail cases).
2. $\xi_\beta(x) = (\ln(x+1))^\beta (\ln(\ln(x+1)+1))^\alpha$, $x \geq 0$, $\alpha \in \mathbb{R}$. This second case is a scale refinement of the first case. It is not studied here.

These functions satisfy the necessary condition (2.2.1) if $\beta < 2$ and for a large c ². Furthermore, since for any $c > 0$, $\xi_\beta\left(\frac{x}{c}\right) < \epsilon(\ln x - \mu)^2$ for any $\epsilon > 0$ for x large enough, $\mathbb{E} \left[\exp \left\{ \left(\left| \xi_\beta \left(\frac{|Y|}{c} \right) \right| \right) \right\} \right] < +\infty$.

2.2.2.2 Other distributions satisfying $\mathbb{E} \left[\exp \left\{ \left(\xi_\beta \left(\frac{\|Y\|_B}{c} \right) \right) \right\} \right] < +\infty$

The associated Orlicz function $\Psi_\beta^{\text{HT}}(x) := \exp(\xi_\beta(x)) - 1$ is adapted to other distributions than just the log-normal distribution. For any random variable X admitting finite α -exponential moment with $\alpha > 1$, then Y defined by $\ln(Y) = X$ will admit β -heavy tailed for any $1 < \beta < \alpha$. We refer to 1.2 in the previous chapter for an exhaustive list of distributions admitting α -exponential moments. Here are a few examples:

- The Generalized normal distribution with parameters $c \in \mathbb{R}$, $b > 0$, $\alpha > 0$ has a density $f(x) = c_f e^{-\frac{1}{2} \left(\frac{|x-c|}{b} \right)^\alpha}$ up to a positive normalization constant c_f : it clearly admits a finite α -exponential moment. Hence $Y = \exp(X)$ where X has density f hence admits β -heavy tails for $\beta < \alpha$.

² $\beta = 2$ is possible under restriction on σ : if $\sigma < \frac{1}{\sqrt{2}}$, then $\liminf_{x \rightarrow \infty} \xi_2(x) - \frac{(\ln x)^2}{2\sigma^2} = -\infty$.

- The Skew normal distribution with parameters $b \in \mathbb{R}$, $c \in \mathbb{R}$, $v > 0$ has a density $f(x) = c_f e^{-\frac{(x-c)^2}{2v}} \Phi\left(\frac{b(x-c)}{\sqrt{v}}\right)$, where Φ denotes the standard Gaussian cumulative distribution function and c_f is a positive normalization constant: it admits 2-exponential moment. If X has this density, then $Y = \exp(X)$ has β -heavy tails for $\beta < 2$.
- The Weibull distribution with parameters $\lambda > 0, k > 0$ has a density $f(x) = c_f x^{k-1} e^{-\left(\frac{x}{\lambda}\right)^k} \mathbf{1}_{x \geq 0}$ up to a positive normalization constant c_f : it has finite k -exponential moment. Consequently, $Y = \exp(X)$ where X has the density as above, admits β -heavy tails for $\beta < k$. Such distributions are used, for instance, for earthquake magnitude modelling [HR99].

2.2.3 Ψ_β^{HT} -Orlicz norm: properties and inequalities

We state different properties of the Orlicz function to be used for β -heavy tailed distribution. The proof is postponed to Section 2.3.4.

Proposition 2.2.1. *For $\beta > 0$ define $\Psi_\beta^{\text{HT}} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ by*

$$\Psi_\beta^{\text{HT}}(x) := \exp(\xi_\beta(x)) - 1 \quad \text{with} \quad \xi_\beta(x) := (\ln(1+x))^\beta, \quad x \geq 0.$$

The following properties hold:

1. The application $\beta \mapsto \Psi_\beta^{\text{HT}}$ defines a group isomorphism between $((0, +\infty), \times)$ and $((\Psi_\beta^{\text{HT}} : \beta > 0), \circ)$, and in particular, $(\Psi_\beta^{\text{HT}})^{-1} = \Psi_{1/\beta}^{\text{HT}}$.
2. For $\beta > 0$, Ψ_β^{HT} is an Orlicz function.
3. For $\beta > 1$, Ψ_β^{HT} is convex.
4. For $\beta > 1$ (resp. $\beta < 1$), the limit as $x \rightarrow +\infty$ of $\Psi_\beta^{\text{HT}}(x)/x^k$ equals to $+\infty$ (resp. 0), for any $k > 0$.

As a consequence, the associated Ψ_β^{HT} -Orlicz norm satisfies to the triangle inequality for $\beta > 1$.

Hereafter, we mostly restrict the results to the more interesting case $\beta > 1$. Let us start with the Talagrand inequality (2.1.3) for the Ψ_β^{HT} -Orlicz norm.

Theorem 2.2.2 (Talagrand type inequality). *Let $\beta \in (1, +\infty)$. Then there is a universal constant $\mathcal{K}_{\beta, (2.2.2)}$ s.t. for all independent, mean zero, random variables sequence $(Y_m)_{m \in [M]}$ with $Y_m \in L_{\Psi_\beta^{\text{HT}}}(B)$ for all $m \in [M]$, we have*

$$\left\| \sum_{m \in [M]} Y_m \right\|_{\Psi_\beta^{\text{HT}}} \leq \mathcal{K}_{\beta, (2.2.2)} \left(\left\| \sum_{m \in [M]} Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} \right). \quad (2.2.2)$$

We also establish that the general maximal inequality [vW96, Lemma 2.2.2.] (recalled in Lemma 2.3.7) holds for the Ψ_β^{HT} function:

Theorem 2.2.3 (A Ψ_β^{HT} maximal inequality). *Let $\beta \in (1, +\infty)$. Then there exists a universal constant $\mathcal{C}_{\beta, (2.2.3)}$ s.t. for any random variables Y_1, \dots, Y_M in $L_{\Psi_\beta^{\text{HT}}}(B)$,*

$$\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} \leq \mathcal{C}_{\beta, (2.2.3)} (\Psi_\beta^{\text{HT}})^{-1}(M) \max_{m \in [M]} \|Y_m\|_{\Psi_\beta^{\text{HT}}}. \quad (2.2.3)$$

Recall that $(\Psi_\beta^{\text{HT}})^{-1}(M) = \Psi_{1/\beta}^{\text{HT}}(M)$. As a consequence of the Talagrand inequality (2.2.2) and the maximal inequality (2.2.3), by following the same steps as described in (2.1.4), we can derive the following concentration inequality:

Corollary 2.2.4 (A concentration inequality for sum of independent β -heavy tailed random variables). *Let $\beta \in (1, +\infty)$. Assume that B is an Hilbert space or a Banach space of type 2. Then for any Y_1, \dots, Y_M independent and centered random variables in $L_{\Psi_\beta^{\text{HT}}}(B)$, for any $\epsilon > 0$,*

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{M} \left\| \sum_{m \in [M]} Y_m \right\|_B \geq \epsilon \right) \\ & \leq 2 \exp \left(- \ln \left(1 + \frac{\epsilon M}{\mathcal{K}_{\beta, (2.2.2)} \left(C(2)^{1/2} \mu_2 \sqrt{M} + \mathcal{C}_{\beta, (2.2.3)} \mu_{\Psi_\beta^{\text{HT}}} \Psi_{1/\beta}^{\text{HT}}(M) \right)} \right)^\beta \right), \end{aligned}$$

where $\mu_{\Psi_\beta^{\text{HT}}} := \max_{m \in [M]} \|Y_m\|_{\Psi_\beta^{\text{HT}}}$ and $\mu_2 := \max_{m \in [M]} \|Y_m\|_{L_2(B)}$, $C(2)$ denotes the universal constant in the Pisier inequality [Pis16, Proposition 4.35], $\mathcal{K}_{\beta, (2.2.2)}$ the Talagrand constant in (2.2.2) and $\mathcal{C}_{\beta, (2.2.3)}$ the maximal inequality constant in (2.2.3).

Recall that $\Psi_{1/\beta}^{\text{HT}}(M)$ goes to infinity slower than M^k (for $\beta > 1$, $k > 0$) (Proposition 2.2.1-(4)). Thus, when Y_1, \dots, Y_M are i.i.d. – implying that $\mu_{\Psi_\beta^{\text{HT}}}$ and μ_2 do not depend on M – the above upper bound takes the simple form

$$2 \exp \left(- \left(\ln \left(1 + K \epsilon \sqrt{M} \right) \right)^\beta \right),$$

for some universal constant $K > 0$ (depending on $\mu_{\Psi_\beta^{\text{HT}}}$ and μ_2).

In addition, the above pointwise estimate can be turned into a uniform deviation estimate. On the technical side, the strategy consists in splitting the deviation between truncated functions and their residuals. The residuals are handled using Hoffman-Jorgensen inequality [LT13, Proposition 6.8], following an initial idea from [Ada08], more deeply analysed in Chapter 1. The "truncated part" can be handled using Klein-Rio concentration bounds, together with the Dudley entropy integral bounds. For the latter which is related to the complexity of the space of functions and their related covering numbers, we choose to describe it using its Vapnik-Chervonenkis (VC) dimension (see [GKKW02, Theorem 9.4]). For alternative descriptions, see [vdG00, Sections 2.3 and 2.4] and [NP07]; adaptation of the following result to these other complexity descriptions is somehow direct and left to the reader.

Theorem 2.2.5 (A uniform concentration inequality for β -heavy tailed random variables). *Let $\beta \in (1, +\infty)$. Let (X_1, \dots, X_M) be independent random variables taking values in \mathbb{R}^d and let \mathcal{F} be a countably-generated class of functions $f : \mathbb{R}^d \mapsto \mathbb{R}$ with envelope $F(x) := \sup_{f \in \mathcal{F}} |f(x)|$, such that $F(X_m) \in L_{\Psi_\beta^{\text{HT}}}(\mathbb{R})$ for any $m \in [M]$. Set*

$$\begin{aligned} \mu_{\Psi_\beta^{\text{HT}}} &:= \max_{m \in [M], f \in \mathcal{F}} \|f(X_m)\|_{\Psi_\beta^{\text{HT}}}, \\ \bar{\mu}_{\Psi_\beta^{\text{HT}}} &:= \max_{m \in [M]} \|F(X_m)\|_{\Psi_\beta^{\text{HT}}}, \\ \mu_2 &:= \max_{m \in [M], f \in \mathcal{F}} \|f(X_m)\|_{L_2}. \end{aligned} \tag{2.2.4}$$

Assume that the Vapnik-Chervonenkis dimension $V_{\mathcal{F}^+}$ of $\mathcal{F}^+ := \{(x, t) \in \mathbb{R}^d \times \mathbb{R}, t \leq f(x)\}; f \in \mathcal{F}\}$ is finite. Then, there exist two universal constants K_1, K_2 (depending only on β) such that for any $\epsilon > 0$ satisfying the constraint

$$\epsilon \geq K_1 c \sqrt{\frac{V_{\mathcal{F}^+}}{M}} \tag{2.2.5}$$

with

$$c := \left(K_1 \Psi_{1/\beta}^{\text{HT}}(M) \bar{\mu}_{\Psi_{1/\beta}^{\text{HT}}} \right) \vee \left(\mu_{\Psi_{1/\beta}^{\text{HT}}} \left(\exp \left[\left(2 \ln_+ \left(K_1 \mu_{\Psi_{1/\beta}^{\text{HT}}} / \varepsilon \right) \right]^{1/\beta} \right] - 1 \right) \right), \quad (2.2.6)$$

$$\ln_+(x) := \max(\ln(x), 0),$$

we have

$$\begin{aligned} & \mathbb{P} \left(\sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{m \in [M]} (f(X_m) - \mathbb{E}[f(X_m)]) \geq \varepsilon \right) \\ & \leq 2 \exp \left(- \left(\ln \left(1 + \frac{M\varepsilon}{K_2 \bar{\mu}_{\Psi_{1/\beta}^{\text{HT}}} \Psi_{1/\beta}^{\text{HT}}(M)} \right) \right)^\beta \right) + \exp \left(- \frac{M\varepsilon^2}{K_2(\mu_2^2 + c\varepsilon)} \right). \end{aligned} \quad (2.2.7)$$

A similar bound holds for lower deviations, i.e. replacing the sup and $\geq \varepsilon$ by inf and $\leq -\varepsilon$: it is obtained by changing \mathcal{F} into $-\mathcal{F}$ in the bounds.

If \mathcal{F} is a finite-dimensional vector space, $V_{\mathcal{F}^+} \leq \dim(\mathcal{F}) + 1$ [GKKW02, Theorem 9.5].

For i.i.d. $(X_m)_m$, i.e. the μ -parameters (2.2.4) do not depend on M , both the condition (2.2.6) and the bound (2.2.7) take simple forms in terms of M (without focusing much on the best constants), which makes the application of Theorem 2.2.5 even more easily applicable.

- The bound (2.2.7) becomes $2 \exp \left(- \left(\ln \left(1 + \frac{M\varepsilon}{K \Psi_{1/\beta}^{\text{HT}}(M)} \right) \right)^\beta \right) + \exp \left(- \frac{M\varepsilon^2}{K(1+c\varepsilon)} \right)$ for a positive constant K depending on β and the μ -parameters.
- The equation (2.2.6) becomes simply

$$c := K_1 \Psi_{1/\beta}^{\text{HT}}(M), \quad (2.2.8)$$

with a new constant K_1 , depending on β and the μ -parameters. Indeed, from the first term in the definition (2.2.6) of c , one gets that $c \geq \inf_{M \geq 1} \left(K_1 \Psi_{1/\beta}^{\text{HT}}(M) \bar{\mu}_{\Psi_{1/\beta}^{\text{HT}}} \right) =: c_0 > 0$, which, from (2.2.5), yields the rough lower bound $\varepsilon \geq K_1 c_0 / \sqrt{M}$. This implies in turn (after tedious computations) that the second term in the definition (2.2.6) of c cannot be (up to constant) larger than the first term, hence the equality (2.2.8).

2.3 Proofs

In this section we give the proofs of our main results, namely our adaptation of the Talagrand inequality (Theorem 2.2.2 - Subsection 2.3.1), the maximal inequality (Theorem 2.2.3 - Subsection 2.3.2) and then our uniform concentration result (Theorem 2.2.5 - Subsection 2.3.3).

2.3.1 Proof of Theorem 2.2.2

In this section we start by recalling the tools from the Talagrand's seminal article [Tal89] used to prove our Talagrand inequality, then we give some technical argument and we prove our adapted Talagrand's inequality.

2.3.1.1 Preliminary results

Here, we recall Lemmas 8 and 9 of [Tal89], as well as the "Basic Estimate", which will enable us to prove Theorem 2.2.2. In addition to the independent B -valued random variables $(Y_m)_{m \in [M]}$, we will need to consider extra independent Rademacher random variables. Everything is defined

as follows. Let $(\Omega^M \times \Omega', \Sigma^M \otimes \Sigma', \mathbb{P})$ the basic probability space, where $\mathbb{P} = \mathbb{P} \otimes \mathbb{P}'$ such that the variables Y_m are defined on Ω^M and for $\omega = (\omega_m)_{m \in [M]}$, $Y_m(\omega)$ depends only on ω_m . Let $(\varepsilon_m)_{m \in [M]}$ be a set of random variables defined on Ω' with a Rademacher distribution independent of $(Y_m)_{m \in [M]}$. The following inequalities can be proven independently apart from the context of Orlicz norms.

Lemma 2.3.1 ([Tal89, Lemma 8]). *If $\mathbb{P}(\max_{m \in [M]} \|Y_m\|_B \geq t) \leq \frac{1}{2}$, then*

$$\sum_{m \in [M]} \mathbb{P}(\|Y_m\|_B \geq t) \leq 2\mathbb{P}\left(\max_{m \in [M]} \|Y_m\|_B \geq t\right).$$

Lemma 2.3.2 ([Tal89, Lemma 9]). *Set $X^{(r)}$ the r -th largest term of $(\|Y_m\|_B)_{m \in [M]}$. Then*

$$\mathbb{P}\left(X^{(r)} \geq t\right) \leq \frac{1}{r!} \left(\sum_{m \in [M]} \mathbb{P}(\|Y_m\|_B \geq t) \right)^r.$$

Set

$$\mu := \mathbb{E} \left[\left\| \sum_{m \in [M]} \varepsilon_m Y_m \right\|_B \right],$$

$\mu > 0$ because the Y_m 's are not all zero random variables (to avoid trivial situations). We now recall a key inequality which, combined with the previous lemmas, will enable us to prove the announced theorem.

Theorem 2.3.3 ([Tal89, Equation (2.5)]). *For k, q positive integers s.t. $k \geq q$, $u > 0$ and $u' > 0$, we have*

$$\begin{aligned} \mathbb{P} \left(\left\| \sum_{i \in [M]} \varepsilon_i Y_i \right\|_B \geq 4q\mu + u + u' \right) &\leq 4 \exp\left(-\frac{u^2}{64q\mu^2}\right) + \left(\frac{K_0}{q}\right)^k \\ &\quad + \mathbb{P} \left(\sum_{r \leq k} X^{(r)} > u' \right) \end{aligned}$$

where the constant K_0 is a universal constant.

2.3.1.2 Symmetrisation argument for Ψ convex

In the next subsection, because we rely on Theorem 2.3.3, we are going to prove the inequality (2.2.2) on symmetric random variables first (e.g. variables Y_m s.t. $\varepsilon_m Y_m \stackrel{d}{=} Y_m$). The extension to non-symmetric variables will be direct thanks to Lemma 2.3.5 which establishes an "equivalence in norms" relationship between the Orlicz norm of the sum of random variables and its associated Rademacher average.

Lemma 2.3.4. *Let Ψ be convex Orlicz function and $\|\cdot\|_\Psi$ the associate Orlicz norm. For any mean zero random variable $Z \in L_\Psi(B)$, we have $\|Z\|_\Psi \leq \|Z - Z'\|_\Psi$, with Z' any B -valued random variable such that $\mathbb{E}[Z'|Z] = 0$.*

Proof. Let $c > 0$,

$$\mathbb{E}[\Psi(\|Z\|_B/c)] \stackrel{(a)}{=} \mathbb{E} \left[\Psi \left(\frac{\|\mathbb{E}[Z - Z' | Z]\|_B}{c} \right) \right] \stackrel{(b)}{\leq} \mathbb{E} \left[\Psi \left(\frac{\mathbb{E}[\|Z - Z'\|_B | Z]}{c} \right) \right]$$

$$\stackrel{(c)}{\leq} \mathbb{E} \left[\mathbb{E} \left[\Psi \left(\frac{\|Z - Z'\|_B}{c} \right) \mid Z \right] \right] = \mathbb{E} \left[\Psi \left(\frac{\|Z - Z'\|_B}{c} \right) \right]$$

where in (a) we use Z' has a zero conditional mean, in (b) we use that Ψ is non decreasing and the triangular inequality holds for the $\|\cdot\|_B$, in (c) we apply the Jensen inequality. Hence by taking $c = \|Z - Z'\|_\Psi > 0$, the right hand side is smaller than 1 (using Property ((i)) in Section 2.2.1), and therefore $\|Z\|_\Psi \leq c = \|Z - Z'\|_\Psi$. \square

Lemma 2.3.5. *Let Ψ be as in Lemma 2.3.4. Let $(Y_m)_{m \in [M]}$ be a sequence of independent mean-zero random variables in $L_\Psi(B)$. Let $(\varepsilon_m)_{m \in [M]}$ be independent Rademacher random variables, and let $(Y'_m)_{m \in [M]}$ be an independent copy of the sequence $(Y_m)_{m \in [M]}$. Then*

$$\begin{aligned} \left\| \sum_{m \in [M]} Y_m \right\|_\Psi &\leq \left\| \sum_{m \in [M]} Y_m - \sum_{m \in [M]} Y'_m \right\|_\Psi = \left\| \sum_{m \in [M]} \varepsilon_m (Y_m - Y'_m) \right\|_\Psi \\ &\leq 2 \left\| \sum_{m \in [M]} \varepsilon_m Y_m \right\|_\Psi \leq 4 \left\| \sum_{m \in [M]} Y_m \right\|_\Psi. \end{aligned}$$

Later on, we will apply these inequalities with $\Psi = \Psi_\beta^{\text{HT}}$ and $\Psi(x) = x$ (the associated Orlicz norm corresponds then to the L_1 norm).

Proof. The first inequality comes from the application of Lemma 2.3.4 with $Z = \sum_{m \in [M]} Y_m$ and $Z' = \sum_{m \in [M]} Y'_m$. Since ε_m takes values ± 1 independently of Z, Z' , we have $Y_m - Y'_m \stackrel{d}{=} Y'_m - Y_m \stackrel{d}{=} \varepsilon_m (Y_m - Y'_m)$. Since the sequences are independent in m , we obtain the equality of Lemma 2.3.5. The second inequality is a consequence of the triangular inequality ((iv)) and the previous identities in distribution. The last inequality is a consequence of the application of Lemma 2.3.4 with $Z = \sum_{m \in [M]} \varepsilon_m Y_m$ and $Z' = \sum_{m \in [M]} \varepsilon_m Y'_m$ satisfying

$$\mathbb{E} [Z' | Z] = \mathbb{E} \left[\mathbb{E} \left[\sum_{m \in [M]} \varepsilon_m Y'_m \mid \varepsilon_m, Y_m, m \in [M] \right] \mid Z \right] = 0$$

and of the triangular inequality: $\|Z\|_\Psi \leq \|Z - Z'\|_\Psi = \left\| \sum_{m \in [M]} (Y_m - Y'_m) \right\|_\Psi \leq 2 \left\| \sum_{m \in [M]} Y_m \right\|_\Psi$. \square

2.3.1.3 Completion of the proof of Theorem 2.2.2

We will denote K a positive constant depending only on β , that may vary from line to line. We assume that at least one of the Y_m 's is not zero *a.s.*, otherwise the announced inequality (2.2.2) is obvious.

In view of the inequalities of Lemma 2.3.5, it is enough to do the reasoning and show the inequality (2.2.2) with the variables $(\varepsilon_m Y_m, m \in [M])$ instead of $(Y_m, m \in [M])$.

\triangleright **Rescaling** Note that (2.2.2) is invariant by homogeneous rescaling (see Property ((ii)) of Section 2.2.1), i.e. the inequality remains the same for the random variables $\tilde{Y}_m := \frac{\varepsilon_m Y_m}{C}$ for any $C > 0$. For the choice

$$C := \left\| \sum_{m \in [M]} \varepsilon_m Y_m \right\|_{L_1(B)} + \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} > 0,$$

observe that

$$\left\| \sum_{m \in [M]} \tilde{Y}_m \right\|_{L_1(B)} \leq 1 \quad \text{and} \quad \left\| \max_{m \in [M]} \|\tilde{Y}_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} \leq 1, \quad (2.3.1)$$

therefore the inequality (2.2.2) writes

$$\left\| \sum_{m \in [M]} \tilde{Y}_m \right\|_{\Psi_\beta^{\text{HT}}} \leq 2K.$$

Conversely, if the above holds for some K (independent from the \tilde{Y}_m 's verifying (2.3.1)), then (2.2.2) holds for the Y_m 's. All in all, it means that without loss of generality, we can assume

$$\left\| \sum_{m \in [M]} \varepsilon_m Y_m \right\|_{L_1(B)} \leq 1 \quad \text{and} \quad \left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} \leq 1,$$

and then show, under these assumptions, the existence of $K \in \mathbb{R}$ (independent on Y_m 's) such that

$$\mathbb{E} \left[\exp \left\{ \left(\xi_\beta \left(\frac{\left\| \sum_{m \in [M]} \varepsilon_m Y_m \right\|_B}{K} \right) \right) \right\} \right] \leq 2.$$

▷ **Deviation bounds** By Property ((iii)) of Section 2.2.1 and since we assumed $\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_{\Psi_\beta^{\text{HT}}} \leq 1$,

$$\mathbb{P} \left(\max_{m \in [M]} \|Y_m\|_B \geq t \right) \leq 2 \exp(-\xi_\beta(t)), \quad t \geq 0.$$

The function $\xi_\beta(\cdot) = (\ln(1 + \cdot))^\beta$ being continuously increasing from 0 to $+\infty$, there exists t_0 s.t. $\xi_\beta(t_0) = 2 \ln 2$ and $\forall t \geq t_0$, $2 \exp(-\xi_\beta(t)) \leq 1/2$; for further use, notice the value $t_0 = e^{(2 \ln 2)^{\frac{1}{\beta}}} - 1$. Then applying Lemma 2.3.1, for $t \geq t_0$, we have

$$\sum_{m \in [M]} \mathbb{P}(\|Y_m\|_B \geq t) \leq 2 \mathbb{P} \left(\max_{m \in [M]} \|Y_m\|_B \geq t \right) \leq 4 \exp\{-\xi_\beta(t)\}.$$

Hence Lemma 2.3.2 yields for $r \in \mathbb{N}^*$, $t \geq t_0$

$$\mathbb{P} \left(X^{(r)} \geq t \right) \leq \frac{4^r \exp\{-r \xi_\beta(t)\}}{r!}. \quad (2.3.2)$$

Denote $\tilde{\beta} = \lfloor \beta \rfloor + 1 \geq 2$. Equation (2.3.2) yields for $t \geq (e^{\tilde{\beta}} - 1) r^{\frac{\tilde{\beta}}{\beta}}$ (notice that $t \geq e^2 - 1 \geq t_0$ as requested)

$$\mathbb{P} \left(X^{(r)} \geq t r^{-\frac{\tilde{\beta}}{\beta}} \right) \leq \frac{4^r \exp \left\{ \left(-r \left[\ln \left(1 + t r^{-\frac{\tilde{\beta}}{\beta}} \right) \right]^\beta \right) \right\}}{r!} =: f(r, t).$$

Since $\tilde{\beta}/\beta > 1$, the sequence $(r^{-\frac{\tilde{\beta}}{\beta}})_{r \geq 1}$ is summable. Set $S_\beta := \sum_{r \geq 1} r^{-\frac{\tilde{\beta}}{\beta}} < +\infty$ and $g(t) := (t/(e^{\tilde{\beta}} - 1))^{\beta/\tilde{\beta}}$. From the inclusion $\{\sum_{r \leq g(t)} X^{(r)} \geq t S_\beta\} \subset \bigcup_{r \leq g(t)} \{X^{(r)} \geq t r^{-\frac{\tilde{\beta}}{\beta}}\}$ and writing a union bound, we get

$$\mathbb{P} \left(\sum_{r \leq g(t)} X^{(r)} \geq t S_\beta \right) \leq \sum_{r \leq g(t)} f(r, t). \quad (2.3.3)$$

We claim that for all $1 \leq r \leq g(t)$

$$r^{1/\beta} \ln\left(1 + tr^{-\frac{\tilde{\beta}}{\beta}}\right) \geq \ln(1+t). \quad (2.3.4)$$

This is a consequence of the above lemma applied with $\rho = r^{\frac{1}{\beta}} \geq 1$ and $\tau = tr^{-\frac{\tilde{\beta}}{\beta}} \geq e^{\tilde{\beta}} - 1$.

Lemma 2.3.6. *For all $\rho \geq 1$ and $\tau \geq e^{\tilde{\beta}} - 1$, we have $\rho \ln(1 + \tau) \geq \ln(1 + \tau \rho^{\tilde{\beta}})$.*

Proof. The function $f(\rho) := \rho \ln(1 + \tau) - \ln(1 + \tau \rho^{\tilde{\beta}})$ vanishes at $\rho = 1$, let us prove that it is non-decreasing in ρ provided that $\tau \geq e^{\tilde{\beta}} - 1$. Indeed,

$$f'(\rho) = \ln(1 + \tau) - \frac{\tilde{\beta} \rho^{\tilde{\beta}-1} \tau}{1 + \rho^{\tilde{\beta}} \tau}.$$

Since $\rho^{\tilde{\beta}} \tau \geq 0$ and $\rho \geq 1$, we have $\frac{\tilde{\beta} \rho^{\tilde{\beta}-1} \tau}{1 + \rho^{\tilde{\beta}} \tau} \leq \frac{\tilde{\beta}}{\rho} \leq \tilde{\beta}$. Hence, $f'(\rho) \geq \ln(1 + \tau) - \tilde{\beta} \geq 0$. We are done. \square

Plugging (2.3.4) into (2.3.3) yields

$$\mathbb{P}\left(\sum_{r \leq g(t)} X^{(r)} \geq tS_\beta\right) \leq \sum_{r \leq g(t)} \frac{4^r}{r!} \exp\left\{-[\ln(1+t)]^\beta\right\} \leq \exp(4) \exp(-\xi_\beta(t)).$$

Let us recall that $\mu = \left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_{L_1(B)} \leq 1$. We are now at the point to apply Theorem 2.3.3 with $q = \lceil eK_0 \rceil$, $u = t$, $u' = tS_\beta$, $2q\mu \leq t$ and $k = \lfloor g(t) \rfloor$:

$$\begin{aligned} & \mathbb{P}\left(\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B \geq t(S_\beta + 3)\right) \\ & \leq \mathbb{P}\left(\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B \geq 4q\mu + u + u'\right) \\ & \leq 4 \exp\left\{-\frac{t^2}{64q\mu^2}\right\} + \exp\{-\lfloor g(t) \rfloor\} + \mathbb{P}\left(\sum_{r \leq g(t)} X^{(r)} \geq tS_\beta\right) \\ & \leq 4 \exp\left\{-\frac{t^2}{64q\mu^2}\right\} + \exp\{-\lfloor g(t) \rfloor\} + \exp\{(4 - \xi_\beta(t))\}. \end{aligned}$$

The above inequality is valid for any $t \geq t_0 \vee (2\lceil eK_0 \rceil \mu)$. Besides, in the above upper bound, the last third is asymptotically the largest one, therefore there exists $K > 0$ such that

$$\mathbb{P}\left(\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B \geq Kt\right) \leq K \exp\{-\xi_\beta(t)\}, \quad t \geq 0.$$

▷ **Orlicz norm bounds** The estimate implies for all $c > 0$:

$$\mathbb{E}\left[\exp\left\{\left(\xi_\beta\left(\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B / (cK)\right)\right)\right\}\right] - 1$$

$$\begin{aligned}
&= \int_0^\infty \exp\{(\xi_\beta(t))\} \xi'_\beta(t) \mathbb{P}\left(\frac{\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B}{cK} \geq t\right) dt \\
&\leq K \int_0^\infty \xi'_\beta(t) \exp\{(\xi_\beta(t) - \xi_\beta(ct))\} dt.
\end{aligned} \tag{2.3.5}$$

Let us check that the above integral is finite for $c > 1$. Only the integrability at $t \rightarrow +\infty$ is questionable. Write

$$\begin{aligned}
\xi_\beta(t) - \xi_\beta(ct) &= (\ln(1+t))^\beta \left(1 - \left[1 + \frac{\ln\left(\frac{1+ct}{1+t}\right)}{\ln(1+t)}\right]^\beta\right) \\
&\approx_{t \rightarrow +\infty} -\beta(\ln(1+t))^{\beta-1} \ln(c).
\end{aligned}$$

Therefore, the function to integrate is bounded for t large by (up to constant)

$$g(t) := \frac{(\ln(1+t))^{\beta-1}}{(1+t)} e^{-\frac{1}{2}\beta(\ln(1+t))^{\beta-1} \ln(c)}.$$

We easily check that $\int_0^{+\infty} g(t) dt = \int_0^{+\infty} y^{\beta-1} e^{-\frac{1}{2}\beta y^{\beta-1} \ln(c)} dy < +\infty$ since $\beta > 1$ and $c > 1$.

Furthermore, by monotone convergence theorem, the bound (2.3.5) converges to 0 as $c \rightarrow +\infty$, consequently

$$\mathbb{E} \left[\exp \left(\xi_\beta \left(\frac{\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_B}{cK} \right) \right) \right] \leq 2$$

for a $c = c_\beta$ large enough. We have proved that $\left\|\sum_{m \in [M]} \varepsilon_m Y_m\right\|_{\Psi_\beta^{\text{HT}}} \leq c_\beta K$. \square

2.3.2 Proof of Theorem 2.2.3

We start by recalling the general maximal inequality on which our proof is based.

Lemma 2.3.7 ([vW96, Lemma 2.2.2]). *Let Ψ be a convex Orlicz function satisfying*

$$\limsup_{x, y \rightarrow +\infty} \Psi(x)\Psi(y)/\Psi(c_\Psi xy) < +\infty \tag{2.3.6}$$

for some constant $c_\Psi > 0$. Then, there is a constant $K > 0$ such that for any B -valued random variables Y_1, \dots, Y_M ,

$$\left\| \max_{m \in [M]} \|Y_m\|_B \right\|_\Psi \leq K \Psi^{-1}(M) \max_{m \in [M]} \|Y_m\|_\Psi.$$

For $\beta > 1$, Ψ_β^{HT} is a convex Orlicz function, thus it remains to establish (2.3.6) to get Theorem 2.2.3. We prove that one can take $c_\Psi = 1$. Let $c \geq 1^3$ s.t. $\Psi_\beta^{\text{HT}}(c^2) \geq 1$. Let x, y st. $x \geq c$ and $y \geq c$: then

$$\begin{aligned}
\Psi_\beta^{\text{HT}}(x)\Psi_\beta^{\text{HT}}(y) &\leq e^{(\ln(1+x))^\beta} e^{(\ln(1+y))^\beta} \\
&\leq e^{(\ln(x))^\beta + (\ln(y))^\beta - (\ln(xy))^\beta} e^{(\ln(1+xy))^\beta} \underbrace{e^{2 \sup_{z \geq c} (\ln(1+z))^\beta - (\ln z)^\beta}}_{:=C(c)}.
\end{aligned}$$

³one can take $c = \sqrt{e^{(\ln 2)^{1/\beta}} - 1} \geq 1$ for which $\Psi_\beta^{\text{HT}}(c^2) = 1$.

- $C(c)$ is finite: indeed, by standard equivalents, we have that

$$(\ln(1+z))^\beta - (\ln z)^\beta = (\ln z)^\beta \left(\left[1 + \frac{\ln(1+z^{-1})}{\ln(z)} \right]^\beta - 1 \right) \underset{z \rightarrow \infty}{\sim} \beta \frac{(\ln z)^{\beta-1}}{z}$$

which converges to 0 at infinity.

- Notice that $(\ln(xy))^\beta - (\ln(x))^\beta \geq (\ln(y))^\beta$ for any $x, y \geq 1$. Indeed, setting $u = \ln x \geq 0$, $v = \ln y \geq 0$,

$$(u+v)^\beta - u^\beta = \int_u^{u+v} \beta z^{\beta-1} dz \geq \int_0^v \beta z^{\beta-1} dz = v^\beta$$

(because $z \mapsto \beta z^{\beta-1}$ is increasing since $\beta > 1$).

- Last, since $e^{(\ln(1+xy))^\beta} = \Psi_\beta^{\text{HT}}(xy) + 1 \geq \Psi_\beta^{\text{HT}}(c^2) + 1 \geq 2$, one has

$$e^{(\ln(1+xy))^\beta} = \frac{e^{(\ln(1+xy))^\beta}}{e^{(\ln(1+xy))^\beta} - 1} \Psi_\beta^{\text{HT}}(xy) \leq 2\Psi_\beta^{\text{HT}}(xy).$$

All in all, we conclude that $\Psi_\beta^{\text{HT}}(x)\Psi_\beta^{\text{HT}}(y) \leq 2C(c)\Psi_\beta^{\text{HT}}(xy)$, for any $x, y \geq c$. We are done. \square

2.3.3 Proof of Theorem 2.2.5

We follow the strategy of Chapter 1 by truncating the unbounded functions f by a threshold c , whose impact is analyzed using the Hoffman-Jorgensen inequality [LT13, Proposition 6.8] and the Talagrand inequality of Theorem 2.2.2. The deviation probability related to the newly bounded random variables is quantified thanks to the Klein-Rio inequalities [KR05] together with the Dudley entropy integral bound.

Here are the notations used along this proof. We denote by K a positive constant that may change from line to line in the computations: this generic constant K may depend on universal constants and β , but it does not depend on the sample X_1, \dots, X_M , its size M , nor the class of functions \mathcal{F} , neither ε . For ease of notations, we write $a \leq_K b$ when $a \leq Kb$.

For a given $c > 0$, set

$$\begin{aligned} \mathcal{R}_c f &:= f - \mathcal{T}_c f \quad , \text{ where } \quad \mathcal{T}_c f := -c \vee f \wedge c, \\ \mathcal{T}_c \mathcal{F} &:= \{\mathcal{T}_c f : f \in \mathcal{F}\}, \\ \mathcal{T}_c^m f(\cdot) &:= \mathcal{T}_c f(\cdot) - \mathbb{E}[\mathcal{T}_c f(X_m)], \\ Z_c &:= \sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{m \in [M]} \mathcal{T}_c^m f(X_m). \end{aligned}$$

Note that the function $\mathcal{T}_c^m f$ is centered w.r.t. the distribution of X_m , and bounded by $2c$. Assume that $c > 0$ and $\varepsilon > 0$ are such that

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m \in [M]} \mathbb{E}[\mathcal{R}_c f(X_m)] \right| \leq \varepsilon/4, \tag{2.3.7}$$

$$\mathbb{E}[Z_c] \leq \varepsilon/4. \tag{2.3.8}$$

By writing $f = \mathcal{R}_c f + \mathcal{T}_c f$ and using the sub-additivity of the supremum, we easily get

$$\sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{m \in [M]} (f(X_m) - \mathbb{E}[f(X_m)])$$

$$\leq Z_c - \mathbb{E}[Z_c] + \sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m \in [M]} \mathcal{R}_c f(X_m) \right| + \varepsilon/2.$$

Hence, the probability of deviation in Theorem 2.2.5 is bounded by

$$\mathbb{P} \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m \in [M]} \mathcal{R}_c f(X_m) \right| \geq \varepsilon/4 \right) + \mathbb{P}(Z_c - \mathbb{E}[Z_c] \geq \varepsilon/4) =: (\star) + (\star\star).$$

▷ **Term** (\star) Owing to the deviation inequality ((iii)) from Section 2.2.1, it is bounded by

$$\begin{aligned} & \mathbb{P} \left(\sum_{m \in [M]} \sup_{f \in \mathcal{F}} |\mathcal{R}_c f(X_m)| \geq M\varepsilon/4 \right) \\ & \leq 2 \exp \left(- \left(\ln \left(\frac{M\varepsilon/4}{\left\| \sum_{m \in [M]} \sup_{f \in \mathcal{F}} |\mathcal{R}_c f(X_m)| \right\|_{\Psi_\beta^{\text{HT}}}} + 1 \right) \right)^\beta \right). \end{aligned}$$

Using the Talagrand inequality of Theorem 2.2.2 and the Hoffman-Jorgensen inequality [LT13, Proposition 6.8], and following line by line the arguments of Chapter 1 Section 1.5.5 equations (1.5.21) and (1.5.22), we can show that the above $\|\cdot\|_{\Psi_\beta^{\text{HT}}}$ norm is bounded by

$$K \left\| \max_{m \in [M]} F(X_m) \right\|_{\Psi_\beta^{\text{HT}}},$$

provided that $c \geq 8\mathbb{E}[\max_{m \in [M]} F(X_m)]$. The above arguments are crucial to deal both with the truncation in c and the sup in f . Furthermore, the maximal inequality (2.2.3) with $Y_m := F(X_m)$ gives that

$$\left\| \max_{m \in [M]} F(X_m) \right\|_{\Psi_\beta^{\text{HT}}} \leq \mathcal{C}_{\beta, (2.2.3)} \Psi_{1/\beta}^{\text{HT}}(M) \bar{\mu}_{\Psi_\beta^{\text{HT}}}. \quad (2.3.9)$$

All in all, we have

$$c \geq 8\mathbb{E} \left[\max_{m \in [M]} F(X_m) \right] \implies (\star) \leq 2 \exp \left(- \left(\ln \left(\frac{M\varepsilon}{K \bar{\mu}_{\Psi_\beta^{\text{HT}}} \Psi_{1/\beta}^{\text{HT}}(M)} + 1 \right) \right)^\beta \right).$$

The above condition on the left hand side is met as soon as

$$c \geq K \Psi_{1/\beta}^{\text{HT}}(M) \bar{\mu}_{\Psi_\beta^{\text{HT}}},$$

where we have used $\mathbb{E}[\cdot] \leq K \|\cdot\|_{\Psi_\beta^{\text{HT}}}$ and (2.3.9).

▷ **Term** $(\star\star)$ Apply the Klein-Rio inequality [KR05, Theorem 1.1] (we shall use the form presented in Chapter 1 Theorem 1.3 which directly fits our setting), it shows that $(\star\star)$ is bounded by

$$\exp \left(- \frac{M(\varepsilon/4)^2}{2(\sigma^2 + 4c\mathbb{E}(Z_c)) + 6c(\varepsilon/4)} \right)$$

where $\sigma^2 := \sup_{f \in \mathcal{F}} \frac{1}{M} \max_{m \in [M]} \mathbb{E} [(\mathcal{T}_c^m f)^2(X_m)]$. Observe that

$$\sigma^2 \leq \sup_{f \in \mathcal{F}} \max_{m \in [M]} \text{Var} [\mathcal{T}_c^m f(X_m)] \leq \sup_{f \in \mathcal{F}} \max_{m \in [M]} \mathbb{E} [\mathcal{T}_c f^2(X_m)] \leq \mu_2^2.$$

Using in addition the bound (2.3.8) on $\mathbb{E}(Z_c)$, we get

$$(\star\star) \leq \exp\left(-\frac{M\varepsilon^2}{K(\mu_2^2 + c\varepsilon)}\right)$$

where K is a universal constant.

▷ **Condition (2.3.7)** From $\mathcal{R}_c f(x) = (f(x) - c)_+ - (f(x) + c)_-$, we easily get

$$\begin{aligned} |\mathbb{E} [\mathcal{R}_c f(X_m)]| &\leq \int_c^{+\infty} \mathbb{P}(|f(X_m)| \geq z) dz \\ &\leq 2 \int_c^{+\infty} \exp\left(-(\ln(z/\lambda + 1))^\beta\right) dz \\ &= 2\lambda \int_{c/\lambda}^{+\infty} \exp\left(-(\ln(z + 1))^\beta\right) dz =: 2\lambda \mathcal{I}(c/\lambda) \end{aligned}$$

where $\lambda := \mu_{\Psi_\beta^{\text{HT}}}$. A standard calculus shows that

$$\mathcal{I}(y) \sim_{y \rightarrow +\infty} \frac{y}{\beta(\ln(y + 1))^{\beta-1}} \exp\left(-(\ln(y + 1))^\beta\right),$$

and thus

$$\mathcal{I}(y) \leq_K \exp\left(-(\ln(y + 1))^\beta/2\right) =: \mathcal{J}(y), \quad \forall y \geq 0.$$

This gives

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{M} \sum_{m \in [M]} \mathbb{E} [\mathcal{R}_c f(X_m)] \right| \leq_K \lambda \exp\left(-(\ln(c/\lambda + 1))^\beta/2\right).$$

Therefore, to ensure (2.3.7) it is enough to take

$$c \geq \mu_{\Psi_\beta^{\text{HT}}} \left(\exp\left((2 \ln_+(K \mu_{\Psi_\beta^{\text{HT}}}/\varepsilon))^{1/\beta}\right) - 1 \right)$$

for some constant $K > 0$. Observe that the use of $\ln_+(\cdot)$ guarantees that for a deviation ε large enough, the above lower bound is zero, meaning that any value of $c \geq 0$ ensures that (2.3.7) holds, as it is expected (for large ε).

▷ **Condition (2.3.8)** Deriving a bound on the expectation of the supremum follows a standard routine using Dudley entropy integral bound. For sake of brevity, we closely follow the arguments of 1.5.3.2, term $\mathbb{E}[\bar{Z}^{\mathcal{T}_c}]$ bounding. It gives that

$$\mathbb{E}[Z_c] \leq 2\mathbb{E} \left[\frac{C_D}{\sqrt{M}} \int_0^\infty \sqrt{\ln(\mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F}))} dz \right] \quad (2.3.10)$$

where $d_{\mathcal{F}}(f, g) := \left(\frac{1}{M} \sum_{m=1}^M |f(X_m) - g(X_m)|^2 \right)^{1/2}$ and $\mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F})$ is the covering number of $\mathcal{T}_c \mathcal{F}$ with respect to the distance $d_{\mathcal{F}}$ with balls of radius z (see [GKKW02, Definition 9.3]). Actually, since functions in $\mathcal{T}_c \mathcal{F}$ are bounded by c , $\mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F}) = 1$ for $z \geq 2c$ and therefore,

the above integral can be restricted to $[0, 2c]$ without modification. In addition, we have the following universal upper bound in terms of VC dimension:

$$0 < z < 2c/4 \implies \mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F}) \leq 3 \left[2e \left(\frac{2c}{z} \right)^2 \ln \left(3e \left(\frac{2c}{z} \right)^2 \right) \right]^{V_{\mathcal{F}^+}}.$$

Indeed, the above estimate follows from [GKKW02, Lemma 9.2, Theorem 9.4 with $B = 2c$ and $p = 2$, $V_{(\mathcal{T}_c \mathcal{F})^+} \leq V_{\mathcal{F}^+}$ in the proof of Theorem 9.6]. See [vW96, Theorem 2.6.7] for a variant of this upper bound. Since $\mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F})$ is non-decreasing in z , and since we do not pay much attention to universal constants, we can simply write

$$0 < z \leq 2c \implies \mathcal{N}_2(z, d_{\mathcal{F}}, \mathcal{T}_c \mathcal{F}) \leq \left(\frac{Kc}{z} \right)^{3V_{\mathcal{F}^+}},$$

for a universal constant K . Plugging this into (2.3.10) readily leads to

$$\mathbb{E}[Z_c] \leq Kc \frac{\sqrt{V_{\mathcal{F}^+}}}{\sqrt{M}}.$$

▷ **Conclusion** Gathering all the estimates and conditions leads to the statement of Theorem 2.2.5. \square

2.3.4 Proof of Proposition 2.2.1

Item 1. Observe that $\Psi_1^{\text{HT}}(x) = x$ and $\Psi_{\beta_1}^{\text{HT}}(\Psi_{\beta_2}^{\text{HT}}(x)) = \Psi_{\beta_1 \beta_2}^{\text{HT}}(x)$ for any $x \geq 0$; the property of group isomorphism readily follows.

Items 2 and 4 are straightforward to verify.

Item 3. Ψ_{β}^{HT} is a C^∞ -function on $(0, \infty)$, with a second derivative equal to

$$\begin{aligned} \Psi_{\beta}^{\text{HT}''}(x) &= \frac{\exp\left((\ln(1+x))^{\beta}\right) (\ln(1+x))^{\beta-2}}{(1+x)^2} \\ &\quad \times \beta \times \underbrace{\left(\beta (\ln(1+x))^{\beta} + (\beta-1) - \ln(1+x) \right)}_{=: g(\ln(1+x))}. \end{aligned}$$

The function g is continuously differentiable on \mathbb{R}^+ , strictly positive at 0 ($g(0) = \beta - 1 > 0$) and goes to infinity at infinity (since $\beta > 1$); the critical points of g' are solutions to $\beta^2 y^{\beta-1} - 1 = 0$, therefore it is unique (equal to $y_{\beta} := \beta^{-\frac{2}{\beta-1}}$) and corresponds to the minimum of g . Let us evaluate the sign of g at the minimum:

$$\begin{aligned} g(y_{\beta}) &= \beta y_{\beta}^{\beta} + (\beta-1) - y_{\beta} = \frac{y_{\beta}}{\beta} + (\beta-1) - y_{\beta} \\ &= (\beta-1) \left(1 - \frac{y_{\beta}}{\beta} \right) = (\beta-1) \left(1 - \frac{1}{\beta^{\frac{\beta+1}{\beta-1}}} \right) > 0. \end{aligned}$$

All in all, we have proved that $\Psi_{\beta}^{\text{HT}''}(x) > 0$ for any $x > 0$. \square

2.4 Conclusion

To conclude, we have extended the Talagrand inequality for an Orlicz norm adapted to variables with β -heavy tails (Proposition 2.2.1 and Theorem 2.2.2). We have also shown that a maximal inequality holds (Theorem 2.2.3), which, in combination with the Talagrand inequality, allows for

a concentration inequality for the sum of independent centered β -heavy tailed random variables (Corollary 2.2.4). Then we have extended this inequality to supremum of functions of random variables with β -heavy tails (Theorem 2.2.5), by combining previous results with the Hoffman-Jorgensen, Klein-Rio and Dudley entropy integral inequalities.

Part II

Covariance estimation for minimum variance portfolio

CHAPTER 3

ASYMPTOTIC ANALYSIS OF DIFFERENT COVARIANCE MATRICES ESTIMATION FOR MINIMUM VARIANCE PORTFOLIO

Note. This chapter is based on an article written jointly with E. Gobet and J.-P. Lemor. A shortened version has been submitted to *Frontiers of Mathematical Finance*.

Abstract. In dynamic minimum variance portfolio, we study the impact of the sequence of covariance matrices taken in inputs, on the realized variance of the portfolio computed along a sample market path. The allocation of the portfolio is adjusted on a regular basis (every H days) using an updated covariance matrix estimator. In a modelling framework where the covariance matrix of the asset returns evolves as an ergodic process, we quantify the probability of observing an underperformance of the optimal dynamic covariance matrix compared to any other choice. The bounds depend on the tails of the returns, on the adjustment period H , and on the total number of rebalancing times N . These results provide asset managers with new insights into the optimality of their choice of covariance matrix estimators, depending on the depth of the backtest NH and the investment period H . Experiments based on GARCH modelling support our theoretical results

3.1 Introduction

The mean-variance efficient portfolio theory by Markowitz [Mar52] has had a profound impact on modern finance. The Markowitz portfolio selection requires estimates of the covariance matrix and of the average expected returns. The covariance matrix can be either estimated in a non-parametric way, using sample-based empirical estimation, or in a parametric way, using factor models for example. In both cases, we may require sliding moving averages on the historical data.

Yet, the selected allocations are very sensitive to the values of the covariance matrix and expected returns used for its computation, and small changes in the inputs can lead to large changes in the allocations [Mic89].

On the other hand, it is well known that financial data exhibit heteroscedasticity, that is to say time dependent conditional covariance. This statistical property, referred to as stylized fact in the financial data setting, has been largely documented in the literature ([?], [EP07]). This implies that the expected covariance matrix in the near future can be very different from the average of the expected covariance matrix over a long time horizon. In this work, we aim to address the problem of the covariance matrix choice only.

In this chapter, we are interested in the forecast of the *realized covariance* for a specific time and period of investment. The minimum variance portfolio is the portfolio taking as input a covariance matrix and giving as output the allocation minimizing the associated variance. The key quantity is then the realized covariance over the period of investment. We consider a model for the returns of type

$$\mathbf{r}_{t|\mathcal{F}_{t-1}} \sim \mathcal{N}(0, V_t), \quad (3.1.1)$$

denoting t the time of investment and H the period of investment, the realized covariance corresponds to the future matrix $\frac{1}{H} \sum_{k=1}^H \mathbf{r}_{t+k} \mathbf{r}_{t+k}^\top$. Its best estimation at time t is the *conditional realized covariance*:

$$\frac{1}{H} \mathbb{E} \left[\sum_{k=1}^H \mathbf{r}_{t+k} \mathbf{r}_{t+k}^\top | \mathcal{F}_t \right] = \frac{1}{H} \sum_{k=1}^H \mathbb{E} [V_{t+k} | \mathcal{F}_t].$$

For $H = 1$, it coincides with the conditional covariance V_{t+1} ; when the period of investment exceeds the period of observation of the returns ($H > 1$), this quantity can still be estimated at time t . In practice, t and H are measured in days and $H = 21$ would correspond to an investment over a month (in business days) for example.

Usually, the asset manager might also consider a historical based covariance $\frac{1}{T} \sum_{k=1}^T \mathbf{r}_{t-k} \mathbf{r}_{t-k}^\top$ based on the past returns. When the backtest size T goes to infinity, this estimator converges to the *stationary covariance matrix* V_∞ .

The purpose of this chapter is to study the impact of the choice of the covariance matrix on the performance of the strategy.

In particular, we would like to give optimality guarantees under the form of concentration of measure inequalities of the outperformance of the portfolio based on the conditional realized covariance, versus any other covariance estimate V_{ref} . In the context of minimum variance investment, the natural portfolio metric is the realized variance, also called out-of-sample variance of the portfolio. Our result takes the form of a high probability event over the sums of the realized variance (RV) for N rebalancing dates of the portfolios:

$$\mathbb{P} \left(\sum_{n=1}^N \text{RV} \left(\frac{1}{H} \mathbb{E} \left[\sum_{k=1}^H \mathbf{r}_{t_n+k} \mathbf{r}_{t_n+k}^\top | \mathcal{F}_{t_n} \right] \right) \leq \sum_{n=1}^N \text{RV} (V_{\text{ref}}) \right) \approx 1 - c \left(\frac{H}{N} \right)^q$$

where $c > 0$ is a constant independent from H and N and q is a positive constant depending on the integrability of the process.

It means that the realized variance of the portfolio based on the conditional realized covariance will on the long term be lower than with other covariance with a high probability. The probability is even higher as the number N of times the portfolio is rebalanced grows, and decreases with the period of investment H , with a convergence rate increasing with the integrability of the process. The fact that the probability decreases with H is due to the fact that as H goes to infinity, the conditional realized covariance tends to the stationary covariance, which becomes optimal. But as we may illustrate later, the half-life time associated to financial returns is often larger than the usual order of magnitude for H used by practitioners. Hence, considering non infinite H as we do is meaningful.

3.1.1 Literature background

Our article falls within the line of the sensitivity analysis for Markowitz allocation articles.

In Guigues [Gui11] and Kan et al [KZ07], the authors give bounds on the portfolio risk as function of the variation of the input. In their approach, the metric is the mean-variance utility: $U(\mathbf{w}, \boldsymbol{\mu}, \Sigma) = t\mathbf{w}'\boldsymbol{\mu} - \frac{1}{2}\mathbf{w}'\Sigma\mathbf{w}$. Using perturbation analysis results from Bonnans and Shapiro [BS00], Guigues shows that for two pairs of inputs $(\boldsymbol{\mu}_1, \Sigma_1)$ and $(\boldsymbol{\mu}_2, \Sigma_2)$, the optimal utilities difference can be bounded by the norm of the inputs differences:

$$|U(\mathbf{w}_2^*, \boldsymbol{\mu}_2, \Sigma_2) - U(\mathbf{w}_1^*, \boldsymbol{\mu}_1, \Sigma_1)| \leq \frac{1}{2}|\Sigma_2 - \Sigma_1|_\infty + k|\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1|_\infty,$$

where \mathbf{w}^* is the portfolio which maximizes the mean-variance utility in \mathbf{w} . In Kan and Zhou's article, the authors assume a multivariate Gaussian model for the returns and provide bounds on the expected difference between the utility of the mean variance portfolio with sample based estimates for $\boldsymbol{\mu}$ and Σ versus the optimal utility with the true parameters $\boldsymbol{\mu}$ and Σ (also called population parameters). The bound is linear in the number of assets and proportional to the inverse of the time of estimation.

Their bound holds on the expected value of the risk whereas in our approach, it is on the empirical risk, which is closer to practitioners needs.

A first attempt could be to use concentration of measures results for matrices (Tropp [Tro12], El Karoui [EK18]), and combine them with the aforementioned sensitivity results. But usually, this approach relies on independence properties between the matrices, which we don't have since V_t is a stochastic process, and on high order integrability (like sub-Gaussian tails) which we don't have since usually, asset returns have heavy tails.

The optimal horizon-adapted covariance matrix is similar to the one proposed in De Nard et al [DNLW18]. In this article, the authors compare empirically the performance of the minimum variance portfolio with different dynamic and static covariance estimation methods, including the GARCH-Dynamic Conditional Correlation (GARCH-DCC) models for returns or for residuals of a static factor models. Their experiments show that these two models outperform 10 other parametric estimation models in term of realized out-of-sample variance, which is in line with the theoretical results we have obtained in this chapter.

3.1.2 Contribution and outline of this chapter

In this chapter:

- We give guarantees of the optimality of the conditional realized covariance in the minimum variance portfolio setting in the form of a large deviation inequality, with a convergence rate polynomially decreasing in the number of rebalancing times of the portfolio,
- We give a scheme for estimating this covariance matrix in the specific GARCH-Constant Conditional Correlation (GARCH-CCC) model,
- We display results of numerical experiments to illustrate our statements.

Section 3.2 provides a presentation of the problem, introduces the notations and states the main result of the chapter. In Section 5.A, the proof of the main result is given. In Section 3.4, we introduce the GARCH-CCC model and verify that it satisfies the conditions of Section 3.2. Section 5.4 presents the result of the numerical experiments.

3.2 Formulation of the problem

In this section we start by formulating our problem and we introduce our model assumptions before stating our main result.



Figure 3.1 – Illustration of quantities time-dependence in our setting: since r_t and V_{t+1} are \mathcal{F}_t measurables, they are the known quantities at each time t . The red dots dates correspond to the rebalancing times.

3.2.1 Problem setup

In this section, we introduce the setting of the problem and we define the variance quantities of interest in our study.

The model: Consider a pool of d assets, with daily centered returns processes $r_{1,t}, \dots, r_{d,t}$. Denote by \mathbf{r}_t the vector of the returns processes from day $t-1$ and t . Let $\{\eta_t\}_{t \geq 0}$ i.i.d., $\eta_t \sim \mathcal{N}(0, I_d)$ the *innovations* process and $\mathcal{F}_t = \sigma\{\eta_s, 0 \leq s \leq t\}$ the associated filtration. We assume that \mathbf{r}_t is given by:

$$\mathbf{r}_t = V_t^{1/2} \eta_t, \quad \eta_t \sim \mathcal{N}(0, I_d), \quad t \geq 1, \quad (3.2.1)$$

where $V_t \in \mathbb{R}^{d \times d}$ is a \mathcal{F}_{t-1} measurable, positive definite matrix. It means that $\mathbf{r}_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, V_t)$. The initial condition V_1 is deterministic.

We assume that $\{V_t\}_{t > 1}$ is square integrable, and we denote $p_{\max} \geq 2$ s.t. $\mathbb{E}[|V_t|^{p_{\max}}]$ is finite for $t > 1$.

Covariance notation: A given portfolio with the allocation vector \mathbf{w} at time T_{reb} and a holding period of H has the realized variance:

$$\mathbf{w}^\top \left(\sum_{k=1}^H \mathbf{r}_{T_{reb}+k} \mathbf{r}_{T_{reb}+k}^\top \right) \mathbf{w} = \mathbf{w}^\top RC_{H, T_{reb}} \mathbf{w} \quad (3.2.2)$$

with $RC_{H, T_{reb}}$ denoting the *realized covariance* over the period of investment. We seek for the portfolio which minimizes this realized variance.

At the time of the investment, the best estimation of $RC_{H, T_{reb}}$ is:

$$cRC_{H, T_{reb}} := \mathbb{E}[RC_{H, T_{reb}} | \mathcal{F}_{T_{reb}}] = \sum_{k=1}^H \mathbb{E}[V_{T_{reb}+k} | \mathcal{F}_{T_{reb}}]$$

which we call the *conditional realized covariance*. Given a covariance matrix C which we assume to be definite positive, we will consider the following risk optimization under constraints

$$\text{mv}(C) := \arg \min_{\mathbf{w} \in \mathcal{W}} \mathbf{w}^\top C \mathbf{w},$$

where \mathcal{W} is the set of constraints containing at least the budget constraint and a maximal allocation constraint: $\left\{ \mathbf{w} = \{w_i\}_{i \in \{1, \dots, d\}} \in \mathbb{R}^d : \sum_{i=1}^d w_i = 1, \text{ and } |w_i| \leq c_w, i \in \{1, \dots, d\} \right\} \subset \mathcal{W}$, $c_w > 0$. We assume that \mathcal{W} is closed and convex, which ensures the existence and uniqueness of $\text{mv}(C)$. We aim at showing that $\text{mv}(cRC_{H, T_{reb}})$ is a good $\mathcal{F}_{T_{reb}}$ -measurable allocation for minimizing the realized variance (3.2.2).

We will consider multiple rebalancing times of the portfolio: $T_{reb} = t_1, \dots, t_N$, $t_{n+1} - t_n = H$, $t_0 = 0$.

Let us introduce the following processes:

$$\begin{aligned}
R_{N,H} &:= \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top RC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\
cRV_{N,H} &:= \sum_{n=1}^N \text{mv}(cRC_{H,t_n})^\top cRC_{H,t_n} \text{mv}(cRC_{H,t_n}), \\
R_{N,H}^{\text{ref}} &:= \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top RC_{H,t_n} \text{mv}(V_{\text{ref}}), \\
cRV_{N,H}^{\text{ref}} &:= \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top cRC_{H,t_n} \text{mv}(V_{\text{ref}}),
\end{aligned} \tag{3.2.3}$$

with V_{ref} is a deterministic, positive definite covariance matrix which we take as the benchmark covariance the asset manager considers for his optimization.

- $R_{N,H}$ (resp. $R_{N,H}^{\text{ref}}$) is the realized variance of the cRC_{H,t_n} -based portfolios (resp. the V_{ref} -based portfolio);
- $cRV_{N,H}$ (resp. $cRV_{N,H}^{\text{ref}}$) denotes the sum of conditional realized variances of the cRC_{H,t_n} -based portfolios (resp. the V_{ref} -based based portfolio).

For $p_{\max} \geq 2$, our main result Theorem 3.2.2 takes the form:

$$\mathbb{P} \left(R_{N,H} < R_{N,H}^{\text{ref}} \right) \geq 1 - C \left(\frac{H}{N} \right)^{\frac{p_{\max}}{2}}.$$

This inequality is of the form of the probability of the realized variance being lower (e.g. better) for the conditional realized covariance-based portfolio than for the reference covariance-based portfolio.

The probability bound is polynomially decreasing in N , the number of rebalancing times of the portfolio and polynomially increasing in H , the investment horizon, with an exponent $\frac{p_{\max}}{2}$ equal to a quarter of the maximum finite moment of the returns. This relatively slow convergence (polynomial rather than exponential) stems from the low integrability on the process \mathbf{r}_t .

The main message of this result is that, on the long-run (N going to infinity), the conditional realized covariance-based portfolio is outperforming with high probability. For very large H though, this bound might become loose, since the infinite-horizon estimate of $\frac{RC_{H,T_{\text{reb}}}}{H}$ coincides with the stationary covariance V_∞ .

3.2.2 Model assumptions

We recall that we are under the model (3.2.1) for the returns. We are now going to specify the assumptions on the conditional covariance matrix V_t .

We denote $\mathcal{S} \subset \mathcal{S}_+$ the state space on which V_t takes its values, where \mathcal{S}_+ is the set of symmetric positive definite matrices.

H_{stat} . $\{V_t\}_{t \in \mathbb{N}^*}$ possesses a strictly stationary and ergodic distribution μ with at least L_2 moment and $V_\infty := \int_{\mathcal{S}} v \mu(dv)$.

$H_{\mathcal{S}}$. $\{V_t\}_{t \in \mathbb{N}}$ is a time homogeneous, aperiodic Lebesgue-irreducible Markov chain¹ on the state-space \mathcal{S} .

¹We refer the reader to paragraph 3.3.2.1 for reminders on Markov chains elements of language.

$H_{\mathcal{L}}$. There exist some constant $\delta \in (0, 1)$, $b \in \mathbb{R}$ and a measurable function $\mathcal{L} : \mathcal{S} \rightarrow [1, +\infty)$ s.t. $\lim_{|x| \rightarrow \infty} \mathcal{L}(x) = +\infty$, and an accessible small set $\mathcal{C} \subset \mathcal{B}(\mathcal{S})$, such that for all $V_1 \in \mathcal{S}$,

$$\mathbb{E} [\mathcal{L}(V_2) | V_1] \leq \delta \mathcal{L}(V_1) + b \mathbf{1}_{\mathcal{C}}(V_1).$$

$H_{p_{\max}}$. Under $H_{\mathcal{L}}$, the growth of \mathcal{L} at infinity is polynomial of order p_{\max} : $\exists c_{\mathcal{L}}, C_{\mathcal{L}} > 0$, $c_{\mathcal{L}} |x|^{p_{\max}} \leq \mathcal{L}(x) \leq C_{\mathcal{L}}(1 + |x|^{p_{\max}})$.

H_x . $\{\mathbf{x}_t\}_{t \in \mathbb{N}^*}$ denotes a portfolio allocation process, e.g. a \mathcal{F}_t -measurable process with values in \mathbb{W} , hence bounded: $\mathbb{P}(|x_{i,t}| > c_w) = 0$ for $i \in \{1, \dots, d\}$.

As will be shown later in the proof, these assumptions are compatible with the existence of moment $p_{\max} < \infty$.

3.2.3 Main results

We will call *performance gap* the quantity ℓ_H defined by the expected value of $\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N}$ under the stationary law:

$$\ell_H := \mathbb{E} \left[\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N} \middle| V_1 \sim \mu \right].$$

It is a key quantity since it can be interpreted as a performance gap between the benchmark and the estimated conditional realized covariance portfolio.

Let us first state a result on the sign of ℓ_H .

Proposition 3.2.1 (Non-negativity of the performance gap). *Assume H_{stat} , $H_{\mathcal{S}}$, $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Then, ℓ_H is finite, deterministic, and non-negative.*

Our main result is the following:

Theorem 3.2.2. *Assume that $\{V_t, t \in \mathbb{N}^*\}$ satisfies assumptions H_{stat} , $H_{\mathcal{S}}$, $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Assume that ℓ_H is strictly positive, then for any $N, H \in \mathbb{N}^*$, the processes $R_{N,H}, R_{N,H}^{\text{ref}}$ defined in (3.2.3) satisfy:*

$$\mathbb{P} \left(R_{N,H} > R_{N,H}^{\text{ref}} \right) \leq \frac{C}{\ell_H^{p_{\max}}} \frac{H^{\frac{p_{\max}}{2}}}{N^{\frac{p_{\max}}{2}}},$$

where $C > 0$ depends on H , p_{\max} , d and \mathcal{L} .

Comments:

- This inequality is of the form of an upper bound on the probability that the realized variance is higher (e.g. worse) for the conditional covariance-based portfolio than for the reference covariance-based portfolio. It is a probability bound on the *underperformance* of the estimated covariance based-portfolio versus the reference portfolio.
- It is polynomially decreasing in N , the number of rebalancing times of the portfolio, and in ℓ_H , the performance gap, and at least polynomially increasing in H , the investment horizon. As we show in the proof of Theorem 3.2.2, the constant C is linear in a quantity $C_{FM}^{(H)}$ on which we provide a bound in Proposition 3.3.6 which has an exponential growth in H and p_{\max} .

Interpretations:

- As mentioned before, ℓ_H can be interpreted as a performance gap between the benchmark and the estimated conditional realized covariance portfolio. The higher the ℓ_H , the more discriminant the impact of using the estimated realized covariance.

- When $CH/(N\ell_H^2)$ is large, the bound is uninformative. If $N \ll CH/\ell_H^2$, there is not enough observations to statistically distinguish which covariance matrix brings the best performance.

As we show in Lemma 3.2.3, the average performance gap ℓ_H/H , when using the estimated realized covariance versus the stationary covariance V_∞ , goes to zero when H goes to infinity.

It means that in the $(NH) \ll C(H/\ell_H)^2$ regime, the asset manager should not bother much using a sophisticated estimation for V_t : a good approximation of V_∞ is enough.

- When N tends to infinity, the probability goes to zero: this is a concentration of measure effect, and since the expected realized variance difference ℓ_H is non-negative (see Proposition 3.2.1,) the probability of having a negative empirical difference goes quickly towards zero. It is coherent with the intuition than with more data, the historical measure of the realized variance difference will be more likely to be of the same sign than its expected value.

When H goes to infinity, the average conditional realized covariance converges towards the mean-value of the process: $\frac{cRC_{H,T_{reb}}}{H} \rightarrow V_\infty$. If $V_{ref} = V_\infty$, we expect that $\lim_{H \rightarrow \infty} \frac{\ell_H}{H} = 0$. This is what is stated in the next Lemma.

Lemma 3.2.3 (Convergence of $\frac{\ell_H^\infty}{H}$ to zero). *Assume that $\{V_t, t \in \mathbb{N}^*\}$ satisfies assumptions H_{stat} , H_S , H_C and H_{pmax} , and let $V_{ref} = V_\infty$. In that case, denote ℓ_H^∞ the performance gap. Then*

$$\lim_{H \rightarrow \infty} \frac{\ell_H^\infty}{H} = 0.$$

3.3 Proofs and auxiliary results

In this section we start by stating auxiliary concentration results needed to prove our main Theorem 3.2.2 and we detail the proof of the theorem. Then the proofs of the auxiliary results are provided.

3.3.1 Proof of Theorem 3.2.2

To show this concentration of measures result, we will need the following auxiliary results (whose proofs are postponed to Subsections 3.3.2, 3.3.3 and 3.3.4).

Let us first give a hint of the result by rewriting the realized variance difference $R_{N,H}^{ref} - R_{N,H}$. From (3.2.3), the realized variance difference breaks down into:

$$R_{N,H}^{ref} - R_{N,H} = \underbrace{R_{N,H}^{ref} - cRV_{N,H}^{ref}}_{D_{N,H}^{ref}} - \underbrace{(R_{N,H} - cRV_{N,H})}_{D_{N,H}} + \underbrace{cRV_{N,H}^{ref} - cRV_{N,H}}_{\mathcal{E}_{N,H}}. \quad (3.3.1)$$

From the processes definition, it is easy to see that:

- by the definition of $mv(\cdot)$: $(mv(V_{ref})^\top cRC_{H,t_n} mv(V_{ref})) \geq (mv(cRC_{H,t_n})^\top cRC_{H,t_n} mv(cRC_{H,t_n}))$ so $\mathcal{E}_{N,H} = cRV_{N,H}^{ref} - cRV_{N,H} \geq 0$ almost surely,
- by the definition of cRC_{H,t_n} , $\mathbb{E}[RC_{H,t_n} | \mathcal{F}_{t_n}] = cRC_{H,t_n}$, so $D_{N,H}^{ref} = R_{N,H}^{ref} - cRV_{N,H}^{ref}$ and $D_{N,H} = R_{N,H} - cRV_{N,H}$ are \mathcal{F}_{t_n} -centered martingales which concentrate around zero.

The following result gives a concentration of measure result on $\mathcal{E}_{N,H}$:

Proposition 3.3.1 (Concentration of measure for the ergodic conditional realized variance). *Assume that $\{V_t, t \in \mathbb{N}^*\}$ satisfies assumptions H_{stat} , H_S , $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Then, for any $2 \leq q \leq p_{\max}$, there exists $C_{q,\mathcal{L},H,d} > 0$ such that:*

$$\mathbb{E}[|\mathcal{E}_{N,H} - N\ell_H|^q] \leq C_{q,\mathcal{L},H,d}(NH)^{q/2}.$$

A direct application via the Markov inequality gives, for any $a > 0$:

$$\mathbb{P}\left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H > a\right) \leq C_{q,\mathcal{L},H,d} \left(\frac{H}{Na^2}\right)^{\frac{q}{2}}$$

and

$$\mathbb{P}\left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H < -a\right) \leq C_{q,\mathcal{L},H,d} \left(\frac{H}{Na^2}\right)^{\frac{q}{2}}. \quad (3.3.2)$$

We can also show that ℓ_H is non-negative and finite, and as a consequence of the previous Proposition, $\frac{\mathcal{E}_{N,H}}{N}$ can be shown to converge towards ℓ_H :

Proposition 3.3.2 (Convergence of the average conditional realized variance). *Assume H_{stat} , H_S , $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Then, for $\ell_H = \mathbb{E}\left[\frac{cRV_{N,H}^{ref} - cRV_{N,H}}{N} \middle| V_1 \sim \mu\right]$,*

- $\frac{cRV_{N,H}^{ref} - cRV_{N,H}}{N}$ converges in $L_{p_{\max}}$ norm towards ℓ_H :

$$\frac{cRV_{N,H}^{ref} - cRV_{N,H}}{N} \xrightarrow[N \rightarrow \infty]{L_{p_{\max}}} \ell_H.$$

- if $p_{\max} > 2$, $\frac{cRV_{N,H}^{ref} - cRV_{N,H}}{N}$ converges almost surely towards ℓ_H :

$$\frac{cRV_{N,H}^{ref} - cRV_{N,H}}{N} \xrightarrow[N \rightarrow \infty]{a.s.} \ell_H.$$

We now state the concentration of measure results for the martingale processes of type $D_{N,H}$ and $D_{N,H}^{ref}$.

Lemma 3.3.3. *Assume that $\{V_t, t \in \mathbb{N}^*\}$ satisfies (3.2.1), H_{stat} , H_S , $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Let $\{\mathbf{x}_{t_n}\}_{\substack{n \in \mathbb{N}^*, \\ t_{n+1} - t_n = H}}$ be a portfolio allocation satisfying H_x and $D_{N,H}(\mathbf{x}) :=$*

$$\sum_{n=1}^N \mathbf{x}_{t_n}^\top (RC_{H,t_n} - cRC_{H,t_n}) \mathbf{x}_{t_n}.$$

Then, there exists $C_{p_{\max},d,\mathcal{L}} > 0$ such that:

$$\mathbb{E}[|D_{N,H}(\mathbf{x})|^{p_{\max}}] \leq C_{p_{\max},d,\mathcal{L}}(NH)^{\frac{p_{\max}}{2}}.$$

A direct application of the Markov inequality gives, for any $a > 0$:

$$\mathbb{P}\left(\frac{D_{N,H}(\mathbf{x})}{N} > a\right) \leq C_{p_{\max},d,\mathcal{L}} \left(\frac{H}{Na^2}\right)^{\frac{p_{\max}}{2}}, \quad (3.3.3)$$

and

$$\mathbb{P}\left(\frac{D_{N,H}(\mathbf{x})}{N} < -a\right) \leq C_{p_{\max},d,\mathcal{L}} \left(\frac{H}{Na^2}\right)^{\frac{p_{\max}}{2}}. \quad (3.3.4)$$

We can now move to the proof of our main result.

Proof of Theorem 3.2.2. We aim at showing that $R_{N,H} < R_{N,H}^{\text{ref}}$ with high probability.

From equation (3.3.1), we see that the realized variance difference $R_{N,H}^{\text{ref}} - R_{N,H}$ boils down to the sum of the martingales $D_{N,H}^{\text{ref}} = R_{N,H}^{\text{ref}} - cRV_{N,H}^{\text{ref}}$ and $-D_{N,H} = cRV_{N,H} - R_{N,H}$ and of the ergodic term $\mathcal{E}_{N,H} = cRV_{N,H}^{\text{ref}} - cRV_{N,H}$:

$$\begin{aligned} R_{N,H}^{\text{ref}} - R_{N,H} &= D_{N,H}^{\text{ref}} - D_{N,H} + \mathcal{E}_{N,H} \\ &= N \left(\frac{D_{N,H}^{\text{ref}}}{N} - \frac{D_{N,H}}{N} + \left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H \right) + \ell_H \right). \end{aligned}$$

From this equality, we see that if each of the three first terms is higher than $-\frac{\ell_H}{3}$, then the sum plus ℓ_H is non-negative:

$$\left. \begin{array}{l} \frac{D_{N,H}^{\text{ref}}}{N} \geq -\frac{\ell_H}{3} \\ -\frac{D_{N,H}}{N} \geq -\frac{\ell_H}{3} \\ \frac{\mathcal{E}_{N,H}}{N} - \ell_H \geq -\frac{\ell_H}{3} \end{array} \right\} \Rightarrow R_{N,H}^{\text{ref}} - R_{N,H} \geq 0.$$

This translates into the following inclusion of events:

$$\left\{ \left\{ \frac{D_{N,H}}{N} \leq \frac{\ell_H}{3} \right\} \cap \left\{ -\frac{D_{N,H}^{\text{ref}}}{N} \leq \frac{\ell_H}{3} \right\} \cap \left\{ -\left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H \right) \leq \frac{\ell_H}{3} \right\} \right\} \subset \left\{ R_{N,H} \leq R_{N,H}^{\text{ref}} \right\}.$$

Taking the complementary, we see that the event $\left\{ R_{N,H} > R_{N,H}^{\text{ref}} \right\}$ is included in an union of low probability events:

$$\left\{ R_{N,H} > R_{N,H}^{\text{ref}} \right\} \subset \left\{ \left\{ \frac{D_{N,H}}{N} > \frac{\ell_H}{3} \right\} \cup \left\{ \frac{D_{N,H}^{\text{ref}}}{N} < -\frac{\ell_H}{3} \right\} \cup \left\{ \left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H \right) < -\frac{\ell_H}{3} \right\} \right\}.$$

Hence we can bound the probability of $\left\{ R_{N,H} > R_{N,H}^{\text{ref}} \right\}$ by the sum of the probability of the three events, on which we know explicit bounds via Lemma 3.3.3 and Proposition 3.3.1.

Bound on the long horizon martingale term: By Lemma 3.3.3, taking $a = \frac{\ell_H}{3}$ (which is positive by assumption) in equation (3.3.3), with $\mathbf{x}_{t_n} = \text{mv}(cRC_{H,t_n})$, we have:

$$\mathbb{P} \left(\frac{D_{N,H}}{N} > \frac{\ell_H}{3} \right) \leq C_{p_{\max}, d, \mathcal{L}} \left(\frac{9H}{N\ell_H^2} \right)^{\frac{p_{\max}}{2}}.$$

From equation (3.3.4), with $\mathbf{x}_{t_n} = \text{mv}(V_{\text{ref}})$,

$$\mathbb{P} \left(\frac{D_{N,H}^{\text{ref}}}{N} < -\frac{\ell_H}{3} \right) \leq C_{p_{\max}, d, \mathcal{L}} \left(\frac{9H}{N\ell_H^2} \right)^{\frac{p_{\max}}{2}}.$$

Bound on the long horizon ergodic term: From Proposition 3.3.1, replacing a by $\frac{\ell_H}{3}$ in equation (3.3.2) we have:

$$\mathbb{P} \left(\frac{\mathcal{E}_{N,H}}{N} - \ell_H < -\frac{\ell_H}{3} \right) \leq C_{p_{\max}, \mathcal{L}, H, d} \left(\frac{9H}{N\ell_H^2} \right)^{\frac{p_{\max}}{2}}.$$

By union bound, we conclude:

$$\mathbb{P} \left(R_{N,H} > R_{N,H}^{\text{ref}} \right) \leq 3^{p_{\max}} (2C_{p_{\max}, d, \mathcal{L}} + C_{p_{\max}, \mathcal{L}, H, d}) \left(\frac{H}{N\ell_H^2} \right)^{\frac{p_{\max}}{2}}.$$

□

3.3.2 Proof of Proposition 3.3.1

The non-trivial part of our Proposition 3.3.1 lies in the fact that we want to highlight both the N and H dependence. From the conditional realized variance definition, there is a nested concentration effect, both in N , the number of rebalancing times, and H , the number of days on which the variance is measured, that we can exploit.

3.3.2.1 Preparatory results

In this paragraph, we state the concentration of measure result for ergodic Markov process that we will adapt to show our Proposition. We recall also some Markov chain elements of vocabulary.

Concentration of measure for irreducible aperiodic Markov chain: We recall here the concentration of measure result for irreducible aperiodic Markov chain as stated in Fort and Moulines article [FM03, Proposition 2].

Proposition 3.3.4 ([FM03][Proposition 2]). *Let $\{V_t\}_{t \in \mathbb{N}^*}$ be a ϕ -irreducible aperiodic Markov chain on \mathcal{S} , and let $\mathcal{C} \in \mathcal{B}(\mathcal{S})$ be an accessible petite set. Assume that there exist some constants $\delta \in (0, 1)$, $b < \infty$ and a measurable $\mathcal{L} : \mathcal{S} \rightarrow [1, +\infty)$, bounded on \mathcal{C} , such that*

$$\mathbb{E}[\mathcal{L}(V_2)|V_1] \leq \delta \mathcal{L}(V_1) + b \mathbf{1}_{\mathcal{C}}(V_1), \quad \forall V_1 \in \mathcal{S}. \quad (3.3.5)$$

Let $q \geq 2$. Choose $M > \sup_{\mathcal{C}} \mathcal{L} \vee b / (1 - \delta^{1/q})^q$. Then the set $\{\mathcal{L} \leq M\}$ is ν_m -small with minorizing constant $\epsilon > 0$ and for any Borel function $g : \mathcal{S} \rightarrow \mathbb{R}^{d'}$, $|g| \leq \mathcal{L}^{1/q}$, it holds that for all $V_1 \in \mathcal{S}$, $H \geq 1$,

$$\mathbb{E} \left[\left| \sum_{k=1}^H (g(V_{k+1}) - \mu(g)) \right|^q \middle| V_1 \right] \leq C_{FM} \mathcal{L}(V_1) H^{q/2},$$

where $C_{FM} = C \left(\frac{m+1}{\epsilon} \right)^{q+1} \frac{M^2}{A^{2q}}$, $A = (1 - \delta)^{1/q} - (b/M)^{1/q}$ and C is a constant which depends only on q , and μ is the invariant measure associated to $\{V_t\}_{t \in \mathbb{N}^*}$.

Remark 3. A few remarks on this Proposition and how we will apply it:

- For simplicity, we have denoted C_{FM} the constant $C \left(\frac{m+1}{\epsilon} \right)^{q+1} \frac{M^2}{A^{2q}}$. This constant depends on the Lyapunov condition (3.3.5) (so on \mathcal{L}, δ, b), on the petite set \mathcal{C} , on q and on ν_m and ϵ and on the dynamic of $\{V_t\}_{t \in \mathbb{N}}$. It is hard to quantify it explicitly.
- The proposition is stated for functions g such that $|g| \leq \mathcal{L}^{1/q}$. It can be extended to functions which are bounded in $\mathcal{L}^{1/q}$ norm, e.g. such that: $\left(\sup_{V \in \mathcal{S}} \frac{|g(V)|^q}{\mathcal{L}(V)} \right)^{1/q} < \infty$. Indeed, denoting $\tilde{g} = g / \left(\sup_{V \in \mathcal{S}} \frac{|g(V)|^q}{\mathcal{L}(V)} \right)^{1/q}$, then for any $V' \in \mathcal{S}$,

$$\frac{|\tilde{g}(V')|^q}{\mathcal{L}(V')} = \frac{|g(V')|^q / \mathcal{L}(V')}{\sup_{V \in \mathcal{S}} |g(V)|^q / \mathcal{L}(V)} \leq 1$$

so $|\tilde{g}| \leq \mathcal{L}^{1/q}$. We can apply the proposition on \tilde{g} and express it in terms of g :

$$\mathbb{E} \left[\left| \sum_{k=1}^H (g(V_{k+1}) - \mu(g)) \right|^q \middle| V_1 \right] \leq C_{FM} \left(\sup_{V \in \mathcal{S}} \frac{|g(V)|^q}{\mathcal{L}(V)} \right) \mathcal{L}(V_1) H^{q/2}.$$

Markov chain elements of vocabulary: We recall the definition of *time of first return*, *irreducibility*, *petite set*, *small set*, *aperiodicity* and *accessibility*, which can be found in the Meyn and Tweedie's Book [MT09, pages 71, 82, 117, 102, 114 and 86]:

Let $\{X_n, n \in \mathbb{N}\}$ be a Markov chain on a state space \mathcal{S} , P its transition probability. Let $A \subset \mathcal{B}(\mathcal{S})$. We denote $\tau_A := \min\{n \geq 1 : X_n \in A\}$ the *first return* time on A and for $x \in \mathcal{S}$, $L(x, A) := \mathbb{P}(\tau_A < \infty | X_0 = x)$ the probability to access A from a specific x . Given ϕ a Borelian measure, we say that X_n is ϕ -irreducible if $\phi(A) > 0 \Rightarrow L(y, A) > 0$ for any $y \in \mathcal{S}$.

A *petite set* is a set $\mathcal{C} \in \mathcal{B}(\mathcal{S})$ such that there is a probability distribution a on \mathbb{N} and a non-trivial measure ν_a such that $\forall x \in \mathcal{C}, \forall A \subset \mathcal{B}(\mathcal{S}), \sum_{n=0}^{\infty} a(n)P^n(x, A) \geq \nu_a(A)$.

A *small set* is a particular case of a petite set in which a only charges a specific $m \in \mathbb{N}$: the definition becomes: $P^m(x, A) \geq \nu_m(A)$. When there exists a small set with $m = 1$ and $\nu_1(\mathcal{C}) > 0$, then the chain is called *strongly aperiodic*.

A set $A \subset \mathcal{B}(\mathcal{S})$ is said *accessible* if it can be accessed from any $x \in \mathcal{S}$: $L(x, A) > 0, \forall x \in \mathcal{S}$.

Concentration of measure for function of lagged Markov chain: motivation Notice that since $\{V_{t+1}, t \in \mathbb{N}^*\}$ is Markovian, the random variable $cRC_{H,t} = \mathbb{E}\left[\sum_{k=1}^H V_{t+k} | \mathcal{F}_t\right]$ can be seen as a function of the \mathcal{F}_t -measurable V_{t+1} : $cRC_{H,t} = \phi(V_{t+1})$. Hence $\mathcal{E}_{N,H}$ can be identically written:

$$\begin{aligned} \mathcal{E}_{N,H} &= \sum_{n=1}^N \text{mv}(V_{\text{ref}})^\top cRC_{H,t_n} \text{mv}(V_{\text{ref}}) - \text{mv}(cRC_{H,t_n})^\top cRC_{H,t_n} \text{mv}(cRC_{H,t_n}) \\ &= \sum_{n=1}^N [\text{mv}(V_{\text{ref}}) - \text{mv}(cRC_{H,t_n})]^\top cRC_{H,t_n} [\text{mv}(V_{\text{ref}}) + \text{mv}(cRC_{H,t_n})] \\ &= \sum_{n=1}^N \mathbf{x}_{t_n}^\top \phi(V_{t_n+1}) \mathbf{y}_{t_n}, \end{aligned} \quad (3.3.6)$$

where

$$\begin{aligned} \mathbf{x}_{t_n} &= \text{mv}(V_{\text{ref}}) - \text{mv}(cRC_{H,t_n}), \\ \mathbf{y}_{t_n} &= \text{mv}(V_{\text{ref}}) + \text{mv}(cRC_{H,t_n}). \end{aligned} \quad (3.3.7)$$

In the following result, we show that the concentration of measure result Proposition 3.3.4 remains valid for function of the lagged chain: $g(V_k) \rightarrow g(V_{kH})$.

Proposition 3.3.5. [Fort Moulines proposition extension to lagged-Markov chains]

Assume that $\{V_{t+1}, t \in \mathbb{N}^*\}$ satisfies assumptions H_{stat} , $H_{\mathcal{S}}$, $H_{\mathcal{L}}$ and $H_{p_{\max}}$. Then for any $q \geq 2$, for any Borel function $g : \mathcal{S} \rightarrow \mathbb{R}^{d \times d}$ bounded in $\mathcal{L}^{1/q}$ -norm, $T \in \mathbb{N}^*$, $H \in \mathbb{N}^*$, we have:

$$\mathbb{E} \left[\left| \sum_{t=1}^T (g(V_{tH+1}) - \mu(g)) \right|^q \middle| V_1 \right] \leq C_{FM}^{(H)} \left(\sup_{\mathcal{S}} \frac{|g|^q}{\mathcal{L}} \right) \mathcal{L}(V_1) T^{q/2}, \quad (3.3.8)$$

where $C_{FM}^{(H)}$ is a constant on which we provide a bound in H and q in Proposition 3.3.6.

Proof. Let $H \in \mathbb{N}^*$. First, let us show that the lagged Markov chain $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ satisfies the assumptions of Proposition 3.3.4.

1. Irreducibility and aperiodicity of the lagged Markov chain

$\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ is a Lebesgue-irreducible aperiodic Markov chain on \mathcal{S} by application of Proposition 3.C.3 ([MT09, Proposition 5.4.5]: extension of irreducibility and aperiodicity of irreducible and aperiod chains) to the chain $\{V_{t+1}\}_{t \in \mathbb{N}^*}$ which is Lebesgue-irreducible and aperiodic by assumption H_{stat} .

2. Lyapunov condition on the lagged Markov chain

In Meyn and Tweedie's book, Theorem 3.C.1 ([MT09, Theorem 15.3.4]) states that if $\{V_{t+1}\}_{t \in \mathbb{N}^*}$ satisfies a drift condition $\mathbf{H}_{\mathcal{L}}$ with a Lyapunov function \mathcal{L} and a petite set \mathcal{C} , then $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ also satisfies a drift condition with the same Lyapunov function \mathcal{L} and some set $\mathcal{C}^{(H)}$ which is petite for the H -skeleton.

In our Proposition 3.C.2, we give a quantitative assessment of the Lyapunov condition for $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ by giving explicitly the constants d_H and b_H s.t.:

$$\mathbb{E} [\mathcal{L}(V_{H+1}) | V_1] \leq d_H \mathcal{L}(V_1) + b_H \mathbb{1}_{\mathcal{C}^{(H)}}(V_1).$$

Our computation yields:

$$d_H = \frac{1 + \delta^H}{2}, \quad \mathcal{C}^{(H)} = \mathcal{C} \cup \left\{ x \in \mathcal{S} \mid |x| \leq R^{(H)} \right\},$$

$$b_H = \sup_{x \in \mathcal{C}^{(H)}} \left(b \frac{1 - \delta^H}{1 - \delta} - \frac{1 - \delta^H}{2} \mathcal{L}(x) \right),$$

and $R^{(H)} > 0$ s.t. for every $x \in \mathcal{S}$ s.t. $|x| > R^{(H)}$, $\frac{1 - \delta^H}{2} \mathcal{L}(x) \geq b \frac{1 - \delta^{H-1}}{1 - \delta}$.

3. Smallness of the set $\mathcal{C}^{(H)}$

By irreducibility and aperiodicity, petite sets are also small sets [MT09, Theorem 5.5.7] Since \mathcal{C} is included in the new one $\mathcal{C}^{(H)}$ (whether in Meyn and Tweedie's book [MT09, Lemma 14.2.8] and in our Proposition 3.C.2), it will be enough to show that \mathcal{C} is a small set for $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$. A key assumption to prove this result will be that we can consider a measure associated to \mathcal{C} for $\{V_{t+1}\}_{t \in \mathbb{N}^*}$ satisfying $\nu(\mathcal{C}) > 0$.

Indeed, let us denote ν_m and $\epsilon > 0$ the measure and minorizing constant s.t. $\nu_m(\mathcal{C}) > 0$ (w.l.o.g. we can assume that $\nu_m(\mathcal{C}) = 1$) and $P^m(x, A) \geq \epsilon \nu_m(A)$ for any $x \in \mathcal{C}$ and $A \subset \mathcal{B}(\mathcal{S})$. They exist by application of [MT09, Proposition 5.2.4 - (iii)] (existence of a measure positive on the small set for irreducible chains).

Let us show that \mathcal{C} is a small set for $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$. Indeed, we can show recursively that $P^{mk}(x, A) \geq \epsilon^k \nu_m(A)$ for any $k \in \mathbb{N}^*$. It is satisfied for $k = 1$. Let us assume the property for $k \in \mathbb{N}^*$. Then for $k + 1$:

$$\begin{aligned} P^{(k+1)m}(x, A) &= \int_{\mathcal{S}} P^{mk}(x, dy) P^m(y, A) \geq \epsilon^k \int_{\mathcal{S}} \nu_m(dy) P^m(y, A) \\ &\geq \epsilon^k \int_{\mathcal{C}} \nu_m(dy) \underbrace{P^m(y, A)}_{\substack{\geq \epsilon \nu_m(A) \\ \text{since } y \in \mathcal{C}}} \quad (3.3.9) \\ &\geq \epsilon^{k+1} \underbrace{\nu_m(\mathcal{C})}_{=1} \nu_m(A) = \epsilon^{k+1} \nu_m(A). \end{aligned}$$

So for $k = H$, \mathcal{C} is a small set for $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$. To show that $\mathcal{C}^{(H)}$ is a small set for $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$, it suffices to take $\nu_m^{(H)} = \nu_m$ on \mathcal{C} and $\nu_m^{(H)} = 0$ on $\mathcal{C}^{(H)} \setminus \mathcal{C}$.

4. Accessibility of the set

Since $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ is Lebesgue-irreducible and since $\lambda(\mathcal{C}^{(H)}) > 0$ (because we have $\{x \leq R^{(H)}\} \subset \mathcal{C}^{(H)}$ with $R^{(H)} > 0$ so $\lambda(\mathcal{C}^{(H)}) \geq \lambda(\{|x| \leq R^{(H)}\}) > 0$) then by the irreducibility definition, $L(x, \mathcal{C}^{(H)}) > 0$ for any $x \in \mathcal{S}$ so $\mathcal{C}^{(H)}$ is accessible.

5. Boundedness of the Lyapunov function

\mathcal{L} is bounded on $\mathcal{C}^{(H)}$ because $\mathcal{C}^{(H)} \subsetneq \mathcal{S}$ is bounded and we have assumed \mathcal{L} of polynomial growth.

The assumptions of proposition 3.3.4 are verified so by application of the proposition to the lagged chain $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$ the announced inequality (3.3.8) is true. \square

Parameters dependence in H : In this paragraph, we are going to exhibit the dependence of $C_{FM}^{(H)}$ in H and q . From Proposition 3.3.4,

$$C_{FM}^{(H)} = C \left(\frac{m_H + 1}{\epsilon_H} \right)^{q+1} \frac{M_H^2}{A_H^{2q}},$$

where A_H, M_H, m_H and ϵ_H depend on the Lyapunov constants and set d_H, b_H and $\mathcal{C}^{(H)}$.

Alternative choice of d_H, b_H and $\mathcal{C}^{(H)}$ uniform in H : Let us show that we can find d_{\max}, b_{\max} and \mathcal{C}^{\max} independent from H s.t.

$$\mathbb{E} [\mathcal{L}(V_{H+1}) | V_1 = x] \leq d_{\max} \mathcal{L}(x) + b_{\max} \mathbb{1}_{x \in \mathcal{C}^{\max}}.$$

From Proposition 3.C.2's proof, $\mathbb{E} [\mathcal{L}(V_{H+1}) | V_1 = x]$ satisfies the inequality (3.C.1):

$$\mathbb{E} [\mathcal{L}(V_{H+1}) | V_1 = x] \leq \delta^H \mathcal{L}(x) + b \frac{1 - \delta^{H-1}}{1 - \delta} + \delta^{H-1} b \mathbb{1}_{\mathcal{C}}(x).$$

We can bound this inequality uniformly in H :

$$\begin{aligned} \mathbb{E} [\mathcal{L}(V_{H+1}) | V_1 = x] &\leq \delta \mathcal{L}(x) + \frac{b}{1 - \delta} + b \mathbb{1}_{\mathcal{C}}(x) \\ &= \frac{1 + \delta}{2} \mathcal{L}(x) - \frac{1 - \delta}{2} \mathcal{L}(x) + \frac{b}{1 - \delta} + b \mathbb{1}_{\mathcal{C}}(x). \end{aligned}$$

Taking R^{\max} s.t. $\frac{b}{1 - \delta} + b < \frac{1 - \delta}{2} \mathcal{L}(x)$ for $|x| > R^{\max}$ and $\mathcal{C}^{\max} = \mathcal{C} \cup \{x \in \mathcal{S} \mid |x| \leq R^{\max}\}$,

$$\mathbb{E} [\mathcal{L}(V_{H+1}) | V_1 = x] \leq \underbrace{\frac{1 + \delta}{2} \mathcal{L}(x)}_{d^{\max}} + \underbrace{\sup_{x \in \mathcal{C}^{\max}} \left(\frac{b}{1 - \delta} + b - \frac{1 - \delta}{2} \mathcal{L}(x) \right)}_{b^{\max}} \mathbb{1}_{x \in \mathcal{C}^{\max}}.$$

Choice of M_H and A_H uniform in H : By definition,

$$\begin{aligned} M_H &> \sup_{\mathcal{C}^{(H)}} \mathcal{L} \vee \left(b_H / \left(1 - d_H^{\frac{1}{q}} \right)^q \right), \\ A_H &= (1 - d_H)^{\frac{1}{q}} - \left(\frac{b_H}{M_H} \right)^{\frac{1}{q}}. \end{aligned} \tag{3.3.10}$$

We can replace b_H, d_H and $\mathcal{C}^{(H)}$ by b_{\max}, d_{\max} and \mathcal{C}^{\max} in (3.3.10) to have M_H and A_H uniform in H .

Behavior of m_H and ϵ_H : By application of Proposition 3.3.4 to $\{V_{t+1}, t \in \mathbb{N}\}$, $\{\mathcal{L} \leq M\}$ is a ν_m -small set with minorizing constant ϵ . W.l.o.g., we can assume that $M_{\max} \geq M$. Then, as we have shown in the proof of Proposition 3.3.5 equation (3.3.9), a ν_m -small set with minorizing constant ϵ for $\{V_{t+1}, t \in \mathbb{N}\}$ is a ν_m -small set with minorizing constant $\epsilon_H := \epsilon^H$ for $\{V_{tH+1}, t \in \mathbb{N}\}$. As we did in the proof of Proposition 3.3.5, we can take $\nu_{m_H} = \nu_m$ on $\{\mathcal{L} \leq M\}$ and null on $\{\mathcal{L} \leq M^{\max}\} \setminus \{\mathcal{L} \leq M\}$ so that $\{\mathcal{L} \leq M^{\max}\}$ is a ν_{m_H} small set for $\{V_{t+1}, t \in \mathbb{N}\}$ with minorizing constant ϵ^H .

We give our conclusion in the following proposition:

Proposition 3.3.6. *Under the assumption of Proposition 3.3.5, the constant $C_{FM}^{(H)}$ can be upper bounded by:*

$$C \left(\frac{m+1}{\epsilon^H} \right)^{q+1} \frac{M_{\max}(q)^2}{A_{\max}(q)^{2q}}$$

where $M_{\max}(q) > \sup_{\mathcal{C}(m_{ax})} \mathcal{L} \vee \left(b_{\max} / \left(1 - d_{\max}^{\frac{1}{q}} \right)^q \right)$, and $A_{\max}(q) := (1 - d_{\max})^{\frac{1}{q}} - \left(\frac{b_{\max}}{M_{\max}(q)} \right)^{\frac{1}{q}}$, R^{\max} s.t. $\frac{b}{1-\delta} + b < \frac{1-\delta}{2} \mathcal{L}(x)$, $\mathcal{C}^{\max} = \mathcal{C} \cup \{x \in \mathcal{S} \mid |x| \leq R^{\max}\}$, $d^{\max} = \frac{1+\delta}{2}$, $b^{\max} = \sup_{x \in \mathcal{C}^{\max}} \left(\frac{b}{1-\delta} + b - \frac{1-\delta}{2} \mathcal{L}(x) \right)$ and ϵ, m are given by the application of Proposition 3.3.4 on $\{V_{t+1}, t \in \mathbb{N}\}$, and C depends on q only.

Remark 4. This bound goes to infinity when H and q go to infinity, because

- $\epsilon \in (0, 1)$ so $1/\epsilon^H$ goes to infinity when H goes to infinity,
- $\left(1 - d_{\max}^{\frac{1}{q}} \right)$ goes to 0 when q goes to infinity, so $M_{\max}(q)$ goes to infinity when q goes to infinity.

It means that the control becomes loose when H and q are too large. However, this bound is a uniform bound which can be far from the minimal constant one could get with a more refined analysis.

3.3.2.2 Completion of Proof of Proposition 3.3.1

We can now prove our Proposition 3.3.1.

Proof. Let $q \in [2, p_{\max}]$, let $\mathcal{E}_{N,H}$ be defined by (3.3.6):

$$\mathcal{E}_{N,H} = \sum_{n=1}^N \mathbf{x}_{t_n}^\top \phi(V_{t_n+1}) \mathbf{y}_{t_n} = \sum_{n=1}^N g(V_{t_n+1}),$$

$$\phi(V_{t_n+1}) = \mathbb{E} \left[\sum_{k=1}^H V_{t_n+k} \mid \mathcal{F}_{t_n} \right], \quad g(V_{t_n}) = \mathbf{x}_{t_n}^\top \phi(V_{t_n+1}) \mathbf{y}_{t_n},$$

with \mathbf{x}_{t_n} and \mathbf{y}_{t_n} defined in (3.3.7) by $\mathbf{x}_{t_n} = \text{mv}(V_{\text{ref}}) - \text{mv}(cRC_{H,t_n})$ and $\mathbf{y}_{t_n} = \text{mv}(V_{\text{ref}}) + \text{mv}(cRC_{H,t_n})$. We want to control $\mathbb{E} [|\mathcal{E}_{N,H} - N\ell_H|^q]$ in N and H , where $\ell_H = \mathbb{E} \left[\frac{\mathcal{E}_{N,H}}{N} \mid V_1 \sim \mu \right]$.

In order to apply Proposition 3.3.4 on $\mathcal{E}_{N,H}$, we have to verify that $g(\cdot)$ is bounded in $\mathcal{L}^{1/q}$ norm. From $H_{p_{\max}}$, a function is bounded in $\mathcal{L}^{1/q}$ norm if it can be bounded by a polynomial function of order $\frac{p_{\max}}{q} \geq 1$ (because \mathcal{L} is larger than one and of polynomial growth of order p_{\max} at infinity). We are going to show that $g(\cdot)$ is sub-linear. Hence $g(\cdot)$ will be bounded in $\mathcal{L}^{1/q}$ norm.

Sublinearity of g : By triangle inequality, using that $|x_{i,t}| \leq 2c_w$ and $|y_{i,t}| \leq 2c_w$ by their definition as sum and difference of portfolios in \mathcal{W} , for $t \in \mathbb{N}^*$,

$$|g(V_{t+1})| \leq (2c_w)^2 d \sum_{k=1}^H \mathbb{E} [|V_{t+k}| \mid \mathcal{F}_t].$$

To show that $g(\cdot)$ is sub-linear, we will show that each $\mathbb{E} [|V_{t+k}| \mid V_{t+1}]$ is sub-linear (e.g. linearly bounded with respect to $|V_{t+1}|$). By Markov property, $\mathbb{E} [|V_{t+k}| \mid \mathcal{F}_t] = \mathbb{E} [|V_{t+k}| \mid V_{t+1}]$, the sub-linearity of $g(\cdot)$ will ensue.

Let $k \in \mathbb{N}^*$. By the Jensen inequality in (*) and $\mathbf{H}_{p_{\max}}$:

$$\mathbb{E} [|V_{t+k}| | V_{t+1}] \stackrel{(*)}{\leq} \left(\mathbb{E} [|V_{t+k}|^{p_{\max}} | V_{t+1}] \right)^{\frac{1}{p_{\max}}} \stackrel{\mathbf{H}_{p_{\max}}}{\leq} \left(\frac{1}{c_{\mathcal{L}}} \mathbb{E} [\mathcal{L}(V_{t+k}) | V_{t+1}] \right)^{\frac{1}{p_{\max}}}.$$

By the extended Lyapunov condition (3.C.1) and $\mathbf{H}_{p_{\max}}$:

$$\begin{aligned} \mathbb{E} [\mathcal{L}(V_{t+k}) | V_{t+1}] &\leq \delta^{k-1} \mathcal{L}(V_{t+1}) + b \underbrace{\frac{1 - \delta^{k-2}}{1 - \delta}}_{b^{(k)}} + \delta^{k-2} b \\ &\leq \delta^{k-1} (C_{\mathcal{L}}(1 + |V_{t+1}|^{p_{\max}})) + b^{(k)}. \end{aligned}$$

Combining these inequalities and applying the inequality $(|x| + |y|)^{\frac{1}{p_{\max}}} \leq |x|^{\frac{1}{p_{\max}}} + |y|^{\frac{1}{p_{\max}}}$,

$$\begin{aligned} \mathbb{E} [|V_{t+k}| | \mathcal{F}_t] &\leq \left(\frac{1}{c_{\mathcal{L}}} (\delta^{k-1} C_{\mathcal{L}} |V_{t+1}|^{p_{\max}}) + \delta^{k-1} C_{\mathcal{L}} + b^{(k)} \right)^{\frac{1}{p_{\max}}} \\ &\leq \left(\frac{\delta^{k-1} C_{\mathcal{L}}}{c_{\mathcal{L}}} \right)^{\frac{1}{p_{\max}}} |V_{t+1}| + \left(\frac{b^{(k)} + \delta^{k-1} C_{\mathcal{L}}}{c_{\mathcal{L}}} \right)^{\frac{1}{p_{\max}}}. \end{aligned}$$

Hence, $\mathbb{E} [|V_{t+k}| | V_{t+1}]$ is sub-linear, and by linear combination, so is $g(\cdot)$.

Proposition 3.3.5 application: dependence in N : Since $g(\cdot)$ is bounded in $\mathcal{L}^{1/q}$ norm, we can apply the concentration of measure result for lagged chain Proposition 3.3.5 on $\mathcal{E}_{N,H}$:

$$\mathbb{E} \left[|\mathcal{E}_{N,H} - N \ell_H|^q \middle| V_1 \right] \leq C_{FM}^{(H)} \left(\sup_{\mathcal{S}} \frac{|g - \mu(g)|^q}{\mathcal{L}} \right) \mathcal{L}(V_1) N^{q/2}. \quad (3.3.11)$$

Bound on $\sup_{\mathcal{S}} \frac{|g - \mu(g)|^q}{\mathcal{L}}$ in terms of H : Let $v \in \mathcal{S}$. By $g(\cdot)$ definition:

$$\begin{aligned} g(v) - \mu(g) &= \text{mv}(V_{\text{ref}})^{\top} \phi(v) \text{mv}(V_{\text{ref}}) - \mathbb{E}_{\mu} \left[\text{mv}(V_{\text{ref}})^{\top} \phi(V) \text{mv}(V_{\text{ref}}) \right] \\ &\quad - \left(\text{mv}(\phi(v))^{\top} \phi(v) \text{mv}(\phi(v)) - \mathbb{E}_{\mu} \left[\text{mv}(V)^{\top} \phi(V) \text{mv}(V) \right] \right). \end{aligned}$$

Let us denote a the first term:

$$a = \text{mv}(V_{\text{ref}})^{\top} (\phi(v) - \mathbb{E}_{\mu} [\phi(V)]) \text{mv}(V_{\text{ref}}).$$

Using that $|\text{mv}(V_{\text{ref}})_i| \leq c_{\mathbf{w}}$, $i \in \{1, \dots, d\}$,

$$|a|^q \leq (c_{\mathbf{w}}^2 d)^q |\phi(v) - \mathbb{E}_{\mu} [\phi(V)]|^q. \quad (3.3.12)$$

We can write the second term as:

$$b = \text{mv}(\phi(v))^{\top} \phi(v) \text{mv}(\phi(v)) - \mathbb{E}_{\mu} \left[\text{mv}(\phi(V))^{\top} \phi(V) \text{mv}(\phi(V)) \right].$$

Notice that, by definition of mv , for any $C, \tilde{C} \in \mathcal{S}_+^d$,

$$|\text{mv}(C)^{\top} C \text{mv}(C) - \text{mv}(\tilde{C})^{\top} \tilde{C} \text{mv}(\tilde{C})| \leq (c_{\mathbf{w}}^2 d)^q |C - \tilde{C}|. \quad (3.3.13)$$

Let $C, \tilde{C} \in \mathcal{S}_+^d$. Let us assume that $\text{mv}(C)^{\top} C \text{mv}(C) - \text{mv}(\tilde{C})^{\top} \tilde{C} \text{mv}(\tilde{C}) \geq 0$. Then

$$\text{mv}(C)^{\top} C \text{mv}(C) - \text{mv}(\tilde{C})^{\top} \tilde{C} \text{mv}(\tilde{C})$$

$$\begin{aligned}
&= \underbrace{\text{mv}(C)^\top C \text{mv}(C) - \text{mv}(\tilde{C})^\top C \text{mv}(\tilde{C})}_{\leq 0} + \underbrace{\text{mv}(\tilde{C})^\top C \text{mv}(\tilde{C}) - \text{mv}(\tilde{C})^\top \tilde{C} \text{mv}(\tilde{C})}_{=\text{mv}(\tilde{C})^\top (C - \tilde{C}) \text{mv}(\tilde{C})} \\
&\leq \text{mv}(\tilde{C})^\top (C - \tilde{C}) \text{mv}(\tilde{C}) \leq (c_{\mathbf{w}}^2 d)^q |C - \tilde{C}|.
\end{aligned}$$

Conversely, if $\text{mv}(\tilde{C})^\top \tilde{C} \text{mv}(\tilde{C}) - \text{mv}(C)^\top C \text{mv}(C) \geq 0$, then we can do the same reasoning inverting C and \tilde{C} :

$$\text{mv}(\tilde{C})^\top \tilde{C} \text{mv}(\tilde{C}) - \text{mv}(C)^\top C \text{mv}(C) \leq \text{mv}(C)^\top (\tilde{C} - C) \text{mv}(C) \leq (c_{\mathbf{w}}^2 d)^q |C - \tilde{C}|.$$

Applying this inequality with $C = \phi(v)$ and $\tilde{C} = \phi(V)$, we can bound $|b|^q$:

$$|b|^q \leq (c_{\mathbf{w}}^2 d)^q |\phi(v) - \mathbb{E}_\mu[\phi(V)]|^q.$$

Hence we arrived at the same bound as (3.3.12).

By double conditioning, and invariance of the stationary law: $V_t | V_1 \sim \mu \stackrel{(d)}{=} \mu$,

$$\mathbb{E}_\mu[\phi(V)] = \mathbb{E}_\mu \left[\mathbb{E} \left[\sum_{k=1}^H V_{t+k} | V_{t+1} = V \right] \right] = \mathbb{E}_\mu \left[\sum_{k=1}^H V_{t+k} \right] = HV_\infty.$$

By the former remark and applying the Jensen inequality,

$$\begin{aligned}
|\phi(v) - \mathbb{E}_\mu[\phi(V)]|^q &= \left| \mathbb{E} \left[\sum_{k=1}^H V_{t+k} - HV_\infty | V_{t+1} = v \right] \right|^q \\
&\leq \mathbb{E} \left[\left| \sum_{k=1}^H V_{t+k} - HV_\infty \right|^q \middle| V_{t+1} = v \right].
\end{aligned}$$

We can separate $\mathbb{E} \left[\left| \sum_{k=1}^H V_{t+k} - HV_\infty \right|^q \middle| V_{t+1} = v \right]$ between the V_{t+1} measurable term and a term of the form sum of function -here the identity function- of V_{t+2}, \dots, V_{t+H} conditioned on V_{t+1} , i.e. in the right form to apply Proposition 3.3.4 :

$$\mathbb{E} \left[\left| \sum_{k=1}^H (V_{t+k} - V_\infty) \right|^q \middle| V_{t+1} = v \right] \leq 2^{q-1} \left(|v - V_\infty|^q + \mathbb{E} \left[\left| \sum_{k=2}^H (V_{t+k} - V_\infty) \right|^q \middle| V_{t+1} = v \right] \right).$$

To summarize, we obtain:

$$|g(v) - \mu(g)|^q \leq 2^{2q-1} (c_{\mathbf{w}}^2 d)^q \left(|v - V_\infty|^q + \mathbb{E} \left[\left| \sum_{k=2}^H (V_{t+k} - V_\infty) \right|^q \middle| V_{t+1} = v \right] \right).$$

Proposition 3.3.4 application, dependence in H : Since $\mathcal{L}(x) \geq c_{\mathcal{L}} |x|^{p_{\max}}$ for large x , and since $\mathcal{L}(x) > 1$ for any $x \in \mathcal{S}$, $\sup_{x \in \mathcal{S}} \frac{|x|^q}{\mathcal{L}(x)}$ is bounded. We can then apply the concentration of measure result for standard chain Proposition 3.3.4 on the function $\tilde{g}(V_{t+k}) = V_{t+k}$ and $\mu(\tilde{g}) = \mathbb{E}_\mu[V_{t+k}] = V_\infty$:

$$\mathbb{E} \left[\left| \sum_{k=2}^H (V_{t+k} - V_\infty) \right|^q \middle| V_{t+1} = v \right] \leq C_{FM} \sup_{v' \in \mathcal{S}} \frac{|v'|^q}{\mathcal{L}(v')} \mathcal{L}(v) (H-1)^{q/2}. \quad (3.3.14)$$

Hence, there exists a constant $c_{q,d}$ depending on q , d and on the Lyapunov condition $\mathbf{H}_{\mathcal{L}}$ such that, for any $H \geq 1$:

$$\begin{aligned} \frac{|g(v) - \mu(g)|^q}{\mathcal{L}(v)} &\leq 2^{2q-1}(c_{\mathbf{w}}^2 d)^q \left(C_{FM} \sup_{v' \in \mathcal{S}} \frac{|v'|^q}{\mathcal{L}(v')} (H-1)^{q/2} + \frac{|v - V_{\infty}|^q}{\mathcal{L}(v)} \right) \\ &\leq c_{q,d} H^{q/2}. \end{aligned}$$

For example, $c_{q,d} = 2^{2q}(c_{\mathbf{w}}^2 d)^q \left(C_{FM} \sup_{v' \in \mathcal{S}} \frac{|v'|^q}{\mathcal{L}(v')} + \sup_{v' \in \mathcal{S}} \frac{|v' - V_{\infty}|^q}{\mathcal{L}(v')} \right)$.

Combining this upper bound with (3.3.11), we obtain:

$$\mathbb{E}[|\mathcal{E}_{N,H} - N\ell_H|^q] \leq C_{q,\mathcal{L},H,d} (NH)^{q/2}.$$

□

3.3.3 Proof of Propositions 3.2.1 and 3.3.2

In this subsection, we are interested in the characterization of $\ell_H = \mathbb{E}_{\mu} \left[\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N} \right]$ and of the limit value of $\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N}$ when N goes to infinity.

Proof. The proof goes as follows: first we show the finiteness and non-negativity of ℓ_H . The convergence result will follow from the concentration of measure result on $\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N} = \frac{\mathcal{E}_{N,H}}{N}$.

- **Finiteness**

ℓ_H is defined as the expectation under the stationary law of a linear combination of quantities of type $\text{mv}(V_{\text{ref}})^{\top} \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}] \text{mv}(V_{\text{ref}})$ and $\text{mv}(cRC_{H,t_n})^{\top} \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}] \text{mv}(cRC_{H,t_n})$, for $k \in \{1, \dots, H\}$ and $n \in \{1, \dots, N\}$.

Since the $\text{mv}(\cdot)$ operator is bounded, and since the $\{V_t\}_{t \in \mathbb{N}^*}$ admit L_1 moment by \mathbf{H}_{stat} , ℓ_H is well defined and finite.

- **Non-negativity**

By definition of the $\text{mv}(\cdot)$ mapping,

$$\text{mv}(V_{\text{ref}})^{\top} cRC_{H,t_n} \text{mv}(V_{\text{ref}}) \geq \text{mv}(cRC_{H,t_n})^{\top} cRC_{H,t_n} \text{mv}(cRC_{H,t_n}) \quad \text{a.s.}$$

Hence, summing on n from 1 to N , we have that $cRV_{N,H}^{\text{ref}} \geq cRV_{N,H}$ almost surely.

Taking the expectation: $0 \leq \mathbb{E}_{\mu} \left[\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N} \right] := \ell_H$.

- **Convergence of $\frac{cRV_{N,H}^{\text{ref}} - cRV_{N,H}}{N}$**

From the moment inequality in Proposition 3.3.1,

$$\mathbb{E} \left[\left| \frac{\mathcal{E}_{N,H}}{N} - \ell_H \right|^{p_{\max}} \right] \leq C_{q,\mathcal{L},H,d} \left(\frac{H}{N} \right)^{\frac{p_{\max}}{2}}.$$

Hence, $\mathbb{E} \left[\left| \frac{\mathcal{E}_{N,H}}{N} - \ell_H \right|^{p_{\max}} \right] \xrightarrow{N \rightarrow +\infty} 0$ hence the convergence in $L_{p_{\max}}$ -norm.

From the Markov inequality in Proposition 3.3.1,

$$\mathbb{P} \left(\left| \frac{\mathcal{E}_{N,H}}{N} - \ell_H \right| > a \right) \leq \frac{2C_q H^{\frac{p_{\max}}{2}}}{N^{\frac{p_{\max}}{2}} a^{p_{\max}}}. \quad (3.3.15)$$

If $p_{\max} > 2$, we can apply the Borel-Cantelli Lemma to show that $\frac{\varepsilon_{N,H}}{N}$ converges almost surely towards ℓ_H . Indeed, since $p_{\max} > 2$, for every $a > 0$, from (3.3.15), $\sum_{N=1}^{\infty} \mathbb{P} \left(\left| \frac{\varepsilon_{N,H}}{N} - \ell_H \right| > a \right) < \infty$ hence $\frac{\varepsilon_{N,H}}{N} \xrightarrow[N \rightarrow +\infty]{a.s.} \ell_H$.

□

3.3.4 Proof of Lemma 3.3.3

Proof. Let $D_{N,H}(\mathbf{x}) := \sum_{n=1}^N \mathbf{x}_{t_n}^\top (RC_{H,t_n} - cRC_{H,t_n}) \mathbf{x}_{t_n}$, where \mathbf{x} satisfies \mathbf{H}_x . To alleviate the notations, we just write $D_{N,H}$ instead of $D_{N,H}(\mathbf{x})$. We want to control $\mathbb{E}[|D_{N,H}|^q]$ in N and H . The difficulty stems from the fact that we want to exhibit the dependence both in N and H . We are going to rely on a double arguments of martingale: $D_{N,H}$ can be seen as a nested martingale, in the long time period N , and in the short time period H .

Dependence in N : Denote $X_n := \mathbf{x}_{t_n}^\top (RC_{H,t_n} - cRC_{H,t_n}) \mathbf{x}_{t_n}$. Then X_n is $\mathcal{F}_{t_{n+1}}$ -measurable and $D_{N,H} := \sum_{n=1}^N X_n$. By cRC_{H,t_n} definition, $\mathbb{E}[RC_{H,t_n} | \mathcal{F}_{t_n}] = cRC_{H,t_n}$, hence since the \mathbf{x}_{t_n} are \mathcal{F}_{t_n} -measurable, $\mathbb{E}[X_n | \mathcal{F}_{t_n}] = 0$, so $D_{N,H}$ is a martingale w.r.t. the filtration $\{\mathcal{F}_{t_n}\}_{n=1}^{N+1}$.

Let us apply the Burkholder inequality (3.C.3) on $D_{N,H}$, martingale with increments X_n : denoting $C_{p_{\max}}^B$ the Burkholder constant (depending on p_{\max} only), we have:

$$\mathbb{E}[|D_{N,H}|^{p_{\max}}] \leq C_{p_{\max}}^B \mathbb{E} \left[\left| \sum_{n=1}^N X_n^2 \right|^{p_{\max}/2} \right].$$

By convexity inequality on $\left| \sum_{n=1}^N X_n^2 \right|^{p_{\max}/2}$,

$$\mathbb{E}[|D_{N,H}|^{p_{\max}}] \leq C_{p_{\max}}^B N^{\frac{p_{\max}}{2}-1} \sum_{n=1}^N \mathbb{E}[|X_n|^{p_{\max}}].$$

Dependence in H : We can decompose $X_n = \sum_{k=1}^H \mathbf{x}_{t_n}^\top (\mathbf{r}_{t_n+k} \mathbf{r}_{t_n+k}^\top - \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}]) \mathbf{x}_{t_n}$ in a martingale part $Y_n = \sum_{k=1}^H Y_k^{t_n}$ and a remaining part $Z_n = \sum_{k=1}^H Z_k^{t_n}$, where:

$$\begin{aligned} Y_k^{t_n} &:= \mathbf{x}_{t_n}^\top (\mathbf{r}_{t_n+k} \mathbf{r}_{t_n+k}^\top - V_{t_n+k}) \mathbf{x}_{t_n}, \\ Z_k^{t_n} &:= \mathbf{x}_{t_n}^\top (V_{t_n+k} - \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}]) \mathbf{x}_{t_n}, \quad 1 \leq k \leq H. \end{aligned}$$

Notice that $Z_1^{t_n} = \mathbf{x}_{t_n}^\top (V_{t_n+1} - \underbrace{\mathbb{E}[V_{t_n+1} | \mathcal{F}_{t_n}]}_{=V_{t_n+1}}) \mathbf{x}_{t_n} = 0$.

By convexity inequality

$$\mathbb{E}[|X_n|^{p_{\max}}] = \mathbb{E}[|Y_n + Z_n|^{p_{\max}}] \leq 2^{p_{\max}-1} (\mathbb{E}[|Y_n|^{p_{\max}}] + \mathbb{E}[|Z_n|^{p_{\max}}]).$$

Bound on $\mathbb{E}[|Y_n|^{p_{\max}}]$: By the the Burkholder inequality (3.C.3) on Y_n and by convexity:

$$\begin{aligned} \mathbb{E}[|Y_n|^{p_{\max}}] &\leq C_{p_{\max}}^B \mathbb{E} \left[\left| \sum_{k=1}^H (Y_k^{t_n})^2 \right|^{p_{\max}/2} \right] \\ &\leq C_{p_{\max}}^B H^{\frac{p_{\max}}{2}-1} \sum_{k=1}^H \mathbb{E}[|Y_k^{t_n}|^{p_{\max}}]. \end{aligned}$$

Bound on $\mathbb{E}[|Z_n|^{p_{\max}}]$: Since the \mathbf{x}_{t_n} components are bounded by $c_{\mathbf{w}}$ from assumption \mathbf{H}_x ,

$$\mathbb{E}[|Z_n|^{p_{\max}}] \leq (c_{\mathbf{w}}^2 d)^{p_{\max}} \mathbb{E} \left[\left| \sum_{k=2}^H V_{t_n+k} - \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}] \right|^{p_{\max}} \right].$$

Making appear $\mathbb{E}_{\mu}[V_t] = \mathbb{E}[V_{t_n+k} | V_{t_n} \sim \mu] = V_{\infty}$, by convexity inequality:

$$\mathbb{E}[|Z_n|^{p_{\max}}] \leq 2^{p_{\max}-1} (c_{\mathbf{w}}^2 d)^{p_{\max}} \left(\mathbb{E} \left[\left| \sum_{k=2}^H (V_{t_n+k} - V_{\infty}) \right|^{p_{\max}} \right] + \mathbb{E} \left[\left| \sum_{k=2}^H (V_{\infty} - \mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}]) \right|^{p_{\max}} \right] \right).$$

Notice that the second term is bounded by the first one, since the conditional expectation is non-expansive in $L_{p_{\max}}$ -norm,

$$\mathbb{E} \left[\left| \sum_{k=2}^H (\mathbb{E}[V_{t_n+k} | \mathcal{F}_{t_n}] - V_{\infty}) \right|^{p_{\max}} \right] = \mathbb{E} \left[\left| \sum_{k=2}^H (V_{t_n+k} - V_{\infty}) \right|^{p_{\max}} \right]. \quad (3.3.16)$$

Hence

$$\mathbb{E}[|Z_n|^{p_{\max}}] \leq 2^{p_{\max}} (c_{\mathbf{w}}^2 d)^{p_{\max}} \mathbb{E} \left[\left| \sum_{k=2}^H (V_{t_n+k} - V_{\infty}) \right|^{p_{\max}} \right].$$

As we have done in (3.3.14), we can apply Fort-Moulines Proposition extension (Proposition 3.3.5)) to get the dependence in H :

$$\mathbb{E} \left[\left| \sum_{k=2}^H (V_{t_n+k} - V_{\infty}) \right|^q \middle| V_{t_n+1} \right] \leq C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^q}{\mathcal{L}(x)} \mathcal{L}(V_{t_n+1}) (H-1)^{q/2}.$$

So,

$$\mathbb{E}[|Z_n|^{p_{\max}}] \leq (2c_{\mathbf{w}}^2 d)^{p_{\max}} C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^q}{\mathcal{L}(x)} \mathbb{E}[\mathcal{L}(V_{t_n+1}) | V_1] H^{q/2}.$$

Combining the obtained bounds on Z_n and Y_n , we obtain:

$$\begin{aligned} \mathbb{E}[|X_n|^{p_{\max}}] &\leq 2^{p_{\max}-1} \left(C_{p_{\max}}^B m_{Y,n} + (2c_{\mathbf{w}}^2 d)^{p_{\max}} C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^q}{\mathcal{L}(x)} m_{\mathcal{L},n} \right) H^{\frac{p_{\max}}{2}}, \\ \mathbb{E}[|D_{N,H}|^{p_{\max}}] &\leq C_{p_{\max},d,\mathcal{L}} (NH)^{\frac{p_{\max}}{2}}, \end{aligned}$$

where

$$\begin{aligned} m_{Y,n} &:= \frac{\sum_{k=1}^H \mathbb{E}[|Y_k^{t_n}|^{p_{\max}}]}{H}, \quad m_{\mathcal{L},n} := \frac{\sum_{k=1}^H \mathbb{E}[|Y_k^{t_n}|^{p_{\max}}]}{H} \quad \text{and} \\ C_{p_{\max},d,\mathcal{L}} &:= 2^{p_{\max}-1} C_{p_{\max}}^B \left(C_{p_{\max}}^B \frac{\sum_{n=1}^N m_{Y,n}}{N} + (2c_{\mathbf{w}}^2 d)^{p_{\max}} C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^q}{\mathcal{L}(x)} \frac{\sum_{n=1}^N m_{\mathcal{L},n}}{N} \right) \end{aligned}$$

are quantities which can be bounded independently from N and H .

□

3.3.5 Auxiliary result

In what follows, we prove the Lemma 3.2.3, e.g. that when $V_{\text{ref}} = V_\infty$, $\lim_{H \rightarrow \infty} \frac{\ell_H^\infty}{H} = 0$.

Proof. Let us first show that $\frac{cRC_{H,T_{\text{reb}}}}{H}$ converges to V_∞ . Let $T_{\text{reb}} \geq 2$.

As we did in Lemma 3.3.3 equation (3.3.16), by the non-expansivity of the conditional expectation in L_p -norm,

$$\begin{aligned} \mathbb{E} \left[\left| cRC_{H,T_{\text{reb}}} - HV_\infty \right|^{p_{\text{max}}} \right] &\leq \mathbb{E} \left[\left| \sum_{k=1}^H V_{T_{\text{reb}}+k} - HV_\infty \right|^{p_{\text{max}}} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{k=1}^H V_{T_{\text{reb}}+k} - HV_\infty \right|^{p_{\text{max}}} \middle| \mathcal{F}_{T_{\text{reb}}-1} \right] \right]. \end{aligned}$$

And as done in (3.3.14),

$$\mathbb{E} \left[\left| \sum_{k=1}^H V_{T_{\text{reb}}+k} - HV_\infty \right|^{p_{\text{max}}} \middle| V_{T_{\text{reb}}} \right] \leq C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^{p_{\text{max}}}}{\mathcal{L}(x)} \mathcal{L}(V_{T_{\text{reb}}}) H^{\frac{p_{\text{max}}}{2}}.$$

Taking the expectation, $\mathbb{E}[\mathcal{L}(V_{T_{\text{reb}}})|V_1]$ is finite by the Lyapunov condition drift $\mathbf{H}_{\mathcal{L}}$.

So dividing by $H^{p_{\text{max}}}$:

$$\mathbb{E} \left[\left| \frac{cRC_{H,T_{\text{reb}}}}{H} - V_\infty \right|^{p_{\text{max}}} \middle| V_{T_{\text{reb}}} \right] \leq C_{FM} \sup_{x \in \mathcal{S}} \frac{|x|^{p_{\text{max}}}}{\mathcal{L}(x)} \mathbb{E}[\mathcal{L}(V_{T_{\text{reb}}})|V_1] \frac{1}{H^{\frac{p_{\text{max}}}{2}}}. \quad (3.3.17)$$

When H goes to infinity, the bound in (3.3.17) goes to zero so $\frac{cRC_{H,T_{\text{reb}}}}{H} \xrightarrow[H \rightarrow +\infty]{L_{p_{\text{max}}}} V_\infty$.

By homogeneity and continuity of the mv operator which is Lipschitz from equation (3.3.13), and passage to the limit, since $\frac{cRC_{H,T_{\text{reb}}}}{H} \xrightarrow[H \rightarrow +\infty]{L_{p_{\text{max}}}} V_\infty$ with $p_{\text{max}} \geq 2$, we have convergence in L_1 norm and we can infer:

$$\underbrace{\mathbb{E}_{V_1 \sim \mu} \left[\text{mv}(V_\infty)^\top \frac{cRC_{H,T_{\text{reb}}}}{H} \text{mv}(V_\infty) - \text{mv} \left(\frac{cRC_{H,T_{\text{reb}}}}{H} \right)^\top \frac{cRC_{H,T_{\text{reb}}}}{H} \text{mv} \left(\frac{cRC_{H,T_{\text{reb}}}}{H} \right) \right]}_{= \frac{\ell_H^\infty}{H}} \xrightarrow[H \rightarrow +\infty]{} 0.$$

□

3.4 Model specification

We want to specify models satisfying the assumptions (3.2.1), \mathbf{H}_{stat} , $\mathbf{H}_{\mathcal{S}}$, $\mathbf{H}_{\mathcal{L}}$ and $\mathbf{H}_{p_{\text{max}}}$.

3.4.1 Motivation: GARCH-CCC model

The GARCH-CCC model [Bol90] is one of the simplest extension of GARCH model to the multidimensional case. For an exhaustive review of multidimensional GARCH, see [BLR06]. In this chapter, we elaborate our results with the simple GARCH-CCC, presented in [Car01].

This model assumes the following structure for the centered returns \mathbf{r}_t :

$$\begin{aligned} \mathbf{r}_t &= D_t \Gamma^{1/2} \eta_t = D_t \tilde{\eta}_t, \\ \tilde{\eta}_t &= \Gamma^{1/2} \eta_t, \\ V_t &= D_t \Gamma D_t, \end{aligned} \quad (3.4.1)$$

where

- $D_t = \text{Diag}(\sigma_{1,t}, \dots, \sigma_{d,t})$,
- $\{\sigma_{i,t}\}_{1 \leq i \leq d}$ are one-dimensional GARCH volatilities,
- $\Gamma = \{\rho_{ij}\}_{1 \leq i, j \leq d}$ is a positive definite matrix (the Constant Conditional Correlation matrix),
- η_t is a d -dimensional vector with independent components, $\mathbb{E}[\eta_{i,t}\eta_{j,t}] = 0 \forall i \neq j$, $\mathbb{E}[\eta_{i,t}^2] = 1$, and η_t independent from D_t .

The original GARCH-CCC assumes a simple GARCH(1,1) volatility recursion for the $\sigma_{i,t}$:

$$\sigma_{i,t}^2 = w_i + \alpha_i r_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2, \quad i = 1, \dots, d, \quad (3.4.2)$$

where $w_i, \alpha_i, \beta_i \in \mathbb{R}$.

The advantages of the "initial" GARCH-CCC model is its parsimony: it requires $\frac{d(d+5)}{2}$ parameters versus $\frac{d(5d+1)}{2}$ for the interdependent GARCH-CCC with $p = q = 1$ (3.4.4). The disadvantage is the strong assumption of constant conditional correlation.

In what follows, we will assume that $\{\eta_t\}_{t \in \mathbb{N}}$ is a sequence of independent Gaussian vectors.

3.4.2 Stationarity, ergodicity and application of results

In this subsection, we show that the GARCH-CCC satisfies the general model equation (3.2.1) and the model assumptions \mathbf{H}_{stat} , $\mathbf{H}_{\mathcal{S}}$, $\mathbf{H}_{\mathcal{L}}$ and $\mathbf{H}_{p_{\max}}$, under the following assumptions and definition on the parameters:

H_{param} . The GARCH-CCC parameters $\{w_i, \alpha_i, \beta_i\}_{i=1}^d$ and Γ are deterministic and satisfy:

- $\Gamma \in \mathcal{S}_+$ is a correlation matrix.
- $w_i > 0$, $\alpha_i > 0$ and $\beta_i > 0$ for all $i \in \{1, \dots, d\}$.
- $3\alpha_i^2 + \beta_i^2 + 2\alpha_i\beta_i < 1$ for all $i \in \{1, \dots, d\}$,

$q_{\max} \in \mathbb{R}_+$ is defined as:

$$q_{\max} := \min_{1 \leq i \leq d} \arg \max_{p \in \mathbb{R}_+} \{ \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p] < 1 \}, \quad \tilde{\eta} \sim \mathcal{N}(0, \Gamma). \quad (3.4.3)$$

Remark 5. A few remarks on $\mathbf{H}_{\text{param}}$:

- The condition $\mathbf{H}_{\text{param}}$ (iii) is equivalent to the condition $\mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^2] < 1$ (using that $\mathbb{E} [\tilde{\eta}_i^4] = 3$ and $\mathbb{E} [\tilde{\eta}_i^2] = 1$). Hence, it implies that $q_{\max} > 2$. This condition ensures in particular $\alpha_i + \beta_i < 1$, hence the existence of a stationary and ergodic solution to the model, as shown in the section 3.A.2.
- Continuity and evolution in p of $\mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p]$:
 - By L_p norms growth in p , if $q > 0$ is s.t. $|\alpha_i \tilde{\eta}_i^2 + \beta_i|_q^q = \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^q] < 1$, then for $0 \leq p \leq q$,

$$|\alpha_i \tilde{\eta}_i^2 + \beta_i|_p \leq |\alpha_i \tilde{\eta}_i^2 + \beta_i|_q < 1.$$
 - The function $g : p \mapsto \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p]$ is continuous (for example, by dominated convergence, for any sequence p_n converging to p , $g(p_n)$ converges to $g(p)$).
 - The function $g(\cdot)$ goes to infinity when p goes to infinity (because $\alpha_i > 0$), hence we are sure to have p s.t. $g(p) \geq 1$.

- Since $g(\cdot)$ is continuous, the quantity $\max_{p \in \mathbb{R}_+} \{ \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p] : \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p] < 1 \}$ is equal to 1. To prove the Lyapunov drift condition, we need a p s.t. $\mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p] < 1$. We denote p_{\max} such a p . From $\mathbf{H}_{\text{param}}$ (iii) and the first remark, we can take $p_{\max} \geq 2$.

Theorem 3.4.1. *Under $\mathbf{H}_{\text{param}}$, the GARCH-CCC variances process $\{\sigma_t^2\}_{t \in \mathbb{N}}$ satisfies the assumption of Theorem 3.2.2. \mathbf{H}_{stat} , $\mathbf{H}_{\mathcal{S}}$, $\mathbf{H}_{\mathcal{L}}$ and $\mathbf{H}_{p_{\max}}$.*

In what follows, we are going to prove the properties \mathbf{H}_{stat} , $\mathbf{H}_{\mathcal{S}}$, $\mathbf{H}_{\mathcal{L}}$ and $\mathbf{H}_{p_{\max}}$ on the GARCH-variances process σ_t^2 , on the stable state-space $\mathcal{S} = \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, +\infty \right)$.

Proof. Let $\{\sigma_t^2\}_{t \in \mathbb{N}}$ be a GARCH-CCC variances process.

\mathbf{H}_{stat} : Stationarity and ergodicity: From Theorem 3.A.2 explicit condition (3.A.2), a sufficient condition for stationarity and ergodicity for σ_t^2 is implied by our assumption $\mathbf{H}_{\text{param}}$ (iii).

$\mathbf{H}_{\mathcal{S}}$ on σ_t^2 : By definition, $\{\sigma_t^2\}_{t \in \mathbb{N}^*}$ is Markovian and time homogeneous. We show in our Corollary 3.A.4 that $\{\sigma_t^2\}_{t \in \mathbb{N}^*}$ is Lebesgue-irreducible on $\mathcal{S} = \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, +\infty \right)$. The aperiodicity will follow from the existence of a small set of positive measure shown below.

$\mathbf{H}_{\mathcal{L}}$ and $\mathbf{H}_{p_{\max}}$ on σ_t^2 : For $\mathbf{x} \in \mathcal{S}$, define

$$\mathcal{L}(\mathbf{x}) = 1 + \sum_{i=1}^d |x_i - w_i|^{p_{\max}},$$

and $\rho = \max_{1 \leq i \leq d} \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^{p_{\max}}]$ and $\delta = \frac{1+\rho}{2}$.

Then, if $\sigma_{t-1}^2 = \mathbf{x}$, $\mathbb{E}_{\mathbf{x}} \mathcal{L}(\sigma_t^2) = 1 + \mathbb{E} \sum_{i=1}^d |(\alpha_i \tilde{\eta}_i^2 + \beta_i) x_i|^{p_{\max}} \leq 1 + \rho (\sum_{i=1}^d |x_i|^{p_{\max}})$.

Since $\lim_{|\mathbf{x}| \rightarrow \infty} \frac{\mathbb{E}_{\mathbf{x}} \mathcal{L}(\sigma_t^2)}{\mathcal{L}(\mathbf{x})} \leq \rho$, the drift condition $\mathbb{E}_{\mathbf{x}} \mathcal{L}(\sigma_t^2) \leq \delta \mathcal{L}(\mathbf{x})$ is verified for \mathbf{x} large enough.

Let $R > \max_{1 \leq i \leq d} \frac{w_i}{1-\beta_i}$ s.t. it is verified for $|\mathbf{x}| > R$, $\mathcal{C} := \{\mathbf{x} \in \mathcal{S} \mid |\mathbf{x}| \leq R\}$ and $b := \sup_{\mathbf{x} \in \mathcal{C}} |\mathbb{E}_{\mathbf{x}} \mathcal{L}(\sigma_t^2) - \delta \mathcal{L}(\mathbf{x})|$, then

$$\mathbb{E}_{\mathbf{x}} \mathcal{L}(\sigma_t^2) \leq \delta \mathcal{L}(\mathbf{x}) + b \mathbf{1}_{\mathbf{x} \in \mathcal{C}}, \quad \mathbf{x} \in \mathcal{S}.$$

By definition, $\sigma_t^2 | \sigma_{t-1}^2 = \mathbf{x}$ follows the same law as $\{w_i + (\alpha_i \tilde{\eta}_i^2 + \beta_i) x_i\}$, $\tilde{\eta} \sim N(0, \Gamma)$. It admits a density $g_{\mathbf{x}}$ derived explicitly in Subsection 3.A.4 equation (3.A.4).

Properties of $g_{\mathbf{x}}$:

- $\mathbf{v} \mapsto g_{\mathbf{x}}(\mathbf{v})$ is continuous in \mathbf{v} for $v_i > w_i + \beta_i x_i$,
- For $(\mathbf{x}, \mathbf{v}) \in \underbrace{\left(B(0, R) \cup \prod_{i=1}^d \left[\frac{w_i}{1-\beta_i}, +\infty \right) \right)}_{\mathcal{C}} \times \left(\prod_{i=1}^d \left[\frac{w_i + \beta_i R}{2} + \frac{R}{2}, +\infty \right) \right) =: D$, we have

$$\begin{aligned} \frac{v_i - w_i - \beta_i x_i}{\alpha_i x_i} &\geq \frac{\frac{w_i + \beta_i R}{2} + \frac{R}{2} - w_i - \beta_i R}{\alpha_i R} \\ &= \frac{1}{2\alpha_i} \underbrace{\left(1 - \beta_i - \frac{w_i}{R} \right)}_{>0 \text{ by } R \text{ choice}} > 0. \end{aligned}$$

This function restricted to D is hence continuous and positive.

In particular, $\inf_{\mathbf{x} \in \mathcal{C}} g_{\mathbf{x}}(\mathbf{v}) > 0$ for each $\mathbf{v} \in \prod_{i=1}^d \left[\frac{w_i + \beta_i R}{2} + \frac{R}{2}, +\infty \right)$.

Let

$$\nu(A) := \int_A \inf_{\mathbf{x} \in \mathcal{C}} g_{\mathbf{x}}(\mathbf{v}) \mathbb{1}_{\left\{v_i > \frac{w_i + \beta_i R}{2} + \frac{R}{2}\right\}_{i=1}^d} d\mathbf{v}, \quad A \in \mathcal{B}(\mathcal{S}).$$

For any $\mathbf{x} \in \mathcal{C}$, $P(\mathbf{x}, A) = \int_A g_{\mathbf{x}}(\mathbf{v}) d\mathbf{v} \geq \nu(A)$ with ν non-null measure with density, hence \mathcal{C} is a petite set.

Since $\frac{w_i + \beta_i R}{2} + \frac{R}{2} < R$ (because $R > \frac{w_i}{1 - \beta_i}$), $\prod_{i=1}^d \left[\frac{w_i + \beta_i R}{2} + \frac{R}{2}, R \right]$ is non empty and has a positive Lebesgue measure and is included in \mathcal{C} . Hence:

$$\nu(\mathcal{C}) \geq \int_{\left\{ \frac{w_i + \beta_i R}{2} + \frac{R}{2} < v_i < R \right\}_{i=1}^d} \inf_{\mathbf{x} \in \mathcal{C}} g_{\mathbf{x}}(\mathbf{v}) d\mathbf{v} > 0$$

so \mathcal{C} is accessible and as a consequence of [MT09, Aperiodicity definition - page 114], the chain is aperiodic. Hence the Lyapunov condition is satisfied, with \mathcal{L} polynomial in p_{\max} .

H_{stat}: Existence of L_2 -moment: From the Lyapunov drift condition, we are ensured to have finite p_{\max} moments. Since we assumed $p_{\max} \geq 2$, it implies the square integrability of the process $\{\sigma_t^2\}_{t \in \mathbb{N}^*}$ for any initial condition and under the stationary law. □

Bijection between $\{V_t\}_{t \in \mathbb{N}}$ and σ_t^2 : Since $\sigma_{i,t}^2 > 0$ on \mathcal{S} , there exists a function Φ_{Γ} s.t.

$$(V_t)_{i,j} = \begin{cases} \rho_{ij} \sigma_{i,t} \sigma_{j,t} & \text{if } i \neq j, \\ \sigma_{i,t}^2 & \text{else} \end{cases} =: \Phi_{\Gamma}(\sigma_t^2).$$

- Φ_{Γ} is bijective,
- Φ_{Γ} is sub-linear: $|\Phi_{\Gamma}(\sigma_t^2)| \leq |\Gamma| |\sigma_t^2|$.

By Theorem 3.A.1, if σ_t^2 admits an ergodic and stationary solution, so is $\Phi_{\Gamma}(\sigma_t^2) = V_t$.

Results application:

- By definition of the GARCH-CCC model (3.4.1), $\{V_t = \Phi_{\Gamma}(\sigma_t^2)\}_{t \in \mathbb{N}^*}$ satisfies (3.2.1).
- If $\{\sigma_t^2\}_{t \in \mathbb{N}^*}$ satisfies **H_{stat}**, **H_S**, **H_L** and **H_{p_{max}}**, then the ergodic concentration results (Propositions 3.3.5 and 3.3.4) apply to any function $g(\Phi_{\Gamma}(\sigma_t^2))$, $g \circ \Phi_{\Gamma}$ bounded in $L^{1/q}$ norm, $q \geq 2$.

Since Φ_{Γ} is sub-linear, our Proposition 3.3.1, Lemma 3.3.3 and main result Theorem 3.2.2 apply without any adaptation.

Hence, under parameters condition **H_{param}**, by Theorem 3.4.1 and these remarks, the main result Theorem 3.2.2 is satisfied with p_{\max} defined in equation (3.4.3).

3.4.3 Additional models

In this subsection, we list additional GARCH models on which it is possible to extend our study. We first list one-dimensional GARCH models which can be extended to multidimensional model with a constant correlation matrix. Then we refer to a multidimensional model with inter-dependence between the GARCH volatilities.

3.4.3.1 One-dimensional GARCH models

The Threshold GARCH (T-GARCH) and the asymmetric power-GARCH are affine models, hence we can directly apply our results on them:

- T-GARCH:

$$\sigma_t = w + [\alpha_+(\eta_{t-1})_+ - \alpha_-(\eta_{t-1})_- + \beta] \sigma_{t-1},$$

where $w, \alpha_+, \alpha_-, \beta > 0$.

- Power GARCH:

$$\sigma_t^\delta = w + [\alpha(|\eta_{t-1}| - \zeta \eta_{t-1})^\delta + \beta] \sigma_{t-1}^\delta,$$

where $w, \alpha, \beta, \delta > 0, |\zeta| \leq 1$.

Indeed, we can adapt the assumptions verification taking the same Lyapunov function defined on σ_t instead of σ_t^2 for the T-GARCH, and σ_t^δ for the power GARCH.

The following models are of the form $\sigma_t \leq a(\eta_{t-1})\sigma_{t-1}^\gamma$ with $\gamma < 1$ if $\alpha + \beta < 1$. We can adapt our results on these models by taking a linear Lyapunov function.

- Exponential GARCH:

$$\begin{aligned} \ln \sigma_t^2 &= w + \alpha g(\eta_{t-1}) + \beta \ln \sigma_{t-1}, \\ g(\eta_{t-1}) &= \theta \eta_{t-1} + \zeta (|\eta_{t-1}| - \mathbb{E}[|\eta_{t-1}|]), \\ \sigma_t^2 &= e^w e^{\alpha g(\eta_{t-1})} (\sigma_{t-1}^2)^\beta, \end{aligned}$$

with $\alpha, \beta > 0, \beta < 1$ and $-\zeta < \theta < \zeta$ for g to be increasing in $|\eta_{t-1}|$ (and $\theta < 0$ for negative innovation to have more impact than positive ones).

In this model, in the Gaussian innovation case, moments exist at any order [FZ19, p. 79].

- Log-GARCH: the recursion holds on the logarithm of variances:

$$\begin{aligned} \ln \sigma_t^2 &= w + \alpha \ln r_{t-1} + \beta \ln \sigma_{t-1} \\ &= w + \alpha \ln \eta_{t-1}^2 + (\alpha + \beta) \ln \sigma_{t-1}, \\ \sigma_t^2 &= e^w (\eta_{t-1})^\alpha \sigma_{t-1}^{\alpha+\beta}, \end{aligned}$$

where $w, \alpha, \beta \in \mathbb{R}$.

3.4.3.2 Multidimensional GARCH models

In Francq and Zakoian's book [FZ19, page 280], the authors define the interdependent GARCH volatilities model:

$$\sigma_t^2 = \mathbf{w} + A \mathbf{r}_{t-1}^2 + B \sigma_{t-1}^2, \quad (3.4.4)$$

where $A, B \in \mathbb{R}^{d \times d}, \mathbf{w} \in \mathbb{R}^d$.

It is possible to adapt our results to this setting.

3.5 Numerical experiments

In this section, we confront our theoretical results to the real probabilities of better performance of the realized covariance based portfolio. To do so, we base ourselves on GARCH fitted parameters on real financial data to have realistic range of parameters (Subsection 3.5.1). Then we simulate GARCH-CCC based returns and compute the corresponding cRC_{H,t_n} and mv matrices and portfolio, as described in Subsection 3.5.2. Finally, we exhibit the evolution of $\hat{\mathbb{P}}(R_{N,H} < R_{N,H}^\infty)$ for multiple values of H and N .

3.5.1 Realistic GARCH values, fitting procedure

We have fitted GARCH-CCC parameters on real financial time series, using a two steps procedure:

1. fit of the one-dimensional GARCH(1,1) models by maximum likelihood (where the explicit likelihood for one-dimensional GARCH can be found in [EB86] for example) on the recentered returns (we assumed a fixed drift),
2. fit of the constant conditional correlation on the reconstructed residuals $\tilde{\eta}_{i,t}$.

The financial time series consist in the 12 components of BNPP QIS Momentum Strategy, composed of 4 indices consisting in index futures rolling (indices being EuroStoxx50, S&P 500, Nikkei, HSCEI), 3 indices of rolling of futures of bonds (German Bund, US government bond, Japanese government bond) and 5 indices of rolling of commodity (on gold, Brent, S&P GSCI excess return, Goldman Sachs US industrial metal ER) fx hedged in euros, from 12/03/1993 until 02/02/2017 (6046 dates).

	Minimum	Maximum	Average	Standard-deviation
σ_∞^2	0.060	5.15	1.77	1.44
$\sigma_{\infty,an}$	4.04	36.02	18.91	9.47
w	0.00032	0.04320	0.0135	0.0137
α	0.0329	0.0955	0.0622	0.0242
β	0.885	0.966	0.931	0.0271
$\alpha + \beta$	0.981	0.998	0.993	0.00182
ρ	-22%	82%	10%	22%
p_{\max}	1.83	5.72	3.34	1.22
$t_{1/2}$	36.2	437.2	135.8	98.52

Table 3.1 – Summary of variance (σ_∞^2), GARCH(1,1) parameters (α, β, w), constant conditional correlation (ρ), highest moment (p_{\max}) and half-life time ($t_{1/2}$, as defined in (3.5.2)) minimum, maximum, average and standard-deviation of the calibrated parameters. The variance and w parameters are expressed in basis points of daily variances, $\sigma_{\infty,an}$ denotes the annualized volatility and is expressed in percents, $t_{1/2}$ is expressed in days. α, β and p_{\max} have no dimension. What we call standard-deviation is the standard-deviation over the estimated quantities.

We give the minimum, maximum, average and standard-deviation of the obtained values in Table 3.1. In the same Table, we also give the minimum, maximum, average and standard-deviation of additional quantities computed on the parameters:

- the maximum moment $p_{\max,i}$, computed on the parameters with formula (3.4.3)

$$p_{\max,i} := \arg \max_{p \in \mathbb{R}_+} \{ \mathbb{E} [|\alpha_i \tilde{\eta}_i^2 + \beta_i|^p] < 1 \}, \quad (3.5.1)$$

- the half-life $t_{1/2}$, defined as the average time for the gap between the variance level to its long-term level to be reduced by one half:

$$t_{1/2,i} := -\frac{\ln 2}{\ln(\alpha_i + \beta_i)}, \quad (3.5.2)$$

- $\sigma_{i,\infty}^2$ the empirical variance measured on the data.

We display in Figure 3.2 the obtained values for α , β and p_{\max} .

Comments and interpretation

- The average α is relatively small (0.062) and the average β is relatively high (0.931), and the sum $\alpha + \beta$ is very close to 1. According to Campbell [CLM97, page 483], the α parameter measures how much a previous shock (high value of an innovation) will propagate to the future.
- The half-life times are large (because the $\alpha + \beta$ are close to 1): this phenomena corresponds to the persistence of shocks over time. The average half-life time is 135 days hence approximately six months. It means that before this horizon, the instantaneous variance can be far from its expectation under the stationary distribution.
- The w_i values are very close to zero since by stationarity condition:

$$\sigma_{i,\infty}^2 = \mathbb{E}[\epsilon_{i,t}^2] = \mathbb{E}[\sigma_{i,t}^2] = \frac{w_i}{1 - \alpha_i - \beta_i},$$

they should be of the same order of magnitude than $\sigma_{i,\infty}^2(1 - \alpha_i - \beta_i)$, and $\alpha_i + \beta_i \approx 1$.

- We notice that low values of α are associated to high values of β (as shown in Figure 3.2) and conversely, in such a way that the stationary condition $\alpha + \beta < 1$ is always enforced.
- The average p_{\max} is equal to 3.34, and the minimum is slightly smaller than 2. As shown in Figure 3.2, most p_{\max} values are higher than 2 which means that the time series have finite kurtosis.

The estimated parameters show that our approach of considering portfolio with horizon H not too large (for example: 21 days) and $p_{\max} \geq 2$ but of small order is relevant since the average half-time is very large (more than a month) and the existence of order-2 moments is almost always verified.

3.5.2 Simulation procedure

In what follows, we describe how we choose the GARCH-CCC parameters and how we estimate the benchmark covariance V_∞ and the realized conditional covariance $cRC_{H,T_{reb}}$.

- **$\{\alpha, \beta, \Gamma\}$ choice**

The $\{\alpha_i, \beta_i, \Gamma_{i,j}\}_{i,j=1}^d$ are simulated uniformly in the ranges indicated in Table 3.1. The (α_i, β_i) are sorted such that the smallest α_i are associated to the largest β_i . (We enforce Γ to be definite positive by ensuring or capping its eigen-values to 10^{-2} and transforming the matrix to retrieve a 1-diagonal.)

- **V_∞ estimation**

We estimate V_∞ as the empirical covariance of the returns over a long range of time, typically three times the largest half-life time. It amounts to consider a backtest of three times the largest half-life time. Given the half-life times observed range, the largest life-time is around 400 business days, hence 1.6 years: 3 times this period corresponds to a backtest size of 4.8 years.

- **$cRC_{H,T_{reb}}$ estimation**

We evaluate the $cRC_{H,T_{reb}}$ via Monte Carlo:

$$\widehat{cRC_{H,T_{reb}}} = \sum_{k=1}^H \widehat{\mathbb{E}}[V_{T_{reb}+k} | \mathcal{F}_{T_{reb}}] = \frac{1}{N_{mc}} \sum_{n=1}^{N_{mc}} \sum_{k=1}^H V_{T_{reb}+k}^{(n)} | \mathcal{F}_{T_{reb}}$$

with




Figure 3.2 – $(\alpha, \beta, \rho_{\max})$ parameters fitted on 12 financial time series. In orange, we draw the finite variance area (condition H_{param} (iii)), in yellow, we draw the stationary but not finite variance area. A dot corresponds to one (α_i, β_i) fitted parameter. Its color indicates its associated ρ_{\max} value (as defined in $H_{\text{param}}(\rho_{\max})$). The color scale to the right corresponds to ranges of observed values for ρ_{\max} : lower than 2 values are in dark blue, $\rho_{\max} \in [2, 3]$ is in cyan, higher values are in green, orange and brown. Most of our fitted parameters are colored in cyan, i.e. are associated to $\rho_{\max} \in [2, 3]$ values.

- $V_{T_{reb}}^{(n)} = V_{T_{reb}}$,
- $V_{T_{reb}+k}^{(n)}$ obtained by application of the recursion formula on $V_{T_{reb}+k}^{(n)}$ and a generated $r_{T_{reb}+k}^{(n)} \sim \mathcal{N}(0, V_{T_{reb}+k-1}^{(n)})$, for $k = 1, \dots, H$.

We took $N_{mc} = 100$.

We consider the minimum variance portfolio without constraints: $\mathcal{W} = \{\mathbf{w} = \{w_i\}_{i=1}^d \in \mathbb{R}^d : \mathbf{1}_d^\top \mathbf{w} = 1\}$, where $\mathbf{1}_d$ is the d -dimensional vector of ones. Then $\text{mv}(\cdot)$ is explicit: for $C \in \mathbb{R}^{d \times d}$ positive definite matrix, $\text{mv}(C) = \frac{C^{-1} \mathbf{1}_d}{\mathbf{1}_d^\top C^{-1} \mathbf{1}_d}$. Since we deal with a non-degenerate GARCH-CCC model, we know that our allocation weights will be bounded during the length of the experiment. This is why we did not enforce a bound on their norm.

We will place in the following settings:

- $d \in \{10, 50\}$,
- $H = 1$ (daily rebalancing), $H = 5$ (weekly), $H = 21$ (monthly),
- $N = 1, \dots, 6 \times 21/H$.

For a daily rebalancing period, we will consider up to $N = 126$ rebalancing times of the portfolio, and for monthly time period, up to $N = 6$ rebalancing times.

For one set of d -dimensional GARCH parameters, we reproduce the following experiment multiple times ($N_{MC} = 10^4$ times):

- We let the multidimensional GARCH evolve during 3 times the maximal half-life time associated to the d sets of GARCH parameters. We do so in order to reach plausible GARCH values beyond the burn-in phase of the process. We precise that our results are valid whatever the starting point and that it is not necessary to reach the stationary regime for the results to apply.
- **Computation of the probabilities: trajectory approach (less independent) versus by (H, N) approach (more independent).**

To do our experiments, we first considered a trajectory approach, consisting in the following steps:

We start the allocation procedure by initializing our daily, weekly and monthly portfolios: $\text{mv}(V_1)$, $\text{mv}(\widehat{cRC}_{5,1})$ and $\text{mv}(\widehat{cRC}_{21,1})$. Then for $t = 1, \dots, 126$,

- We update the daily realized variance $R_{t,1}$ and we update the portfolio $\text{mv}(V_t)$, and the realized variance.
- If t is zero modulo 5, we update the weekly realized variance $R_{t,5}$ and portfolio with $\widehat{cRC}_{5,t}$.
- If t is zero modulo 21, we update the monthly realized variance $R_{t,21}$ and portfolio with $\widehat{cRC}_{21,t}$.

Another more independent (but more time consuming) approach consists in generating a new GARCH trajectory for each (H, N) pair considered. This is the approach we have finally taken in the following illustrations (except for the $d = 50$ experiment).

Hence, we can compute iteratively the empirical probabilities:

$$\widehat{\mathbb{P}}(R_{N,H} < R_{N,H}^\infty) = \frac{1}{N_{MC}} \sum_{N=1}^{N_{MC}} \mathbf{1}_{R_{N,H} < R_{N,H}^\infty}, \quad H \in \{1, 5, 21\}, \quad N = 1, \dots, \frac{126}{H}.$$

3.5.3 Empirical probability: impact of p_{\max} and d

We display in Figures 3.3 and 3.4 the evolution of the empirical $\mathbb{P}\left(R_{N,H} > R_{N,H}^{\infty}\right)$ with H and N , for $d = 10$ and small versus large p_{\max} (Figure 3.3), and in the other figures, we compare large dimension ($d = 50$) versus small dimension ($d = 10$) (Figure 3.4). The time is in business days. We compare daily (red dots), weekly (blue) and monthly (green) rebalancing times of the portfolios. On all the experiments, the probabilities decreased towards 0.

3.5.3.1 Impact of p_{\max}




Figure 3.3 – Empirical probabilities $\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^{\infty})$, for $H = 1$ (red), $H = 5$ (blue) and $H = 21$ (green) as a function of the number of rebalancing times N , in log-log scale, for a set of GARCH parameters such that $d = 10$, $p_{\max} = 2.8$ (left) and $p_{\max} = 7$ (right). In black dots, we plotted the tendency line fitted on $\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^{\infty})$ for $H = 5$ and $N \geq 5$.

In Figure 3.3 we compare at fixed dimension $d = 10$ the impact of the integrability. For the figure on the left, it requires more than 7 rebalancing of the monthly or weekly portfolio to reach a 1% probability level that the GARCH-covariance will outperform the benchmark, whereas in the large p_{\max} case, the 1% level is reached at the 5th rebalancing of the portfolio. Practically speaking, it means that for a period of 5 weeks (weakly portfolio) or 5 months (monthly portfolio), we have a 99% probability that the GARCH-based portfolio has a lowest realized variance than the benchmark. Since a higher p_{\max} implies lighter tails, the deviation are less important and it is easier to see the performance gain. For $N \geq 4$, the tendency lines have a slope of -3.4 (small p_{\max} case) versus -3.2 (large p_{\max} case).

Link with Theorem 3.2.2: We observe as expected a decreasing shape of the probability $\mathbb{P}\left(R_{N,H} > R_{N,H}^{\infty}\right)$, but with a concave evolution and not the linear, with a $-\frac{p_{\max}}{2}$ slope, expected shape. Our results are not conclusive for N small (the rate of decline is lower than expected). For N large, we even observe a faster convergence than expected in the small p_{\max} case (slope of -3.4 instead of $-\frac{p_{\max}}{2} = -1.4$). This is an indication that there can be several decay regimes depending on N value. This is hard to give a clear, quantitative explanation of this phenomena. It can be argued that our model consists in a vectorized version of multiple unidimensional GARCH processes, each of them being associated to different tail thickness. Since we deal at the portfolio variance level, it is possible than GARCH with lighter tails are associated to higher components in the minimum variance allocation, which could explain this faster than expected convergence rate. Nonetheless, these results are not in contradiction with our theoretical results. Indeed, our results can be interpreted as a "worst-case" bound on the probability, so it is not

surprising to do better. And the probability levels are lower for the large p_{\max} case, which is coherent with the bound behavior.

3.5.3.2 Impact of the dimension

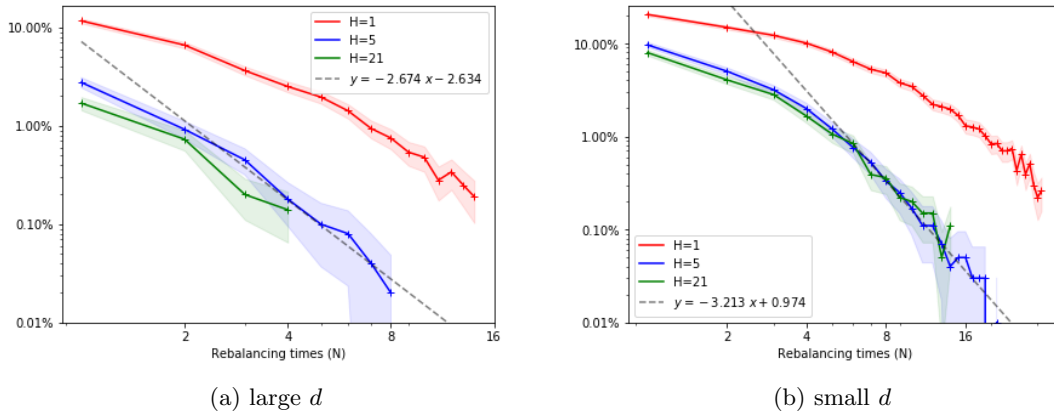


Figure 3.4 – Empirical probabilities $\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^{\infty})$, for $H = 1$ (red), $H = 5$ (blue) and $H = 21$ (green) as a function of the number of rebalancing times N , in log-log scale, for a set of GARCH parameters such that $d = 50$, $p_{\max} = 5.05$ (left) and $d = 10$, $p_{\max} = 7.03$ (right).

In Figure 3.4, we display the evolution of the empirical $\mathbb{P}(R_{N,H} > R_{N,H}^{\infty})$ for large p_{\max} and very different dimensions: $d = 50$ (Figure 3.4 (a)) and $d = 10$ (Figure 3.4 (b)).

Comparing 3.4 (a) and (b), the dimension does not seem to have a strong impact, or might be compensated by the stronger impact of the large integrability.

In the large dimension case, the tendency line has a slope of -2.6 which is lower than the theoretical minimum convergence rate of $-\frac{p_{\max}}{2} = -2.025$ (hence the convergence is faster than the theoretical bound).

3.5.3.3 Impact of the number of rebalancing times

In Figure 3.5, we display the histogram of the realized variance difference for a monthly portfolio with 10 assets and different N . We see that for $N = 1$, the realized variance difference is almost centered around 0 but for $N \geq 4$, all the distributions are significantly centered on positive values. Indeed, in our setting (initial condition following approximately the stationary law), the average realized variance difference is equal to the renormalized performance gap $\frac{\ell_H^{\infty}}{H}$ which is positive and deterministic as shown in Proposition 3.2.1.




Figure 3.5 – Histograms of $\left\{ \frac{R_{N,H}^\infty - R_{N,H}}{NH} \right\}_{N \in \{1,6,9,12\}}$ with $H = 21$ for a set of GARCH parameters such that $d = 10$, $p_{\max} = 2.8$. Here the variance differences are expressed as the difference in the realized variance over N rebalancing times of the portfolio, in basis points of daily variance. $N = 6$ corresponds to a variance difference over a six months, between the benchmark portfolio and a monthly re-updated portfolio. In black dots, we plot the average of $\frac{R_{N,H}^\infty - R_{N,H}}{NH}$ for $N = 6$ and $H = 21$.

The x scale is a daily variance measure in basis points. We see that in average, the monthly portfolio implies a drop of daily variance of 0.02 bps. This variance level has to be put in perspective: we consider a range of daily variance of $[0.06, 5.2]$ bps in our experiments.

In terms of associated annualized volatility, we can upper bound the volatility difference by the square-root of the variance difference, which translates to a drop of $\sqrt{0.02 \times 252/10^4} = 2.2\%$.

3.5.3.4 Impact of the investment period

As mentioned in the comments and interpretations paragraph following our main result, if we are in the $(NH) \ll C(H/\ell_H)^2$ regime, the asset manager should not bother much using a sophisticated estimation for V_t : a good approximation of V_∞ is enough. In this Subsubsection, we aim at illustrating this comment. In particular, we aim at answering the question: *starting from which investment period H do we have $R_{N,H} > R_{N,H}^\infty$ in half the cases?* For this H threshold, we are indifferent between the stationary and the conditional realized covariance matrices.

We are interested in the evolution of $\mathbb{P}(R_{N,H} < R_{N,H}^\infty)$ for increasing H values. For sake of simplicity, we consider only one rebalancing of the portfolio ($N = 1$). For larger N , the probability levels will be even lower according to the previous experiments and theoretical results, so $N = 1$ can be seen as a worst case-scenario.

In Figure 3.6, we display the evolution of $\mathbb{P}(R_{1,H} < R_{1,H}^\infty)$ with H in months and in log-log scale, for a universe of 10 assets associated to an integrability of order $p_{\max} = 2.8$ and a half-life

of 39 days hence approximately 2 months. As previously observed in Figure 3.3 (same set of parameters), for $H = 21$ days (so one month) and $N = 1$, the probability is 90%. Then the probability is decreasing with the number of months and it is approximately equal to $\frac{1}{2}$ when H is equal to one-year and a half (18 months). The investment period must therefore be as long as 9 times the half-life of the process for the stationary covariance to be as efficient as the GARCH-based covariance.




Figure 3.6 – Empirical probabilities $\widehat{\mathbb{P}}(R_{1,H} < R_{1,H}^\infty)$ as a function of the investment period H for a set of GARCH parameters such that $d = 10$, $p_{\max} = 2.8$, $t_{1/2} = 39$ days.

For investment period lower than a year, the GARCH covariance is hence relevant and efficient.

3.5.4 Tail exponent Hill estimation

From (3.3.3), there is a $C > 0$ independent from N s.t.

$$\bar{F}_{N,H}(y) = \mathbb{P}\left(\frac{D_{N,H}}{\sqrt{NH}} > y\right) \leq \frac{C}{y^{p_{\max}}}.$$

In what follows, we are going to compare our estimation of p_{\max} via numerical method on equation (3.5.1) and the one obtained via estimation of the tail exponent of $\frac{D_{N,H}}{\sqrt{NH}}$. To do so, we use the Hill estimator.

Hill estimator: The Hill estimator estimates the distribution tail exponent, e.g. the exponent γ such that its complementary cumulative distribution function \bar{F} verifies:

$$\bar{F}(x) = \frac{\ell(x)}{x^{1/\gamma}}$$

where $\lim_{x \rightarrow \infty} \frac{\ell(tx)}{\ell(x)} = 1$ for any $t > 0$.

For X_1, \dots, X_n random variables of cumulative distribution function F , denoting $X_{1,n} < X_{2,n} < \dots < X_{n,n}$ the ordered variables, the Hill estimator of order $k \in \mathbb{N}$ is defined by:

$$\hat{\gamma}(k) = \frac{1}{k} \sum_{i=0}^{k-1} (\ln X_{n-i,n} - \ln X_{n-k,n}).$$

It converges in probability towards γ if k increases with n with $k/n \xrightarrow[n \rightarrow \infty]{} 0$ (cf [DHF06, Theorem 3.2.4]) and converges asymptotically under a second order condition, as detailed in Appendix 3.B.

Reasonable choice of k : Assuming a Fréchet domain of attraction (this is the natural attraction domain for Pareto-like distributions), we can get an explicit formulation of the asymptotic variance and minimize it in k . It gives the following optimal k for $\gamma \neq 1$:

$$k_{opt}(\gamma, n) = 2 \left(\frac{\gamma}{1 - \gamma} n \right)^{\frac{2}{3}}. \quad (3.5.3)$$

The details on how we get this result is postponed to Appendix 3.B. If we expect a certain value of γ , we can plug it in (3.5.3) to estimate k_{opt} .




Figure 3.7 – Histogram of $\frac{D_{N,H}}{\sqrt{N}}$ for $H = 1$ and varying N (left) and Hill estimation of $\frac{1}{p_{\max}}$ on 10^4 simulations of $\frac{D_{N,H}}{\sqrt{N}}$ for $H = 1$ and $N = 40$, (right) with $p_{\max} = 2.8$ and $d = 10$ assets (both figures).

Experiment: We display the histogram of $\frac{D_{N,1}}{\sqrt{N}}$ for multiple N and the Gaussian density (black dots) and the Hill estimator $\hat{\gamma}(k_{opt}(1/p_{\max}, n))$ in Figure 3.7. We have computed the Hill estimator on an increasing number of samples, $n \in \{100, 200, \dots, 10\,000\}$ of the renormalized martingale $\frac{D_{N,1}}{\sqrt{N}}$, $N = 40$.

Comments:

- From Figure 3.7[left], we see that \sqrt{N} is a good renormalization for the $D_{N,1}$ martingale. Since we have finite variance, a central limit theorem seems to be verified.
- From Figure 3.7[right], the Hill estimator is overall close to $1/p_{\max}$. We plotted the confidence interval associated to the central limit theorem for the Hill estimator (3.B.1). $1/p_{\max}$ is inside the confidence interval.

3.6 Conclusion, perspectives

In this chapter, assuming a time-dependent conditional model on the returns of the form $r_{t|\mathcal{F}_{t-1}} \sim \mathcal{N}(0, V_t)$, we give the optimal covariance for a given period and time of investment for the minimum variance problem. Using a decomposition between a martingale and a positive ergodic

term, we can show that our covariance-based portfolio has a lower realized variance than any other benchmark covariance with a high probability, which is increasing with the number of rebalancing of the portfolio.

We give an explicit recursion scheme for the computation of this covariance matrix for the specific GARCH-CCC model. This recursion scheme could be adapted to many other models. We empirically illustrate our result by computing the empirical probability that the realized variance of our optimal covariance portfolio is smaller than the one with the stationary covariance matrix. The experiment results are not in contradiction with our theoretical analysis: we verify the convergence of the GARCH superperformance probability with the number of rebalancing, at a rate at least equal to $\frac{p_{\max}}{2}$ in the large N regime.

The question the practitioner may ask is *(when) is it relevant to use this more sophisticated covariance matrix, rather than a simple empirical covariance?* Under the assumption that the model is GARCH with known parameters, we have shown that the performance gap goes to zero when the investment period goes to infinity and the threshold of the investment period for both covariances to perform equally well seems to be around 10 times the half-life of the process. But we do not tackle in this work the fact that in practice, the models have to be fitted, leading to estimation error on the parameters, not to mention model error. We can advocate that estimation error can still be tackled by taking into account the Gaussian uncertainty on the parameters, stemming from the central limit behavior of the parameters when the backtest size is large enough.

APPENDICES

3.A GARCH properties

In this section we recall general and external ergodicity and Markov chain properties, which we specify to the GARCH-CCC model.

3.A.1 Ergodicity and stationarity

In this section we recall a general ergodicity property used to prove the GARCH covariance ergodicity 3.4.2.

Theorem 3.A.1 ([FZ19, Theorem A.1 page 367]). *If $(Z_t)_{t \in \mathbb{Z}}$ is an ergodic strictly stationary sequence and if $(Y_t)_{t \in \mathbb{Z}}$ is defined by*

$$Y_t = f(\dots, Z_{t-1}, Z_t, Z_{t+1}, \dots),$$

where f is a measurable function from \mathbb{R}^∞ to \mathbb{R} , then $(Y_t)_{t \in \mathbb{Z}}$ is also an ergodic strictly stationary sequence.

3.A.2 GARCH-CCC ergodicity

In this section, we give a simple sufficient condition for the GARCH-CCC model to be stationary and ergodic.

Before recalling Francq and Zakoian's result on the GARCH-CCC stationarity, let us introduce the definition of the *top Lyapunov exponent*.

The GARCH-CCC model (3.4.1) can be put under the following vector form.

Let $\mathbf{z}_t = \begin{pmatrix} \mathbf{r}_t^2 \\ \sigma_t^2 \end{pmatrix} \in \mathbb{R}^{2d}$ and let $E_t = \text{Diag}(\tilde{\eta}_t^2)$. Then \mathbf{z}_t satisfies the recursion:

$$\begin{aligned} \mathbf{z}_t &= \mathbf{b}_t + A_t \mathbf{z}_{t-1}, \\ \text{where } \mathbf{b}_t &= \begin{pmatrix} E_t \mathbf{w} \\ \mathbf{w} \end{pmatrix}, \quad A_t = \begin{pmatrix} E_t A & E_t B \\ A & B \end{pmatrix}, \end{aligned} \tag{3.A.1}$$

$\mathbf{w} = \{w_i\}_{i=1}^d$, $A = \text{Diag}(\{\alpha_i\}_{i=1}^d)$, $B = \text{Diag}(\{\beta_i\}_{i=1}^d)$.

The top Lyapunov exponent is defined by [FZ19, Theorem 2.3]:

$$\gamma := \lim_{t \rightarrow \infty} \text{a.s.} \frac{1}{t} \ln |A_t A_{t-1} \dots A_1|.$$

Theorem 3.A.2 (Strict stationarity of the CCC model [FZ19, Theorem 10.6]). *A necessary and sufficient condition for the existence of a strict stationary and non-anticipative solution process for the model (3.4.1) is $\gamma < 0$, where γ is the top Lyapunov exponent of the sequence $\{A_t, t \in \mathbb{Z}\}$ defined in (3.A.1). This stationary and non-anticipative solution, when $\gamma < 0$, is unique and ergodic.*

Explicit γ formulation and condition: Notice that γ does not depend on the chosen matricial norm, because norms are equivalent on the finite dimension space of matrices considered, and if $\frac{1}{K}|\cdot|_2 \leq |\cdot|_1 \leq K|\cdot|_2$, $K > 0$,

$$-\frac{\ln(K)}{t} + \frac{1}{t} \ln|A_t A_{t-1} \dots A_1|_2 \leq \frac{1}{t} \ln|A_t A_{t-1} \dots A_1|_1 \leq \frac{\ln(K)}{t} + \frac{1}{t} \ln|A_t A_{t-1} \dots A_1|_2$$

and by passage to the limit, the term $\frac{\ln(K)}{t}$ vanishes when t goes to infinity.

In what follows, we are going to specify an explicit sufficient condition s.t. $\gamma < 0$. W.l.o.g., we consider as matricial norm the infinite norm: $|A| = \max_{1 \leq i, j \leq d} |A_{i,j}|$, $A \in \mathbb{R}^{d \times d}$.

We can specify the top Lyapunov condition by noticing that A_t can be written as: $A_t = \begin{pmatrix} E_t \\ I_d \end{pmatrix} (A \ B)$, and for $1 \leq s \leq t$, $(A \ B) \begin{pmatrix} E_s \\ I_d \end{pmatrix} = \text{Diag} \left(\left\{ \alpha_i \tilde{\eta}_{s,i}^2 + \beta_i \right\}_{i=1}^d \right)$.

It follows, using the sub-additivity of the infinite norm, that:

$$\begin{aligned} \ln|A_t A_{t-1} \dots A_1| &= \ln \left| \begin{pmatrix} E_t \\ I_d \end{pmatrix} \prod_{s=1}^{t-1} \text{Diag} \left(\left\{ \alpha_i \tilde{\eta}_{s,i}^2 + \beta_i \right\}_{i=1}^d \right) (A \ B) \right| \\ &\leq \ln \left| \begin{pmatrix} E_t \\ I_d \end{pmatrix} \right| + \ln |(A \ B)| + \sum_{s=1}^{t-1} \ln \left| \text{Diag} \left(\left\{ \alpha_i \tilde{\eta}_{s,i}^2 + \beta_i \right\}_{i=1}^d \right) \right|, \end{aligned}$$

and by the strong law of large numbers, $\gamma \leq \mathbb{E} \ln \left| \text{Diag} \left(\left\{ \alpha_i \tilde{\eta}_i^2 + \beta_i \right\}_{i=1}^d \right) \right|$, $\tilde{\eta} \sim \mathcal{N}(0, \Gamma)$.

Taking as matricial norm the infinite norm, the condition

$$\mathbb{E} \ln \max_{1 \leq i \leq d} (\alpha_i \tilde{\eta}_i^2 + \beta_i) < 0$$

is a sufficient condition for $\gamma < 0$. It can be verified via numerical evaluation for example.

Easier to verify sufficient conditions: Since for fixed $\tilde{\eta}_i^2$: $\ln \max_{1 \leq i \leq d} (\alpha_i \tilde{\eta}_i^2 + \beta_i) \leq \sum_{i=1}^d \ln(\alpha_i \tilde{\eta}_i^2 + \beta_i)$, and taking the expectation:

$$\gamma \leq \mathbb{E} \ln \max_{1 \leq i \leq d} (\alpha_i \tilde{\eta}_i^2 + \beta_i) \leq \sum_{i=1}^d \mathbb{E} \ln(\alpha_i \tilde{\eta}_i^2 + \beta_i) \leq d \max_{1 \leq i \leq d} \mathbb{E} \ln(\alpha_i \tilde{\eta}_i^2 + \beta_i),$$

the following condition is a sufficient condition for $\gamma < 0$:

$$\max_{1 \leq i \leq d} \mathbb{E} \ln(\alpha_i \tilde{\eta}_i^2 + \beta_i) < 0.$$

The following condition ensures also the ergodicity condition, as a consequence of the Jensen inequality on $\mathbb{E} \ln(\alpha_i \tilde{\eta}_i^2 + \beta_i)$:

$$\alpha_i + \beta_i < 1, \quad i \in \{1, \dots, d\}. \quad (3.A.2)$$

3.A.3 Irreducibility and auxiliary results on the GARCH-CCC

In this section, we show the irreducibility of the GARCH-CCC process. A similar result is available in the Chapters 3 "Mixing properties of univariate GARCH" and 10 "Multivariate GARCH" of Francq and Zakoian's book [FZ19]. Compared to [FZ19, Chapters 3], we will give very explicit arguments which are the main adding of this study.

In this subsection, we assume H_{param} . The following lemma enables us to prove that $\mathcal{S} = \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, +\infty \right)$ is a stable state space, e.g. that from any point $\mathbf{x} \in \mathcal{S}$, we can reach any point as close from $\frac{w_i}{1-\beta_i}$ as wanted in a finite number of steps.

Lemma 3.A.3. *Let $\mathbf{x} \in \mathcal{S} = \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, +\infty \right)$. For all $c > 0$, there exists a finite $N_c \in \mathbb{N}$ such that:*

$$\mathbb{P} \left(\sigma_{i,N_c}^2 < \frac{w_i}{1-\beta_i} + c \mid \sigma_0^2 = \mathbf{x} \right) > 0, \quad \forall i = 1, \dots, d.$$

Proof. Applying recursively equation (3.4.2) on $\sigma_{i,N}^2$, for $\sigma_0^2 = \mathbf{x} \in \mathcal{S}$, $N \in \mathbb{N}^*$, with $\tilde{\eta}_n \sim \mathcal{N}(0, \Gamma)$, $n = 1, \dots, N-1$,

$$\begin{aligned} \sigma_{i,N}^2 &= w_i + (\alpha_i \tilde{\eta}_{i,N-1}^2 + \beta_i) \left(w_i + (\alpha_i \tilde{\eta}_{i,N-2}^2 + \beta_i) (w_i + \dots + (\alpha_i \tilde{\eta}_{i,0}^2 + \beta_i) \sigma_{i,0}^2) \right) \\ &= w_i \left(1 + (\alpha_i \tilde{\eta}_{i,N-1}^2 + \beta_i) + \dots + \prod_{n=0}^{N-2} (\alpha_i \tilde{\eta}_{i,N-1-n}^2 + \beta_i) \right) \\ &\quad + \prod_{n=0}^{N-1} (\alpha_i \tilde{\eta}_{i,N-1-n}^2 + \beta_i) x_i. \end{aligned} \tag{3.A.3}$$

Case $\forall i \in \{1, \dots, d\}$, $\frac{w_i}{1-\beta_i} + c > w_i + \beta_i x_i$: By definition, $\sigma_{i,1}^2 = w_i + (\alpha_i \tilde{\eta}_i^2 + \beta_i) x_i$, $\tilde{\eta} \sim \mathcal{N}(0, \Gamma)$: $\sigma_{i,1}^2$ can be seen as a function of $\tilde{\eta}_i^2 \geq 0$. For small values of $\tilde{\eta}_i^2$, with non null probability, $w_i + \beta_i x_i \leq \sigma_{i,1}^2 < \frac{w_i}{1-\beta_i} + c$: $N_c = 1$ works.

Case $\exists i \in \{1, \dots, d\}$, $\frac{w_i}{1-\beta_i} + c < w_i + \beta_i x_i$: Let $\epsilon > 0$ s.t. $\alpha_i \epsilon^2 + \beta_i < 1$. Let us consider the event $\{|\tilde{\eta}_{i,n}| < \epsilon, n = 0, \dots, N-1, i = 1, \dots, d\}$. From (3.A.3),

$$\begin{aligned} \sigma_{i,N}^2 &\leq w_i \sum_{n=0}^{N-1} (\alpha_i \epsilon^2 + \beta_i)^n + (\alpha_i \epsilon^2 + \beta_i)^N x_i \\ &= w_i \frac{1 - (\alpha_i \epsilon^2 + \beta_i)^N}{1 - (\alpha_i \epsilon^2 + \beta_i)} + (\alpha_i \epsilon^2 + \beta_i)^N x_i. \end{aligned}$$

This quantity goes to $\frac{w_i}{1-(\alpha_i \epsilon^2 + \beta_i)}$ when N goes to infinity. Given $c > 0$, let us choose ϵ such that $\frac{w_i}{1-(\alpha_i \epsilon^2 + \beta_i)} < \frac{w_i}{1-\beta_i} + \frac{c}{2}$. We can take N_c finite such that $(\alpha_i \epsilon^2 + \beta_i)^{N_c} x_i < \frac{c}{2}$. Then $w_i \frac{1-(\alpha_i \epsilon^2 + \beta_i)^{N_c}}{1-(\alpha_i \epsilon^2 + \beta_i)} + (\alpha_i \epsilon^2 + \beta_i)^{N_c} x_i < \frac{w_i}{1-\beta_i} + c$ and then

$$\mathbb{P}_{\mathbf{x}} \left(\sigma_{i,N_c}^2 \in \left(\frac{w_i}{1-\beta_i}, \frac{w_i}{1-\beta_i} + c \right) \right) \geq \mathbb{P} (|\tilde{\eta}_{i,n}| < \epsilon, n = 1, \dots, N_c, i = 1, \dots, d) > 0.$$

□

Corollary 3.A.4 (Irreducibility). *The Markov process $\{\sigma_t^2\}_{t \in \mathbb{Z}}$ taking values in \mathcal{S} is Lebesgue-irreducible [MT09, Proposition 4.2.1 (ii)], e.g. for any $A \subset \mathcal{B}(\mathcal{S})$ with strictly positive Lebesgue measure, for any $\mathbf{x} \in \mathcal{S}$*

$$\mathbb{P}_{\mathbf{x}}[\arg \min\{n \geq 1 \mid \sigma_n^2 \in A\} < \infty] > 0.$$

Proof. Let $A \subset \mathcal{B}(\mathcal{S})$, $\lambda(A) > 0$ (where $\lambda(\cdot)$ denotes the Lebesgue measure) and $\mathbf{x} \in \mathcal{S}$.

We define $\tau_A := \arg \min\{n \geq 1 \mid \sigma_n^2 \in A\}$.

Let $c > 0$ such that $\lambda\left(A \cap \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i} + c, +\infty\right)\right) > 0$. It exists since

- $c \mapsto \lambda\left(A \cap \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i} + c, +\infty\right)\right)$ is increasing when c decreases,
- $\lim_{c \rightarrow 0} \lambda\left(A \cap \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i} + c, +\infty\right)\right) = \lambda(A) > 0$.

Let ϵ and N_c defined in lemma 3.A.3 s.t. $\mathbb{P}_{\mathbf{x}}\left(\sigma_{N_c}^2 \in \left(\frac{w_i}{1-\beta_i}, \frac{w_i}{1-\beta_i} + c\right)\right) > 0$. Then we can show that $\mathbb{P}_{\mathbf{x}}(\tau_A = N_c + 1) > 0$.

Indeed,

$$\begin{aligned} \mathbb{P}_{\mathbf{x}}(\sigma_{N_c+1}^2 \in A) &\geq \mathbb{P}_{\mathbf{x}}\left(\sigma_{N_c+1}^2 \in A \cap \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i} + c, +\infty\right), \sigma_{N_c}^2 \in \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, \frac{w_i}{1-\beta_i} + c\right)\right) \\ &= \mathbb{E}\left[\underbrace{\mathbb{P}\left(\sigma_{N_c+1}^2 \in A \cap \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i} + c, +\infty\right) \mid \sigma_{N_c}^2\right)}_{p(\sigma_{N_c}^2)} \mathbb{1}_{\sigma_{N_c}^2 \in \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, \frac{w_i}{1-\beta_i} + c\right)}\right], \end{aligned}$$

where $p(\sigma_{N_c}^2) > 0$ for any $\sigma_{N_c}^2$, and $\left\{\sigma_{N_c}^2 \in \prod_{i=1}^d \left(\frac{w_i}{1-\beta_i}, \frac{w_i}{1-\beta_i} + c\right)\right\}$ is of non null probability.

Hence $\mathbb{P}_{\mathbf{x}}(\tau_A = N_c + 1) > 0$ and $\{\sigma_t^2\}_{t \in \mathbb{Z}}$ is Lebesgue-irreducible. \square

3.A.4 GARCH-CCC density

Let ψ be a test function. Let us denote g_{Γ} the density of $\tilde{\boldsymbol{\eta}} \sim \mathcal{N}(0, \Gamma)$. We want to compute the density of the vector $\tilde{\boldsymbol{\eta}}^2 = \{\tilde{\eta}_i^2\}_{i=1}^d$:

$$\begin{aligned} \mathbb{E}\psi(\tilde{\boldsymbol{\eta}}^2) &= \int_{\mathbb{R}^d} \psi\left(\{x_i^2\}_{i=1}^d\right) g_{\Gamma}\left(\{x_i\}_{i=1}^d\right) dx_1 \dots dx_d \\ &= \sum_{\substack{s_i = \pm \\ 1 \leq i \leq d}} \int_{\mathbb{R}_{s_1} \dots \mathbb{R}_{s_d}} \psi\left(\{x_i^2\}_{i=1}^d\right) g_{\Gamma}\left(\{x_i\}_{i=1}^d\right) dx_1 \dots dx_d. \end{aligned}$$

We denote $s_i = \pm$ sign symbol denoting the part of \mathbb{R} interval on which x_i is integrated. We can make a change of variable: denoting

$$z_i = \begin{cases} x_i & \text{if } s_i = +, \\ -x_i & \text{if } s_i = -, \end{cases}$$

and $\phi_{\mathbf{s}}$ s.t. $\mathbf{x} = \phi_{\mathbf{s}}(\mathbf{z})$, with $\mathbf{s} = (s_1, \dots, s_d)$ then

$$\mathbb{E}\psi(\tilde{\boldsymbol{\eta}}^2) = \sum_{\substack{s_i = \pm \\ 1 \leq i \leq d}} \int_{(\mathbb{R}_+)^d} \psi\left(\{z_i^2\}_{i=1}^d\right) g_{\Gamma}(\phi_{\mathbf{s}}(\mathbf{z})) dz_1 \dots dz_d.$$

Taking $u_i = z_i^2$, by change of variable, we retrieve $\tilde{\boldsymbol{\eta}}^2$ density:

$$\mathbb{E}\psi(\tilde{\boldsymbol{\eta}}^2) = \sum_{\substack{s_i = \pm \\ 1 \leq i \leq d}} \int_{(\mathbb{R}_+)^d} \psi\left(\{u_i\}_{i=1}^d\right) g_{\Gamma}\left(\phi_{\mathbf{s}}\left(\{\sqrt{u_i}\}_{i=1}^d\right)\right) \left(\prod_{i=1}^d \frac{1}{2\sqrt{u_i}}\right) du_1 \dots du_d.$$

Denoting p the density of $\tilde{\boldsymbol{\eta}}^2$,

$$p(\mathbf{u}) = \sum_{\substack{s_i = \pm \\ 1 \leq i \leq d}} g_{\Gamma} \left(\phi_{\mathbf{s}} \left(\{\sqrt{u_i}\}_{i=1}^d \right) \right) \left(\prod_{i=1}^d \frac{1}{2\sqrt{u_i}} \mathbb{1}_{u_i > 0} \right).$$

Properties: p is continuous and positive on $(\mathbb{R}_+^*)^d$.

Let us derive the density of $\sigma_t^2 | \sigma_{t-1}^2 = \mathbf{x}$. By definition of σ_t^2 , it is equal to the density of $\{w_i + (\alpha_i \tilde{\eta}_i^2 + \beta_i) x_i\}_{i=1}^d$, for $\tilde{\boldsymbol{\eta}} \sim \mathcal{N}(0, \Gamma)$.

Denoting ϕ a test function, by the change of variable $v_i = w_i + (\alpha_i u_i + \beta_i) x_i$,

$$\begin{aligned} \mathbb{E} [\phi(\sigma_t^2) | \sigma_{t-1}^2 = \mathbf{x}] &= \int_{\mathbb{R}^d} \phi(\{w_i + (\alpha_i u_i + \beta_i) x_i\}_{i=1}^d) p(\mathbf{u}) du_1 \dots du_d \\ &= \int_{\mathbb{R}^d} \phi(\mathbf{v}) \underbrace{p \left(\left\{ \frac{v_i - w_i - \beta_i x_i}{\alpha_i x_i} \right\} \right) \prod_{i=1}^d \frac{1}{\alpha_i x_i} \mathbb{1}_{(w_i + \beta_i x_i, +\infty)}(v_i)}_{g_{\mathbf{x}}(\mathbf{v})} dv_1 \dots dv_d. \end{aligned}$$

Hence

$$g_{\mathbf{x}}(\mathbf{v}) := p \left(\left\{ \frac{v_i - w_i - \beta_i x_i}{\alpha_i x_i} \right\} \right) \prod_{i=1}^d \frac{1}{\alpha_i x_i} \mathbb{1}_{(w_i + \beta_i x_i, +\infty)}(v_i). \quad (3.A.4)$$

3.B Hill estimator

Second order condition: Denoting $U(y) = \overline{F}^{-1}(1/y)$, we say that U satisfies a second order condition if there exists $\gamma > 0$, $\rho \leq 0$ and A of constant sign, $\lim_{t \rightarrow \infty} A(t) = 0$ s.t.

$$\frac{1}{A(t)} \left[\frac{U(tx)}{U(t)} - x^\gamma \right] \xrightarrow{t \rightarrow \infty} x^\gamma \frac{x^\rho - 1}{\rho} \quad \forall x > 0.$$

Under second order condition, if additionally $\sqrt{k}A(n/k) \xrightarrow{n \rightarrow \infty} \lambda < \infty$, $\lim_{n \rightarrow \infty} \frac{k}{n} = \frac{k(n)}{n} = 0$ and $\lim_{n \rightarrow \infty} k = +\infty$, then the Hill estimator satisfies a central limit theorem [DHF06, Theorem 3.2.5]:

$$\sqrt{k}(\hat{\gamma}_n - \gamma) \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{N}(\lambda/(1 - \rho), \gamma^2). \quad (3.B.1)$$

Choice of k to minimize the asymptotic MSE: From [DHF06, equation 3.2.13 p.77], the asymptotic mean-squared error is given by

$$\text{asMSE} = \frac{\gamma^2}{k} + \frac{A^2(n/k)}{(1 - \rho)^2}. \quad (3.B.2)$$

Fraga Alves et al [FAGdHN07] propose an explicit function A and constant ρ in the case of Fréchet attraction domain in their Example 4.3, for $\gamma \neq 1$:

$$A(t) = -\frac{1 - \gamma}{2t} \quad \text{and} \quad \rho = -1. \quad (3.B.3)$$

Minimizing the asymptotic MSE for these functions, we get:

$$k_{opt} = 2 \left(\frac{\gamma}{1 - \gamma} n \right)^{\frac{2}{3}}. \quad (3.B.4)$$

In particular, for $k = k_{opt}$, $\sqrt{k}A(n/k) = -\sqrt{2\gamma} < \infty$, so the central limit theorem (3.B.1) holds.

Proof. Plugging (3.B.3) in the asymptotic MSE equation (3.B.2), we get:

$$\text{asMSE} = \frac{\gamma^2}{k} + \frac{1}{4} \left(\frac{1-\gamma}{2n/k} \right)^2 = \frac{\gamma^2}{k} + \left(\frac{1-\gamma}{4n} \right)^2 k^2.$$

Its derivative in k is:

$$\frac{\partial \text{asMSE}}{\partial k} = -\frac{\gamma^2}{k^2} + 2 \left(\frac{1-\gamma}{4n} \right)^2 k = -\frac{\gamma^2}{k^2} + \frac{1}{2^3} \left(\frac{1-\gamma}{n} \right)^2 k$$

which is positive for $k \geq k_{opt}$ and negative for $k \leq k_{opt}$, where k_{opt} zero of $\frac{\partial \text{asMSE}}{\partial k}$ is given by (3.B.4). \square

3.C Additional properties

Extension of the Lyapunov condition to the lagged chain: This paragraph is dedicated to the extension of the Lyapunov condition $\mathbf{H}_{\mathcal{L}}$ to lagged Markov chains, e.g. the fact that if we have a Lyapunov condition $\mathbf{H}_{\mathcal{L}}$ for $\{V_t\}_{t \geq 1}$ then we have $\mathbf{H}_{\mathcal{L}}$ for lagged chain $\{V_{tH}\}_{t \geq 1}$, $H \in \mathbb{N}^*$.

In Meyn and Tweedie's book, the statement that the extension is possible can be found in [MT09, Theorem 15.3.4] page 383.

First, let us notice that the notion of H -skeleton and H -lagged Markov chain coincide. Skeletons are defined page 62 in [MT09] as the chain $V^{(H)}$ with transition law:

$$\mathbb{P}(V_{tH+1} \in A | V_1 = x) = P^{tH}(x, A), \quad A \subset \mathcal{B}(\mathcal{S}).$$

Hence it is the same Markov chain than $\{V_{tH+1}\}_{t \in \mathbb{N}^*}$.

Let us recall Meyn and Tweedie's result:

Theorem 3.C.1 (Extension of $\mathbf{H}_{\mathcal{L}}$ to lagged processes [MT09, Theorem 15.3.4]). *Suppose that V is a Φ -irreducible and aperiodic. If V satisfied $\mathbf{H}_{\mathcal{L}}$ with a petite set \mathcal{C} then for any H -skeleton, the function \mathcal{L} also satisfies $\mathbf{H}_{\mathcal{L}}$ for some set \mathcal{C}' which is petite for the H -skeleton.*

The Meyn and Tweedie's result is a qualitative result. In what follows, we give a quantitative result on the extension of the Lyapunov condition to lagged Markov chains, in which the Lyapunov drift condition constants for the lagged process are explicit. In this way, we can quantify explicitly the dependence of the constant in the lag period H .

Proposition 3.C.2 (Lyapunov condition for lagged Markov chain). *Assume that the Lyapunov drift criteria $\mathbf{H}_{\mathcal{L}}$ is verified for $\{V_t, t \in \mathbb{N}^*\}$. Then, for any $H \in \mathbb{N}^*$, the H -lagged chain $\{V_{nH+1}, n \in \mathbb{N}^*\}$ also satisfies a drift condition:*

$$\mathbb{E}[\mathcal{L}(V_{H+1}) | V_1 = x] \leq d_H \mathcal{L}(x) + b_H \mathbf{1}_{x \in \mathcal{C}^{(H)}}.$$

where

$$d_H = \frac{1 + \delta^H}{2}, \quad \mathcal{C}^{(H)} = \mathcal{C} \cup \left\{ x \in \mathcal{S} \mid |x| \leq R^{(H)} \right\},$$

$$b_H = \sup_{x \in \mathcal{C}^{(H)}} \left(b \frac{1 - \delta^H}{1 - \delta} - \frac{1 - \delta^H}{2} \mathcal{L}(x) \right),$$

and $R^{(H)} > 0$ s.t. for every $x \in \mathcal{S}$ s.t. $|x| > R^{(H)}$, $\frac{1 - \delta^H}{2} \mathcal{L}(x) \geq b \frac{1 - \delta^{H-1}}{1 - \delta}$.

Proof. Let $H \in \mathbb{N}^*$ be fixed. Let P denotes the transition kernel associated to $\{V_t\}$: using standard Markov chain notations, $\mathbb{E}[g(V_2) | V_1 = x] = Pg(x)$, $\mathbb{E}[g(V_{H+1}) | V_1 = x] = P^H g(x)$. We want to show that there is a $d_H \in (0, 1)$, a constant b_H and a set $\mathcal{C}^{(H)}$ s.t. for all $x \in \mathcal{S}$:

$$P^H \mathcal{L}(x) \leq d_H \mathcal{L}(x) + b_H \mathbf{1}_{\mathcal{C}^{(H)}}(x).$$

Let's apply recursively the initial drift criteria equation $\mathbf{H}_{\mathcal{L}}$: $P\mathcal{L}(x) \leq \delta\mathcal{L}(x) + b\mathbb{1}_{\mathcal{C}}(x)$. We get, for $x \in \mathcal{S}$:

$$\begin{aligned}
P^2\mathcal{L}(x) &\leq P(\delta\mathcal{L}(x) + b\mathbb{1}_{\mathcal{C}}(x)) = \delta P\mathcal{L}(x) + bP(x, \mathcal{C}) \\
&\leq \delta^2\mathcal{L}(x) + \delta b\mathbb{1}_{\mathcal{C}}(x) + bP(x, \mathcal{C}) \\
&\leq \delta^2\mathcal{L}(x) + b + \delta b\mathbb{1}_{\mathcal{C}}(x), \\
P^3\mathcal{L}(x) &\leq \delta^3\mathcal{L}(x) + b + \delta b + \delta^2 b\mathbb{1}_{\mathcal{C}}(x), \\
&\vdots \\
P^H\mathcal{L}(x) &\leq \delta^H\mathcal{L}(x) + b\sum_{k=0}^{H-2}\delta^k + \delta^{H-1}b\mathbb{1}_{\mathcal{C}}(x) \\
&= \delta^H\mathcal{L}(x) + b\frac{1-\delta^{H-1}}{1-\delta} + \delta^{H-1}b\mathbb{1}_{\mathcal{C}}(x). \tag{3.C.1}
\end{aligned}$$

Let $d_H = \frac{1+\delta^H}{2} \in (0, 1)$ since $\delta \in (0, 1)$. Let $R^{(H)} > 0$ s.t. for all $x \in \mathcal{S}$, ($|x| > R^{(H)} \Rightarrow \left(b\frac{1-\delta^{H-1}}{1-\delta} \leq \frac{1-\delta^H}{2}\mathcal{L}(x)\right)$ (it exists since $\lim_{|x| \rightarrow \infty} \mathcal{L}(x) = +\infty$). Let $\mathcal{C}^{(H)} = \mathcal{C} \cup \{x \in \mathcal{S} \mid |x| \leq R^{(H)}\}$. Then, outside $\mathcal{C}^{(H)}$, using the trick $\delta^H = \frac{1+\delta^H}{2} - \frac{1-\delta^H}{2}$, since $|x| > R^{(H)}$,

$$\begin{aligned}
\delta^H\mathcal{L}(x) + b\frac{1-\delta^{H-1}}{1-\delta} &= \frac{1+\delta^H}{2}\mathcal{L}(x) + b\underbrace{\frac{1-\delta^{H-1}}{1-\delta} - \frac{1-\delta^H}{2}}_{\leq 0}\mathcal{L}(x) \\
&\leq \frac{1+\delta^H}{2}\mathcal{L}(x) := d_H\mathcal{L}(x).
\end{aligned}$$

For $x \in \mathcal{C}^{(H)}$,

$$\begin{aligned}
P^H\mathcal{L}(x) &\leq \delta^H\mathcal{L}(x) + b\frac{1-\delta^{H-1}}{1-\delta} + \delta^{H-1}b \\
&= d_H\mathcal{L}(x) + (\delta^H - d_H)\mathcal{L}(x) + b\frac{1-\delta^{H-1}}{1-\delta} + \delta^{H-1}b.
\end{aligned}$$

Making the simplification: $b\frac{1-\delta^{H-1}}{1-\delta} + \delta^{H-1}b = b\frac{1-\delta^H}{1-\delta}$ and $\delta^H - d_H = \delta^H - \frac{1+\delta^H}{2} = -\frac{1-\delta^H}{2}$, denoting

$$b_H = \sup_{x \in \mathcal{C}^{(H)}} \left(b\frac{1-\delta^H}{1-\delta} - \frac{1-\delta^H}{2}\mathcal{L}(x) \right),$$

then

$$P^H\mathcal{L}(x) \leq d_H\mathcal{L}(x) + \left(b\frac{1-\delta^H}{1-\delta} - \frac{1-\delta^H}{2}\mathcal{L}(x) \right) \leq d_H\mathcal{L}(x) + b_H. \tag{3.C.2}$$

Thus the wanted inequality $P^H\mathcal{L}(x) \leq d_H\mathcal{L}(x) + b_H\mathbb{1}_{\mathcal{C}^{(H)}}(x)$ is satisfied. \square

The following result gives the irreducibility and the aperiodicity of the lagged chain without giving explicitly a Lyapunov condition as we did in (3.C.2).

Proposition 3.C.3 (Extension of $\mathbf{H}_{\mathcal{S}}$ to lagged processes [MT09][Proposition 5.4.5 - (iii) page 114].) *Suppose that V is a Φ -irreducible Markov chain. If V is aperiodic then every skeleton is Φ -irreducible and aperiodic.*

Martingale concentration We recall here the Burkholder's inequality.

Theorem 3.C.4 (Burkholder's inequality [HH80, Theorem 2.10]). *Let $\{S_i, \mathcal{F}_i, 1 \leq i \leq n\}$ be a real-valued martingale, and $1 < p < \infty$. Denoting $X_1 = S_1$ and $X_i = S_i - S_{i-1}, 2 \leq i \leq n$, then there exist constants $c(p)$ and C_p^B depending only on p such that*

$$c(p) \mathbb{E} \left[\left| \sum_{i=1}^n X_i^2 \right|^{p/2} \right] \leq \mathbb{E} \left[\left| S_n \right|^p \right] \leq C_p^B \mathbb{E} \left[\left| \sum_{i=1}^n X_i^2 \right|^{p/2} \right] \quad (3.C.3)$$

where suitable constants are given by $c(p)^{-1} = (18p^{1/2}q)^p$ and $C_p^B = (18pq^{1/2})^p$, where $p^{-1} + q^{-1} = 1$.

Part III

Uncertain quantification for portfolio

CHAPTER 4

UNCERTAIN QUANTIFICATION FOR PORTFOLIO SHARPE RATIO AND ALLOCATION

Abstract. In this chapter, application of uncertainty quantification (UQ) techniques via stochastic approximation (SA) is investigated in the financial portfolio context. We want to take advantage of the algorithm developed in [CFGS20] which enables one to compute efficiently the Polynomial Chaos Expansion (PCE) of the solution of a zero-expectation problem. Two main applications are considered: UQSA for the Sharpe ratio of the portfolio and UQSA for the portfolio allocation. In our applications, we apply this approach to express these quantities in terms of the correlation considered uncertain.

4.1 Introduction

In this chapter, we are going to apply tools from Uncertainty Quantification (UQ) to the financial portfolio context. Three main kinds of uncertainty can be mentioned:

- *the statistical uncertainty* arising from the statistical error inherent in measurements on a finite number of data;
- *the parametric uncertainty* which corresponds to parameters misspecification assuming a parametric model (this is the type of uncertainty we will focus on);
- *the structural uncertainty* or *model error*.

In the general UQ setting, we are interested in the propagation of the parametric uncertainty, modeled by θ taking values in Θ and following a distribution π , on the output of the model $z^*(\theta)$. z^* and θ are linked through the model function h : $h(z^*(\theta), \theta) = 0$, π -a.e. More specifically, we might be interested in the distribution of the output $z^*(\theta)$ or of a function of the output, $\text{QoI}(z^*(\theta))$ (where QoI stands for Quantity of Interest):

$$\theta \sim \pi \Rightarrow z^*(\theta), \text{QoI}(z^*(\theta)).$$

In practice, we have access to observable data D_θ and we can assume that they follow a specific parametric model, so that the variability of the data is encompassed through the distribution of the data $\mu(dv, \theta)$. The parameter θ can help us to account for parametric error as well as measurement error. For example, in finance, data can consist in financial returns

$\{R_t(\theta), 1 \leq t \leq T\}$. θ can represent a stochastic mean for the returns and account for the estimation error (for example, by taking $\theta = \hat{\mu} + \frac{\sigma}{\sqrt{T}}\mathcal{N}(0, 1)$ if we assume a Gaussian distribution for the returns), or following the Black Litterman approach [BL91], can incorporate the asset manager views in the estimation of the trend (then, $\theta = \mathbb{E}[\hat{\mu}|\text{investor views}] + \text{Gaussian noise}$). The output can be a portfolio allocation and h would represent the gradient of the objective function that the investor seeks to minimize. The Quantity of Interest could be the Sharpe ratio or the Value-at-Risk associated to the portfolio.

4.1.1 UQ techniques

To quantify uncertainties, a global sensitivity analysis approach can be used [SRA⁺08], with Sobol indices for independent model inputs [Sob01, Owe14], or Shapley values when the inputs are dependent [OP17]. Another strategy is to use the Chaos analysis of the optimisation problem with respect to the model inputs [BGP20, CFGS20].

Spectral methods for UQ (see [LMK10] for theory and applications) corresponds to the search for a spectral projection of the problem (on a stochastic basis). They rely on the fact that when a function of θ is squared integrable with regards to π , it admits a decomposition on an orthogonal basis of functions of θ , via Karhunen Loève expansion for example. Polynomial Chaos Expansion (PCE) corresponds to the case when the basis is formed of orthonormal polynomials $\{B_i\}_{i \in \mathbb{N}}$:

$$z^*(\theta) = \sum_{i \geq 0} u_i B_i(\theta), \quad \forall \theta \in \Theta \subset \mathbb{R}^p, \quad \text{and} \quad \int_{\Theta} B_i(\theta) B_j(\theta) \pi(d\theta) = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{else.} \end{cases}$$

Spectral methods are adapted for the case when $z^*(\theta)$ is a smooth function of θ , since we are looking for a projection of z^* on a basis of polynomial functions. In our application cases, we will consider cases where $z^*(\theta)$ is a smooth function of θ , so the spectral method is relevant.

Using this fact, the aim is to find a good approximation of this expansion. In practice, one evaluates a truncated version of this infinite sum representation. Then $z^*(\theta) \approx \sum_{i=0}^m u_i B_i(\theta)$ and $\text{QoI}(z^*(\theta)) \approx \text{QoI}(\sum_{i=0}^m u_i B_i(\theta))$.

The evaluation can be done using non-intrusive techniques (using the underlying model as a *black box* and finding the optimal truncated function minimizing the distance to the true output on specific values of $\theta \sim \pi$ for which the code has been used), or by taking advantage of the form of the problem (via Galerkin method for example, by finding specific equations/models satisfied by the coefficients of the expansion -this method is often applied for problems involving PDEs, such as in [LMRN⁺02, BGP20]).

When h takes the form of an expectation,

$$h(z, \theta) = \mathbb{E}[H(z, V_\theta, \theta)] = \int_{\mathcal{V}} H(z, v, \theta) \mu(\theta, dv), \quad \pi - \text{a.e.}, \quad (4.1.1)$$

then $z^*(\theta)$ can be computed using a Stochastic Approximation (SA) approach. It has been done by [CFGS20] who have proposed an algorithm to compute efficiently the coefficients of the projected output.

4.1.2 UQSA setting

We call Uncertainty Quantification Stochastic Approximation (UQSA) setting the class of problems which can be formulated as a zero-search problem of type $\mathbb{E}[H(z, V)] = 0$, as in usual Stochastic Approximation (SA) setting, but with random variables $V \sim \mu(dv, \theta)$ depending on an uncertain parameter, and an objective function H which can also depend on θ . In this setting, $z^*(\theta)$ is the zero of (4.1.1).

The UQSA algorithm [CFGS20] (described in Appendix 4.B) is an iterative procedure to find the decomposition $z^*(\theta) = \sum_i u_i B_i(\theta)$ for problems of the form (4.1.1). It is a convenient formulation for problems corresponding to minimization of convex functions, typically in portfolio optimization when the objective function is convex.

4.1.3 Explored financial applications

In this chapter, we explore the three following UQ applications for the financial context:

1. UQSA for the Sharpe ratio: the Sharpe ratio is one of the main profitability indices looked at by investors to measure the *interest* of an investment. Denoting R the excess return of the portfolio, the Sharpe ratio is defined as

$$\text{SR} = \frac{\mathbb{E}[R]}{\sqrt{\text{Var}[R]}}.$$

In the case when we assume an uncertainty on R distribution, the Sharpe ratio is also uncertain. In this case, $z^*(\theta) = \text{SR}(\theta)$ and we are interested in the distribution of $\text{SR}(\theta)$. A difficulty of the approach is to find a way to express the Sharpe ratio as the zero of an expectation. Once this difficulty has been overcome, it is possible to see the impact of parameters uncertainty on the distribution of the Sharpe ratio. We give illustrations with uncertain correlation following non-linear function of a uniform distribution.

2. UQSA for portfolio allocation with constraints: here we are interested in the distribution of the portfolio allocation when the objective function is known up to uncertain parameters.

We are interested in portfolio optimization of the form:

$$\mathbf{w} \in \mathcal{W} \text{ solution of } \mathbb{E}[F(\mathbf{w}, V)] = 0,$$

where V denotes the assets returns, $\mathcal{W} \subset \mathbb{R}^d$ denotes the set of constraints, and can correspond for example to the case: $\mathcal{W} = \left\{ \mathbf{w} \in (\mathbb{R}_+)^d : \mathbf{1}_d^T \mathbf{w} = 1 \right\}$ (budget and positivity constraints). Usual SA algorithms produce sequences in \mathbb{R}^d , which are not guaranteed to stay in \mathcal{W} . In order to force the iteration outputs to stay in \mathcal{W} , we could have considered doing projections on \mathcal{W} , using à la Chen SA algorithm [CZ86]. For example, in the non-uncertain context, the update in the SA approach with projection would consist in

$$\mathbf{w}_{k+1} = \Pi_{\mathcal{W}}(\mathbf{w}_k - \gamma_{k+1} H(\mathbf{w}_k, V_{k+1})),$$

where $\Pi_{\mathcal{W}}$ stands for the Euclidean projection on \mathcal{W} . The UQSA algorithm allows for projection of the expansion coefficients on some convex set

$$u_i^{(k+1)} = \Pi \left(u_i^{(k)} - \gamma_{k+1} \langle H(\mathbf{w}_k, V_{k+1}, \cdot), B_i(\cdot) \rangle_{\pi} \right)$$

(see Sections 4.2.1 and 4.2.2 to see UQSA notations and iterations equations). But even if we project the u_i on a convex set, we have no guarantee that the whole PCE approximation $\sum_i u_i^{(k)} B_i(\theta)$ will remain in \mathcal{W} . We therefore prefer to stick to an approach without projection and apply a change of variable $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ s.t. $\mathbf{w} = \ell(\mathbf{y})$. This way, we benefit from the convergence results of the UQSA algorithm. $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ s.t. $\mathbf{w} = \ell(\mathbf{y})$.

In this application case, the output of the problem will be $z^*(\theta) = \mathbf{y}(\theta)$ and the function of interest is the portfolio allocation $\mathbf{w}(\theta) = \ell(\mathbf{y}(\theta)) = \text{QoI}(z^*(\theta))$.

4.2 UQSA algorithm: principle and main assumptions

In this section, we describe the algorithm used in this chapter (Sections 4.2.1 and 4.2.2) and we give the underlying assumptions under a which a stability and convergence result holds (Section 4.2.3). At the end of this section, we provide the main metrics which will serve to judge the convergence of the algorithm.

4.2.1 Principle and notations

When dealing with a problem of type (4.1.1), it is possible to find an approximation $z^*(\theta) = \sum_{i \leq m_k} u_i^{(k)} B_i(\theta)$ using the UQSA algorithm [CFGS20] (described in appendix 4.B). In what follows, we detail the general notations and PCE theory basics.

Notations: Let \mathcal{V} be a metric space endowed with its Borel σ -field, Θ be a subset of \mathbb{R}^p , and $H : \mathbb{R}^q \times \mathcal{V} \times \Theta \rightarrow \mathbb{R}^q$. In our applications, $q = 1$ (Sharpe ratio) or $q = d - 1$ (portfolio allocation). Let π be a probability distribution and μ be a transition kernel from Θ to \mathcal{V} . We define the scalar product induced by π by

$$\langle f; g \rangle_\pi := \int_{\Theta} f(\theta)g(\theta)\pi(d\theta), \quad (4.2.1)$$

for any measurable functions $f, g : \Theta \rightarrow \mathbb{R}$. By extension, for measurable functions $f = (f_1, \dots, f_q) : \Theta \rightarrow \mathbb{R}^q$ and $g : \Theta \rightarrow \mathbb{R}$, we write in vector form

$$\langle f; g \rangle_\pi := \begin{pmatrix} \langle f_1; g \rangle_\pi \\ \dots \\ \langle f_q; g \rangle_\pi \end{pmatrix}.$$

We denote by $L_{2,q}^\pi$ the Hilbert space of functions $f : \Theta \rightarrow \mathbb{R}^q$ such that the norm $\|f\|_\pi := \sqrt{\sum_{i=1}^q \langle f_i; f_i \rangle_\pi}$ is finite.

We consider the following problem:

$$\text{Finding } \phi^* \text{ in } L_{2,q}^\pi \text{ such that } \int_{\mathcal{V}} H(\phi^*(\theta), v, \theta) \mu(\theta, dv) = 0, \quad \pi - \text{a.e.} \quad (4.2.2)$$

Polynomial Chaos Expansion: Let $\{\theta \rightarrow B_i(\theta)\}$ be an orthonormal basis of $L_{2,1}^\pi$ for the scalar product (4.2.1). We denote by $l_{2,q}$ the normed vector space of the \mathbb{R}^q -valued sequences $\{u_i, i \in \mathbb{N}\}$ with $\sum_{i \geq 0} |u_i|^2 < \infty$. As is well known, given an orthonormal basis $\{B_i, i \in \mathbb{N}\}$ in L_2^π of $L_{2,1}^\pi$, any function $\phi \in L_{2,q}^\pi$ is characterized by a sequence $\{u_i, i \in \mathbb{N}\}$ in $l_{2,q}$ such that $\phi = \sum_{i \geq 0} u_i B_i$.

There is a natural isomorphism $\mathbf{Is} : l_{2,q} \rightarrow L_{2,q}^\pi$ given by

$$\phi = \mathbf{Is}(u) = \sum_{i \geq 0} u_i B_i, \quad \text{i.e. } u_i = \langle \phi; B_i \rangle_\pi \text{ for each } i \in \mathbb{N}, \quad (4.2.3)$$

and a corresponding isometry $\|\phi\|_\pi = \|u\|_{l_{2,q}}$.

In this view, the problem (4.2.2) can be restated on $l_{2,q}$ as

$$\text{Finding } u^* \text{ in } \mathbf{T}^* \text{ where}$$

$$\mathbf{T}^* := \left\{ u \in l_{2,q}; \quad \int_{\mathcal{V}} H \left(\sum_{i \geq 0} u_i B_i(\theta), v, \theta \right) \mu(\theta, dv) = 0, \quad \pi - \text{a.e.} \right\}.$$

4.2.2 UQSA algorithm

The SA algorithm consists in the following update of the $\{u_i^{(k)}\}_{i \leq m_k}$:

$$\hat{\phi}^{(k)}(\theta) = \sum_{i=1}^{m_k} u_i^{(k)} B_i(\theta), \quad (4.2.4)$$

$$u_i^{(k+1)} = u_i^{(k)} - \frac{\gamma_{k+1}}{M_{k+1}} \sum_{s=1}^{M_{k+1}} H(\hat{\phi}^{(k)}(\theta_s^{(k+1)}), V_s^{(k+1)}) B_i(\theta_s^{(k+1)}), \quad i \leq m_{k+1},$$

where m_k and M_k are nondecreasing natural sequences, γ_k -the step size- is a decreasing natural sequence, and $(\theta_s^{(k+1)}, V_s^{(k+1)})$, $s = 1, \dots, M_{k+1}$ are independent variables sampled under $\pi(d\theta)\mu(\theta, dv)$. We let m_k go to infinity while γ_k goes to zero when k goes to infinity.

4.2.3 Main assumptions and convergence results

For the UQSA algorithm detailed in Appendix 4.B to converge towards ϕ^* , we need the following assumptions on the functions H , h and on the sequences γ_k , m_k and M_k used in the algorithm. H_{T^*} . The set T^* of zeros is compact and non-empty.

H_{m_k, M_k} . $\{m_k, k \geq 1\}$ and $\{M_k, k \geq 1\}$ are deterministic sequences of positive integers; $\{\gamma_k, k \geq 1\}$ is a deterministic sequence of positive real numbers such that, for some $\kappa > 0$,

$$\sum_{k \geq 1} \gamma_k = \infty, \quad \sum_{k \geq 1} \gamma_k^{1+\kappa} < \infty, \quad \sum_{k \geq 1} \gamma_k^2 \frac{Q_{m_k}}{M_k} < \infty, \quad \sum_{k \geq 1} \gamma_k^{1-\kappa} q_{m_k} < \infty$$

where the sequences $\{q_m, m \in \mathbb{N}\}$ and $\{Q_m, m \in \mathbb{N}\}$ are defined by

$$q_m = \sup_{u^* \in T^*} \sum_{i > m} |u_i|^2, \quad Q_m = \sup_{\theta \in \Theta} \sum_{i \leq m} |B_i(\theta)|^2.$$

H_{int} . $\forall z \in \mathbb{R}^q$,

$$\int_{\Theta} \int_{\mathcal{V}} |H(z, v, \theta)| \mu(\theta, dv) \pi(d\theta) < \infty.$$

$\forall z \in \mathbb{R}^q, \theta \in \Theta$, $h(z, \theta) = \int_{\mathcal{V}} H(z, v, \theta) \mu(dv, \theta)$ exists. For any $\phi \in L_{2,q}^{\pi}$, the mapping $h(\phi(\cdot), \cdot) : \theta \mapsto h(\phi(\theta), \theta)$ is in $L_{2,q}^{\pi}$. The mapping $\phi \mapsto h(\phi(\cdot), \cdot)$ from $L_{2,q}^{\pi}$ into itself is continuous.

H_{sep} . For π -almost every θ , for any $z_{\theta}, z_{\theta}^* \in \mathbb{R}^q$ such that $h(z_{\theta}) \neq 0$ and $h(z_{\theta}^*) = 0$,

$$\langle z_{\theta} - z_{\theta}^*, h(z_{\theta}, \theta) \rangle > 0.$$

H_{quad} . $\exists C_H > 0, \forall z \in \mathbb{R}^q$,

$$\sup_{\theta \in \Theta} \int_{\mathcal{V}} |H(z, v, \theta)|^2 \mu(\theta, dv) \leq C_H (1 + |z|^2).$$

H_{coerc} . $\forall B > 0, \exists C_B > 0, \forall (\phi, \phi^*) \in L_{2,q}^{\pi} \times \mathbf{Is}(T^*)$ such that $\|\phi - \phi^*\|_{\pi} \leq B$,

$$\int_{\Theta} (\phi - \phi^*)(\theta) \cdot h(\phi(\theta), \theta) \pi(d\theta) \geq C_B \min_{\phi \in T^*} \|\phi - \bar{\phi}\|_{\pi}^2.$$

Remarks

- Assumptions H_{T^*} , H_{int} , H_{sep} and H_{quad} are similar to the one required for standard SA algorithm to converge (see [RM51] for a reference on the Robbins Monro algorithm), as well as the assumption H_{m_k, M_k} for the γ_k sequence. These assumptions allow to apply the Robbins-Siegmund Lemma [RS71] to show the convergence of the algorithm. Typically, these assumptions are satisfied when $h(\cdot, \theta)$ can be expressed as the gradient of a convex function, with appropriate choice of sequences.
- The coercivity assumption H_{coerc} will be satisfied for non-degenerated quadratic minimization problem.
- Assumption H_{m_k, M_k} adds some restriction on the choice of the basis $\{B_i(\cdot)\}_{i \in \mathbb{N}}$ (if we deal with unbounded Θ , unbounded functions, such as Hermite polynomials, will be prohibited).

- As we will see later in our applications, condition H_{m_k, M_k} can be made explicit by taking γ_k, m_k and M_k in the form k^α , $\alpha \in \mathbb{R}$.

Under the aforementioned assumptions, the following stability and convergence theorem applies:

Theorem 4.2.1 ([CFGS20] Theorem 1). *Assume assumptions H_{T^*} , H_{m_k, M_k} , H_{int} , H_{sep} and H_{quad} . Let $\{(\theta_s^{(k)}, V_s^{(k)}), 1 \leq s \leq M_k, k \geq 1\}$ be i.i.d. random variables with distribution $\pi(d\theta)\mu(\theta, dv)$. Let $u^{(k)}$ and $\hat{\phi}^{(k)}$ be the outputs of the USA Algorithm (cf. (4.2.4)).*

Stability. For any $\phi^ \in \mathbf{Is}(T^*)$, $\lim_{k \rightarrow +\infty} \|\hat{\phi}^{(k)} - \phi^*\|_\pi$ exists, is finite a.s., and we have*

$$\sup_{k \geq 0} \mathbb{E} \left[\|\hat{\phi}^{(k)} - \phi^*\|_\pi^2 \right] < +\infty.$$

Convergence. In addition, under the additional assumption H_{coerc} , there exists a random variable ϕ^ taking values in $\mathbf{Is}(T^*)$ such that*

$$\lim_{k \rightarrow +\infty} \|\hat{\phi}^{(k)} - \phi^*\|_\pi = 0 \text{ a.s. and, for any } p \in (0, 2), \lim_{k \rightarrow +\infty} \mathbb{E} \left[\|\hat{\phi}^{(k)} - \phi^*\|_\pi^p \right] = 0.$$

4.2.4 Applications: convergence metrics

In practice, in our applications, we will monitor the following error metrics. We will consider situations when we know ϕ^* explicitly. For the sake of checking the theoretical results. At iteration k , the UQSA algorithm provides us with an approximation $\hat{\phi}^{(k)} = \sum_{i=1}^{m_k} u_i^{(k)} B_i$ of ϕ^* (m_k coefficients are evaluated, where $\{m_k\}_{k \in \mathbb{N}^*}$ is a sequence satisfying H_{m_k, M_k}).

The Mean Squared Error (MSE) at iteration k is defined as: $MSE^{(k)} = \mathbb{E} \left[\|\hat{\phi}^{(k)} - \phi^*\|_\pi^2 \right]$.

It can be decomposed as the sum of two errors: $MSE^{(k)} = \left(\mathcal{E}_{SA}^{(k)} \right)^2 + \left(\mathcal{E}_{TR}^{(k)} \right)^2$, where

- \mathcal{E}_{SA} is the stochastic approximation error: $\left(\mathcal{E}_{SA}^{(k)} \right)^2 = \mathbb{E} \left[\sum_{i=0}^{m_k} |u_i^{(k)} - u_i^*|^2 \right]$,
- \mathcal{E}_{TR} is the truncation error: $\left(\mathcal{E}_{TR}^{(k)} \right)^2 = \sum_{i > m_k} |u_i^*|^2$.

In practice, we will evaluate the MSE (and analogously the SA and truncation errors) on multiple macro-runs of the algorithm: for $N_{MSE} \in \mathbb{N}^*$ (typically, $N_{MSE} = 50$),

$$\widehat{MSE}^{(k)} = \frac{1}{N_{MSE}} \sum_{n=1}^{N_{MSE}} \|\hat{\phi}_n^{(k)} - \phi^*\|_\pi^2.$$

We call *total error* the square root of the MSE.

Practical computations:

- The term $\|\hat{\phi}_n^{(k)} - \phi^*\|_\pi^2$ can be approximated by Monte Carlo simulations: for $K_\Theta \in \mathbb{N}^*$ large, $\theta_1, \dots, \theta_{K_\Theta} \stackrel{\text{i.i.d.}}{\sim} \pi$,

$$\|\hat{\phi}_n^{(k)} - \phi^*\|_\pi^2 \approx \frac{1}{K_\Theta} \sum_{l=1}^{K_\Theta} |\hat{\phi}_n^{(k)}(\theta_l) - \phi^*(\theta_l)|^2.$$

- When the $\{u_i^*\}_{i \in \mathbb{N}}$ cannot be computed explicitly but ϕ^* can, we can take advantage of its definition (4.2.3) and approximate them by numerical integration or by Monte Carlo approximation. For $K_\Theta \in \mathbb{N}^*$ large, $\theta_1, \dots, \theta_{K_\Theta} \stackrel{\text{i.i.d.}}{\sim} \pi$, $i \in \mathbb{N}$,

$$u_i^* = \langle \phi^*, B_i \rangle_\pi = \int_{\Theta} \phi^*(\theta) B_i(\theta) \pi(d\theta) \approx \frac{1}{K_\Theta} \sum_{l=1}^{K_\Theta} \phi^*(\theta_l) B_i(\theta_l).$$

- Whether the u_i^* are exactly known or approximated, it is possible to evaluate the SA error via $\left(\widehat{\mathcal{E}}_{SA}^{(k)}\right)^2 = \frac{1}{N_{MSE}} \sum_{n=1}^{N_{MSE}} \sum_{i=0}^{m_k} |u_{i,n}^{(k)} - u_i^*|^2$.
- The truncation error can be obtained by evaluation of $\|\phi^{*(k)} - \phi^*\|_{\pi}^2$, where $\phi^{*(k)}$ denotes the truncated true function: $\phi^{*(k)} = \sum_{i \leq m_k} u_i^* B_i$, via numerical integration or Monte Carlo approximation.

4.3 UQSA for the Sharpe ratio

In this section, we propose an SA scheme to evaluate the Sharpe ratio, and we provide illustration assuming a log-normal distribution for the portfolio returns. We start by recalling the definition of the Sharpe ratio and by motivating the UQSA application to this portfolio metric.

4.3.1 Definition and motivation

We place ourselves in the following setting:

- Let us consider a universe of d assets. Let R_i denote the return of asset i between time 0 and 1.
- The strategy allocation $\{w_i\}_{1 \leq i \leq d}$ is fixed, given by an optimization routine (we could consider that we do not have access to the algorithm and that everything is done in a **black box**). In this Section, there is no uncertainty analysis on \mathbf{w} . We look at the uncertainty on the Sharpe ratio only. We consider \mathbf{w} that depends on the uncertainty in the next section. In our experiments, we will focus on a particular choice of \mathbf{w} : we will take \mathbf{w} as the *tangent portfolio* as defined below in Subsubsection 4.3.4.1.
- The return of the portfolio is $R = \sum_{i=1}^d w_i R_i$.

The Sharpe ratio [Sha66] is one of the main profitability indices that investors monitor to evaluate an investment. It measures the profitability gap between a risky portfolio and a non-risky portfolio, by unity of risk (where the risk is namely the standard deviation of the portfolio). We give a mathematical definition of the Sharpe Ratio.

Definition 4.3.1 (Sharpe ratio). Denoting r the risk free rate, the Sharpe ratio is defined as the difference between the expected portfolio return $\mathbb{E}[R]$ and the risk free rate, divided by the standard deviation of the portfolio return:

$$\text{SR} = \frac{\mathbb{E}[R] - r}{\sqrt{\text{Var}[R]}}.$$

To simplify the notation, we assume that we deal with risk-free rate adjusted portfolio returns " $R = R - r$ ". In practice, the Sharpe ratio is estimated on realized values of the portfolio returns, as the ratio of the portfolio return empirical mean over the empirical standard deviation. In this section, a single investment period is considered. Portfolio returns are calculated from portfolio values as $R = \frac{V_1 - V_0}{V_0}$, where V_1 denotes the portfolio value at time 1.

We are interested in the distribution of $\text{SR}(\theta)$ when the returns distribution is known up to an uncertain parameter $\theta \in \Theta$ (in practice, we will assume an uncertain distribution on the asset returns R_i , and with fixed allocation \mathbf{w} , so that the portfolio return distribution is inferred from the assets returns distribution). Let π be a probability distribution over Θ , μ be a transition kernel from Θ to \mathcal{V} s.t. $\theta \sim \pi$ and $R_\theta \sim \mu(\theta, dv)$. This θ parameter enables to take into account the model error/measurement uncertainty/sensitivity to small perturbation.

The Sharpe ratio is then an uncertain function:

$$\text{SR}(\theta) = \frac{\mathbb{E}[R_\theta]}{\sqrt{\text{Var}[R_\theta]}}$$

on which we want to estimate a projection $\text{SR}(\theta) \approx \sum_i u_i B_i(\theta)$.

Motivation for a study on the Sharpe ratio uncertainty in the fixed \mathbf{w} setting: Here we assume that the portfolio weights \mathbf{w} are fixed, given by a black-box for example, and we are interested in the fluctuations of the Sharpe ratio due to the uncertainty on R_θ . This is a coherent approach since in practice, an investor fixes its allocation using a value of θ he believes is "correct" (a historical θ for example) but in reality, the true θ might deviate from the assumed θ . It is an important question for the investor to know if his estimated Sharpe ratio is not too optimistic.

This question has been tackled in the literature for example in [Bro93]. In this article, Brodie illustrates how much the estimated efficient frontier (the one we can wrongly assumed is correct) is far away and too optimistic compared to the true efficient frontier. The efficient frontier corresponds to the set of attainable portfolio mean and standard-deviation (e.g. such that there is an allocation \mathbf{w} realizing this mean and standard-deviation) which are efficient, meaning we cannot increase the portfolio expected return without increasing its variance. In [BEKL18], a similar experiment is done. Both experiments assumed a Gaussian distribution for the returns.

Many authors tried to define the distribution of the Sharpe ratio under specific assumption for the assets returns. In [Lo02], assuming autocorrelated returns, the asymptotic distribution (when the backtest size goes to infinity) is derived. More recently, under Gaussian distribution, [Ben19] shows that the Sharpe ratio calculated on T Gaussian returns follows a non centered Student distribution, with non-centrality parameter proportional to the true Sharpe ratio:

$$\sqrt{T}\widehat{\text{SR}} \sim \text{Student}_{\sqrt{T}\text{SR}}(d-1).$$

Using these results, it is possible to test for the significativity of the Sharpe ratio. Since in practice, multiple Sharpe ratio can be evaluated on the same data sets, classical tests would be biased (it is the data snooping effect, occurring when "*a given set of data is used more than once for purposes of inference or model selection*" [Whi00]). This is why approaches based on multiple testing [RW05, HL14] and bootstrap [BBLdPZ16] are advocated in the literature.

Our approach is more general in the sense that we can assume any parametric distribution for the assets returns and have access to the distribution of the Sharpe ratio.

Motivation for a UQSA approach: We advocate that a UQSA approach for the Sharpe ratio is relevant for the following reasons:

- The Sharpe ratio is not necessarily explicit if the moment of the portfolio returns are not available in closed form, or hard to compute explicitly: an SA scheme for the Sharpe ratio would be relevant when we can simulate the asset returns R_i but not compute explicitly their moments. For example, in the GARCH model, which is a model on the conditional distribution of the returns, the moments can be hard to get in closed form.
- Here we are interested in the distribution of the Sharpe ratio assuming an uncertainty on the returns model. The relationship between the Sharpe ratio and θ can be complex. This is why a PCE and more particularly a UQSA approach is relevant.

4.3.2 SA scheme for the Sharpe ratio

Naive approach: A naive approach would consist in estimating separately the mean and the variance. For example, one can estimate sample estimates $\hat{\mu}_{\theta_n}, \hat{\sigma}_{\theta_n}^2$ for multiple values of $\theta \in \{\theta_1, \dots, \theta_N\}$ and fit a PCE approximation of these quantities via regression:

- Using that $\mu_\theta = \mathbb{E}[R^\theta]$, $\sigma_\theta^2 = \mathbb{E}[(R^\theta - \mathbb{E}[R^\theta])^2]$, given $\{R_t^\theta, t \leq T\}, \theta \in \{\theta_1, \dots, \theta_N\}$, for each $\theta \in \{\theta_1, \dots, \theta_N\}$, we have a sample estimate of μ_θ and σ_θ :

$$\hat{\mu}_\theta = \frac{1}{T} \sum_{t=1}^T R_t^\theta, \quad \hat{\sigma}_\theta^2 = \frac{1}{T} \sum_{t=1}^T (R_t^\theta - \hat{\mu}_\theta)^2.$$

- Then to have an estimation for any θ , we can do a regression.

For a given loss function ℓ , minimization in u_i, v_i of:

$$\frac{1}{N} \sum_{n=1}^N \ell \left(\hat{\mu}_{\theta_n} - \sum_i u_i B_i(\theta_n) \right), \quad \frac{1}{N} \sum_{n=1}^N \ell \left(\hat{\sigma}_{\theta_n}^2 - \sum_i v_i B_i(\theta_n) \right).$$

Then, under the condition that $\sum_i v_i B_i(\theta) > 0$, an approximation for the Sharpe ratio would be: $\frac{\sum_i u_i B_i(\theta)}{\sqrt{\sum_i v_i B_i(\theta)}}$.

This approach is a batch, non-sequential method: all quantities are calculated once and for all, and it can be computationally heavy if the sample size is large. Alternatively, a sequential approach may be preferred such that at any time, an approximation of the Sharpe ratio is available and can be improved, while being guaranteed to converge. Such a sequential approach is given by UQSA.

One can also use UQSA algorithm twice, separately on the mean and the variance. These approaches require a two-steps estimation, multiplying the source of approximation error and making it more difficult to quantify. This was not the preferred approach. An application of UQSA for a single function H whose zero would give the Sharpe ratio was preferred.

SA formulation for the squared Sharpe ratio: Denoting $\mu_\theta = \mathbb{E}[R_\theta]$ and $\sigma_\theta = \sqrt{\text{Var}[R_\theta]}$, notice that the Sharpe ratio is the solution of the equation $x\sigma_\theta - \mu_\theta = 0$. Since the standard-deviation is not elicitable¹ (cannot be written as an expectation), the Sharpe ratio cannot be expressed directly as the zero of the expectation of a function. But the squared Sharpe ratio can, as shown in the following lemma:

Lemma 4.3.1. *Let R and \tilde{R} be independent and identically distributed portfolio returns. Denote*

$$H(x, R, \tilde{R}) = x \frac{R^2 + \tilde{R}^2}{2} - (1+x)R\tilde{R}.$$

Then the unique solution x^ such that $f(x^*) = \mathbb{E}[H(x^*, R, \tilde{R})] = 0$ is the squared Sharpe ratio.*

Proof. Let's notice that if R is a scalar random variable and \tilde{R} is an independent copy of R :

¹See [Tas14, Slide 23] for a reference on the non-elicitability of the standard-deviation. The variance is also not formally elicitable [Bre17, Example 1.18] but it is *conditionally elicitable* [EKT15, Lemma 3.5] or *jointly elicitable* with the mean [Bre17, Example 1.23], because there exists a bijection between the two first moments of a random variable and the pair (mean, variance).

- $\text{Var}[R] = \mathbb{E}\left[\frac{(R-\tilde{R})^2}{2}\right],$
- $\mathbb{E}[R]^2 = \mathbb{E}[R\tilde{R}].$

Using expectation formulation of the variance and the squared mean, we get the announced H function:

$$\begin{aligned} x = \text{SR}^2 = \frac{\mathbb{E}[R]^2}{\text{Var}[R]} &\Leftrightarrow x\text{Var}[R] - \mathbb{E}[R]^2 = 0, \\ &\Leftrightarrow x\mathbb{E}\left[\frac{(R-\tilde{R})^2}{2}\right] - \mathbb{E}[R\tilde{R}] = 0, \\ &\Leftrightarrow \mathbb{E}\left[\left(x\frac{R^2 + \tilde{R}^2}{2} - (1+x)R\tilde{R}\right)\right] = 0. \end{aligned}$$

□

The Sharpe ratio is then: $\text{SR} = \sqrt{x^*(\theta)}\text{sign}(\mu_\theta).$

In what follows, by economic rational, we will assume that $\mu_\theta > \mathbf{0}$ and that we deal with positive Sharpe ratios.

4.3.3 UQSA assumptions verification

Let

$$\begin{aligned} H(x, R_\theta, \tilde{R}_\theta) &= x\frac{(R_\theta - \tilde{R}_\theta)^2}{2} - R_\theta\tilde{R}_\theta, \\ h(x, \theta) &= \mathbb{E}_{R_\theta, \tilde{R}_\theta}\left[H(x, R_\theta, \tilde{R}_\theta)\right] = x\text{Var}_{R_\theta}[R_\theta] - (\mathbb{E}_{R_\theta}[R_\theta])^2. \end{aligned} \quad (4.3.1)$$

We will use the following notations:

- \mathbb{E}_{R_θ} denotes the expectation under the law of R_θ at fixed θ : $\mathbb{E}_{R_\theta}f(R_\theta) = \int_{\mathcal{V}} f(v)\mu(\theta, dv).$
- $\mathbb{E}_{R_\theta, \theta}$ denotes the expectation under $\mu(\theta, dv) \otimes \pi$: $\mathbb{E}_{R_\theta, \theta}f(R_\theta) = \int_{\Theta} \int_{\mathcal{V}} f(v)\mu(\theta, dv)\pi(d\theta).$

Under the following assumptions on R_θ , we are going to show that the UQSA algorithm is in the good conditions to converge with the function H :

$$H_{unif,2}. \mathbb{E}_{R_\theta, \theta}[R_\theta^2] < \infty \text{ and } \inf_{\theta \in \Theta} \text{Var}_{R_\theta}[R_\theta] \geq v_{\text{inf}}, v_{\text{inf}} > 0,$$

$$H_{unif,4}. \mathbb{E}_{R_\theta}[R_\theta^4] < \infty \text{ and } \sup_{\theta} \mathbb{E}_{R_\theta}[R_\theta^4] < \infty.$$

Proposition 4.3.2. *Let $H(x, R_\theta, \tilde{R}_\theta) = x\frac{(R_\theta - \tilde{R}_\theta)^2}{2} - R_\theta\tilde{R}_\theta$. Under assumptions $H_{unif,2}$ and $H_{unif,4}$, under appropriate choice of sequence γ_k, m_k and M_k satisfying H_{m_k, M_k} , the UQSA stability and convergence theorem applies.*

The approximate solution converges in π -norm to SR_θ^2 :

$$\|x^k - \text{SR}^2\|_\pi \rightarrow 0 \text{ a.e.} \quad (4.3.2)$$

Remark 6. This result is just the application of UQSA stability theorem 4.2.1. For k large enough, $x_k(\theta) \geq 0$, and we can take $y^k(\theta) = \sqrt{x_k(\theta)}$.

The convergence result (4.3.2) implies that $x^k \rightarrow \text{SR}^2$ π -a.e. along a subsequence. Under the assumption, $\text{sign}(\mu_\theta) > 0$, y^k converges to the Sharpe ratio, still along a subsequence.

Proof. We are going to show that the assumptions of the stability and convergence theorem are satisfied.

H_{T*} From Lemma 4.3.1, the zero of h is unique hence **H_{T*}** is satisfied.

H_{int} Since $\mathbb{E}_{R_\theta, \theta} [R_\theta^2] < \infty$ by assumption **H_{unif,2}**, and since H is quadratic in $R_\theta, \tilde{R}_\theta$, H is integrable, and h exists.

From assumption **H_{unif,4}**, $\sup_\theta \mathbb{E}_{R_\theta} R_\theta^2 \leq (\sup_\theta \mathbb{E}_{R_\theta} R_\theta^4)^{1/2} < \infty$, so h is uniformly bounded in θ :

$$\begin{aligned} h(x, \theta) &= x \text{Var}_{R_\theta} [R_\theta] - (\mathbb{E}_{R_\theta} [R_\theta])^2, \\ |h(x, \theta)| &\leq 2 \sup_\theta \mathbb{E}_{R_\theta} R_\theta^2 (1 + |x|). \end{aligned}$$

So, h is linear in x , with a coefficient uniformly bounded in θ , so for any $\phi \in L_{2,q}^\pi$, the mapping $h(\phi(\cdot), \cdot) : \theta \mapsto h(\phi(\theta), \theta)$ is in $L_{2,q}^\pi$ and is continuous.

H_{sep} Let $x^*(\theta) = \left(\frac{\mu_\theta}{\sigma_\theta}\right)^2$ the zero of h defined in (4.3.1). Notice that h can be rewritten as: for $x \in \mathbb{R}$, $h(x, \theta) = \sigma_\theta^2(x - x^*(\theta))$, where $\sigma_\theta^2 \geq v_{\text{inf}} > 0$ from assumption **H_{unif,2}**. It is then easy to see that, for $x \neq x^*(\theta)$ the separation condition $(x - x_\theta^*)h(z, \theta) = \sigma_\theta^2(x - x_\theta^*)^2 > 0$ is satisfied.

H_{quad} By convexity inequality, we have

$$\begin{aligned} &\sup_{\theta \in \Theta} \int_{\mathcal{V} \times \mathcal{V}} |H(x, v, \tilde{v})|^2 \mu(\theta, dv) \mu(\theta, d\tilde{v}) = \sup_\theta \mathbb{E}_{R_\theta, \tilde{R}_\theta} \left| x \frac{(R_\theta - \tilde{R}_\theta)^2}{2} - R_\theta \tilde{R}_\theta \right|^2 \\ &\leq 2 \sup_{\theta \in \Theta} \mathbb{E}_{R_\theta, \tilde{R}_\theta} \left| \frac{(R_\theta - \tilde{R}_\theta)^2}{2} \right|^2 |x|^2 + 2 \sup_\theta \mathbb{E}_{R_\theta, \tilde{R}_\theta} |R_\theta \tilde{R}_\theta|^2 \\ &\leq 4 \sup_{\theta \in \Theta} \mathbb{E}_{R_\theta} [R_\theta^4] |x|^2 + 2 \sup_{\theta \in \Theta} \mathbb{E}_{R_\theta} [R_\theta^4] \\ &\leq C(1 + |x|^2), \end{aligned}$$

with $C = 6 \sup_{\theta \in \Theta} \mathbb{E}_{R_\theta} [R_\theta^4] < \infty$ from assumption **H_{unif,4}**.

H_{coerc} For $\phi \in L_{2,1}^\pi$, $h(\phi, \theta) = \sigma_\theta^2(\phi(\theta) - x^*(\theta))$ hence

$$\begin{aligned} \int_{\Theta} (\phi - x^*)(\theta) \cdot h(\phi(\theta), \theta) \pi(d\theta) &= \int_{\Theta} \sigma_\theta^2(\phi - x^*)^2(\theta) \pi(d\theta) \\ &\stackrel{(*)}{\geq} v_{\text{inf}} \int_{\Theta} (\phi - x^*)^2(\theta) \pi(d\theta) = v_{\text{inf}} \|\phi - x^*\|_\pi^2, \end{aligned}$$

where we have used assumption **H_{unif,2}** in (*). So **H_{coerc}** is satisfied. □

The assumption **H_{m_k, M_k}** will depend on the specific choices of π and μ .

In practice, we choose the sequences under the form: $\gamma_k \sim k^{-a}$, $m_k \sim k^b$, $M_k \sim k^p$,

In general, we can find δ, Δ s.t. $q_m \sim m^{-\delta}$, $Q_m \sim m^\Delta$. The assumption **H_{m_k, M_k}** gives conditions on a, b and p :

$$0 < a \leq 1, \quad 2 - \delta b < 2a, \quad b\Delta + 1 < 2a + p \quad (4.3.3)$$

Example:

- When dealing with uniform $\theta \sim \mathcal{U}[-1, 1]$, the adapted polynomials are the Legendre polynomials $\{P_i\}_{i \in \mathbb{N}}$ [CFG20] (then $\int_{[-1, 1]} P_i P_j d\pi = \delta_{i, j} \frac{2}{2i+1}$). The orthonormal Legendre polynomials $\{B_i\}_{i \in \mathbb{N}}$ are defined as $B_i = \sqrt{\frac{2i+1}{2}} P_i$, where $|P_i| \leq 1$ so $Q_m = \sup_{\theta \in \Theta} \sum_{i \leq m} |B_i(\theta)|^2 = \mathcal{O}(m^2)$: $\Delta = 2$.
- δ depends of the same order than the regularity of $x^*(\theta)$.
- Constraints on a, b, p with very regular $z^*(\theta)$ and Legendre basis:

$$0 < a \leq 1, \quad 2b + 1 < 2a + p.$$

4.3.4 Applications: log-normal returns and specific weights

In this subsection, we are going to detail a use case of UQSA for the Sharpe ratio. We place ourselves in the case when the assets values follow a Black Scholes model, and are correlated with a constant correlation ρ .

The sub-section is organized as follows:

1. First, we recall the optimal (or tangent) portfolio formula. This is the fixed portfolio we will consider in our experiments. Under a condition on the returns first moments (" $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} > 0$ "), we can ensure that the Sharpe ratio is positive in the tangent portfolio.
2. Secondly, we remind the formulation of the returns moments under Black Scholes model. This allows us to have an explicit formulation of the Sharpe ratio, which will help us to compute the error metrics in our experiments.
3. Then we give the results of our experiments: assuming the correlation uncertain, we fix the allocation as the tangent portfolio for a specific correlation level and see the evolution of the uncertain Sharpe ratio with this fixed allocation.

Setting: We are interested in the value of a portfolio starting at value 1 after one period. We consider a universe of d assets, whose price S_1^i at period 1 follows a Black Scholes model with **unique** uncertain correlation $\rho = f(\theta)$, $\theta \sim \pi$, $f : \Theta \rightarrow [-1, 1]$. The portfolio value with given allocation w_i and asset value S_1^i is:

$$V_1 = \sum_{i=1}^d w_i S_1^i, \quad S_1^i \stackrel{(d)}{=} e^{\mu_i + \frac{\sigma_i^2}{2} - \sigma_i W_1^i}, \quad (4.3.4)$$

where $W_1 \sim \mathcal{N}(0, \Gamma)$, $\Gamma_{i, j} = \begin{cases} \rho & \text{if } i \neq j, \\ 1 & \text{else.} \end{cases}$

The Sharpe ratio is then computed on the period 1 return $\frac{V_1 - V_0}{V_0} = V_1 - 1$, playing the role of R_θ in the previous section.

4.3.4.1 Tangent (or maximal Sharpe ratio) portfolio

In this subsection, we recall the formula for the tangent portfolio and we give the condition under which it is associated to a positive Sharpe ratio.

Definition 4.3.2 (Tangent portfolio). Let a vector $\boldsymbol{\mu} \in \mathbb{R}^d$ (mean vector) and a positive definite matrix $C \in \mathbb{R}^{d \times d}$ (covariance matrix) such that $\text{sign}(\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}) > 0$ be given. The tangent portfolio is defined by

$$\mathbf{w}^{tgt}(\boldsymbol{\mu}, C) := \arg \max_{\mathbf{w} \in \mathbb{R}^d, \mathbf{w}^\top \mathbf{1}_d = 1} \frac{\mathbf{w}^\top \boldsymbol{\mu}}{\sqrt{\mathbf{w}^\top C \mathbf{w}}} = \frac{C^{-1} \boldsymbol{\mu}}{\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}}.$$

The tangent Sharpe ratio is then:

$$\text{SR}^{tgt} := \max_{\mathbf{w}, \mathbf{w}^\top \mathbf{1}_d = 1} \frac{\mathbf{w}^\top \boldsymbol{\mu}}{\sqrt{\mathbf{w}^\top C \mathbf{w}}} = \sqrt{\boldsymbol{\mu}^\top C^{-1} \boldsymbol{\mu}}.$$

A derivation of and discussion around the tangent portfolio is given in [Cha11] for example. The tangent portfolio is named after the fact that it corresponds to the intersection between the tangent line above the efficient frontier (EF) passing through the origin. The minimum variance portfolio corresponds to the left-most extremity of the efficient frontier. Its return has the same sign as $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}$. When this return is negative, there is no "positive" tangent line to the EF.

This portfolio corresponds to the portfolio with maximum ratio mean/volatility when $\text{sign}(\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}) > 0$. The condition $\text{sign}(\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}) > 0$ is necessary for the return of the tangent portfolio to be positive. In what follows, so that the approach is meaningful (to fit our $\text{sign}(\mu_\theta) > 0$ assumption and by economic rationale), we will assume that $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} > 0$.

Remark 7. What we see from the tangent portfolio formula is that the allocation will take more extreme position as $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}$ tends to zero. We explicit this phenomenon in dimension 2.

For $C = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sqrt{1-\rho^2}\sigma_2^2 \end{pmatrix}$, then

$$C^{-1} = \frac{1}{(1-\rho^2)\sigma_1^2\sigma_2^2} \begin{pmatrix} \sigma_2^2 & -\rho\sigma_1\sigma_2 \\ -\rho\sigma_1\sigma_2 & \sigma_1^2 \end{pmatrix} \quad (4.3.5)$$

and

$$\mathbf{w}^{tgt} = \frac{C^{-1} \boldsymbol{\mu}}{\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}} = \frac{1}{\mu_1\sigma_2^2 + \mu_2\sigma_1^2 - \rho(\mu_1 + \mu_2)\sigma_1\sigma_2} \begin{pmatrix} \mu_1\sigma_2^2 - \rho\mu_2\sigma_1\sigma_2 \\ \mu_2\sigma_1^2 - \rho\mu_1\sigma_1\sigma_2 \end{pmatrix}. \quad (4.3.6)$$

From C^{-1} definition (4.3.5), $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} = \frac{1}{(1-\rho^2)\sigma_1^2\sigma_2^2} (\mu_1\sigma_2^2 + \mu_2\sigma_1^2 - \rho(\mu_1 + \mu_2)\sigma_1\sigma_2)$.

Let ρ_{lim} value of ρ for which $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} = 0$.

$$\begin{aligned} \mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} = 0 &\Leftrightarrow \frac{1}{(1-\rho_{\text{lim}}^2)\sigma_1^2\sigma_2^2} (\mu_1\sigma_2^2 + \mu_2\sigma_1^2 - \rho_{\text{lim}}(\mu_1 + \mu_2)\sigma_1\sigma_2) = 0 \\ &\Leftrightarrow \mu_1\sigma_2^2 + \mu_2\sigma_1^2 - \rho_{\text{lim}}(\mu_1 + \mu_2)\sigma_1\sigma_2 = 0, \quad \rho_{\text{lim}} \neq \pm 1 \\ &\Leftrightarrow \rho_{\text{lim}} = \frac{\mu_1\sigma_2^2 + \mu_2\sigma_1^2}{(\mu_1 + \mu_2)\sigma_1\sigma_2}, \quad \rho_{\text{lim}} \neq \pm 1. \end{aligned} \quad (4.3.7)$$

If $\rho_{\text{lim}} \in (-1, 1)$, at $\rho = \rho_{\text{lim}}$, the denominator in (4.3.6) goes to zero so the tangent portfolio takes "infinite" values. This is why in practice, constraints on allocation are added to prevent from these extreme positions. From (4.3.7), we see that if $c(\rho) := \mu_1\sigma_2^2 + \mu_2\sigma_1^2 - \rho(\mu_1 + \mu_2)\sigma_1\sigma_2$ is different from 0 for any $\rho \in (-1, 1)$, then $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu}$ is different from 0. In particular, since $c(\rho)$ is a decreasing function of ρ , if $\lim_{\rho \rightarrow 1} c(\rho) = c(1) > 0$, then $\mathbf{1}_d^\top C^{-1} \boldsymbol{\mu} \neq 0$ for any $\rho \in (-1, 1)$.

4.3.4.2 Benchmark: explicit Sharpe ratio under Black Scholes model

In this subsection, we make explicit R two first moments in the Black Scholes model and derive the explicit Sharpe ratio. This will serve as a benchmark in our experiments.

Lemma 4.3.3 (Sharpe ratio formula - one period Black Scholes model). *Let us denote $\bar{\boldsymbol{\mu}} \in \mathbb{R}^d$ and $\bar{C} \in \mathbb{R}^{d \times d}$ the vector and matrices defined by*

$$\bar{\mu}_i = (e^{\mu_i} - 1), \quad \bar{C}_{i,j} = \begin{cases} e^{\mu_i + \mu_j} (e^{\rho\sigma_i\sigma_j} - 1) & \text{if } i \neq j, \quad i, j \in [d], \\ e^{2\mu_i} (e^{\sigma_i^2} - 1) & \text{else.} \end{cases} \quad (4.3.8)$$

Then $\mathbb{E}[R] = \mathbf{w}^\top \bar{\boldsymbol{\mu}}$ and $\text{Var}[R] = \mathbf{w}^\top \bar{C} \mathbf{w}$, and the Sharpe ratio is given by $SR(\rho) = \frac{\mathbf{w}^\top \bar{\boldsymbol{\mu}}}{\sqrt{\mathbf{w}^\top \bar{C} \mathbf{w}}}$.

Proof. Let us denote $\Gamma \in \mathbb{R}^{d \times d}$, $\Gamma_{i,j} = \begin{cases} \rho & \text{if } i \neq j, \\ 1 & \text{else,} \end{cases}$ the correlation matrix. In the Black Scholes model (4.3.4), for $W_1 \sim \mathcal{N}(0, \Gamma)$, we have:

$$V_1 = \sum_{i=1}^d w_i S_1^i = \sum_{i=1}^d w_i e^{\mu_i - \frac{\sigma_i^2}{2} + \sigma_i W_1^i}, \quad V_0 = 1, \quad R \stackrel{(d)}{=} \frac{V_1 - V_0}{V_0} = V_1 - 1.$$

Using the generating moment of the Gaussian distribution ($\mathbb{E}[e^{\lambda G}] = e^{\frac{\lambda^2}{2}}$, $\forall \lambda \in \mathbb{R}$, $G \sim \mathcal{N}(0, 1)$),

$$\mathbb{E}[V_1] = \mathbb{E}\left[\sum_{i=1}^d w_i e^{\mu_i - \frac{\sigma_i^2}{2} + \sigma_i W_1^i}\right] = \sum_{i=1}^d w_i e^{\mu_i}.$$

To compute the variance, we first compute $\mathbb{E}[V_1^2]$:

$$\mathbb{E}[V_1^2] = \sum_{i,j=1}^d w_i w_j e^{\mu_i + \mu_j - \frac{\sigma_i^2 + \sigma_j^2}{2}} \mathbb{E}\left[e^{\sigma_i W_1^i + \sigma_j W_1^j}\right],$$

where (W_1^i, W_1^j) is a Gaussian vector of covariance $\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ if $i \neq j$, hence, using the generating moment of the Gaussian vector $\mathbb{E}[e^{\lambda^\top X}] = e^{\frac{\lambda^\top C \lambda}{2}}$, $\forall \lambda \in \mathbb{R}^2$, $X \sim \mathcal{N}(0, C)$, $C \in \mathbb{R}^{2 \times 2}$,

$$\mathbb{E}\left[e^{\sigma_i W_1^i + \sigma_j W_1^j}\right] = \begin{cases} e^{\frac{\sigma_i^2 + \sigma_j^2 + 2\rho\sigma_i\sigma_j}{2}} & \text{if } i \neq j, \\ e^{2\sigma_i^2} & \text{else.} \end{cases}$$

After simplification:

$$\begin{aligned} \mathbb{E}[V_1^2] &= \sum_{\substack{i,j \in [d] \\ i \neq j}} w_i w_j e^{\mu_i + \mu_j + \rho\sigma_i\sigma_j} + \sum_{i \in [d]} w_i^2 e^{2\mu_i + \sigma_i^2}, \\ \text{Var}[V_1] &= \sum_{\substack{i,j \in [d] \\ i \neq j}} w_i w_j e^{\mu_i + \mu_j} (e^{\rho\sigma_i\sigma_j} - 1) + \sum_{i \in [d]} w_i^2 e^{2\mu_i} (e^{\sigma_i^2} - 1). \end{aligned}$$

Using that $R \stackrel{(d)}{=} \frac{V_1 - V_0}{V_0} = V_1 - 1$, we retrieve $\mathbb{E}[R] = \mathbf{w}^\top \bar{\boldsymbol{\mu}}$ and $\text{Var}[R] = \mathbf{w}^\top \bar{C} \mathbf{w}$ with the $\bar{\boldsymbol{\mu}}$ and \bar{C} defined in (4.3.8). □

4.3.5 Experiments

To apply the UQSA for Sharpe ratio technique, we need:

- $\frac{V_1^\theta - V_0}{V_0}$ to have finite 2 and 4 order moments,
- finite second order moment in θ .

Since V_1^θ is a sum of geometric Brownian motions, with $\rho = f(\theta)$ bounded, this is verified.

In this experiments section, we take $d = 2$, $S_0^1 = S_0^2 = 1$, $\mu_1 = 3.5\%$, $\mu_2 = 5\%$, $\sigma_1 = 8\%$ and $\sigma_2 = 10\%$. Let $\bar{\boldsymbol{\mu}}$ and \bar{C} as defined in (4.3.8). By (4.3.7), $\mathbf{1}^\top \bar{C}^{-1} \bar{\boldsymbol{\mu}}$ is a decreasing function of ρ . For $\rho = 1$, $\mathbf{1}^\top \bar{C}^{-1} \bar{\boldsymbol{\mu}} \approx 6.8 > 0$ so for any $\rho \in (-1, 1)$, $\mathbf{1}^\top \bar{C}^{-1} \bar{\boldsymbol{\mu}} > 0$ and the tangent portfolio is the maximum Sharpe ratio portfolio for any $\rho \in (-1, 1)$.

Sharpe ratio with fixed allocation: We place in the case when the investor considers a high correlation ($\rho = 90\%$) and allocates his portfolio along the associated tangent portfolio. The allocation can then take extreme position. In our setting, $\mathbf{w}(90\%) \approx (-0.8, 1.8)$. The UQ for Sharpe ratio allows to quantify how much the Sharpe ratio will move if the assumed correlation $\rho = 90\%$ is mis-specified.

We take the following approach: the investor has a good confidence on his assumed correlation $\rho = 90\%$ but is still interested in the evolution of the Sharpe ratio in a neighborhood $[\rho_{\min}, \rho_{\max}]$, $\rho_{\min} \leq \rho \leq \rho_{\max}$ close to ρ . In our following experiments, we have taken: $\rho_{\min} = 80\%$, $\rho_{\max} = 99\%$.

4.3.5.1 Choice of uncertainty function

In our setting, the Sharpe ratio is a rather quadratic, smooth function of ρ , as shown in Figure 4.1: taking a uniform uncertainty $\theta = \rho$ would be relatively simple.




Figure 4.1 – Plot of $SR(\rho)$ as a function of the correlation ρ .

An empirical based approach for estimating the uncertain function: We want to find a sensible function f such that $\rho \stackrel{(d)}{=} f(\theta)$, $\theta \sim \mathcal{U}[-1, 1]$.

If ρ follows a certain distribution of cumulative distribution function F , then ρ has the same distribution than $q\left(\frac{1+\theta}{2}\right)$ where $q := F^{-1}$ is the quantile function and $\frac{1+\theta}{2} \sim \mathcal{U}[0, 1]$.

To get an idea of what would be a good quantile function for ρ , an empirical study was conducted on financial data (BNPIMAD time series, $d = 9$ assets, $n = 19$ years of data). The correlation between each pair of assets was estimated on $T = 5$ years of data, with a sliding window $W = 6$ months (first estimation at time 0 on the first 5 years of data, then estimation at time 6 months on 5 years of data, ... until the last estimation on the last 5 years of data, hence $n_W = \frac{n-T}{W} = 29$ estimations of each $\rho_{i,j}$, $i, j \in \{1, \dots, d\}$). We plot in Figure 4.2 the empirical quantile functions, defined on $\{\rho^{(1)} \leq \dots \leq \rho^{(n_W)}\}$ as:

$$q(x) = \frac{1}{n_W} \sum_{k=1}^{n_W} \rho^{(k)} \mathbf{1}_{\frac{k-1}{n_W} \leq x < \frac{k}{n_W}}, \quad x \in [0, 1].$$

The quantile function plot is equivalent to a Q-Q plot of the correlation distribution against the uniform distribution on $[0, 1]$.

Some empirical quantile plots are given in Figure 4.2. The shape of the graph can vary but there are three main shapes:

1. almost linear shape (meaning that all the correlation range is equiprobable),
2. hill shape (multiple values are possible but clear plateau on the most likely values),

3. succession of slightly increasing, almost plateau portions of graph, followed by strong growth shape: there are multiple likely values and unlikely range of values.

We choose to focus on this third, more challenging shape. In order to reproduce this shape, we rely on convex combination of incomplete Beta functions:

$$F_{\alpha,\beta}(x) = \frac{1}{B(\alpha,\beta)} \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt, \quad x \in [0, 1],$$

$$F_{\alpha,\beta}^{l,s}(x) = F_{\alpha,\beta}\left(\frac{x-l}{s}\right) \quad x \in [l, l+s], \text{ (rescaled incomplete Beta)}$$
(4.3.9)

where B denotes the Beta function. These functions have the property to take their values between 0 and 1 and could be interpreted as cumulative distribution functions. This class of functions was chosen because they allow for great flexibility of shape, as illustrated in Figure 4.3.

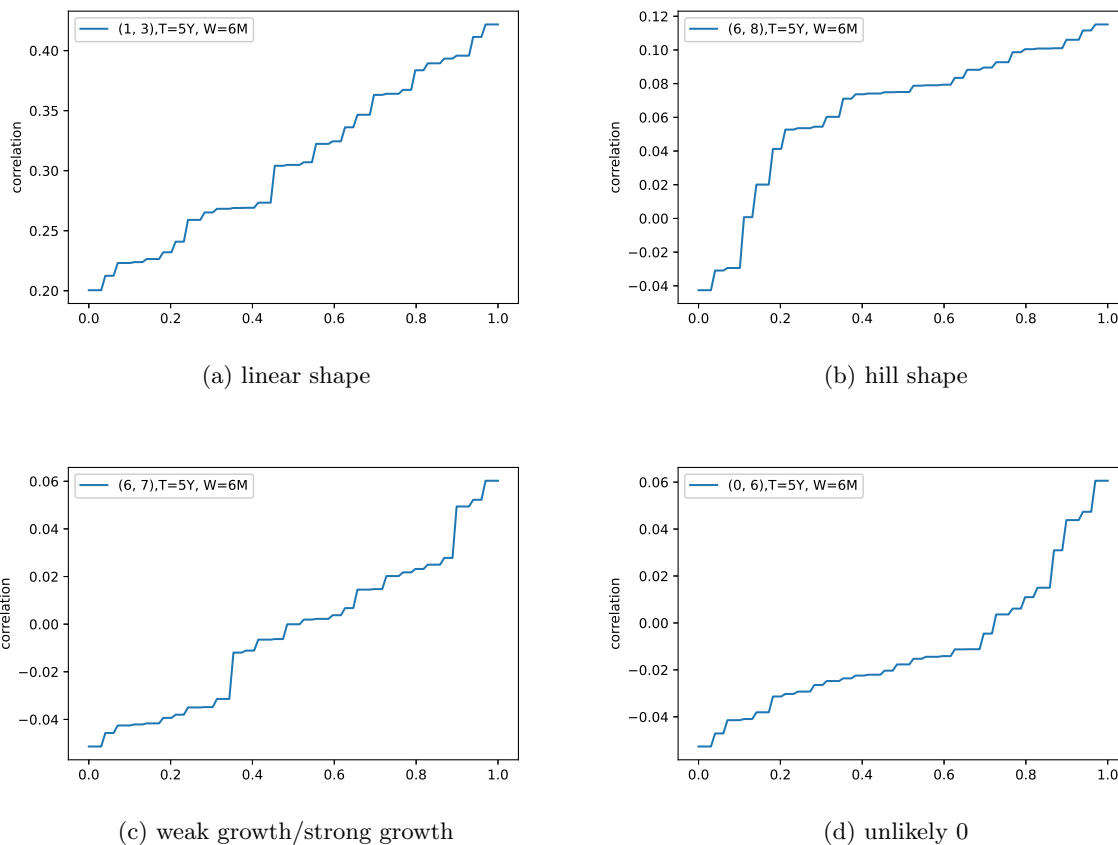


Figure 4.2 – Empirical quantile function for 4 pairs of correlation estimated on $T = 5$ years of data, with a sliding window $W = 6$ months.




Figure 4.3 – Plot of Beta incomplete function, for $(\alpha, \beta) = (5, 2)$ (a) and $(\alpha, \beta) = (2, 5)$ (b). We see that for large α (respectively β), the function is flat near zero (resp. 1).

Our uncertainty function takes the form of a convex combination of (possibly rescaled) Beta incomplete functions $\{F_{\alpha_i, \beta_i}^{l_i, s_i}\}_{i=1}^{n_B}$ such that:

$$\rho \stackrel{(d)}{\approx} f_{\beta}(\theta) = \rho_{\min} + (\rho_{\max} - \rho_{\min}) \sum_{i=1}^{n_B} \lambda_i F_{\alpha_i, \beta_i}^{l_i, s_i} \left(\frac{1 + \theta}{2} \right),$$

$$\sum_{i=1}^{n_B} \lambda_i = 1, \quad \theta \sim \mathcal{U}[-1, 1].$$

In Figure 4.4, we present a convex combination of incomplete Beta function that fits well the Figure 4.2 weak growth/strong growth shape.




Figure 4.4 – Plot of convex combination of Beta incomplete functions with parameters indicated in the plot legend.

4.3.5.2 Evolution and approximation of the uncertain Sharpe ratio

In this subsection, we illustrate how the Sharpe ratio evolves with the chosen uncertainty function, and how the truncation error decreases with the number of coefficients.

Sharpe ratio approximation and truncation error convergence rate - Beta combination uncertainty: In this paragraph, we discuss the convergence rate of the truncation error.

For the Beta combination uncertainty presented in Figure 4.4, the Sharpe ratio is a rather smooth function of θ , as seen from the Figure 4.5.




Figure 4.5 – Plot of $\theta \mapsto f_\beta(\theta)$ (a) and $\theta \mapsto SR(f_\beta(\theta))$ (b).

In Figure 4.6 (a), we show how fast we can approximate the function $SR^2(f_\beta(\theta))$ with an increasing number of coefficients in the Legendre basis. The coefficients used for the approximation are computed via numerical integration: $u_i^* = \int_{[-1,1]} SR^2(f_\beta(\theta))B_i(\theta)\pi(d\theta)$. The obtained coefficients are shown in the plot of Figure 4.6 (b). We see that we need at least 20 coefficients to properly approximate the function.

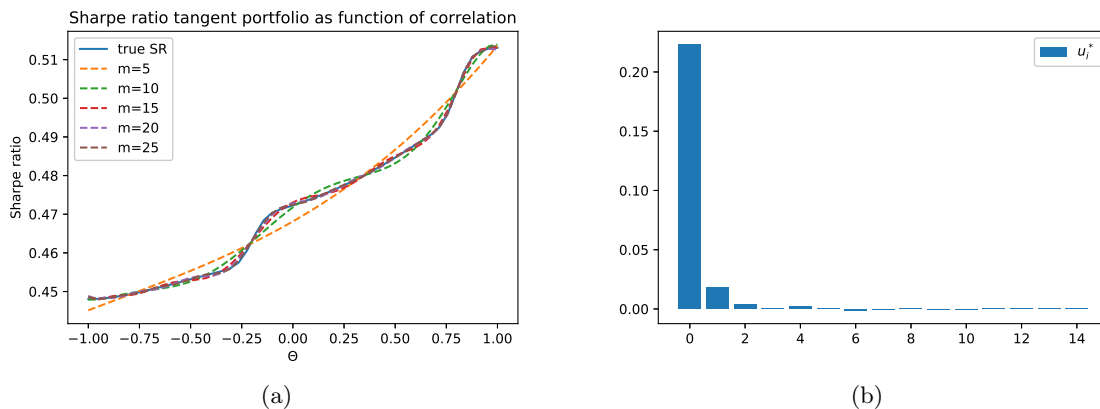


Figure 4.6 – (a): Evolution of $SR(f(\theta))$ approximation for different values of m in $\{5, 10, 15, 20, 25\}$ (with coefficients computed via integration). (b): Evolution of the coefficients u_i^* , $i \in \{0, \dots, 15\}$.




Figure 4.7 – (a): Truncation error $\mathcal{E}_{TR} := (\sum_{i>m} (u_i^*)^2)^{1/2} = \|\phi^* - \sum_{i \leq m} u_i^* B_i\|_\pi$ computed via integration as a function of the number of coefficients m in log-log scale, and the associated tendency line. (b): a and b frontier, for the value of δ obtained via truncation error convergence rate estimation.

In Figure 4.7 (a), we display the truncation error \mathcal{E}_{TR} for increasing m_k as a function of the number of iterations k . The regression slope of $\log(\mathcal{E}_{TR})$ on $\log k$ is equal to -1.8 , which corresponds to a relatively slow convergence speed.

Remark 8. This rate of convergence was to be expected given the regularity of the function f_β . For projection on Legendre polynomials, the convergence rate of the truncation error of a function r -differentiable is of order $\mathcal{O}(m^r)$. This is shown for example in [WX12]. Here, the incomplete Beta functions (4.3.9) are defined as primitive of polynomial functions in $t^{\alpha-1}$, $t^{\beta-1}$, defined on $[0, 1]$. Hence, they are continuously differentiable until the order $\min(\alpha-1, \beta-1)$ and piecewise differentiable at the order $\min(\alpha, \beta)$. For our choice of function shown in Figure 4.4, the minimum α/β parameter is equal to 2, hence we could expect a convergence rate between 1 and 2.

Frontier estimation: This rate gives the δ parameter needed in equations (4.3.3) to derive the limiting values for a , b and p . Fixing p equals to 0 (constant inner Monte Carlo number of simulations), we display the corresponding a, b frontiers in Figure 4.7 (b) for $\delta = 1.5$ (absolute regression slope) and $\Delta = 2$ (Legendre basis increase rate). The frontiers correspond to the values of b between the lines $x \mapsto \frac{2a-p-1}{\Delta} = \frac{2a-1}{2}$ and $x \mapsto \frac{2(1-a)}{\delta} = \frac{4(1-a)}{3}$.

Numerical precision threshold: On the Figure 4.7 (a), we display the truncation error until the 32th coefficient approximation. We can see that the truncation error starts exploding from the 28/30-th coefficient. This is due to the fact that the estimation error is important for coefficients of high order. Indeed, from component 20 onward, the value of the coefficient falls below 10^{-8} which is the numerical precision associated to the integration method we used (scipy nquad function). In our experiments, we will rely only the first 20 estimated coefficients.

4.3.5.3 Impact of hyperparameters

In the numerical experiments that follow, we compare the performance of the algorithms with different choices of m_k , M_k and γ_k sequences. The comparison relies on root-mean-square errors, approximated with 15 independent runs of the algorithms.

When running the UQSA algorithm, we take $M_k = M = 100$ simulations at each iteration.

Impact of increasing dimension: In this paragraph, we discuss the role of the sequence $\{m_k, k \in \mathbb{N}\}$ and we illustrate the convergence rate of the total error, when the sequence $\{m_k, k \in \mathbb{N}\}$ is increasing.

In Figure 4.8, we plot the evolution of the total error for different possibilities of fixed dimensions $m_k = m \in \{10, 30\}$ versus the evolution with slowly increasing dimension $m_k = \lfloor k + 1 \rfloor^b$, $b = 0.3$. For $b = 0.3$ and $K = 10^4$, we obtain $m_K \approx K^b \approx 16$ coefficients. Here, the step size is taken equal to 0.95.




Figure 4.8 – In the case of increasing m_k (solid blue line) and $m_k = m \in \{10, 30\}$ (dotted line) the total error $\|\hat{\phi}^{(k)} - \phi^*\|_\pi$ as a function of the number of iterations k .

We observe that the total error is much larger with fixed $m_k = 10$ or 30 in the burn-in phase. The total error is greater the larger the fixed m . By trying to fit simultaneously all the m coefficients, the performance is worse than with the sequential, increasing dimension approach.

The first iterations are associated to larger values of the step size γ_k . With the increasing dimension approach, only the first coefficients, which are also the largest, are estimated, and can benefit from the large values of the step size.

Additionally to the total error, we plot in the following figures the truncation error levels as a function of the iteration

$$\mathcal{E}_{TR} = \left\| \phi^* - \sum_{i=1}^{m_k} u_i^* B_i \right\|_\pi,$$

previously computed and graphed (as a function of m_k) in Figure 4.7. For fixed $m_k = m$, it is then a constant. We plot also the case $m = 3$ and $b = 0.15$ (m_k slowly increasing). For the fixed m approach, the total error must stay above this fixed level of truncation error, while for the increasing m_k approach, the truncation error will keep decreasing.




(a) $m \in \{3, 10, 30\}$ (b) $b = 0.3$ (c) $b = 0.15$

Figure 4.9 – In the case of increasing m_k (solid blue line) and $m_k = m \in \{10, 30\}$ (dotted line) the total error $\|\hat{\phi}^{(k)} - \phi^*\|_\pi$ and truncation error, fixed $m_k = m \in \{3, 10, 30\}$ (up) or increasing $m_k = \lfloor k + 1 \rfloor^b$, $b \in \{0.15, 0.3\}$ (down) as a function of the number of iterations k .

We can make the following comments:

- For fixed m : the error is further from the truncation error the larger the fixed m .
- For increasing m_k
 - In the smaller $b = 0.15$ case, the total error is closer to its truncation limit than in the case with faster m_k growth.
 - Each time a new coefficient starts being approximated by the algorithm, the total error makes a small jump, and the statistical error (difference between the total error and the truncation error) makes a big jump. For example, on Figure 4.9[(c)], we can see such a jump at iteration 100: the total error increases then decreases again.
 - The total error has more variance (we filled the confidence intervals with the same color and more transparency in Figure 4.8 and 4.9[(b) and (c)]) but converges better. This is in line with the stability theorem of [CFG20][Theorem 1] since we are in a case when the assumptions of the theorem are satisfied.

In the following Figures 4.11 and ??, we display how the Sharpe ratio can be approximated via both methods, at iteration $k \in \{1000, 5000, 10\,000\}$, for small b and fixed m in Figure 4.11, larger b and m in Figure ??.

- In our experiments, we chose to consider a polynomial basis of functions to approximate the function $\text{SR}^2(f_\beta(\theta))$, hence the approximation are under the form of linear combinations




(a) $k = 1000$ (b) $k = 5000$ (c) $k = 10\ 000$

Figure 4.10 – Approximation of the Sharpe ratio - mixture uncertainty, increasing number of iterations, $a = 0.95$, $b = 0.15$, $m_k = k^b$ versus $m_k = 3$.




(a) $k = 1000$ (b) $k = 5000$ (c) $k = 10\ 000$

Figure 4.11 – Approximation of the Sharpe ratio - mixture uncertainty, increasing number of iterations, $a = 0.95$, $b = 0.3$, $m_k = k^b$ versus $m_k = 3$.

of increasing order polynomials. For the approximation with small b and m , since the approximation is a linear combination of polynomials of order smaller or equal to 3, it has a rather smooth shape but does not capture the changes in convexity of the true function.

- Conversely, for large m and b , the approximation is of the form of a linear combination of polynomials with much higher order. For $b = 0.15$, $m_k = 7$ for $k = 1000$. This explains the oscillations in the representation of the approximations.
- The oscillation are far less present in the increasing m_k approximation of the function (even if the number of coefficients is equal to 7, 12 then 15 for the final iteration). This reflects the interest of considering a progressive estimation of the coefficients, and shows that in our specific example the convergence works even in infinite norm, and not only in $L_2(\pi)$ norm as guaranteed by the stability theorem.




Figure 4.12 – The total error \mathcal{E} as a function of the number of iterations, for different values of a in $\{0.85, 0.9, 0.95\}$.

Impact of the step size parameter a : In this experiment, we set $b = 0.3$ and $p = 0$. The admissible values of a are in the range $(0.775, 1)$. We have tested $a \in \{0.85, 0.9, 0.95\}$. We display in Figure 4.12 the total error as a function of iterations k for different values of a . We can see that the convergence improves with increasing a . This might be due to the fact that too low value for a leads to more instable estimation, and since the function is rather smooth, high values of a are preferable.

4.3.6 Conclusion to UQ for the Sharpe ratio section

In this section, we proposed a SA formulation for the Sharpe ratio, in order to apply the UQSA algorithm to quantify parameter uncertainty on the Sharpe ratio. For the sake of clarity, while using a realistic example, we have settled on an allocation of type tangent portfolio and portfolio values following a Black Scholes model with uncertain correlation. A realistic example of a distribution on correlations was sought. The impact of the choice of UQSA hyperparameters was explored and the efficiency of the sequential algorithm considered in its increasing dimension approach has been highlighted.

In the next section, we will focus on the impact of uncertainty on portfolio allocation. For our future experiments, we will only consider examples with increasing dimension approach.

4.4 UQSA for the portfolio allocation

This section is devoted to our applications of the UQSA algorithm to portfolio allocation.

4.4.1 Portfolio optimization as SA problem

In this section, we consider portfolio optimization problem of type $\arg \min_{\mathbf{w} \in \mathcal{W}} U(\mathbf{w})$, where \mathcal{W} denotes the convex set of constraints. We will deal with U convex, which has a unique solution \mathbf{w}^* . For example, for the minimum-variance problem, given a covariance matrix $C \in \mathbb{R}^{d \times d}$,

$$U^{MV}(\mathbf{w}) := \frac{1}{2} \mathbf{w}^\top C \mathbf{w}. \quad (4.4.1)$$

We are interested in the following set of constraints:

$$\begin{aligned} \mathcal{W}_1 &= \left\{ \mathbf{w} \in \mathbb{R}^d : \mathbf{1}_d^\top \mathbf{w} = 1 \right\} && \text{(budget constraint),} \\ \mathcal{W}_1^{\geq 0} &= \left\{ \mathbf{w} \in (\mathbb{R}_+)^d : \mathbf{1}_d^\top \mathbf{w} = 1 \right\} && \text{(budget and positivity constraints).} \end{aligned} \quad (4.4.2)$$

We say that the solution $\mathbf{w}^* := \arg \min_{\mathbf{w} \in \mathcal{W}_1^{\geq 0}} U(\mathbf{w})$ is *unsaturated* if it does not saturate the positivity constraint, e.g. for any $i \in \{1, \dots, d\}$, $w_i^* > 0$. In this case, the solution on $\mathcal{W}_1^{\geq 0}$ and \mathcal{W}_1 are the same. Noticing that the budget constraint amounts to take $w_d = 1 - \sum_{i=1}^{d-1} w_i$, \mathbf{w}^* is then the zero of the gradient of U in (w_1, \dots, w_{d-1}) . In our proposed approach, we are going to focus on cases when \mathbf{w}^* is unsaturated.

Reparametrization of the problem. We would like to apply the UQSA algorithm on the problem solution $\mathbf{w}^*(\theta)$ subject to small perturbations modeled through θ . It is not possible to consider directly a UQSA approach on $\mathbf{w}^*(\theta)$ because when computing iteratively the PCE expansion of $\mathbf{w}^*(\theta)$, there is no guarantee that the expansion stays in \mathcal{W} . This is why we propose to consider a parameterization $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}_1^{\geq 0}$ such that for any $\mathbf{w} \in \mathcal{W}$, there is a $\mathbf{y} \in \mathbb{R}^{d-1}$ such that $\mathbf{w} = \ell(\mathbf{y})$, and do the UQSA approach on the unconstrained vector \mathbf{y} . The problem then becomes:

$$\mathbf{y}^* \in \mathbb{R}^{d-1} \text{ solution of } \mathbb{E}[H(\mathbf{y}^*, R)] = \mathbf{0}_{d-1}, \quad \text{where } \mathbb{E}[H(\mathbf{y}^*, R)] = \nabla U(\ell(\mathbf{y}^*)).$$

R denotes the returns random variables whose distribution depends on an uncertain parameter θ . The approach we propose in this section relies on the following assumption: $\mathbf{w}^*(\theta)$ is unsaturated for any $\theta \in \Theta$.


Many mappings have already been designed in the literature. The field focusing on the representation of vectors in the simplex is referred to as Compositional Data Analysis (CDA) [PGB11]. For example, the log-ratio transformation offers a one-to-one mapping from \mathbb{R}^{d-1} to $\mathcal{W}_1^{\geq 0} = \mathcal{W} \cap (\mathbb{R}_+^*)^d$, by taking $w_i = \frac{e^{y_i}}{1 + \sum_{j=1}^{d-1} e^{y_j}}$, for $i = 1, \dots, d-1$, $w_d = \frac{1}{1 + \sum_{j=1}^{d-1} e^{y_j}}$.

We propose a new transformation allowing for 0 component w_i , that we call "integral-transformation" and relying in the following idea: given $p : \mathbb{R} \rightarrow \mathbb{R}_+$ such that $p(0) = 0$ and $f : \mathbb{R} \rightarrow \mathbb{R}_+$ such that $\int_{\mathbb{R}} f = 1$, then one can interpret the w_i as portions of the integral of f . There exists $\mathbf{y} \in \mathbb{R}^{d-1}$ such that:

$$w_1 = \int_{-\infty}^{y_1} f(t) dt,$$

$$w_i = \int_{y_1 + p(y_2) + \dots + p(y_{i-1})}^{y_1 + p(y_2) + \dots + p(y_i)} f(t) dt,$$

$$w_d = \int_{y_1 + p(y_2) + \dots + p(y_d)}^{\infty} f(t) dt.$$



In the following subsection, we give a condition on ℓ to guarantee that \mathbf{y}^* can be defined as a zero gradient problem and correspond to \mathbf{w}^* up to transformation and we give specific integral transformation examples for which this condition is satisfied. Finally, we show illustration of our integral-based approach and we advocate for the use of polynomial-based transformation rather than sigmoid-based transformations.

4.4.2 Allocation model

Before expanding on possible transformation functions ℓ , we are going to give the condition that our transformation function must satisfy in order to preserve the solution existence and

convexity of the problem.

4.4.2.1 Convexity preservation

In this subsection, we want to show under which condition a convex objective function remains convex after transformation. We say that ℓ is a transformation function if for any $\mathbf{w} \in \mathcal{W}_1^{\geq 0}$, there exists a $\mathbf{y} \in \mathbb{R}^{d-1}$ such that $\mathbf{y} = \ell(\mathbf{w})$.

We start by giving the gradient and Hessian of the transformed problem.

Lemma 4.4.1 (Explicit gradient and Hessian). *Let $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ the transformation function and $U : \mathcal{W} \rightarrow \mathbb{R}_+$ the objective function.*

The gradient of $U(\ell(\mathbf{y}))$ with regards to \mathbf{y} is given by

$$\nabla_{\mathbf{y}} U(\ell(\mathbf{y})) = D(\ell) \overline{\nabla_x U},$$

where $D(\ell)$ is the ℓ Jacobian and $\overline{\nabla_x U}$ is the "lagged" original gradient:

$$D(\ell) = \left\{ \frac{\partial \ell_j(\mathbf{y})}{\partial y_i} \right\}_{\substack{1 \leq j \leq d-1, \\ 1 \leq i \leq d-1}}, \quad \overline{\nabla_x U} = \begin{pmatrix} \partial_{x_1} U - \partial_{x_d} U \\ \vdots \\ \partial_{x_{d-1}} U - \partial_{x_d} U \end{pmatrix}.$$

The Hessian of $U(\ell(\mathbf{y}))$ with regards to \mathbf{y} is given by

$$\mathcal{H}(y) = \{\partial_{y_i y_j} U\}_{i,j=1}^{d-1} = D^2(\ell) \overline{\nabla_x U} + D(\ell) \overline{\mathcal{H}_U} D(\ell)^\top,$$

where $D^2(\ell) = \{\partial_{y_i y_j} \ell\}_{i,j=1}^{d-1}$ and $\overline{\mathcal{H}_U} = \{\partial_{x_i x_j} U - \partial_{x_i x_d} U - \partial_{x_d x_j} U + \partial_{x_d x_d} U\}_{j,j=1}^{d-1}$ is the "lagged" original Hessian.

Proof. By Chain rule, for $i \in \{1, \dots, d\}$, using the fact that the last component of ℓ is equal to $1 - \sum_{i=1}^{d-1} \ell_i$, $\partial_{y_i} U = \sum_{j=1}^{d-1} \frac{\partial \ell_j(\mathbf{y})}{\partial y_i} \partial_{x_j} U - \sum_{j=1}^{d-1} \frac{\partial \ell_j(\mathbf{y})}{\partial y_i} \partial_{x_d} U$ where $\partial_{x_j} U$ denotes the partial derivative of U w.r.t. the j -th component. Denoting $D(\ell) = \left\{ \frac{\partial \ell_j(\mathbf{y})}{\partial y_i} \right\}_{\substack{1 \leq j \leq d-1, \\ 1 \leq i \leq d-1}}$, the ℓ Jacobian, we get the announced formula. Likewise, we obtain the Hessian formula. \square

Corollary 4.4.2. *Let $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ be a transformation function and $U : \mathcal{W} \rightarrow \mathbb{R}_+$ a convex objective function. Assume that there exists $\mathbf{w}^* := \arg \min_{\mathbf{w} \in \mathcal{W}_1^{\geq 0}} U(\mathbf{w}) = \arg \min_{\mathbf{w} \in \mathcal{W}_1} U(\mathbf{w})$ and that $D(\ell)$ is invertible in \mathbf{y} such that $\mathbf{w}^* = \ell(\mathbf{y}^*)$. Then,*

$$\mathbf{y}^* \text{ s.t. } \overline{\nabla_x U}(\ell(\mathbf{y}^*)) = \mathbf{0}_{d-1} \Leftrightarrow \ell(\mathbf{y}^*) = \mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}_1^{\geq 0}} U(\mathbf{w})$$

and $U(\ell(\cdot))$ is convex around \mathbf{y}^* .

Proof. Denoting $\mathcal{H}_U = \{\partial_{x_i x_j} U\}_{i,j=1}^d$ the Hessian of U , it is easy to see that:

- If \mathcal{H}_U is positive semi-definite, then the Lagged Hessian $\overline{\mathcal{H}_U}$ defined in Lemma 4.4.1 is positive semi-definite. Indeed, we can notice that for $\mathbf{x} \in \mathbb{R}^{d-1}$,

$$\mathbf{x}^\top \overline{\mathcal{H}_U} \mathbf{x} = (x_1, \dots, x_{d-1}, -\sum_{i=1}^{d-1} x_i) \mathcal{H}_U (x_1, \dots, x_{d-1}, -\sum_{i=1}^{d-1} x_i)^\top > 0.$$

- If $D(\ell)$ is invertible and $\nabla_{\mathbf{y}} U(\ell(\mathbf{y}^*)) = \mathbf{0}_{d-1}$, since $\nabla_{\mathbf{y}} U(\ell(\mathbf{y}^*)) = D(\ell) \overline{\nabla_x U}$, then $\overline{\nabla_x U}(\ell(\mathbf{y}^*)) = \mathbf{0}_{d-1}$ and the Hessian simplifies to $\mathcal{H}(\mathbf{y}^*) = D(\ell(\mathbf{y}^*)) \overline{\mathcal{H}_U} D(\ell(\mathbf{y}^*))^\top$. So if $\mathcal{H}_U(\ell(\mathbf{y}^*))$ is positive semi-definite, then $H(\ell(\mathbf{y}^*))$ is positive semi-definite.

Equality of \mathbf{w}^* and $\ell(\mathbf{y}^*)$: In the unsaturated case, the Lagrangian of the problem is: $U(\mathbf{w}) - \lambda_0(\mathbf{w}^\top \mathbf{1}_d - 1)$. Deriving with regards to \mathbf{w} , we get: $\nabla_x U = \lambda_0 \mathbf{1}_d$ which is equivalent to the condition $\overline{\nabla_x U} = \mathbf{0}_{d-1}$.

Since we deal with a convex problem, $\overline{\nabla_x U}(\mathbf{w}) = \mathbf{0}_{d-1}$ is a necessary and sufficient condition for \mathbf{w} to be the unique minimizer of U .

By uniqueness of \mathbf{w}^* satisfying $\overline{\nabla_x U} = \mathbf{0}_{d-1}$,

- if \mathbf{y}^* satisfies $\overline{\nabla_x U}(\ell(\mathbf{y}^*)) = \mathbf{0}_{d-1}$, then $\mathbf{w}^* = \ell(\mathbf{y}^*)$ is U unique minimizer.
- Conversely, for \mathbf{w}^* unique minimizer of U , denoting \mathbf{y}^* its pre-image by ℓ , $\overline{\nabla_x U}(\ell(\mathbf{y}^*)) = \overline{\nabla_x U}(\mathbf{w}^*) = \mathbf{0}_{d-1}$.

□

4.4.2.2 Transformation functions

In what follows, we are going to present two main transformation functions, reproducing strictly positive weights $\mathcal{W}_1^{>0} = \{\mathbf{w} \in (\mathbb{R}_+^*)^d : \mathbf{1}_d^\top \mathbf{w} = 1\}$ or non-negative allocations in $\mathcal{W}_1^{\geq 0}$, and check the invertibility of the associated Jacobian matrix.

Log-ratio transform We recall the log-ratio transformation presented in the previous subsection:

$$\begin{cases} w_i &= \ell_i(y) = \frac{e^{y_i}}{1 + \sum_{j=1}^{d-1} e^{y_j}}, \text{ for } i = 1, \dots, d-1, \\ w_d &= 1 - \sum_{j=1}^{d-1} \ell_j(y) = \frac{1}{1 + \sum_{j=1}^{d-1} e^{y_j}}. \end{cases}$$

It is called log-ratio transformation because the invert transformation is defined by: $y_i = \ln\left(\frac{w_i}{w_d}\right)$, $i = 1, \dots, d-1$. We can verify the invertibility of $D(\ell)$ by computing its associated determinant. The proof is postponed to Appendix 4.A.

Lemma 4.4.3 (Invertibility of log-ratio derivatives matrix). *In the log-ratio transformation, the transformation derivatives matrix $D(\ell)$ is defined by*

$$D(\ell) = \begin{pmatrix} \ell_1 - \ell_1^2 & -\ell_1 \ell_2 & \dots & -\ell_1 \ell_{d-1} \\ -\ell_1 \ell_2 & \ell_2 - \ell_2^2 & \dots & -\ell_2 \ell_{d-1} \\ \vdots & & \ddots & \vdots \\ -\ell_1 \ell_{d-1} & \dots & \dots & \ell_{d-1}(1 - \ell_{d-1}) \end{pmatrix}. \quad (4.4.3)$$

The matrix $D(\ell)$ is invertible, as a consequence of the strict positivity of its determinant given by:

$$\det(D(\ell)) = \prod_{i=1}^{d-1} \ell_i \left(1 - \sum_{i=1}^{d-1} \ell_i\right).$$

A limit of the log-ratio transformation is that it does not allow zero component. This is why we proposed the "integral approach". Indeed, we are interested in cases when the allocation is almost saturated. In our experiment section, we will compare both approach (0 component allowed by the transformation or not) in an almost saturated case.

A particular and novel CDA approach: the integral transformation. Let $\mathbf{w} = (w_1, \dots, w_d) \in \mathcal{W}_1^{\geq 0}$ such that

$$H_{\forall} < 1. \quad \forall i \in \{1, \dots, d\}, w_i < 1.$$

$$H_{\exists} \neq 0. \quad \exists i, j \in \{1, \dots, d\}, i \neq j, w_i \neq 0 \text{ and } w_j \neq 0.$$

Definition 4.4.1 (Integral approach). Let $F : \mathbb{R} \rightarrow [0, 1]$ strictly increasing surjective, continuously differentiable function, let $\mathbb{1} = \{x \in \mathbb{R} : F'(x) \neq 0\}$ ("support of the distribution"), and let $B = [0, \sup I - \inf I]$ if $\mathbb{1}$ is bounded, $B = \mathbb{R}_+$ else. Let $p : \mathbb{R} \rightarrow \mathbb{R}_+$ surjective on B , attaining 0 ($\exists t, p(t) = 0$) be given. Then for any $\mathbf{w} \in \mathcal{W}_1^{\geq 0}$ satisfying assumptions $H_{\forall} < 1$ and $H_{\exists} \neq 0$, there exist $y_1, y_2, \dots, y_{d-1} \in \mathbb{R}^{d-1}$ s.t.:

$$\begin{cases} w_1 &= F(y_1), \\ w_i &= F(y_1 + p(y_2) + \dots + p(y_i)) - F(y_1 + p(y_2) + \dots + p(y_{i-1})), \quad i = 2, \dots, d-1, \\ w_d &= 1 - F(y_1 + p(y_2) + \dots + p(y_{d-1})). \end{cases}$$

Introducing the notation $z_i = y_1 + p(y_2) + \dots + p(y_i)$, so that $z_i = z_{i-1} + p(y_i)$, \mathbf{w} is written in a more compact way:

$$\begin{cases} w_1 &= F(z_1) \\ w_i &= F(z_i) - F(z_{i-1}), \quad i = 2, \dots, d-1, \\ w_d &= 1 - F(z_d). \end{cases}$$

Remark 9. F satisfying the conditions of Definition 4.4.1 can be seen as the primitive of a function of integral one. As illustrated in Figure 4.13, the allocation \mathbf{w} (positive vector summing to one) can be seen as portions of an integral summing to one.

The positivity of the function p enforces to have increasing cumulative weights: $F(z_{i-1}) = w_1 + \dots + w_{i-1} \leq w_1 + \dots + w_i = F(z_{i-1} + p(y_i))$. If $z_{i-1} \in \mathbb{1}$, a necessary condition for $z_{i-1} + p(y_i) \in \mathbb{1}$ is to have $\inf \mathbb{1} + p(y_i) \leq \sup \mathbb{1}$, which is enforced if p is surjective on B .

This representation allows for zero-valued components while enforcing $\sum_{i=1}^d w_i = 1$.

We start by giving an example with F taken as the sigmoid function: when $d = 2$, this transformation coincides with the log-ratio transform. For general d , contrary to the log-ratio transform, the intermediary components $w_i, i \in \{2, \dots, d-1\}$ can be equal to 0. This transformation corresponds to a distribution with unbounded support.

Example (Transformation function F with unbounded support: sigmoid transformation). The couple of functions:

- $F(t) = \frac{1}{1+e^{-t}}$ (sigmoid function, represented in Figure 4.13),
- $p(t) = t^{2n}, n \in \mathbb{N}^*$,

enable to represent $\mathcal{W}_1^{\geq 0}$ with non null first and last components.

The problem which can arise when using exponential-based/unbounded support transformation is that the \mathbf{y} can take very large values, especially when w_i is close to 0 or 1: $y_1 = F^{-1}(w_1) = \ln\left(\frac{w_1}{1-w_1}\right)$ goes to $-\infty$ (resp. $+\infty$) when w_1 goes to 0 (resp. 1). It might be difficult to approximate the associated $\mathbf{y}^*(\theta)$ via UQSA.

This is why we propose the following bounded-support, polynomial-based transformation.

Example (Transformation function F with bounded support: Polynomial (Beta) Transformation). Let $F'(t) = 6t(1-t)\mathbb{1}_{0 \leq t \leq 1}$ the Beta(2, 2) density. Then



(a) $F'(x)$ plot(b) $F(x)$ plot

Figure 4.13 – Illustration of the integral approach: the allocation can be seen as portions of the area under the curve of a function with integral equal to one. Hence taking the primitive F of the integral-one function enables to have a nice representation for the allocation. The plots are done with F corresponding to the sigmoid function.

$$\bullet F(t) = \begin{cases} 3t^2 - 2t^3 & \text{if } 0 \leq t \leq 1, \\ 0 & \text{if } t < 0, \\ 1 & \text{if } t > 1. \end{cases}$$

We can enforce $p(y_i) \in I = [0, 1)$ by taking the following p function, surjective from \mathbb{R} to I .

$$\bullet p : \begin{cases} \mathbb{R} \rightarrow [0, 1] \\ t \mapsto \tanh(t^{2n}) \end{cases} .$$

We can take $w_1 = F(p(y_1))$ (instead of $F(y_1), y_1 \in \mathbb{R}$) to enforce $w_1 \in (0, 1)$.

In the integral transformation case, $D(\ell)$ is an upper triangular matrix given by:

$$D(\ell) = \begin{pmatrix} F'(z_1) & F'(z_2) - F'(z_1) & \dots & F'(z_{d-1}) - F'(z_{d-2}) \\ 0 & p'(y_2)F'(z_2) & \dots & p'(y_2)(F'(z_{d-1}) - F'(z_{d-2})) \\ \vdots & & \ddots & \vdots \\ 0 & \dots & p'(y_{d-2})F'(z_{d-2}) & p'(y_{d-2})(F'(z_{d-1}) - F'(z_{d-2})) \\ & & 0 & p'(y_{d-1})F'(z_{d-1}) \end{pmatrix} .$$

Indeed, retaking the notation $z_i = y_1 + p(y_2) + \dots + p(y_i)$, and $w_i = F(z_i) - F(z_{i-1})$,

- for $j > i > 1$, $\frac{\partial w_j}{\partial y_i} = \frac{\partial F(z_j) - F(z_{j-1})}{\partial y_i} = p'(y_i)(F'(z_j) - F'(z_{j-1}))$,
- for $j = i > 1$, $\frac{\partial w_j}{\partial y_i} = \frac{\partial F(z_j) - F(z_{j-1})}{\partial y_i} = p'(y_i)F'(z_i)$,
- for $j > i = 1$, $\frac{\partial w_j}{\partial y_1} = \frac{\partial F(z_j) - F(z_{j-1})}{\partial y_1} = F'(z_j) - F'(z_{j-1})$, and for $i = j = 1$, it gives $F'(y_1)$,
- for $j < i$, $\frac{\partial w_j}{\partial y_i} = 0$.

If we take $w_1 = F(p(y_1))$, we have to multiply the first line by $p'(y_1)$.

Invertibility of $D(\ell)$: the eigenvalues of $D(\ell)$ are $\{F'(z_1), p'(y_2)F'(z_2), \dots, p'(y_{d-1})F'(z_{d-1})\}$ so since $F' > 0$, $D(\ell)$ is invertible as long as $p'(y_i) \neq 0$. For $\mathbf{y}^* = \ell^{-1}(\mathbf{w}^*)$, this is well verified since we deal with unsaturated solutions. **Non-uniqueness of \mathbf{y} representation:** we deal with positive function p , that can take the value 0 in a finite t . With this definition, there are multiple $\mathbf{y} \in \mathbb{R}^{d-1}$ corresponding to the same \mathbf{w} through ℓ . For example, for the symmetric function p considered in our examples, if $\mathbf{w} = \ell(\mathbf{y})$, replacing y_i by $-y_i$ for $i \in \{2, \dots, d-1\}$ leads to the same \mathbf{w} .

4.4.3 Application

This section is devoted to numerical experiments on PCE and UQSA for portfolio allocation.

4.4.3.1 UQSA assumptions verification

In this subsection, we focus on the quadratic objective function $U^{MV}(\mathbf{w})$ given in (4.4.1). We place ourselves in the small perturbation, unsaturated setting:

H_{pert} . for any $\theta \in \Theta$, the solution $\mathbf{w}(\theta) := \arg \min_{\mathbf{w} \in \mathcal{W}_{\mathbb{I}}^{\geq 0}} U_{\theta}^{MV}(\mathbf{w}) = \frac{1}{2} \mathbf{w}^{\top} C_{\theta} \mathbf{w}$ is unsaturated.

We will consider a transformation function ℓ satisfying invertibility of its Jacobian $D(\ell)$.

We start by specifying the H^{MV} function such that $\mathbb{E} \left[H^{MV}(\mathbf{y}(\theta), R_{\theta}, \tilde{R}_{\theta}) \right] = \mathbf{0}_{d-1}$ for $\ell(\mathbf{y}_{\theta}) = \arg \min_{\mathbf{w} \in \mathcal{W}_{\mathbb{I}}^{\geq 0}} U_{\theta}^{MV}(\mathbf{w})$ and we check the validity of UQSA assumptions under mild assumptions on R_{θ} .

Gradient of the transformed quadratic objective. In the following lemma, we specify the minimum-variance objective lagged gradient. In the unsaturated convex setting, a zero of this lagged gradient in \mathbf{y} gives the minimum in $\mathbf{w} = \ell(\mathbf{y})$ of the objective so this lagged-gradient could be used to build our H^{MV} function.

Lemma 4.4.4 (Lagged gradient for the minimum-variance objective). *Let $\Gamma = \{C_{i,j} - C_{i,d} - C_{d,j} + C_{d,d}\}_{i,j=1}^{d-1}$ and $\tilde{C} = \{(C_{i,d} - C_{d,d})\}_{i=1}^{d-1}$. Denote $\bar{\ell}(\mathbf{y}) = (\ell_1(\mathbf{y}), \dots, \ell_{d-1}(\mathbf{y}))$. The lagged gradient $\overline{\nabla_x U^{MV}} = \{\partial_i U^{MV} - \partial_d U^{MV}\}_{i=1}^{d-1}$ is*

$$\overline{\nabla_x U^{MV}} = \Gamma \bar{\ell} + \tilde{C},$$

and the lagged Hessian is the positive definite matrix Γ .

Proof. It is easy to see that $\nabla_x U^{MV} = C\mathbf{w}$. The lagged gradient is then $(C\mathbf{w})_i - (C\mathbf{w})_d$. Replacing \mathbf{w} by $\ell(\mathbf{y})$ and using that $w_d = 1 - \sum_{i=1}^{d-1} w_i = 1 - \sum_{i=1}^{d-1} \ell_i(\mathbf{y})$, we get the announced formula. \square

As a consequence, by Lemma 4.4.1, $\nabla_y U^{MV} = D(\ell) [\Gamma \bar{\ell} + \tilde{C}]$ and $U^{MV}(\ell(\mathbf{y}))$ is convex around \mathbf{y}^* such that $\ell(\mathbf{y}^*) = \mathbf{w}^*$.

H^{MV} function choice for the ℓ transformation: Let R, \tilde{R} i.i.d. vector of asset returns of covariance matrix C . Let $i, j \in \{1, \dots, d\}$. Using the fact that $C_{i,j} = \mathbb{E} [R_i R_j - R_i \tilde{R}_j]$, we denote $\Gamma(R, \tilde{R})$ and $\tilde{C}(R, \tilde{R})$ the matrix and vector replacing $C_{i,j}$ by $R_i R_j - R_i \tilde{R}_j$ in Γ and \tilde{C} defined in Lemma 4.4.4. Then $\mathbb{E} [\Gamma(R, \tilde{R})] = \Gamma$ and $\mathbb{E} [\tilde{C}(R, \tilde{R})] = \tilde{C}$.

We take as function H^{MV} :

$$H^{MV}(\mathbf{y}, R, \tilde{R}) = D(\ell) \left[\Gamma(R, \tilde{R}) \bar{\ell}(\mathbf{y}) + \tilde{C}(R, \tilde{R}) \right],$$

where $\bar{\ell}(\mathbf{y}) = (\ell_1(\mathbf{y}), \dots, \ell_{d-1}(\mathbf{y}))$.

Assumptions verification: In the general $D(\ell)$ -invertible case, we can show that the UQSA assumptions are satisfied at least at the neighborhood of the optimal solution, under some mild assumptions on R_θ and ℓ . $\|\cdot\|$ will denote indifferently the Euclidean norm on \mathbb{R}^{d-1} or the Euclidean-subordinated matrix norm.

$$H_{unif R}. \mathbb{E}_{R_{i,\theta}} [R_{i,\theta}^4] < \infty \text{ and } \sup_{\theta} \mathbb{E}_{R_{i,\theta}} [R_{i,\theta}^4] < \infty, 1 \leq i \leq d.$$

$H_{unif \Gamma}$. For any $\theta \in \Theta$, $\Gamma_\theta := \mathbb{E}_{R_\theta, \tilde{R}_\theta} \Gamma(R_\theta, \tilde{R}_\theta)$ is positive definite, with minimum eigenvalue uniformly bounded from below in $\theta \in \Theta$ by $\gamma > 0$.

$H_{D(\ell)}$. The Jacobian of the transformation function is at most linear in \mathbf{y} : there exists C_ℓ such that:

$$\|D(\ell(\mathbf{y}))\| \leq C_\ell (1 + \|\mathbf{y}\|).$$

$H_{\varepsilon, \mathbf{y}^*}$. There exists $\varepsilon > 0$, $\theta_0 \in \Theta$ and $\mathbf{y}^*(\theta_0)$, $\mathbf{y}^*(\theta_0) = \ell(\mathbf{w}^*(\theta_0))$, such that

- For any $\theta \in \Theta$, there is a unique $\mathbf{y}^*(\theta) \in \mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$ s.t. $\mathbf{y}^*(\theta) = \ell(\mathbf{w}^*(\theta))$ (small perturbation setting).
- For any $\mathbf{y} \in \mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$, $D(\ell(\mathbf{y}))$ is invertible.
- For any $\mathbf{y} \in \mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$, for any $\theta \in \Theta$, $\langle \nabla_y U_\theta^{MV}(\ell(\mathbf{y})), \mathbf{y} - \mathbf{y}^* \rangle_\pi \geq c \|\mathbf{y} - \mathbf{y}^*\|_\pi^2$ (strong convexity) where $\nabla_y U_\theta^{MV} = D(\ell) (\Gamma_\theta \bar{\ell}(\mathbf{y}) - \tilde{C}_\theta)$.

Remark 10. Assumption $H_{\varepsilon, \mathbf{y}^*}$ interpretation:

- In practice, there might be multiple $\mathbf{y} \in \mathbb{R}^{d-1}$ such that $\mathbf{y} = \ell(\mathbf{w}^*(\theta))$. $\mathbf{H}_{\varepsilon, \mathbf{y}^*}$ assumes that there is only one such \mathbf{y} in $\mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$.
- $D(\ell(\mathbf{y}^*(\theta)))$ invertibility is rather a consequence of \mathbf{H}_{pert} : since $\mathbf{w}^*(\theta)$ is unsaturated, by ℓ choice, $D(\ell(\mathbf{y}^*(\theta)))$ is invertible.
- Under assumptions \mathbf{H}_{pert} and $\mathbf{H}_{\text{unif } \Gamma}$, the local strong convexity condition can be seen as a consequence of the strict positivity of $D(\ell)\Gamma_\theta D(\ell)^\top$ eigenvalues. Indeed, under these assumptions, $D(\ell)$'s eigenvalues are non-null and Γ_θ 's are strictly positive, so $D(\ell)\Gamma_\theta D(\ell)^\top$ is symmetric positive definite. In what follows, we will assume the strongest assumption $\mathbf{H}_{\varepsilon, \mathbf{y}^*}$ only (and not $\mathbf{H}_{\text{unif } \Gamma}$).

Let us assume \mathbf{H}_{pert} , \mathbf{H}_{int} , $\mathbf{H}_{D(\ell)}$ and $\mathbf{H}_{\varepsilon, \mathbf{y}^*}$. In what follows, we are going to show that the UQSA assumptions hold at least on $\mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$.

\mathbf{H}_{T^*} For any $\theta \in \Theta$, by convexity of the objective function $U_\theta^{MV}(\mathbf{w}) = \frac{1}{2}\mathbf{w}^\top C_\theta \mathbf{w}$, the solution $\mathbf{w}^*(\theta)$ exists, is unique and unsaturated from assumption \mathbf{H}_{pert} . From $\mathbf{H}_{\varepsilon, \mathbf{y}^*}$, there is a unique $\mathbf{y}^*(\theta) \in \mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$ such that $\mathbb{E} \left[H^{MV}(\mathbf{y}^*(\theta), R_\theta, \tilde{R}_\theta) \right] = 0$ π -a.e. so T^* is compact.

\mathbf{H}_{int} By $\mathbf{H}_{\text{unif } R}$, at fixed $\theta \in \Theta$, $\mathbb{E}_{R_{i,\theta}} R_{i,\theta}^2$ is finite. Since H^{MV} is quadratic in $R_{i,\theta}$, H^{MV} is integrable, and $h^{MV}(\mathbf{y}, \theta) = \mathbb{E}_{R_\theta, \tilde{R}_\theta} H^{MV}(\mathbf{y}, R_\theta, \tilde{R}_\theta)$ exists. Denoting $\Gamma_\theta := \mathbb{E}_{R_\theta, \tilde{R}_\theta} \Gamma(R_\theta, \tilde{R}_\theta)$ and $\tilde{C}_\theta := \mathbb{E}_{R_\theta, \tilde{R}_\theta} \tilde{C}(R_\theta, \tilde{R}_\theta)$, we have, for any arbitrary $\mathbf{y} \in \mathbb{R}^{d-1}$

$$h^{MV}(\mathbf{y}, \theta) = D(\ell) \mathbb{E}_{R_\theta, \tilde{R}_\theta} \left(\Gamma(R, \tilde{R}) \bar{\ell}(\mathbf{y}) + \tilde{C}(R, \tilde{R}) \right) = D(\ell) \left(\Gamma_\theta \bar{\ell}(\mathbf{y}) + \tilde{C}_\theta \right).$$

By $\mathbf{H}_{D(\ell)}$, $\|h^{MV}(\mathbf{y}, \theta)\|$ is at most linear in $\|\mathbf{y}\|$. From assumption $\mathbf{H}_{\text{unif } R}$, $\mathbb{E}_{R_{i,\theta}} R_{i,\theta}^2$ is uniformly bounded in θ , so $\|\Gamma_\theta\|$ and $\|\tilde{C}_\theta\|$ are uniformly bounded in θ . By ℓ definition, $\|\ell\| \leq 1$. So there exists a $C_{\Theta, \ell} > 0$,

$$\|h^{MV}(\mathbf{y}, \theta)\| \leq C_\ell \sup_{\theta \in \Theta} \|\Gamma_\theta\| (1 + \|\mathbf{y}\|) + \sup_{\theta \in \Theta} \|\tilde{C}_\theta\| \leq C_{\Theta, \ell} (1 + \|\mathbf{y}\|).$$

So $\|h^{MV}(\mathbf{y}, \theta)\|$ is bounded by a linear function in $\|\mathbf{y}\|$ with a coefficient uniformly bounded in θ , so $\phi \mapsto h^{MV}(\phi(\cdot), \cdot)$ is a continuous function from $L_{2,d-1}^\pi$ to $L_{2,d-1}^\pi$.

\mathbf{H}_{quad} Similarly, by $\mathbf{H}_{\text{unif } R}$ and by $\mathbf{H}_{D(\ell)}$, we can find a constant $C_{\Theta, \ell} > 0$ such that:

$$\sup_{\theta \in \Theta} \mathbb{E}_{R_\theta, \tilde{R}_\theta} \|H^{MV}(\mathbf{y}, R_\theta, \tilde{R}_\theta)\|^2 \leq C_{\Theta, \ell} (1 + \|\mathbf{y}\|).$$

\mathbf{H}_{sep} , $\mathbf{H}_{\text{coerc}}$ The separability and coercivity assumptions consist in verifying that $\langle h^{MV}(\mathbf{y}, \cdot), \mathbf{y} - \mathbf{y}^* \rangle_\pi \geq c \|\mathbf{y} - \mathbf{y}^*\|_\pi^2$ for some constant $c > 0$. With our choice of h^{MV} function, $h^{MV}(\mathbf{y}, \theta) = \nabla_{\mathbf{y}} U_\theta^{MV}$. By the strict convexity assumption $\mathbf{H}_{\varepsilon, \mathbf{y}^*}$ [Bullet 3], such a constant exists and we verify well the coercitivity and separation assumptions.

Remark 11. We could have rigorously quantified an ε value and specified a use case in which the strict convexity holds on $\mathcal{B}(\mathbf{y}^*(\theta_0), \varepsilon)$. It has not been done but the heuristic arguments given above show that UQSA assumptions are locally verified.

4.4.3.2 Experiments

In this subsection, we apply the transformation technique on U^{MV} . Our aim is twofold:

1. we want to know which transformation function is the most relevant/efficient to do the PCE of the transformed function $\ell^{-1}(\mathbf{w})$,

2. we would like a transformation which deals well with almost saturated solutions.

We focus on the uncertainty on the correlation. In the three assets case, the covariance matrix writes: $C = \text{Diag}(\sigma_1, \sigma_2, \sigma_3) \begin{pmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{pmatrix} \text{Diag}(\sigma_1, \sigma_2, \sigma_3)$. In our applications, we took $(\sigma_1, \sigma_2, \sigma_3) = [10, 11, 12]\%$ and $R \sim \mathcal{N}(\mathbf{0}_d, C)$.

Allocation, choice of admissible ρ values: By focusing on the quadratic objective, the optimal allocation takes the explicit form: $\mathbf{w}^* = \frac{C^{-1}\mathbf{1}}{\mathbf{1}^T C^{-1}\mathbf{1}}$. We focused on the case $\rho \in [20, 77]\%$, for which $w_3(77\%) \approx 0$ (almost saturated solution). We show in Figure 4.18 the evolution on \mathbf{w} on this range of correlation.




Figure 4.14 – Evolution of allocation $\mathbf{w}^* = (w_1, w_2, w_3)$ components as a function of ρ for $\rho \in [20, 77]\%$

Choice of transformation functions: We focus on the following transformation functions:

- log-ratio transformation;
- integral approach transformations:
 - sigmoid-polynomial: $F(z) = \frac{1}{1+e^{-z}}$, $p(t) = t^2$;
 - beta-polynomial: $F(z) = 3z^2 - 2z^3$, $p(t) = t^2$;
 - beta-hyperbolic-tangent: $F(z) = (3z^2 - 2z^3)$ if $z \in [0, 1]$, 0 if $z < 0$, 1 else, $p(t) = \tanh(t^2)$;

Remark 12. These choices of p function satisfy $H_{D(\ell)}$ assumption, because $p'(t)$ is at most linear in t and the other terms appearing in $D(\ell)$ are bounded.


- sum transformation ($\mathbf{y} = (w_1, \dots, w_{d-1})$).

Choice of uncertainty function: We retake the incomplete Beta function combination f_β as taken in the Sharpe ratio study:


$$\rho = f_\beta(\theta) = \rho_{\min} + (\rho_{\max} - \rho_{\min}) \sum_{i=1}^{n_B} \lambda_i F_{\alpha_i, \beta_i}^{l_i, s_i} \left(\frac{1 + \theta}{2} \right),$$

with $n_B, \lambda_i, l_i, s_i, \alpha_i$ and β_i as given in Figure 4.4 and $\rho_{\min} = 20\%$, $\rho_{\max} = 77\%$.


PCE components and truncation error study: In the following figures, we plot the coefficients estimation as well as the truncation error for the different transformation functions.




(a) log-ratio




(b) sigmoid-polynomial



(c) beta-polynomial



(d) beta-hyperbolic tangent



(e) sum

Figure 4.15 – Evolution of the coefficients $\mathbf{u}_i^* = (u_{i,1}^*, u_{i,2}^*)$ of $\ell(\mathbf{y}^*(\theta))$ projection on the Legendre basis for different transformation functions.

We can make the following comments:

- As seen in Figure 4.15[(a) and (b)], in the log-ratio and the sigmoid transformations, the coefficients are of different signs and the significant coefficients are more numerous than in the beta or the sum transformation. With the log-ratio and the sigmoid transformations, the UQSA algorithm will have difficulties in estimating the largest order coefficients, because their estimation starts when the value of the step size is relatively small (increasing

dimension approach) whereas the u_i^* profiles for the beta-based and sum transformations are more suited to the UQSA algorithm.

- As shown on Figure 4.16, the beta transformations and the sum transformation have a smaller truncation error (but a higher variance for the coefficients of order higher than 30). Less coefficients are needed to estimate \mathbf{y} with these transformations. The convergence rate of the truncation error is equal to 2.2 which is approximately in line with the expected 2 convergence rate (order of differentiability of the uncertainty function).






Figure 4.16 – Truncation error $\mathcal{E}_{TR} := (\sum_{i>m} \|\mathbf{u}_i^*\|_2^2)^{1/2} = \|\mathbf{y}^* - \sum_{i \leq m} \mathbf{u}_i^* B_i\|_\pi$ computed via integration, as a function of the number of coefficients m in log-log scale, and the associated tendency line, for $\mathbf{y}^*(\theta) = \ell^{-1}(\mathbf{w}^*(\theta))$, for various choices of ℓ .

In Figure 4.17, we display the approximation of the true allocation function for the log-ratio and the beta-polynomial transformations. We observe that approximately 20 coefficients are needed to approximate well the true function. We see that in the log-ratio transformation, the second component of \mathbf{w} is less well approximated than with the beta transformation.



(a) log-ratio



(b) beta-polynomial

Figure 4.17 – Evolution of $\mathbf{w}^*(\theta)$ approximation as a function of $\rho = f_\beta(\theta)$ for different values of m in $\{5, 10, 15, 20, 25\}$ (with coefficients computed via integration) for the log-ratio and the beta polynomial transformation function.




Figure 4.18 – True function approximation with UQSA, obtained after transformation on \mathbf{y}^* approximation, $a = 0.9$, $b = 0.2$, $K = 5000$.




Figure 4.19 – Total error evolution with UQSA, obtained after transformation on \mathbf{y}^* approximation, $a = 0.9$, $b = 0.2$, $K = 5000$.




Figure 4.20 – $\mathbf{y}^*(\theta)$ approximation with UQSA, $a = 0.9$, $b = 0.2$, $K = 5000$.

UQSA first experiment - increasing dimension and beta-polynomial transformation:

In Figure 4.19, we represent the evolution of the total error for one launch of the algorithm. The approximation of the true function with the coefficients after 5000 iterations are plotted in Figure 4.18. We see that the shape of the true function is well captured except for the second component, which is under-estimated for θ close to 1. The output of UQSA algorithm is displayed in Figure 4.20 against the true $\mathbf{y}^*(\theta)$ function. We have a better convergence in infinite norm for \mathbf{y} than for its transformation \mathbf{w} . The UQSA results with sigmoid transformation (not displayed) are far more instable.

4.4.4 Conclusion to UQ for portfolio analysis section

In this section, we have explored an approach to tackle the problem of portfolio allocation with uncertainty using UQSA algorithm based on change of variable techniques. We have empirically illustrated the use of different changes of variable, or "transformation functions". We propose a new CDA transformation that we call the integral approach. Within this class of transformations, bounded support, polynomial-based transformations seem more suited to our problem than sigmoid based transformations.

In the current setting, our approach has a limited scope of applicability since we can guarantee the optimality of the solution found by the transformed problem only in cases when the optimal solution lies in the interior of the simplex. When the solution is saturated and lies in the boundary of the simplex, then defining a smooth transformation $\ell(\cdot)$ appears hopeless, while keeping a unique y^* . Such a reparametrization approach is presumably not appropriate for saturated problem and we leave to further research the investigation of alternative parametrizations.

APPENDICES

In this appendices section, we give the proof that the log-ratio transformation is well defined (Section 4.A) and we give additional explanations on the UQSA algorithm (Section 4.B).

4.A Proof of Lemma 4.4.3

In the log-ratio transformation, $\ell_j(\mathbf{y}) = \frac{e^{y_j}}{1 + \sum_{i=1}^{d-1} e^{y_i}}$ for $j \in \{1, \dots, d-1\}$ so

$$\frac{\partial \ell_j}{\partial y_i} = \begin{cases} \frac{e^{y_i}}{1 + \sum_{k=1}^{d-1} e^{y_k}} - \frac{e^{y_i}}{(1 + \sum_{k=1}^{d-1} e^{y_k})^2} e^{y_i} = \ell_i - \ell_i^2 & \text{if } i = j, \\ -\frac{e^{y_j}}{(1 + \sum_{k=1}^{d-1} e^{y_k})^2} e^{y_i} = -\ell_i \ell_j & \text{else,} \end{cases}$$

so $D(\ell)$ is given by (4.4.3). We denote by $\det(D(\ell))$ the determinant of the $D(\ell)$ matrix. We start by simplifying a little $\det(D(\ell))$: multiplying each column by $\frac{1}{\ell_i}$, which amounts to multiply the determinant by $\prod_{i=1}^{d-1} \ell_i$, we get:

$$\det(D(\ell)) = \prod_{i=1}^{d-1} \ell_i \begin{vmatrix} 1 - \ell_1 & -\ell_1 & \dots & -\ell_1 \\ -\ell_2 & 1 - \ell_2 & \dots & -\ell_2 \\ \vdots & \ddots & \ddots & \vdots \\ -\ell_{d-1} & \ddots & -\ell_{d-1} & 1 - \ell_{d-1} \end{vmatrix}.$$

Multiplying each row by $\frac{1}{\ell_i}$, which amounts to multiply the determinant by $\prod_{i=1}^{d-1} \ell_i$, we get:

$$\det(D(\ell)) = \prod_{i=1}^{d-1} \ell_i^2 \begin{vmatrix} \frac{1-\ell_1}{\ell_1} & -1 & \dots & -1 \\ -1 & \frac{1-\ell_2}{\ell_2} & \dots & -1 \\ \vdots & \ddots & \ddots & \vdots \\ -1 & \ddots & -1 & \frac{1-\ell_{d-1}}{\ell_{d-1}} \end{vmatrix}.$$

We proceed by induction. We denote $\frac{\det(D(\ell))}{\prod_{i=1}^{d-1} \ell_i^2} =: \mathcal{D}(\ell_1, \dots, \ell_{d-1})$. Our induction hypothesis writes:

$$\mathcal{D}(\ell_1, \dots, \ell_{d-1}) = \frac{1 - \sum_{i=1}^{d-1} \ell_i}{\prod_{i=1}^{d-1} \ell_i}. \quad (H_d)$$

- For $d = 3$, $\mathcal{D}(\ell_1, \ell_2) = \begin{vmatrix} \frac{1-\ell_1}{\ell_1} & -1 \\ -1 & \frac{1-\ell_2}{\ell_2} \end{vmatrix} = \frac{1-\ell_1-\ell_2}{\ell_1\ell_2}$.
- Let $d \in \mathbb{N}^*$, $d \geq 3$, be given. Let us assume (H_d) and let us show (H_{d+1}) .
Developing along the last column:

$$\mathcal{D}(\ell_1, \dots, \ell_d) = \frac{1-\ell_d}{\ell_d} \mathcal{D}(\ell_1, \dots, \ell_{d-1}) + 1\Delta_{d-1,d} - 1\Delta_{d-2,d} + \dots + (-1)^{d+2}\Delta_{1,d}$$

where $\Delta_{i,j}$ denotes the (i, j) minor associated to the determinant $\mathcal{D}(\ell_1, \dots, \ell_d)$.

Each minor $\Delta_{d-i,d}$ consists in a last row containing only "−1", a diagonal containing $\left\{ \frac{1-\ell_k}{\ell_k} \right\}_{1 \leq k \leq d-i-1}$, and for $i > 1$, terms just above the diagonal in $\left\{ \frac{1-\ell_k}{\ell_k} \right\}_{d-i+1 \leq k \leq d-1}$ and −1 everywhere else.

Subtracting the last row to the other rows (which amounts to add +1 on every row except the last) makes appear "0" and makes the calculation straightforward. For example, for $i = 1$, subtracting the last row in (*):

$$\Delta_{d-1,d} = \begin{vmatrix} \frac{1-\ell_1}{\ell_1} & -1 & \dots & -1 \\ -1 & \frac{1-\ell_2}{\ell_2} & \dots & -1 \\ \vdots & \ddots & & \vdots \\ & & \frac{1-\ell_{d-2}}{\ell_{d-2}} & -1 \\ -1 & \ddots & -1 & -1 \end{vmatrix} \stackrel{(*)}{=} \begin{vmatrix} \frac{1}{\ell_1} & 0 & \dots & 0 \\ 0 & \frac{1}{\ell_2} & \dots & 0 \\ \vdots & \ddots & & \vdots \\ & & \frac{1}{\ell_{d-2}} & 0 \\ -1 & \ddots & -1 & -1 \end{vmatrix} = - \prod_{i=1}^{d-2} \frac{1}{\ell_i}.$$

Likewise, we find: $\Delta_{d-2,d} = \prod_{i=1, i \neq d-2}^{d-1} \frac{1}{\ell_i}$, $\Delta_{d-3,d} = - \prod_{i=1, i \neq d-3}^{d-1} \frac{1}{\ell_i}$, ...

Putting under the same denominator:

$$\begin{aligned} \mathcal{D}(\ell_1, \dots, \ell_d) &= \frac{(1-\ell_d)(1 - \sum_{i=1}^{d-1} \ell_i) - \ell_d \ell_{d-1} - \ell_d \ell_{d-2} - \dots - \ell_d \ell_1}{\prod_{i=1}^d \ell_i} \\ &= \frac{1 - \sum_{i=1}^d \ell_i}{\prod_{i=1}^d \ell_i}. \end{aligned}$$

4.B UQSA algorithm: a projection approach

In this subsection, we give the intuition underlying the UQSA algorithm.

Note that the problem (4.1.1) is equivalent to finding $\phi^* \in L_{2,q}^\pi$ such that

$$\int_{\Theta} \left(\int_{\mathcal{V}} H \left(\sum_{j \geq 0} u_j B_j(\theta), v, \theta \right) \mu(\theta, dv) \right) B_i(\theta) \pi(d\theta) = \mathbf{0}_q, \quad \forall i \in \mathbb{N}.$$

Note also that the SA approach for a single θ corresponds to the update:

$$\hat{\phi}^{(k+1)}(\theta) = \hat{\phi}^{(k+1)}(\theta) - \gamma_{k+1} H(\hat{\phi}^{(k)}(\theta), V_{k+1}, \theta).$$

Integrating in θ and V the equation above gives the following update equation of the $\hat{\phi}^{(k+1)}$ components, $u_i^{(k+1)}$:

$$u_i^{(k+1)} = u_i^{(k)} - \gamma_{k+1} \int_{\Theta \times \mathcal{V}} H(\hat{\phi}^{(k)}(\theta), v, \theta) B_i(\theta) \mu(\theta, dv) \pi(d\theta), \quad i \leq m.$$

The SA algorithm consists in the following update of the sequence, replacing the integral by a Monte Carlo approximation:

$$u_i^{(k+1)} = u_i^{(k)} - \frac{\gamma_{k+1}}{M_{k+1}} \sum_{s=1}^{M_{k+1}} H(\hat{\phi}^{(k)}(\theta_s^{(k+1)}, V_s^{(k+1)}) B_i(\theta_s^{(k+1)}), \quad i \leq m_{k+1},$$

where m_k and M_k are nondecreasing integer sequences converging to ∞ and $(\theta_s^{(k+1)}, V_s^{(k+1)}, s = 1, \dots, M_{k+1})$ independent variables sampled under $\pi(d\theta)\mu(\theta, dv)$.

Part IV

Portfolio optimization with kernel and optimal transport (K.O.T.) divergence

CHAPTER 5

K.O.T. PORTFOLIO OPTIMIZATION WITH TARGET DISTRIBUTION

Note. This chapter corresponds to a work in progress with Z. Szabó.

Abstract. Portfolio optimization is one of the most important task in finance. The aim is to create an investment portfolio matching the investor's preference. Minimum divergence portfolio consists in finding the portfolio which fits the best the ideal target portfolio distribution the investor has in mind. Previous applications focused on the Kullback-Leibler or f -divergence measures, with model assumptions on the historical returns possibly erroneous and incompatible with the targeted distribuion. In this chapter, we investigate the minimum kernel and optimal transport (K.O.T.) divergences-based portfolios, with parametric distributions allowing to target light tails or positively skewed distributions, which are generally appealing properties from the investor's point of view. Our approach is model-free on the returns. In this chapter, we compute explicit forms of these divergence metrics in terms of the target distribution. In the kernel-based Maximum Mean Discrepancy (MMD), we establish new convergence guarantees in case of bounded kernels as well as for unbounded exponential kernels, when the MMD estimator is explicit in the target distribution. The approach is illustrated through numerical experiments.

5.1 Introduction

Portfolio optimization [FKPF12] is among the most fundamental and important tasks in finance. In this chapter we show how kernel techniques and optimal transport can be leveraged to design efficient portfolio optimization schemes, and present novel theoretical results reducing the uncertainty in the estimates of the objectives which can also be of independent interest.

In portfolio optimization the goal is to find an allocation in line with the preferences of an investor. In practice, the task can be formulating for instance by optimizing a concave utility function encoding the investor's priorities [Mer71]. However, oftentimes there is a significant gap between the investor's targeted and the realized portfolio distributions. Traditional portfolio optimization schemes (such as mean-variance optimization; [Mar52]) take into account the first two moments of the portfolio distribution which can be insufficient since financial returns are generally not Gaussian and exhibit non-zero excess skewness and kurtosis [Con01]. Approaches considering the smoothing of the mean-variance utility function [Jor86, LW03, KFCF14] and the inclusion of higher order moments [MZ10] have become more and more dominant. Though

these techniques relax the mean-variance paradigm, they are restricted to considering finite many moments of the portfolio distribution only.

These encouraging directions to match the investors' preferences have been recently superseded by the so-called target distribution technique [CW12, Las19], where the goal is to find a portfolio weight $\mathbf{w} \in \mathcal{W} \subseteq \mathbb{R}^d$ by minimizing the discrepancy (measured in the sense of a divergence D) between the associated portfolio return distribution $P = \mathbf{w}^\top R \sim \mathbb{P}_{\mathbf{w}}$ and the investor's target distribution of portfolio returns \mathbb{P}_T :

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} D(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T). \quad (5.1.1)$$

The domain \mathcal{W} expresses the constraints on the portfolio weights; for instance budget (\mathcal{W}_1), or budget with non-negativity constraints ($\mathcal{W}_1^{\geq 0}$) can be formulated as

$$\mathcal{W}_1 = \left\{ \mathbf{w} \in \mathbb{R}^d : \sum_{j=1}^d w_j = 1 \right\}, \quad \mathcal{W}_1^{\geq 0} = \left\{ \mathbf{w} \in (\mathbb{R}^{\geq 0})^d : \sum_{j=1}^d w_j = 1 \right\}. \quad (5.1.2)$$

[CW12] made use of ϕ -divergences ($D(\mathbb{P}, \mathbb{Q}) = D_\phi(\mathbb{P}, \mathbb{Q}) = \int_{\mathbb{R}} \phi\left(\frac{d\mathbb{P}}{d\mathbb{Q}}(x)\right) dx$ where \mathbb{P} is assumed to be absolutely continuous with respect to \mathbb{Q} ($\frac{d\mathbb{P}}{d\mathbb{Q}}$ denotes the corresponding Radon-Nikodym derivative), with specific focus on power divergences.¹ Particularly, the authors proved and used the dual representation

$$D_\phi(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T) = \sup_{\mathbb{P} : \text{supp}(\mathbb{P}) = \text{supp}(\mathbb{P}_T)} d_\phi(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T, \mathbb{P}), \quad (5.1.3)$$

$$d_\phi(\mathbb{P}_{\mathbf{w}}, \mathbb{P}_T, \mathbb{P}) = \int_{\mathbb{R}} \phi\left(\frac{d\mathbb{P}}{d\mathbb{P}_T}(x)\right) d\mathbb{P}_T(x) + \int_{\mathbb{R}} \phi'\left(\frac{d\mathbb{P}}{d\mathbb{P}_T}(x)\right) d\mathbb{P}_{\mathbf{w}}(x) - \int_{\mathbb{R}} \phi'\left(\frac{d\mathbb{P}}{d\mathbb{P}_T}(x)\right) d\mathbb{P}(x)$$

to determine the optimal portfolio weights with supp denoting the support of its argument.

[Las19] specialized the ϕ -divergence framework to Kullback-Leibler divergence when $\phi(x) = x \log(x) - x + 1$, and assumed \mathbb{P}_T to be a generalized normal (GN) distribution with pdf $f_T(t) = \frac{2^{-\frac{\gamma+1}{\gamma}} \gamma}{\beta \Gamma(1/\gamma)} e^{-\frac{1}{2} \left(\frac{t-\alpha}{\beta}\right)^\gamma}$ ($\alpha \in \mathbb{R}$, $\beta > 0$, $\gamma > 0$). With this specific choice, he showed (Proposition 5.1) that (5.1.1) is equivalent to

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{2\beta^\gamma} \mathbb{E}_{P \sim \mathbb{P}_{\mathbf{w}}} [|P - \alpha|^\gamma] - H(P), \quad (5.1.4)$$

where $H(P)$ denotes the Shannon differential entropy of P . By further assuming that P is a mixture of Gaussian distribution (MoG) $\mathbb{E}_{P \sim \mathbb{P}_P} [|P - \alpha|^\gamma]$ takes a closed form expressed via the Kummer's confluent hypergeometric function [Las19, Proposition 5.4].

Despite the pioneering nature of these contributions, they unfortunately suffer from serious computational and modelling bottlenecks: the first approach [CW12] requires ad-hoc heuristics for encoding the constraints in (5.1.3), the second work [Las19] needs strong parametric assumptions (GN) in addition to the slight incompatibility of the GN and MoG assumptions. It is also worth mentioning the related approach of [GLLN20] who proposed an optimal-transport based technique to reach a target terminal wealth distribution in the continuous time setting. Their goal was to find the portfolio allocation path \mathbf{w}_t , from time $t = 0$ to some horizon H , by minimizing the integral over the time horizon of a given cost function f augmented with a penalty term of the form $D(\mathbb{P}_{\mathbf{w}_H}, \mathbb{P}_T)$. The penalization aims at steering the terminal distribution of the portfolio $\mathbb{P}_H = \mathbf{w}_H^\top R_H$ towards the pre-defined target distribution \mathbb{P}_T . Similarly to [Las19], one of the discrepancy measures D was chosen to be the Kullback-Leibler divergence. The returns

¹In this case $\phi(x) = \frac{x^\gamma - 1 - \gamma(x-1)}{\gamma(\gamma-1)}$ where $\gamma \in \mathbb{R}$.

were assumed to follow the Black-Scholes model with time-dependent drift and volatility, which again falls under the umbrella of parametric portfolio return distributions.

In order to mitigate these severe shortcomings, in this chapter we propose to use alternative divergence measures based on kernels and optimal transport: Maximum Mean Discrepancy (MMD, [SGSS07, GBR⁺12]), Kernel Stein Discrepancy (KSD, [CSG16, LLJ16]) and Wasserstein distance (WAD, [Vil08, PC19]). There are various reasons which motivate the choice of these discrepancy measures. They are flexible in terms of the target distribution: Stein discrepancy requires the knowledge of the pdfs up to normalizing constant only, the Wasserstein distance can be written explicitly as a function of the inverse cdf (see (5.2.6)), the mean embedding (which form the basis of MMD) can be computed analytically for different target distribution and kernel pairs (as it will be detailed in Section 5.3.1). Successful applications demonstrating the practicality of these discrepancy measures include for instance model criticism [LDG⁺14, KKK16], two-sample [HBM07, GBR⁺12], independence [GFT⁺08, PBSP18] and goodness-of-fit testing [LLJ16, CSG16, JXS⁺17, GM17, HM18, GWJ⁺20, GLHL20, BLY21], portfolio valuation [BF19], statistical inference of generative models [BBDG19] and post selection inference [YUFT18], causal discovery [MPJ⁺16, PBSP18], generative adversarial networks [DRG15, LSZ15, BSAG18], assessing and tuning MCMC samplers [GM15, GM17, HM18, GRM20], or designing Monte-Carlo control functionals for variance reduction [OGC17, OCBG19, SKN⁺20], among many others.

Our **contributions** are two-fold:

1. on the theoretical side: (i) we compute analytically the mean embedding (the underlying representation of probability measures used in MMD) for various target distribution - kernel pairs (Section 5.3.1). (ii) We show (Theorem 5.2) that such analytical knowledge leads to better concentration properties of MMD estimators, (iii) extend the result to the case of unbounded kernels (Theorem 5.3; recently motivated in finance for instance by [BF19]), and (iv) present minimax lower bounds (Theorem 5.4).
2. on the practical front: we demonstrate that the flexibility of considered measures result in more efficient portfolios in terms of skewness and kurtosis (Section 5.4-5.5), which provides a novel application of these divergence measures.

The chapter is organized as follows: in Section 5.2 we formulate our problem after introducing a few notations. Section 5.3 is devoted to our main theoretical contributions with focus on explicit mean-embedding functions and concentration results. The numerical efficiency of our approach is demonstrated on simulated data (Section 5.4) and on real-world financial benchmarks (Section 5.5). Conclusions are drawn in Section 5.6. The proofs of our major results are collected in Section 5.A. In Section 5.B tools which are useful for practical implementation are gathered. Section 5.C contains external statements used in our proofs.

5.2 Problem formulation and optimization

In this section we formulate our problem: we define the divergence measures used in our portfolio optimization objective functions, we describe their estimators and the optimization applied.

Notations: Natural numbers are denoted by $\mathbb{N} = \{0, 1, \dots\}$. We use the shorthand $\mathbb{N}^* = \{1, 2, \dots\}$ for positive integers. For $N \in \mathbb{N}^*$, $[N] = \{1, \dots, N\}$. Positive reals are denoted by \mathbb{R}^+ . The vector of ones in \mathbb{R}^d is $\mathbf{1}_d$; the transpose of a vector \mathbf{v} is denoted by \mathbf{v}^\top . The beta function for $\alpha \in \mathbb{R}^+$ and $\beta \in \mathbb{R}^+$ is defined as $B(\alpha, \beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1} dt$. The minimum of $a, b \in \mathbb{R}$ is denoted by $a \wedge b$; their maximum is $a \vee b$. Given $(L_s)_{s \in [S]} \subset \mathbb{R}$, the associated order statistics are $L_{(1)} \leq \dots \leq L_{(S)}$. The indicator function of a set S is $\mathbf{1}_S$: $\mathbf{1}_S(s) = 1$ if $s \in S$, $\mathbf{1}_S(s) = 0$ otherwise. Let $m_{\mathbb{Q}}$ and $\sigma_{\mathbb{Q}}$ denote the expectation and the standard deviation of a random variable with distribution \mathbb{Q} . The skewness is defined as the standardized third moment $\zeta_{\mathbb{Q}} = \mathbb{E}_{x \sim \mathbb{Q}} \left[\left(\frac{x - m_{\mathbb{Q}}}{\sigma_{\mathbb{Q}}} \right)^3 \right]$. The excess kurtosis is the

standardized fourth moment $\kappa_{\mathbb{Q}} = \mathbb{E}_{x \sim \mathbb{Q}} \left[\left(\frac{x - m_{\mathbb{Q}}}{\sigma_{\mathbb{Q}}} \right)^4 \right] - 3$. The notation $b_n = \mathcal{O}(a_n)$ (resp. $b_n = o(a_n)$) means that $\left(\frac{b_n}{a_n} \right)_{n \in \mathbb{N}}$ is bounded (resp. $\lim_{n \rightarrow \infty} \frac{b_n}{a_n} = 0$). For random variables $X_n = \mathcal{O}_{a.s.}(a_n)$ (resp. $X_n = o_{a.s.}(a_n)$) means that $\left(\frac{X_n}{a_n} \right)_{n \in \mathbb{N}}$ is bounded (resp. converges to zero) almost surely. Φ refers to the cdf of the standard normal distribution: $\Phi(x) = \int_{-\infty}^x e^{-\frac{t^2}{2}} dt$. The uniform distribution on the interval $[a, b]$ is denoted by $\mathcal{U}[a, b]$. Let \mathcal{X} be a non-empty set. A function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is called kernel if there exists a feature map φ from \mathcal{X} to a Hilbert space \mathcal{H} such that $K(x, y) = \langle \varphi(x), \varphi(y) \rangle_{\mathcal{H}}$ for all $x, y \in \mathcal{X}$. While the feature map and the Hilbert space might not be unique one can always take the reproducing kernel Hilbert space (RKHS) \mathcal{H}_K associated to K . \mathcal{H}_K is a Hilbert space of $\mathcal{X} \rightarrow \mathbb{R}$ functions characterized by two properties: $K(x, \cdot) \in \mathcal{H}_K$ ($\forall x \in \mathcal{X}$) and $f(x) = \langle f, K(x, \cdot) \rangle_{\mathcal{H}_K}$ ($\forall x \in \mathcal{X}, f \in \mathcal{H}_K$).² The first property describes the basic elements of \mathcal{H}_K , the second one is called the reproducing property; combining the two properties makes the canonical feature map and feature space explicit: $K(x, y) = \langle \varphi(x), \varphi(y) \rangle_{\mathcal{H}_K}$ where $\varphi(x) = K(\cdot, x) \in \mathcal{H}_K$. The closed unit ball of \mathcal{H}_K is denoted by $\mathcal{B}_K = \left\{ f \in \mathcal{H}_K : \|f\|_{\mathcal{H}_K} \leq 1 \right\}$. Throughout the chapter the kernel K is assumed to be measurable, and $\mathcal{X} = \mathbb{R}$. Let $\mathcal{M}_1^+(\mathbb{R})$ denote the set of Borel probability measures on \mathbb{R} . A divergence is a mapping $D : \mathcal{M}_1^+(\mathbb{R}) \times \mathcal{M}_1^+(\mathbb{R}) \rightarrow \mathbb{R}^{\geq 0}$ measuring the discrepancy between two probability distributions \mathbb{P} and $\mathbb{Q} \in \mathcal{M}_1^+(\mathbb{R})$. For two vectors $\mathbf{a} \in \mathbb{R}^{d_1}$ and $\mathbf{b} \in \mathbb{R}^{d_2}$, $[\mathbf{a}; \mathbf{b}] \in \mathbb{R}^{d_1+d_2}$ denotes their concatenation.

Having introduced these notations, let us now define our divergences of interest.

- **Maximum Mean Discrepancy (MMD)** [SGSS07, GBR⁺12]: For a given kernel K with associated RKHS \mathcal{H}_K , let

$$\mu_K(\mathbb{P}) = \int_{\mathbb{R}} K(x, \cdot) d\mathbb{P}(x) \in \mathcal{H}_K$$

denote the mean embedding [BTA04, SGSS07] of the probability distribution $\mathbb{P} \in \mathcal{M}_1^+(\mathcal{X})$; the integral is meant in Bochner sense. The MMD of two distributions $\mathbb{P}, \mathbb{Q} \in \mathcal{M}_1^+(\mathcal{X})$ is a semi-metric defined by

$$\begin{aligned} \text{MMD}_K(\mathbb{P}, \mathbb{Q}) &= \|\mu_K(\mathbb{P}) - \mu_K(\mathbb{Q})\|_{\mathcal{H}_K} = \sup_{f \in \mathcal{B}_K} [\mathbb{E}_{x \sim \mathbb{P}} f(x) - \mathbb{E}_{y \sim \mathbb{Q}} f(y)], \\ &= \sqrt{\mathbb{E}_{x \sim \mathbb{P}} \mathbb{E}_{x' \sim \mathbb{P}} K(x, x') + \mathbb{E}_{y \sim \mathbb{Q}} \mathbb{E}_{y' \sim \mathbb{Q}} K(y, y') - 2\mathbb{E}_{x \sim \mathbb{P}} \mathbb{E}_{y \sim \mathbb{Q}} K(x, y)}, \end{aligned} \quad (5.2.1)$$

where the second form ($\sup_{f \in \mathcal{B}_K}$) encodes that the discrepancy of two probability distributions is measured by their maximal mean discrepancy over \mathcal{B}_K . It also shows that MMD belongs to the class of integral probability metrics [Zol83, Mül97].

The mean embedding and MMD are well-defined when $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} < \infty$. MMD is metric if and only if the kernel is characteristic [FGSS08, SGF⁺10]; examples of characteristic kernels include the Gaussian, Laplacian, Matérn, inverse multiquadrics or the B-spline kernel.

- **Kernel Stein Discrepancy (KSD)**: The Kernel Stein Discrepancy (KSD, [CSG16, LLJ16]) is defined for probability distributions admitting continuously differentiable pdfs. Assume that \mathbb{P} and \mathbb{Q} have pdf-s p and q , and we are given a kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Let the Stein operator be defined as $(S_q f)(x) : \mathbb{R} \rightarrow \mathbb{R}$, $(S_q f)(x) = \frac{[q(x)f(x)]'}{q(x)} = [\log q(x)]' f(x) + f'(x) = \langle f, \xi_q(x, \cdot) \rangle_K$ with

$$\xi_q(x, \cdot) = [\log q(x)]' K(x, \cdot) + \frac{\partial K(x, \cdot)}{\partial x}.$$

²The shorthand $K(\cdot, x)$ stands for the function $y \in \mathcal{X} \mapsto K(y, x) \in \mathbb{R}$ while keeping $x \in \mathcal{X}$ fixed.

KSD is defined as

$$\mathbb{T}_K(\mathbb{P}, \mathbb{Q}) = \sup_{f \in \mathcal{B}_K} [\mathbb{E}_{x \sim \mathbb{Q}}(S_q f)(x) - \mathbb{E}_{x \sim \mathbb{P}}(S_q f)(x)] = \sup_{f \in \mathcal{B}_K} \left\langle f, \underbrace{\mathbb{E}_{x \sim \mathbb{P}} \xi_q(x, \cdot)}_{=: g_{\mathbb{P}, \mathbb{Q}}} \right\rangle_{\mathcal{H}_K} = \|g_{\mathbb{P}, \mathbb{Q}}\|_{\mathcal{H}_K}, \quad (5.2.2)$$

where the first form shows the close resemblance to MMD (f is changed to $S_q f$), and the 2nd equality follows from the fact that $\mathbb{E}_{x \sim \mathbb{Q}}(S_q f)(x) = 0$ and by the symmetry of \mathcal{B}_K ($f \in \mathcal{B}_K \Rightarrow -f \in \mathcal{B}_K$). Let us define the Stein kernel

$$h_q(x, y) = s_q(x)s_q(y)K(x, y) + s_q(y)\partial_x K(x, y) + s_q(x)\partial_y K(x, y) + \frac{\partial^2 K(x, y)}{\partial x \partial y}. \quad (5.2.3)$$

If K is c_0 -universal [SFL11], $\mathbb{E}_{x \sim \mathbb{P}} \mathbb{E}_{y \sim \mathbb{P}} h_q(x, y) < \infty$, and $\mathbb{E}_{x \sim \mathbb{Q}} [(\log p(x))' - (\log q(x))']^2 < \infty$, $\mathbb{T}_K(\mathbb{P}, \mathbb{Q}) = 0$ iff $\mathbb{P} = \mathbb{Q}$. Moreover, the squared KSD can be reformulated in terms of expectations

$$\mathbb{T}_K^2(\mathbb{P}, \mathbb{Q}) = \mathbb{E}_{x \sim \mathbb{P}} \mathbb{E}_{y \sim \mathbb{P}} h_q(x, y) \quad (5.2.4)$$

similarly to MMD in (5.2.1).

- **Finite set Stein discrepancy (FSSD):** Let us assume that we are given J locations $\mathcal{V} = \{v_j\}_{j \in [J]} \subset \mathbb{R}$, consider the Stein witness function $g_{\mathbb{P}, \mathbb{Q}}$ defined in (5.2.2), and let $\mathcal{X} \subset \mathbb{R}$. By changing the RKHS norm to $\|\cdot\|_{L^2(\nu)}$ one gets the FSSD measure [JXS⁺17] of $\mathbb{P}, \mathbb{Q} \in \mathcal{M}_1^+(\mathcal{X})$:

$$\text{FSSD}_K(\mathbb{P}, \mathbb{Q}) = \|g\|_{L^2(\nu)} = \sqrt{\frac{1}{J} \sum_{j \in [J]} g_{\mathbb{P}, \mathbb{Q}}^2(v_j)}. \quad (5.2.5)$$

If the v_j locations are sampled from a measure ν which is absolutely continuous w.r.t. the Lebesgue measure, \mathcal{X} is a connected open set, k is real analytic, in addition to the standard KSD requirements, then it is known [JXS⁺17, Theorem 1] that for any $J \in \mathbb{N}^*$, ν -almost surely $\text{FSSD}_K(\mathbb{P}, \mathbb{Q}) = 0$ iff $\mathbb{P} = \mathbb{Q}$.

- **Wasserstein distance (WAD):** Let $p \geq 1$. The Wasserstein distance [PC19] of the probability measures $\mathbb{P}, \mathbb{Q} \in \mathcal{M}_1^+(\mathbb{R})$ is defined as

$$\begin{aligned} W_p(\mathbb{P}, \mathbb{Q}) &= \left[\inf_{\mathbb{J} \in C(\mathbb{P}, \mathbb{Q})} \int_{\mathbb{R}^2} |x - y|^p d\mathbb{J}(x, y) \right]^{1/p} = \left(\int_0^1 |F_{\mathbb{P}}^{-1}(t) - F_{\mathbb{Q}}^{-1}(t)|^p dt \right)^{1/p} \\ &= \left\| F_{\mathbb{P}}^{-1} - F_{\mathbb{Q}}^{-1} \right\|_{L^p([0, 1])}, \end{aligned} \quad (5.2.6)$$

where $C(\mathbb{P}, \mathbb{Q})$ denotes the set of all joint measures (so-called couplings) whose marginals are \mathbb{P} and \mathbb{Q} , $F_{\mathbb{P}}^{-1}$ and $F_{\mathbb{Q}}^{-1}$ are the inverse cdfs of \mathbb{P} and \mathbb{Q} , and $L^p([0, 1])$ refers to the real-valued p -power Lebesgue-integrable functions on $[0, 1]$.

Having defined the divergence measures to the minimum-divergence portfolio optimization (5.1.1), we now turn to their empirical estimation. In our application, one estimates the divergence between $\mathbb{P} = \mathbb{P}_{\mathbf{w}, N}$ (empirical portfolio returns) and $\mathbb{Q} = \mathbb{P}_T$ (explicit target distribution). Particularly, one has access to samples $\{\mathbf{r}_n\}_{n \in [N]} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_R$, and the **goal** is to solve

$$\arg \min_{\mathbf{w} \in \mathcal{W}} \hat{D}(\mathbb{P}_{\mathbf{w}, N}, \mathbb{P}_T),$$

where $\mathbb{P}_{\mathbf{w}, N}$ denotes the empirical measure associated to the samples $\{\mathbf{w}^\top \mathbf{r}_n\}_{n \in [N]} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_{\mathbf{w}}$, and \hat{D} is the estimated divergence with $D = \text{MMD}_K, \mathbb{T}_K, \text{FSSD}_K$ or W_p . The empirical versions of

Divergence D	Estimator \hat{D}
MMD (5.2.1)	$\widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) = \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{P}_T)(\mathbf{w}^\top \mathbf{r}_i)$
KSD (5.2.4)	$\hat{\text{T}}_{K,V}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) = \frac{1}{N^2} \sum_{i,j \in [N]} h_{\mathbb{P}_T}(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j)$
FSSD (5.2.5)	$\widehat{\text{FSSD}}_K^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) = \frac{2}{N(N-1)} \sum_{i < j} \Delta(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j)^a$
WAD (5.2.6)	$\widehat{\text{W}}_p(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) = \left[\frac{1}{N} \sum_{j \in [N]} \left (\mathbf{w}^\top \mathbf{r})_{(j)} - F_{\mathbb{P}_T}^{-1} \left(\frac{j}{N} \right) \right ^p \right]^{\frac{1}{p}}$

^a Δ is defined as $\Delta(a, b) = \tau(a)^\top \tau(b)$, $\tau(a) = [\xi_a(a, v_j)]_{j \in [J]} / \sqrt{J} \in \mathbb{R}^J$.

Table 5.1 – Semi-explicit MMD, KSD, FSD and WAD estimators.

Divergence D	Estimator \hat{D}	Complexity
MMD (5.2.1)	$\widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$	$\mathcal{O}(N^2)$
KSD (5.2.4)	$\hat{\text{T}}_{K,V}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$	$\mathcal{O}(N^2)$
FSSD (5.2.5)	$\widehat{\text{FSSD}}_K^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$	$\mathcal{O}(N)$
WAD (5.2.6)	$\widehat{\text{W}}_p(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$	$\mathcal{O}(N \log(N))$

Table 5.2 – Semi-explicit MMD, KSD, FSD and WAD estimators and their complexity.

(5.2.1), (5.2.4), (5.2.5) and (5.2.6) are summarized in Table 5.1, in the semi-explicit setting (i.e. having access to $\mu_K(\mathbb{P}_T)$, $h_{\mathbb{P}_T}$, Δ and $F_{\mathbb{P}_T}^{-1}$ respectively).

Let us consider the case of MMD in more detail. The discrete version of (5.2.1) lead to the U and V-statistics based estimators which we elaborate in the following.

Given i.i.d. samples $\{x_i\}_{i \in [N]} \sim \mathbb{P}$ and $\{y_i\}_{i \in [M]} \sim \mathbb{Q}$, one can estimate the squared MMD by using V- or U-statistics as

$$\begin{aligned} \widehat{\text{MMD}}_{K,V}^2(\mathbb{P}_N, \mathbb{Q}_M) &= \|\mu_K(\mathbb{P}_N) - \mu_K(\mathbb{Q}_M)\|_K^2 = \|\mu_K(\mathbb{P}_N)\|_{\mathcal{H}_K}^2 + \|\mu_K(\mathbb{Q}_M)\|_{\mathcal{H}_K}^2 \\ &\quad - 2 \langle \mu_K(\mathbb{P}_N), \mu_K(\mathbb{Q}_M) \rangle_{\mathcal{H}_K} \\ &= \frac{1}{N^2} \sum_{i,j \in [N]} K(x_i, x_j) + \frac{1}{M^2} \sum_{i,j \in [M]} K(y_i, y_j) - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} K(x_i, y_j), \end{aligned} \quad (5.2.7)$$

$$\begin{aligned} \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_N, \mathbb{Q}_M) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j) + \frac{1}{M(M-1)} \sum_{\substack{i,j \in [M] \\ i \neq j}} K(y_i, y_j) \\ &\quad - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} K(x_i, y_j), \end{aligned} \quad (5.2.8)$$

where $\mathbb{P}_N = \frac{1}{N} \sum_{n \in [N]} \delta_{x_n}$ and $\mathbb{Q}_M = \frac{1}{M} \sum_{m \in [M]} \delta_{y_m}$ denote the empirical measures. The estimator $\widehat{\text{MMD}}_{K,V}^2(\mathbb{P}_N, \mathbb{Q}_M)$ is non-negative, $\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_N, \mathbb{Q}_M)$ is unbiased; both have computational complexity $\mathcal{O}((N+M)^2)$. This gives rise in our context to the estimators

$$\widehat{\text{MMD}}_{K,V}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M}) = \frac{1}{N^2} \sum_{i,j \in [N]} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \frac{1}{M^2} \sum_{i,j \in [M]} K(t_i, t_j)$$

$$\begin{aligned}
& - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} K(\mathbf{w}^\top \mathbf{r}_i, t_j), \\
\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M}) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \frac{1}{M(M-1)} \sum_{\substack{i,j \in [M] \\ i \neq j}} K(t_i, t_j) \\
& - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} K(\mathbf{w}^\top \mathbf{r}_i, t_j). \tag{5.2.9}
\end{aligned}$$

If the mean embedding $\mu_K(\mathbb{Q})$ can be computed in closed-form, one can alternatively estimate the squared MMD using the plugin idea of (5.2.7), or that of (5.2.8) as

$$\begin{aligned}
\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}_N, \mathbb{Q}) &= \|\mu_K(\mathbb{P}_N) - \mu_K(\mathbb{Q})\|_K^2 \\
&= \|\mu_K(\mathbb{P}_N)\|_{\mathcal{H}_K}^2 + \|\mu_K(\mathbb{Q})\|_{\mathcal{H}_K}^2 - 2\langle \mu_K(\mathbb{P}_N), \mu_K(\mathbb{Q}) \rangle_{\mathcal{H}_K} \\
&= \frac{1}{N^2} \sum_{i,j \in [N]} K(x_i, x_j) + \mathbb{E}_{y \sim \mathbb{Q}} \mu_K(\mathbb{Q})(y) - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{Q})(x_i) \tag{5.2.10} \\
\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}_N, \mathbb{Q}) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j) + \mathbb{E}_{y \sim \mathbb{Q}} \mu_K(\mathbb{Q})(y) \\
& - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{Q})(x_i). \tag{5.2.11}
\end{aligned}$$

We call these estimators semi-explicit MMD estimators. Particularly, in our application with $x_i = \mathbf{w}^\top \mathbf{r}_i$ for the U-statistic variant (which we will use in our numerical experiments; we dropped the U subscript) this means

$$\begin{aligned}
\widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) \\
& - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{P}_T)(\mathbf{w}^\top \mathbf{r}_i). \tag{5.2.12}
\end{aligned}$$

There are multiple motivations to use the semi-explicit MMD estimators: (i) they reduce the computational time from $\mathcal{O}((N+M)^2)$ to $\mathcal{O}(N^2)$, and (ii) they give rise to better concentration properties (detailed in the next section).

To **optimize** the divergence objectives we tailor the cross-entropy method (CEM; [RK04]) to the task. Generally, the true (5.1.1) and similarly the estimated objective functions might not be convex. In order to tackle this challenge, we use CEM for the optimization

$$\max_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w}) := -\hat{D}(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T), \text{ or } \max_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w}) := -\hat{D}(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M}). \tag{5.2.13}$$

The CEM technique is a zero-order optimization approach constructing a sequence of pdfs $f(\cdot; \boldsymbol{\theta}^{(t)})$ which gradually concentrates around the optimum as $t \rightarrow \infty$. The idea of the CEM method is generating samples, followed by adaptively updating $f(\cdot; \boldsymbol{\theta}^{(t)})$ based on maximum likelihood estimate (MLE) relying on the top ρ -percent of the samples (elite in L -sense), and smoothing; for details see Alg. 1.

To generate portfolio weights in $\mathcal{W} = \mathcal{W}_1$, one can apply the normal distribution $\mathcal{N}(\mathbf{m}, \boldsymbol{\Sigma})$ in \mathbb{R}^{d-1} as a parametric distribution family in CEM. In this case $\boldsymbol{\theta} = (\mathbf{m}, \boldsymbol{\Sigma})$. In accordance with the constraint \mathcal{W}_1 , (i) CEM estimates $\hat{\mathbf{w}}_{-d} := [\hat{w}_j]_{j \in [d-1]}$, (ii) the final estimate is $\hat{\mathbf{w}} = [\hat{\mathbf{w}}_{-d}; 1 - \mathbf{1}_{d-1}^\top \hat{\mathbf{w}}_{-d}]$ and (iii) during the optimization the goodness of a sample \mathbf{w} is evaluated

Algorithm 1 Maximization of L with CEM

-
- 1: **Input:** Initial value $\boldsymbol{\theta}^{(0)} > \mathbf{0}$, quantile parameter $\rho > 0$, smoothing parameter $\beta \in (0, 1]$, sample size $S \in \mathbb{N}^*$.
 - 2: **for all** $t = 1, 2, \dots, T$ **do**
 - 3: Generate samples: $(\mathbf{w}_s)_{s \in [S]} \stackrel{\text{i.i.d.}}{\sim} f(\cdot; \boldsymbol{\theta}^{(t-1)})$.
 - 4: Evaluate performance: $L_s = L(\mathbf{w}_s)$, $s \in [S]$.
 - 5: Set level: $\gamma_t = L_{([\lceil (1-\rho)S \rceil])}$ // $(1-\rho)$ -quantile of $\{L(\mathbf{w}_s)\}_{s \in [S]}$.
 - 6: Estimate new parameter: $\tilde{\boldsymbol{\theta}}^{(t)} = \arg \max_{\boldsymbol{\theta} \in \Theta} \frac{1}{S} \sum_{s \in [S]} \mathbb{1}_{\{L(\mathbf{w}_s) \geq \gamma_t\}} \log [f(\mathbf{w}_s; \boldsymbol{\theta})]$ // MLE on the elite.
 - 7: Smoothing: $\boldsymbol{\theta}^{(t)} = (1-\beta)\boldsymbol{\theta}^{(t-1)} + \beta\tilde{\boldsymbol{\theta}}^{(t)}$.
 - 8: **Output:** $\hat{\mathbf{w}} = \mathbb{E}_{\mathbf{w} \sim f(\cdot; \boldsymbol{\theta}^{(T)})} \mathbf{w}$.
-

(Line 4) via $[\mathbf{w}; 1 - \mathbf{1}_{d-1}^\top \mathbf{w}] \in \mathbb{R}^d$. Line 6 in Alg. 1 takes the form of the empirical mean and covariance matrix of the elite:

$$\tilde{\mathbf{m}}^{(t)} = \frac{\sum_{s \in [S]} \mathbb{1}_{\{L(\mathbf{w}_s) \geq \gamma_t\}} \mathbf{w}_s}{\sum_{j \in [S]} \mathbb{1}_{\{L(\mathbf{w}_j) \geq \gamma_t\}}}, \quad \tilde{\boldsymbol{\Sigma}}^{(t)} = \frac{\sum_{s \in [S]} \mathbb{1}_{\{L(\mathbf{w}_s) \geq \gamma_t\}} (\mathbf{w}_s - \tilde{\mathbf{m}}^{(t)}) (\mathbf{w}_s - \tilde{\mathbf{m}}^{(t)})^\top}{\sum_{j \in [S]} \mathbb{1}_{\{L(\mathbf{w}_j) \geq \gamma_t\}}}.$$

Remark: In case of other constraints ($\mathcal{W} \neq \mathcal{W}_1$), one can similarly apply the CEM technique, by either relying on a parametric transformation $\ell : \mathbb{R}^{d-1} \rightarrow \mathcal{W}$ such that $\mathbf{w} = \ell(\mathbf{y})$, or by using distributions adapted directly to \mathcal{W} (for example, when $\mathcal{W} = \mathcal{W}_1^{\geq 0}$, the Dirichlet distribution can be applied).

5.3 Results

This section is dedicated to our theoretical results. In Section 5.3.1 we prove analytical mean embeddings for various kernel-distribution pairs; in Section 5.4 we will illustrate them numerically. We show improved concentration results for MMD estimators using this analytical knowledge and extend the analysis to unbounded kernels in Section 5.3.2.

5.3.1 Analytical formulas for mean embedding

In this section we show how the mean embedding can be computed in closed-form for the $(K, \mathbb{Q}) = (\text{Gaussian-exponentiated}, \text{Gaussian})$ and $(\text{Matérn}, \text{beta})$ kernel-distribution pairs, followed by a discussion to existing works. The proofs of our results are available in Section 5.A. The considered kernels are defined in Table 5.3 with their relation illustrated in Fig. 5.1. The target distributions are defined in Table 5.4; their relation is depicted in Fig. 5.2. The studied kernels generalize the widely-used Gaussian, Laplacian and exponential ones; the beta distribution extends the uniform one (in which case $\alpha = \beta = 1$).

Let us recall the U-statistic based MMD estimator from (5.2.12):

$$\begin{aligned} & \widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) \\ &= \underbrace{\frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j)}_{\text{can be estimated from } \{\mathbf{r}_n\}_{n \in [N]} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_R} + \underbrace{\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t)}_{\substack{\text{independent of } \mathbf{w}, \\ \text{it can be discarded}}} - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{P}_T)(\mathbf{w}^\top \mathbf{r}_i). \end{aligned}$$

The analytical knowledge of the mean embedding $\mu_K(\mathbb{P}_T)(x) = \mathbb{E}_{y \sim \mathbb{P}_T} K(x, y)$ can be leveraged in the 3rd term of this estimator, this is what we focus on in the sequel. Our results are as follows.

Kernel	$K(x, y)$	Parameters
Gaussian-exponentiated	$e^{-a(x-y)^2+bx}$	$a \geq 0, b \geq 0$
Matérn	$\sigma_0^2 e^{-\frac{\sqrt{2p+1} x-y }{\sigma}} \frac{p!}{(2p)!}$ $\times \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1} x-y }{\sigma}\right)^{p-i}$	$\sigma_0 \in \mathbb{R}^+, \sigma \in \mathbb{R}^+, p \in \mathbb{N}$
Gaussian	$e^{-\frac{(x-y)^2}{2c^2}}$	$c \in \mathbb{R}^+$
Laplacian	$e^{-\lambda x-y }$	$\lambda \in \mathbb{R}^+$
exponential	e^{bxy}	$b \in \mathbb{R}^+$

Table 5.3 – Kernel definitions; their relation is illustrated in Fig. 5.1.




Figure 5.1 – Relation of the kernels in Table 5.3.

Distribution	$q(x)$	Parameters
skew Gaussian	$2(2\pi v)^{-1/2} e^{-\frac{(x-m)^2}{2v}} \Phi\left(\frac{s(x-m)}{\sqrt{v}}\right)$	$s \in \mathbb{R}, m \in \mathbb{R}, v \in \mathbb{R}^+$
Gaussian	$\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-m)^2}{2\sigma^2}}$	$m \in \mathbb{R}, \sigma \in \mathbb{R}^+$
beta	$\frac{1}{B(\alpha,\beta)} x^{\alpha-1} (1-x)^{\beta-1} \mathbb{1}_{[0,1]}(x)$	$\alpha \in \mathbb{R}^+, \beta \in \mathbb{R}^+$
uniform	$\mathbb{1}_{[0,1]}(x)$	

Table 5.4 – Target distributions; their relation is illustrated in Fig. 5.2.




Figure 5.2 – Relation of the target distributions in Table 5.4.

Lemma 5.3.1 (Mean embedding: Gaussian target - Gaussian-exponentiated kernel). *Let the target distribution be Gaussian $q(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-m)^2}{2\sigma^2}}$, the kernel be Gaussian-exponentiated $K(x, y) = e^{-a(x-y)^2 + bxy}$ where $m \in \mathbb{R}$, $\sigma \in \mathbb{R}^+$, $a \geq 0, b \geq 0$. Then the mean embedding $\mu_K(\mathbb{Q})$ can be computed analytically as*

$$\mu_K(\mathbb{Q})(x) = \frac{e^{-\frac{a(x-m)^2}{1+2a\sigma^2} + \frac{2bmx + b(b+4a)\sigma^2 x^2}{2(1+2a\sigma^2)}}}{\sqrt{1+2a\sigma^2}}. \quad (5.3.1)$$

Lemma 5.3.2 (Mean embedding: beta target - Matérn kernel). *Let the target distribution be beta $q(x) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1} \mathbf{1}_{[0,1]}(x)$ with $\alpha \in \mathbb{R}^+$, $\beta \in \mathbb{R}^+$, and let the kernel be Matérn with half-integer ν ($\nu = p + \frac{1}{2}$, $p \in \mathbb{N}$), $\sigma_0 \in \mathbb{R}^+$, $\sigma \in \mathbb{R}^+$*

$$K(x, y) = \sigma_0^2 e^{-\frac{\sqrt{2p+1}|x-y|}{\sigma}} \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}|x-y|}{\sigma} \right)^{p-i}.$$

Then the mean embedding $\mu_K(\mathbb{Q})$ can be analytically computed as

$$\begin{aligned} \mu_K(\mathbb{Q})(x) &= \frac{\sigma_0^2}{B(\alpha, \beta)} \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}}{\sigma} \right)^{p-i} \times \\ &\quad \sum_{k=0}^{p-i} \binom{p-i}{k} x^k \left[(-1)^{p-i-k} e^{-\frac{\sqrt{2p+1}x}{\sigma}} E_1 \frac{\sqrt{2p+1}}{\sigma} ((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1) \right. \\ &\quad \left. + (-1)^k e^{\frac{\sqrt{2p+1}x}{\sigma}} E_2 \frac{\sqrt{2p+1}}{\sigma} ((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1) \right], \end{aligned}$$

where for $a, b > -1$, E_1^λ and E_2^λ are defined as

$$E_1^\lambda(z, a, b) = \int_0^z y^a (1-y)^b e^{\lambda y} dy, \quad E_2^\lambda(z, a, b) = \int_z^1 y^a (1-y)^b e^{-\lambda y} dy$$

and can be evaluated using Lemma 5.B.1.

Table 5.5 summarizes our results, which we complement with a few **remarks**:

- Relation to previous work:
 - The mean embedding of the Gaussian distribution with Gaussian kernel is [SZS⁺08]

$$\mu_K(\mathbb{Q})(x) = \frac{c}{\sqrt{\sigma^2 + c^2}} e^{-\frac{(x-m)^2}{2(\sigma^2 + c^2)}}, \quad (5.3.2)$$

which follows as a specific case of Lemma 5.3.1 by choosing $b = 0$, $a = \frac{1}{2c^2}$, and noticing that $\frac{1}{1+2a\sigma^2} = \frac{1}{1+\frac{\sigma^2}{c^2}} = \frac{c^2}{\sigma^2 + c^2}$ and $\frac{a}{1+2a\sigma^2} = \frac{1}{2(\sigma^2 + c^2)}$.

- The mean-embedding for the skew Gaussian distribution with the Gaussian kernel $K(x, y) = e^{-a(x-y)^2}$ is [Ken98, Section 9.2]

$$\mu_K(\mathbb{Q})(x) = \frac{2}{\sqrt{1+2av}} \Phi \left(\frac{2as\sqrt{v}(x-m)}{(1+2av)\sqrt{1+2av+s^2}} \right) e^{-\frac{a(x-m)^2}{1+2av}}. \quad (5.3.3)$$

Choosing $s = 0$ and $v = \sigma^2$ in the skew Gaussian distribution, and reparameterizing the Gaussian kernel as $a = \frac{1}{2c^2}$, again gives the mean embedding for the Gaussian target and Gaussian kernel.

\mathbb{Q}	K	$\mu_K(\mathbb{Q})$
Gaussian	Gaussian-exponentiated	Lemma 5.3.1
beta	Matérn	Lemma 5.3.2
Gaussian	Gaussian	Lemma 5.3.1 \Rightarrow (5.3.2): [SZS ⁺ 08] \Leftarrow [Ken98]
skew Gaussian	Gaussian	(5.3.3): [Ken98]
beta	Laplacian	Lemma 5.3.2 \Rightarrow (5.3.4)

Table 5.5 – Summary of obtained analytical mean embedding results.

- Specific case of beta target - Laplacian kernel: The Matérn kernel specializes to the Laplacian one with $\sigma_0 = 1$ and $\sigma = \frac{1}{\lambda}$ (see Fig. 5.1). Hence, Lemma 5.3.2 implies that the mean embedding of the beta target ($\alpha \in \mathbb{R}^+$, $\beta \in \mathbb{R}^+$) and Laplacian kernel ($\lambda \in \mathbb{R}^+$) is

$$\mu_K(\mathbb{Q})(x) = \frac{1}{B(\alpha, \beta)} \left[e^{-\lambda x} E_1^\lambda((0 \vee x) \wedge 1, \alpha - 1, \beta - 1) + e^{\lambda x} E_2^\lambda((0 \vee x) \wedge 1, \alpha - 1, \beta - 1) \right], \quad (5.3.4)$$

where E_1^λ and E_2^λ can be evaluated using Lemma 5.B.1.

In the next section we show how the analytic knowledge of the mean embedding results in better concentration properties of the MMD estimators.

5.3.2 Concentration of semi-explicit MMD

In this section, we show that explicit mean embedding, in case of both bounded and unbounded kernels, leads to better concentration properties of the MMD estimator; the proofs are detailed in Section 5.A. We start by recalling the concentration of the classical U-statistic based MMD estimator for bounded kernels (Theorem 5.1), followed by presenting our result (Theorem 5.2) for the semi-explicit MMD. We generalize the statement to unbounded kernels in Theorem 5.3 and show lower bounds in Theorem 5.4.

Theorem 5.1 (Concentration of $\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}, \mathbb{Q})$, bounded kernel [GBR⁺12, Theorem 10]). *Assume that $0 \leq K(x, x') \leq B$ for all (x, x') , and let $\epsilon > 0$. Then*

$$\mathbb{P} \left(\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \epsilon \right) \leq e^{-\frac{\lfloor \frac{N}{2} \rfloor \epsilon^2}{8B^2}},$$

and the same bound holds for the deviation of $-\epsilon$ below.

Using the analytical knowledge of $\mu_K(\mathbb{Q})$ leads to tighter concentration properties as it is shown by our next result.

Theorem 5.2 (Concentration of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$, bounded kernel). *Assume that $A \leq K(x, x') \leq B$ for all (x, x') , and let $\epsilon > 0$. Then*

$$\mathbb{P} \left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \epsilon \right) \leq e^{-\frac{\lfloor \frac{N}{2} \rfloor \epsilon^2}{2(B-A)^2}} + e^{-\frac{N\epsilon^2}{8(B-A)^2}}$$

and the same bound holds for the deviation below with $-\epsilon$.

Remarks:

- The proof relies on rewriting the difference $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$ as a sum of two U-statistics, followed by applying twice the Hoeffding inequality for U-statistics and union bounding.

- Compared to the bound in Theorem 5.1 for $\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}, \mathbb{Q})$ with $A = 0$, we gain in terms of constant in front of $N\varepsilon^2$ in the exponent: we have $\frac{1}{4B^2}$ and $\frac{1}{8B^2}$ instead of $\frac{1}{16B^2}$. This means that the estimator using the analytical knowledge of $\mu_K(\mathbb{Q})$ brings a factor of 2 improvement in the exponent.

In Theorem 5.2 the deviation of the estimator $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$ was captured for bounded kernels. Our next theorem extends the result to the unbounded exponential kernel.

Theorem 5.3 (Concentration of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$, exponential kernel). *Let us consider the exponential kernel $K(x, y) = e^{bxy}$ ($b > 0$) with probability measures \mathbb{P} and \mathbb{Q} satisfying*

$$\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} < \infty \quad \text{and} \quad \mathbb{E}_{x \sim \mathbb{Q}} e^{\lambda x^2} < \infty \quad \forall \lambda \in \mathbb{R}. \quad (5.3.5)$$

Let the number of samples N taken from \mathbb{P} be even. Then for any $p \geq 2$, there exists a universal constant $C = C_{p,\mathbb{P},K} > 0$ such that for any $\varepsilon > 0$

$$\mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon\right) \leq \frac{C}{\varepsilon^p N^{p/2}},$$

and the same bound holds for the deviation below with $-\varepsilon$.

Remarks:

- In Theorem 5.3, we showed convergence guarantees for the semi-explicit MMD estimator with the exponential kernel. The proof relies on combining concentration results for U-statistics and martingales. One could use similar ideas to cover the two-sample MMD estimator $\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}, \mathbb{Q})$ for the exponential kernel.
- **Assumption (5.3.5):** This condition holds for instance for generalized normal distributions (see Table 5.10) with parameter $\gamma > 2$. Indeed, for $\gamma = 2 + \kappa$, $\kappa > 0$, the pdf of the generalized normal distribution with parameters α , β and γ is proportional to $e^{-\frac{1}{2} \frac{|x-\alpha|^{2+\kappa}}{\beta}}$. Hence, $\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} < \infty$ for any $\lambda \in \mathbb{R}$ and (5.3.5) is satisfied. Indeed,
 - $\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} \leq 1$ for any $\lambda \leq 0$ as for $\lambda = 0$ one has $\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} = 1$ and $\lambda \in (-\infty, 0] \mapsto \mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2}$ is increasing in λ .
 - For $\lambda > 0$, for x large enough, there exists a constant $c > 0$ such that $\lambda x^2 - \frac{1}{2} \frac{|x-\alpha|^{2+\kappa}}{\beta} \leq -c|x|^{2+\kappa}$. Taking the exponential, this quantity is integrable so

$$\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} \propto \int_{\mathbb{R}} e^{\lambda x^2 - \frac{1}{2} \frac{|x-\alpha|^{2+\kappa}}{\beta}} dx \leq \int_{\mathbb{R}} e^{-c|x|^{2+\kappa}} dx < \infty.$$

- **Convergence of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$:** Theorem 5.2 means the $\mathcal{O}_{a.s.}\left(\frac{1}{\sqrt{N}}\right)$ convergence of the estimator $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$ for bounded kernels. Theorem 5.3 implies the same convergence for the unbounded exponential kernel. Indeed, for any $\kappa > 0$, one can find p such that $\kappa p > 2$. Taking $\varepsilon_N = \left(\frac{1}{\sqrt{N}}\right)^{1-\kappa}$ in the Borel-Cantelli lemma, using Theorem 5.3 and that in this case $\frac{1}{\varepsilon_N^p N^{p/2}} = \frac{N^{\frac{1}{2}(1-\kappa)p}}{N^{p/2}} = N^{-\frac{\kappa p}{2}}$, one arrives at

$$\sum_{N \in \mathbb{N}^*} \mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon_N\right) \leq \sum_{N \in \mathbb{N}^*} \frac{C_p}{N^{\frac{\kappa p}{2}}} < \infty.$$

- **Convergence of $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q})$:** In order to understand the convergence behaviour of $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q})$, let us rewrite $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$ in terms of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$. By the definition of $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q})$ and

$\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}_N, \mathbb{Q})$ [see (5.2.10)-(5.2.11)], the two estimators only differ in their first terms which we denote as

$$T_1^V := \frac{1}{N^2} \sum_{i,j \in [N]} K(x_i, x_j), \quad T_1^U := \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j).$$

These two terms are closely related; let us write T_1^V in terms of T_1^U

$$\begin{aligned} T_1^V &= \frac{1}{N^2} \left(\sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j) + \sum_{i \in [N]} K(x_i, x_i) \right) \\ &= \underbrace{\frac{N(N-1)}{N^2}}_{1 - \frac{1}{N}} \underbrace{\left(\frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j) \right)}_{T_1^U} + \frac{1}{N^2} \sum_{i \in [N]} K(x_i, x_i) \end{aligned}$$

which means that $T_1^V = (1 - \frac{1}{N}) T_1^U + \frac{1}{N^2} \sum_{i \in [N]} K(x_i, x_i)$. Denoting the the second and third common terms of $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q})$ and $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$ by $T_2 := \mathbb{E}_{y \sim \mathbb{Q}} \mu_K(\mathbb{Q})(y)$ and $T_3 := -\frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{Q})(x_i)$, we hence have

$$\begin{aligned} &\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \\ &= T_1^V + T_2 + T_3 - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \\ &= \left(1 - \frac{1}{N}\right) T_1^U + T_2 + T_3 - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) + \frac{1}{N^2} \sum_{i \in [N]} K(x_i, x_i) \\ &= \left(1 - \frac{1}{N}\right) \underbrace{[T_1^U + T_2 + T_3 - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})]}_{\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})} + \frac{T_2}{N} + \frac{T_3}{N} - \frac{\text{MMD}_K^2(\mathbb{P}, \mathbb{Q})}{N} + \frac{1}{N^2} \sum_{i \in [N]} K(x_i, x_i). \end{aligned}$$

This implies that

$$\begin{aligned} &\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \\ &= \left(1 - \frac{1}{N}\right) \left[\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \right] + o_{a.s.} \left(\frac{1}{\sqrt{N}} \right) \end{aligned} \quad (5.3.6)$$

as T_2 and T_3 are constants, and $\frac{1}{N} \sum_{i \in [N]} K(x_i, x_i)$ converge to a constant by the law of large numbers.

Since $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) = \mathcal{O}_{a.s.} \left(\frac{1}{\sqrt{N}} \right)$, (5.3.6) means that $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) = \mathcal{O}_{a.s.} \left(\frac{1}{\sqrt{N}} \right)$ also holds.

- **Convergence of $\widehat{\text{MMD}}_{K,e,V}(\mathbb{P}, \mathbb{Q})$:** Considering the estimator without square, since

$$\begin{aligned} \mathcal{O}_{a.s.} \left(\frac{1}{\sqrt{N}} \right) &\stackrel{(*)}{=} \left| \widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \right| \\ &= \left| \widehat{\text{MMD}}_{K,e,V}(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K(\mathbb{P}, \mathbb{Q}) \right| \underbrace{\left[\widehat{\text{MMD}}_{K,e,V}(\mathbb{P}, \mathbb{Q}) + \text{MMD}_K(\mathbb{P}, \mathbb{Q}) \right]}_{\substack{\geq 0 \text{ from (5.2.10)} \\ \geq 0}} \end{aligned}$$

$$\geq \left| \widehat{\text{MMD}}_{K,e,V}(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K(\mathbb{P}, \mathbb{Q}) \right| \text{MMD}_K(\mathbb{P}, \mathbb{Q})$$

one gets that $\left| \widehat{\text{MMD}}_{K,e,V}(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K(\mathbb{P}, \mathbb{Q}) \right| = \mathcal{O}_{a.s.} \left(\frac{1}{\sqrt{N}} \right)$ by using in (*) the previously established convergence $\widehat{\text{MMD}}_{K,e,V}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) = \mathcal{O}_{a.s.} \left(\frac{1}{\sqrt{N}} \right)$.

It is known [TSS16, Theorem 2] that the rate $\frac{1}{\sqrt{N}}$ for bounded continuous radial kernels of the form $K(x, y) = \int_0^\infty e^{-t(x-y)^2} d\nu(t)$ with ν bounded non-negative measure is optimal in the two-sample setting for the class of probability measures with infinitely differentiable density. We prove that a similar result holds for the considered one-sample setting and unbounded exponential kernel.

Theorem 5.4 (Minimax rate for one-sample MMD estimators, exponential kernel). *Let us consider the exponential kernel $K(x, y) = e^{bxy}$ ($b > 0$). Let \mathcal{P} be the set of all Borel probability measures \mathbb{P} on \mathbb{R} for which $\mu_k(\mathbb{P})$ is well-defined, in other words $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} < \infty$ for all $\mathbb{P} \in \mathcal{P}$. Let $\{x_n\}_{n \in [N]} \stackrel{i.i.d.}{\sim} \mathbb{P}$. Then*

$$\inf_{\widehat{\text{MMD}}_N} \sup_{\mathbb{P}, \mathbb{Q} \in \mathcal{P}} \mathbb{P} \left(\left| \widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}, \mathbb{Q}) \right| \geq \frac{c}{\sqrt{N}} \right) \geq \max \left(\frac{e^{-\frac{a^2 b}{2}}}{4}, \frac{1 - \sqrt{\frac{a^2 b}{2}}}{2} \right)$$

for some finite constant $c > 0$, arbitrary constant $a > 0$, and $\widehat{\text{MMD}}_N$ runs over all the estimators using the samples $\{x_n\}_{n \in [N]}$.

Remarks:

- Theorem 5.4 shows that $\text{MMD}_K(\mathbb{P}, \mathbb{Q})$ with K being the exponential kernel cannot be estimated at a rate faster than $\frac{1}{\sqrt{N}}$ by any $\widehat{\text{MMD}}_N$ estimator for all $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$. The fact that the rate $\frac{1}{\sqrt{N}}$ is achievable was shown in Theorem 5.3.
- The proof relies on the Le Cam's method (as [TSS16]). The main technical difference and challenge which were resolved are that using the unbounded exponential kernel one requires a dedicated MMD computation (resolved separately in Lemma 5.A.2), and with this need the parameter dependence of MMD becomes more intricate.
- The condition $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} < \infty$ appearing in the definition of \mathcal{P} can only be milder than (5.3.5), since the former is a specific case of (5.3.5) with $\lambda = \frac{b}{2}$. In fact, (5.3.5) is more restrictive as it can be seen for instance for Gaussian distributions $\mathbb{P} = \mathcal{N}(m, \sigma^2)$. Indeed, in this case a standard calculation shows that $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} = e^{\frac{m^2}{1-b\sigma^2}} \frac{1}{\sqrt{1-\sigma^2 b}}$ which is finite (or equivalently $\mathbb{P} \in \mathcal{P}$) iff $\sigma^2 < 1/b$. However, $\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} \propto \frac{1}{\sqrt{1-2\lambda\sigma^2}}$ which is finite iff $\lambda < \frac{1}{2\sigma^2}$; in other words the Gaussian distributions $\mathcal{N}(m, \sigma^2)$ do not obey (5.3.5). This suggests that Theorem 5.3 might also hold under milder conditions.

In this section, we have shown the theoretical advantages of using explicit mean embedding when available. In the next sections, we demonstrate the improved convergence of the semi-explicit MMD on simulated (Section 5.4) and on real-world (Section 5.5) portfolio benchmarks.

5.4 Experiments on simulated data

In this section we demonstrate the efficiency of the investigated portfolio optimization scheme on simulated data. We designed **two sets** of experiments:

1. Advantage of semi-explicit MMD: Recall that our objective (5.2.13) can be written as divergence minimization over the portfolio weights

$$\min_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w}) := \hat{D}(\mathbb{P}_{\mathbf{w}, N}, \mathbb{P}_T), \text{ or } \min_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w}) := \hat{D}(\mathbb{P}_{\mathbf{w}, N}, \mathbb{P}_{T, M}).$$

\mathbb{P}_T	K	$\mu_K(\mathbb{P}_T)$
Gaussian	Gaussian	(5.3.2)
Gaussian	exponential	(5.3.1)
beta	Laplacian	(5.3.4)

Table 5.6 – Summary of the investigated target distribution (\mathbb{P}_T) - kernel (K) pairs, with the associated analytical mean embeddings.

In the first three experiments (Section 5.4.1), we investigate the practical benefits of the explicit knowledge of the mean embedding on the objective L . Our goal is two-fold: to numerically illustrate the improved concentration properties (Theorem 5.2) of the semi-explicit MMD estimators on the objective L ; the experiments also enable one to focus on the best performing (target distribution, kernel) pairs on the real-world benchmarks (Section 5.5). The investigated (target distribution, kernel) pairs with the underlying explicit mean embeddings are summarized in Table 5.6. The rationale behind these selected pairs are as follows:

- the Gaussian-Gaussian case corresponds to a standard choice,
- the exponential kernel is a popular example of unbounded kernels,
- the first two experiments relying on Gaussian target (and hence unbounded support) is complemented with the beta target distribution having bounded support.

In addition to MMD, in our first experiment we also include the KSD measure for illustration. The experiments are well-specified in the sense that there exists an optimal \mathbf{w}^* such that $D(\mathbb{P}_{\mathbf{w}^*}, \mathbb{P}_T) = 0$.⁴

2. Misspecified setting: The goal of this experiment (Section 5.4.2) is to go one step further, and study the impact of misspecification, in other words when $D(\mathbb{P}_{\mathbf{w}^*}, \mathbb{P}_T) \neq 0$.

In these experiments, we will consider two-dimensional returns ($d = 2$) with budget constraints $\mathcal{W} = \mathcal{W}_1$ defined in (5.1.2).

5.4.1 Advantage of semi-explicit MMD

This section is dedicated to the illustration of the advantage of the closed-form knowledge of mean embedding. We will consider two different settings: in the first case the return $R \in \mathbb{R}^2$ will be normally distributed, in the second case R will consist of beta random variables. In both cases, the components of the random variable R are independent, either normal or beta, and R can be specified by its mean and its covariance. In order to guarantee that the problem is well-specified, we assume that there exists an optimal portfolio $\mathbf{w}^* \in \mathcal{W}$, and the target distribution \mathbb{P}_T is chosen to be equal to the distribution of $(\mathbf{w}^*)^\top R \in \mathbb{R}$, in other words $\mathbb{P}_T = \mathbb{P}_{\mathbf{w}^*}$. Given \mathbf{w}^* and R , the resulting distribution of the portfolio return $(\mathbf{w}^*)^\top R$ is normal in the first case, linear combination of beta in the second case. Denoting by \mathbf{m}_R and Σ_R the mean and covariance of R , the mean and the variance of the $(\mathbf{w}^*)^\top R$ distribution can be computed as

$$m = \mathbf{w}^{*\top} \mathbf{m}_R \quad \text{and} \quad \sigma^2 = \mathbf{w}^{*\top} \Sigma_R \mathbf{w}^*. \quad (5.4.1)$$

In our experiments, we chose

$$\mathbf{m}_R = \begin{bmatrix} 0.02 \\ 0.04 \end{bmatrix}, \quad \Sigma_R = \begin{bmatrix} 0.01 & 0 \\ 0 & 0.03 \end{bmatrix}, \quad \mathbf{w}^* = \begin{bmatrix} 0.2 \\ 0.8 \end{bmatrix}.$$

Gaussian returns, Gaussian kernel: In this experiment, we assume that the distribution of $R \in \mathbb{R}^2$ is Gaussian, and we consider the Gaussian kernel $K(x, y) = e^{-\frac{(x-y)^2}{2c^2}}$ with $c > 0$. Let us detail the estimators used: the two-sample MMD, the semi-explicit MMD and the KSD one.

- Two-sample MMD: Let us recall the formula of the two-sample MMD estimator (5.2.11):

$$\begin{aligned} \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M}) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \frac{1}{M(M-1)} \sum_{\substack{i,j \in [M] \\ i \neq j}} K(t_i, t_j) \\ &\quad - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} K(\mathbf{w}^\top \mathbf{r}_i, t_j). \end{aligned}$$

Replacing K by its analytical form, one obtains that

$$\begin{aligned} \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M}) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} e^{-\frac{(\mathbf{w}^\top \mathbf{r}_i - \mathbf{w}^\top \mathbf{r}_j)^2}{2c^2}} + \frac{1}{M(M-1)} \sum_{\substack{i,j \in [M] \\ i \neq j}} e^{-\frac{(t_i - t_j)^2}{2c^2}} \\ &\quad - \frac{2}{NM} \sum_{i \in [N]} \sum_{j \in [M]} e^{-\frac{(\mathbf{w}^\top \mathbf{r}_i - t_j)^2}{2c^2}}. \end{aligned}$$

- Semi-explicit MMD: The semi-explicit MMD estimator (5.2.12) is recalled here:

$$\begin{aligned} \widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(\mathbf{w}^\top \mathbf{r}_i, \mathbf{w}^\top \mathbf{r}_j) + \mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) \\ &\quad - \frac{2}{N} \sum_{i \in [N]} \mu_K(\mathbb{P}_T)(\mathbf{w}^\top \mathbf{r}_i). \end{aligned}$$

We can replace K by its analytical value and from (5.3.2) it known that $\mu_K(\mathbb{P}_T)(x) = \frac{c}{\sqrt{\sigma^2 + c^2}} e^{-\frac{(x-m)^2}{2(\sigma^2 + c^2)}}$. One can obtain the second term of the squared MMD $\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t)$ via standard integration. Indeed, since \mathbb{P}_T is normal with mean m and standard deviation σ , by definition $\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) = \int_{\mathbb{R}} \mu_K(\mathbb{P}_T)(t) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(t-m)^2}{2\sigma^2}} dt$ which can be simplified by recognizing the integral of a normal distribution³:

$$\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) = \frac{c}{\sqrt{c^2 + 2\sigma^2}}.$$

The explicit MMD then takes the form

$$\begin{aligned} \widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} e^{-\frac{(\mathbf{w}^\top \mathbf{r}_i - \mathbf{w}^\top \mathbf{r}_j)^2}{2c^2}} + \frac{c}{\sqrt{c^2 + 2\sigma^2}} \\ &\quad - \frac{2}{N} \sum_{i \in [N]} \frac{c}{\sqrt{\sigma^2 + c^2}} e^{-\frac{(\mathbf{w}^\top \mathbf{r}_i - m)^2}{2(\sigma^2 + c^2)}}. \end{aligned}$$

³This formula can be obtained by noticing that $-\frac{(t-m)^2}{2(\sigma^2 + c^2)} - \frac{(t-m)^2}{2\sigma^2} = -\frac{(t-m)^2(c^2 + 2\sigma^2)}{2\sigma^2(\sigma^2 + c^2)}$. Denoting by $s^2 = \frac{\sigma^2(\sigma^2 + c^2)}{c^2 + 2\sigma^2}$, one gets that $\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) = \frac{c}{\sqrt{\sigma^2 + c^2}} \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(t-m)^2}{2s^2}} dt = \frac{c}{\sqrt{c^2 + 2\sigma^2}} \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi s^2}} e^{-\frac{(t-m)^2}{2s^2}} dt = \frac{c}{\sqrt{c^2 + 2\sigma^2}}$.

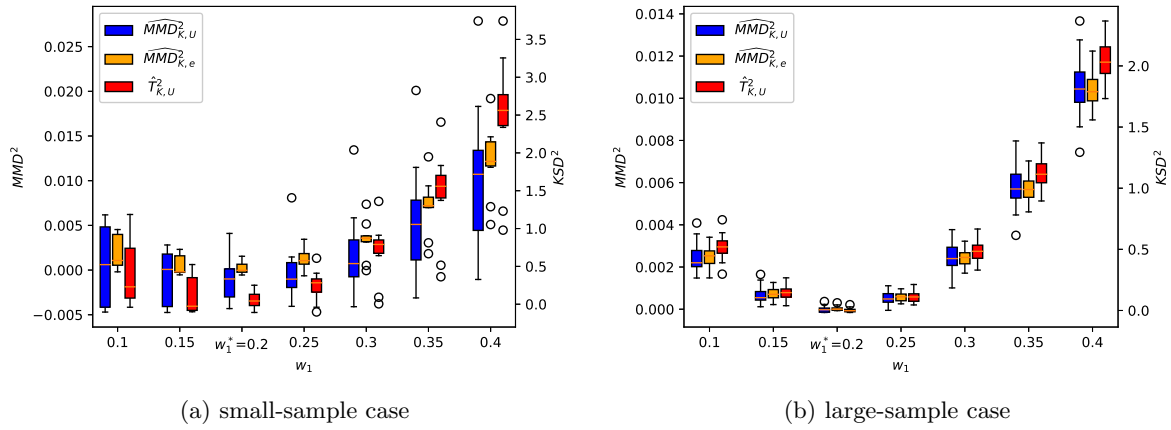


Figure 5.3 – Boxplots of the objective values for Gaussian returns (\mathbb{P}_R) and Gaussian kernel (K). Plotted: $w_1 \mapsto \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M})$, $\widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$, $\widehat{T}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$, with portfolio weight $\mathbf{w} = [w_1; 1 - w_1]$. The optimal portfolio weight is $\mathbf{w}^* = [w_1^*; 1 - w_1^*]$ with $w_1^* = 0.2$. Left: $N = 100$ samples. Right: $N = 5 \times 10^3$ samples.

- KSD: The KSD estimator can be computed using the partial derivatives of K (given in Table 5.11), the score function of \mathbb{P}_T in the h_q formula (5.2.3), followed by the form of $\widehat{T}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$ in (5.2.4); see Lemma 5.A.1 for the detailed computation. The resulting estimator is

$$\begin{aligned} & \widehat{T}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T) \\ &= \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} \left[\frac{(\mathbf{w}^\top \mathbf{r}_i - m)(\mathbf{w}^\top \mathbf{r}_j - m)}{\sigma^4} - \frac{(\mathbf{w}^\top \mathbf{r}_i - \mathbf{w}^\top \mathbf{r}_j)^2 (\sigma^2 + c^2)}{\sigma^2 c^4} + \frac{1}{c^2} \right] e^{-\frac{(\mathbf{w}^\top \mathbf{r}_i - \mathbf{w}^\top \mathbf{r}_j)^2}{2c^2}}. \end{aligned}$$

We illustrate the evolution of the various objectives as a function of $\mathbf{w} = [w_1; 1 - w_1]$ by varying the value of w_1 in the set $\{0.1, 0.15, 0.2, \dots, 0.4\}$. We consider two different situations: one is the large-sample case ($N = M = 5000$), the other is the small-sample case ($N = M = 100$). We evaluate the objectives on samples $\{\mathbf{r}_i\}_{i \in [N]} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{m}_R, \Sigma_R)$ and $\{t_i\}_{i \in [M]} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(m, \sigma^2)$ where m and σ^2 are defined in (5.4.1). The kernel bandwidth c was taken to be equal to the standard deviation of $\{\mathbf{w}^\top \mathbf{r}_i\}_{i \in [N]}$ for each $\mathbf{w} = [w_1; 1 - w_1]$. In order to investigate the robustness of the estimators, we report the median values, quartiles and extreme values of the estimators as a function of the weight w_1 using 15 repetitions. The corresponding summary statistics are available in Fig. 5.3. As the figure shows the estimated MMD and Stein objectives follow a U-shape as a function of w_1 and take their minimum at $w_1^* = 0.2$. Notice that the scale on the y-axis is different for the MMD estimators (l.h.s.) and for the Stein estimator (r.h.s.); the latter increases more rapidly as we move away from $w_1^* = 0.2$. In accordance with Theorem 5.2 the semi-explicit MMD estimator shows improved concentration properties, and this behaviour is particularly emphasized in the small-sample regime.

Gaussian returns, exponential kernel: In this experiment, we follow the same assumption of Gaussian returns, but we consider the unbounded exponential kernel $K(x, y) = e^{bxy}$ with $b > 0$. The two-sample MMD estimator is calculated via (5.2.11). The semi-explicit MMD objective (5.2.12) is obtained by using the explicit mean embedding $\mu_K(\mathbb{P}_T)(x)$ given in Lemma 5.3.1 and by the relation

$$\mathbb{E}_{t \sim \mathbb{P}_T} \mu_K(\mathbb{P}_T)(t) = \frac{e^{\frac{bm^2}{1-b\sigma^2}}}{\sqrt{1 - (b\sigma^2)^2}},$$

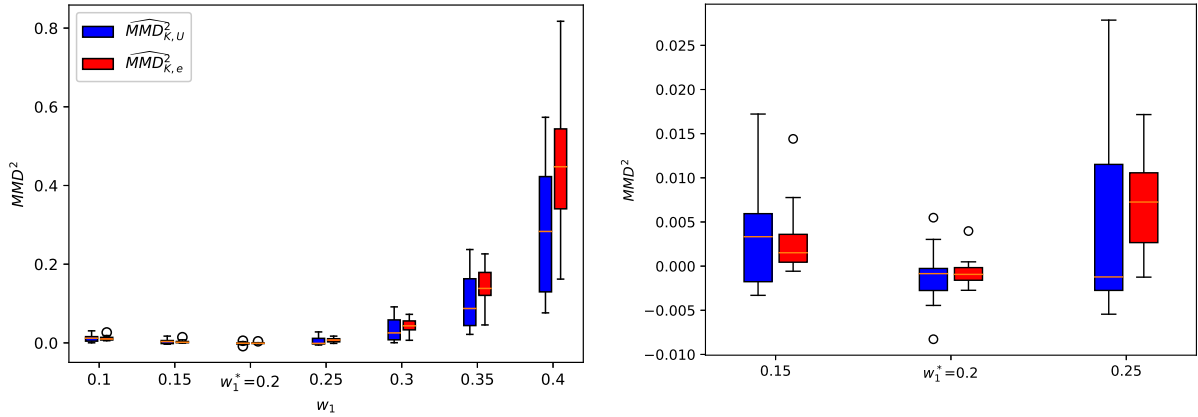


Figure 5.4 – Boxplots of the objective values for Gaussian return (\mathbb{P}_R) and exponential kernel (K). Plotted: $w_1 \mapsto \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M})$ and $w_1 \mapsto \widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$, with portfolio weight $\mathbf{w} = [w_1; 1 - w_1]$. The optimal portfolio weight is $\mathbf{w}^* = [w_1^*; 1 - w_1^*]$ with $w_1^* = 0.2$. Right: left zoomed in.




Figure 5.5 – Comparison of the histogram of the portfolio returns $(\mathbf{w}^*)^\top R$ (with R having beta components) and one-dimensional beta target distribution with the same mean and variance on 10^4 samples. Right: left zoomed in.

which is derived later in (5.A.3). We chose $b = \frac{1}{2\hat{\sigma}^2}$ with $\hat{\sigma}$ being the standard deviation of $\{\mathbf{w}^\top \mathbf{r}_i\}_{i \in [N]}$ for each $\mathbf{w} = [w_1; 1 - w_1]$. The evaluation is made on $N = M = 10^4$ samples, with 50 repetitions for each w_1 .

The resulting median values, quartiles and extreme values of the estimators are summarized in Fig. 5.4. The figure shows that (i) again the semi-explicit MMD estimator has better concentration properties than the two-sample MMD one, (ii) the U-shape of the objective values is less pronounced with the exponential kernel, (iii) by zooming in to the estimated objectives curves they still take their minimum around the optimal weight $w_1^* = 0.2$.

beta returns, Laplacian kernel: In this experiment the return variable $R \in \mathbb{R}^2$ has components with beta distribution, and the kernel is the Laplacian $K(x, y) = e^{-\lambda|x-y|}$ with $\lambda > 0$. The parameters of the beta target distribution are set by moment matching (recalled in Table 5.12) applied to (5.4.1).⁴

In the $\widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M})$ and semi-explicit $\widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$ estimators [(5.2.9) and (5.2.12)], we discard the 2nd terms which are independent of \mathbf{w} ; hence the objective values can

⁴Notice that in this case, we are *not* in the well-specified setting since the sum of two beta distributions is not a beta distribution, but the approximation resulting from moment matching is quite accurate; see Fig. 5.5 for an illustration.




Figure 5.6 – Boxplots of the objective values for beta return (\mathbb{P}_R) and Laplacian kernel (K). Plotted: $w_1 \mapsto \widehat{\text{MMD}}_{K,U}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_{T,M})$ and $w_1 \mapsto \widehat{\text{MMD}}_{K,e}^2(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$, with portfolio weight $\mathbf{w} = [w_1; 1 - w_1]$. The optimal portfolio weight is $\mathbf{w}^* = [w_1^*; 1 - w_1^*]$ with $w_1^* = 0.2$.

easily take negative values. We chose $\lambda = \frac{1}{2\hat{\sigma}}$ with $\hat{\sigma}$ being the standard deviation of $\{\mathbf{w}^\top \mathbf{r}_i\}_{i \in [N]}$ for each $\mathbf{w} = [w_1; 1 - w_1]$. The evaluation was carried out using $N = M = 10^4$ samples. The boxplots of the resulting objective values as a function of the portfolio weight w_1 derived from 15 repetitions are summarized in Fig. 5.6. As it can be seen, the analytical knowledge of the mean embedding is again beneficial: the objective values obtained using the semi-explicit MMD estimator follow a more emphasized U-shape behaviour.

These experiments demonstrate the efficiency of the proposed technique. Our theoretical and numerical results show that the semi-explicit MMD has improved concentration property compared to the two-sample MMD estimator. The Gaussian kernel leads to more pronounced U-shape of the objectives than the unbounded kernel, and suggests easier optimization. The unbounded support of Gaussian targets is better motivated in finance compared to the bounded one of the beta target. Hence we keep the Gaussian target for our subsequent numerical studies. The next subsection is dedicated to the exploration of the efficiency of our technique in the misspecified setting.

5.4.2 Misspecified setting

In this section, we study the impact of misspecification, in other words when $D(P_{\mathbf{w}^*}, P_T) \neq 0$. In particular, we investigate the case when $d = 2$, the target distribution is Gaussian, and the first coordinate R_1 of the returns $R = [R_1; R_2] \in \mathbb{R}^2$ is generalized normal⁵, its second coordinate R_2 is normal. We assume that there exists an initial portfolio $\mathbf{w}_0 = [x; 1 - x] \in \mathcal{W}$. The target distribution \mathbb{P}_T is chosen to be equal to the distribution of $(\mathbf{w}_0)^\top R \in \mathbb{R}$ with $\gamma = 2$. Our goal is to study the impact on the estimated optimal portfolio

$$\mathbf{w}^* := [x^*; 1 - x^*] := \arg \min_{\mathbf{w} \in \mathcal{W}} D(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$$

as a function of the heavy-tailness parameter $\gamma < 2$ when the underlying \mathbf{w} is changed slightly, and to explore if x^* turns into smaller than x (less weight is allocated on the heavy-tailed component) or if x^* becomes higher than x (more weight is allocated on the heavy-tailed component). Denoting by $\varepsilon > 0$ the measure of local perturbation, and by $\mathbf{w}_0^- = [(x - \varepsilon); 1 - (x - \varepsilon)]$ and $\mathbf{w}_0^+ = [x + \varepsilon; 1 - (x + \varepsilon)]$ the corresponding portfolio weights, let

$$\mathbf{w}^{\min} := [x^{\min}; 1 - x^{\min}] := \arg \min_{\mathbf{w} \in \{\mathbf{w}_0^-, \mathbf{w}_0, \mathbf{w}_0^+\}} D(\mathbb{P}_{\mathbf{w},N}, \mathbb{P}_T)$$

⁵Generalized normal distributions include Gaussians with $\gamma = 2$ and they correspond to heavy-tailed distributions when $\gamma < 2$.

be the weight with the lowest divergence among these 3 allocations. The sign of the optimal allocation variation is $\text{sign}(\Delta x) = \text{sign}(x^* - x) = \text{sign}(x^{\min} - x)$. In our experiments, the heavy-tailness parameter was chosen to be $\gamma = 1$ (which gives rise to the Laplace distribution), we used the semi-explicit MMD as divergence estimator with sample size $N = 10^4$, and we set $\varepsilon = 0.1$. The mean $\mathbf{m}_R = \begin{bmatrix} m_1 \\ m_2 \end{bmatrix}$ and the covariance matrix $\Sigma_R = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}$ of R was generated randomly from $m_i \sim \mathcal{U}[0.01, 0.05]$ and $\sigma_i^2 \sim \mathcal{U}[0.01, 0.06]$. The initial portfolio weight $\mathbf{w} = [x; 1 - x]$ was varied by changing $x \in \{0, 0.1, 0.2, \dots, 0.9, 1\}$. Let $[\Delta m, \Delta \sigma^2] := [m_2 - m_1; \sigma_2^2 - \sigma_1^2]$ and $\Delta \gamma = 1 - 2 = -1$. Fig. 5.7 illustrates 10^4 pairs of $(\Delta m, \Delta \sigma^2)$ colored by the associated sign of $\frac{\Delta x}{\Delta \gamma}$; each subfigure corresponds to a specific choice of weight x . Positive sign of $\frac{\Delta x}{\Delta \gamma}$ means that the optimal allocation turns to be less concentrated on the heavy-tailed component ($\Delta x < 0$) when the initial return distribution has one component with heavier tails than the target ($\Delta \gamma < 0$). Conversely, negative sign of $\frac{\Delta x}{\Delta \gamma}$ shows that the new allocation becomes more concentrated on the heavy-tailed component.

As it can be seen in the figure, there is no systematic decrease of the allocation associated to the heavy-tailed component, rather there is an interplay between the variance variation $\Delta \sigma^2$ and the initial value x , without Δm affecting the results. For low values of x ($x \leq 0.4$) heavy-tailness has limited impact, for "balanced" initial portfolios ($0.5 \leq x \leq 0.7$) and $\Delta \sigma^2 > 0$ we retrieve the expected lower weight on heavy-tailed component, for larger x and negative $\Delta \sigma^2$ the new optimal portfolio becomes more concentrated on the heavy-tailed component. These experiment show that larger variance or concentration lead to more weight on the heavy-tailed component. One can explain this phenomena using the fact that heavy-tailed distributions with high variance have more extreme values, which affect the divergence matching.

5.5 Experiments on financial time series

In this section, we demonstrate the efficiency of the investigated target distribution technique for portfolio optimization. The considered datasets, rebalancing, performance measures and target distributions are summarized in Section 5.5.1, followed by the numerical studies in Section 5.5.2.

5.5.1 Datasets, rebalancing, performance measures, target distributions

This section summarizes the necessary ingredients of the numerical experiments.

Datasets: We selected 2 representative portfolio datasets (see Table 5.7) from the Kenneth French's library [Fre21] (also used in the recent extensive numerical study carried out by [Las19]). The 6BTM benchmark is based on the size, book-to-market and operating profitability performance indicators. The dataset components are built as the linear combination of stock returns from the NYSE market equity sorted by two main indicators. For example, the first component of 6BTM corresponds to a linear combination of small size (or market equity) and high book-to-market ratio (or low growth), small meaning belonging to the the first half (size criteria) and 70 – 100% percentile (book-to-market ratio criteria) bucket; see Fig. 5.8 for an illustration of the price evolution of 2 components of this benchmark. For this dataset the portfolio size is $d = 6$. Our second portfolio dataset is referred to as 10Ind, and it corresponds to the linear combination of stock returns by sector of activity, with portfolio size $d = 10$. The variance of the components of the two datasets is illustrated in Fig. 5.9. As the figure shows the variance is more balanced in the 6BTM than on the 10Ind dataset; in case of the latter the higher variance components correspond to Energy, HighTech and Durables. According to our misspecification study (Section 5.4.2), the target distribution technique is expected to perform better on the 6BTM dataset. The average excess kurtosis is 15.1 for the 6BTM and 14.5 for the 10Ind benchmark. The average skewnesses are negative, equal to -0.58 and -0.37 , respectively. These statistics corresponds to a heavy-tailed and negatively-skewed setting. In Fig. 5.10, we

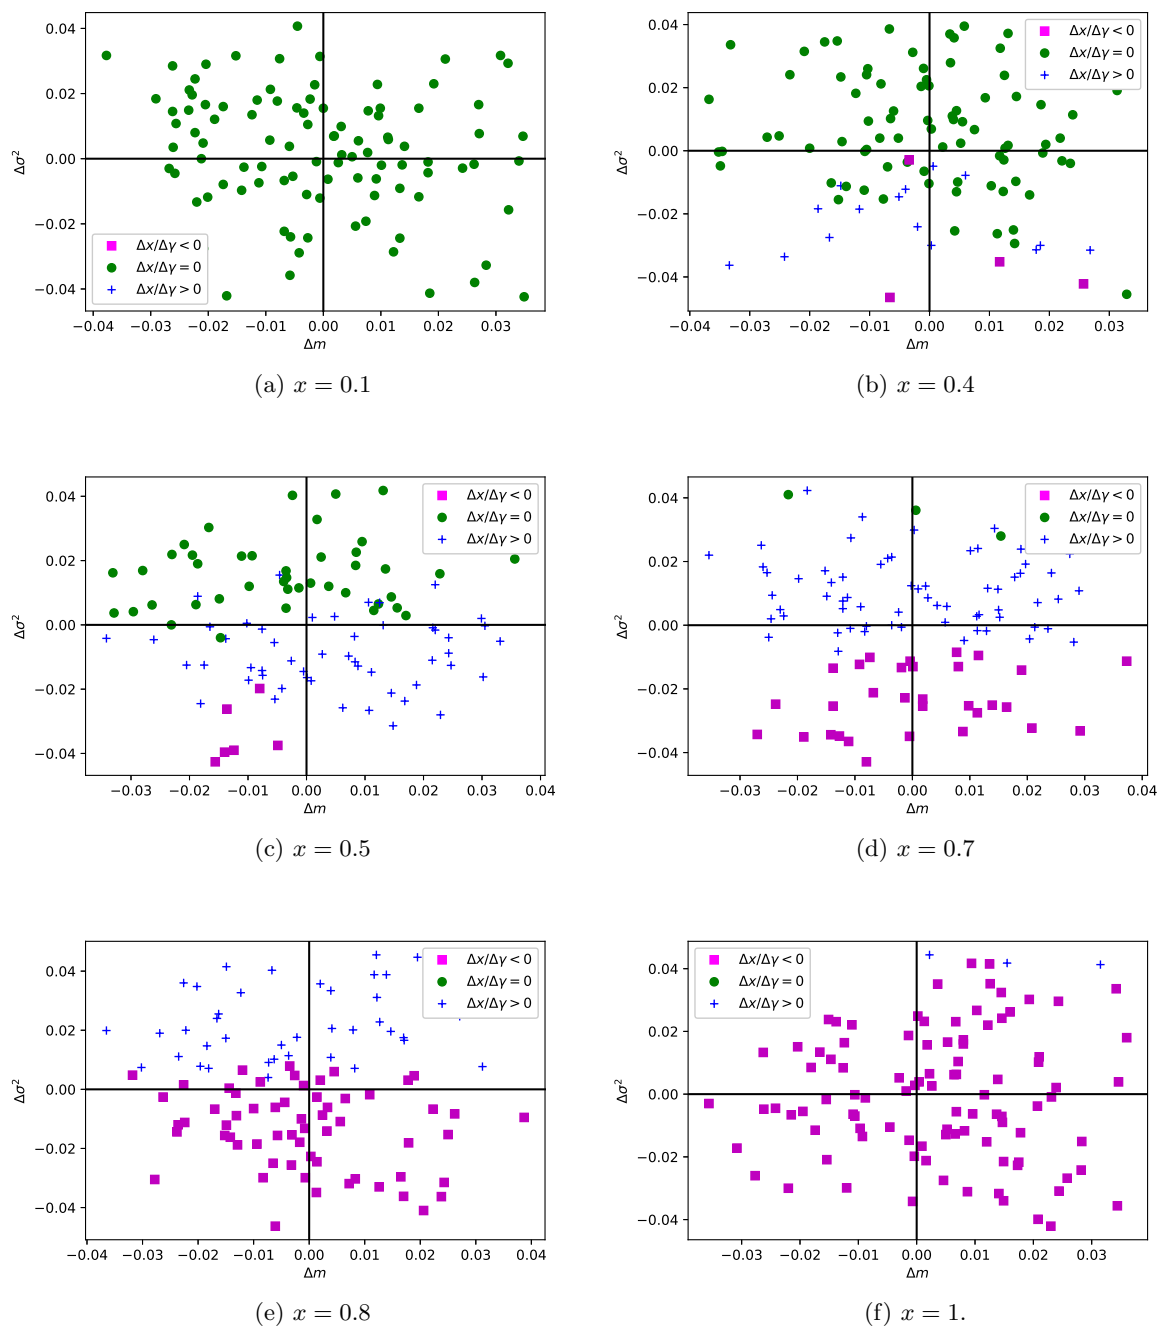


Figure 5.7 – Plots of randomly simulated $(\Delta\mu, \Delta\sigma)$ pairs colored by their associated sign of $\frac{\Delta x}{\Delta\gamma}$. Magenta: $\frac{\Delta x}{\Delta\gamma} < 0$. Blue: $\frac{\Delta x}{\Delta\gamma} > 0$. Green: $\frac{\Delta x}{\Delta\gamma} = 0$. The subfigures correspond to different x weights.

Dataset	Abbreviation	d
6 Fama-French portfolios of firms sorted by size and book-to-market	6BTM	6
10 industry portfolios representing the US stock market	10Ind	10

Table 5.7 – Financial datasets used in our empirical study. They are composed of daily value-weighted arithmetic returns from January 1980 to December 2020. Total number of samples: $N = 10318$.




Figure 5.8 – Price evolution associated to the first and last component of the 6BTM datasets. The prices (P_t) are calculated from returns (r_t) data as $P_t = P_{t-1}(1 + r_t)$ where we fixed $P_0 = 1$.

display the evolution of the returns of the 6BTM and 10Ind portfolios. One can notice periods of high and low volatility; the periods of high volatility occur at the same time points in the two datasets. To study the dependence of the target distribution technique on market conditions (low-volatility and high-volatility periods, corresponding to normal economic conditions and potential crisis periods), the datasets were split in two sub-samples. The considered quantitative criteria to distinguish these two regimes consists in comparing the semi-annual average variance to the average variance over the 41 years of data, as displayed in Fig. 5.10(c). When the annual variance is larger than 1.5 times the long-term average variance, we say that we are in a high-volatility/crisis period. There are 7 buckets out of 71 (specified by the rolling window approach detailed below) that this criteria considers to be highly volatile.

Rebalancing, performance measures: For a fair numerical comparison we use the same rebalancing methodology as [Las19] in his recent extensive study. An estimation window of 5 years and a rolling window (interval on which the output statistics are computed) of six months were considered. On our 41 years of data, it makes 71 buckets on which the out-of-sample performance of the portfolio is evaluated in terms of kurtosis and skewness. Our benchmark portfolio allocation strategies with their abbreviations are listed in Table 5.8. They correspond to the best-performing traditional and minimum divergence portfolio of the recent report [Las19, Section 5.7.3].

Target distributions: In practice, financial returns exhibit positive excess kurtosis and negative skewness. These properties are referred to as stylized facts [Con01]. Hence the returns cannot be considered Gaussian or elliptical (and hence described solely by their mean and variance). Considering standard concave and increasing utility maximisation framework [Mer71] it can be showed that investors prefer higher skewness and lower kurtosis.⁶ Targeting lighter-

⁶More generally, the investor's preference direction is positive (negative) for positive (negative) values of every




Figure 5.9 – Variance of the component returns of the 6BTM (a) and the 10Ind dataset (b).

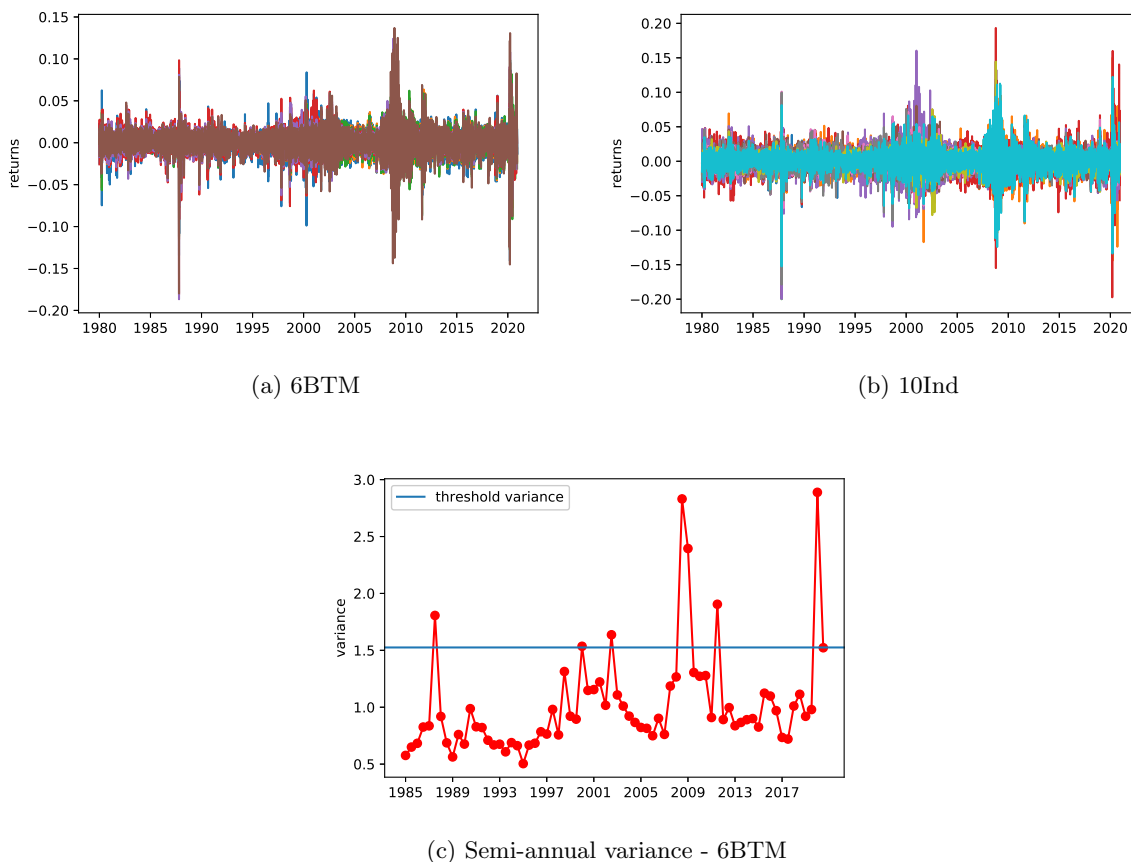


Figure 5.10 – Superposition of the component returns of the 6BTM (a) and the 10Ind dataset (b). In (c) the semi-annual variance of the 6BTM returns is shown. There are volatility clusters in 1988, 2000, 2007 and 2020 (due to financial crisis/inflation/Covid-19) in both datasets.

Portfolio	Abbreviation	Reference
sample minimum variance portfolio	MV	[JM03]
Bayes-Stein maximum Sharpe ratio portfolio	MSR	[Jor86]
minimum Kullback-Leibler portfolio with Gaussian target	KL ^a	[Las19] (5.1.4)

^aIn case of using the KL divergence with Gaussian target, the objective function simplifies to (see [Las19, Prop. 5.1] with $\gamma = 2$) $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{2\sigma_T^2} \mathbb{E}_{P \sim \mathbb{P}_{\mathbf{w}}} [(P - m_T)^2] - H(P)$ where denotes the Shannon entropy of P , which can be estimated using m-estimator, and $\frac{1}{2\sigma_T^2} \mathbb{E}_{P \sim \mathbb{P}_{\mathbf{w}}} [(P - m_T)^2]$ can be estimated using P sample moments.

Table 5.8 – State-of-the-art portfolios [Las19].

Divergence (D)	Target distribution (\mathcal{Q})
semi-explicit MMD ^a	Gaussian
MMD with low-rank approximation ^b	Gaussian, generalized normal, skew Gaussian
KSD	Gaussian
FSSD	Gaussian, generalized normal, skew Gaussian
WAD with $p = 1$	Gaussian, generalized normal, skew Gaussian
WAD with $p = 2$	Gaussian, generalized normal, skew Gaussian
KL	Gaussian

^aWe use the shorthand 'MMD explicit' for this estimator in the figures.

^bThe low-rank approximation is implemented using incomplete Cholesky decomposition.

Table 5.9 – Tested divergence - target distribution pairs.

tailed distribution (via for example generalized normal distribution) or more positively skewed distributions (for instance with skew Gaussian target) should have a favorable impact regarding the out-of-sample moments of the minimum-divergence portfolio compared to traditional mean-variance portfolios. In our numerical studies, we are going to focus on three families of target distributions: (i) the Gaussian distribution for which explicit mean embedding is available for various kernel choices (see Table 5.5), (ii) the generalized normal distribution with parameter γ allowing to target negative excess kurtosis when chosen $\gamma > 2$, (iii) the skew Gaussian distribution with parameter c allowing to target positive skewness with $c > 0$. In the latter two cases, we chose light-tailed generalized normal ($\gamma = 4$), and positively skewed skew Gaussian distribution ($s = 1$, in this case the excess kurtosis is 0.0617 by Table 5.12, hence very close to 0) as targets. We fixed the target mean and standard deviation as given by the maximum Sharpe ratio portfolio (in sample mean and standard-deviation associated to the maximum Sharpe ratio allocation). We summarize in Table 5.9 our choices of target distributions and divergences, noting that the linear-time FSSD provides a faster alternative to the quadratic-time KSD divergence.

5.5.2 Numerical studies

In this section, the out-of-sample kurtosis and skewness of our divergence-based portfolios versus the traditional minimum variance and maximum Sharpe ratio portfolios are compared, across different datasets and economic conditions. Kurtosis is investigated in Section 5.5.2.1, we focus on skewness in Section 5.5.2.2; in these sections the low and high-volatility regimes are studied separately. Section 5.5.2.3 provides an overall assessment with low and high-volatility periods merged.

In our experiments, we made use of the KGOF and ITE toolbox packages [Sza14, Sza14]

odd central moment and negative for every even central moment [SH80].




Figure 5.11 – OOS kurtosis values in the low-volatility regime. Left: baselines (MV, MSR). Right: proposed minimum divergence portfolios with skew Gaussian (1st column) and generalized normal (2nd column) target. 1st row: 6BTM dataset. 2nd row: 10Ind benchmark.

when computing the FSSD and the two-samples MMD divergence measures.

5.5.2.1 Kurtosis impact

In this section, the impact on the portfolio kurtosis is investigated. We report the out-of-sample (OOS) excess kurtosis values for the different divergence objectives considered, and compare them to the traditional minimum variance portfolio (MV) and the maximum Sharpe ratio (MSR) portfolio. Lower kurtosis is an indication of less heavy tails, and this is what is usually advantageous from financial point of view, in other words more negative (or at least closer to zero) excess kurtosis corresponds to better performing portfolio.

Normal economic conditions: The OOS kurtosis values with generalized normal and skew Gaussian targets are summarized in Fig. 5.11. For both datasets, the MMD-based divergences give lower median kurtosis than the WAD-based portfolios and the traditional portfolios, and the lower quartile for most divergence based portfolios is lower than that of the traditional portfolios.

In Fig. 5.12, we provide the barplots of the number of buckets (time periods) on which a given objective-target pair achieved the lowest kurtosis across all the objectives; they are sorted by increasing efficiency and we show the 5 winners. On the 6BTM dataset, the 5 winners correspond to light-tailed target distributions (skew or generalized normal) with similar performance. On the 10Ind benchmark, MMD with skew Gaussian target performs the best, followed by the minimum KL approach using Gaussian target (which has heavier tails than the skew Gaussian distribution). This (i) means that light-tailed targets are not predominant here, and the reason for this can be that the 10Ind dataset is a more heterogeneous than 6BTM as it can be seen




Figure 5.12 – The five top objective-target pairs in kurtosis sense. Regime: low-volatility. Ordering: from worst (left) to best (right). (a): 6BTM dataset. (b): 10Ind benchmark.

from the variance barplots in Fig. 5.9.

High-volatility conditions: In high-volatility scenarios on both datasets the results are more diverse, see Fig. 5.13. For the 6BTM benchmark, MMD with skew Gaussian target and FSSD with generalized normal target lead to relatively small kurtosis. On the 10Ind dataset, WAD ($p = 1$) with generalized normal target also seems to perform well.

Considering the Gaussian target (which gives rise to analytical divergence computations and targets with 0 kurtosis which is lower than the average 14 – 15 kurtosis value of the datasets), the OOS kurtosis values obtained for various divergences are summarized in Fig. 5.14. In the high-volatility regime Fig. 5.14(b), the semi-explicit MMD and the KSD portfolios give more concentrated kurtosis values compared to the low-rank MMD approximation; this phenomenon can be explained due to the lack of sampling needed from the target distribution. These effects are less pronounced in the low-volatility regime, see Fig. 5.14(a).

5.5.2.2 Skewness impact

In this section, the impact on the portfolio skewness is investigated. We report the OOS excess skewness values for the different divergence objectives considered, and compare them to the traditional minimum variance portfolio (MV) and the maximum Sharpe ratio (MSR) portfolio. Higher skewness is an indication of less negatively skewed distribution, and this is what is usually beneficial from financial perspective, in other words more positive (or at least closer to zero) excess skewness is interpreted as better performing portfolio.

Normal economic conditions: The OOS skewness values with generalized normal and skew Gaussian targets are summarized in Fig. 5.15. For both datasets, the skew Gaussian target with either MMD or FSSD divergence give the highest median skewness and the lower quartile for most divergence based portfolios is higher than that of the traditional portfolios. The latter behaviour is particularly apparent on the 6BTM benchmark with the generalized normal target.

In Fig. 5.16, we provide the barplots of the number of buckets on which a given objective-target pair achieved the highest skewness across all the objectives; we show the 5 winners.

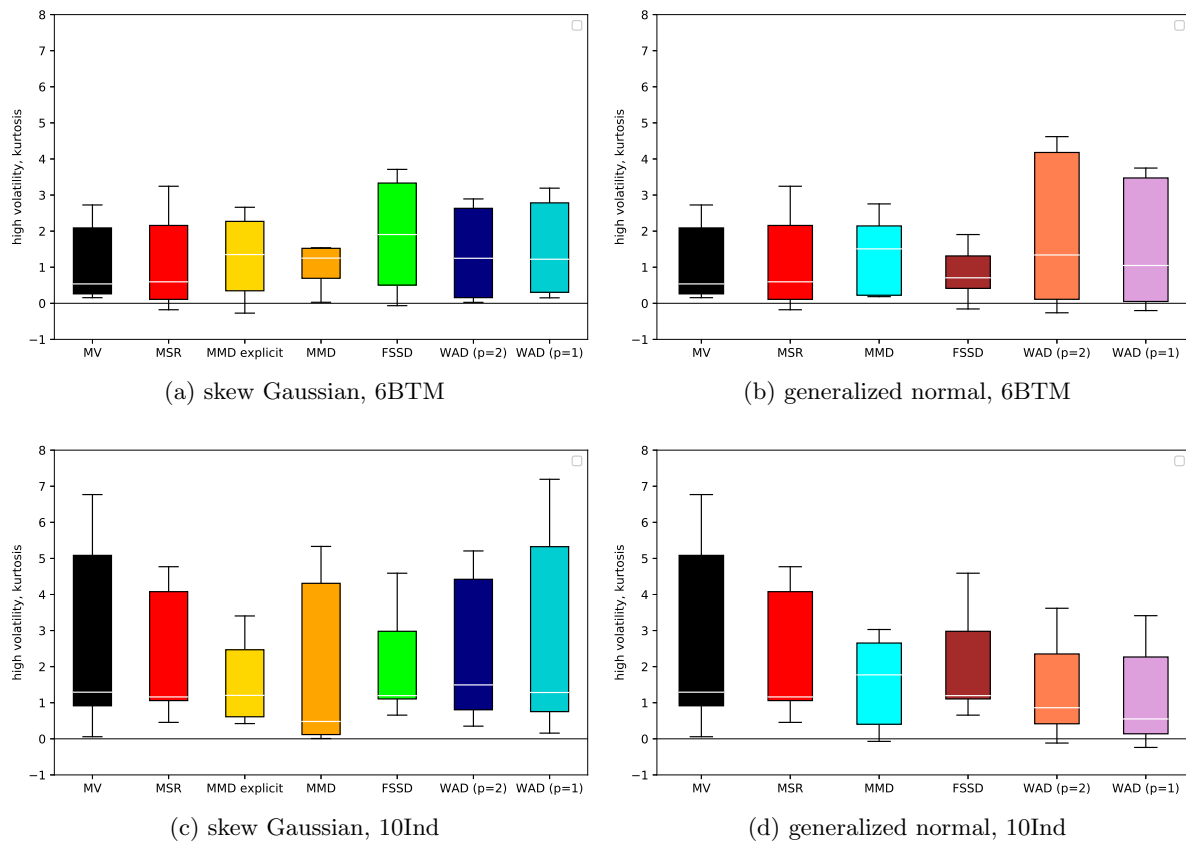


Figure 5.13 – OOS kurtosis boxplots in the high-volatility regime. Left: baselines (MV, MSR). Right: proposed minimum divergence portfolios with skew Gaussian (1st column) and generalized normal (2nd column) target. 1st row: 6BTM benchmark. 2nd row: 10Ind dataset.

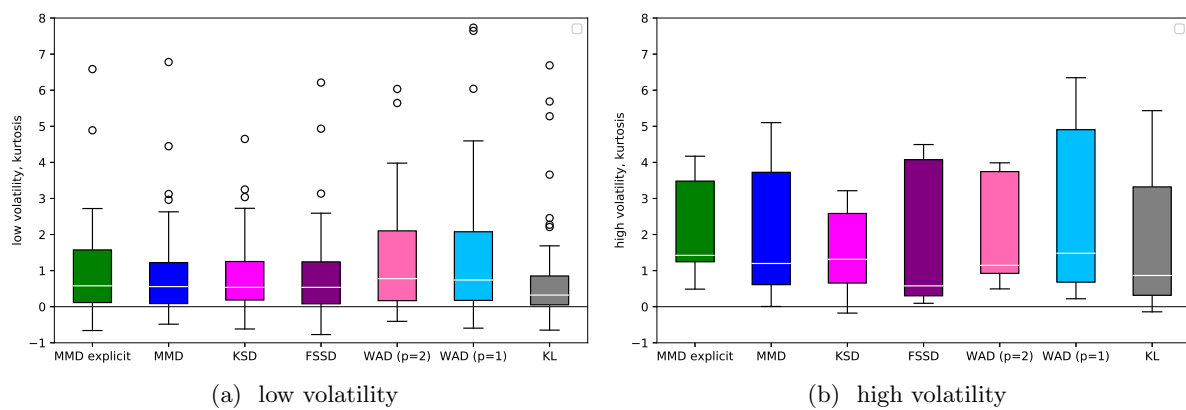


Figure 5.14 – OOS kurtosis values for minimum divergence portfolios with the Gaussian target. Period: low-volatility (a), high-volatility (b). Dataset: 10Ind.

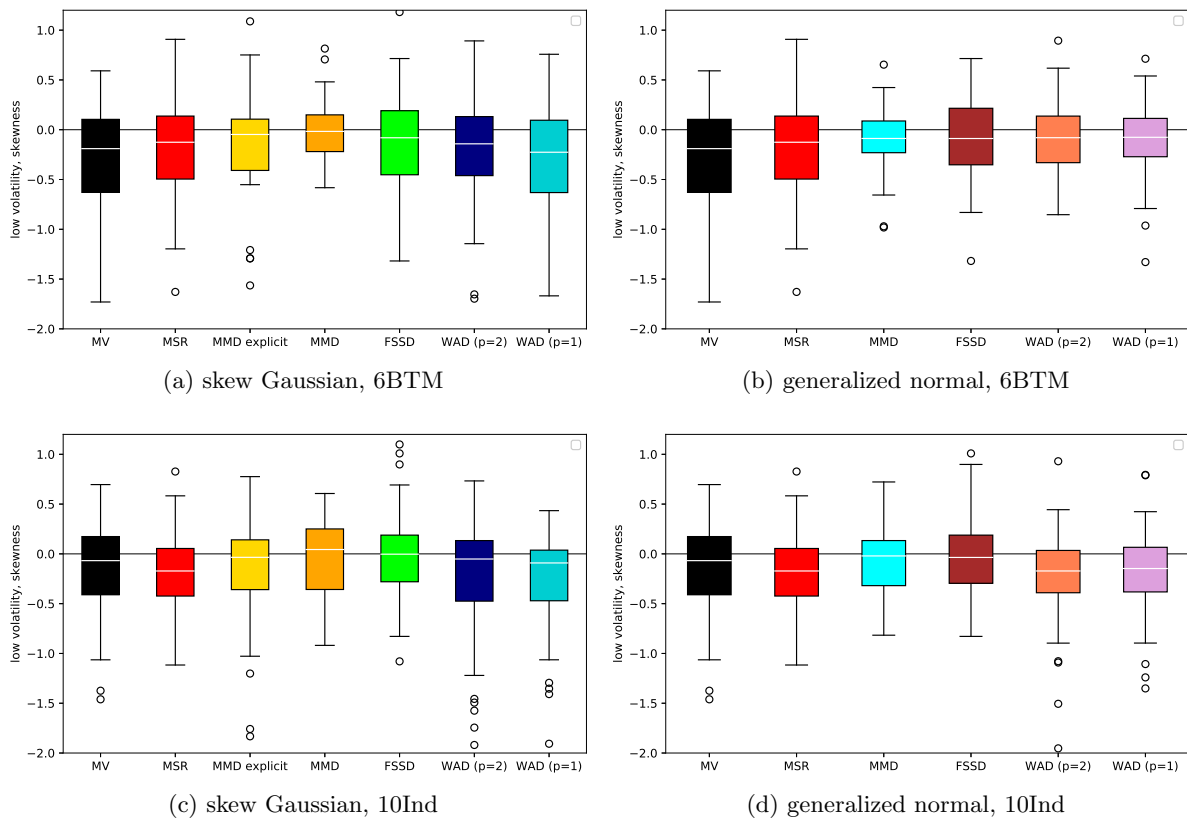


Figure 5.15 – OOS skewness values in the low-volatility regime. Left: baselines (MV, MSR). Right: proposed minimum divergence portfolios with skew Gaussian (1st column) and generalized normal (2nd column) target. 1st row: 6BTM dataset. 2nd row: 10Ind benchmark.




Figure 5.16 – The five top objective-target pairs in skewness sense. Regime: low-volatility. Ordering: from worst (left) to best (right). (a): 6BTM dataset. (b): 10Ind benchmark.

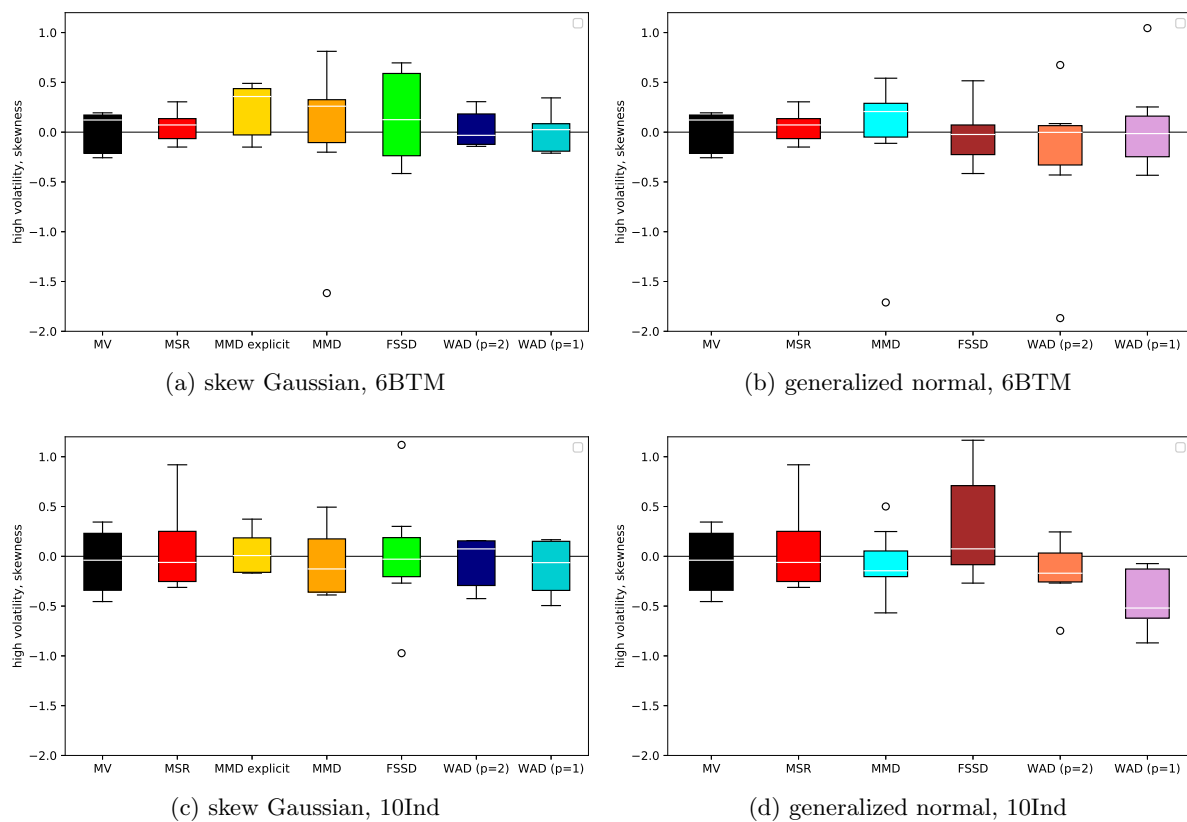


Figure 5.17 – OOS skewness boxplots in the high-volatility regime. Left: baselines (MV, MSR). Right: proposed minimum divergence portfolios with skew Gaussian (1st column) and generalized normal (2nd column) target. 1st row: 6BTM benchmark. 2nd row: 10Ind dataset.

On both datasets, the top performers correspond to light-tailed target distributions (skew or generalized normal). We also see that the WAD objective with $p = 1$ is performing well on the 6BTM dataset.

High-volatility conditions: In high-volatility scenarios on both datasets the results are more diverse, see Fig. 5.17. The skewness of traditional portfolios is not necessarily negative. As for the kurtosis, the target-distribution portfolio approach seems to have more impact in normal economic conditions. For the 6BTM benchmark, MMD-based objectives give a larger median skewness than the traditional portfolios. On the 10Ind dataset, only FSSD with generalized normal target succeeds in beating the traditional portfolios, but with quite a large margin in terms of the upper quartile.

5.5.2.3 Merged results

In this section, we comment the barplots of Fig. 5.18 corresponding to the number of buckets on which a given objective-target pair achieved the highest kurtosis or skewness across all the objectives, across all the buckets (with low and high volatility scenarios merged).

For both benchmarks, the clear winners are

- MMD with skew Gaussian target (orange, present in all 4 subfigures),
- FSSD with generalized normal target (dark red, present in 3 subfigures, 1st/2nd place in 2 of them),
- Wasserstein distance ($p = 1$) with generalized normal target (mauve, present in 3 subfigures, 1st/2nd position in 2 cases),
- Kullback-Leibler divergence with Gaussian target (grey), especially on the 10Ind dataset.

The good performance of the WAD divergence (with $p = 1$) might be due to the outlier-robust character of L_1 -based objective functions. Overall, the best performing portfolios are in majority associated to light-tail distribution targets, except for the KL-Gaussian which performs well on the 10Ind dataset. The final barplots in Fig. 5.18 show more homogeneity on the 6BTM datasets (in terms of kurtosis), many objective functions are performing similarly.


5.6 Conclusion

In this chapter, we focused on kernel-based divergence measures and their application in financial portfolio optimization using the target distribution framework. We showed that prior knowledge available on target distribution leads to analytical forms for the mean embedding (Lemma 5.3.1 and Lemma 5.3.2) and improved MMD concentration properties for bounded kernels (Theorem 5.2). Motivated by recent financial studies relying on unbounded kernels [BF19], using the Burkholder inequality we proved that the $1/\sqrt{N}$ rate (with N denoting the sample size) of MMD estimators can be extended to unbounded kernels (Theorem 5.3); we illustrated the idea for the exponential kernel. We showed matching minimax lower bounds (Theorem 5.4) under slightly more restrictive conditions.

We illustrated the effectiveness of the resulting semi-explicit MMD estimator over the classical two-sample MMD estimator not relying on the knowledge of mean embedding on simulated examples. We demonstrated the proposed kernel-based divergence framework in real-world portfolio optimization problems. According to our numerical experiences, the suggested technique


- performs favorably compared to classical portfolio optimization scheme on balanced portfolios in normal economic conditions,
- MMD and KSD with light-tailed targets often improves the out-of-sample skewness and kurtosis of the portfolio.

Our preliminary results on alternative statistics show that the information theoretical technique can have slightly higher turnover than traditional schemes, which makes it an ideal candidate as a regularizer added to the global portfolio optimization objective.



(a) skewness, 6BTM

(b) kurtosis, 6BTM



(c) skewness, 10Ind

(d) kurtosis, 10Ind

Figure 5.18 – The five top objective-target pairs. Regime: low and high-volatility merged. Ordering: from worst to best in skewness sense (1st column) and kurtosis sense (2nd column). 1st row: 6BTM benchmark. 2nd row: 10Ind dataset.

APPENDICES

5.A Proofs

This section contains the proofs of theoretical results. In Section 5.A.1 we give the proofs of the two analytical mean embedding results (Lemma 5.3.1 and Lemma 5.3.2), accompanied by two additional lemmas used in Section 5.4.1 and in the proof of Theorem 5.4, respectively. Section 5.A.2 is dedicated to the proof of our concentration results (Theorem 5.2, Theorem 5.3, Theorem 5.4).

5.A.1 Proofs of auxiliary statements

In this section we prove Lemma 5.3.1 and Lemma 5.3.2. These proofs are accompanied by two additional auxiliary statements (Lemma 5.A.1 and Lemma 5.A.2). Lemma 5.A.1 is on the Stein kernel for the Gaussian distribution - Gaussian kernel, Lemma 5.A.2 is about MMD for the Gaussian distribution and the exponential kernel. The former is used in the KSD computation in Section 5.4.1; latter is applied in the reasoning of Theorem 5.4.

Proof. (**Lemma 5.3.1**; mean embedding: Gaussian target - Gaussian-exponentiated kernel)
Our target quantity is

$$\mu_K(\mathbb{Q})(x) = \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-m)^2}{2\sigma^2}} e^{-a(x-y)^2 + bxy} dy.$$

Completing the square, one gets

$$\begin{aligned} \frac{(y-m)^2}{2\sigma^2} + a(x-y)^2 - bxy &= \frac{(y-m)^2 + 2\sigma^2 a(x-y)^2 - 2\sigma^2 bxy}{2\sigma^2} \\ &= \frac{(1+2a\sigma^2)y^2 - 2(m+2a\sigma^2x+b\sigma^2x)y + 2a\sigma^2x^2 + m^2}{2\sigma^2} \\ &= \underbrace{\frac{(1+2a\sigma^2)}{2\sigma^2}}_{=:\frac{1}{2\sigma^{*2}}} \left(y^2 - 2 \underbrace{\frac{m+2a\sigma^2x+b\sigma^2x}{1+2a\sigma^2}}_{=:m^*} y \right) + \frac{2a\sigma^2x^2 + m^2}{2\sigma^2} \\ &= \frac{1}{2\sigma^{*2}} \left[(y-m^*)^2 - m^{*2} \right] + \frac{2a\sigma^2x^2 + m^2}{2\sigma^2} \\ &= \frac{1}{2\sigma^{*2}} (y-m^*)^2 - \underbrace{\frac{(m+2a\sigma^2x+b\sigma^2x)^2}{2\sigma^2(1+2a\sigma^2)}}_{=:-c^*} + \frac{2a\sigma^2x^2 + m^2}{2\sigma^2}. \end{aligned}$$

Bringing the two terms in c^* to common denominator, after simplification one arrives at

$$\begin{aligned} c^* &= \frac{(m + 2a\sigma^2x + b\sigma^2x)^2 - (2a\sigma^2x^2 + m^2)(1 + 2a\sigma^2)}{2\sigma^2(1 + 2a\sigma^2)} \\ &= \frac{4a\sigma^2mx - 2a\sigma^2(m^2 + x^2) + b\sigma^2x(b\sigma^2x + 4a\sigma^2x + 2m)}{2\sigma^2(1 + 2a\sigma^2)} \\ &= \frac{-2a\sigma^2(m - x)^2 + b\sigma^2x(b\sigma^2x + 4a\sigma^2x + 2m)}{2\sigma^2(1 + 2a\sigma^2)} \\ &= -\frac{a(x - m)^2}{1 + 2a\sigma^2} + \frac{2bm x + b\sigma^2(b + 4a)x^2}{2(1 + 2a\sigma^2)}. \end{aligned}$$

This means that our target quantity can be rewritten as

$$\begin{aligned} \mu_K(\mathbb{Q})(x) &= \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^{*2}}(y-m^*)^2 + c^*} dy = \frac{\sigma^*}{\sigma} e^{c^*} \underbrace{\int_{\mathbb{R}} \frac{1}{\sqrt{2\pi(\sigma^*)^2}} e^{-\frac{1}{2\sigma^{*2}}(y-m^*)^2} dy}_{=1} \\ &= \frac{e^{-\frac{a(x-m)^2}{1+2a\sigma^2} + \frac{2bm x + b(b+4a)\sigma^2 x^2}{2(1+2a\sigma^2)}}}{\sqrt{1 + 2a\sigma^2}}. \end{aligned}$$

□

Proof. (**Lemma 5.3.2**; mean embedding: beta target - Matérn kernel) Our target quantity is

$$\begin{aligned} &\mu_K(\mathbb{Q})(x) \\ &= \frac{\sigma_0^2}{B(\alpha, \beta)} \int_0^1 e^{-\frac{\sqrt{2p+1}|x-y|}{\sigma}} \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}|x-y|}{\sigma} \right)^{p-i} y^{\alpha-1} (1-y)^{\beta-1} \mathbf{1}_{y \in [0,1]} dy \\ &\stackrel{(*)}{=} \frac{\sigma_0^2}{B(\alpha, \beta)} \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}}{\sigma} \right)^{p-i} \left[\int_0^{(0 \vee x) \wedge 1} (x-y)^{p-i} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(x-y)}{\sigma}} dy \right. \\ &\quad \left. + \int_{(0 \vee x) \wedge 1}^1 (y-x)^{p-i} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(y-x)}{\sigma}} dy \right]. \end{aligned}$$

where in (*) we applied the decomposition trick: for $f : \mathbb{R} \rightarrow \mathbb{R}$ and $x \in \mathbb{R}$,

$$\int_0^1 f(|x-y|) dy = \int_0^{(0 \vee x) \wedge 1} f(x-y) dy + \int_{(0 \vee x) \wedge 1}^1 f(y-x) dy.$$

Using the binomial formula $(x-y)^{p-i} = \sum_{k=0}^{p-i} \binom{p-i}{k} (-1)^{p-i-k} x^k y^{p-i-k}$ and $(y-x)^{p-i} = \sum_{k=0}^{p-i} \binom{p-i}{k} y^{p-i-k} (-x)^k$, one can rewrite the two integrals as

$$\begin{aligned} &\int_0^{(0 \vee x) \wedge 1} (x-y)^{p-i} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(x-y)}{\sigma}} dy = \\ &= \sum_{k=0}^{p-i} \binom{p-i}{k} (-1)^{p-i-k} x^k \int_0^{(0 \vee x) \wedge 1} y^{p-i-k} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(x-y)}{\sigma}} dy \\ &= \sum_{k=0}^{p-i} \binom{p-i}{k} (-1)^{p-i-k} x^k e^{-\frac{\sqrt{2p+1}x}{\sigma}} \underbrace{\int_0^{(0 \vee x) \wedge 1} y^{p-i-k+\alpha-1} (1-y)^{\beta-1} e^{\frac{\sqrt{2p+1}y}{\sigma}} dy}_{E_1^{\frac{\sqrt{2p+1}}{\sigma}}((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1)}, \end{aligned}$$

and

$$\begin{aligned}
& \int_{(0 \vee x) \wedge 1}^1 (y-x)^{p-i} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(y-x)}{\sigma}} dy = \\
&= \sum_{k=0}^{p-i} \binom{p-i}{k} (-x)^k \int_0^{(0 \vee x) \wedge 1} y^{p-i-k} y^{\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}(y-x)}{\sigma}} dy \\
&= \sum_{k=0}^{p-i} \binom{p-i}{k} (-x)^k e^{\frac{\sqrt{2p+1}x}{\sigma}} \underbrace{\int_{(0 \vee x) \wedge 1}^1 y^{p-i-k+\alpha-1} (1-y)^{\beta-1} e^{-\frac{\sqrt{2p+1}y}{\sigma}} dy}_{E_2^{\frac{\sqrt{2p+1}}{\sigma}}((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1)},
\end{aligned}$$

Hence,

$$\begin{aligned}
\mu_K(\mathbb{Q})(x) &= \frac{\sigma_0^2}{B(\alpha, \beta)} \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}}{\sigma} \right)^{p-i} \times \\
& \sum_{k=0}^{p-i} \binom{p-i}{k} x^k \left[(-1)^{p-i-k} e^{-\frac{\sqrt{2p+1}x}{\sigma}} E_1((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1) \right. \\
& \left. + (-1)^k e^{\frac{\sqrt{2p+1}x}{\sigma}} E_2((0 \vee x) \wedge 1, p-i-k+\alpha-1, \beta-1) \right].
\end{aligned}$$

□

Lemma 5.A.1 (Stein kernel: Gaussian target - Gaussian kernel). *Let the target distribution and the kernel be Gaussian: $q(x) = \frac{1}{\sqrt{2\pi}\sigma^2} e^{-\frac{(x-m)^2}{2\sigma^2}}$, $K(x, y) = e^{-\frac{(x-y)^2}{2c^2}}$ with $\sigma \in \mathbb{R}^+$ and $c \in \mathbb{R}^+$. Then*

$$\begin{aligned}
\xi_q(x, y) &= \left(\frac{m-x}{\sigma^2} + \frac{y-x}{c^2} \right) e^{-\frac{(x-y)^2}{2c^2}}, \\
h_q(x, y) &= \left[\frac{(x-m)(y-m)}{\sigma^4} - \frac{(x-y)^2(\sigma^2 + c^2)}{\sigma^2 c^4} + \frac{1}{c^2} \right] e^{-\frac{(x-y)^2}{2c^2}}.
\end{aligned}$$

Proof. (**Lemma 5.A.1**; Gaussian target - Gaussian kernel) Using that

$$\begin{aligned}
s_q(x) &= [\log q(x)]' = \left[-\frac{(x-m)^2}{2\sigma^2} \right]' = -\frac{x-m}{\sigma^2}, \\
\frac{\partial K(x, y)}{\partial x} &= -e^{-\frac{(x-y)^2}{2c^2}} \frac{x-y}{c^2} = -\frac{\partial K(x, y)}{\partial y} \text{ by symmetry,} \\
\frac{\partial^2 K(x, y)}{\partial x \partial y} &= \frac{\partial}{\partial y} \left(-e^{-\frac{(x-y)^2}{2c^2}} \frac{x-y}{\sigma^2} \right) = e^{-\frac{(x-y)^2}{2c^2}} \left[-\frac{(x-y)^2}{c^4} + \frac{1}{c^2} \right]
\end{aligned}$$

we obtain the statement

$$\begin{aligned}
\xi_q(x, y) &= s_q(x)K(x, y) + \frac{\partial K(x, y)}{\partial x} = -\frac{x-m}{\sigma^2} e^{-\frac{(x-y)^2}{2c^2}} - e^{-\frac{(x-y)^2}{2c^2}} \frac{x-y}{c^2} \\
&= \left(\frac{m-x}{\sigma^2} + \frac{y-x}{c^2} \right) e^{-\frac{(x-y)^2}{2c^2}}, \\
h_q(x, y) &= s_q(x)s_q(y)K(x, y) + s_q(y)\partial_x K(x, y) + s_q(x)\partial_y K(x, y) + \frac{\partial^2 K(x, y)}{\partial x \partial y} \\
&= \left[\frac{(x-m)(y-m)}{\sigma^2} + \underbrace{\left(-\frac{y-m}{\sigma^2} + \frac{x-m}{\sigma^2} \right)}_{=\frac{x-y}{\sigma^2}} \left(-\frac{x-y}{c^2} \right) - \frac{(x-y)^2}{c^4} + \frac{1}{c^2} \right] e^{-\frac{(x-y)^2}{2c^2}}
\end{aligned}$$

$$= \left[\frac{(x-m)(y-m)}{\sigma^4} - \frac{(x-y)^2(\sigma^2 + c^2)}{\sigma^2 c^4} + \frac{1}{c^2} \right] e^{-\frac{(x-y)^2}{2c^2}}.$$

□

Lemma 5.A.2 (Explicit MMD: Gaussian target - exponential kernel). *Let $\mathbb{P} = \mathcal{N}(m_0, \sigma^2)$ and $\mathbb{Q} = \mathcal{N}(m_1, \sigma^2)$ with $m_0, m_1 \in \mathbb{R}$ and $\sigma > 0$. Let $K(x, y) = e^{bxy}$ with $0 < b < \frac{1}{2\sigma^2}$. Then $\text{MMD}_K(\mathbb{P}, \mathbb{Q})$ can be computed analytically as*

$$\text{MMD}_K(\mathbb{P}, \mathbb{Q}) = \left(\frac{1}{\sqrt{1-r^2}} \left[e^{\frac{bm_0^2}{1-r}} + e^{\frac{bm_1^2}{1-r}} - 2e^{\frac{br(m_0^2+m_1^2)+2bm_0m_1}{2(1-r^2)}} \right] \right)^{1/2},$$

where $r = b\sigma^2$.

Proof. (**Lemma 5.A.2**) By definition

$$[\text{MMD}_K(\mathbb{P}, \mathbb{Q})]^2 = \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{P})(x) + \mathbb{E}_{x \sim \mathbb{Q}} \mu_K(\mathbb{Q})(x) - 2\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x). \quad (5.A.1)$$

Below we compute the value of the individual terms.

- **Term $\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)$:** The value of $\mu_K(\mathbb{Q})(x)$ can be obtained by Lemma 5.3.1 with $a = 0$:

$$\mu_K(\mathbb{Q})(x) = e^{\frac{2bm_1x + b^2\sigma^2x^2}{2}}.$$

Hence, by using the notation $r = b\sigma^2$ we have

$$\begin{aligned} \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x) &= \int_{\mathbb{R}} e^{\frac{2bm_1x + b^2\sigma^2x^2}{2}} \frac{e^{-\frac{(x-m_0)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma^2}} dx = \int_{\mathbb{R}} e^{-\frac{x^2}{2} \left[\frac{1}{\sigma^2} - (b\sigma)^2 \right] + x \left(bm_1 + \frac{m_0}{\sigma^2} \right) - \frac{m_0^2}{2\sigma^2}} \frac{1}{\sqrt{2\pi\sigma^2}} dx \\ &= \int_{\mathbb{R}} e^{-\frac{x^2}{2\sigma^2}(1-r^2) + \frac{x}{\sigma^2}(rm_1+m_0) - \frac{m_0^2}{2\sigma^2}} \frac{1}{\sqrt{2\pi\sigma^2}} dx. \end{aligned}$$

Completing the square, we have

$$\begin{aligned} x^2(1-r^2) - 2x(rm_1+m_0) &= (1-r^2) \left(x - \frac{rm_1+m_0}{1-r^2} \right)^2 - \frac{(rm_1+m_0)^2}{1-r^2}, \\ \frac{(rm_1+m_0)^2}{1-r^2} - \frac{m_0^2}{2\sigma^2} &= \frac{1}{1-r^2} [(rm_1+m_0)^2 - m_0^2(1-r^2)] \\ &= \frac{1}{1-r^2} [(rm_0)^2 + (rm_1)^2 + 2rm_0m_1]. \end{aligned}$$

This means that our target quantity can be rewritten as

$$\begin{aligned} \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x) &= \frac{e^{\frac{(rm_0)^2 + (rm_1)^2 + 2rm_0m_1}{2\sigma^2(1-r^2)}}}{\sqrt{1-r^2}} \underbrace{\int_{\mathbb{R}} e^{-\frac{(1-r^2)\left(x - \frac{rm_1+m_0}{1-r^2}\right)^2}{2\sigma^2}} \frac{\sqrt{1-r^2}}{\sqrt{2\pi\sigma^2}} dx}_{=1} \\ &\stackrel{(*)}{=} \frac{1}{\sqrt{1-r^2}} e^{\frac{(b\sigma^2 m_0)^2 + (b\sigma^2 m_1)^2 + 2b\sigma^2 m_0 m_1}{2\sigma^2(1-r^2)}} \\ &\stackrel{(**)}{=} \frac{1}{\sqrt{1-r^2}} e^{\frac{br(m_0^2+m_1^2)+2bm_0m_1}{2(1-r^2)}} \end{aligned} \quad (5.A.2)$$

where in (*) and (**), we have used $r = b\sigma^2$, factorized by σ^2 then replaced $b\sigma^2$ by r .

- **Term** $\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{P})(x)$: This term can be computed by specializing (5.A.2) to $m_0 = m_1$ and using that $1 - r^2 = (1 - r)(1 + r)$:

$$\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{P})(x) = \frac{1}{\sqrt{1 - r^2}} e^{\frac{2brm_0^2 + 2bm_0^2}{2(1-r^2)}} = \frac{1}{\sqrt{1 - r^2}} e^{\frac{2bm_0^2(r+1)}{2(1-r^2)}} = \frac{1}{\sqrt{1 - r^2}} e^{\frac{bm_0^2}{1-r}}. \quad (5.A.3)$$

- **Term** $\mathbb{E}_{x \sim \mathbb{Q}} \mu_K(\mathbb{Q})(x)$: Analogously to $\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{P})(x)$ we have

$$\mathbb{E}_{x \sim \mathbb{Q}} \mu_K(\mathbb{Q})(x) = \frac{1}{\sqrt{1 - r^2}} e^{\frac{bm_1^2}{1-r}}. \quad (5.A.4)$$

Substituting the obtained analytical forms (5.A.2), (5.A.3) and (5.A.4) to (5.A.1) we arrive at

$$[\text{MMD}_K(\mathbb{P}, \mathbb{Q})]^2 = \frac{1}{\sqrt{1 - r^2}} \left[e^{\frac{bm_0^2}{1-r}} + e^{\frac{bm_1^2}{1-r}} - 2e^{\frac{br(m_0^2 + m_1^2) + 2bm_0m_1}{2(1-r^2)}} \right].$$

□

5.A.2 Proof of concentration results (Theorem 5.2, Theorem 5.3, Theorem 5.4)

In the section we prove Theorem 5.2, Theorem 5.3, and Theorem 5.4.

Proof. (**Theorem 5.2**; concentration of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$, bounded kernel) By the definition of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$ and $\text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$, they have the term $\mathbb{E}_{y \sim \mathbb{Q}} \mu_K(\mathbb{Q})(y)$ in common, hence their difference writes as

$$\begin{aligned} & \widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \\ & \stackrel{(a)}{=} \underbrace{\frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j)}_{=: U_2} - \underbrace{\mathbb{E}_{x,x' \sim \mathbb{P}} K(x, x')}_{\mathbb{E}U_2} - 2 \left[\underbrace{\frac{1}{N} \sum_{i \in [N]} \mu_K(\mathbb{Q})(x_i)}_{=: U_1} - \underbrace{\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)}_{\mathbb{E}U_1} \right] \\ & = U_2 - \mathbb{E}U_2 - 2(U_1 - \mathbb{E}U_1) \end{aligned} \quad (5.A.5)$$

by using in (a) that U_1 and U_2 are U-statistics. The kernel of U_1 is $h_1(x) = \mu_K(\mathbb{Q})(x)$, the kernel of U_2 is $h_2(x, x') = K(x, x')$. Since by assumption the kernel K is lower bounded by A and upper bounded by B , the same property holds for h_1 and h_2 . Hence applying the Hoeffding bound for U-statistics (Theorem 5.1), for any $t > 0$

$$\begin{aligned} \mathbb{P}(U_1 - \mathbb{E}U_1 < -t) &= \mathbb{P}((-U_1) - \mathbb{E}(-U_1) > t) \leq e^{-\frac{2Nt^2}{(B-A)^2}}, \\ \mathbb{P}(U_2 - \mathbb{E}U_2 > t) &\leq e^{-\frac{2\lfloor \frac{N}{2} \rfloor t^2}{(B-A)^2}}. \end{aligned} \quad (5.A.6)$$

Returning to our target quantity $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$, for any $\varepsilon > 0$

$$\begin{aligned} \left\{ \widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon \right\} &\stackrel{(5.A.5)}{=} \{U_2 - \mathbb{E}U_2 - 2(U_1 - \mathbb{E}U_1) > \varepsilon\} \\ &\stackrel{(a)}{\subseteq} \left\{ U_2 - \mathbb{E}U_2 > \frac{\varepsilon}{2} \right\} \cup \left\{ U_1 - \mathbb{E}U_1 < -\frac{\varepsilon}{4} \right\}, \end{aligned} \quad (5.A.7)$$

where the inclusion $A \subseteq B \cup C$ in (a) is equivalent to $\bar{B} \cap \bar{C} \subseteq \bar{A}$; the latter holds as $\{U_2 - \mathbb{E}U_2 \leq \frac{\varepsilon}{2}\} \cap \{-2(U_1 - \mathbb{E}U_1) \leq \frac{\varepsilon}{2}\} \subseteq \{U_2 - \mathbb{E}U_2 - 2(U_1 - \mathbb{E}U_1) \leq \varepsilon\}$. Using (5.A.7) and the bound (5.A.6) with $t = \frac{\varepsilon}{2}$ for U_2 and $t = \frac{\varepsilon}{4}$ for U_1 one arrives at

$$\begin{aligned} \mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon\right) &\leq \mathbb{P}\left(U_2 - \mathbb{E}U_2 > \frac{\varepsilon}{2}\right) + \mathbb{P}\left(U_1 - \mathbb{E}U_1 < -\frac{\varepsilon}{4}\right) \\ &\leq e^{-\frac{2\lfloor \frac{N}{2} \rfloor \varepsilon^2}{2^2(B-A)^2}} + e^{-\frac{2N\varepsilon^2}{4^2(B-A)^2}} = e^{-\frac{\lfloor \frac{N}{2} \rfloor \varepsilon^2}{2(B-A)^2}} + e^{-\frac{N\varepsilon^2}{8(B-A)^2}}. \end{aligned}$$

To establish the bound in $< -\varepsilon$, we can use a similar union bounding argument as in (5.A.7):

$$\begin{aligned} &\left\{\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) < -\varepsilon\right\} = \{U_2 - \mathbb{E}U_2 - 2(U_1 - \mathbb{E}U_1) < -\varepsilon\} \\ &\stackrel{(b)}{\subseteq} \left\{U_2 - \mathbb{E}U_2 < -\frac{\varepsilon}{2}\right\} \cup \left\{-2(U_1 - \mathbb{E}U_1) < -\frac{\varepsilon}{2}\right\} \\ &= \left\{-U_2 - \mathbb{E}U_2 > \frac{\varepsilon}{2}\right\} \cup \left\{-(U_1 - \mathbb{E}U_1) < -\frac{\varepsilon}{4}\right\}, \end{aligned} \tag{5.A.8}$$

where the inclusion $A \subseteq B \cup C$ in (b) is equivalent to $\bar{B} \cap \bar{C} \subseteq \bar{A}$; the latter holds as $\{U_2 - \mathbb{E}U_2 \geq -\frac{\varepsilon}{2}\} \cap \{-2(U_1 - \mathbb{E}U_1) \geq -\frac{\varepsilon}{2}\} \subseteq \{U_2 - \mathbb{E}U_2 - 2(U_1 - \mathbb{E}U_1) \geq -\varepsilon\}$. Using (5.A.8) and the bound (5.A.6) replacing U_1 by $-U_1$ and U_2 by $-U_2$, with $t = \frac{\varepsilon}{2}$ for U_2 and $t = \frac{\varepsilon}{4}$ for U_1 we arrived at

$$\begin{aligned} &\mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon\right) \\ &\leq \mathbb{P}\left((-U_2) - \mathbb{E}(-U_2) > \frac{\varepsilon}{2}\right) + \mathbb{P}\left((-U_1) - \mathbb{E}(-U_1) < -\frac{\varepsilon}{4}\right) \\ &\leq e^{-\frac{2\lfloor \frac{N}{2} \rfloor \varepsilon^2}{2^2(B-A)^2}} + e^{-\frac{2N\varepsilon^2}{4^2(B-A)^2}} = e^{-\frac{\lfloor \frac{N}{2} \rfloor \varepsilon^2}{2(B-A)^2}} + e^{-\frac{N\varepsilon^2}{8(B-A)^2}}. \end{aligned}$$

□

Proof. (**Theorem 5.3**; concentration of $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$, exponential kernel)

$\text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$ is well-defined since $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} = \mathbb{E}_{x \sim \mathbb{P}} e^{\frac{bx^2}{2}}$ and $\mathbb{E}_{x \sim \mathbb{Q}} \sqrt{K(x, x)} = \mathbb{E}_{x \sim \mathbb{Q}} e^{\frac{bx^2}{2}}$ are finite by assumption (5.3.5), using it with $\lambda = \frac{b}{2}$.

Similarly to the proof of Theorem 5.2, we write the difference between $\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q})$ and $\text{MMD}_K^2(\mathbb{P}, \mathbb{Q})$ in terms of two U -statistics

$$\begin{aligned} &\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) \\ &= \underbrace{\frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j)}_{=: T_2} - \underbrace{\mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x')}_{\mathbb{E}T_2} - 2 \left[\underbrace{\frac{1}{N} \sum_{i \in [N]} \mu_K(\mathbb{Q})(x_i)}_{=: T_1} - \underbrace{\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)}_{\mathbb{E}T_1} \right] \\ &= T_2 - \mathbb{E}T_2 - 2(T_1 - \mathbb{E}T_1) \end{aligned}$$

with T_1 having the kernel $h_1(x) = \mu_K(\mathbb{Q})(x)$ and T_2 using the kernel $h_2(x, x') = K(x, x')$. We establish a concentration result on $T_1 - \mathbb{E}T_1$ and $T_2 - \mathbb{E}T_2$ separately, and combine them with the union bound:

$$\begin{aligned} &\left\{\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon\right\} = \{T_2 - \mathbb{E}T_2 - 2(T_1 - \mathbb{E}T_1) > \varepsilon\} \\ &\subseteq \left\{T_2 - \mathbb{E}T_2 > \frac{\varepsilon}{2}\right\} \cup \left\{-2(T_1 - \mathbb{E}T_1) > \frac{\varepsilon}{2}\right\} \\ &= \left\{T_2 - \mathbb{E}T_2 > \frac{\varepsilon}{2}\right\} \cup \left\{T_1 - \mathbb{E}T_1 < -\frac{\varepsilon}{4}\right\}. \end{aligned} \tag{5.A.9}$$

Below let $p \geq 2$ denote a fixed constant. By assumption $\tilde{N} := \frac{N}{2} \in \mathbb{N}^*$.

- **Bound on T_2 :** Let us introduce the notation V for the sum of independent processes

$$V(x_1, x_2, \dots, x_N) = K(x_1, x_2) + K(x_3, x_4) + \dots + K(x_{N-1}, x_N).$$

With this notation our target quantity T_2 can be rewritten [Pit17] as

$$T_2 = \frac{1}{N(N-1)} \sum_{\substack{i,j \in [N] \\ i \neq j}} K(x_i, x_j) = \frac{2}{N} \left(\frac{1}{N!} \sum_{\sigma \in S_N} V(x_{\sigma_1}, \dots, x_{\sigma_N}) \right), \quad (5.A.10)$$

where S_N denotes the set of permutations of $[N]$. Then for any $t > 0$

$$\begin{aligned} \mathbb{P}(T_2 - \mathbb{E}T_2 > t) &= \mathbb{P}\left(\frac{1}{N(N-1)} \sum_{i \neq j} K(x_i, x_j) - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') > t\right) \\ &\stackrel{(a)}{\leq} \frac{\mathbb{E} \left| \frac{1}{N(N-1)} \sum_{i \neq j} K(x_i, x_j) - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') \right|^p}{t^p} \\ &\stackrel{(b)}{\leq} \frac{\mathbb{E} \left| \frac{2}{N} \frac{1}{N!} \sum_{\sigma \in S_N} V(x_{\sigma_1}, \dots, x_{\sigma_N}) - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') \right|^p}{t^p} \\ &\stackrel{(c)}{\leq} \left(\frac{2}{Nt} \right)^p \mathbb{E} \left| \frac{1}{N!} \sum_{\sigma \in S_N} \left[V(x_{\sigma_1}, \dots, x_{\sigma_N}) - \frac{N}{2} \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') \right] \right|^p \\ &\stackrel{(d)}{\leq} \left(\frac{2}{Nt} \right)^p \frac{1}{N!} \sum_{\sigma \in S_N} \mathbb{E} \left| \underbrace{V(x_{\sigma_1}, \dots, x_{\sigma_N}) - \frac{N}{2} \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x')}_{M_{\tilde{N}}^\sigma} \right|^p \end{aligned} \quad (5.A.11)$$

(a) comes from the generalized Markov's inequality (Lemma 5.C.1) by choosing $\phi(x) := |x|^p$ and $I = \mathbb{R}$. In (b) we applied (5.A.10). Pulling out $(\frac{2}{N})^p$ gives (c). (d) follows from the Jensen inequality by applying it to the argument of the expectation with the convex function $x \mapsto |x|^p$.

Let us introduce the notation $M_{\tilde{N}}^\sigma = V(x_{\sigma_1}, \dots, x_{\sigma_N}) - \frac{N}{2} \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x')$ in (5.A.11). One can expand $M_{\tilde{N}}^\sigma$ as a sum of centered independent processes:

$$\begin{aligned} M_{\tilde{N}}^\sigma &= \underbrace{K(x_{\sigma_1}, x_{\sigma_2}) + K(x_{\sigma_3}, x_{\sigma_4}) + \dots + K(x_{\sigma_{N-1}}, x_{\sigma_N})}_{\tilde{N}=N/2 \text{ terms}} - \tilde{N} \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') \\ &= \sum_{k \in [\tilde{N}]} \underbrace{\left[K(x_{\sigma_{2k-1}}, x_{\sigma_{2k}}) - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x') \right]}_{=: Y_k}. \end{aligned}$$

Similarly, let us denote $M_n^\sigma = \sum_{k \in [n]} Y_k$, $n \in [\tilde{N}]$. By definition $\mathbb{E}Y_k = 0$ for all $k \in [n]$, which implies that M_n^σ is a martingale w.r.t. the filtration $\mathcal{F}_n = \sigma((Y_k)_{k \in [n]})$:

$$\mathbb{E}(M_n^\sigma | \mathcal{F}_{n-1}) = \mathbb{E}\left(\underbrace{M_{n-1}^\sigma}_{\mathcal{F}_{n-1}\text{-measurable}} + \underbrace{Y_n}_{\text{independent from } \mathcal{F}_{n-1}} \mid \mathcal{F}_{n-1} \right) = M_{n-1}^\sigma + \underbrace{\mathbb{E}(Y_n)}_0 = M_{n-1}^\sigma.$$

Hence, we can apply the Burkholder's inequality (Theorem 5.2) on the martingale $\{(M_n^\sigma, \mathcal{F}_n)\}_{n \in \tilde{N}}$; it ensures the existence of a constant $C_p > 0$ such that

$$\mathbb{E} \left| M_{\tilde{N}}^\sigma \right|^p \leq C_p \mathbb{E} \left(\sum_{k \in \tilde{N}} Y_k^2 \right)^{p/2} = C_p (\tilde{N})^{p/2} \mathbb{E} \left(\frac{1}{\tilde{N}} \sum_{k \in \tilde{N}} Y_k^2 \right)^{p/2}$$

$$\stackrel{(a)}{\leq} C_p \left(\tilde{N}\right)^{p/2} \mathbb{E} \left(\underbrace{\frac{1}{\tilde{N}} \sum_{k \in [\tilde{N}]} |Y_k|^p}_{=: m_p = \mathbb{E}_{x, x' \sim \mathbb{P}} |K(x, x') - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x')|^p} \right), \quad (5.A.12)$$

where in (a) we applied the Jensen inequality to the argument of the expectation with the convex function $x \mapsto x^{p/2}$. Moreover, m_p is finite since by assumption (5.3.5) with $\lambda = bp$, $\mathbb{E}_{x \sim \mathbb{P}} |K(x, x)|^p = \mathbb{E}_{x \sim \mathbb{P}} e^{pbx^2} < \infty$ and $\mathbb{E}_{x, x' \sim \mathbb{P}} |K(x, x')|^p \leq \sqrt{\mathbb{E}_{x \sim \mathbb{P}} |K(x, x)|^p \mathbb{E}_{x' \sim \mathbb{P}} |K(x', x')|^p}$. By (5.A.12) we arrive at

$$\mathbb{E} \left| \sum_{k \in \tilde{N}} [K(x_{\sigma_k}, x_{\sigma_{k+1}}) - \mathbb{E}_{x, x' \sim \mathbb{P}} K(x, x')] \right|^p = \mathbb{E} |M_{\tilde{N}}^\sigma|^p \leq C_p m_p \left(\tilde{N}\right)^{p/2}$$

which combined with (5.A.11) gives that for any $t > 0$

$$\mathbb{P}(T_2 - \mathbb{E}T_2 > t) \leq \left(\frac{2}{Nt}\right)^p \frac{1}{N!} \sum_{\sigma \in S_N} C_p m_p \left(\frac{N}{2}\right)^{p/2} = C_p m_p \left(\frac{2}{Nt^2}\right)^{p/2}. \quad (5.A.13)$$

Note: The same bound holds for $-t$

$$\mathbb{P}(T_2 - \mathbb{E}T_2 < -t) \leq C_p m_p \left(\frac{2}{Nt^2}\right)^{p/2}, \quad (5.A.14)$$

by changing K to $-K$ in all the previous steps.

- **Bound on T_1 :** Applying the generalized Markov's inequality (Lemma 5.C.1) with $\phi(x) := |x|^p$, one can bound the probability $\mathbb{P}(T_1 - \mathbb{E}T_1 > t)$ in terms of $\mathbb{E}|T_1 - \mathbb{E}T_1|^p$ for any $t > 0$ as

$$\begin{aligned} \mathbb{P}(T_1 - \mathbb{E}T_1 > t) &= \mathbb{P} \left(\frac{1}{N} \sum_{n \in [N]} \mu_K(\mathbb{Q})(x_n) - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x) > t \right) \\ &\leq \frac{\mathbb{E} \left| \frac{1}{N} \sum_{n \in [N]} \mu_K(\mathbb{Q})(x_n) - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x) \right|^p}{t^p} \\ &= \frac{1}{(Nt)^p} \mathbb{E} \left| \underbrace{\sum_{n \in [N]} [\mu_K(\mathbb{Q})(x_n) - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)]}_{=: S_N} \right|^p. \end{aligned} \quad (5.A.15)$$

S_N is a sum of centered independent random variables and $T_1 - \mathbb{E}T_1 = \frac{1}{N} S_N$. By introducing the notation $Z_k = \mu_K(\mathbb{Q})(x_k) - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)$, $S_n = \sum_{k \in [n]} Z_k$ is a martingale w.r.t. the filtration $\mathcal{F}_n := \sigma((Z_k)_{k \in [n]})$. Hence, one can apply the Burkholder's inequality (Theorem 5.2) on $\{(S_n, \mathcal{F}_n)\}_{n \in [N]}$: it ensures the existence of a constant $C_p > 0$ such that

$$\begin{aligned} \mathbb{E} |S_N|^p &\leq C_p \mathbb{E} \left(\sum_{n \in [N]} Z_n^2 \right)^{p/2} = C_p N^{p/2} \mathbb{E} \left(\frac{1}{N} \sum_{n \in [N]} Z_n^2 \right)^{p/2} \\ &\stackrel{(a)}{\leq} C_p N^{p/2} \mathbb{E} \left(\frac{1}{N} \sum_{n \in [N]} |Z_n|^p \right) = C_p N^{p/2} \mathbb{E} |Z_N|^p \\ &= C_p N^{p/2} \underbrace{\mathbb{E}_{x' \sim \mathbb{P}} |\mu_K(\mathbb{Q})(x') - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)|^p}_{m_p'}, \end{aligned} \quad (5.A.16)$$

where in (a) we applied the Jensen inequality with the convex function $\phi(x) := x^{p/2}$. Let us show the finiteness of m'_p .

Proof (finiteness of m'_p):

- Let us first notice that assumption (5.3.5) (i.e., $\mathbb{E}_{x \sim \mathbb{P}} e^{\lambda x^2} < \infty$ and $\mathbb{E}_{x \sim \mathbb{Q}} e^{\lambda x^2} < \infty$ for all $\lambda \in \mathbb{R}$) implies that

$$\mathbb{E}_{x \sim \mathbb{P}, y \sim \mathbb{Q}} e^{\lambda xy} < \infty \quad \forall \lambda \in \mathbb{R}^+. \quad (5.A.17)$$

Indeed, taking $\lambda' \in \mathbb{R}^+$ and using the inequality $xy \leq \frac{x^2 + y^2}{2}$ for any $x, y \in \mathbb{R}$, one gets

$$\mathbb{E}_{x \sim \mathbb{P}, y \sim \mathbb{Q}} e^{\lambda' xy} \leq \mathbb{E}_{x \sim \mathbb{P}, y \sim \mathbb{Q}} e^{\frac{\lambda'}{2}(x^2 + y^2)}. \quad (5.A.18)$$

By the independence of $x \sim \mathbb{P}$ and $y \sim \mathbb{Q}$, the r.h.s. of (5.A.18) equals to $\mathbb{E}_{x \sim \mathbb{P}} e^{\frac{\lambda'}{2}x^2} \mathbb{E}_{y \sim \mathbb{Q}} e^{\frac{\lambda'}{2}y^2}$ which is finite by using (5.3.5) with $\lambda = \frac{\lambda'}{2}$.

- Let us show that $m'_p \leq 2^p \mathbb{E}_{x \sim \mathbb{P}, y \sim \mathbb{Q}} e^{pbxy}$. Indeed,

$$\begin{aligned} m'_p &= \mathbb{E}_{x' \sim \mathbb{P}} |\mu_K(\mathbb{Q})(x') - \mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)|^p \stackrel{(a)}{\leq} 2^{p-1} \mathbb{E}_{x' \sim \mathbb{P}} (|\mu_K(\mathbb{Q})(x')|^p + |\mathbb{E}_{x \sim \mathbb{P}} \mu_K(\mathbb{Q})(x)|^p) \\ &\stackrel{(b)}{\leq} 2^p \mathbb{E}_{x' \sim \mathbb{P}} |\mu_K(\mathbb{Q})(x')|^p \stackrel{(c)}{=} 2^p \mathbb{E}_{x' \sim \mathbb{P}} |\mathbb{E}_{x \sim \mathbb{Q}} K(x, x')|^p \stackrel{(d)}{\leq} 2^p \mathbb{E}_{x' \sim \mathbb{P}} \mathbb{E}_{x \sim \mathbb{Q}} |K(x, x')|^p \\ &\stackrel{(e)}{=} 2^p \mathbb{E}_{x' \sim \mathbb{P}, x \sim \mathbb{Q}} e^{pbxx'}, \end{aligned} \quad (5.A.19)$$

where (a) follows from the convexity inequality $|a + b|^p \leq 2^{p-1}(|a|^p + |b|^p)$ for any $a, b \in \mathbb{R}$, $p \geq 1$ and the linearity of the integral, in (b) we applied the Jensen inequality with the convex function $\phi(x) := x^p$, in (c) the definition of $\mu_K(\mathbb{Q})(x)$ was used, (d) follows from the Jensen inequality, (e) is implied by the fact that $K(x, x') = e^{bx'}$. Applying (5.A.17) with $\lambda = bp > 0$ (as $b > 0$ and $p \geq 2$) implies that $\mathbb{E}_{x' \sim \mathbb{P}, x \sim \mathbb{Q}} e^{pbxx'} < \infty$ which guarantees the finiteness of m'_p by (5.A.19).

Substituting the bound (5.A.16) to (5.A.15), one gets that for any $t > 0$

$$\mathbb{P}(T_1 - \mathbb{E}T_1 > t) \leq \frac{1}{(Nt)^p} \mathbb{E} |S_N|^p \leq C_p m'_p \frac{1}{(Nt^2)^{p/2}}. \quad (5.A.20)$$

Note: the same bound holds for the deviation below with $-t$

$$\mathbb{P}(T_1 - \mathbb{E}T_1 < -t) \leq C_p m'_p \frac{1}{(Nt^2)^{p/2}} \quad (5.A.21)$$

by changing T_1 to $-T_1$ in the reasoning above.

Using in (5.A.9) the bound (5.A.21) for T_1 with $t = \frac{\varepsilon}{4}$ and the bound (5.A.13) for T_2 with $t = \frac{\varepsilon}{2}$ one gets

$$\begin{aligned} \mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) > \varepsilon\right) &\leq \mathbb{P}\left(T_2 - \mathbb{E}T_2 > \frac{\varepsilon}{2}\right) + \mathbb{P}\left(T_1 - \mathbb{E}T_1 < -\frac{\varepsilon}{4}\right) \\ &\leq C_p m_p \underbrace{\left(\frac{2^2 \times 2}{N\varepsilon^2}\right)^{\frac{p}{2}}}_{\frac{(2\sqrt{2})^p}{\varepsilon^p N^{p/2}}} + C_p m'_p \frac{4^p}{\varepsilon^p N^{p/2}} \leq \frac{1}{\varepsilon^p N^{p/2}} \underbrace{\left[C_p m_p (2\sqrt{2})^p + C_p m'_p 4^p\right]}_{=: C_{p,\mathbb{P},K}}. \end{aligned}$$

The lower deviation bound with $-\varepsilon$ follows by using the bound (5.A.14) for T_2 with $t = \frac{\varepsilon}{2}$ and the bound (5.A.20) for T_1 with $t = \frac{\varepsilon}{4}$

$$\begin{aligned} \mathbb{P}\left(\widehat{\text{MMD}}_{K,e,U}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}_K^2(\mathbb{P}, \mathbb{Q}) < -\varepsilon\right) &\leq \mathbb{P}\left(T_2 - \mathbb{E}T_2 < -\frac{\varepsilon}{2}\right) + \mathbb{P}\left(T_1 - \mathbb{E}T_1 > \frac{\varepsilon}{4}\right) \\ &\leq \frac{C_{p,\mathbb{P},K}}{\varepsilon^p N^{p/2}}. \end{aligned}$$

□

Proof. (**Theorem 5.4**) Let $D = (x_n)_{n \in [N]} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P} \in \mathcal{P}$ and let $\widehat{\text{MMD}}_N$ denote any estimator of $\text{MMD}_K(\mathbb{P}, \mathbb{Q})$ based on D . We are interested in the worst-case error (among all $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$) of the best estimator $\widehat{\text{MMD}}_N$, in other words our target quantity is

$$\inf_{\widehat{\text{MMD}}_N} \sup_{\mathbb{P}, \mathbb{Q} \in \mathcal{P}} \mathbb{P}^N \left(|\widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}, \mathbb{Q})| \geq s \right), \quad s > 0,$$

where \mathbb{P}^N denotes the N -times product measure of \mathbb{P} . Particularly, our goal is to show that $s = \frac{c}{\sqrt{N}}$ is a possible (hence optimal) rate, with some finite constant $c > 0$. Let us define a parametric class of distributions \mathcal{P}_Θ , domain \mathcal{X} , and functional F

$$\begin{aligned} \mathcal{P}_\Theta &:= \{[\mathcal{N}(m, \sigma^2)]^N : (m, \sigma) \in \Theta\}, & \Theta &:= \left\{ (m, \sigma) \in \mathbb{R} \times \left(0, \frac{1}{\sqrt{b}}\right) \right\}, \\ \mathcal{X} &= \mathbb{R}^N, & F(\theta) &= \text{MMD}_K(\mathbb{P}_\theta, \mathbb{Q}) \end{aligned}$$

which we will use to invoke Theorem 5.3. Here $\mathbb{P}_\theta := \mathcal{N}(m_\mathbb{P}, \sigma^2)$ where $\theta := (m_\mathbb{P}, \sigma) \in \Theta$ and a fixed $\mathbb{Q} = \mathcal{N}(m_\mathbb{Q}, \sigma^2)$ is taken with $\theta_\mathbb{Q} = (m_\mathbb{Q}, \sigma) \in \Theta$.⁷ First, let us notice that $\mathcal{P}_\Theta \subset \mathcal{P}$ since $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} < \infty$ means that $\sigma^2 < \frac{1}{b}$.⁸ Using this inclusion one gets the following lower bound (which translated to a lower bound on the target quantity by taking the infimum over $\widehat{\text{MMD}}_N$)

$$\begin{aligned} &\sup_{\mathbb{P}, \mathbb{Q} \in \mathcal{P}} \mathbb{P}^N \left(|\widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}, \mathbb{Q})| \geq s \right) \geq \sup_{\theta, \theta_\mathbb{Q} \in \Theta} \mathbb{P}_\theta^N \left(|\widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}_\theta, \mathbb{Q}_{\theta_\mathbb{Q}})| \geq s \right) \\ &\geq \sup_{\theta \in \Theta} \mathbb{P}_\theta^N \left(|\widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}_\theta, \mathbb{Q}_{\theta_\mathbb{Q}})| \geq s \right), \quad \forall \theta_\mathbb{Q} \in \Theta, \end{aligned} \quad (5.A.22)$$

which means that for any fixed $\theta_\mathbb{Q} \in \Theta$ we are in the realm of Theorem 5.3. To apply the theorem, one needs (i) an upper bound on $D_{\text{KL}}(\mathbb{Q}_{\theta_\mathbb{Q}}^{\otimes N}, \mathbb{P}_\theta^{\otimes N})$, and (ii) a lower bound on $|F(\theta) - F(\theta_\mathbb{Q})|$. This is what we compute in the following.

- **Upper bound on $D_{\text{KL}}(\mathbb{Q}_{\theta_\mathbb{Q}}^{\otimes N}, \mathbb{P}_\theta^{\otimes N})$:** Let p and q denote the pdf of \mathbb{P}_θ and $\mathbb{Q}_{\theta_\mathbb{Q}}$. Then the Kullback-Leibler divergence can be computed as

$$\begin{aligned} D_{\text{KL}}(\mathbb{Q}_{\theta_\mathbb{Q}}^{\otimes N}, \mathbb{P}_\theta^{\otimes N}) &= \int_{\mathbb{R}^N} \log \left(\frac{\prod_{n \in [N]} q(x_n)}{\prod_{n \in [N]} p(x_n)} \right) \prod_{j \in [N]} q(x_j) \, dx_1 \dots dx_N \\ &= \sum_{n \in [N]} \underbrace{\int_{\mathbb{R}^N} \log \left(\frac{q(x_n)}{p(x_n)} \right) \prod_{j \in [N]} q(x_j) \, dx_1 \dots dx_N}_{D_{\text{KL}}(\mathbb{Q}_{\theta_\mathbb{Q}}, \mathbb{P}_\theta)} = \sum_{n \in [N]} \int_{\mathbb{R}} \log \left(\frac{q(x_n)}{p(x_n)} \right) dq(x_n) \prod_{j \in [N], j \neq n} \underbrace{\int_{\mathbb{R}} q(x_j) dx_j}_{=1} \\ &\stackrel{(a)}{=} N \frac{(m_\mathbb{P} - m_\mathbb{Q})^2}{2\sigma^2} = \frac{a^2}{2\sigma^2} =: \alpha, \end{aligned} \quad (5.A.23)$$

where in (a) we used Lemma 5.C.2, and in (b) we assumed that

$$m_\mathbb{P} = m_\mathbb{Q} + \frac{a}{\sqrt{N}} \quad (5.A.24)$$

for some $a > 0$.

⁷Notice that the variance parameter of \mathbb{P}_θ and $\mathbb{Q}_{\theta_\mathbb{Q}}$ are chosen to be identical.

⁸A standard calculation shows that if $\mathbb{P} = \mathcal{N}(m, \sigma)$, $\mathbb{E}_{x \sim \mathbb{P}} \sqrt{K(x, x)} = e^{\frac{m^2}{1-b\sigma^2}} \frac{1}{\sqrt{1-\sigma^2 b}}$ which is finite iff $\sigma^2 < 1/b$.

- **Lower bound on $|F(\theta) - F(\theta_{\mathbb{Q}})|$:** Since $F(\theta_{\mathbb{Q}}) = \text{MMD}_K(\mathbb{Q}, \mathbb{Q}) = 0$, it is sufficient to compute $F(\theta) = \text{MMD}_K(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta_{\mathbb{Q}}})$. By Lemma 5.A.2, one has

$$[F(\theta)]^2 = \frac{1}{\sqrt{1-r^2}} \left[e^{\frac{bm_{\mathbb{P}}^2}{1-r}} + e^{\frac{bm_{\mathbb{Q}}^2}{1-r}} - 2e^{\frac{bc(m_{\mathbb{P}}^2+m_{\mathbb{Q}}^2)+2bm_{\mathbb{P}}m_{\mathbb{Q}}}{2(1-r^2)}} \right] \quad (5.A.25)$$

where $r = b\sigma^2$. We are going to show that

$$[\text{MMD}_K(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta_{\mathbb{Q}}})]^2 \geq \frac{(2K)^2}{N} \quad (5.A.26)$$

for some constant $K > 0$. Let $x = \frac{a}{\sqrt{N}}$, in other words $m_{\mathbb{P}} = m_{\mathbb{Q}} + x$ (in accordance with (5.A.24)); we are going to rewrite the squared MMD in (5.A.26) as a function of $x = m_{\mathbb{P}} - m_{\mathbb{Q}}$. To do so we will apply a Taylor expansion of the squared MMD around $x = 0$. By introducing the notation

$$f_1(x) := e^{\frac{b(m_{\mathbb{Q}}+x)^2}{1-r}} = e^{\frac{bm_{\mathbb{P}}^2}{1-r}}, \quad f_2(x) := e^{\frac{2b(1+c)xm_{\mathbb{Q}}+bcx^2}{2(1-r^2)}},$$

our target quantity writes as

$$[\text{MMD}_K(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta_{\mathbb{Q}}})]^2 = \frac{1}{\sqrt{1-r^2}} \left[f_1\left(\frac{a}{\sqrt{N}}\right) + f_1(0) - 2f_1(0)f_2\left(\frac{a}{\sqrt{N}}\right) \right]. \quad (5.A.27)$$

Indeed, by substituting $m_{\mathbb{P}} = m_{\mathbb{Q}} + x$ in the third term of (5.A.25), one gets

$$\begin{aligned} e^{\frac{bc(m_{\mathbb{P}}^2+m_{\mathbb{Q}}^2)+2bm_{\mathbb{P}}m_{\mathbb{Q}}}{2(1-r^2)}} &= e^{\frac{bc[(m_{\mathbb{Q}}+x)^2+m_{\mathbb{Q}}^2]+2b(m_{\mathbb{Q}}+x)m_{\mathbb{Q}}}{2(1-r^2)}} \\ &= \underbrace{e^{\frac{2bcm_{\mathbb{Q}}^2+2bm_{\mathbb{Q}}^2}{2(1-r^2)}}}_{e^{\frac{2b(c+1)m_{\mathbb{Q}}^2}{2(1-r)(1+c)}}} e^{\frac{bc(2xm_{\mathbb{Q}}+x^2)+2bxm_{\mathbb{Q}}}{2(1-r^2)}} = \underbrace{e^{\frac{bm_{\mathbb{Q}}^2}{1-r}}}_{f_1(0)} \underbrace{e^{\frac{2b(1+c)xm_{\mathbb{Q}}+bcx^2}{2(1-r^2)}}}_{f_2(x)}. \end{aligned}$$

Let us first form the second-order Taylor expansion of f_1 and f_2 around $x = 0$; for this approximation the derivatives are

$$\begin{aligned} f_1'(x) &= f_1(x) \left[\frac{2b}{1-r}(m_{\mathbb{Q}}+x) \right], & f_1''(x) &= f_1(x) \left[\frac{2b}{1-r} + \left[\frac{2b}{1-r}(m_{\mathbb{Q}}+x) \right]^2 \right], \\ f_2'(x) &= f_2(x) \left(\frac{bm_{\mathbb{Q}}}{1-r} + \frac{bcx}{1-r^2} \right), & f_2''(x) &= f_2(x) \left[\frac{bc}{1-r^2} + \left(\frac{bm_{\mathbb{Q}}}{1-r} + \frac{bcx}{1-r^2} \right)^2 \right], \end{aligned}$$

which means that for $x = 0$ one has

$$\begin{aligned} f_1'(0) &= f_1(0) \frac{2bm_{\mathbb{Q}}}{1-r}, & f_1''(0) &= f_1(0) \left[\frac{2b}{1-r} + \left(\frac{2b}{1-r}m_{\mathbb{Q}} \right)^2 \right], \\ f_2'(0) &= \underbrace{f_2(0)}_{=1} \frac{bm_{\mathbb{Q}}}{1-r}, & f_2''(0) &= \underbrace{f_2(0)}_{=1} \left[\frac{bc}{1-r^2} + \left(\frac{bm_{\mathbb{Q}}}{1-r} \right)^2 \right]. \end{aligned}$$

Consequently, the 2nd-order Taylor expansion of f_1 and f_2 takes the form

$$f_1\left(\frac{a}{\sqrt{N}}\right) = f_1(0) + f_1(0) \frac{2bm_{\mathbb{Q}}}{1-r} \frac{a}{\sqrt{N}} + f_1(0) \left[\frac{2b}{1-r} + \left(\frac{2b}{1-r}m_{\mathbb{Q}} \right)^2 \right] \frac{a^2}{2N} + o\left(\frac{a^2}{2N}\right),$$

$$f_2\left(\frac{a}{\sqrt{N}}\right) = 1 + \frac{bm_{\mathbb{Q}}}{1-r} \frac{a}{\sqrt{N}} + \left[\frac{bc}{1-r^2} + \left(\frac{bm_{\mathbb{Q}}}{1-r}\right)^2 \right] \frac{a^2}{2N} + o\left(\frac{a^2}{2N}\right).$$

Using these expansions in (5.A.27), the $f_1(0)$ and $f_1(0) \frac{2bm_{\mathbb{Q}}}{1-r} \frac{a}{\sqrt{N}}$ terms simplify and one gets

$$\begin{aligned} & [\text{MMD}_K(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta_{\mathbb{Q}}})]^2 \\ &= \frac{1}{\sqrt{1-r^2}} f_1(0) \left(\underbrace{\left[\frac{2b}{1-r} + \left(\frac{2bm_{\mathbb{Q}}}{1-r}\right)^2 \right]}_{= \frac{2b}{1-r} - \frac{2rb}{1-r^2} + 2\left(\frac{bm_{\mathbb{Q}}}{1-r}\right)^2} - 2 \left[\frac{rb}{1-r^2} + \left(\frac{bm_{\mathbb{Q}}}{1-r}\right)^2 \right] \right) \frac{a^2}{2N} + o\left(\frac{a^2}{2N}\right) \\ &= \frac{1}{\sqrt{1-r^2}} f_1(0) \left[\frac{2b}{1-r} + 2 \left(\frac{bm_{\mathbb{Q}}}{1-r}\right)^2 \right] \frac{a^2}{2N} + o\left(\frac{a^2}{2N}\right). \end{aligned}$$

This means that the term in $\frac{a^2}{2N}$ will be smaller than the remaining term $o\left(\frac{a^2}{2N}\right)$ for large enough N . Hence there exists a constant $K > 0$ such that

$$[\text{MMD}_K(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta_{\mathbb{Q}}})]^2 \geq \frac{(2K)^2}{N}.$$

Hence we have that

$$|F(\theta) - F(\theta_{\mathbb{Q}})| = |F(\theta)| \geq \frac{2K}{\sqrt{N}} := 2s. \quad (5.A.28)$$

By using the derived bounds (5.A.23) and (5.A.28), Theorem 5.3 can be applied with $\alpha = \frac{a^2}{2\sigma^2}$ and $s = \frac{K}{\sqrt{N}}$, and the bound (5.A.22) implies that

$$\inf_{\widehat{\text{MMD}}_N} \sup_{\mathbb{P}, \mathbb{Q} \in \mathcal{P}} \mathbb{P} \left(|\widehat{\text{MMD}}_N - \text{MMD}_K(\mathbb{P}, \mathbb{Q})| \geq \frac{K}{\sqrt{N}} \right) \geq \max \left(\frac{e^{-\frac{a^2}{2\sigma^2}}}{4}, \frac{1 - \sqrt{\frac{a^2}{2\sigma^2}}}{2} \right).$$

Since the bound is valid for any σ for which $\sigma^2 < \frac{1}{b}$, by continuity one can also take the limit $\sigma^2 = \frac{1}{b}$ for which the lower bound is maximized and writes as

$$\max \left(\frac{e^{-\frac{a^2 b}{2}}}{4}, \frac{1 - \sqrt{\frac{a^2 b}{2}}}{2} \right).$$

□

5.B Implementation tools

This section is about tools which are useful from practical perspective. Section 5.B.1 is about the evaluation of E_i^λ , Section 5.B.3 contains the moments of various distributions, Section 5.B.2 contains score functions and kernel derivatives.

5.B.1 Truncated evaluation of E_i^λ

Lemma 5.B.1 (Infinite sum formulation of E_1^λ and E_2^λ , and truncation error control). *For $z \in [0, 1]$, $a > -1$, $b > -1$, let $E_1^\lambda(z, a, b) = \int_0^z y^a(1-y)^b e^{\lambda y} dy$ and $E_2^\lambda(z, a, b) = \int_z^1 y^a(1-y)^b e^{-\lambda y} dy$. Then we have the following infinite sum formulation:*

$$E_1^\lambda(z, a, b) = \sum_{k=0}^{\infty} \frac{\lambda^k}{k!} B_{inc}(a+k+1, b+1, z),$$

$$E_2^\lambda(z, a, b) = \sum_{k=0}^{\infty} \frac{(-\lambda)^k}{k!} [B(a+k+1, b+1) - B_{inc}(a+k+1, b+1, z)].$$

Let $K \in \mathbb{N}$ be fixed and let us denote the truncated E_1^λ and E_2^λ by

$$E_1^{\lambda, tr} = \sum_{k=0}^K \frac{\lambda^k}{k!} B_{inc}(a+k+1, b+1, z),$$

$$E_2^{\lambda, tr} = \sum_{k=0}^K \frac{(-\lambda)^k}{k!} [B(a+k+1, b+1) - B_{inc}(a+k+1, b+1, z)].$$

Then the following bounds hold for the truncation errors

$$E_1^\lambda - E_1^{\lambda, tr} \leq \frac{\lambda^{K+1} e^{\lambda z}}{(K+1)!} B_{inc}(a+K+2, b+1, z) := \mathcal{E}_1^{\lambda, tr},$$

$$E_2^\lambda - E_2^{\lambda, tr} \leq \frac{(-\lambda)^{K+1} e^{-\lambda z}}{(K+1)!} [B(a+K+2, b+1) - B_{inc}(a+K+2, b+1, z)] := \mathcal{E}_2^{\lambda, tr}.$$

Remarks:

- Since the first two arguments of the beta and incomplete beta functions are positive, these functions are bounded by 1 and the ratio $\frac{\lambda^{K+1}}{(K+1)!}$ drives the error bounds \mathcal{E}_1^{tr} and \mathcal{E}_2^{tr} . By the Stirling formula $(K+1)! = \mathcal{O}\left((K+1)^{(K+\frac{3}{2})}\right)$. This means that to control \mathcal{E}_1^{tr} and \mathcal{E}_2^{tr} , K should be chosen at least of the same order of magnitude than λ .
- We show in Fig. 5.19 the evolution of $E_1^{\lambda, tr}$ and $E_2^{\lambda, tr}$ for different values of λ and a as a function of the level of truncation K . As anticipated in the previous remark, the sum starts to converge for K of the same order of magnitude of λ . The function $E_2^{\lambda, tr}$ takes a few more iterations to converge.

Proof. (**Lemma 5.B.1**) For $z \in [0, 1]$, $a > -1$ and $b > -1$, let

$$E_1^\lambda(z, a, b) = \int_0^z y^a(1-y)^b e^{\lambda y} dy, \quad E_2^\lambda(z, a, b) = \int_z^1 y^a(1-y)^b e^{-\lambda y} dy.$$

The infinite sum formulations follow from the exponential series expansions $e^{\lambda y} = \sum_{k=0}^{\infty} \frac{(\lambda y)^k}{k!}$ and $e^{-\lambda y} = \sum_{k=0}^{\infty} \frac{(-\lambda y)^k}{k!}$:

$$E_1^\lambda(z, a, b) = \int_0^z y^a(1-y)^b e^{\lambda y} dy = \int_0^z y^a(1-y)^b \sum_{k=0}^{\infty} \frac{(\lambda y)^k}{k!} dy = \sum_{k=0}^{\infty} \frac{\lambda^k}{k!} \underbrace{\int_0^z y^{a+k}(1-y)^b dy}_{B_{inc}(a+k+1, b+1, z)},$$

$$E_2^\lambda(z, a, b) = \int_z^1 y^a(1-y)^{b+1} e^{-\lambda y} dy = \sum_{k=0}^{\infty} \frac{(-\lambda)^k}{k!} \int_z^1 y^{a+k}(1-y)^b dy$$




Figure 5.19 – Evolution of $E_i^{\lambda, tr}$ for $i \in [2]$ and for various truncation levels (K).

$$= \sum_{k=0}^{\infty} \frac{(-\lambda)^k}{k!} \left[\underbrace{\int_0^1 y^{a+k}(1-y)^b dy}_{B(a+k+1, b+1)} - \underbrace{\int_0^z y^{a+k}(1-y)^b dy}_{B_{\text{inc}}(a+k+1, b+1, z)} \right].$$

Let us now fix $K \in \mathbb{N}$, and truncate $E_1^\lambda(z, a, b)$ and $E_2^\lambda(z, a, b)$ to the first $K+1$ terms:

$$E_1^{\lambda, tr} = \sum_{k=0}^K \frac{\lambda^k}{k!} B_{\text{inc}}(a+k+1, b+1, z),$$

$$E_2^{\lambda, tr} = \sum_{k=0}^K \frac{(-\lambda)^k}{k!} [B(a+k+1, b+1) - B_{\text{inc}}(a+k+1, b+1, z)].$$

By the Taylor-Lagrange theorem, in case of

- E_1^λ : for any $y \in [0, z]$ there is a $y_K \in (0, y)$ such that

$$e^{\lambda y} = \sum_{k=0}^K \frac{(\lambda y)^k}{k!} + e^{\lambda y_K} \frac{(\lambda y)^{K+1}}{(K+1)!}.$$

- E_2^λ : for any $y \in [z, 1]$ there is a $y'_K \in (0, y)$ such that

$$e^{-\lambda y} = \sum_{k=0}^K \frac{(-\lambda y)^k}{k!} + e^{-\lambda y'_K} \frac{(-\lambda y)^{K+1}}{(K+1)!}.$$

Hence, the truncation errors can be bounded as

$$E_1^\lambda - E_1^{\lambda, tr} = \int_0^z y^a (1-y)^b e^{\lambda y_K} \frac{(\lambda y)^{K+1}}{(K+1)!} dy \stackrel{(a)}{\leq} \int_0^z y^a (1-y)^{-1/2} e^{\lambda z} \frac{(\lambda y)^{K+1}}{(K+1)!} dy$$

$$= \frac{\lambda^{K+1} e^{\lambda z}}{(K+1)!} \underbrace{\int_0^z y^{a+K+1} (1-y)^b dy}_{B_{\text{inc}}(a+K+2, b+1, z)},$$

$$E_{21}^\lambda - E_{21}^{\lambda, tr} = \int_z^1 y^a (1-y)^b e^{-\lambda y'_K} \frac{(-\lambda y)^{K+1}}{(K+1)!} dy \stackrel{(b)}{\leq} \int_z^1 y^a (1-y)^b e^{-\lambda z} \frac{(-\lambda y)^{K+1}}{(K+1)!} dy$$

$$= \frac{(-\lambda)^{K+1}}{(K+1)!} e^{-\lambda z} \int_z^1 y^{a+K+1} (1-y)^b dy$$

$$= \frac{(-\lambda)^{K+1}}{(K+1)!} e^{-\lambda z} \left[\underbrace{\int_0^1 y^{a+K+1} (1-y)^b dy}_{B(a+K+2, b+1)} - \underbrace{\int_0^z y^{a+K+1} (1-y)^b dy}_{B_{\text{inc}}(a+K+2, b+1, z)} \right]$$

$$= \frac{(-\lambda)^{K+1}}{(K+1)!} e^{-\lambda z} [B(a+K+2, b+1) - B_{\text{inc}}(a+K+2, b+1, z)],$$

where in (a) we used that $e^{\lambda y_K} \leq e^{\lambda z}$ and in (b) that $e^{-\lambda y'_K} \leq e^{-\lambda z}$. \square

5.B.2 Score functions and kernel derivatives

The KSD and FSSD measures (see (5.2.4) and (5.2.5)) and their empirical counterparts rely on the functions ξ_q and h_q functions which can be readily obtained from the score function $s_q(x)$, kernel $K(x, y)$ and the associated derivatives $\frac{\partial K(x, y)}{\partial x}$ and $\frac{\partial^2 K(x, y)}{\partial x \partial y}$. Table 5.10 and Table 5.11 provide these ingredients.

Distribution	$q(x)$	$s_q(x)$	Parameters
Gaussian	$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-m)^2}{2\sigma^2}}$	$-\frac{x-m}{\sigma^2}$	$m \in \mathbb{R}, \sigma \in \mathbb{R}^+$
GN ⁹	$\frac{2^{-\frac{\gamma+1}{\gamma}} \gamma}{\beta \Gamma(1/\gamma)} e^{-\frac{1}{2} \left(\frac{ x-\alpha }{\beta}\right)^\gamma}$	$-\frac{\gamma x-\alpha ^{\gamma-1}}{2\beta\gamma} \text{sign}(x-\alpha)$	$\beta, \gamma \in \mathbb{R}^+, \alpha \in \mathbb{R}$
SG ¹⁰	$2(2\pi v)^{-1/2} e^{-\frac{(x-m)^2}{2v}} \Phi\left(\frac{s(x-m)}{\sqrt{v}}\right)$	$-\frac{x-m}{v} + \frac{e^{-\frac{s(x-m)}{\sqrt{v}}}}{\Phi\left(\frac{s(x-m)}{\sqrt{v}}\right)} \frac{s}{\sqrt{2\pi v}}$	$s, m \in \mathbb{R}, v \in \mathbb{R}^+$
beta	$\frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1} \mathbb{1}_{[0,1]}(x)$	$\left(\frac{\alpha-1}{x} + \frac{\beta-1}{x-1}\right) \mathbb{1}_{[0,1]}(x)$	$\alpha, \beta \in \mathbb{R}^+$

Table 5.10 – Distributions and their score functions.

Kernel	$K(x, y)$	$\frac{\partial K(x, y)}{\partial x}$	$\frac{\partial^2 K(x, y)}{\partial x \partial y}$	Parameters
Gaussian	$e^{-\frac{(x-y)^2}{2c^2}}$	$e^{-\frac{(x-y)^2}{2c^2}} \left(-\frac{x-y}{c^2}\right)$	$e^{-\frac{(x-y)^2}{2c^2}} \left[-\frac{(x-y)^2}{c^4} + \frac{1}{c^2}\right]$	$c \in \mathbb{R}^+$
Laplacian ^a	$e^{-\lambda x-y }$	$e^{-\lambda x-y } [-\lambda \text{sign}(x-y)]$	$-\lambda^2 e^{-\lambda x-y }$	$\lambda \in \mathbb{R}^+$

Table 5.11 – Kernels and their derivatives.

^aIn columns 2-3, it is assumed that $x \neq y$.

5.B.3 Moments of various distributions

Table 5.12 contains moments expressed as a function of the parameters for various target distributions.

5.C External statements

This section contains external statements used in our proofs.

Lemma 5.C.1 (Generalized Markov’s inequality; (2.1) in [BLM13]). *Let ϕ denote a nondecreasing and nonnegative function defined on $I \subseteq \mathbb{R}$ and let Y denote a random variable taking*

Distribution (\mathbb{Q})	$m_{\mathbb{Q}}$	$\sigma_{\mathbb{Q}}^2$	$\zeta_{\mathbb{Q}}$	$\kappa_{\mathbb{Q}}$
Gaussian	m	σ^2	0	0
GN	α	$\beta^2 4^{1/\gamma} \frac{\Gamma(3/\gamma)}{\Gamma(1/\gamma)}$	0	$\frac{\Gamma(5/\gamma)\Gamma(1/\gamma)}{\Gamma(3/\gamma)^2} - 3^a$
SG ^b	$m + \delta \sqrt{\frac{2v}{\pi}}$	$v \left(1 - \frac{2\delta^2}{\pi}\right)$	$\frac{4-\pi}{2} \frac{(\delta\sqrt{2/\pi})^3}{(1-2\delta^2/\pi)^{3/2}}$ ^c	$2(\pi-3) \frac{(\delta\sqrt{2/\pi})^4}{(1-2\delta^2/\pi)^2}$
beta	$\frac{\alpha}{\alpha+\beta}$	$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$	$2 \frac{(\beta-\alpha)\sqrt{\alpha+\beta+1}}{(\alpha+\beta+2)\sqrt{\alpha\beta}}$	$6 \frac{(\beta-\alpha)^2(\alpha+\beta+1) - \alpha\beta(\alpha+\beta+2)}{\alpha\beta(\alpha+\beta+2)(\alpha+\beta+3)}$

^aIn the GN case, the excess kurtosis $\kappa_{\mathbb{Q}}$ is negative if $\gamma > 2$.

^bWe denote $\delta = \frac{s}{\sqrt{1+s^2}}$ where s is one of the parameters of the SG distribution; its sign indicates if the distribution has a positive or a negative excess skewness.

^cIn the SG case, the excess skewness $\zeta_{\mathbb{Q}}$ is positive if $s > 0$.

Table 5.12 – Moments of the target distributions.

values in I . Then Markov's inequality implies that for every $t \in I$ with $\phi(t) > 0$

$$\mathbb{P}(Y \geq t) \leq \frac{\mathbb{E}\phi(Y)}{\phi(t)}.$$

Theorem 5.1 (Hoeffding inequality for U-statistic; [Hoe63], [Pit17]). Assume we have n i.i.d. samples $\{X_i\}_{i \in [n]} \sim \mathbb{P}$. Let I_m^n be the set m -tuples chosen without repetition from $[n]$. Suppose that $h : \mathbb{R}^m \rightarrow \mathbb{R}$ is bounded: $a \leq h(x_1, \dots, x_m) \leq b$ for all (x_1, \dots, x_m) . We denote $m_h = \mathbb{E}h(X_1, \dots, X_m)$ and its U-statistic based estimator $U_n = \frac{1}{\binom{n}{m}} \sum_{(i_1, \dots, i_m) \in I_m^n} h(X_{i_1}, \dots, X_{i_m})$. Then, for any $\varepsilon > 0$

$$\mathbb{P}(U_n - m_h > \varepsilon) \leq e^{-\frac{2\lfloor \frac{n}{m} \rfloor \varepsilon^2}{(b-a)^2}},$$

and the same deviation bound holds for $-\varepsilon$ below, i.e. $\mathbb{P}(U_n - m_h < -\varepsilon) \leq e^{-\frac{2\lfloor \frac{n}{m} \rfloor \varepsilon^2}{(b-a)^2}}$.

Theorem 5.2 (Burkholder's inequality; Theorem 2.10 in [HH80]). Assume that $\{(S_i, \mathcal{F}_i)\}_{i \in [n]}$ is a martingale sequence and its filtration, $1 < p < \infty$. Let the associated martingale increments be denoted by $X_1 = S_1$ and $X_i = S_i - S_{i-1}$, $2 \leq i \leq n$. Then there exist a constant C_p depending on p such that

$$\mathbb{E} \left| S_n \right|^p \leq C_p \mathbb{E} \left| \sum_{i=1}^n X_i^2 \right|^{p/2}.$$

Lemma 5.C.2 (Kullback-Leibler divergence for univariate Gaussian variables; page 13 in [Duc07]). Let $\mathbb{P} = \mathcal{N}(m_{\mathbb{P}}, \sigma_{\mathbb{P}})$, $\mathbb{Q} = \mathcal{N}(m_{\mathbb{Q}}, \sigma_{\mathbb{Q}})$, $m_{\mathbb{P}}, m_{\mathbb{Q}} \in \mathbb{R}$, $\sigma_{\mathbb{P}}, \sigma_{\mathbb{Q}} \in \mathbb{R}^{>0}$. Then

$$D_{KL}(\mathbb{P}, \mathbb{Q}) = \log \left(\frac{\sigma_{\mathbb{Q}}}{\sigma_{\mathbb{P}}} \right) + \frac{\sigma_{\mathbb{P}}^2 + (m_{\mathbb{Q}} - m_{\mathbb{P}})^2}{2\sigma_{\mathbb{Q}}^2} - \frac{1}{2}.$$

Theorem 5.3 (Theorem 2.2 in [Tsy09]). Let \mathcal{X} and Θ denote two measurable spaces. Let $F : \Theta \rightarrow \mathbb{R}$ be a functional. Let $\mathcal{P}_{\Theta} = \{\mathbb{P}_{\theta} : \theta \in \Theta\}$ be a class of probability measures on \mathcal{X} indexed by Θ . We observe the data D distributed according $\mathbb{P}_{\theta} \in \mathcal{P}_{\Theta}$ with some unknown θ . The goal is to estimate $F(\theta)$. Let $\hat{F} := \hat{F}(D)$ be an estimator of $F(\theta)$ based on D . Assume that there exist $\theta_0, \theta_1 \in \Theta$ such that $|F(\theta_0) - F(\theta_1)| \geq 2s > 0$ and $D_{KL}(\mathbb{P}_{\theta_1}, \mathbb{P}_{\theta_0}) \leq \alpha$ with $0 < \alpha < \infty$. Then

$$\inf_{\hat{F}} \sup_{\theta \in \Theta} \mathbb{P}_{\theta} \left(|\hat{F} - F(\theta)| \geq s \right) \geq \max \left(\frac{e^{-\alpha}}{4}, \frac{1 - \sqrt{\alpha/2}}{2} \right).$$

Note: Typically \mathcal{X} , D and \mathbb{P}_{θ} depend on the sample size N .

LIST OF FIGURES

1	Plot of daily returns of the S&500 and the CAC40 and the Gaussian distribution of the same mean and variance.	11
2	Quantile-quantile plots of the daily returns of the S&500 and the CAC40.	12
3	Autocorrelation plot on daily returns of the S&P500 and the CAC40 in absolute value.	12
4	Plot of volatility measured by sliding window on the daily returns of the S&P500 and the CAC40.	13
5	Higher moment p_{\max} on the variance of returns calibrated with the GARCH model.	13
6	Quantile quantile plot of correlation calibrated by sliding window between two assets.	14
7	$\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^{\infty})$, $d = 50$ and $p_{\max} = 5.05$	18
8	Evolution of the total error for the Sharpe ratio approximation.	20
9	Distribution of portfolio returns based on traditional approaches (minimum variance, maximum Sharpe ratio) versus minimum MMD approach with skew Gaussian target.	23
10	Histogrammes des rendements journaliers du S&P500 et du CAC40 et de la distribution Gaussienne de même moyenne et variance.	25
11	Diagrammes quantile-quantile des rendements journaliers du S&P500 et du CAC40.	26
12	Autocorrélogramme des rendements journaliers en valeur absolue du S&P500 et du CAC40.	26
13	Tracé de la volatilité mesurée par fenêtre glissante sur les rendements journaliers du S&P500 et du CAC40.	27
14	Plus haut moment p_{\max} sur la variance de rendements calibrés avec le modèle GARCH.	28
15	Diagramme quantile quantile: quantiles de la corrélation calibrée par fenêtre glissante entre deux actifs en fonction des quantiles de la loi uniforme.	28
16	$\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^{\infty})$, $d = 50$ and $p_{\max} = 5.05$	33
17	Évolution de l'erreur totale sur l'approximation du ratio de Sharpe	34
18	Distribution des rendements de portefeuilles basés sur des approches traditionnelles versus approche MMD minimum avec cible skew Gaussienne.	36
1.1	Ψ_{α} for different α values.	45
1.2	Relation of the spectral density examples of Table 1.2.	50
3.1	r_t and V_{t+1} time-dependence	89

3.2	GARCH(1,1) parameters stationary and finite variance area.	112
3.3	$\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^\infty)$, $p_{\max} = 2.8$ versus $p_{\max} = 7$	114
3.4	$\widehat{\mathbb{P}}(R_{N,H} > R_{N,H}^\infty)$, $d = 50$ versus $d = 10$	115
3.5	Histograms of $\left\{ \frac{R_{N,H}^\infty - R_{N,H}}{NH} \right\}_{N \in \{1,6,9,12\}}$	116
3.6	$\widehat{\mathbb{P}}(R_{1,H} < R_{1,H}^\infty)$ as a function of H	117
3.7	Histogram of $\frac{D_{N,H}}{\sqrt{N}}$ and Hill estimation of $\frac{1}{p_{\max}}$	118
4.1	Plot of $SR(\rho)$ as a function of the correlation ρ	144
4.2	Empirical quantile function for 4 pairs of correlation estimated on $T = 5$ years of data, with a sliding window $W = 6$ months.	145
4.3	Beta incomplete function plot	146
4.4	Plot of convex combination of Beta incomplete functions	146
4.5	Plot of $\theta \mapsto f_\beta(\theta)$ and $\theta \mapsto SR(f_\beta(\theta))$	147
4.6	Evolution of $SR(f(\theta))$ approximation for different values of m and of the projection coefficients.	147
4.7	Truncation error and the associated tendency line, and parameters frontier.	148
4.8	Total error, increasing versus fixed m_k	149
4.9	Total error and truncation error, increasing versus fixed m_k	150
4.10	Approximation of the Sharpe ratio - mixture uncertainty, increasing number of iterations, $b = 0.15$	151
4.11	Approximation of the Sharpe ratio - mixture uncertainty, increasing number of iterations, $b = 0.3$	152
4.12	Total error as a function of the number of iterations, for different values of a	153
4.13	Illustration of the integral approach.	158
4.14	Evolution of allocation \mathbf{w}^* components as a function of ρ	162
4.15	Evolution of the coefficients \mathbf{u}_i^* of $\ell(\mathbf{y}^*(\theta))$ projection on the Legendre basis for different transformation functions.	163
4.16	Truncation error of portfolio integral transformation.	164
4.17	Evolution of $\mathbf{w}^*(\theta)$ approximation as a function of $\rho = f_\beta(\theta)$	165
4.18	True function approximation with UQSA, obtained after transformation on \mathbf{y}^* approximation.	166
4.19	Total error evolution with UQSA, obtained after transformation on \mathbf{y}^* approximation.	166
4.20	$\mathbf{y}^*(\theta)$ approximation with UQSA.	167
5.1	Relation of the kernels in Table 5.3.	181
5.2	Relation of the target distributions in Table 5.4.	181
5.3	Boxplots of the objective values for Gaussian returns (\mathbb{P}_R) and Gaussian kernel (K).	189
5.4	Boxplots of the objective values for Gaussian return (\mathbb{P}_R) and exponential kernel (K).	190
5.5	Comparison of the histogram of the portfolio returns $(\mathbf{w}^*)^\top R$ (with R having beta components) and one-dimensional beta target distribution.	190
5.6	Boxplots of the objective values for beta return (\mathbb{P}_R) and Laplacian kernel (K).	191
5.7	Plots of randomly simulated $(\Delta\mu, \Delta\sigma)$ pairs colored by their associated sign of $\frac{\Delta x}{\Delta \gamma}$	193
5.8	Price evolution associated to the first and last component of the 6BTM datasets.	194
5.9	Variance of the component returns of the 6BTM (a) and the 10Ind dataset (b).	195
5.10	Superposition of the component returns of the 6BTM and the 10Ind dataset.	195
5.11	OOS kurtosis values in the low-volatility regime.	197
5.12	The five top objective-target pairs in kurtosis sense.	198

5.13	OOS kurtosis boxplots in the high-volatility regime.	199
5.14	OOS kurtosis values for minimum divergence portfolios with the Gaussian target.	199
5.15	OOS skewness values in the low-volatility regime.	200
5.16	The five top objective-target pairs in skewness sense.	200
5.17	OOS skewness boxplots in the high-volatility regime.	201
5.18	The five top objective-target pairs. Regime: low and high-volatility merged.	203
5.19	Evolution of $E_i^{\lambda, tr}$ for $i \in [2]$ and for various truncation levels (K).	217

LIST OF TABLES

1.1	Summary of RFF guarantees on kernel values and derivatives.	43
1.2	Kernel spectrum examples in one dimension ($d = 1$) obeying the α -exponential Orlicz assumption.	49
1.3	Kernel examples for the spectral densities given in Table 1.2.	51
1.4	Definition of special functions.	67
3.1	GARCH(1,1) parameters summary.	110
5.1	Semi-explicit MMD, KSD, FSD and WAD estimators.	178
5.2	Semi-explicit MMD, KSD, FSD and WAD estimators and their complexity.	178
5.3	Kernel definitions; their relation is illustrated in Fig. 5.1.	181
5.4	Target distributions; their relation is illustrated in Fig. 5.2.	181
5.5	Summary of obtained analytical mean embedding results.	183
5.6	Summary of the investigated target distribution (\mathbb{P}_T) - kernel (K) pairs, with the associated analytical mean embeddings.	187
5.7	Financial datasets used in our empirical study.	194
5.8	State-of-the-art portfolios [Las19].	196
5.9	Tested divergence - target distribution pairs.	196
5.10	Distributions and their score functions.	219
5.11	Kernels and their derivatives.	219
5.12	Moments of the target distributions.	219

BIBLIOGRAPHY

- [Ada08] R. Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electronic Journal of Probability*, 13:1000–1034, 2008.
- [Alz97] H. Alzer. On some inequalities for the incomplete Gamma function. *Mathematics of Computation of the American Mathematical Society*, 66(218):771–778, 1997.
- [AM15] A. E. Alaoui and M. Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Advances in Neural Information Processing Systems (NIPS)*, pages 775–783, 2015.
- [Aro50] N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68:337–404, 1950.
- [Asm03] S. Asmussen. *Applied probability and queues*. Springer, second edition, 2003.
- [Bac00] L. Bachelier. Théorie de la spéculation. *Annales Scientifiques de l’École Normale Supérieure. Troisième Série*, 17:21–86, 1900.
- [Bac17] F. Bach. On the equivalence between kernel quadrature rules and random feature expansions. *Journal of Machine Learning Research*, 18:1–38, 2017.
- [Bar13] P. Bartlett. Theoretical statistics. Available at <https://www.stat.berkeley.edu/~bartlett/courses/2013spring-stat210b/>, 2013.
- [BBD⁺19] A. Barp, F.-X. Briol, A. B. Duncan, M. Girolami, and L. Mackey. Minimum Stein discrepancy estimators. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 12964–12976, 2019.
- [BBDG19] F.-X. Briol, A. Barp, A. B. Duncan, and M. Girolami. Statistical inference for generative models with maximum mean discrepancy. Technical report, 2019. (<https://arxiv.org/abs/1906.05944>).
- [BBLdPZ16] D. H. Bailey, J. Borwein, M. López de Prado, and Q. J. Zhu. The probability of backtest overfitting. *Journal of Computational Finance*, 20:39–69, 2016.
- [BCC⁺17] M. Bojarski, A. Choromanska, K. Choromanski, F. Fagan, C. Gouy-Pailler, A. Morvan, N. Sakr, T. Sarlós, and J. Atif. Structured adaptive and random spinners for fast machine learning computations. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 1020–1029, 2017.

- [BD16] P. J. Brockwell and R. A. Davis. *Introduction to time series and forecasting*. Springer, third edition, 2016.
- [BEKL18] G.-Y. Ban, N. El Karoui, and A. E. B. Lim. Machine learning and portfolio optimization. *Management Science*, 64(3):1136–1154, 2018.
- [Ben19] E. Benhamou. Connecting Sharpe ratio and Student t-statistic, and beyond. Technical report, 2019.
- [BF19] L. Boudabsa and D. Filipović. Machine learning with kernels for portfolio valuation and risk management. Technical report, EPFL, 2019. (Swiss Finance Institute Research Paper No. 19-34; <https://dx.doi.org/10.2139/ssrn.3401539>).
- [BG91] M. J. Best and R. R. Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342, 1991.
- [BGP20] F. Bernal, E. Gobet, and J. Printems. Volatility uncertainty quantification in a stochastic control problem applied to energy. *Methodology and Computing in Applied Probability*, 22(1):135–159, 2020.
- [BJ02] F. Bach and M. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- [BL91] F. Black and R. Litterman. Asset allocation: combining investor views with market equilibrium. *The Journal of Fixed Income*, 1(2):7–18, 1991.
- [BLM13] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013.
- [BLR06] L. Bauwens, S. Laurent, and J. V. K. Rombouts. Multivariate GARCH models: a survey. *Journal of Applied Econometrics*, 21(1):79–109, 2006.
- [BLY21] K. Balasubramanian, T. Li, and M. Yuan. On the optimality of kernel-embedding based goodness-of-fit tests. *Journal of Machine Learning Research*, 22(1):1–45, 2021.
- [BNMR01] O. E. Barndorff-Nielsen, T. Mikosch, and S. I. Resnick. *Lévy Processes – Theory and Applications*. Springer, 2001.
- [Bol90] T. Bollerslev. Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *The Review of Economics and Statistics*, 72(3):498–505, 1990.
- [Bou02] O. Bousquet. *Concentration inequalities and empirical processes theory applied to the analysis of learning algorithms*. PhD thesis, École Polytechnique: Department of Applied Mathematics Paris, France, 2002.
- [Bre17] J. Brehmer. Elicitability and its Application in Risk Management. Master’s thesis, School of Business Informatics and Mathematics, July 2017.
- [Bro93] M. Broadie. Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, 45(1):21–58, 1993.
- [BS73] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

- [BS00] J. F. Bonnans and A. Shapiro. *Perturbation analysis of optimization problems*. Springer, 2000.
- [BSAG18] M. Binkowski, D. Sutherland, M. Arbel, and A. Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations (ICLR)*, 2018.
- [BTA04] A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer, 2004.
- [Car01] A. Carol. *Market models : a guide to financial data analysis*. John Wiley & Sons, 2001.
- [CFG20] S. Crépey, G. Fort, E. Gobet, and U. Stazhynski. Uncertainty quantification for stochastic approximation limits using chaos expansion. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3):1061–1089, 2020.
- [CGL21a] L. Chamakh, E. Gobet, and J.-P. Lemor. Asymptotic analysis of different covariance matrices estimation for minimum variance portfolio. (https://hal.archives-ouvertes.fr/hal-03207061/file/GARCH_min_var_final_v0.pdf), April 2021.
- [CGL21b] L. Chamakh, E. Gobet, and W. Liu. Orlicz norms and concentration inequalities for β -heavy tailed random variables. (https://hal.archives-ouvertes.fr/hal-03175697v2/file/Orlicz_beta_heavy_tail_versionEJS_v3_HAL.pdf), April 2021.
- [CGS20] L. Chamakh, E. Gobet, and Z. Szabó. Orlicz random Fourier feature. *Journal of Machine Learning Research*, 21:1–37, 2020.
- [Cha11] N. Chapados. *Portfolio choice problems: An introductory survey of single and multiperiod models*. Springer, 2011.
- [CLM97] J. Y. Campbell, A. W. Lo, and A. C. MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, 1997.
- [CMS11] K. Chaudhuri, C. Monteleoni, and A. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.
- [Con01] R. Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.
- [CP09] R. Ciunara and V. Preda. The Weibull-logarithmic distribution in lifetime analysis and its properties. In *International Conference Applied Stochastic Models and Data Analysis (ASMDA)*, pages 395–399. 2009.
- [CRR18] L. Carratino, A. Rudi, and L. Rosasco. Learning with SGD and random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 10192–10203, 2018.
- [CSG16] K. Chwialkowski, H. Strathmann, and A. Gretton. A kernel test of goodness of fit. In *International Conference on Machine Learning (ICML)*, pages 2606–2615, 2016.
- [CT83] S. Csörgö and V. Totik. On how long interval is the empirical characteristic function uniformly consistent? *Acta Scientiarum Mathematicarum*, 45:141–149, 1983.

- [CW12] Y. Chalabi and D. Wuertz. Portfolio optimization based on divergence measures. Technical report, 2012. (<https://mpira.ub.uni-muenchen.de/43332/>).
- [CZ86] H. F. Chen and Y. M. Zhu. Stochastic approximation procedures with randomly varying truncations. *Scientia Sinica. Series A. Mathematical, Physical, Astronomical & Technical Sciences*, 29(9):914–926, 1986.
- [DHB⁺17] C. Downey, A. Hefny, B. Boots, B. Li, and G. Gordon. Predictive state recurrent neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6053–6064, 2017.
- [DHF06] L. De Haan and A. Ferreira. *Extreme value theory: an introduction*. Springer, 2006.
- [DM05] P. Drineas and M. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005.
- [DNLW18] G. De Nard, O. Ledoit, and M. Wolf. Factor models for portfolio selection in large dimensions: the good, the better and the ugly. Technical report, the University of Zurich, 2018. (Journal of Financial Econometrics Working Paper No. 290; https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3194492).
- [DRG15] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 258–267, 2015.
- [DTPL14] R. Duijkers, R. Tóth, D. Piga, and V. Laurain. Shrinking complexity of scheduling dependencies in LS-SVM based LPV system identification. In *IEEE Conference on Decision and Control (CDC)*, pages 2561–2566, 2014.
- [Duc07] J. Duchi. Derivations for linear algebra and optimization. Technical report, The University of Stanford, 2007. (https://web.stanford.edu/~jduchi/projects/general_notes.pdf).
- [DXH⁺14] B. Dai, B. Xie, N. He, Y. Liang, A. Raj, M.-F. Balcan, and L. Song. Scalable kernel methods via doubly stochastic gradients. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3041–3049, 2014.
- [EB86] R. F. Engle and T. Bollerslev. Modelling the persistence of conditional variances. *Econometric Reviews*, 5(1):1–87, 1986. With comments and a reply by the authors.
- [EK18] N. El Karoui. Random matrices and high-dimensional statistics: beyond covariance matrices. In *International Congress of Mathematicians (ICM)*, pages 2875–2894, 2018.
- [EKT15] S. Emmer, M. Kratz, and D. Tasche. What is the best risk measure in practice? a comparison of standard measures. *Journal of Risk*, 18(2):31–60, 2015.
- [EP07] R. F. Engle and A. J. Patton. What good is a volatility model? In *Forecasting volatility in the financial markets*, pages 47–63. Elsevier, 2007.
- [FAGdHN07] M. I. Fraga Alves, M. I. Gomes, L. de Haan, and C. Neves. A note on second order conditions in extreme value theory: linking general and heavy tail conditions. *REVSTAT Statistical Journal*, 5(3):285–304, 2007.

- [Fam71] E. F. Fama. Efficient capital markets: a review of theory and empirical work. In *Frontiers of quantitative economics (Invited Papers, Econometric Soc. Winter Meetings, New York, 1969)*, pages 309–361. Contributions to Economic Analysis, Vol. 71. 1971. With comments by William F. Sharpe and Robert A. Schwartz.
- [FGSS08] K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems (NIPS)*, pages 498–496, 2008.
- [FKPF12] F. J. Fabozzi, P. N. Kolm, D. Pachamanova, and S. M. Focardi. *Robust portfolio optimization and management*. John Wiley & Sons, 2012.
- [FM03] G. Fort and E. Moulines. Convergence of the Monte Carlo expectation maximization for curved exponential families. *The Annals of Statistics*, 31(4):1220–1259, 2003.
- [Fre21] K. R. French. Data Library, 2021. (https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).
- [FZ19] C. Francq and J.-M. Zakoïan. *GARCH models*. John Wiley & Sons, second edition, 2019. Structure, statistical inference and financial applications.
- [GBR⁺12] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- [GFT⁺08] A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. Smola. A kernel statistical test of independence. In *Advances in Neural Information Processing Systems (NIPS)*, pages 585–592, 2008.
- [GHZ13] J. Green, J. R. Hand, and X. F. Zhang. The supraview of return predictive signals. *Review of Accounting Studies*, 18(3):692–730, 2013.
- [GKKW02] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer, 2002.
- [GLHL20] W. Gong, Y. Li, and J. M. Hernández-Lobato. Sliced kernelized Stein discrepancy. Technical report, 2020. (<https://arxiv.org/abs/2006.16531>).
- [GLLN20] I. Guo, N. Langrené, G. Loeper, and W. Ning. Portfolio optimization with a prescribed terminal wealth distribution. Technical report, 2020. (<https://arxiv.org/abs/2009.12823>).
- [GM15] J. Gorham and L. Mackey. Measuring sample quality with stein’s method. In *Advances in Neural Information Processing Systems (NIPS)*, pages 226–234, 2015.
- [GM17] J. Gorham and L. Mackey. Measuring sample quality with kernels. In *International Conference on Machine Learning (ICML)*, pages 1292–1301, 2017.
- [GRM20] J. Gorham, A. Raj, and L. Mackey. Stochastic Stein discrepancies. Technical report, 2020. (<https://arxiv.org/abs/2007.02857>).
- [GTS18] A. Gilbert, A. Tewari, and Y. Sung. But how does it work in theory? Linear SVM with random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3379–3388, 2018.
- [Gui11] V. Guigues. Sensitivity analysis and calibration of the covariance matrix for stable portfolio selection. *Computational Optimization and Applications*, 48(3):553–579, 2011.

- [GWJ+20] W. Grathwohl, K.-C. Wang, J.-H. Jacobsen, D. Duvenaud, and R. Zemel. Learning the Stein discrepancy for training and evaluating energy-based models without sampling. In *International Conference on Machine Learning (ICML)*, 2020.
- [HBM07] Z. Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel Fisher discriminant analysis. In *Advances in Neural Information Processing Systems (NIPS)*, pages 609–616, 2007.
- [HH80] P. Hall and C. C. Heyde. *Martingale limit theory and its application*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1980.
- [HL14] C. R. Harvey and Y. Liu. Evaluating trading strategies. *The Journal of Portfolio Management*, 40(5):108–118, 2014.
- [HM18] J. H. Huggins and L. Mackey. Random feature Stein discrepancies. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1899–1909, 2018.
- [Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [HR99] T. Huillet and H.-F. Raynaud. Rare events in a log-Weibull scenario-Application to earthquake magnitude data. *The European Physical Journal B - Condensed Matter and Complex Systems*, 12(3):457–469, 1999.
- [JGH+15] W. Jitkrittum, A. Gretton, N. Heess, A. Eslami, B. Lakshminarayanan, D. Sejdinovic, and Z. Szabó. Kernel-based just-in-time learning for passing expectation propagation messages. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 405–414, 2015.
- [JM03] R. Jagannathan and T. Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1683, 2003.
- [Jor86] P. Jorion. Bayes-Stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis*, 21(3):279–292, 1986.
- [JXS+17] W. Jitkrittum, W. Xu, Z. Szabó, K. Fukumizu, and A. Gretton. A linear-time kernel goodness-of-fit test. In *Advances in Neural Information Processing Systems (NIPS)*, pages 261–270, 2017.
- [KBGP18] N. Keriven, A. Bourrier, R. Gribonval, and P. Pérez. Sketching for large-scale learning of mixture models. *Information and Inference: A Journal of the IMA*, 7:447–508, 2018.
- [Ken98] M. Kennedy. Bayesian quadrature with non-normal approximating functions. *Statistics and Computing*, 8(4):365–375, 1998.
- [KFCF14] W. C. Kim, F. J. Fabozzi, P. Cheridito, and C. Fox. Controlling portfolio skewness and kurtosis without directly optimizing third and fourth moments. *Economics Letters*, 122(2):154–158, 2014.
- [KKK16] B. Kim, R. Khanna, and O. Koyejo. Examples are not enough, learn to criticize! criticism for interpretability. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2280–2288, 2016.
- [KR61] M. Krasnoselskii and Y. Rutickii. *Convex functions and Orlicz spaces*. P. Noordhoff Ltd., Groningen, 1961.

- [KR05] T. Klein and E. Rio. Concentration around the mean for maxima of empirical processes. *The Annals of Probability*, 33(3):1060–1077, 2005.
- [KZ07] R. Kan and G. Zhou. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 42(3):621–656, 2007.
- [Las19] N. Lassance. *Information-theoretic approaches to portfolio selection*. PhD thesis, Louvain School of Management, 2019.
- [LDG⁺14] J. R. Lloyd, D. Duvenaud, R. Grosse, J. Tenenbaum, and Z. Ghahramani. Automatic construction and natural-language description of nonparametric regression models. In *AAAI Conference on Artificial Intelligence*, pages 1242–1250, 2014.
- [LJD⁺18] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.
- [LLB12] F. Lauer, V. L. Le, and G. Bloch. Learning smooth models of nonsmooth functions via convex optimization. In *International Workshop on Machine Learning for Signal Processing (IEEE-MLSP)*, 2012.
- [LLJ16] Q. Liu, J. Lee, and M. Jordan. A kernelized Stein discrepancy for goodness-of-fit tests. In *International Conference on Machine Learning (ICML)*, pages 276–284, 2016.
- [LMK10] O. P. Le Maître and O. M. Knio. *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*. Scientific Computation. Springer, 2010.
- [LMRN⁺02] O. P. Le Maître, M. T. Reagan, H. N. Najm, R. G. Ghanem, and O. M. Knio. A stochastic projection method for fluid flow. II. Random process. *Journal of Computational Physics*, 181(1):9–44, 2002.
- [Lo02] A. W. Lo. The statistics of Sharpe ratios. *Financial Analysts Journal*, 58(4):36–52, 2002.
- [LPMST15] D. Lopez-Paz, K. Muandet, B. Schölkopf, and I. Tolstikhin. Towards a learning theory of cause-effect inference. In *International Conference on Machine Learning (ICML)*, pages 1452–1461, 2015.
- [LSS13] Q. Le, T. Sarlós, and A. Smola. Fastfood - computing Hilbert space expansions in loglinear time. In *International Conference on Machine Learning (ICML)*, volume 28, pages 244–252, 2013.
- [LSZ15] Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. In *International Conference on Machine Learning (ICML)*, pages 1718–1727, 2015.
- [LT13] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer, 2013.
- [LTOS19] Z. Li, J.-F. Ton, D. Oglic, and D. Sejdinovic. A unified analysis of random Fourier features. In *International Conference on Machine Learning (ICML)*, pages 3905–3914, 2019.
- [LW03] O. Ledoit and M. Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5):602–621, 2003.

- [Man63] B. Mandelbrot. The Variation of Certain Speculative Prices. volume 36, pages 394–394. 1963.
- [Mar52] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.
- [Mar17] A. Marchina. Concentration inequalities for suprema of unbounded empirical processes. *hal-01545101*, 2017. (hal.archives-ouvertes.fr/hal-01545101).
- [Mer71] R. C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971.
- [Mic89] R. O. Michaud. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42, 1989.
- [MPJ⁺16] J. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing cause from effect using observational data: Methods and benchmarks. *Journal of Machine Learning Research*, 17:1–102, 2016.
- [MRT10a] S. Maillard, T. Roncalli, and J. Teiletche. The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management*, 36(4):60–70, 2010.
- [MRT10b] S. Maillard, T. Roncalli, and J. Teiletche. The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management*, 36(4):60–70, 2010.
- [MT09] S. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, Cambridge, second edition, 2009.
- [Mül97] A. Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29:429–443, 1997.
- [MZ10] L. Martellini and V. Ziemann. Improved estimates of higher-order comoments and implications for portfolio selection. *The Review of Financial Studies*, 23(4):1467–1502, 2010.
- [NP07] R. Nickl and B. Pötscher. Bracketing metric entropy rates and empirical central limit theorems for function classes of Besov- and Sobolev-type. *Journal of Theoretical Probability*, 20(2):177–199, 2007.
- [OCBG19] C. J. Oates, J. Cockayne, F.-X. Briol, and M. Girolami. Convergence rates for a class of estimators based on Stein's method. *Bernoulli*, 25(2):1141–1159, 2019.
- [OGC17] C. J. Oates, M. Girolami, and N. Chopin. Control functionals for Monte Carlo integration. *Journal of the Royal Statistical Society: Series B*, 79(3):695–718, 2017.
- [ONP⁺15] J. Oliva, W. Neiswanger, B. Póczos, E. Xing, H. Trac, S. Ho, and J. Schneider. Fast function to function regression. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 38, pages 717–725, 2015.
- [OP17] A. B. Owen and C. Prieur. On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002, 2017.
- [Osb59] M. F. M. Osborne. Brownian motion in the stock market. *Operations Research*, 7:145–173, 1959.
- [Owe14] A. B. Owen. Sobol' indices and Shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.

- [PBSP18] N. Pfister, P. Bühlmann, B. Schölkopf, and J. Peters. Kernel-based tests for joint independence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(1):5–31, 2018.
- [PC19] G. Peyré and M. Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [PGB11] V. Pawlowsky-Glahn and A. Buccianti. *Compositional data analysis: Theory and applications*. John Wiley & Sons, 2011.
- [Pis16] G. Pisier. *Martingales in Banach spaces*, volume 155. Cambridge University Press, 2016.
- [Pit17] Y. Pitcan. A Note on Concentration Inequalities for U-Statistics. Technical report, University of Berkeley, 2017. (<https://arxiv.org/abs/1712.06160>).
- [PN10] T. K. Pogány and S. Nadarajah. On the characteristic function of the generalized normal distribution. *Comptes Rendus Mathématique*, 348(3-4):203–206, 2010.
- [PR16] V. I. Paulsen and M. Raghupathi. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*. Cambridge University Press, 2016.
- [RCR17] A. Rudi, L. Carratino, and L. Rosasco. FALKON: An optimal large scale kernel method. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3891–3901, 2017.
- [Rio17] E. Rio. About the constants in the Fuk-Nagaev inequalities. *Electronic Communications in Probability*, 22, 2017.
- [RK04] R. Y. Rubinstein and D. P. Kroese. *The cross-entropy method: A unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer, 2004.
- [RM51] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- [RR07] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1177–1184, 2007.
- [RR08] A. Rahimi and B. Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1313–1320, 2008.
- [RR17] A. Rudi and L. Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3218–3228, 2017.
- [RS71] H. Robbins and D. Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics (Proc. Sympos., Ohio State Univ., Columbus, Ohio, 1971)*, pages 233–257, 1971.
- [RSM⁺10] L. Rosasco, M. Santoro, S. Mosci, A. Verri, and S. Villa. A regularization approach to nonlinear variable selection. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 653–660, 2010.
- [Rud90] W. Rudin. *Fourier Analysis on Groups*. John Wiley & Sons, 1990.

- [RVM⁺13] L. Rosasco, S. Villa, S. Mosci, M. Santoro, and A. Verri. Nonparametric sparsity and regularization. *Journal of Machine Learning Research*, 14:1665–1714, 2013.
- [RW05] J. P. Romano and M. Wolf. Stepwise multiple testing as formalized data snooping. *Econometrica. Journal of the Econometric Society*, 73(4):1237–1282, 2005.
- [SC08] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, 2008.
- [SFG⁺17] B. Sriperumbudur, K. Fukumizu, A. Gretton, A. Hyvärinen, and R. Kumar. Density estimation in infinite dimensional exponential families. *Journal of Machine Learning Research*, 18(57):1–59, 2017.
- [SFL11] B. Sriperumbudur, K. Fukumizu, and G. Lanckriet. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12:2389–2410, 2011.
- [SGF⁺10] B. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- [SGSS07] A. Smola, A. Gretton, L. Song, and B. Schölkopf. A Hilbert space embedding for distributions. In *Algorithmic Learning Theory (ALT)*, pages 13–31, 2007.
- [SGZ10] L. Shi, X. Guo, and D.-X. Zhou. Hermite learning with gradient data. *Journal of Computational and Applied Mathematics*, 233:3046–3059, 2010.
- [SH80] R. C. Scott and P. A. Horvath. On the direction of preference for moments of higher order than the variance. *The Journal of finance*, 35(4):915–919, 1980.
- [Sha66] W. F. Sharpe. Mutual fund performance. *The Journal of Business*, 39(1):119–138, 1966.
- [SKN⁺20] L. F. South, T. Karvonen, C. Nemeth, M. Girolami, and C. J. Oates. Semi-exact control functionals from Sard’s method. Technical report, 2020. (<https://arxiv.org/abs/2002.00033>).
- [Sob01] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1):271–280, 2001.
- [SRA⁺08] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.
- [SS15a] B. K. Sriperumbudur and Z. Szabó. Optimal rates for random Fourier features. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1144–1152, 2015.
- [SS15b] D. Sutherland and J. Schneider. On the error of random Fourier features. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 862–871, 2015.
- [SS18] B. Sriperumbudur and N. Sterge. Approximate kernel PCA using random features: Computational vs. statistical trade-off. Technical report, Pennsylvania State University, 2018. (<https://arxiv.org/abs/1706.06296>).
- [SS19] Z. Szabó and B. K. Sriperumbudur. On kernel derivative approximation with random Fourier features. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 827–836, 2019.

- [SSL⁺15] H. Strathmann, D. Sejdinovic, S. Livingstone, Z. Szabó, and A. Gretton. Gradient-free Hamiltonian Monte Carlo with efficient kernel exponential families. In *Advances in Neural Information Processing Systems (NIPS)*, pages 955–963, 2015.
- [Sza14] Z. Szabó. Information theoretical estimators toolbox. *Journal of Machine Learning Research*, 15(9):283–287, 2014.
- [SZS⁺08] L. Song, X. Zhang, A. Smola, A. Gretton, and B. Schölkopf. Tailoring density estimation via reproducing kernel moment matching. In *International Conference on Machine Learning (ICML)*, page 992–999, 2008.
- [SZV19] E. V. Strobl, K. Zhang, and S. Visweswaran. Approximate kernel-based conditional independence tests for fast non-parametric causal discovery. *Journal of Causal Inference*, 7(1):20180017, 2019. (eISSN 2193-3685; <https://doi.org/10.1515/jci-2018-0017>).
- [Tal89] M. Talagrand. Isoperimetry and integrability of the sum of independent Banach-space valued random variables. *The Annals of Probability*, 17(4):1546–1570, 1989.
- [Tas14] D. Tasche. Expected Shortfall is not elicitable. So what? (https://www.insurance.uni-hannover.de/fileadmin/house-of-insurance/News_and_Events/Events/2014/Colloquium/2014.01.23/Talk_Tasche.pdf), 2014.
- [TC04] J. S. Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [Tro12] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [TSS16] I. Tolstikhin, B. Sriperumbudur, and B. Schölkopf. Minimax estimation of maximal mean discrepancy with radial kernels. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1930–1938, 2016.
- [Tsy09] A. B. Tsybakov. *Introduction to nonparametric estimation*. Springer, 2009.
- [UMMA18] E. Ullah, P. Mianjy, T. Marinov, and R. Arora. Streaming kernel PCA with $\tilde{O}(\sqrt{n})$ random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7311–7321, 2018.
- [van16] R. van Handel. Probability in high dimension., 2016. (<https://web.math.princeton.edu/~rvan/APC550.pdf>).
- [Vap00] V. N. Vapnik. *The nature of statistical learning theory*. Springer, second edition, 2000.
- [vdG00] S. van de Geer. *Empirical Processes in M-estimation*. Cambridge University Press, 2000.
- [vdGL13] S. van de Geer and J. Lederer. The Bernstein-Orlicz norm and deviation inequalities. *Probability Theory and Related Fields*, 157(1-2):225–250, 2013.
- [Vil08] C. Villani. *Optimal Transport: Old and New*. Springer, 2008.
- [vW96] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical processes: with Applications to Statistics*. Springer, 1996.
- [Wel17] J. Wellner. The Bennett-Orlicz norm. *Sankhya A. The Indian Journal of Statistics*, 79(2):355–383, 2017.

- [Whi00] H. White. A reality check for data snooping. *Econometrica. Journal of the Econometric Society*, 68(5):1097–1126, 2000.
- [WS01] C. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, pages 682–688, 2001.
- [WX12] H. Wang and S. Xiang. On the convergence rates of Legendre approximation. *Mathematics of Computation*, 81(278):861–877, 2012.
- [YPW17] Y. Yang, M. Pilanci, and M. Wainwright. Randomized sketches for kernels: Fast and optimal non-parametric regression. *The Annals of Statistics*, 45:991–1023, 2017.
- [YSAM14] J. Yang, V. Sindhwani, H. Avron, and M. Mahoney. Quasi-Monte Carlo feature maps for shift-invariant kernels. *International Conference on Machine Learning (ICML)*, 32:485–493, 2014.
- [YSC⁺16] F. X. Yu, A. T. Suresh, K. Choromanski, D. Holtmann-Rice, and S. Kumar. Orthogonal random features. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1975–1983, 2016.
- [YUFT18] M. Yamada, Y. Umezū, K. Fukumizu, and I. Takeuchi. Post selection inference with kernels. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 152–160, 2018.
- [YWC12] Y. Ying, Q. Wu, and C. Campbell. Learning the coordinate gradients. *Advances in Computational Mathematics*, 37:355–378, 2012.
- [ZFGS17] Q. Zhang, S. Filippi, A. Gretton, and D. Sejdinovic. Large-scale kernel methods for independence testing. *Statistics and Computing*, pages 1–18, 2017.
- [Zho08] D.-X. Zhou. Derivative reproducing properties for kernel methods in learning theory. *Journal of Computational and Applied Mathematics*, 220:456–463, 2008.
- [ZMDR19] J. Zhang, A. May, T. Dao, and C. Ré. Low-precision random Fourier features for memory-constrained kernel approximation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 89, pages 1264–1274, 2019.
- [Zol83] V. Zolotarev. Probability metrics. *Theory of Probability and its Applications*, 28:278–302, 1983.

Titre : Quantification des incertitudes en gestion d'actifs: méthodes à noyaux et fluctuations statistiques

Mots clés : Queues épaisses ; Inégalité de concentration ; Méthodes à noyau ; Estimation de la matrice de covariance ; Théorie du portefeuille ; Quantification d'incertitudes.

Résumé : Le traitement des incertitudes est un problème fondamental dans le contexte financier. Les variables étudiées sont souvent dépendantes du temps, avec des queues de distribution épaisses. Dans cette thèse, on s'intéresse à des outils permettant de prendre en compte les incertitudes sous ses formes principales: incertitudes statistiques, paramétriques et erreur de modèle, tout en gardant en tête qu'on souhaite les appliquer à ce contexte.

La première partie est consacrée à l'établissement d'inégalités de concentration dans le cadre de variables à queues épaisses. L'objectif de ces inégalités est de quantifier quelle confiance on peut donner à un estimateur basé sur une taille finie d'observations. Dans cette thèse, nous établissons de nouvelles inégalités de concentration, qui couvrent notamment le cas d'estimateur à distribution log-normale.

Dans la seconde partie, on traite de l'impact de l'erreur de modèle pour l'estimation de la matrice de covariance sur des rendements boursiers, sous hypothèse qu'il existe un processus de covariance instantanée entre les rendements dont la valeur présente

dépend de sa valeur passée. On peut alors construire explicitement la meilleure estimée de la matrice de covariance pour un instant et un horizon d'investissement donnés, et montrer qu'elle fournit la variance réalisée la plus faible avec grande probabilité dans le cadre du portefeuille minimum variance.

Dans la troisième partie, on propose une approche pour estimer le ratio de Sharpe et l'allocation de portefeuille lorsqu'ils dépendent de paramètres jugés incertains. Notre approche passe par l'adaptation d'une technique d'approximation stochastique pour le calcul de la décomposition en polynômes du chaos de la quantité d'intérêt.

Enfin, dans la dernière partie de cette thèse, on s'intéresse à l'optimisation de portefeuille avec distribution cible. Cette technique peut être formalisée sans avoir recours à aucune hypothèse de modèle sur les rendements. Nous proposons de trouver ces portefeuilles en minimisant des mesures de divergence basées sur les fonctions noyau et la théorie du transport optimal.

Title : Quantifying uncertainty in asset management: Kernel methods and statistical fluctuations

Keywords : Heavy tails; Concentration inequality; Kernel methods; Covariance matrix estimation; Portfolio theory; Uncertainty Quantification.

Abstract : The treatment of uncertainties is a fundamental problem in the financial context, and more precisely in portfolio optimisation. The variables studied are often time dependent, with heavy tails. In this thesis, we are interested in tools allowing to take into account uncertainties in its main forms: statistical uncertainties, parametric uncertainties and model error, keeping in mind that we wish to apply them to the financial context.

The first part is devoted to the establishment of concentration inequalities for variables with heavy tailed distributions. The objective of these inequalities is to quantify the confidence that can be given to an estimator based on observations of finite size. In this thesis, we establish new concentration inequalities which include the case of estimators with log-normal distribution.

In the second part, we discuss the impact of the model error for the estimation of the covariance matrix on stock returns, under the assumption that there is an instantaneous covariance process between the returns whose present value depends on its past values.

One can then explicitly construct the best estimate of the covariance matrix for a given time and investment horizon, and we show that this estimate gives the best performance with high probability in the minimum variance portfolio framework.

In the third part, we propose an approach to estimate the Sharpe ratio and the portfolio allocation when they depend on parameters considered uncertain. Our approach involves the adaptation of a stochastic approximation technique for the computation of the polynomial decomposition of the quantity of interest.

Finally, in the last part of this thesis, we focus on portfolio optimization with target distribution. This technique can be formalised without the need for any model assumptions on returns. We propose to find these portfolios by minimizing divergence measures based on kernels or optimal transport. Since these divergence measures can be unbounded and have not been studied much yet in the unbounded kernel case, we establish new convergence guarantees based on concentration inequalities.