Robust Optimal Growth from an analytic and learning perspective

Josef Teichmann
(joint work with David Itkin, Benedikt Koch, and Martin Larsson,
and with Florian Krach and Hanna Wutte)

ETH Zürich

BFS Online Seminar September 2023

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Optimal Portfolio Selection

- Keynes' approach: analyze fundamentals.
- Markowitz approach: *observe* the market and solve a mean-variance problem.
- Utility Optimization: *observe* the market and solve an expected utility optimization problem.
- Robust Utility Optimization: *observe* the market and *include* uncertainty into the expected utility optimization problem in a zero sum game way.
- Data driven approaches using learning techniques as, e.g., signature transforms.
- Stochastic Portfolio Theory (SPT): only rely on observables (to be defined) and build robustly growing portfolios due to normative assumptions by means of a generating function.

Critique

- What are pathwise observables? Prices, quadratic covariation, portfolios, etc, but ...
- ... neither jump structures nor drifts can be considered observable with high accuracy.
- A quantitative knowledge of preferences is ad hoc and cannot easily be based on observables.
- On the other hand: portfolios soley build upon observables and soley relying on functional generation as in SPT might be too restriced a class?

Critique

- What are pathwise observables? Prices, quadratic covariation, portfolios, etc, but ...
- ... neither jump structures nor drifts can be considered observable with high accuracy.
- A quantitative knowledge of preferences is ad hoc and cannot easily be based on observables.
- On the other hand: portfolios soley build upon observables and soley relying on functional generation as in SPT might be too restriced a class?

Critique

- What are pathwise observables? Prices, quadratic covariation, portfolios, etc, but ...
- ... neither jump structures nor drifts can be considered observable with high accuracy.
- A quantitative knowledge of preferences is ad hoc and cannot easily be based on observables.
- On the other hand: portfolios soley build upon observables and soley relying on functional generation as in SPT might be too restriced a class?

Critique

- What are pathwise observables? Prices, quadratic covariation, portfolios, etc, but ...
- ... neither jump structures nor drifts can be considered observable with high accuracy.
- A quantitative knowledge of preferences is ad hoc and cannot easily be based on observables.
- On the other hand: portfolios soley build upon observables and soley relying on functional generation as in SPT might be too restriced a class?

In an ideal world

- Robust Utility Optimization takes all data and uncertainty into account appears as cleanest approach.
- Robust Utility Optimization is a zero sum game (in case of fixed utility): the investor against an adverserial market. Analytical techniques are, however, challenging.
- Machine learning techniques in the spirit if generative adversarial networks can be applied and are helpful.
- Problems: high complexity of training of path dependent trading and adversarial strategies. Additionally we face difficulties to determine reasonable adversarial markets. A priori there is no reason why path dependence should disappear.

In an ideal world

- Robust Utility Optimization takes all data and uncertainty into account appears as cleanest approach.
- Robust Utility Optimization is a zero sum game (in case of fixed utility): the investor against an adverserial market. Analytical techniques are, however, challenging.
- Machine learning techniques in the spirit if generative adversarial networks can be applied and are helpful.
- Problems: high complexity of training of path dependent trading and adversarial strategies. Additionally we face difficulties to determine reasonable adversarial markets. A priori there is no reason why path dependence should disappear.

In an ideal world

- Robust Utility Optimization takes all data and uncertainty into account appears as cleanest approach.
- Robust Utility Optimization is a zero sum game (in case of fixed utility): the investor against an adverserial market. Analytical techniques are, however, challenging.
- Machine learning techniques in the spirit if generative adversarial networks can be applied and are helpful.
- Problems: high complexity of training of path dependent trading and adversarial strategies. Additionally we face difficulties to determine reasonable adversarial markets. A priori there is no reason why path dependence should disappear.

In an ideal world

- Robust Utility Optimization takes all data and uncertainty into account appears as cleanest approach.
- Robust Utility Optimization is a zero sum game (in case of fixed utility): the investor against an adverserial market. Analytical techniques are, however, challenging.
- Machine learning techniques in the spirit if generative adversarial networks can be applied and are helpful.
- Problems: high complexity of training of path dependent trading and adversarial strategies. Additionally we face difficulties to determine reasonable adversarial markets. A priori there is no reason why path dependence should disappear.

Message from SPT

- If SPT makes sense, then path dependence should at least in case logarithmic utilities vanish, i.e. trading strategies can be chosen from soley price dependent classes (in particular no dependence on wealth processes), however complicated markets are.
- When does that actually happen and how complicated can we choose the market?
- If path dependence vanishes the learning problem sees a spectacular decrease of complexity (a game is replace by an optimization problem and path dependence is replaced by price dependence).

Message from SPT

- If SPT makes sense, then path dependence should at least in case logarithmic utilities vanish, i.e. trading strategies can be chosen from soley price dependent classes (in particular no dependence on wealth processes), however complicated markets are.
- When does that actually happen and how complicated can we choose the market?
- If path dependence vanishes the learning problem sees a spectacular decrease of complexity (a game is replace by an optimization problem and path dependence is replaced by price dependence).

Message from SPT

- If SPT makes sense, then path dependence should at least in case logarithmic utilities vanish, i.e. trading strategies can be chosen from soley price dependent classes (in particular no dependence on wealth processes), however complicated markets are.
- When does that actually happen and how complicated can we choose the market?
- If path dependence vanishes the learning problem sees a spectacular decrease of complexity (a game is replace by an optimization problem and path dependence is replaced by price dependence).

Goal of the talk

- Showcase when path dependence vanishes following seminal work of Kostas Kardaras and Scott Robertson.
- Reduce the zero sum game to an optimization problem.
- Introduce a finite time horizon zero sum game of robust growth optimization type with a similar phenomenon.
- Highlight on learning techniques which show the phenomenon beyond logarithmic utility.

Goal of the talk

- Showcase when path dependence vanishes following seminal work of Kostas Kardaras and Scott Robertson.
- Reduce the zero sum game to an optimization problem.
- Introduce a finite time horizon zero sum game of robust growth optimization type with a similar phenomenon.
- Highlight on learning techniques which show the phenomenon beyond logarithmic utility.

Goal of the talk

- Showcase when path dependence vanishes following seminal work of Kostas Kardaras and Scott Robertson.
- Reduce the zero sum game to an optimization problem.
- Introduce a finite time horizon zero sum game of robust growth optimization type with a similar phenomenon.
- Highlight on learning techniques which show the phenomenon beyond logarithmic utility.

Goal of the talk

- Showcase when path dependence vanishes following seminal work of Kostas Kardaras and Scott Robertson.
- Reduce the zero sum game to an optimization problem.
- Introduce a finite time horizon zero sum game of robust growth optimization type with a similar phenomenon.
- Highlight on learning techniques which show the phenomenon beyond logarithmic utility.

Local covariance Setting

We look at a market, appropriately discounted, where sufficiently many observed trading instruments modelled by continuous semi-martingales X on path space (and therefore with some extrapolation the instantaneous covariance function c) as well as a strictly positive invariant measure $p(x)dx$ are available. Let Π be *the* set of measures such that ...

- X is a semimartingale and $[X, X] = \int_0^\cdot c(X_s)ds$, for some invertible symmetric matrix c for $P \in \Pi$,
- $\frac{1}{T}\int_0^T f(X_t)dt \longrightarrow \int_{\mathbb{R}^d} f(x)p(x)dx$ almost surely for all $P \in \Pi$.
- For every $\epsilon > 0$ there is a compact K such that $P(X_t \notin K) < \epsilon$ for all $t \geq 0$ holds for *all* $P \in \Pi$.

Notice that providing an invariant measure is less than providing a drift.

Themes

- Why are functionally generated portfolios a robust (data driven) choice for portfolio selection?
- Understand the geometry of all drifts (given the covariance structure), which lead to the same invariant measure.

Local Covariance Setting

The growth rate of an investment θ in the market (value processes are assumed be positive here) is defined

$$g(\theta, P) := \sup \left\{ \gamma \in \mathbb{R} \mid \lim_{T\to\infty} P\big(\log \mathcal{E}((\theta \bullet X))_T \geq \gamma T \big) = 1 \right\}$$

What can we say about the *robust growth rate* of an optimal investment into this market, i.e.

$$\sup_{\theta} \inf_{P\in\Pi} g(\theta, P)$$

and how could we possibly calculate it and how does a robust growth portfolio look like?

Heuristics

We can easily establish a class of strategies $\Theta_0 \subset \Theta$ that have the following *growth rate invariance* property:

$$g(\theta; \mathbb{P}) \text{ is independent of } \mathbb{P} \in \Pi \text{ for every } \theta \in \Theta_0.$$

We define

$$\mathcal{D} = \left\{ \varphi \in C^2(E) : \int_E \left| \frac{\mathrm{Tr}\big(c(x)\, \nabla^2 e^{\varphi(x)}\big)}{e^{\varphi(x)}} \right| p(x)\, dx \, < \infty \right\}.$$

Heuristics

The class

$$\Theta_0 = \{\theta^{\varphi}_{\cdot} = \nabla\varphi(X_{\cdot}) : \varphi \in \mathcal{D}\}$$

of *functionally generated portfolios* has the growth rate invariance property. Indeed, denoting by $\theta^{\varphi} \in \Theta_0$ the strategy determined by φ we see by Itô's formula that under any measure $\mathbb{P} \in \Pi$,

$$\log V_T^{\theta^{\varphi}} = \varphi(X_T) - \varphi(X_0) - \frac{1}{2}\int_0^T \frac{\mathrm{Tr}(c(X_t)\nabla^2 e^{\varphi(X_t)})}{e^{\varphi(X_t)}}\, dt,$$

for instantaneous covariance structure c.

Heuristics

By tightness of the laws of $\{X_T\}_{T\geq 0}$ we have that $\varphi(X_T)/T \to 0$ in probability as $T \to \infty$.

Hence, by the ergodic property it follows that

$$g(\theta;\mathbb{P}) = -\frac{1}{2}\int_F \frac{\mathrm{Tr}(c(x)\nabla^2 e^{\varphi(x)})}{e^{\varphi(x)}}p(x)\,dx\,.$$

It is clear that we should optimize over the right hand side. The delicate question which remains is the following: can we find at least one $P \in \Pi$ such that the optimal θ^{φ^*} is growth-optimal overall. This is actually a question of non-explosion of a formally given model.

Results by Karadaras, Robertson AAP 2022

Consider the optimization problem

$$u \mapsto \int_{\mathbb{R}^d} \frac{L^c u}{u}(x) p(x) dx$$

with $L^c(u)(x) = \mathrm{Tr}(c(x)\,\text{Hessian}\,u(x)$, whose *strictly positive* solution u^* (existence!) defines a strategy $\theta_t := \frac{\nabla u^*(X_t)}{u^*(X_t)}$, which realizes the robust growth rate and allows also to define a worst case model, where this very strategy is optimal.

In particular the robust growth-optimal strategy is functionally generated.

Question: can the setting also be made more realistic, i.e. does this also work with a factor dependence of the covariance structure? (this is of considerably higher complexity since we are loosing a sort of Markovianity)

Setting

We fix integers $d, m \geq 1$ and a non-empty open, connected sets $E \subset \mathbb{R}^d$, $D \subset \mathbb{R}^m$. Set $F = E \times D$. We will generically denote elements of E by x, elements of D by y and set $z = (x, y)$ for elements of F. For a function u we write $\nabla_x u$ for $(\partial_{x_1} u, \ldots, \partial_{x_d} u)$, $\mathrm{div}_x\, u$ for $\sum_{i=1}^d \partial_{x_i} u$, and use $\nabla_y u$ and $\mathrm{div}_y\, u$ analogously.

The set E will serve as the state space for the d-dimensional asset process X, while the set D serves as the state space for a m-dimensional factor process Y. Consequently the joint $(d+m)$-dimensional process $Z = (X, Y)$ has state space F.

We take as input a matrix valued function $c_X : F \to \mathbb{S}^d_{++}$ and a positive function $p : F \to (0, \infty)$, which will serve as the instantaneous covariance matrix for X (which also depends on Y), and the joint invariant density of X and Y. We impose the following regularity assumptions on the inputs.

Regularity assumptions

For a fixed $\gamma \in (0,1]$,

1. D is a bounded convex set,
2. $c_X \in W^{1,\infty}_{\mathrm{loc}}(F;\mathbb{S}^d_{++}) \cap C(F;\mathbb{S}^d_{++})$ and for every $y \in D$, $c_X(\cdot,y) \in C^{2,\gamma}(E;\mathbb{S}^d_{++})$.
3. $p \in W^{2,\infty}_{\mathrm{loc}}(F;(0,\infty)) \cap C(F;(0,\infty))$ and for every $y \in D$, $p(\cdot,y) \in C^{2,\gamma}(E;(0,\infty))$. Additionally it is assumed that $\int_F p = 1$.

Regularity Assumptions

Given inputs c_X and p we denote the averaged instantaneous covariance with

$$A^{ij}(x) = \int_D c_X^{ij}(x,y)p(x,y)\,dy; \qquad i,j = 1,\ldots,d, \quad x \in E.$$

We will need to allow ourselves to modify the input c_X outside of a compact set $K \subset E$, but without changing the average value A. Let c_X and p satisfying the regularity assumptions. Define

$$\mathcal{C} = \left\{ \tilde{c}_X \in W^{1,\infty}_{\mathrm{loc}}(F;\mathbb{S}^d_{++}) \cap C(F;\mathbb{S}^d_{++}) : \int_D \tilde{c}_X(.,y)p(.,y)\,dy = A(.) \right\}.$$

For a compact set $K \subset E$ we define the class

$$\mathcal{C}_K = \{\tilde{c}_X \in \mathcal{C} : \tilde{c}_X = c_X \text{ on } K\}.$$

We say an element $\tilde{c}_X$ of $\mathcal{C}_K$ is a K-*modification* of c_X.

Regularity Assumptions

For a compact set $K \subset E$ we denote by Π_K the set of all measures P on $(\Omega, \mathcal{F})$ such that the following conditions hold:

1. X is a P-semimartingale with covariation process $[X, X] = \int_0^\cdot \tilde{c}_X(Z_t)\, dt$ for some $\tilde{c}_X \in \mathcal{C}_K$,
2. For any locally bounded $h \in L^1(F, \mu)$, where $\mu(dz) := p(z)dz$,
$$\lim_{T\to\infty} \frac{1}{T} \int_0^T h(Z_t)\, dt = \int_F hp; \quad \mathbb{P}\text{-a.s.},$$
3. The laws of $\{X_t\}_{t\geq 0}$ under P are tight.

We also define the classes $\Pi_\cup = \bigcup_{K\subset E} \Pi_K$ and $\Pi_\cap = \bigcap_{K\subset E} \Pi_K$, where the union and intersection are taken over compact sets K.

Note that if $K_1 \subset K_2$ then $\mathcal{C}_{K_2} \subset \mathcal{C}_{K_1}$ so that $\Pi_{K_2} \subset \Pi_{K_1}$. Additionally any measure $P \in \Pi_\cap$ has $[X, X] = \int_0^\cdot c_X(Z_t)\, dt$, rather than a K-modification appearing as the instantaneous covariance matrix.

Important Remark

To consider Y is a stochastic factor driving stochastic covariance is natural, but not the only interpretation. Actually to interpret Y is a factor representing uncertainty in the knowledge of c_X is equally important. It can also be a mixture of both.

In the latter case (uncertainty) a prescribed diffusion dynamics of Y is of less importance, but the invariant measure matters. In the former case (stochastic covariance) a prescribed generic diffusion dynamics is of importance.

Trading

Let Θ be the set of all predictable processes that are X-integrable with respect to every $P \in \Pi_\cup$. For any $\theta \in \Theta$ we define the investor's wealth process by

$$V^\theta = \mathcal{E}\left(\int_0^\cdot \theta_t^\top dX_t\right),$$

where $\mathcal{E}$ denotes the Doléans-Dade exponential of a semimartingale. The goal is to maximize the investor's asymptotic growth rate over our admissible class of measures, which is defined as follows:

For a strategy $\theta \in \Theta$ and a measure $P \in \Pi_\cup$ we define the *asymptotic growth rate* of θ to be

$$g(\theta; P) = \sup\left\{\gamma : \lim_{T\to\infty} P(T^{-1}\log V_T^\theta \geq \gamma) = 1\right\}.$$

Furthermore the robust optimal asymptotic growth rate λ_K is defined as $\sup_\theta \inf_{P\in\Pi_K} g(\theta, P)$.

Stochastic Factor Case (Itkin, Koch, Larsson, JT)

One can consider *two* auxiliary maximization problem, whose solutions define a factor independent strategy $\theta_t = \frac{\nabla u^*(X_t)}{u^*(X_t)}$ and an overall worst case model, where the constructed strategy is growth optimal (overall).

The resulting strategy is *again (and surprisingly)* functionally generated and does not have a factor dependence. However, the worst case model is considerably more involved and needs only recently developed analytical tools (non-symmetric Dirichlet forms).

The following theorems make the results precise – be aware that additional assumptions are needed which are a bit technical.

Theorem 1

There exists a unique (up to additive constant) $\hat{\varphi}$ satisfying

$$\hat{\varphi} = \arg\min_{\varphi \in \mathcal{D}} \frac{1}{2} \int_E \frac{\mathrm{Tr}(A(x)\nabla^2 e^{\varphi(x)})}{e^{\varphi(x)}}\, dx.$$

Define

$$\lambda = \frac{1}{2} \int_E \nabla\hat{\varphi}^\top(x) A(x) \nabla\hat{\varphi}(x) dx$$

and the strategy

$$\hat{\theta}_t := \nabla\hat{\varphi}(X_t); \quad t \geq 0.$$

Then for every compact set $K \subset E$ we have that $\lambda_K = \lambda$. Moreover, $g(\hat{\theta}; \mathbb{P}) = \lambda$ for every $\mathbb{P} \in \Pi_K$, so that $\hat{\theta}$ is robust growth-optimal.

Theorem 2

For every compact set $K \subset E$ and every $(x, y) \in F$ there exists a measure $\widehat{P}^K_{(x,y)}$ which is a weak solution to the stochastic differential equation

$$dX_t = \tilde{c}_X(X_t, Y_t)\nabla\hat{\varphi}(X_t)\,dt + \tilde{c}_X^{1/2}(X_t, Y_t)\,dW_t^X$$
$$dY_t = c_Y(X_t, Y_t)\nabla_y\hat{v}(X_t, Y_t)\,dt + c_Y^{1/2}(X_t, Y_t)\,dW_t^Y$$

and satisfies $\widehat{P}^K_{(x,y)}(X_0 = x, Y_0 = y) = 1$. Here $W := (W^X, W^Y)$ is a standard $(d+m)$-dimensional Brownian motion, c_Y satisfies unstated assumptions, $\hat{\varphi}$, $\tilde{c}_X$ and $\hat{v}$ are given via certain optimization problems.

We additionally have that μ is an invariant measure for (X, Y) and for every locally bounded $h \in L^1(F, \mu)$

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T h(X_t, Y_t)\,dt = \int_F hp; \quad \widehat{P}^K_{(x,y)}\text{-a.s.}$$

Consequently, the laws of $\{X_t\}_{t\geq 0}$ are tight under $\widehat{P}^K_{(x,y)}$ and so we have that $\widehat{P}^K_{(x,y)} \in \Pi_K$ for every $(x, y) \in F$.

Conclusion

- Functionally generated portfolios appear as asymptotic growth optimal choice even in the (robust) presence of exogenous stochastic factors. Path dependence appears to be asymptotically of less interest.
- Even though the heuristic argument is simple, the proof is surprisingly involved.
- It is open to find a simpler proof and to ask which additional information on Y would change the picture. So far a joint invariant law is given for (X, Y), the instantaneous covariance for X is given, and one can choose the instantaneous covariance for Y in a relatively rich class of functions. $[X, Y]$ is assumed to vanish for the worst case model but this can also be generalized.

Conclusion and Critique

- Surprisingly the robust growth optimal is functionally generated and does not depend on the factor process.
- Its calculation soley depends on p and c.
- Is p really an observable?
- What happens in a setting of $T < \infty$?

$T < \infty$

1. X is a P-semimartingale with covariation process $[X, X] = \int_0^{\cdot} \tilde{c}_X(Z_t)\, dt$ for some $\tilde{c}_X \in \mathcal{C}_K$,
2. The laws of Z_t equal $\mu(dz) := p(z)dz$ for $t \in [0, T]$ (invariant measure case).

Then completely analogous conclusions to Theorem 1 and 2 hold true in case of logarithmic utility and averaging over initial prices, in particular the robust growth portfolio is functionally generated and there is again a worst case model.

Robust Utility Optimization in a stochastic factor setting

We consider now a finite time horizon (general) utility optimization problem

$$\sup_{\theta} \inf_{P \in \Pi} E\left[u(V_T^{x,\theta})\right],$$

where the uncertainty class is pinned down by constraining $\frac{1}{T}\int_0^T h(X_s, Y_s)ds$ for certain h, which is less information than knowing p. The actor chooses θ and the critic chooses P.

Before Learning we learn a bit from ILKT2022

- We still know that asymptotic growth rates are realized by functionally generated portfolios in X.
- It seems that at least in ergodic setting with logarithmic utility path dependence evaporates.
- Idea: model strategies of the generator by (random) recurrent networks (following work of Ivan Guo et al (2022)).
- Loss function: utility plus penalizations for deviations of observations of time averages.
- Discriminator: fully generic discriminator dynamics satisfying appropriate constraints.

Results

(joint work with Florian Krach and Hanna Wutte)

- GAN-based (robust) utility optimization works well,
- applicable in very general settings,
- strategy recovers analytic solution,
- approximates solution well in cases without analytic solution,
- FFNN and RNN nearly same (i.e. no strong path dependence).

References

- Guo, I., Langrene, N., Loeper, G., and Ning, W.: *Robust utility maximization under model uncertainty via a penalization approach*, Mathematics and Financial Economics (2022).
- D. Itkin, B. Koch, M. Larsson, J. Teichmann: *Ergodic robust maximization of asymptotic growth under stochastic volatility*, Arxiv, 2022.
- K. Kardaras, S. Robertson: *Ergodic robust maximization of asymptotic growth*, AAP 2022.
- H. Lee, W. Stannat, and G. Trutnau: *Analytic Theory of Itô–Stochastic Differ- ential Equations with Non-smooth Coefficients*. SpringerBriefs in Probability and Mathematical Statistics, 2022.