

Generalization Bounds for Regularized Portfolio Selection with Market Side Information

Thierry Bazier-Matte^a and Erick Delage^b

^aCaisse de dépôt et placement du Québec, Canada.

^bHEC Montreal, Canada.

ARTICLE HISTORY

Compiled February 20, 2018

ABSTRACT

Drawing on statistical learning theory, we derive out-of-sample and suboptimal guarantees about the investment strategy obtained from a regularized portfolio optimization model which attempts to exploit side information about the financial market in order to reach an optimal risk-return tradeoff. This side information might include for instance recent stock returns, volatility indexes, financial news indicators, etc. In particular, we demonstrate that an investment policy that linearly combines this side information in a way that is optimal from the perspective of a random sample set is guaranteed to perform also relatively well (*i.e.*, within a perturbing factor of $O(1/\sqrt{n})$) with respect to the unknown distribution that generated this sample set. Finally, we evaluate the sensitivity of these results in a high dimensional regime where the size of the side information vector is of an order that is comparable to the sample size.

KEYWORDS

Portfolio optimization; Generalization bound; Utility maximization; learning theory

1. Introduction

There is no doubt that modern portfolio management theory has been dramatically affected by two important historical events. First, Markowitz in 1952 highlighted in his seminal paper Markowitz (1952) how investment decisions needed to inherently trade-off between risk (typically measured using variance) and returns (in the form of expected returns). This was later reinterpreted as a special case of characterizing risk aversion using expected utility theory von Neumann and Morgenstern (1944). The flexibility of such a theory has since then been demonstrated in many occasions regarding the wide diversity of investors' risk aversion that it can represent (see Ingersoll (1987) and reference therein for an overview of the type of attitudes that can be modeled).

The second turning point of this theory can be considered to have occurred with the financial crisis of 2008 which provided strong evidence that the use of statistics such as variance and value-at-risk, and of distribution models that are calibrated using

CONTACT T. Bazier-Matte at tbaziermatte@cdpq.ca and E. Delage at erick.delage@hec.ca. E. Delage gratefully acknowledges the support of the Canadian Natural Sciences and Engineering Research Council [Grant 386416-2010].

historical data could provide a false sense of security Salmon (2009). In an attempt to address some of these new challenges, researchers have proposed using more robust statistical estimators Madan et al. (1998); Goldfarb and Iyengar (2003); Olivares-Nadal and DeMiguel (2018) while others encouraged the use of robust portfolio management models that are designed to produce out-of-samples guarantees by exploiting the use of a confidence region for the distribution of future returns Delage and Ye (2010); Huang et al. (2010); Esfahani and Kuhn (2017); Bertsimas and Van Parys (2017).

In this work, we draw on statistical learning theory to establish what are the out-of-sample guarantees that can be obtained when using regularization in an expected utility model that allows to exploit side information about the financial markets (see Brandt et al. (2009) where non-regularized version was introduced). This side information could consist of fundamental analysis (as was famously done in Fama and French (1993)), but also of technical analysis, of financial news, etc. Overall, we consider our contribution to be three-fold.

- (1) We derive a lower bound on the out-of-sample performance of the investment strategy returned by this regularized model. In this respect, our results differ from the usual statistical learning and stability theory results in the sense that our guarantees will not be in terms of quality of fit of a model (*e.g.*, expected squared loss, hinge loss, etc.), but rather in terms of the actual performance perceived by the investor (through the notion of a certainty equivalent).
- (2) We derive an upper bound on the suboptimality of the investment strategy when compared to the optimal strategy that would be derived using the full knowledge of the sample distribution. Note that such guarantees have not been established for data-driven or distributionally robust optimization.
- (3) Considering that nowadays a growing amount of side-information can be exploited by individuals to make their investments, we establish precisely how these bounds are affected at a high-dimensional (or “big data”) regime.

It is worth mentioning that the above contributions are similar in spirit to those of Rudin and Vahn (2015) who applied stability theory to provide generalization bounds for a newsvendor problem. There are however a number of distinctions regarding how stability theory needs to be articulated for the two applications. For example, our paper deals with a more general performance function which is non-linear and possibly unbounded on both sides, and we focus on generalization and sub-optimality bounds which, in our opinion, are more informative than the measure that is bounded in Rudin and Vahn (2015). Moreover, we attempt to be more precise in our analysis on characterizing the effect of data dimensionality on the out-of-sample performance.

The rest of the paper is divided as follows. First, we formally introduce our model and assumptions in Section 2. Section 3 then presents what kind of out-of-sample guarantees can be provided on the certainty equivalent (CE) of the investor using a sample of market returns and side information when assuming a stationary market distribution. We then proceed in Section 4 to show that the same kind of guarantees can also be derived for the CE suboptimality, before showing in Section 5 what kind of behaviour can be expected in “big-data” situation. We then conclude in Section 6. All proofs have been pushed to the appendix section. Additional discussions on this topic can also be found in Bazier-Matte (2017), *i.e.*, the thesis from which these results are drawn.

2. Model and Assumptions

We consider a classical financial portfolio selection problem involving a risky asset with random return rate R and a risk-free asset with return rate of 0% for simplicity. We also suppose that the investor's risk aversion can be characterized using expected utility theory using a strictly increasing concave utility function u , and that the investor has access to side information regarding the returns. This information might be the result of processing the most recent financial or economic news, etc. We let this information be described as a vector of p normalized random features $[X_1, X_2, \dots, X_p]$. In this context, if the the distribution F of the pair (X, R) of side information and return is known, a linear investment policy that exploits the side information optimally for this investor can be obtained by solving the following optimization problem:

$$\underset{q \in \mathbb{R}^p}{\text{maximize}} \quad E_F[u(R \cdot q^T X)], \quad (1)$$

where it is assumed that short-selling is permitted.

In practice however, the exact distribution describing the relation between X and R is not available at the time of designing the investment policy and one might instead need to exploit a sample set $s_n := \{(x_i, r_i)\}_{i=1}^n$ drawn independently and identically from F . Unfortunately, when the sample size n is relatively small compared to p , it is well known that the problem (1) using the empirical distribution \hat{F} obtained from sample s_n can suffer from severe overfitting and produce investment policies that perform badly out of sample. This is for instance illustrated in the following example.

Example. Consider for instance a case where $n = p$ and each feature X_i is independently and identically drawn from a Gaussian distribution. Given that it is well known that the probability that the random matrix $\Xi := [X_1 \ X_2 \ \dots \ X_n]$ be singular is null, then one can easily establish that problem (1) with \hat{F} is unbounded. Indeed, one can verify that $r_i \bar{q}^T x_i = 1$ for all $i = 1, \dots, n$ when \bar{q} is set to $\Xi^{-1T} [1/r_1 \ 1/r_2 \ \dots \ 1/r_n]^T$. Hence, one can achieve an arbitrarily large empirical expected utility by investing according to $\alpha \bar{q}$ for $\alpha > 0$.

To prevent issues associated to overfitting, one might instead seek the optimal solution of the following regularized empirical expected utility maximization problem:

$$\underset{q \in \mathbb{R}^p}{\text{maximize}} \quad E_{\hat{F}}[u(R \cdot q^T X)] + \lambda \|q\|^2. \quad (2)$$

Note that when it exists, we will refer to the optimal solution of this problem as \hat{q} .

3. Out-of-sample performance bounds

The question remains of understanding what guarantees does one have regarding out of sample performance of the portfolio investment policy obtained from such a regularized problem. In particular, since utility functions are expressed in units without any physical meaning for the investor, any guarantees derived using learning theory should be reinterpreted in terms of a guarantee on the certainty equivalent¹ (in percent of return) of the risky investment produced by $\hat{q}^T X$. In other words, we will be interested

¹The fact that c is the certainty equivalent of a random return R implies that the investors is indifferent between being exposed to the risk of R or getting involved in a risk free investment that has a return rate of c .

in bounding how different the in-sample certainty equivalent performance of \hat{q} might be compared to the out-of-sample certainty equivalent performance.

In order to shed some light on this question, we first make the following assumptions.

Assumption 1. *The random return R is supported on a bounded interval $\mathcal{S}_R \subseteq [-\bar{r}, \bar{r}]$ such that $P(|R| \leq \bar{r}) = 1$.*

Assumption 2. *The random vector of side-information X is supported on bounded set \mathcal{S}_X such that $P(\|X\| \leq \xi) = 1$.*

Assumption 3. *The utility function is normalized such that $u(0) = 0$ and $\lim_{r \rightarrow 0^+} u'(r) = 1$. Furthermore, it is Lipschitz continuous with a Lipschitz constant of γ , i.e., for any $r_1 \in \mathbb{R}$ and $r_2 \in \mathbb{R}$, we have that $|u(r_1) - u(r_2)| \leq \gamma|r_1 - r_2|$.*

The first assumption is relatively realistic given that one can usually assess from historical data a large enough interval of returns which could be assumed to contain R with probability one. For instance, when looking at the last 35 years of daily returns for an index such as S&P 500, this interval can legitimately be set to $[-25\%, 25\%]$ daily returns. If some side information are not known to be bounded, the second assumption might require one to pre-process the vector of side information in order to rely on the results that will be presented. This could typically be done by projecting this vector on the surface of a ball of radius ξ when $\|X\| > \xi$, which is as simple as replacing X with $(\xi/\|X\|) \cdot X$. This assumption will be further studied in Section 5. Finally, while the last assumption is fairly common for establishing generalization bounds and can certainly accommodate any piecewise linear utility function (often used by numerical optimization methods), it is important to mention that it is not one that is commonly made in modern portfolio theory. If, for instance, an investor expresses an absolute risk aversion uniformly equal to α , this suggests the use of $u(r) := (1/\alpha)(1 - \exp(-\alpha r))$ which is not Lipschitz continuous. Fortunately, the theory that will be used only exploits the fact that the function is Lipschitz continuous on the interval $[-\bar{r}^2\xi^2/(2\lambda), \bar{r}^2\xi^2/(2\lambda)]$.

We are now in a position to exploit a well-known learning theory result to establish a bound on the out-of-sample portfolio performance of \hat{q} based its in-sample estimation:

Theorem 1. *Given that assumptions 1, 2 and 3 are satisfied, the certainty equivalent of the out-of-sample performance is at most $O(1/\sqrt{n})$ worse than the in-sample one. Specifically,*

$$\text{CE}(\hat{q}; F) \geq \text{CE}(\hat{q}; \hat{F}) - \Omega_1 / \lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; \hat{F}) + \epsilon) ,$$

where

$$\begin{aligned} \text{CE}(\hat{q}; F) &:= u^{-1}(E_F[u(R \cdot \hat{q}^T X)]) , \\ \text{CE}(\hat{q}; \hat{F}) &:= u^{-1}(n^{-1} \sum_{i=1}^n u(r_i \hat{q}^T x_i)) , \end{aligned}$$

and where

$$\Omega_1 := \frac{\bar{r}^2\xi^2}{2\lambda} \left(\frac{\gamma^2}{n} + \frac{(2\gamma^2 + \gamma + 1)\sqrt{\log(1/\delta)}}{\sqrt{2n}} \right)$$

with probability $1 - \delta$,

Our proof of Theorem 1 proceeds as follow. First, borrowing from the terminology introduced by Bousquet and Elisseeff (2002), we show that the algorithm generates \hat{q} from the sample set is β -stable. We then show that for any \hat{q} generated from a sample of F , the amount of utility generated from implementing the \hat{q} decision necessarily lies on an interval of bounded size. Given that these two conditions are satisfied, we can then rely on Bousquet-Ellisseeff's out-sample error bound theorem (typically used for inference problems) in order to establish out-of-sample guarantees in terms of expected utility. By exploiting the concavity of $u(\cdot)$, we are finally able to describe the implications in terms of certainty equivalent that are expressed in our theorem.

4. Suboptimality performance bounds

We now turn our attention to the suboptimality of the problem, *i.e.*, we would like to understand the behaviour of the performance of the empirical investment policy \hat{q} compared to the optimal policy $q^* := \arg \max_q E_F[u(R \cdot q^T X)]$. It is important to realize that in general, there are situations in which the optimal performance according of (1) could be unbounded. Thus, if one wishes to establish a bound on the suboptimality of an investment policy, it is necessary to impose additional assumptions on the class of problem that he is facing. The two following examples motivate these assumptions.

Example. Consider for instance a risk neutral investor, *i.e.*, such that $u(r) = r$ and suppose $EX_i = 0$. The expected utility simply becomes

$$E_F[u(R \cdot q^T X)] = \sum_{i=1}^n q_i \text{Cov}(R, X_i).$$

If we simply let $\bar{q}_i = \text{Cov}(R, X_i)$, it follows immediately that the expected utility of $\alpha \bar{q}$ can become arbitrarily large as α goes to infinity.

Example. Consider another example in which there exist a j for which feature X_j induces arbitrage over F , namely that $P\{RX_j > 0\} = 1$. In such a case, if we let $\bar{q}_i = 1$ only when $i = j$ and otherwise zero, then the expected utility of $\alpha \bar{q}$ can once again take an arbitrarily large value as α goes to infinity.

Given those two examples, we now introduce two new assumptions that will ensure that problem (1) is bounded, *i.e.*, it has a finite optimal solution.

Assumption 4. *The utility function is sublinear, i.e., $u(r) = o(r)$.*

Assumption 5. *The side information X induces no arbitrage opportunities, that is, for all X_i , $P\{RX_i < 0\} > 0$ and $P\{RX_i > 0\} > 0$.*

In a financial context, assumption 4 is certainly realistic since a financial investor behaviour is usually taken to be risk averse, thus implying assumption 4. As for assumption 5, this notion of arbitrage relates directly to the notion of market efficiency, and in particular to the semi-strong version of it, which states that it should be impossible for an investor to constantly beat the market using publicly available information. See Malkiel and Fama (1970) and Fama (1991) for more details.

Theorem 2. *Given that assumptions 1, 2, 3 are satisfied, the suboptimality of the*

policy \hat{q} can be expressed with confidence $1 - \delta$ by

$$\text{CE}(\hat{q}; F) \geq \text{CE}(q^*; F) - \Omega_2 / \lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; F) + \epsilon),$$

where

$$\Omega_2 = \lambda \|q^*\|^2 + \frac{8\gamma^2 \xi^2 (32 + \log(1/\delta))}{n\lambda} + \frac{2\gamma \bar{r} \xi^2}{\lambda} \sqrt{\frac{32 + \log(1/\delta)}{n}}.$$

Furthermore, if assumptions 4 and 5 are satisfied, then $\text{CE}(q^*; F)$ is finite.

The first term in Ω_2 shows that, unless the regularization constant λ is brought to zero as n increases, the empirical maximization problem (2) will asymptotically converge toward a constant suboptimality bound based on the particular market distribution F and on λ . The two other terms in Ω_2 show that this bound will be reached at a $O(1/\sqrt{n})$ rate in the same fashion as with Theorem 1. Therefore, the best suboptimality performance that can be hoped to be reached is at most $-\lambda \|q^*\|^2 / \lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; F) + \epsilon)$.

5. Big Data Phenomenon

In this section, we question how realistic assumption 2 is in a big data context. In particular, we expose two sets of natural conditions for the generation of the side information vector X that leads to motivating the use of a support set which diameter grow proportionally to the square root of p .

Example. Consider a case where every terms of X are independant from each other, while each X_i has a mean $E[X_i] = 0$, a variance $\text{Var}[X_i] = 1$, and are supported on their respective intervals $P(X_i \in [-\nu, \nu]) = 1$ for all i . By Hoeffding's inequality, one can establish that

$$P\left(\left|\|X\|^2 - \sum_{i=1}^p E[X_i^2]\right| \leq \sqrt{2p \log(\delta/2) \nu^2}\right) \geq 1 - \delta$$

so that $\|X\|^2 \in [p - \sqrt{2p \log(\delta/2) \nu^2}, p + \sqrt{2p \log(\delta/2) \nu^2}]$ with probability $1 - \delta$. Hence, any ball of fixed radius ξ will contain X with a probability that asymptotically converges to zero as p increases, more specifically $P(\|X\|^2 \leq \xi^2) \leq 2 \exp(-2p(1 - \xi^2/\sqrt{p})^2/\nu^2)$. On the other hand, this inequality somehow also prescribes that the diameter of the support \mathcal{S}_X should increase proportionally to \sqrt{p} in order to still contain X with high probability as p increases.

Example. Consider a similar case as above but where the independence assumption is dropped. In this context, although we might not have as much of a strong argument to discredit the use of a constant diameter for \mathcal{S}_X , there is still a good motivation for employing a radius that grows proportionally to \sqrt{p} . Namely, if each X_i has a mean $E[X_i] = 0$ and a variance $\text{Var}[X_i] = 1$ then the random variable $Z := \|X\|^2$ is necessarily positive with an expected value of p . Based on Markov inequality, this implies that with probability $1 - \delta$, we have that $\|X\| \leq \sqrt{p/\delta}$.

Since we believe these two examples provide strong arguments for replacing assump-

tion 2 with the assumption that it is within a ball of radius $\xi\sqrt{p}$, we reformulate our previous two results as follows.

Corollary 1. *Given that assumptions 1 and 3 are satisfied, and that $P(\|X\| \leq \xi\sqrt{p}) = 1$, the certainty equivalent of the out-of-sample performance is at most $O(p/\sqrt{n})$ worse than the in-sample one. Specifically, with probability $1 - \delta$,*

$$\text{CE}(\hat{q}; F) \geq \text{CE}(\hat{q}; \hat{F}) - \Omega_3 / \lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; \hat{F}) + \epsilon) ,$$

where

$$\Omega_3 := \frac{\gamma\bar{r}^2\xi^2}{\lambda} \left(\frac{\gamma p}{2n} + \frac{(1 + \gamma)p\sqrt{\log(1/\delta)}}{\sqrt{2n}} \right) .$$

Likewise, the suboptimality of the decision \hat{q} will reach a constant bound due to regularization at a rate of at most $O(p/\sqrt{n})$:

$$\text{CE}(\hat{q}; F) \geq \text{CE}(q^*; F) - \Omega_4 / \lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; F) + \epsilon) ,$$

where

$$\Omega_4 = \lambda\|q^*\|^2 + \frac{8\gamma^2 p \xi^2 (32 + \log(1/\delta))}{n\lambda} + \frac{2\gamma\bar{r}p\xi^2}{\lambda} \sqrt{\frac{32 + \log(1/\delta)}{n}} ,$$

with probability $1 - \delta$.

Note that assumption 2 was inspired by an early version of Rudin and Vahn (2015) who also studied asymptotic properties of a regularized decision problem in its Big data regime, *i.e.*, when n and p go to infinity simultaneously. Our analysis indicate that the convergence in accuracy that is reported with such an assumption can be misleading for many problems, *e.g.*, when the features can be considered independent from each other. In particular, Corollary 1 states that asymptotic convergence in accuracy is only guaranteed to occur when $p = o(\sqrt{n})$ and $\lambda \rightarrow 0$.

However, it is important to understand that Corollary 1 serves as a worst case scenario and that we don't necessarily expect to observe downgrading performances as soon as $n = o(p^2)$. Still, no matter what, there is a cost to pay in pouring more and more features into such a portfolio selection problem, and this cost is directly exhibited through ξ^2 and the loosening of the guarantees bound. One might therefore wish to be prudent when facing such high-risk regimes.

6. Discussion

As a conclusion, we would like to review the main messages we hope to deliver from this paper. First off, it is possible to use side information from the market, such as market news, financial indicators, economic variables and so on in order to build a portfolio with actual performance guarantees on the out-of-sample. Second, it is also possible to obtain performance bounds on the suboptimality of the empirical decision in comparison to what might have been the best decision, given full knowledge of the market. Third, there might be a cost to pay for increasing the number of side

information treated by the model when the sample size on which the decision is based on is not large enough.

That said, for “small-data” situations where $p = o(\sqrt{n})$, we believe our framework can be particularly well suited for aggregating and treating market side information in order to make a sound investment decision that is guaranteed to appeal to the investor’s perception of risk.

References

- Bazier-Matte T. 2017. Gestion de portefeuille par apprentissage de marché: Quelques notes sur les propriétés algorithmiques et statistiques d’une politique d’investissement obtenue par maximisation d’utilité espérée régularisée [master’s thesis]. HEC Montréal.
- Bertsimas D, Van Parys B. 2017. Bootstrap robust prescriptive analytics. Working draft.
- Bousquet O, Elisseeff A. 2002. Stability and generalization. *The Journal of Machine Learning Research*. 2:499–526.
- Brandt MW, Santa-Clara P, Valkanov R. 2009. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*. 22(9):3411–3447.
- Delage E, Ye Y. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*. 58(3):595–612.
- Esfahani P, Kuhn D. 2017. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. Working paper.
- Fama EF. 1991. Efficient capital markets: II. *The Journal of Finance*. 46(5):1575–1617.
- Fama EF, French KR. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*. 33(1):3–56.
- Goldfarb D, Iyengar G. 2003. Robust portfolio selection problems. *Mathematics of Operations Research*. 28(1):1–38.
- Huang D, Zhu S, Fabozzi F, Fukushima M. 2010. Portfolio selection under distributional uncertainty: A relative robust CVaR approach. *European Journal of Operations Research*. 203(1):185–194.
- Ingersoll J. 1987. *Theory of financial decision making*. Rowman & Littlefield Publishers.
- Madan DB, Carr PP, Chang EC. 1998. The variance gamma process and option pricing. *European Finance Review*. 2(1):79–105.
- Malkiel BG, Fama EF. 1970. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*. 25(2):383–417.
- Markowitz H. 1952. Portfolio selection. *The Journal of Finance*. 7(1):77–91.
- Mohri M, Rostamizadeh A, Talwalkar A. 2012. *Foundations of machine learning*. MIT press.
- Olivares-Nadal AV, DeMiguel V. 2018. Technical note—a robust perspective on transaction costs in portfolio optimization. Working draft.
- Rudin C, Vahn GY. 2015. The big data newsvendor: Practical insights from machine learning. Working draft.
- Salmon F. 2009. Recipe for disaster: the formula that killed Wall Street. *WIRED Magazine*.
- Sridharan K, Shalev-Shwartz S, Srebro N. 2009. Fast rates for regularized objectives. In: *Advances in Neural Information Processing Systems*. p. 1545–1552.
- von Neumann J, Morgenstern O. 1944. *Theory of games and economic behavior*. Princeton University Press.

7. Appendix

7.1. Proof of Theorem 1

In this proof, we will employ a theorem made famous by Bousquet-Elisseef to analyse relevant asymptotic statistical properties of the following estimator.

Definition 7.1. Let $\hat{\mathbf{q}} : \mathbb{R}^{(p+1) \times n} \rightarrow \mathbb{R}^p$ be the procedure that generates the optimal solution of problem (2) based on a sample set $\{(x_i, r_i)\}_{i=1}^n$.

We start by presenting two lemmas that establish some important properties of problem (2).

Lemma 7.2. *When assumptions 1 and 3 are satisfied, the estimator $\hat{\mathbf{q}}(\cdot)$ has β -stability with $\beta = \frac{(\gamma\bar{r}\xi)^2}{2\lambda n}$. Namely, for any two sample sets $s_n^1 := \{(x_i^1, r_i^1)\}_{i=1}^n$ and $s_n^2 := \{(x_i^2, r_i^2)\}_{i=1}^n$ that are exactly identical except for the j -th sample, i.e., $(x_i^1, r_i^1) = (x_i^2, r_i^2)$ for all $i \neq j$, the following holds:*

$$|u(r \hat{\mathbf{q}}(s_n^1)^T x) - u(r \hat{\mathbf{q}}(s_n^2)^T x)| \leq \beta, \forall x \in \mathcal{S}_X, \forall r \in \mathcal{S}_R.$$

Proof. First, following the terminology presented in Bousquet and Elisseeff (2002) (see Definition 19), we can establish that $\hat{\mathbf{q}}(\cdot)$ has σ -admissibility of $\gamma\bar{r}$. This is simply done by exploiting the fact that \mathcal{S}_R is bounded and that $u(\cdot)$ is Lipschitz continuous. The detailed derivations consider that for any pair $(q_1, q_2) \in \mathbb{R}^p \times \mathbb{R}^p$, one has that

$$|u(r q_1^T x) - u(r q_2^T x)| \leq \gamma |r q_1^T x - r q_2^T x| \leq \gamma\bar{r} |q_1^T x - q_2^T x|, \forall r \in \mathcal{S}_R, \forall x \in \mathcal{S}_X.$$

The β -stability of $\hat{\mathbf{q}}(\cdot)$ then follows directly from Theorem 22 in Bousquet and Elisseeff (2002). \square

Lemma 7.3. *When assumptions 1, 2 and 3 are satisfied, the maximum difference in amount of utility attained by implementing two investment strategies obtained using different sample sets s_n^1 and s_n^2 is bounded by*

$$|u(r \hat{\mathbf{q}}(s_n^1)^T x) - u(r \hat{\mathbf{q}}(s_n^2)^T x)| \leq u_{range} := \frac{(\gamma + 1)\xi^2 \bar{r}^2}{2\lambda}, \forall x \in \mathcal{S}_X, \forall r \in \mathcal{S}_R.$$

Proof. This proof relies mostly on demonstrating that $\|\hat{\mathbf{q}}(s_n)\| \leq B$ for some $B > 0$ with probability one for all possible sample sets s_n . Indeed, when this is the case, then we have that

$$|u(r \hat{\mathbf{q}}(s_n^1)^T x) - u(r \hat{\mathbf{q}}(s_n^2)^T x)| \leq u(\bar{r}\xi B) - u(-\bar{r}\xi B) \leq (\gamma + 1)\bar{r}\xi B.$$

In order to show that $\hat{\mathbf{q}}(s_n)$ is bounded, we reformulate problem (2) as follows

$$\begin{aligned} & \underset{s \in \mathbb{R}, v \in \mathbb{R}^p}{\text{maximize}} && \frac{1}{n} \sum_{i=1}^n u(s R_i X_i^T v) - \lambda s^2 \\ & \text{s. t.} && s \geq 0, \|v\| = 1, \end{aligned}$$

such that $\hat{\mathbf{q}}(s_n) = s^* \cdot v^*$ when (s^*, v^*) is the pair of optimal assignments for this optimization problem. It is therefore clear that $s^* = \|\hat{\mathbf{q}}(s_n)\|$ and our proof reduces

to establishing an upper bound for s^* . By recognizing that $s^* = \arg \max_{s \geq 0} g(s) := \frac{1}{n} \sum_{i=1}^n u(sR_i X_i^T v^*) - \lambda s^2$ and that $g(s)$ is a concave function, then it is necessarily the case that if there exists a $\bar{s} \geq 0$ such that $g(\cdot)$ is non-increasing at \bar{s} then $s^* \leq \bar{s}$. We can actually show that this is the case for $\bar{s} := \bar{r}\xi/(2\lambda)$ by upper bounding the impact of taking a step of $\Delta > 0$:

$$\begin{aligned}
g(\bar{s} + \Delta) - g(\bar{s}) &= \frac{1}{n} \sum_{i=1}^n (u((\bar{s} + \Delta)R_i X_i^T v^*) - u(\bar{s}R_i X_i^T v^*)) - \lambda((\bar{s} + \Delta)^2 - \bar{s}^2) \\
&\leq \frac{1}{n} \sum_{i=1}^n (u((\bar{s} + \Delta)|R_i X_i^T v^*|) - u(\bar{s}|R_i X_i^T v^*|)) - \lambda((\bar{s} + \Delta)^2 - \bar{s}^2) \\
&\leq \frac{1}{n} \sum_{i=1}^n \Delta R_i X_i^T v^* - \lambda(2\bar{s}\Delta + \Delta^2) \\
&\leq \Delta \bar{r}\xi - 2\lambda\bar{s}\Delta - \Delta^2 = -\Delta^2 \leq 0,
\end{aligned}$$

where we first used the fact that $u(\cdot)$ is increasing, next that $u(y + \Delta) \leq u(y) + \Delta$ when $\Delta \geq 0$ since it is a concave function with a subgradient of one at zero. Finally, we exploited assumptions 1 and 2. This completes our proof. \square

Having established the above properties, the following theorem follows directly from Theorem 3. While we omit to describe the original theorem in this article for sake of compactness, we refer interested readers to the form presented in Theorem 11.1 of Mohri et al. (2012) for more details.

Theorem 3 (Bousquet-Ellisseeff Outsample Error Theorem). *Given that assumptions 1, 2, and 3 are satisfied, then one has with confidence of $1 - \delta$ that*

$$E_F[u(R \cdot \hat{\mathbf{q}}(s_n)^T X)] \geq E_{\hat{F}}[u(R \cdot \hat{\mathbf{q}}(s_n)^T X)] - \beta - (2n\beta + \hat{u}_{abs}) \sqrt{\frac{\log(1/\delta)}{2n}},$$

where β refers to the β -stability of $\hat{\mathbf{q}}$ and \hat{u}_{abs} refers to a uniform bound $P(|u(R\hat{\mathbf{q}}(s_n)X)| \leq \hat{u}_{abs}) = 1$. Overall, this reduces to $E_F[u(R \cdot \hat{\mathbf{q}}(s_n)^T X)] \geq E_{\hat{F}}[u(R \cdot \hat{\mathbf{q}}(s_n)^T X)] - \Omega_1$. Hence, the out-of-sample performance in terms of expected utility of the investment policy $\hat{\mathbf{q}}(s_n)$ is at most $O(1/\sqrt{n})$ worse than the in-sample one.

We conclude this section by demonstrating how Theorem 1 follows from Theorem 3. In particular, by concavity of the utility function, we have that

$$u(\text{CE}(\hat{q}; F)) \leq u(\text{CE}(\hat{q}; \hat{F})) + (\text{CE}(\hat{q}; F) - \text{CE}(\hat{q}; \hat{F})) \partial u(\text{CE}(\hat{q}; \hat{F})),$$

where $\partial u(r)$ denotes any supergradient of $u(\cdot)$ at r . In particular, since $u(\cdot)$ is an increasing concave function, it follows that $\lim_{\epsilon \rightarrow 0^-} u'(\text{CE}(\hat{q}; \hat{F}) + \epsilon) \geq 0$ is one of the supergradient at $\text{CE}(\hat{q}; \hat{F})$. Combining this inequality with the inequality presented in Theorem 3, we get

$$u(\text{CE}(\hat{q}; \hat{F})) - \Omega_1 \leq u(\text{CE}(\hat{q}; \hat{F})) + (\text{CE}(\hat{q}; F) - \text{CE}(\hat{q}; \hat{F})) \partial u(\text{CE}(\hat{q}; \hat{F}))$$

so that

$$\text{CE}(\hat{q}; F) \geq \text{CE}(\hat{q}; \hat{F}) - \Omega_1 / \partial u(\text{CE}(\hat{q}; \hat{F}))$$

follows since it was assumed that $u(\cdot)$ is strictly increasing. This completes the proof of Theorem 1.

7.2. Proof of Theorem 2

First, in order to tidy up the proof, let us define $EU(q) := E_F(u(R \cdot q^T X))$ and $EU_\lambda(q) := E_F(u(R \cdot q^T X)) - \lambda \|q\|^2$, with $q^* := \arg \min_q EU(q)$ and $q_\lambda^* := \arg \min_q EU_\lambda(q)$.

Theorem 4 (Theorem 1 and surrounding text in Sridharan et al. (2009)). *Given that assumptions 1, 2, and 3 are satisfied, and since the function EU_λ is 2λ -strongly convex, then one has with confidence of $1 - \delta$ that*

$$-\lambda \|\hat{q} - q_\lambda^*\|^2 \geq EU_\lambda(\hat{q}) - EU_\lambda(q_\lambda^*) \geq -\omega,$$

where

$$\omega = \frac{4\gamma^2 \xi^2 (32 + \log(1/\delta))}{\lambda n}.$$

Notice that Theorem 4 implies with confidence $1 - \delta$ that

$$EU(\hat{q}) - EU(q_\lambda^*) \geq \lambda(\|\hat{q}\|^2 - \|q_\lambda^*\|^2) - \omega \geq -\lambda(\|\hat{q} - q_\lambda^*\|^2 + 2\|\hat{q}\|\|\hat{q} - q_\lambda^*\|) - \omega.$$

As shown in Section 7.1, $\|\hat{q}\| \leq \bar{r}\xi/(2\lambda)$. Theorem 4 further implies concerning the same $1 - \delta$ probability outcomes that $\|\hat{q} - q_\lambda^*\|^2 \leq \omega/\lambda$, and therefore $\|\hat{q} - q_\lambda^*\| \leq \sqrt{\omega/\lambda}$, so that we end up with

$$EU(\hat{q}) - EU(q_\lambda^*) \geq -2\omega - \bar{r}\xi \sqrt{\frac{\omega}{\lambda}}.$$

with probability $1 - \delta$. Finally, note that since by the definition of q_λ^* , we have that

$$EU(q^*) - \lambda \|q^*\|^2 \leq EU(q_\lambda^*) - \lambda \|q_\lambda^*\|^2,$$

it follows that

$$EU(q^*) - EU(q_\lambda^*) \leq \lambda(\|q^*\|^2 - \|q_\lambda^*\|^2) \leq \lambda \|q^*\|^2,$$

so that we can bound the suboptimality of the policy \hat{q} with probability $1 - \delta$ in the following fashion:

$$\begin{aligned} EU(\hat{q}) &= EU(q^*) + EU(\hat{q}) - EU(q_\lambda^*) + EU(q_\lambda^*) - EU(q^*) \\ &\geq EU(q^*) - \lambda \|q^*\|^2 - \bar{r}\xi \sqrt{\frac{\omega}{\lambda}} - 2\omega. \end{aligned}$$

This relation can be exploited in a similar way as in the proof of Theorem 1 (see Section 7.1) to derive the relation between certainty equivalents that is presented in our theorem. the same trick as in Section .

Next, we show that $\text{CE}(q^*; F) - \text{CE}(\hat{q}; F)$ is bounded by proving that $\|q^*\|$ is finite. Since the other terms of the upper bound established above are also finite, the second part of Theorem 2 follows.

Instead of optimizing q directly, as was done previously, we can reformulate problem (1) in terms of both an orientation vector and a scale decision variable. This gives us

$$\begin{aligned} & \underset{s \in \mathbb{R}, v \in \mathbb{R}^p}{\text{maximize}} && E[u(sR X^T v)] \\ & \text{s. t.} && s \geq 0, \|v\| = 1. \end{aligned}$$

Based on assumption 5, since no feature induce arbitrage, it follows that, there exists a $\delta > 0$ such that $P\{R X^T v < -\delta\} = \varrho > 0$ for all v with a norm of one. Now, let B be a discrete random variable with two states such that $P\{B = -\delta\} = 1 - P\{B = \bar{r}\xi\} = \varrho$. Since $|R X^T v| < \bar{r}\xi$, we have that $P\{B \geq r\} \geq P\{R X^T v \geq r\}$ for all $r \in \mathbb{R}$, i.e. that B stochastically dominates $R X^T v$, so that it must necessarily follow that $E[u(sB)] \geq E[u(sR X^T v)]$. But, by the sublinearity assumption on u ,

$$\begin{aligned} \lim_{s \rightarrow \infty} E[u(sR X^T v)] &\leq \lim_{s \rightarrow \infty} E[u(sB)] = \lim_{s \rightarrow \infty} (\varrho u(-s\delta) + (1 - \varrho)u(s\bar{r}\xi)) \\ &\leq \lim_{s \rightarrow \infty} -\varrho s\delta + (1 - \varrho)o(s) = -\infty \end{aligned}$$

for all v of norm one, which shows that s^* , and therefore $\|q^*\|$, is bounded.