

# Forecasting with fractional Brownian motion: a financial perspective

Matthieu Garcin\*

May 20, 2021

## Abstract

The fractional Brownian motion (fBm) extends the standard Brownian motion by introducing some dependence between non-overlapping increments. Consequently, if one considers for example that log-prices follow an fBm, one can exploit the non-Markovian nature of the fBm to forecast future states of the process and make statistical arbitrages. We provide new insights into forecasting an fBm, by proposing theoretical formulas for accuracy metrics relevant to a systematic trader, from the hit ratio to the expected gain and risk of a simple strategy. In addition, we answer some key questions about optimizing trading strategies in the fBm framework: Which lagged increments of the fBm, observed in discrete time, are to be considered? If the predicted increment is close to zero, up to which threshold is it more profitable not to invest? We also propose empirical applications on high-frequency FX rates, as well as on realized volatility series, exploring the rough volatility concept in a forecasting perspective.

**Keywords** – fractional Brownian motion, Hurst exponent, systematic trading, rough volatility, foreign-exchange rates

## 1 Introduction

An fBm is a non-Markovian process. A question thus naturally arises when considering the use of this model in finance for describing the dynamic of prices or log-prices: Does the fBm induce pure arbitrages or not? This question has led to numerous articles [7]. In particular, it has been shown that arbitrage opportunities exist when trading in continuous time is allowed [38]. But the answer is different when one trades in discrete time, even at a very high frequency [13]. From the perspective of a financial market practitioner, beyond the strong mathematical interest of the question, the distinction is not so clear between continuous time and discrete time with arbitrary high frequency. Therefore, like others before, we believe that this debate about the existence of arbitrages implied by this model is mainly motivated by theoretical convenience, what is not enough to exclude the fBm as a model of log-prices [17].

A more recent literature deals with the use of fBm in statistical arbitrage [31, 30]. Exploiting the autocovariance structure of the fBm indeed makes it possible to forecast future states of the process [36]. This simple idea has been used since in algorithmic trading to build systematic

---

\*Léonard de Vinci Pôle Universitaire, Research center, 92916 Paris La Défense, France, matthieu.garcin@m4x.org. The author is grateful to Chafic Merhy and Konstantinos Tsianos for valuable discussions and comments.

strategies. Such strategies do not lead to a certain gain as do the pure arbitrages previously mentioned, but they are profitable in average. All the contributions on this subject highlight the asymmetry of the performance of these strategies with respect to  $H = 1/2$ , where  $H$  is the Hurst exponent.

Asymmetry with respect to  $H = 1/2$  is also underlined in the econophysics literature. Indeed, when  $H > 1/2$ , the increments of the fBm are positively correlated and the process also has a long memory. By contrast, when  $H < 1/2$ , non-overlapping increments are negatively correlated and there is no long memory, in the sense that the autocovariance function decreases exponentially. Many researchers equate the notion of market efficiency with the property of long memory, leading them to consider that the market is efficient if and only if  $H > 1/2$  [19, 12, 20, 4, 21, 29]. We indeed acknowledge that when  $H < 1/2$ , the performance of the fBm-based predictor does not always lead to satisfying empirical results [22]. But this practical limitation does not come from a particular property of the fBm itself. It comes instead from the fact that the model is not well specified or from the related difficulties to estimate the parameters properly, in a nutshell from model risk [23, 24]. Besides the identification of market efficiency to long-range dependence, another branch of the econophysics literature considers that the market is efficient for  $H = 1/2$  and that its inefficiency gradually increases as  $H$  gets away from  $1/2$  [33, 34, 32, 10, 5]. Indeed, predictions do not rely on the long-range dependence property but, instead, on the autocovariance of the process, which is different from zero as soon as  $H \neq 1/2$ .

While the fBm is widespread in econophysics and more recently in mainstream quantitative finance, not everything has been said yet on forecasting the fBm in a financial perspective. First, the theoretical evaluation of the quality of the forecast invoked in the literature is rarely very appropriate for systematic traders. For example, the mean squared error (MSE), put forward in several papers [36, 28], albeit statistically relevant, is not related to the performance and risk of a trading strategy. In this perspective, other metrics have to be applied in the fBm setting. Moreover, while many articles focus on the fBm in continuous time, the reality of trading, which is affected by liquidity frictions, is in discrete time, whether for the observation of the price process or for the instants at which one is able to trade. And this limitation has some consequences. For example, determining which lagged price returns should be used as input of the predictor is overriding in order to optimize an fBm-based forecast and related systematic trading strategies.

We answer these questions in the present paper. We introduce several accuracy metrics for covariance-based predictors: a hit ratio, an average gain, a risk defined as a lower semi-deviation, and the resulting risk-adjusted performance. We provide theoretical expressions for all these metrics using the fBm assumption. We find numerically a general expression for the duration of the optimal time lags of price returns used as input variables of the predictor. We also propose two empirical applications highlighting the relevance of predictions based on the fBm in finance. In particular, we maximize an ex-ante risk-adjusted performance of a trading strategy on FX rates, where log-prices follow an fBm. We also investigate another application of the fBm in finance than the sole log-price process and for which a performing forecast method is useful. We indeed study series of realized volatilities, exploring the rough volatility and how one can optimize the forecast of volatility in this context, focusing on hit ratios.

Guasoni and co-authors have recently interestingly contributed on the subject of forecasting the fBm in a financial perspective [31, 30]. But their approach is different from ours. First, they work in continuous time or, asymptotically, with an arbitrary high trading frequency. We work more realistically in discrete time. Second, they consider a mean-variance framework. While the variance is a widespread deviation risk measure, we think that a lower semi-deviation more appropriately depicts downside risk. Despite these two differences, which lead to very different formulas, some of our conclusions are qualitatively similar: the closer the Hurst exponent  $H$  is to  $1/2$ , the worst is the quality of the forecast, and, asymmetrically, the forecast is better for  $H > 1/2$  than for  $1 - H$ .

The new insight we bring is about optimizing a systematic strategy. We do this in two directions. First, we optimally select the input lagged price returns of the predictor. Second, we define an optimal threshold under which we consider that the forecast is too close to zero for an investment to positively impact the risk-adjusted profit. This last idea finally leads to reducing the trading frequency and thus to take into account liquidity constraints. Guasoni and co-authors also deal with this challenge in a different manner, by making a judicious choice of the trading frequency, but still remaining in an asymptotic framework [31, 30].

The rest of the paper is organized as follows. In Section 2, we recall some basic and useful properties of the fBm. Section 3 introduces the art of forecasting an fBm, namely the covariance-based predictor, its accuracy measured by a hit ratio, and the optimal selection of its input. In Section 4, we focus on statistical arbitrage, providing both a simple systematic trading strategy and related accuracy metrics. In Section 5, we present empirical results on realized volatility and FX rates. Section 6 concludes.

## 2 Model description: preliminaries on the fBm

The fBm, introduced by Mandelbort and van Ness in 1968, can follow several equivalent definitions [35]. Among them, we can cite the integral-based definition.

**Definition 1.** *An fBm of Hurst exponent  $H \in (0, 1)$  and volatility parameter  $\sigma > 0$  is a stochastic process  $X_t$  such that,  $\forall t \in \mathbb{R}$ ,*

$$X_t = \frac{\sigma}{\Gamma(H + \frac{1}{2})} \int_{-\infty}^{\infty} \left( (t-s)_+^{H-1/2} - (-s)_+^{H-1/2} \right) dW_s,$$

where  $W_s$  is a standard Brownian motion.

In Definition 1, the fBm is to be seen as a weighted average of a Gaussian white noise. The weights are defined by the kernel  $(s, t) \mapsto \left( (t-s)_+^{H-1/2} - (-s)_+^{H-1/2} \right)$ , whose right part,  $-(-s)_+^{H-1/2}$ , is simply intended to make the process  $X_t$  equal to zero when  $t = 0$ , which is a feature also stated in other definitions of the fBm. One major motivation for introducing the fBm was to define a process with other scaling properties than the standard Brownian motion. This alternative scaling clearly appears in Definition 2, in the expression of the variance of an increment of the process  $X_t$ , but it is not so obvious in the integral-based definition. The idea of Mandelbrot and van Ness to obtain a specific scaling was to use fractional calculus. Indeed, Definition 1 corresponds to the fractional derivative (respectively integral) of order  $1/2 - H$  (resp.  $H - 1/2$ ) of a standard Brownian motion, if  $H < 1/2$  (resp.  $H > 1/2$ ), the case  $H = 1/2$  corresponding to the standard Brownian motion itself. The scaling feature characterized by a given Hurst exponent is thus to be related to a specific autocovariance of the increments of the process in the fBm model.

Since we will work in discrete time in all the paper, we introduce another definition of the fBm, which is equivalent to Definition 1.

**Definition 2.** *An fBm of Hurst exponent  $H \in (0, 1)$  and volatility parameter  $\sigma > 0$  is a stochastic Gaussian process  $X_t$  such that  $X_0 = 0$ ,  $\mathbb{E}\{X_t\} = 0$ , and*

$$\mathbb{E}\{(X_t - X_s)^2\} = \sigma^2 |t - s|^{2H}$$

for all  $s, t \in \mathbb{R}$ .

Thanks to Definition 2, a particular property of scaling, or selfsimilarity, clearly appears for the fBm. This is the purpose of Proposition 1.

**Proposition 1.** *Let  $X_t$  be an fBm of Hurst exponent  $H \in (0, 1)$ . Then, the process  $X_t$  is statistically  $H$ -selfsimilar, that is to say,  $\forall c > 0$ , the random variables  $X_{ct}$  and  $c^H X_t$ , for a given  $t \in \mathbb{R}$ , follow the same distribution.*

*Proof.* According to Definition 2, the two variables  $X_{ct}$  and  $c^H X_t$  are Gaussian, have both a mean equal to zero and a variance equal to  $\sigma^2 c^{2H} |t|^{2H}$ . They thus follow exactly the same distribution.  $\square$

This scaling property of the fBm also leads to a specific autocovariance of the process and of its increments, as exposed in Proposition 2.

**Proposition 2.** *Let  $X_t$  be an fBm of Hurst exponent  $H \in (0, 1)$  and volatility parameter  $\sigma > 0$ . Then, the covariance of  $X_t$  and  $X_s$  is*

$$\mathbb{E}\{X_t X_s\} = \frac{\sigma^2}{2} (|t|^{2H} + |s|^{2H} - |t-s|^{2H}) \quad (1)$$

and the covariance between increments  $X_t - X_s$  and  $X_v - X_u$  is

$$\mathbb{E}\{(X_t - X_s)(X_v - X_u)\} = \frac{\sigma^2}{2} (|u-t|^{2H} + |v-s|^{2H} - |v-t|^{2H} - |u-s|^{2H}). \quad (2)$$

*Proof.* Using Definition 2 and in particular its consequence that  $\mathbb{E}\{X_t^2\} = \mathbb{E}\{(X_t - X_0)^2\} = \sigma^2 |t|^{2H}$ , the proof of Proposition 2 is straightforward:

$$\begin{aligned} \mathbb{E}\{X_t X_s\} &= -\frac{1}{2} (\mathbb{E}\{(X_t - X_s)^2\} - \mathbb{E}\{X_t^2\} - \mathbb{E}\{X_s^2\}) \\ &= -\frac{1}{2} (\sigma^2 |t-s|^{2H} - \sigma^2 |t|^{2H} - \sigma^2 |s|^{2H}) \\ &= \frac{\sigma^2}{2} (|t|^{2H} + |s|^{2H} - |t-s|^{2H}) \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}\{(X_t - X_s)(X_v - X_u)\} &= \mathbb{E}\{X_t X_v\} - \mathbb{E}\{X_t X_u\} - \mathbb{E}\{X_s X_v\} + \mathbb{E}\{X_s X_u\} \\ &= \frac{\sigma^2}{2} (|u-t|^{2H} + |v-s|^{2H} - |v-t|^{2H} - |u-s|^{2H}). \end{aligned}$$

$\square$

From Proposition 2, one sees that the fBm is not a stationary process but its increments are stationary, exactly like the standard Brownian motion. This means that, for making forecasts, we will more easily work with increments than with the process itself.

In all what follows, the process  $X_t$  is an fBm of Hurst exponent  $H$  and volatility parameter  $\sigma > 0$ . In the empirical part of the paper,  $X_t$  will either depict a log-price or the volatility process of a price. We thus cover the two main applications of the fBm in finance. The first one has been put forward by econophysicists and the second one has known a recent interest in mathematical finance with the rise of the literature about rough volatility. For simplification, we will sometimes evoke log-prices in the theoretical developments, but the approaches can also often be applied to volatilities and to any other dynamic described by an fBm, in particular in Section 3. However, Section 4.2, in which we determine the profit and risk of a systematic strategy, is only appropriate for log-prices.

### 3 Forecasting an fBm

In this section, we present the standard formula for the forecast of the process in discrete time along with various accuracy metrics.

### 3.1 Covariance-based forecast

We want to predict the value of the process  $X_t$  for a time horizon  $h$ , that is  $X_{t+h}$ , conditionally to observations at some times  $I = \{t_1, \dots, t_n\}$ , such that  $\forall s \in I, s \leq t$ , where  $t$  is the current time. Writing  $Y = (X_{t_1} \dots X_{t_n})^T$ , the forecast minimizing the MSE is  $\hat{X}_{t+h|Y} = \mathbb{E}\{X_{t+h}|Y\}$ . We recall that  $X_t$  is an fBm. For this reason, we can easily determine the covariance between  $X_s$  and  $X_t$  at any times  $s, t \in \mathbb{R}$ , thanks to equation (1), and we can even build covariance matrices for  $X_{t+h}$  and the vector  $Y$ . This makes it possible to express explicitly  $\hat{X}_{t+h|Y}$  as well as its MSE. This is the purpose of Proposition 3, for which we do not detail the proof since it is a direct consequence of the Gaussian conditioning theorem [36].

**Proposition 3.** *Let  $X_t$  be an fBm,  $h > 0$ ,  $Y = [X_{t_1} \dots X_{t_n}]^T$ , with  $\forall i < j, t_i \neq 0$  and  $t_i < t_j$ . The estimator of  $X_{t+h}$  minimizing the MSE conditionally to  $Y$  is*

$$\hat{X}_{t+h|Y} = \mathbb{E}\{X_{t+h}|Y\} = \Sigma_{XY} \Sigma_Y^{-1} Y$$

and the corresponding MSE is

$$\mathbb{E}\{(\hat{X}_{t+h|Y} - X_{t+h})^2\} = \Sigma_X - \Sigma_{XY} \Sigma_Y^{-1} \Sigma_{XY}^T,$$

where  $\Sigma_Y = \mathbb{E}\{YY^T\}$ ,  $\Sigma_X = \mathbb{E}\{X_{t+h}^2\}$ , and  $\Sigma_{XY} = \mathbb{E}\{X_{t+h}Y^T\}$  can be explicitly expressed thanks to equation (1).

We also stress the fact that using existing results about forecasting an fBm in continuous time in our framework is less relevant. Indeed, the adaptation of the continuous-time formulas to the discrete-time framework needs the discretization of an integral and therefore only leads to an approximation of  $\mathbb{E}\{X_{t+h}|Y\}$ . On the contrary, exploiting the covariance matrix of an fBm is both easier and more accurate.

It is also well known that, with the fBm model, the prediction at time horizon  $h$  is optimal when the only past observation considered has a time lag also equal to  $h$ . In particular, the optimality can be understood as the minimization of the MSE. In other words, for  $n = 2$ , considering  $t_1 = t$  and  $t_2 = t - h$  is the choice of dates which minimizes the MSE in Proposition 3 [36].

### 3.2 Hit ratio

The MSE, while being a useful statistical tool, may be less relevant in finance than other statistics, such as the hit ratio, which is the probability to make a good forecast of the sign of a future price return. Indeed, in order to make an investment decision, a trader forecasts at time  $t$  the price return  $R_{t,t+h}$  between  $t$  and  $t+h$ , defined by

$$R_{t,t+h} = X_{t+h} - X_t, \tag{3}$$

by a specific predictor  $\hat{R}_{t,t+h}$ . If  $\hat{R}_{t,t+h} > 0$  (respectively  $< 0$ ), the trader must be long (resp. short) between  $t$  and  $t+h$  in order to expect a gain at time  $t+h$ . Therefore, the hit ratio is implicitly related to the performance of a trading strategy. Every trader is thus able to interpret the value of a hit ratio and to discard any forecast with a hit ratio lower than 50%. On the contrary, the interpretation of the MSE is not as clear and this metric is only useful for comparing several forecasting methods on the same dataset, not for judging the intrinsic quality of a given predictor.

In our framework, log-prices follow an fBm. Therefore, log-returns are simply increments of the fBm. This will ease the calculation and the analysis of the hit ratio, since increments of the fBm are stationary whereas the fBm itself is not stationary. By the way, the forecast is a weighted sum

of past observations. If one considers prices instead of returns, the weights will depend on  $t$  and not only on the time lags. This is striking for example if  $t = h$ . Indeed, as exposed above, when considering two observations for forecasting, the optimal time lag in the past should be  $h$  and thus one should forecast  $X_{t+h}$  using  $X_t$  and  $X_{t-h} = X_0 = 0$ . But, in this case, the matrix  $\Sigma_Y$  is not invertible and we cannot apply Proposition 3.

Before determining the hit ratio, we need to detail the predictor of the true future log-return  $R_{t,t+h}$ . We forecast  $R_{t,t+h}$  conditionally to a vector  $S$  of adjacent past price returns. The vector  $S$  is built with the help of the set of time lags  $\Delta = \{\delta_0, \dots, \delta_n\}$ , with  $n \geq 1$  and  $\forall i < j, \delta_i < \delta_j$  and  $\delta_0 \geq 0$ . More precisely, we have

$$S = [R_{t-\delta_1, t-\delta_0} \dots R_{t-\delta_n, t-\delta_{n-1}}]^T.$$

For convenience, we set  $\delta_0 = 0$ , because it is natural to consider the current state in our forecast, but for the theoretical results exposed below, this constraint is not mandatory. We note our predictor  $\hat{R}_{t,t+h|\Delta}$ . The extension of Proposition 3 to the increments of the process instead of the process itself is straightforward, since it simply consists in replacing the matrices of covariance of the fBm by the matrices of covariance of its increments.

**Proposition 4.** *Let  $X_t$  be an fBm of Hurst exponent  $H \in (0, 1)$  and volatility parameter  $\sigma > 0$ ,  $h > 0$ ,  $S = [R_{t-\delta_1, t-\delta_0} \dots R_{t-\delta_n, t-\delta_{n-1}}]^T$ , with  $\forall i < j, \delta_i < \delta_j$ ,  $\delta_i \geq 0$ , and  $R_{\cdot, \cdot}$  expressed by equation (3). The estimator of  $R_{t,t+h}$  minimizing the MSE conditionally to  $S$  is*

$$\hat{R}_{t,t+h|\Delta} = \Sigma_{RS} \Sigma_S^{-1} S \tag{4}$$

and the corresponding MSE is

$$\mathbb{E}\{(\hat{R}_{t,t+h|\Delta} - R_{t,t+h})^2\} = \Sigma_R - \Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T,$$

where  $\Sigma_S = \mathbb{E}\{SS^T\}$ ,  $\Sigma_R = \mathbb{E}\{R_{t,t+h}^2\} = \sigma^2 h^{2H}$ , and  $\Sigma_{RS} = \mathbb{E}\{R_{t,t+h} S^T\}$  can be explicitly expressed thanks to equation (2).

It is worth noting that  $\hat{R}_{t,t+h|\Delta}$  does only depend on  $h$ , not on  $t$ . In what follows, we keep  $t$  in most subscripts, but the hit ratios, risk measures, and expected performance displayed in the coming theorems and propositions in fact only depend on the forecast horizon  $h$ .

One can write equation (4) differently, as a weighted sum of past returns:

$$\hat{R}_{t,t+h|\Delta} = \sum_{i=1}^n \beta_i R_{t-\delta_i, t-\delta_{i-1}},$$

where the  $\beta_i$  must be consistent with the covariance approach exposed in Proposition 4:

$$[\beta_1 \dots \beta_n] = \Sigma_{RS} \Sigma_S^{-1}.$$

In the particular case where  $n = 1$ , that is when the predicted return only relies on one past return, namely  $\hat{R}_{t,t+h|\Delta} = \beta_1 R_{t-\delta_1, t}$ , we have the following expression for  $\beta_1$ :

$$\beta_1 = \frac{\text{Cov}(R_{t,t+h}, R_{t-\delta_1, t})}{\text{Var}(R_{t-\delta_1, t})},$$

according to Proposition 4. This is not surprising since one also uses this kind of ratio to estimate the  $\beta$  in the *capital asset pricing model* (CAPM). Nevertheless, the framework is different here. Indeed, in the CAPM, the covariance and variance are usually estimated with their empirical version, whereas the  $\beta_1$  in this work is model-dependent. The fBm assumption will in particular

make it possible to rely on a parsimonious representation of the dependence structure among price returns, even when considering a big number of time lags, and thus to avoid overfitting in the forecast. This specific feature of parsimony will be studied empirically in Section 5.1. Taking into account the fBm assumption, the expression for  $\beta_1$ , when  $n = 1$  is:

$$\beta_1 = \frac{1}{2} \left[ \left( \frac{h}{\delta_1} + 1 \right)^{2H} - \left( \frac{h}{\delta_1} \right)^{2H} - 1 \right], \quad (5)$$

according to equation (2). It is worth noting that  $\beta_1$  is positive when  $H > 1/2$ , equal to zero for  $H = 1/2$ , and negative when  $H < 1/2$ . It is consistent with properties of the fBm: increments are positively correlated when  $H > 1/2$ , non correlated when  $H = 1/2$ , and negatively correlated when  $H < 1/2$ . In particular, we note that when  $\delta_1 = h$ ,  $\beta_1$  tends toward 1 when  $H$  tends toward 1: the persistence is such that the best prediction of the future return is the past return of same duration. Alternatively, in the anti-persistence case, when  $H$  tends toward 0, whatever the value of  $\delta_1$ ,  $\beta_1$  tends toward  $-1/2$ .

In the general case, that is for  $n$  not necessarily equal to 1, we want to determine the theoretical hit ratio of our predictor. As written above, it is the probability to forecast properly the sign of the future price return  $R_{t,t+h}$ . Two definitions are in fact possible, depending on the conditioning of the probability to past returns or not.

**Definition 3.** Let  $\hat{R}_{t,t+h|\Delta}$  be a predictor of  $R_{t,t+h}$  based on the vector  $S$  of past returns, as defined by equation (4). The conditional hit ratio of this predictor is

$$\rho^c(y) = \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} R_{t,t+h} \geq 0 \mid S = y \right]$$

and the non-conditional hit ratio is

$$\rho = \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} R_{t,t+h} \geq 0 \right].$$

In the fBm framework, the vector  $S$  admits a Gaussian density  $g_S$  of mean zero and variance  $\Sigma_S$ . The non-conditional hit ratio is the weighted average of the conditional hit ratio:

$$\rho = \int_{\mathbb{R}} \rho^c(y) g_S(y) dy = \mathbb{E} \{ \rho^c(S) \}.$$

Theorem 1 provides a theoretical expression both for  $\rho^c$  and  $\rho$ , when the price returns are assumed to be increments of an fBm.

**Theorem 1.** Let  $X_t$  be an fBm of Hurst exponent  $H \in (0, 1)$  and volatility parameter  $\sigma > 0$ . Let  $h > 0$  and  $\hat{R}_{t,t+h|\Delta}$  be a predictor of  $R_{t,t+h}$  based on the vector  $S$  of past returns, as defined by equation (4). The conditional hit ratio is

$$\rho^c(y) = N \left( \frac{|\Sigma_{RS} \Sigma_S^{-1} y|}{\sqrt{\sigma^2 h^{2H} - \Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T}} \right)$$

and the non-conditional hit ratio is, for  $H \neq 1/2$ ,

$$\rho = 1 - \frac{1}{\pi} \arctan \left( \sqrt{\frac{\sigma^2 h^{2H}}{\Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T} - 1} \right), \quad (6)$$

where  $\Sigma_S$  and  $\Sigma_{RS}$  are the same as in Proposition 4 and  $N$  is the standard Gaussian cumulative distribution function.

The proof of Theorem 1 is postponed in Appendix B.

The case  $H = 1/2$ , for which the fBm is a standard Brownian motion, is of particular interest. In this case, the  $1 \times 1$  matrix  $\Sigma_R$  contains the element  $\sigma^2 h$ , the  $1 \times n$  matrix  $\Sigma_{RS}$  only contains zeros, and the  $n \times n$  matrix  $\Sigma_S$  is diagonal equal to  $\text{diag}(\sigma^2(\delta_1 - \delta_0), \dots, \sigma^2(\delta_n - \delta_{n-1}))$ . As a consequence, according to Theorem 1, the conditional hit ratio  $\rho^c(y)$  is 0.5, whatever  $y$ . This is consistent with the fact that the standard Brownian motion is a martingale. By the way, any value of  $H$  different from  $1/2$  trivially leads to  $\rho^c > 0.5$  since it is the image by  $N$  of a positive number. Regarding the non-conditional hit ratio  $\rho$ , equation 6 is not defined for  $H = 1/2$ , but we can easily determine that the limit of  $\rho$ , when  $H \rightarrow 1/2$ , is also 0.5. Once, again, any value of  $H$  different from  $1/2$  leads to  $\rho > 0.5$ , that is to more frequent good predictions than bad predictions, as one can see in Figure 1.




Figure 1: Theoretical non-conditional hit ratio for the fBm with one time lag and  $\delta_1 = h$  (continuous line) and for four time lags and  $(\delta_0, \dots, \delta_4) = (0, 0.5, 1, 2, 3)$  (dotted line).

In the rest of the paper, the theoretical hit ratio on which we focus is  $\rho$ , the non-conditional hit ratio. Indeed, in a practical application, we are inclined to choose a specific forecast setting, that is a number of lags and a duration for these lags. The maximization of a theoretical hit ratio ex ante must lead the choice of this setting. This choice is not that simple if one wants to maximize the conditional hit ratio, insofar as one must numerically test the best combination of lagged observations to finally retain the one maximizing  $\rho^c$ . A legitimate question then arises: Why not retaining all the tested lagged observations? If one uses  $\rho$  instead of  $\rho^c$ , the choice of time lags is simpler and will always be the same, provided that the estimated  $H$  is also the same, because the choice of time lags in this case does not depend on the particular value of each past observation but only on the estimated  $H$ .

We see in Figure 1 that Hurst exponents closer to 0 or 1 than to  $1/2$  lead to higher hit ratios. Hurst exponents  $H$  above  $1/2$  also lead to more accurate predictions than when the Hurst exponent is  $1 - H$ . In particular,  $H$  close to 1, that is a very persistent series of returns, leads to the highest possible hit ratio, whereas  $H$  close to 0 only leads to a hit ratio of 0.67 if the forecast is based on one past return ( $n = 1$ ). In other words, it is globally more difficult to make good predictions when  $H < 1/2$  than when  $H > 1/2$  but predictions are however good in average, with a hit ratio



higher than 0.5 in both cases. This contradicts a part of the econophysics literature which tends to consider that predictions are relevant only for  $H > 1/2$ . We can even stress that predictions will be more accurate for  $H = 0.2$  than for  $H = 0.6$ .

In the particular case where  $n = 1$ , for which we have already provided an explicit expression for  $\beta_1$  in equation 5, the hit ratio also has a simpler expression:

$$\rho = 1 - \frac{1}{\pi} \arctan \left( \sqrt{\frac{1}{\beta_1^2} \left( \frac{h}{\delta_1} \right)^{2H} - 1} \right)$$

because  $\Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T$ , in this case, is equal to  $\beta_1 \Sigma_{RS}^T = \beta_1^2 \Sigma_S = \sigma^2 \beta_1^2 \delta_1^{2H}$ .

In a systematic investment perspective, a forecast based on only one past return is often not enough to generate a profitable strategy because of too low hit ratios, in particular because the Hurst exponent in finance is almost never in the interval  $(0.75, 1)$ , where we observe the highest theoretical hit ratios. In order to build a systematic strategy, the prediction is to be improved by including several past returns. Even without optimizing the duration of these past returns, one sees in Figure 1 that adding arbitrary time lags in the set  $\Delta$  indeed increases the hit ratio. This effect is very limited when the Hurst exponent is higher than  $1/2$ , but it is promising for values of  $H$  below  $1/2$ .

In this investment perspective, the question of the best choice for  $n$ , the number of past returns to be included in the predictor, is an important topic. Theoretically, the optimal  $n$  maximizing the theoretical hit ratio is  $+\infty$ . However, considering a predictor based on a big number of past returns leads in general to overfitting. To avoid this drawback, one prefers minimizing information criteria, such as AIC or BIC, instead of simply maximizing an ex-ante hit ratio. Such criteria are based on the likelihood of the model and penalize the number of parameters. Regarding the likelihood, reality is always more complex than any model, so adding time lags will first improve the likelihood but it may also deteriorate it for a too big number of time lags due to a discrepancy between reality and the fBm specification. Regarding the number of parameters, it does not depend on the number of time lags if the model is an fBm. Indeed, the weights  $\beta_i$  applied to past returns all depend on the parameters of the fBm.

The choice of the  $n$  minimizing AIC or BIC thus depends on the dataset. We will see an illustration of this problem in the empirical part of the paper, in Section 5.1. Once  $n$  is selected, the duration of the  $n$  past price returns is to be selected too. We will see in the next subsection how this can be achieved using the theoretical hit ratio.

### 3.3 Optimal duration of time lags

The question of the optimal duration of time lags is to be related to the question of the optimal sampling of high-frequency observations of asset prices, with which the econometric literature has already dealt without any fBm assumption [6, 37]. Observations are equally spaced in time in this econometric approach. This will not be the case in our fBm-based predictor.

We consider  $n \geq 1$  past returns used in the predictor defined in Proposition 4. We write them  $R_{t-\delta_1, t-\delta_0}, \dots, R_{t-\delta_n, t-\delta_{n-1}}$ , with  $\forall i < j, \delta_i < \delta_j, \delta_i \geq 0$ , and  $\delta_0 = 0$ . Given a prediction horizon  $h > 0$ , the vector of time lags  $[\delta_0 \dots \delta_n]^T$  maximizing the hit ratio  $\rho$  of Theorem 1 is noted  $[\delta_0^* \dots \delta_n^*]^T$ . A numerical optimization of  $\rho$  provides these optimal time lags. This numerical study leads to the following observation:

$$\forall i \in \llbracket 1, n \rrbracket, \delta_i^* = \frac{h^2}{\delta_{n+1-i}^*}.$$

We briefly focus on the case  $n = 1$ . Figure 2 displays the value of the hit ratio as a function of  $\delta_1/h$ . The optimum in  $\delta_1/h = 1$  is clear and it also appears that the hit ratio is the same for  $\delta_1/h$  and  $h/\delta_1$ . Therefore, the optimal time lag for the past return used in the forecast is equal to the forecast horizon:  $\delta_1 = h$ . In other words, to forecast the future daily return, one has to consider the last observed daily return. For a monthly return to be forecast, the last observed monthly return is the optimal input of the fBm-based prediction model. This is consistent with the existing literature [36].




Figure 2: Hit ratio with  $H = 0.65$  (black) and  $H = 0.15$  (grey) for the fBm, for various values of the ratio  $\delta_1/h$  and  $n = 1$ .

Other cases than  $n = 1$  provide new insights. Increasing  $n$  of one unit does indeed not mean adding one time lag to the set of time lags optimized in the case  $n - 1$ . For example, if  $n = 2$ ,  $h$  is not one of the optimal time lags. Instead, if we consider  $h = 1$  for simplifying and  $H = 0.65$ , the optimal lags are  $(\delta_1, \delta_2) = (0.29, 3.45)$ . For a monthly forecast, so for a forecast horizon of 22 days, the optimal time lags to be considered are thus roughly 6 and 76 days. For a daily forecast, intraday data are to be taken into account. If  $n$  is odd (respectively even), one has to consider  $(n - 1)/2$  (resp.  $n/2$ ) time lags lower than  $h$ , the same quantity bigger than  $h$ , and one (resp. zero) time lag equal to  $h$ . We gather in Tables 1 and 2 the optimal lags for two values of  $H$  and various  $n$ . We provide them for  $h = 1$ . Multiplying them by  $h$  gives the optimal time lags for a forecast horizon equal to  $h$ .

$n$	$\delta_1^*$	$\delta_2^*$	$\delta_3^*$	$\delta_4^*$	$\delta_5^*$	$\delta_6^*$	$\rho$
1	1.000						57.42%
2	0.289	3.454					58.14%
3	0.127	1.000	7.896				58.72%
4	0.067	0.458	2.185	14.979			58.90%
5	0.039	0.253	1.000	3.949	25.407		58.99%
6	0.025	0.156	0.562	1.780	6.411	39.919	59.05%

Table 1: For various numbers  $n$  of lagged returns, optimal time lags and corresponding theoretical hit ratio. The dynamic is an fBm of Hurst exponent 0.65 and the forecast horizon is  $h = 1$ .

$n$	$\delta_1^*$	$\delta_2^*$	$\delta_3^*$	$\delta_4^*$	$\delta_5^*$	$\delta_6^*$	$\rho$
1	1.000						62.56%
2	0.367	2.726					64.34%
3	0.193	1.000	5.168				65.72%
4	0.120	0.539	1.856	8.365			66.29%
5	0.081	0.341	1.000	2.933	12.347		66.66%
6	0.058	0.236	0.637	1.570	4.241	17.170	66.91%

Table 2: For various numbers  $n$  of lagged returns, optimal time lags and corresponding theoretical hit ratio. The dynamic is an fBm of Hurst exponent 0.15 and the forecast horizon is  $h = 1$ .

Another question is important for defining a systematic trading strategy: what is the best between defining each price return on the one hand between  $t - \delta_i$  and  $t - \delta_{i-1}$  and, on the other hand, between  $t - \delta_i$  and  $t$ ? It is worth noting that if we compare the two vectors  $[R_{t-\delta_1, t-\delta_0} \dots R_{t-\delta_n, t-\delta_{n-1}}]^T$  and  $[R_{t-\delta_1, t} \dots R_{t-\delta_n, t}]^T$ , the weight vector  $[\beta_1 \dots \beta_n]^T$  associated to each vector of price returns is different, but the accuracy of the corresponding predictor, defined as the non-conditional hit ratio, is the same. Consequently, the optimal lags are also the same between the two versions. Whatever the way one splits the time lags among the different returns, provided that there is no redundancy, the non-conditional quality of the forecast will be the same. All the information used by the predictor comes in fact from the list of time lags. One must thus concentrate one's efforts on the selection of  $\Delta$  instead of on its specific division as bounds of all the price returns.

## 4 Statistical arbitrage: to predict or not to predict

The theoretical results in Section 3, about hit ratios for predictors of an fBm, are encouraging. Indeed, many values of the Hurst exponent lead to an acceptable hit ratio in the perspective of building a trading strategy. However a more precise look at the predicted increments show some forecasts close to zero which may affect the performance of a trading strategy for two reasons. First, the smaller the predictions in absolute value, the higher the incertitude about the sign of the future return. The second reason is almost a tautology: the smaller the future price returns of an asset in absolute value, the lower the financial performance a trader can expect by following an investment strategy using this asset. Because of these two reasons, a trader may want to discard predictions which are close to zero. We are thus interested in understanding how the quality of the forecast evolves when one applies a thresholding to the predictor introduced in equation (4). Next, we want not to limit our analysis to hit ratios but to introduce more financial criteria, such as expected performance and risk of a simple trading strategy based on the thresholding concept. This strategy consists in buying the asset if the predicted return is significantly above zero, short selling the asset if it is significantly below, and keeping a neutral position if the prediction is too close to zero for the trader to be confident in the predicted sign of the future return.

### 4.1 Threshold and hit ratio

We still suppose that log-prices follow an fBm. Thanks to this assumption, we can build predictors following equation (4), in which we can input a number  $n$  of lagged observed returns. We want to determine the theoretical forecasting performance of such a predictor when one discards the predicted values which are close to zero. In other words, one defines a threshold  $\theta \geq 0$  and ignores

the predictions lower than  $\theta$  in absolute value. We thus introduce a thresholding function

$$f_\theta : x \in \mathbb{R} \mapsto x \mathbf{1}_{\{|x| \geq \theta\}},$$

so that the considered predictor is now  $f_\theta \left( \hat{R}_{t,t+h|\Delta} \right)$ .

The hit ratio we used in Section 3 is not appropriate in this new framework. Indeed, the hit ratio only supposes two states: on the one hand, the good prediction of the sign of the future return and, on the other hand, the bad prediction. We now use a ternary classification of predictions: good, bad, and zero prediction. The statistics depicting the forecast accuracy of the new predictor can be based on the probability of each of these three states.

**Definition 4.** Let  $\hat{R}_{t,t+h|\Delta}$  be a predictor of  $R_{t,t+h}$  based on the vector  $S$  of past returns, as defined by equation (4). Given  $\theta \geq 0$ , we transform  $\hat{R}_{t,t+h|\Delta}$  in a new predictor  $f_\theta \left( \hat{R}_{t,t+h|\Delta} \right)$ . The non-conditional probability of good sign forecast is

$$p^+(\theta) = \mathbb{P} \left[ f_\theta \left( \hat{R}_{t,t+h|\Delta} \right) R_{t,t+h} > 0 \right],$$

the non-conditional probability of bad sign forecast is

$$p^-(\theta) = \mathbb{P} \left[ f_\theta \left( \hat{R}_{t,t+h|\Delta} \right) R_{t,t+h} < 0 \right],$$

and the non-conditional probability of zero forecast is

$$p^0(\theta) = \mathbb{P} \left[ f_\theta \left( \hat{R}_{t,t+h|\Delta} \right) = 0 \right].$$

We stress the fact that we are only interested here in non-conditional probabilities. Taking into account conditional probabilities may lead to better predictors, in which more appropriate weights are used for lagged returns, and thus to a better performance. Nevertheless, the conditional approach is very specific to the data observed and is a useless refinement in our perspective. The purpose of this paper is indeed not the fine tuning of a trading strategy but, instead, it is about setting general results concerning the relevance of fBm-based predictions in finance.

The formula for the probabilities introduced in Definition 4 are expressed in Theorem 2, in which the fBm assumption for log-prices is used.

**Theorem 2.** Let  $X_t$  be an fBm of Hurst exponent  $H \in (0, 1/2) \cup (1/2, 1)$  and  $\theta \geq 0$ . Let  $h > 0$  and  $\hat{R}_{t,t+h|\Delta}$  be a predictor of  $R_{t,t+h}$  based on the vector  $S$  of past returns, as defined by equation (4). The non-conditional probabilities of sign forecast, as introduced in Definition 4, are provided by a Taylor expansion, when  $\theta$  is in the neighbourhood of 0, for  $p^+(\theta)$ :

$$p^+(\theta) = 1 - N \left( \frac{\theta}{a} \right) + \frac{1}{\pi} \arctan \left( \frac{a}{b} \right) - \frac{\theta^2}{2\pi ab} + \left( \frac{1}{ab^3} + \frac{3}{a^4 b} \right) \frac{\theta^4}{24\pi} + \mathcal{O}(\theta^6)$$

and for  $p^-(\theta)$ :

$$p^-(\theta) = 1 - N \left( \frac{\theta}{a} \right) - \frac{1}{\pi} \arctan \left( \frac{a}{b} \right) + \frac{\theta^2}{2\pi ab} - \left( \frac{1}{ab^3} + \frac{3}{a^4 b} \right) \frac{\theta^4}{24\pi} + \mathcal{O}(\theta^6),$$

as well as by the following exact formula for  $p^0(\theta)$ :

$$p^0(\theta) = -1 + 2N \left( \frac{\theta}{a} \right),$$

where  $a = \sqrt{\Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T}$ ,  $b = \sqrt{\sigma^2 h^{2H} - \Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T}$ , and where  $\Sigma_S$  and  $\Sigma_{RS}$  are the same as in Proposition 4.

The proof of Theorem 2 is postponed in Appendix C.

Theorem 2 provides only Taylor expansions of  $p^+(\theta)$  and  $p^-(\theta)$ , whereas we have an exact formula for  $p^0(\theta)$ . However, beyond the expansions, we can describe how  $p^+(\theta)$  and  $p^-(\theta)$  evolve with  $\theta$ , whatever  $H$ . Indeed, both  $p^+$  and  $p^-$  are trivially decreasing functions. Moreover,  $p^+$  decreases more rapidly than  $p^-$ , as stated in Proposition 5.

**Proposition 5.** *With the assumptions of Theorem 2, for  $\theta \geq 0$  we have:*

$$\frac{dp^+}{d\theta}(\theta) \leq \frac{dp^-}{d\theta}(\theta) \leq 0.$$

*Proof.* We remark from equations (12) and (13), in the proof of Theorem 2, that the derivative of  $p^+$  and  $p^-$  is:

$$\begin{cases} \frac{d}{d\theta}p^+(\theta) &= -\frac{2}{a}N\left(\frac{\theta}{b}\right)g\left(\frac{\theta}{a}\right) \\ \frac{d}{d\theta}p^-(\theta) &= -\frac{2}{a}\left(1 - N\left(\frac{\theta}{b}\right)\right)g\left(\frac{\theta}{a}\right). \end{cases}$$

The two derivatives are negative and the ratio of the second one over the first one is  $N(\theta/b)^{-1} - 1$ , which is in the interval  $[0, 1]$  for  $\theta \geq 0$ , meaning that the derivative of  $p^+$  with respect to  $\theta \geq 0$  is greater in absolute value than the one of  $p^-$ .  $\square$

In Theorem 2, when the threshold is  $\theta = 0$ , the zero forecast probability is  $p^0(0) = 0$ . This particular case corresponds to the binary classification of predictions of Section 3. The non-conditional hit ratio  $\rho$  introduced in Definition 3 corresponds to  $p^+(0)$ , whose formula in Theorem 2 matches the formula of  $\rho$  in Theorem 1.

The greater the threshold  $\theta$ , the greater the value of  $p^0(\theta)$ , whose limit, when  $\theta \rightarrow +\infty$ , is 1. As a consequence, if  $\theta \neq 0$ , we do not have a binary classification between good and bad forecast but a ternary one, the third state being the zero prediction because of the thresholding. The hit ratio is thus inappropriate. One can think of several alternatives exploiting  $p^+(\theta)$ ,  $p^-(\theta)$ , and  $p^0(\theta)$ .

The first natural alternative to the hit ratio consists in considering the proportion of good predictions with respect to the non-zero predictions, that is  $p^+(\theta)/(p^+(\theta) + p^-(\theta))$ . This ratio converges asymptotically towards its maximal value, 1. Indeed, from equations (12) and (13) in the proof of Theorem 2, we have  $p^+(\theta) + p^-(\theta) = 2N(-\theta/a)$  and

$$\begin{aligned} \frac{p^+(\theta)}{p^+(\theta) + p^-(\theta)} &= \frac{1}{2} + \frac{1}{2N(-\theta/a)} \int_{\theta/a}^{+\infty} \left( N\left(\frac{a}{b}u\right) - N\left(-\frac{a}{b}u\right) \right) g(u) du \\ &\stackrel{\theta \rightarrow +\infty}{\sim} \frac{1}{2} + \frac{1}{2N(-\theta/a)} \int_{\theta/a}^{+\infty} g(u) du \\ &= \frac{1}{2} + \frac{1}{2N(-\theta/a)} (1 - N(\theta/a)) \\ &= 1. \end{aligned}$$

In a financial perspective, this accuracy metric is not suitable, because maximizing it leads to selecting  $\theta = +\infty$  and thus to discarding all the predictions.

In order to penalize high values of the zero forecast probability, one can consider the alternative ratio  $p^+(\theta)/(p^+(\theta) + p^-(\theta) + p^0(\theta)) = p^+(\theta)$ , but it is monotonically decreasing in  $\theta$ , as stated by Proposition 5, so that the optimal threshold would be  $\theta = 0$ , whatever  $H$ . This solution thus does not address the problem introduced in this section about discarding the less certain predictions.

Alternatively,  $p^+(\theta) - p^-(\theta)$  is a signed forecast accuracy metric. This metric corresponds to the average number of coins earned in a game in which one earns one coin for a good forecast, loses one for a bad one, and there is no coin transaction if one restrains from playing this forecasting game. According to Proposition 5, this ratio is monotonically decreasing in  $\theta$ , like the previous ratio.

All the natural extensions of the hit ratio we can think of are inappropriate insofar as they lead to an optimal threshold  $\theta$  either equal to 0 or to infinity. We propose in the next subsection an alternative accuracy metric more consistent with a financial application. Indeed, all the extended hit ratios proposed above only rely on probabilities of making a good, a bad, or a small forecast. In particular, they do not consider the amplitude of the predicted price returns. In an investment perspective, forecasting properly large price returns is though more fruitful than forecasting properly small ones.

## 4.2 Optimal threshold and risk-adjusted performance

We now consider the financial performance corresponding to a prediction. A trader is more eager to forecast accurately large price variations than small ones. Therefore, the trader's objective is not minimizing an MSE or maximizing a hit ratio, but maximizing instead a risk-adjusted performance of an investment based on these predictions. A proper evaluation of the predictor must indeed weight the 0-1 contribution of each prediction, which appears in the hit ratio, with the amplitude of the prediction.

The trading strategy we study here is the ternary strategy exposed above: buying a fixed quantity of the asset if one predicts a positive price return, short selling the same quantity if one predicts a negative return, not investing if the expected return of the asset is below the threshold  $\theta$  in absolute value. The trading frequency of the strategy is equal to  $h > 0$ , the forecast horizon. Using the same notations as above and given an amount of money hold by the trader and totally invested in this strategy, the return at horizon  $h$  will thus be:

$$R_{t,t+h}^{\text{strat}}(\theta) = \begin{cases} 0 & \text{if } f_{\theta} \left( \hat{R}_{t,t+h|\Delta} \right) = 0 \\ R_{t,t+h} & \text{if } f_{\theta} \left( \hat{R}_{t,t+h|\Delta} \right) > 0 \\ -R_{t,t+h} & \text{if } f_{\theta} \left( \hat{R}_{t,t+h|\Delta} \right) < 0. \end{cases}$$

This return represents the performance to be maximized. At time  $t$ , one only knows an expectation of this return:

$$\tilde{R}_{t,t+h}(\theta) = \mathbb{E}[R_{t,t+h}^{\text{strat}}(\theta)]. \quad (7)$$

Regarding the risk statistic of this investment, we consider a lower absolute semi-deviation risk measure. It is defined as the average absolute deviation below zero of the return of the strategy, in other words it is the average loss:

$$\tilde{\sigma}_{t,t+h}^{-}(\theta) = -\mathbb{E} \left[ \min(0, R_{t,t+h}^{\text{strat}}(\theta)) \right]. \quad (8)$$

Such a risk measure is less widespread in the asset management industry than the volatility, but it is more relevant since it does not incorporate positive deviations which, by definition, are not risky. This risk measure appears for example at the denominator of the kappa ratio of order one, which is a general risk-performance ratio including the popular Sharpe ratio as a particular case [40].

Theorem 3 provides a formula for  $\tilde{R}_{t,t+h}(\theta)$  and for  $\tilde{\sigma}_{t,t+h}^{-}(\theta)$  when the log-prices follow an fBm.

**Theorem 3.** *Let  $X_t$  be an fBm of parameters  $H \in (0, 1/2) \cup (1/2, 1)$  and  $\sigma > 0$ . Let  $h > 0$  and  $\hat{R}_{t,t+h|\Delta}$  be a predictor of  $R_{t,t+h}$  based on the vector  $S$  of past returns, as defined by equation (4). Given  $\theta \geq 0$ , the average return of the ternary strategy,  $\tilde{R}_{t,t+h}(\theta)$ , and the corresponding lower absolute semi-deviation risk measure,  $\tilde{\sigma}_{t,t+h}^{-}(\theta)$ , defined in equations (7) and (8), are equal to:*

$$\tilde{R}_{t,t+h}(\theta) = 2ag \left( \frac{\theta}{a} \right)$$

and

$$\tilde{\sigma}_{t,t+h}^-(\theta) = -2aN \left( -\frac{\theta}{b} \right) g \left( \frac{\theta}{a} \right) + \sqrt{\frac{2}{\pi}} \sigma h^H N \left( -\theta \sqrt{\frac{1}{a^2} + \frac{1}{b^2}} \right),$$

where  $a = \sqrt{\Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T}$ ,  $b = \sqrt{\sigma^2 h^{2H} - \Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T}$ , and  $g$  is the Gaussian probability density function.

The proof of Theorem 3 is postponed in Appendix D.

In the case where  $\theta = 0$ , that is when the trader does not discard any forecast, the performance and the risk of the strategy are simplified:

$$\begin{cases} \tilde{R}_{t,t+h}(0) &= a \sqrt{\frac{2}{\pi}} \\ \tilde{\sigma}_{t,t+h}^-(0) &= \frac{\sigma h^H - a}{\sqrt{2\pi}}, \end{cases}$$

in which we recognize, for  $\tilde{R}_{t,t+h}(0)$ , the expected value of  $a|U|$ , where  $U \sim \mathcal{N}(0, 1)$ . Both  $\tilde{R}_{t,t+h}(\theta)$  and  $\tilde{\sigma}_{t,t+h}^-(\theta)$  are monotonically decreasing in  $\theta$ : discarding predictions tends to diminish both the expected gain and the risk. At the limit, when  $\theta \rightarrow +\infty$ , both  $\tilde{R}_{t,t+h}(\theta)$  and  $\tilde{\sigma}_{t,t+h}^-(\theta)$  tend to zero: the threshold is so high that no investment decision is made by the trader.

In order to find a good balance between performance and risk, we define an expected risk-adjusted performance, which simply consists in the expected return of the strategy, penalized by the corresponding risk:

$$\tilde{R}_\lambda(\theta) = \tilde{R}_{t,t+h}(\theta) - \lambda \tilde{\sigma}_{t,t+h}^-(\theta). \quad (9)$$

In equation (9), we omit the subscript in  $t$  and  $t+h$  for the risk-adjusted performance  $\tilde{R}_\lambda(\theta)$ , in order not to overload the notations. But one has to keep in mind that  $\tilde{R}_\lambda(\theta)$  depends on  $h$ . The parameter  $\lambda$  in equation (9) plays the role of a risk aversion when its value is positive. A negative value of  $\lambda$  would indicate a risk-seeking behaviour. In particular, if  $\lambda = -1$ ,  $\tilde{R}_\lambda(\theta)$  is simply the upper absolute semi-deviation,  $\mathbb{E}[\max(0, R_{t,t+h}^{\text{strat}}(\theta))]$ . In what follows, we focus on the case  $\lambda \geq 0$ . Because of the particular risk measure we have chosen, the risk-adjusted performance corresponds to the expected performance of the strategy in a distorted probability in which bad outcomes, that is to say negative returns, are overweighted. Indeed,

$$\tilde{R}_\lambda(\theta) = \int_0^1 Q_{R_{t,t+h}^{\text{strat}}(\theta)}(p) d\mu(p),$$

where  $Q_Z(p)$  is the quantile of probability  $p$  of a variable  $Z$  and  $\mu$  is the probability distortion function defined by

$$\mu : x \in [0, 1] \mapsto \begin{cases} \frac{(1+\lambda)x}{1+\lambda p^-(\theta)} & \text{if } x < p^-(\theta) \\ \frac{x+\lambda p^-(\theta)}{1+\lambda p^-(\theta)} & \text{else.} \end{cases}$$

In particular, if  $\lambda = 1$ , the probability distortion doubles the probability of negative returns, relatively to the probability of positive returns.

Figure 3 displays the risk-adjusted performance for various values of risk aversion  $\lambda$  and threshold  $\theta$ . Since both the expected return and the risk measure tend toward zero when  $\theta \rightarrow +\infty$ ,  $\tilde{R}_\lambda(\theta)$  also tends toward zero for high threshold values. When  $\lambda = 0$ , the maximum risk-adjusted performance is reached for  $\theta = 0$ . For  $\lambda > 0$ , we observe a global maximum for a value  $\theta > 0$ . This threshold is important insofar as it constitutes an optimal threshold for a trader whose aim would be the maximization of the risk-adjusted performance:

$$\theta_\lambda^* = \underset{\theta \geq 0}{\operatorname{argmax}} \tilde{R}_\lambda(\theta).$$




Figure 3: Theoretical penalized expected return  $\tilde{R}_\lambda(\theta)$  as a function of the threshold  $\theta$ , for various values of risk aversion  $\lambda$ , namely, from fatter to thinner, 0, 0.25, 0.5, 0.75, 1. The fBm has the parameters  $H = 0.65$  and  $\sigma = 1$ . The forecast horizon is  $h = 1$  and the number of lagged returns in the predictor is  $n = 1$ .

Figure 4 illustrates the link between  $\lambda$  and  $\theta_\lambda^*$ . The higher the risk aversion  $\lambda$ , the higher the optimal threshold  $\theta_\lambda^*$ . By the way, a higher threshold means that the decision not to invest is more frequent. Figure 4 thus shows  $p^0(\theta_\lambda^*)$  as an increasing function of  $\lambda$ .

We also observe that the optimal threshold strongly depends on the value of the Hurst exponent, as displayed in Figure 5. In particular, it is asymmetric with respect to  $H = 1/2$ . Whatever the risk aversion, the highest optimal thresholds are reached for Hurst exponents close to  $1/2$ , but the optimal threshold decreases toward 0 when  $H$  is close to 1, whereas a threshold significantly different from 0 is indicated when  $H$  is close to 0. This shows a relatively higher difficulty to make financially performing predictions when  $H < 1/2$  than when  $H > 1/2$ . Figure 5 also shows  $p^0(\theta_\lambda^*)$  as a function of the Hurst exponent. This zero forecast probability reaches huge values when  $H$  is close to  $1/2$ . For instance, for the risk aversion considered,  $\lambda = 0.1$ ,  $p^0(\theta_\lambda^*)$  is twice as large for  $H = 0.55$  (40%) than for  $H = 0.6$  (20%).

Guasoni and co-authors provide a risk-adjusted performance which differs from our approach for several reasons: their risk is a variance, whereas ours is a lower absolute semi-deviation, and they consider trading in continuous time, whereas we work in discrete time [31, 30]. They notice an asymmetry of their risk-adjusted performance with respect to  $H = 1/2$ . While the shape of our metric is obviously different from theirs, we obtain qualitatively the same asymmetry result, as illustrated by Figure 6: the worst risk-adjusted performance is for  $H = 1/2$  and the improvement of this metric is stronger for  $H > 1/2$  than for  $H < 1/2$ . In the case where  $n = 1$  and with an optimal lag  $\delta_1 = h$ , we obtain, thanks to equation (5),

$$\begin{aligned}
 a &= \sqrt{\Sigma_{RS} \Sigma_S^{-1} \Sigma_{RS}^T} \\
 &= \sqrt{\beta_1^2 \sigma^2 \delta_1^{2H}} \\
 &= |\beta_1 \sigma \delta_1^H| \\
 &= \sigma h^H |2^{2H-1} - 1|,
 \end{aligned}$$






Figure 4: Theoretical optimal threshold  $\theta_{\lambda}^*$  (left) and corresponding non-conditional zero forecast probability  $p^0(\theta_{\lambda}^*)$  (right) as functions of the risk aversion  $\lambda$ . The fBm has the parameters  $\sigma = 1$  and, for  $H$ , respectively from the top curve to the bottom curve, 0.6, 0.7, and 0.8. The forecast horizon is  $h = 1$  and the number of lagged returns in the predictor is  $n = 1$ .




Figure 5: Theoretical optimal threshold  $\theta_{\lambda}^*$  (left) and corresponding non-conditional zero forecast probability  $p^0(\theta_{\lambda}^*)$  (right) as functions of the Hurst exponent  $H$ . The fBm has a volatility parameter  $\sigma = 1$ . The risk aversion is  $\lambda = 0.1$ , the forecast horizon is  $h = 1$ , and the number of lagged returns in the predictor is  $n = 1$ .

and finally the risk-adjusted performance without thresholding small predictions is equal to

$$\begin{aligned}\tilde{R}_\lambda(0) &= a \left( \sqrt{\frac{2}{\pi}} + \frac{\lambda}{\sqrt{2\pi}} \right) - \frac{\lambda \sigma h^H}{\sqrt{2\pi}} \\ &= \frac{\sigma h^H}{\sqrt{2\pi}} [ |2^{2H-1} - 1| (2 + \lambda) - \lambda ].\end{aligned}$$

We see in Figure 6 that introducing the optimal thresholding tends to set a floor at 0 for the risk-adjusted performance. As a consequence, when  $H$  is close to  $1/2$ ,  $\tilde{R}_\lambda(\theta_\lambda^*)$  is close to the performance of the strategy without risk adjustment or thresholding,  $\tilde{R}_0(0)$ . In this case, the thresholding thus tends to erase the risk. However, when  $H$  is close to 0 or 1,  $\tilde{R}_\lambda(\theta_\lambda^*)$  is close to  $\tilde{R}_\lambda(0)$  and the effect of the thresholding is not significant.




Figure 6: Theoretical penalized expected return  $\tilde{R}_0(0)$  (black line),  $\tilde{R}_\lambda(0)$  (continuous grey line), and  $\tilde{R}_\lambda(\theta_\lambda^*)$  (dotted grey line), as a function of the Hurst exponent, for  $\lambda = 0.1$  (left) and  $\lambda = 0.5$  (right). The fBm has a volatility parameter  $\sigma = 1$ , the forecast horizon is  $h = 1$  and the number of lagged returns in the predictor is  $n = 1$ .

## 5 Empirical application

We now apply the theoretical results exposed above to two kinds of financial time series: volatility of stock indices and high-frequency prices of foreign exchange rates.

### 5.1 Rough volatility: fBm vs autoregressive model

The use of the fBm in finance is not limited to log-prices and one of the most popular applications of the fBm is for volatility modelling. Even though accuracy metrics directly related to the performance and risk of a trading strategy, as developed in Section 4.2, are not appropriate in this case, forecasting a volatility is also useful in algorithmic trading, and one can properly use the hit ratio to evaluate the quality of the fBm-based forecast. If the idea of depicting the volatility with the fBm is not new [16], the recent paradigm is rough volatility, that is the fBm used for modelling volatility has a very low Hurst exponent, close to 0.15, so that the series of volatility is strongly antipersistent [3, 28]. This model seems to currently dominate the research on stochastic volatility and leads to a rich literature, regarding for example some simplifications of the model to make it

computationally more efficient [1, 2], or some refinements to make it more realistic [26]. We also stress the fact that the rough framework relies on empirical observations of scaling rules of volatility increments and not on the analysis of the dependence between distant increments of volatility. This analysis of the long-range dependence seems to reveal a slowly decaying correlation [14], so that modelling volatility with a Hurst exponent higher than 1/2 may also be relevant in certain cases [27].

The data used in our analysis are daily realized volatilities computed with a five-minute discretization of prices, imported from the Oxford-Man Institute of Quantitative Finance Realized Library <sup>1</sup>. We focus on the realized volatility of eight stock indices: the AEX index, the CAC 40 index, the FTSE 100 index, the Nasdaq 100 index (IXIC), the Nikkei 225 index (N225), the Oslo Exchange All-share index (OSEAX), the Madrid General index (SMSI), and the S&P 500 index (SPX). The series starts on January 2000, except N225, which starts in February 2000, OSEAX in September 2001 and SMSI in July 2005. The end date of our sample is on the 12th April 2021.

We want to determine if the fBm-based predictor introduced in Proposition 4 is accurate for the time series of the logarithm of the realized volatilities, compared to a predictor based on an autoregressive (AR) model. We want to make a one-day forecast using daily increments of log-volatility, namely  $n$  lagged observations, for  $n$  varying between one up to six days. In particular, we do not try to optimize the time lags, contrary to the developments of Section 3.3, in order to make a fair comparison between fBm and AR models. For each value of  $n$ , the fBm-based forecast is the weighted sum of past increments, as in equation (4). The weights depend both on  $n$  and on  $H$ , the Hurst exponent. For this purpose, we need to estimate a Hurst exponent at each date. We use an estimator exploiting the scaling properties of the fBm in a rolling window of size  $T$ , without regularizing the obtained time-varying Hurst exponent [15, 8, 22]<sup>2</sup>. More precisely, if we note  $v_t$  the series of log-volatilities, supposed to follow an fBm of parameters  $H$  and  $\sigma$  in a specific sub-sample, any increment  $v_t - v_{t-\tau}$  of duration  $\tau > 0$  contained in the sub-sample has a variance  $\sigma^2 |\tau|^{2H}$ . Therefore, comparing the empirical log-variance of increments of two distinct durations,  $\tau_1$  and  $\tau_2$ , leads to a simple estimator of the Hurst exponent using data in the interval  $[t - T, t]$ :

$$\widehat{H}_t = \frac{1}{2 \log(\tau_1/\tau_2)} \log \left( \frac{(T - \tau_2) \sum_{i=0}^{T-\tau_1} (v_{t-i} - v_{t-i-\tau_1})^2}{(T - \tau_1) \sum_{i=0}^{T-\tau_2} (v_{t-i} - v_{t-i-\tau_2})^2} \right). \quad (10)$$

We display the estimated time-varying Hurst exponent of the realized log-volatility of SPX in Figure 7. We observe that, consistently with the literature on rough volatility,  $\widehat{H}_t$  is mainly in the interval  $[0.05, 0.25]$ .

Before presenting the accuracy of the forecast with the fBm approach, we introduce an alternative predictor, as a benchmark. The alternative approach is an AR model with Gaussian residuals and with the same number  $n$  of lagged daily increments of log-volatility as in the fBm predictor. The two models are estimated in the same rolling windows. The fBm and the AR models are very close since both are Gaussian and make linear forecast. The only difference is that the weights are not calculated in the same way. For the AR model, each weight is a parameter. When increasing  $n$ , we add new parameters, which will maximize the in-sample accuracy. For the fBm approach, the  $n$  weights all depend on the parameters  $H$  and  $\sigma$  of the fBm, so that increasing the number of lags will not increase the number of parameters. This last model is thus more parsimonious and less subject to overfitting than the AR approach. In other words, we expect the AR model to be more accurate in sample and the fBm to have better forecasting ability, at least when  $n$  is big enough. The accuracy metric we present below will confirm this intuition.

We determine the likelihood of each model in sample. Logically, this likelihood, calculated on

<sup>1</sup>Available at <http://realized.oxford-man.ox.ac.uk/data/download>.

<sup>2</sup>Precisely, in our empirical analysis,  $T = 504$  business days.



Figure 7: Estimated time-varying Hurst exponent of the realized log-volatility of SPX.

the estimation set, is higher for the AR model. But this is not a relevant indicator of forecasting ability. One often prefers using information criteria, such as AIC and BIC, which take into account, negatively, the log-likelihood and, positively, the number of parameters as well as, for the BIC, the number of data used for the estimation. Since the fBm model has less parameters than the AR model, it is less penalized and we can expect to have a lower AIC and BIC for it.

Empirical results confirm the intuition in the sense that the fBm predictor has always a lower AIC when  $n = 6$ , as one can see in Table 3<sup>3</sup>. But the results are even stronger. Indeed, when  $n = 1$ , that is when the variation of log-volatility is forecast only from the increment observed the last day, the fBm predictor has a slight advantage over the AR one, except for IXIC. We recall that the value of the information criterion cannot be interpreted in absolute value but that it is only useful for comparisons: for a same dataset, the best model is the one with the lowest criterion. For SPX, it seems from Figure 8, that the optimal number of lagged observations for the AR predictor is  $n = 4$ , at least regarding the BIC, whereas it is  $n \geq 6$  for the fBm predictor. Indeed, adding more lagged observations does not increase the number of parameters only in the fBm predictor.

index	fBm(1)	fBm(6)	AR(1)	AR(6)
AEX	875.9	<b>812.9</b>	877.7	823.6
CAC 40	795.0	<b>734.8</b>	796.7	742.1
FTSE 100	886.4	<b>826.8</b>	888.0	837.7
IXIC	920.1	<b>856.6</b>	916.8	861.2
N225	770.9	<b>716.8</b>	772.8	727.4
OSEAX	937.5	<b>838.2</b>	939.1	845.5
SMSI	1025.4	<b>931.1</b>	1027.4	940.6
SPX	933.6	<b>873.9</b>	935.4	880.9

Table 3: AIC for the eight volatility series, for four different predictors: the fBm and the AR model both with 1 and 6 lagged observations. In bold is the lowest AIC for each series.

Information criteria provide in sample a hint of which model will have the higher out-of-sample

<sup>3</sup>The results with BIC, not displayed in a table, are similar.

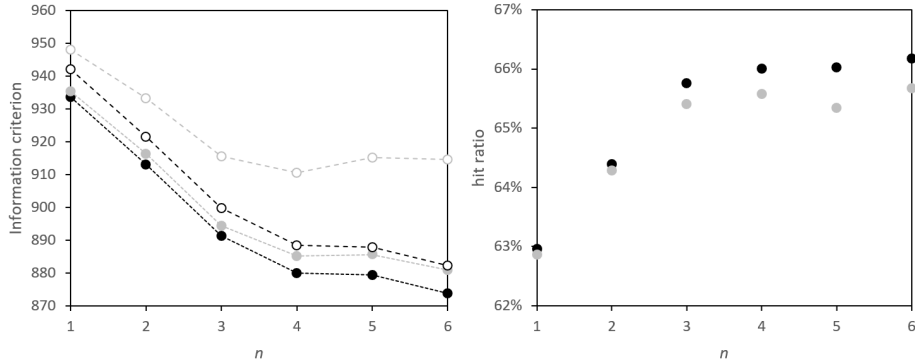


Figure 8: On the left: AIC (filled dots) and BIC (empty dots) for the fBm predictor (black) and the AR one (grey). On the right: out-of-sample hit ratio for the fBm predictor (black) and the AR one (grey). The number of lagged observations in the input of the two predictors is  $n$  and the forecast series is the log-volatility of SPX.

accuracy. In order to check this intuition, we determine the empirical out-of-sample hit ratio for each model. Results for  $n = 1$  and  $n = 6$  are gathered in Table 4. For  $n = 1$ , the hit ratio is always higher for the fBm predictor than for the AR one, except for AEX index. For  $n = 6$ , the fBm is always more accurate than any other model tested in the table. More details are displayed in Figure 8 for the SPX. It shows that the difference of accuracy between fBm and AR predictors becomes stronger from  $n = 3$ . For both models, adding lagged observation has a strong impact on the accuracy until  $n = 3$ . Our recommendation would thus be to use the fBm predictor with at least three lagged observations.

index	fBm(1)	fBm(6)	AR(1)	AR(6)
AEX	62.4%	<b>65.8%</b>	62.5%	65.7%
CAC 40	63.0%	<b>66.8%</b>	62.8%	66.0%
FTSE 100	64.8%	<b>67.5%</b>	64.3%	67.4%
IXIC	62.2%	<b>64.6%</b>	62.1%	64.3%
N225	62.9%	<b>67.4%</b>	62.8%	67.1%
OSEAX	64.0%	<b>67.9%</b>	63.8%	67.6%
SMSI	61.6%	<b>64.0%</b>	61.4%	63.7%
SPX	63.0%	<b>66.2%</b>	62.9%	65.7%

Table 4: Empirical out-of-sample hit ratio for the eight volatility series, for four different predictors: the fBm and the AR model both with 1 and 6 lagged observations. In bold is the highest hit ratio for each series.

This analysis of realized volatility series shows that the fBm provides a good predictor, which is more accurate than an AR model because of its parsimony. It confirms previous findings in the literature, where the superiority of the fBm over the AR model was assessed with the help of another accuracy indicator, namely the MSE [28]. It highlights the relevance of the fBm in finance in a forecasting perspective.

## 5.2 High-frequency FX rates

We now focus on time series of FX prices for three pairs: EURGBP, EURUSD, and GBPUSD. We consider high-frequency observations, with one price every minute. The FX market is continuously open during almost one week. We thus focus on one week in order to limit the effects of closing markets: from the 23rd June 2019 to the 28th June 2019.

We transform the price series in a series of log-prices and we assume it follows an fBm. We do not investigate here the relevance of the fBm for modelling FX log-prices, because many papers already deal with this question [11, 39, 22] and also propose some extensions of the fBm to match other stylized facts such as stationarity [23], high kurtosis [41, 25, 5], multiscaling [18], or even randomness of the Hurst exponent [9, 25]. We understand that several additional parameters may be relevant to perfectly depict the dynamic of FX log-prices, but, if our reference is the standard Brownian motion, the use of the fBm is a good step towards a realistic representation of this financial series.

Using a 12-hour window and equation (10), we estimate the Hurst exponent every minute for each series, as represented in Figure 9. Then, we forecast the evolution of the log-price at a one-hour horizon. We base this forecast on the formulas provided in Section 3. In particular, we take into account two lagged returns in each forecast. We propose either a naive choice of these lagged returns, or an *optimal* one. The naive choice consists in the one-hour and two-hour past returns. Alternatively, the *optimal* choice of lags is theoretically non-conditionally optimal<sup>4</sup> and depends on the Hurst exponent, as exposed in Section 3.3. For example, when  $H = 0.35$ , the lagged returns considered are the last 21-minute and 176-minute returns.




Figure 9: Estimated time-varying Hurst exponent of the log-prices of EURGBP (black), EURUSD (dark grey), GBPUSD (light grey).

As displayed in Table 5, the hit ratios for the EURUSD pair are very close between the naive and the *optimal* approaches. But for the two other pairs, the *optimal* lags outperform the naive lags, for which the hit ratio is close to 50%. It thus sounds that using carefully the lags in the predictor can transform an equiprobable coin toss in a performing strategy. It is worth noting that increasing

<sup>4</sup>This *optimal* threshold is not to be confused with an ex-post optimal threshold.

the number of lags, and not only choosing them properly, should also increase the hit ratio and make this predictor more appealing for practitioners.

FX rate	Method	Return	Risk	Risk-adj. return	$\widehat{p}^+$	$\widehat{p}^-$	$\widehat{p}^0$
EURGBP	Naive lags, $\theta = 0$	-0.51	2.47	-2.97	49.8%	50.2%	0.0%
	<i>Optimal</i> lags, $\theta = 0$	3.66	2.26	1.40	52.9%	47.1%	0.0%
	<i>Optimal</i> lags and $\theta$	3.09	1.50	1.59	30.2%	26.7%	43.1%
EURUSD	Naive lags, $\theta = 0$	2.82	1.85	0.97	54.6%	45.4%	0.0%
	<i>Optimal</i> lags, $\theta = 0$	1.52	1.91	-0.39	54.3%	45.7%	0.0%
	<i>Optimal</i> lags and $\theta$	1.37	1.20	0.17	34.4%	29.6%	36.0%
GBPUSD	Naive lags, $\theta = 0$	1.97	2.42	-0.45	50.6%	49.4%	0.0%
	<i>Optimal</i> lags, $\theta = 0$	4.22	2.31	1.91	52.4%	47.6%	0.0%
	<i>Optimal</i> lags and $\theta$	4.58	1.63	2.96	33.5%	28.8%	37.7%

Table 5: Empirical average one-hour return, risk (lower absolute semi-deviation), and risk-adjusted return (with  $\lambda = 0.1$ ) multiplied respectively by  $10^5$ ,  $10^4$ , and  $10^5$ . The three columns on the right are the corresponding empirical probabilities of good forecast, bad forecast, and zero forecast, following Definition 4. As soon as  $\theta = 0$ ,  $\widehat{p}^+$  is simply the hit ratio.

We then implement the ternary strategy described in Section 4. We consider a threshold  $\theta$  either equal to zero or to the theoretically optimal value determined numerically in Section 4.2 for a given risk aversion. Whatever the FX pair, we observe in Table 5 a sharp decrease of the risk when using the *optimal*  $\theta$ . For EURGBP and GBPUSD, it also leads to the highest risk-adjusted performance.

The EURUSD case is particular. We noticed the close hit ratios for the two kinds of lags, but the naive approach in fact leads to a better risk-adjusted performance. Three causes can explain this superiority of the naive approach for this FX pair. First, the *optimal* choice of lags is intended to maximize a theoretical non-conditional hit ratio, not a risk-adjusted performance. Second, the non-conditional framework used to define optimal lags works in average, but in some cases a conditional approach may be preferable. Third and last, our formulas rely on a model, the fBm, which is only a simplification of reality, so that a wise practical implementation of the proposed predictor should constantly check the relevance of the model and challenge it with other approaches in order to minimize model risk.

Nevertheless for the *optimal* lags, we observe in Table 5, that introducing the *optimal* threshold instead of  $\theta = 0$  improves the risk-adjusted performance of the strategy for the three pairs, including EURUSD. This was precisely the purpose of the thresholding. We also note that the *optimal* threshold leads not to take an investment decision roughly 40% of the time.

## 6 Conclusion

In this paper, we have provided theoretical formulas for accuracy metrics related to forecasts of the fBm in discrete time. We focused on metrics that we consider to be meaningful for a systematic trader, namely the hit ratio, the average performance and risk. We have used these theoretical expressions to derive optimal time lags and an optimal threshold under which the prediction is to be discarded. Finally, we have demonstrated empirically the practical relevance of all these considerations on two kinds of time series: realized volatilities and high-frequency FX rates. We think that systematic traders should find in this contribution some elements for improving their

strategies based on the fBm.

## References

- [1] E. Abi Jaber. Lifting the Heston model. *Quantitative finance*, 19(12):1995–2013, 2019.
- [2] E. Abi Jaber and O. El Euch. Multifactor approximation of rough volatility models. *SIAM journal on financial mathematics*, 10(2):309–349, 2019.
- [3] E. Alòs, J.A. León, and J. Vives. On the short-time behavior of the implied volatility for jump-diffusion models with stochastic volatility. *Finance and stochastics*, 11(4):571–589., 2007.
- [4] J. Alvarez-Ramirez, J. Alvarez, E. Rodriguez, and G. Fernandez-Anaya. Time-varying Hurst exponent for US stock markets. *Physica A: statistical mechanics and its applications*, 387(24):6159–6169, 2008.
- [5] A. Ammy-Driss and M. Garcin. Efficiency of the financial markets during the COVID-19 crisis: time-varying parameters of fractional stable dynamics. *arxiv preprint*, 2020.
- [6] F.M. Bandi and J.R. Russell. Microstructure noise, realized variance, and optimal sampling. *The review of economic studies*, 75(2):339–369, 2008.
- [7] C. Bender, T. Sottinen, and E. Valkeila. Arbitrage with fractional Brownian motion? *Theory of stochastic processes*, 13(1):23–34, 2007.
- [8] S. Bianchi and A. Pantanella. Pointwise regularity exponents and market cross-correlations. *International review of business research papers*, 6(2):39–51, 2010.
- [9] S. Bianchi, A. Pantanella, and A. Pianese. Modeling and simulation of currency exchange rates using multifractional process with random exponent. *International journal of modeling and optimization*, 2(3):309–314, 2012.
- [10] S. Bianchi and A. Pianese. Time-varying Hurst-Hölder exponents and the dynamics of (in)efficiency in stock markets. *Chaos, solitons & fractals*, 109:64–75, 2018.
- [11] M. Bohdalová and M. Greguš. Fractal analysis of forward exchange rates. *Acta polytechnica hungarica*, 7(4):57–69, 2010.
- [12] D.O. Cajueiro and B.M. Tabak. The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient. *Physica A: statistical mechanics and its applications*, 336(3-4):521–537, 2004.
- [13] P. Cheridito. Arbitrage in fractional Brownian motion models. *Finance and stochastics*, 7(4):533–553, 2003.
- [14] A. Chronopoulou and F.G. Viens. Estimation and pricing under long-memory stochastic volatility. *Annals of finance*, 8(2-3):379–403, 2012.
- [15] J.-F. Coeurjolly. Identification of multifractional Brownian motion. *Bernoulli*, 11(6):987–1008, 2005.
- [16] F. Comte and E. Renault. Long memory in continuous-time stochastic volatility models. *Mathematical finance*, 8(4):291–323, 1998.
- [17] R. Cont. Long range dependence in financial markets. In J. Lévy-Véhel and E. Lutton, editors, *Fractals in engineering*, pages 159–179. Springer, 2005.



- [18] T. Di Matteo. Multi-scaling in finance. *Quantitative finance*, 7(1):21–36, 2007.
- [19] T. Di Matteo, T. Aste, and M.M. Dacorogna. Scaling behaviors in differently developed markets. *Physica A: statistical mechanics and its applications*, 324(1-2):183–188, 2003.
- [20] T. Di Matteo, T. Aste, and M.M. Dacorogna. Long-term memories of developed and emerging markets: Using the scaling analysis to characterize their stage of development. *Journal of banking & finance*, 29(4):827–851, 2005.
- [21] C. Eom, S. Choi, G. Oh, and W.S. Jung. Hurst exponent and prediction based on weak-form efficient market hypothesis of stock markets. *Physica A: statistical mechanics and its applications*, 387(18):4630–4636, 2008.
- [22] M. Garcin. Estimation of time-dependent Hurst exponents with variational smoothing and application to forecasting foreign exchange rates. *Physica A: statistical mechanics and its applications*, 483:462–479, 2017.
- [23] M. Garcin. Hurst exponents and delampertized fractional Brownian motions. *International journal of theoretical and applied finance*, 22(5):1950024, 2019.
- [24] M. Garcin. A comparison of maximum likelihood and absolute moments for the estimation of Hurst exponents in a stationary framework. *arxiv preprint*, 2020.
- [25] M. Garcin. Fractal analysis of the multifractality of foreign exchange rates. *Mathematical methods in economics and finance*, 13-14(1):49–73, 2020.
- [26] M. Garcin and M. Grasselli. Long vs short time scales: the rough dilemma and beyond. *arxiv preprint*, 2020.
- [27] J. Garnier and K. Sølna. Option pricing under fast-varying long-memory stochastic volatility. *Mathematical finance*, 29(1):39–83, 2019.
- [28] J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quantitative finance*, 18(6):933–949, 2018.
- [29] M.S. Granero, J.T. Segovia, and J.G. Pérez. Some comments on Hurst exponent and the long memory processes on capital markets. *Physica A: statistical mechanics and its applications*, 387(22):5543–5551, 2008.
- [30] P. Guasoni, Y. Mishura, and M. Rásonyi. High-frequency trading with fractional Brownian motion. *Finance and stochastics*, 25:277–310, 2021.
- [31] P. Guasoni, Z. Nika, and M. Rásonyi. Trading fractional Brownian motion. *SIAM journal on financial mathematics*, 10(3):769–789, 2019.
- [32] L. Kristoufek. On Bitcoin markets (in)efficiency and its evolution. *Physica A: statistical mechanics and its applications*, 503:257–262, 2018.
- [33] L. Kristoufek and M. Vosvrda. Measuring capital market efficiency: Global and local correlations structure. *Physica A: statistical mechanics and its applications*, 392(1):184–193, 2013.
- [34] L. Kristoufek and M. Vosvrda. Gold, currencies and market efficiency. *Physica A: statistical mechanics and its applications*, 449:27–34, 2016.
- [35] B. Mandelbrot and J. van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM review*, 10(4):422–437, 1968.

- [36] C.J. Nuzman and H.V. Poor. Linear estimation of self-similar processes via Lamperti's transformation. *Journal of applied probability*, 37(2):429–452, 2000.
- [37] R.C.A. Oomen. Properties of realized variance under alternative sampling schemes. *Journal of business & economic statistics*, 24(2):219–237, 2006.
- [38] L.C.G. Rogers. Arbitrage with fractional Brownian motion. *Mathematical finance*, 7(1):95–105, 1997.
- [39] D. Surgailis, G. Teyssière, and M. Vaičiulis. The increment ratio statistic. *Journal of multivariate analysis*, 99(3):510–541, 2008.
- [40] W. Van Harlow. Asset allocation in a downside-risk framework. *Financial analysts journal*, 47(5):28–40, 1991.
- [41] A. Weron, K. Burnecki, S. Mercik, and K. Weron. Complete description of all self-similar models driven by Lévy stable noise. *Physical review E*, 71(1):016113, 2005.

## A Lemmas used for the proof of the theorems

Let  $g$  and  $N$  be respectively the probability density function and the cumulative distribution function of a standard Gaussian variable. We introduce the following lemmas.

**Lemma 1.** For all  $\alpha \in \mathbb{R}$ , we have  $\int_0^\infty N(\alpha x)g(x)dx = \frac{1}{4} + \frac{1}{2\pi} \arctan(\alpha)$ .

*Proof.* Let  $f(\alpha) = \int_0^\infty N(\alpha x)g(x)dx$ . Then  $f$  is differentiable and

$$\begin{aligned}
 f'(\alpha) &= \int_0^\infty xg(\alpha x)g(x)dx \\
 &= \int_0^\infty \frac{x}{2\pi} \exp\left(-\frac{x^2}{2}(1+\alpha^2)\right) dx \\
 &= \left[-\frac{1}{2\pi(1+\alpha^2)} \exp\left(-\frac{x^2}{2}(1+\alpha^2)\right)\right]_{x=0}^\infty \\
 &= \frac{1}{2\pi(1+\alpha^2)}.
 \end{aligned}$$

As a consequence,  $f$  can be written in the form  $f(\alpha) = \gamma + \frac{1}{2\pi} \arctan(\alpha)$ , with  $\gamma$  a constant. The constant  $\gamma$  is equal to  $f(0) = 1/4$ . This proves the lemma.  $\square$

**Lemma 2.** For all  $\alpha, a \in \mathbb{R}$ , we have the following Taylor expansion in the neighbourhood of  $a = 0$ :

$$\int_a^\infty N(\alpha x)g(x)dx = \frac{1}{4} + \frac{1}{2\pi} \arctan(\alpha) - \frac{a}{2\sqrt{2\pi}} - \frac{\alpha a^2}{4\pi} + \frac{a^3}{12\sqrt{2\pi}} + \frac{(\alpha^3 + 3\alpha)a^4}{48\pi} - \frac{a^5}{80\sqrt{2\pi}} + \mathcal{O}(a^6).$$

*Proof.* Let  $f(a) = \int_a^\infty N(\alpha x)g(x)dx$ . First, getting the derivatives of  $g$  is straightforward thanks to Hermite polynomials:

$$\begin{cases}
 g(0) &= \frac{1}{\sqrt{2\pi}} \\
 g^{(1)}(0) &= 0 \\
 g^{(2)}(0) &= -\frac{1}{\sqrt{2\pi}} \\
 g^{(3)}(0) &= 0 \\
 g^{(4)}(0) &= \frac{3}{\sqrt{2\pi}}.
 \end{cases}$$

The function  $f$  is differentiable and, using Leibniz formula, we get:

1.  $f^{(1)}(a) = -N(\alpha a)g(a)$  and  $f^{(1)}(0) = -\frac{1}{2\sqrt{2\pi}}$ ,

2.  $f^{(2)}(a) = -\alpha g(\alpha a)g(a) - N(\alpha a)g^{(1)}(a)$  and  $f^{(2)}(0) = -\frac{\alpha}{2\pi}$ ,
3.  $f^{(3)}(a) = -\alpha^2 g^{(1)}(\alpha a)g(a) - 2\alpha g(\alpha a)g^{(1)}(a) - N(\alpha a)g^{(2)}(a)$  and  $f^{(3)}(0) = \frac{1}{2\sqrt{2\pi}}$ ,
4.  $f^{(4)}(a) = -\alpha^3 g^{(2)}(\alpha a)g(a) - 3\alpha^2 g^{(1)}(\alpha a)g^{(1)}(a) - 3\alpha g(\alpha a)g^{(2)}(a) - N(\alpha a)g^{(3)}(a)$  and  $f^{(4)}(0) = \frac{\alpha^3 + 3\alpha}{2\pi}$ ,
5.  $f^{(5)}(a) = -\alpha^4 g^{(3)}(\alpha a)g(a) - 4\alpha^3 g^{(2)}(\alpha a)g^{(1)}(a) - 6\alpha^2 g^{(1)}(\alpha a)g^{(2)}(a) - 4\alpha g(\alpha a)g^{(3)}(a) - N(\alpha a)g^{(4)}(a)$  and  $f^{(5)}(0) = -\frac{3}{2\sqrt{2\pi}}$ .

We conclude with Taylor's theorem applied to  $f$  in  $a = 0$ , using Lemma 1 for  $f(0)$ .  $\square$

**Lemma 3.** For all  $\alpha, a \in \mathbb{R}$ , we have

$$\int_a^{+\infty} uN(\alpha u)g(u)du = N(\alpha a)g(a) + \frac{\alpha}{\sqrt{2\pi(1+\alpha^2)}} \left(1 - N\left(a\sqrt{1+\alpha^2}\right)\right).$$

*Proof.* By noting that  $-g$  is a primitive function of  $u \mapsto ug(u)$ , we integrate by part the following equation:

$$\begin{aligned} \int_a^b uN(\alpha u)g(u)du &= [-N(\alpha u)g(u)]_a^b + \alpha \int_a^b g(\alpha u)g(u)du \\ &= -N(\alpha b)g(b) + N(\alpha a)g(a) + \frac{\alpha}{\sqrt{2\pi}} \int_a^b g(u\sqrt{1+\alpha^2})du \\ &= -N(\alpha b)g(b) + N(\alpha a)g(a) + \frac{\alpha}{\sqrt{2\pi(1+\alpha^2)}} (N(b\sqrt{1+\alpha^2}) - N(a\sqrt{1+\alpha^2})) \end{aligned}$$

after a change of variable between the second and the third line. By doing  $b \rightarrow +\infty$ , we get Lemma 3.  $\square$

## B Proof of Theorem 1

*Proof.* We treat separately the conditional and the non-conditional hit ratios. We begin with the non-conditional case.

1. According to equation (4), the predicted price return is the random variable

$$\hat{R}_{t,t+h|\Delta} = \Sigma_{RS}\Sigma_S^{-1}S.$$

Its mean is zero and its variance is:

$$\begin{aligned} \text{Var}\left(\hat{R}_{t,t+h|\Delta}\right) &= \mathbb{E}\left[\Sigma_{RS}\Sigma_S^{-1}S(\Sigma_{RS}\Sigma_S^{-1}S)^T\right] \\ &= \Sigma_{RS}\Sigma_S^{-1}\mathbb{E}[SS^T](\Sigma_S^T)^{-1}\Sigma_{RS}^T \\ &= \Sigma_{RS}\Sigma_S^{-1}\Sigma_S\Sigma_S^{-1}\Sigma_{RS}^T \\ &= \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T, \end{aligned}$$

where we have used the fact that  $\Sigma_S$  is symmetric. The covariance of  $\hat{R}_{t,t+h|\Delta}$  with  $R_{t,t+h}$  is:

$$\begin{aligned} \text{Cov}(\hat{R}_{t,t+h|\Delta}, R_{t,t+h}) &= \mathbb{E}[\Sigma_{RS}\Sigma_S^{-1}SR_{t,t+h}] \\ &= \Sigma_{RS}\Sigma_S^{-1}\mathbb{E}[SR_{t,t+h}] \\ &= \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T. \end{aligned}$$

Then, the covariance matrix  $\Gamma$  of the Gaussian vector  $[\hat{R}_{t,t+h|\Delta} \ R_{t,t+h}]^T$  is:

$$\Gamma = \begin{pmatrix} \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T & \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T \\ \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T & \sigma^2 h^{2H} \end{pmatrix},$$

which we can also write  $\Gamma = \Sigma\Sigma^T$ , with  $\Sigma$  a lower triangular matrix obtained by Cholesky decomposition:

$$\Sigma = \begin{pmatrix} a & 0 \\ a & b \end{pmatrix},$$

with  $a = \sqrt{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$  and  $b = \sqrt{\sigma^2 h^{2H} - \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$ . Therefore, there exists two Gaussian variables  $U, V \sim \mathcal{N}(0, 1)$ , independent from each other, such that

$$[\hat{R}_{t,t+h|\Delta} \ R_{t,t+h}]^T = \Sigma[U \ V]^T. \quad (11)$$

The non-conditional hit ratio can thus be written as a sum of joint probabilities of the two increments  $\hat{R}_{t,t+h|\Delta}$  and  $R_{t,t+h}$ , which can be expressed as a combination of  $U$  and  $V$ :

$$\begin{aligned} \rho &= \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} \geq 0, R_{t,t+h} \geq 0 \right] + \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} \leq 0, R_{t,t+h} \leq 0 \right] \\ &= \mathbb{P} [aU \geq 0, aU + bV \geq 0] + \mathbb{P} [aU \leq 0, aU + bV \leq 0] \\ &= \mathbb{P} [U \geq 0, V \geq -(a/b)U] + \mathbb{P} [U \leq 0, V \leq -(a/b)U] \\ &= \int_0^{+\infty} \int_{-au/b}^{+\infty} (g(u)g(v)) dv du + \int_{-\infty}^0 \int_{-\infty}^{-au/b} (g(u)g(v)) dv du \\ &= \int_0^{+\infty} (1 - N(-\frac{a}{b}u)) g(u) du + \int_{-\infty}^0 N(-\frac{a}{b}u) g(u) du \\ &= \frac{1}{2} - \int_0^{+\infty} N(-\frac{a}{b}u) g(u) du + \int_0^{+\infty} N(\frac{a}{b}u) g(u) du \\ &= \frac{1}{2} - \frac{1}{4} - \frac{1}{2\pi} \arctan(-\frac{a}{b}) + \frac{1}{4} + \frac{1}{2\pi} \arctan(\frac{a}{b}) \\ &= \frac{1}{2} + \frac{1}{\pi} \arctan(\frac{a}{b}). \end{aligned}$$

where  $g$  and  $N$  are respectively the probability density function and the cumulative distribution function of a standard Gaussian variable and where we used Lemma 1, which is in Appendix A. Moreover, we note that

$$\frac{a}{b} = \frac{\sqrt{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}}{\sqrt{\sigma^2 h^{2H} - \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}} = \left( \frac{\sigma^2 h^{2H}}{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T} - 1 \right)^{-1/2},$$

so that

$$\rho = \frac{1}{2} + \frac{1}{\pi} \arctan \left( \left[ \frac{\sigma^2 h^{2H}}{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T} - 1 \right]^{-1/2} \right) = 1 - \frac{1}{\pi} \arctan \left( \sqrt{\frac{\sigma^2 h^{2H}}{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T} - 1} \right).$$

2. We now deal with the conditional case. Thanks to the above Cholesky decomposition and to equation (11), we know that, conditionally to  $S$ , the variable  $R_{t,t+h}$  is Gaussian of mean  $\mathbb{E}\{R_{t,t+h}|S\} = \hat{R}_{t,t+h|\Delta}$  and of variance  $b^2$ . Therefore, the conditional hit ratio is, where  $V \sim \mathcal{N}(0, 1)$ :

$$\begin{aligned} \rho^c(y) &= \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} \geq 0, \hat{R}_{t,t+h|\Delta} + bV \geq 0 | S = y \right] + \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} < 0, \hat{R}_{t,t+h|\Delta} + bV < 0 | S = y \right] \\ &= \mathbf{1}_{\{\Sigma_{RS}\Sigma_S^{-1}y \geq 0\}} \mathbb{P} [V \geq -\Sigma_{RS}\Sigma_S^{-1}y/b] + \mathbf{1}_{\{\Sigma_{RS}\Sigma_S^{-1}y < 0\}} \mathbb{P} [V < -\Sigma_{RS}\Sigma_S^{-1}y/b] \\ &= \mathbf{1}_{\{\Sigma_{RS}\Sigma_S^{-1}y \geq 0\}} N(\Sigma_{RS}\Sigma_S^{-1}y/b) + \mathbf{1}_{\{\Sigma_{RS}\Sigma_S^{-1}y < 0\}} N(-\Sigma_{RS}\Sigma_S^{-1}y/b) \\ &= N(|\Sigma_{RS}\Sigma_S^{-1}y|/b). \end{aligned}$$

□

## C Proof of Theorem 2

*Proof.* Using the same approach as in the proof of Theorem 1, we can write:

$$\begin{aligned}
p^+(\theta) &= \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} \geq \theta, R_{t,t+h} > 0 \right] + \mathbb{P} \left[ \hat{R}_{t,t+h|\Delta} \leq -\theta, R_{t,t+h} < 0 \right] \\
&= \mathbb{P} [aU \geq \theta, aU + bV > 0] + \mathbb{P} [aU \leq -\theta, aU + bV < 0] \\
&= \mathbb{P} [U \geq \theta/a, V > -(a/b)U] + \mathbb{P} [U \leq -\theta/a, V < -(a/b)U] \\
&= \int_{\theta/a}^{+\infty} \int_{-au/b}^{+\infty} (g(u)g(v))dvdu + \int_{-\infty}^{-\theta/a} \int_{-\infty}^{-au/b} (g(u)g(v))dvdu \\
&= \int_{\theta/a}^{+\infty} (1 - N(-\frac{a}{b}u)) g(u)du + \int_{-\infty}^{-\theta/a} N(-\frac{a}{b}u) g(u)du \\
&= 1 - N(\theta/a) - \int_{\theta/a}^{+\infty} N(-\frac{a}{b}u)g(u)du + \int_{\theta/a}^{+\infty} N(\frac{a}{b}u)g(u)du,
\end{aligned} \tag{12}$$

by symmetry of  $g$ , the Gaussian density, with a change of variable in the second integral, and where  $U, V \sim \mathcal{N}(0, 1)$  are random variables independent from each other,  $a = \sqrt{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$ , and  $b = \sqrt{\sigma^2 h^2 H - \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$ . Using Lemma 2, which is in Appendix A, we obtain:

$$p^+(\theta) = 1 - N(\theta/a) + \frac{1}{\pi} \arctan\left(\frac{a}{b}\right) - \frac{\theta^2}{2\pi ab} + \left(\frac{1}{ab^3} + \frac{3}{a^4 b}\right) \frac{\theta^4}{24\pi} + \mathcal{O}(\theta^6).$$

Similarly, we can write

$$\begin{aligned}
p^-(\theta) &= \mathbb{P} [aU \leq -\theta, aU + bV > 0] + \mathbb{P} [aU \geq \theta, aU + bV < 0] \\
&= \int_{\theta/a}^{+\infty} N(-\frac{a}{b}u)g(u)du + \int_{-\infty}^{-\theta/a} (1 - N(-\frac{a}{b}u))g(u)du \\
&= 1 - N(\theta/a) + \int_{\theta/a}^{+\infty} N(-\frac{a}{b}u)g(u)du - \int_{\theta/a}^{+\infty} N(\frac{a}{b}u)g(u)du.
\end{aligned} \tag{13}$$

Using again Lemma 2, we get:

$$p^-(\theta) = 1 - N(\theta/a) - \frac{1}{\pi} \arctan\left(\frac{a}{b}\right) + \frac{\theta^2}{2\pi ab} - \left(\frac{1}{ab^3} + \frac{3}{a^4 b}\right) \frac{\theta^4}{24\pi} + \mathcal{O}(\theta^6).$$

From equations (12) and (13), we also obtain:

$$p^0(\theta) = 1 - p^+(\theta) - p^-(\theta) = -1 + 2N\left(\frac{\theta}{a}\right).$$

This proves Theorem 2. □

## D Proof of Theorem 3

*Proof.* By definition,

$$\tilde{R}_{t,t+h}(\theta) = \mathbb{E} \left[ R_{t,t+h} \left( \mathbf{1}_{\{\hat{R}_{t,t+h|\Delta} \geq \theta\}} - \mathbf{1}_{\{\hat{R}_{t,t+h|\Delta} \leq -\theta\}} \right) \right].$$

Like in the proof of Theorems 1 and 2, we can write it:

$$\tilde{R}_{t,t+h}(\theta) = \mathbb{E} [(aU + bV) (\mathbf{1}_{\{aU \geq \theta\}} - \mathbf{1}_{\{aU \leq -\theta\}})],$$

where  $U, V \sim \mathcal{N}(0, 1)$  are two independent Gaussian variables,  $a = \sqrt{\Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$ , and  $b = \sqrt{\sigma^2 h^2 H - \Sigma_{RS}\Sigma_S^{-1}\Sigma_{RS}^T}$ . Then

$$\begin{aligned}
\tilde{R}_{t,t+h}(\theta) &= \int_{\theta/a}^{+\infty} \left\{ \int_{\mathbb{R}} (au + bv) g(v)g(u)dv \right\} du - \int_{-\infty}^{-\theta/a} \left\{ \int_{\mathbb{R}} (au + bv) g(v)g(u)dv \right\} du \\
&= \int_{\theta/a}^{+\infty} aug(u)du - \int_{-\infty}^{-\theta/a} aug(u)du \\
&= 2ag(\theta/a),
\end{aligned}$$

using the symmetry of  $g$  and the fact that  $ug(u) = -g'(u)$ . Similarly,

$$\begin{aligned}
\tilde{\sigma}_{t,t+h}^-(\theta) &= -\mathbb{E} \left[ R_{t,t+h} \left( \mathbb{1}_{\{\hat{R}_{t,t+h}|\Delta \geq \theta, R_{t,t+h} < 0\}} - \mathbb{1}_{\{\hat{R}_{t,t+h}|\Delta \leq -\theta, R_{t,t+h} > 0\}} \right) \right] \\
&= -\mathbb{E} \left[ (aU + bV) \left( \mathbb{1}_{\{aU \geq \theta, aU + bV < 0\}} - \mathbb{1}_{\{aU \leq -\theta, aU + bV > 0\}} \right) \right] \\
&= -\int_{\theta/a}^{+\infty} \left\{ \int_{-\infty}^{-au/b} (au + bv) g(v) g(u) dv \right\} du \\
&\quad + \int_{-\infty}^{-\theta/a} \left\{ \int_{-au/b}^{+\infty} (au + bv) g(v) g(u) dv \right\} du \\
&= -\int_{\theta/a}^{+\infty} \left[ auN\left(-\frac{au}{b}\right) - bg\left(-\frac{au}{b}\right) \right] g(u) du \\
&\quad + \int_{-\infty}^{-\theta/a} \left[ au\left(1 - N\left(-\frac{au}{b}\right)\right) + bg\left(-\frac{au}{b}\right) \right] g(u) du \\
&= -2a \int_{\theta/a}^{+\infty} uN\left(-\frac{au}{b}\right) g(u) du + 2b \int_{\theta/a}^{+\infty} g\left(\frac{au}{b}\right) g(u) du,
\end{aligned}$$

after a change of variable  $u \mapsto -u$  in the last line. This leads to the general formula, where we use Lemma 3 (which is postponed in Appendix A) in the first integral, the fact that  $g(au/b)g(u) = g(u\sqrt{1+(a/b)^2})/\sqrt{2\pi}$  and a change of variable  $u \mapsto u\sqrt{1+(a/b)^2}$  in the second integral:

$$\begin{aligned}
\tilde{\sigma}_{t,t+h}^-(\theta) &= -2a \left[ N\left(-\frac{\theta}{b}\right) g\left(\frac{\theta}{a}\right) - \frac{a}{\sqrt{2\pi(a^2+b^2)}} \left(1 - N\left(\theta\sqrt{\frac{1}{a^2} + \frac{1}{b^2}}\right)\right) \right] \\
&\quad + \frac{2b}{\sqrt{2\pi}\sqrt{1+(a/b)^2}} \left(1 - N\left(\theta\sqrt{\frac{1}{a^2} + \frac{1}{b^2}}\right)\right) \\
&= -2aN\left(-\frac{\theta}{b}\right) g\left(\frac{\theta}{a}\right) + \frac{2a^2+2b^2}{\sqrt{2\pi(a^2+b^2)}} \left(1 - N\left(\theta\sqrt{\frac{1}{a^2} + \frac{1}{b^2}}\right)\right) \\
&= -2aN\left(-\frac{\theta}{b}\right) g\left(\frac{\theta}{a}\right) + \sqrt{\frac{2(a^2+b^2)}{\pi}} N\left(-\theta\sqrt{\frac{1}{a^2} + \frac{1}{b^2}}\right).
\end{aligned}$$

By noting that  $a^2 + b^2 = \sigma^2 h^{2H}$ , we obtain the results displayed in Theorem 3.  $\square$