High-Frequency Trading with Fractional Brownian Motion

Yuliya Mishura

Taras Shevchenko National University of Kyiv

2021

Outline

This talk is the result of common research with Paolo Guasoni and Miklos Rasony. It is published in Guasoni, P., Mishura, Y., & Rásonyi, M. (2021). High-frequency trading with fractional Brownian motion. Finance and stochastics, 25(2), 277-310.

In the high-frequency limit, conditional expected increments of fractional Brownian motion converge to a white noise, shedding their dependence on the path history and the forecasting horizon, and making dynamic optimization problems tractable. We find an explicit formula for locally mean-variance optimal strategies and their performance for an asset price that follows fractional Brownian motion. Without trading costs, risk-adjusted profits are linear in the trading horizon and rise asymmetrically as the Hurst exponent departs from Brownian motion, remaining finite as the exponent reaches zero while diverging as it approaches one. Trading costs penalize numerous portfolio updates from short-lived signals, leading to a finite trading frequency, which can be chosen so that the effect of trading costs is arbitrarily small, depending on the required speed of convergence to the high-frequency limit.

Introduction

First proposed as a model of price dynamics by Mandelbrot [17], fractional Brownian motion (fBm) has since puzzled researchers and stirred controversy for its elusive properties, which have confounded both empirical and theoretical work. Long-range dependence in asset prices, the property that originally motivated the use of fBm to describe price dynamics, remains undecided [16, 14, 13, 19, 15, 11, 12, 1]. Arbitrage, which has plagued its adoption in models of optimal investment [10, 9, 8, 7] disappears with frictions [5, 6, 4, 3], leading to finite expected profits [29].

This talk is devoted to locally mean-variance optimal trading strategies in fractional Brownian motion and characterization of their convergence and performance in the high-frequency limit. Our analysis starts from a fixed trading frequency, for which optimal strategies are directly proportional to the (conditional) expected increment and inversely proportional to its variance. The central feature of fractional Brownian motion is that, unlike diffusion models, the conditional expected increment is not proportional to the length of the trading period, but to a power thereof – the Hurst exponent. Because the increment's standard deviation scales with the same power, the average Sharpe ratio is insensitive to the length of the trading period.

The key insight (Theorem 2.3) is that the high-frequency limit of such a forecast (the “latent drift” of fractional Brownian motion) is a white-noise process with a variance depending on the Hurst exponent, but invariant to any scaling of the process (which would affect both expected increments and their variance). This result in turn leads to a cascade of implications for optimal continuous trading of fractional Brownian motion.

First, the optimal mean-variance performance from trading fractional Brownian motion is proportional to the length of the whole trading horizon – as for Brownian motion with drift – in spite of the different scaling of mean and variance on individual periods. The reason is that the *cumulative* performance of high-frequency trading fractional Brownian motion on a finite horizon is essentially equivalent to the *average* performance of a discrete-time model with infinitely many periods and independent, identically-distributed Sharpe ratios. Both performances are deterministic because randomness disappears through ergodicity.

Second, the resulting performance is asymmetric in the Hurst exponent (Figure 1 and Theorem 2.2), remaining bounded as the process approaches a white noise (near $H = 0$), but diverging as it approaches a near-straight line with random drift (near $H = 1$). This result is significant because it does not stem from the autocovariance properties of the strategies' expected returns: indeed, for any Hurst exponent, the instantaneous forecast is a white noise. Instead, the result reflects the magnitude of the variance of the white noise forecast that is extracted from the paths of fBm for different values of H: for H near zero, the weights of the white noise forecast are small and highly concentrated on recent increments, which results in a moderate variance. By contrast, for H near one the forecast's weights are large and reach far into the path's past history, leading to a diverging variance.

Third, and contrary to the intuition from previous results, including our own, we find that such performance is immune to small frictions – such as proportional transaction costs or immediate nonlinear price impact (Theorem 3.1 and Corollary 3.2). Specifically, while frictions detract from performance, their effect vanishes arbitrarily quickly by slowly increasing the trading frequency as trading costs decrease, so that their asymptotic impact vanishes at any required rate. Similarly, holding trading costs constant, their effect also vanishes by increasing the horizon, while appropriately calibrating the trading frequency.

Fourth, we observe that approximations of the latent drift of fractional Brownian motion converge weakly, but not in norm. This observation highlights a qualitative difference between the familiar drifts of diffusions and the their partial analogies for fractional processes. Not only are diffusive drifts of the order of infinitesimal time intervals (informally, dt), while fractional drifts are a power thereof (informally, $(dt)^H$): in addition, diffusive drifts can be understood as close approximations of expected conditional increments over any sufficiently small interval, because such approximations converge (in norm) as random variables. By contrast, fractional drifts are critically dependent on the specific interval: as the interval length declines to zero, the conditional expected returns converge in law, but not as random variables in any reasonable sense.

Finally, it is worthwhile comparing the findings in this paper to the recent results in [29], as both articles study optimal trading strategies for fractional Brownian motion, though in very different settings. The main difference lies in the objective functions considered, here a local mean-variance criterion on a finite interval, while in [29] a risk-neutral target with a long horizon. In particular, the presence of a nonlinear friction is crucial to make the problem in [29] well posed, which would otherwise lead to unbounded expected profits. Vice versa, the present local mean-variance criterion is well-posed even without frictions, as the instantaneous Sharpe ratio remains bounded for any $H \in (0, 1)$, albeit arbitrage is feasible on any interval, because arbitrage profits remain dispersed.

Both [29] and the present paper lead to finite maximal Sharpe ratios that are asymmetric in H, but their skews are reversed, and arise for different reasons: while the asymptotically optimal strategies in [29] have higher Sharpe ratios near zero than near one, they are not necessarily optimal, as the strategies maximize a risk-neutral objective – not the Sharpe ratio.

By contrast, the Sharpe ratios obtained here are indeed optimal, as they maximize the local mean-variance criterion over any finite interval by ergodicity. The rest of the paper is organized as follows. Section 2 describes the model and the main result without frictions, discussing their significance and implications. Section 3 considers frictions, and shows how their impact can be mitigated by a judicious choice of the trading frequency.

Main results

An investor trades a safe and a risky asset. The safe rate is assumed zero to simplify notation, while the price of the risky asset is a multiple of fractional Brownian motion (fBm).

Definition 2.1

Fractional Brownian motion (fBm) with Hurst index $H \in (0,1)$ is a Gaussian process $B^H = \{B_t^H, t \geq 0\}$, defined on a probability space $(\Omega, \mathcal{F}, P)$, with continuous trajectories such that $EB_t^H = 0$, $t \geq 0$ and $E[B_t^H B_s^H] = \frac{1}{2}(t^{2H} + s^{2H} - |t-s|^{2H}),\ t, s \geq 0$.

The case $H = 1/2$ corresponds to usual Brownian motion, henceforth excluded. Thus, $H \in (0,1) \setminus \{1/2\}$ unless stated otherwise. Consider a trading horizon $T > 0$ and a frequency $n \geq 1$, which represents the number of trading periods in the interval $[0, T]$. The corresponding set of strategies, Σ_n, consists of sequences π_s, $s \in \{(Tk)/n :\ 0 \leq k \leq n-1\}$ of random variables such that π_s is $\mathcal{F}_s$-measurable for all such s, where $(\mathcal{F}_s)_{s \geq 0}$ is the augmented natural filtration of B_s^H.

An investor who holds at the beginning of each interval $[kT/n, (k+1)T/n]$ a number of shares equal to $\pi_{\frac{Tk}{n}}$ attains the mean-variance performance

$$\mathrm{MV}(\pi, k, n, S) := E_{\frac{Tk}{n}}[\pi_{\frac{Tk}{n}}(S_{\frac{T(k+1)}{n}} - S_{\frac{Tk}{n}})] - \frac{\gamma}{2}\mathrm{Var}_{\frac{Tk}{n}}(\pi_{\frac{Tk}{n}}(S_{\frac{T(k+1)}{n}} - S_{\frac{Tk}{n}}))$$

where $S_t = \sigma B_t^H$ denotes the risky asset price, while $E_t[X]$ and $\mathrm{Var}_t(X) := E_t[X^2] - (E_t[X])^2$ respectively denote the conditional expectation and conditional variance of a random variable X with respect to $\mathcal{F}_t$. The parameter $\gamma > 0$ represents the investor's aversion to risk (as variance).

Assuming time-additive preferences, the overall performance in the interval $[0, T]$ of a trading strategy is defined as the sum of the per-period performance, weighing each period by its length T/n, i.e.,

$$R_\gamma(\pi, n, S) := \frac{T}{n}\sum_{k=0}^{n-1}\mathrm{MV}(\pi, k, n, S) = \int_0^T \mathrm{MV}(\pi, \lfloor tn \rfloor, n, S)dt.$$

Thus, the high-frequency objective for fractional Brownian motion is

$$V(H,\gamma) := \limsup_{n\to\infty} \sup_{\pi\in\Sigma_n} E[R_\gamma(\pi, n, \sigma B^H)],$$

and represents the maximal performance of a continuous-time strategy that updates the portfolio at arbitrary frequency on the interval $[0, T]$.

With this notation, the main result of this paper is

Theorem 2.2

For each $H \in (0,1)$,

$$V(H,\gamma) = \frac{T}{\gamma}\left[\frac{\Gamma(H+1/2)\Gamma(2-2H)}{2\Gamma(3/2-H)} - \frac{1}{2}\right], \tag{2.1}$$

and the limit superior in the definition of $V(H,\gamma)$ is, in fact, a limit.

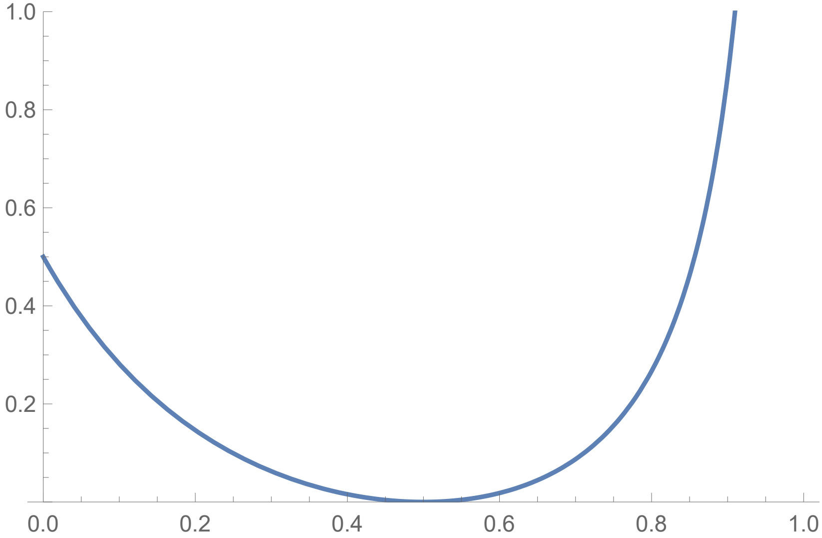

Figure: Profits per unit of risk and and time (vertical), i.e., the expression in (2.1) divided by T/γ, against Hurst exponent (horizontal).

Before discussing the details of this result, it is useful to compare it to the familiar benchmark of Brownian motion with drift, i.e.,

$$S_t = \mu t + \sigma W_t, \tag{2.2}$$

for some Brownian motion W and $\mu, \sigma > 0$. A simple calculation then shows that

$$\sup_{\pi \in \Sigma_n} E[R_\gamma(\pi, n, S)] = E\left[R_\gamma\left(\frac{\mu}{\gamma\sigma^2}, n, S\right)\right] = \frac{\mu^2}{2\gamma\sigma^2} T. \tag{2.3}$$

In other words, both the optimal strategy and its performance do not depend on n, are inversely proportional to the squared volatility σ^2 and risk aversion γ, and are respectively linear and quadratic in the drift. In addition, performance is linear in the investment horizon.

The linear dependence on the drift and the inverse dependence on the volatility is at the heart of the risk-return tradeoff that arises in random-walk models: as returns are serially independent, their randomness is purely a source of risk, and its reduction is unambiguously beneficial.

The fractional high-frequency performance in (2.1) contains surprising features both in its departures and in its analogies with the usual mean-variance performance (2.3). In contrast to (2.3), the performance in (2.1) is independent of volatility.[1] As shown below, the optimal strategy *inversely* depends on variance, but such dependence is lost in performance because the expected return *directly* depends on variance, thereby offsetting its effect.

[1] In fact, the result is also independent of an additional drift. Intuitively, the reason is that, for a short time interval, the conditional expected increment of fBm if of order $(dt)^H$, which makes an ordinary drift of order dt negligible in the mean-variance optimal strategy and its performance.

In analogy to (2.3), the performance in (2.1) is linear in the investment horizon. Upon reflection, also such an analogy is surprising, because the linearity in the horizon of the usual mean-variance performance in (2.3) stems from the independence of increments of Brownian motion and the constant drift. Instead, the dependence in increments of fractional Brownian motion is substantial, and indeed crucial to generate positive returns.

The dependence on the Hurst exponent H, displayed in Figure 1, is similarly puzzling in view of its asymmetry. At one extreme, as H approaches zero and increments increasingly resemble white noise [26, Lemma 4.1], performance converges to a finite limit:

$$V(H,\gamma) = \frac{T}{\gamma}\left[\frac{\Gamma(1/2)\Gamma(2)}{2\Gamma(3/2)} - \frac{1}{2}\right] + O(H) = \frac{T}{2\gamma} + O(H). \tag{2.4}$$

(The last equality follows from the identity $\Gamma\left(\frac{3}{2}\right) = \frac{1}{2}\Gamma\left(\frac{1}{2}\right)$.) As H approaches $1/2$, performance flattens around zero, as the process mimics an ordinary Brownian motion: [2]

$$V(H,\gamma) = \frac{\pi^2}{6}(H-1/2)^2\frac{T}{\gamma} + O\left((H-1/2)^3\right). \tag{2.5}$$

[2] This expansion exploits identities involving the derivatives of the Gamma function, cf. [31].

In particular, this identity confirms the intuition from Figure 1 that performance reaches its unique minimum of zero in the martingale case of $H = 1/2$, while slowly increasing in each direction. Indeed, at the other extreme, as H approaches one and the process resembles a straight line with random slope, performance diverges: [3]

$$V(H,\gamma) = \left(\frac{1}{8(1-H)} + \frac{-3+\log 4}{4} \right) \frac{T}{\gamma} + O\left(1-H\right). \tag{2.6}$$

[3] To obtain the term $1/(8(1-H))$, recall that $\Gamma(x) \sim 1/x$ for x near zero. The term $(-3+\log 4)/4$ follows from more complex higher order asymptotics.

Key to understanding these features is the prediction mechanism at the heart of the problem. As our mean-variance objective is time-additive, the optimal trading strategies maximize performance in the next period. Because for a square-integrable random variable X, the functional $\varphi \to E[\varphi X] - \gamma/2\,\mathrm{Var}(\varphi X)$ attains its maximum $E^2[X]/(2\gamma\,\mathrm{Var}(X))$ at $\phi^* = E[X]/(\gamma\,\mathrm{Var}(X))$, the optimal strategy $\pi(n) \in \Sigma_n$ is

$$\pi_{(Tk)/n}(n) := \frac{E_{(Tk)/n}[B^H_{(T(k+1))/n} - B^H_{(Tk)/n}]}{\gamma\,\mathrm{Var}_{(Tk)/n}(B^H_{(T(k+1))/n} - B^H_{(Tk)/n})},\ 0 \le k \le n-1. \quad (2.7)$$

To investigate the high-frequency limit, it is convenient to extend these strategies by right-continuity to the entire interval $[0, T]$, i.e., setting

$$\pi_t(n) := \pi_{(Tk)/n}(n), \quad t \in [(Tk)/n, T(k+1)/n),\ 0 \le k < n.$$

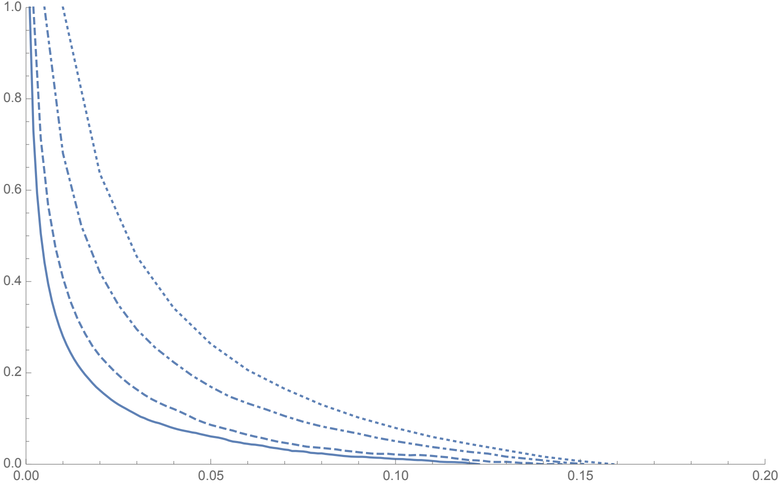

Figure: Autocorrelation (vertical axis) of the strategies $\pi_{(Tk)/n}(n)$ (equivalently, of the expected increments) against time lag (horizontal), as the frequency n increases from 100 (top) to 200, 500, and 1000 (bottom). The autocorrelation converges to the white-noise limit of one at lag zero and zero elsewhere. Each plotted curve is the average of a thousand sample autocorrelograms with $H = 0.6$.

With this notation, the next theorem identifies the limit of such strategies, which is interpreted as the asymptotically optimal strategy in the high-frequency regime:

Theorem 2.3

The sequence $((T/n)^{-H}\pi_t(n))_{t\in[0,T]}^{n\in\mathbb{N}}$ *consists of Gaussian processes that, as n increases, converge in finite-dimensional distributions to a Gaussian process* $(B_t)_{t\in[0,T]}$ *such that*

1. $B_0 = 0$ *a.s.;*
2. $E[B_t] = 0,\ t \in (0, T].$
3. $E[B_t^2] = 1 - \frac{\Gamma(3/2-H)}{\Gamma(H+1/2)\Gamma(2-2H)},\ t \in (0, T];$
4. $E[B_t B_s] = 0$ *for* $t \neq s,\ s, t \in [0, T].$

This result has a striking message: up to a scaling factor, the optimal strategy – hence the expected return over the next period – is essentially a white noise (the exception is $t = 0$, for which the process is conventionally pinned at zero). In other words, regardless of the Hurst exponent H, and regardless of the autocorrelation of increments in fractional Brownian motion, the forecasts of short-term increments are virtually uncorrelated from one instant to the next.

Figure 2 illustrates the convergence result in the theorem by plotting at increasing frequencies the autocorrelation of $\pi_{(Tk)/n}(n)$, which converges to the autocorrelation of a white noise.

The Hurst exponent controls the scale of the strategy: denoting by Δ the length of each trading period, price increments have conditional expectation of order Δ^H and conditional variance of order Δ^{2H}, which implies trading positions of order Δ^{-H}. This feature is in contrast to the Brownian benchmark, in which both the expected return $\mu\Delta$ and its variance $\sigma^2\Delta$ are of the same order. Instead, in the fractional setting the variance has a smaller order, which means that bets become more favorable as the trading frequency increases, and therefore their optimal size increases.

Note, however, that the implied performance on each period is proportional to the conditional expectation Δ^H times the position size Δ^{-H}, hence it is of order 1. As each trading period leads to the same performance (in view of the white-noise property established in the above theorem), trading over an interval of length T generates a performance proportional to T. In particular, the results below show that the optimal trading position is asymptotically

$$\varphi(\Delta) := \frac{\Delta^{-H}}{\gamma} \frac{\Gamma(H + 1/2)\Gamma(2 - 2H)}{\Gamma(3/2 - H)} B_t \tag{2.8}$$

and that on the subsequent interval the expected increment has the asymptotic (conditional) mean and variance

$$m(\Delta) := B_t \Delta^H, \qquad v(\Delta) := \frac{\Gamma(3/2 - H)}{\Gamma(H + 1/2)\Gamma(2 - 2H)} \Delta^{2H}. \tag{2.9}$$

The performance formula in Theorem 2.2 follows from

$$\left(E[\varphi(\Delta)m(\Delta) - \frac{\gamma}{2}\varphi(\Delta)^2 v(\Delta)]\right) T = \frac{T}{2\gamma}\frac{\Gamma(H+1/2)\Gamma(2-2H)}{\Gamma(3/2-H)}E[B_t^2] \tag{2.10}$$

$$= \frac{T}{\gamma}\left[\frac{\Gamma(H+1/2)\Gamma(2-2H)}{2\Gamma(3/2-H)} - \frac{1}{2}\right]. \tag{2.11}$$

This analysis also offers an intuitive explanation for the asymmetric behavior of performance in equations (2.4) and (2.6). For H close to zero, the asset price S_t itself is akin to a white noise, for which the mean and variance in (2.9) are of the same order. Accordingly, performance converges to a finite limit. By contrast, for H close to one the process degenerates to a straight line with random slope, as randomness vanishes from its increments. Thus, the trading strategy generates return with virtually no risk, and performance diverges.

Note that the mean-variance optimal strategies $\pi_t(n)$ are not arbitrage opportunities, as the support of their payoffs is $(-\infty, +\infty)$. Although continuous trading with fBm leads to arbitrage opportunities (e.g., [9, 10]), it is clear that on any finite deterministic grid fBm does not admit arbitrage, because an equivalent martingale measure can be constructed through a backward recursion that aligns all conditional expected increments to zero. (In fact, arbitrage disappears even when a minimal time has to pass between two subsequent transactions [7].)

A deeper question is whether the sequence of strategies $(\pi_t(n))_{n\geq 1}$ yields an arbitrage in some limit sense, and the answer is affirmative. The sequence of discrete-time mean-variance optimal policies offers a statistical arbitrage, in that

$$\lim_{n\to\infty} E[W(n)] = \infty,$$

where $W(n)$ is the wealth process of the strategies $\pi_t(n)$ starting from initial capital 0, that is,

$$W(n) = \sum_{k=0}^{n-1} \pi_{\frac{Tk}{n}}(n)(S_{\frac{T(k+1)}{n}} - S_{\frac{Tk}{n}}).$$

This fact is readily proved by observing that, in mean-variance optimization, the expectation of the optimal strategy is always twice as large as its variance, whence

$$\lim_{n\to\infty} E[W(n)] = \lim_{n\to\infty} \frac{2n}{T} E[R_\gamma(\pi(n), n, S)] = \infty,$$

because $E[R_\gamma(\pi(n), n, S)]$ tends to a finite nonzero limit (for $H \neq 1/2$) by Theorem 2.2.

As Theorem 2.3 establishes that the rescaled strategies essentially converge to a white noise in finite-dimensional distributions, a natural question is whether such a convergence holds in a stronger sense, such as in square-norm, so that its limit can be interpreted as a rescaled asymptotically optimal strategy in continuous time.

The next result provides a negative answer to this question, showing that, even focusing on a sequence of dyadic partitions, the square norm between each discretization and the next remains bounded away from zero.

Theorem 2.4

Let $\Delta_k = T/k$. For all $H \in (0,1) \setminus \{1/2\}$ and $t \in (0, T)$,

$$\lim_{n\to\infty} E\left[\Delta_{2^n}^{-H}\pi_t(2^n) - \Delta_{2^{n+1}}^{-H}\pi_t(2^{n+1})\right]^2 > 0.$$

The significance of this result is that the optimal strategy is extremely sensitive to the trading frequency used, and that optimal strategies at increasing frequencies are not approximations of some underlying continuous-time strategy – which does not exist. In fact, even if such strategy existed, it would be of no use, because the paths of a white-noise process are not even measurable (cf. Revuz and Yor, p. 37).

At a more concrete level, the above results show that, as the frequency increases, the corresponding trading strategies become increasingly variable: thus, in practice their ostensible theoretical performance may be more than offset by the trading costs that such strategies entail. The next section investigates this issue by identifying how the optimal trading frequency depends on the size of trading costs.

Trading Costs

The optimal strategy identified in the previous section implies that asset positions are both large and highly variable, thereby calling into question their robustness to trading costs. To investigate this issue, recall the sequence of strategies $\pi(n)$, $n \geq 1$, defined in (2.7) above.

Assuming that a portfolio changes from θ_1 to θ_2 shares incurs the cost $\lambda|\theta_1 - \theta_2|^\alpha$ for some $\alpha, \lambda > 0$, the local mean variance analysis for a trader applying the strategy $\pi(n)$ leads to the functionals (setting $\pi_t(n) := 0$ for $t < 0$)

$$\tilde{R}(n) := R(n) - \sum_{k=0}^{n-1} \lambda |\pi_{\frac{Tk}{n}}(n) - \pi_{\frac{T(k-1)}{n}}(n)|^\alpha,$$

where the $R(n)$ denotes frictionless performance

$$\tilde{R}(n) := \frac{T}{n} \sum_{k=0}^{n-1} \left[E_{\frac{Tk}{n}} \left[\pi_{\frac{Tk}{n}}(n)(S_{\frac{T(k+1)}{n}} - S_{\frac{Tk}{n}}) \right] \right.$$
$$\left. - \frac{\gamma}{2} \mathrm{Var}_{\frac{Tk}{n}}(\pi_{\frac{Tk}{n}}(n)(S_{\frac{T(k+1)}{n}} - S_{\frac{Tk}{n}})) \right],$$

while the second term in $\tilde{R}(n)$ represents the effect of trading costs.

The next results shows that expected trading costs $E[R(n) - \tilde{R}(n)]$ grow with a superlinear power of the trading frequency n that increases with both the Hurst and the friction exponents. As a result, for fixed transaction costs the objective function arbitrarily deteriorates as the frequency increases, and the optimal trading frequency must be finite.

Theorem 3.1

$E[R(n) - \tilde{R}(n)] = O(n^{1+\alpha H})$ *and therefore* $\lim_{n\to\infty} E[\tilde{R}(n)] = -\infty$.

The next logical step is to understand the effect of small trading costs on the overall objective. Here the above result leads to a surprising implication: with a judicious choice of the trading frequency, the effect of frictions is negligible at any order.

Corollary 3.2

Let $n_\lambda = \lfloor \lambda^{\frac{\beta-1}{1+\alpha H}} \rfloor$ *for* $\beta \in (0,1)$*. Then*

$$E[R(n_\lambda) - \tilde{R}(n_\lambda)] = O(\lambda^\beta) \quad \textit{for} \quad \lambda \to 0, \tag{3.1}$$

that is, trading costs are of order λ^β*.*

Upon reflection, this result is a direct consequence of Theorem 3.1. Yet, its conclusion is counterintuitive when compared to the results for frictions in familiar diffusion models (cf. [32, Theorem 4.1]), whereby the welfare loss is of the order of $\lambda^{\frac{2}{2+\alpha}}$: for example, proportional transaction costs correspond to $\alpha = 1$, leading to a welfare loss of order $2/3$.

Intuitively, the main difference is that in familiar diffusion models the main determinant of optimal portfolios is local drift of the asset price, which is typically smooth. Thus, as the trading frequency increases, smaller and smaller adjustments are required, which means that, holding trading costs constant, the high-frequency limit of portfolio performance is finite.

Vice versa, in the fractional setting considered here, the "latent drift" of the process is highly irregular – in the limit, it is a white noise – hence it entails trading costs that grow with the frequency, as implied by Theorem 3.1. Yet, such irregularity can be harnessed to make trading costs negligible in the high-frequency limit, by choosing the trading frequency n_λ to grow slowly as λ decreases, so that overall costs vanish in the limit. Of course, letting n grow more slowly has the downside that the convergence of the strategy's performance to the optimum in (2.1) is also going to be slower.

The second part of the corollary identifies the maximal speed at which the frequency may grow so that the strategy converges to the optimum. In particular, n_λ may grow at a rate arbitrarily close to $\lambda^{-\frac{1}{1+\alpha H}}$, but not at this exact rate: at such a critical regime, costs would not vanish but converge to a positive finite limit that would be suboptimal.

As trading costs are fixed in applications, the significance of this result is as follows: in practice, the trading cost λ implies that the optimal trading interval should be $1/n_\lambda \approx \lambda^{\frac{1-\beta}{1+\alpha H}}$, where β is close to one. However, the closer β is to one, the larger the trading interval is, which means that the convergence of the strategy to the frictionless limit for a fixed horizon T is slower, and its risk higher. Thus, if the horizon is not long enough to guarantee that the payoff has a sufficiently low risk, one may choose to decrease the value of β to reduce risk further, at the price of an increased trading cost.

Conclusion

This paper finds local mean-variance optimal trading strategies for an asset price that follows fractional Brownian motion, and finds that the average Sharpe ratio is finite, asymmetric in the Hurst exponent, bounded near zero, and unbounded near one. The central result is that conditional expected increments are asymptotically a Gaussian white noise, regardless of the Hurst exponent, but with a variance that depends on that exponent.

The optimal performance is insensitive to small trading frictions, in that their impact can be mitigated arbitrarily well by calibrating the trading frequency appropriately. This phenomenon is in sharp contrast to diffusion models, for which the impact of small frictions has a fixed order of magnitude.

References

[1] Baillie, Richard T.
Long memory processes and fractional integration in econometrics,
Journal of Econometrics, Vol. 73, No.1, pp.5–59, 1996

[2] Bouchaud, Jean-Philippe and Farmer, J Doyne and Lillo, Fabrizio
How markets slowly digest changes in supply and demand,
In: Handbook of Financial Markets: Dynamics and Evolution, eds. T. Hens and K. Schenk-Hoppe, 2008

[3] Czichowsky, Christoph and Peyre, Rémi and Schachermayer, Walter and Yang, Junjian
Shadow prices, fractional Brownian motion, and portfolio optimisation under transaction costs,
Finance and Stochastics, Vol. 22, No. 1, pp. 161–180, 2018

[4] Czichowsky, Christoph and Schachermayer, Walter
Portfolio optimisation beyond semimartingales: shadow prices and fractional Brownian motion,
The Annals of Applied Probability, Vol. 27, No. 3, pp. 1414–1451, 2017

[5] Guasoni, Paolo
No arbitrage under transaction costs, with fractional Brownian motion and beyond,
Mathematical Finance. An International Journal of Mathematics, Statistics and Financial Economics, Vol. 16, No. 3, pp. 569–582, 2006

[6] Guasoni, Paolo and Rásonyi, Miklós and Schachermayer, Walter
Consistent price systems and face-lifting pricing under transaction costs,
The Annals of Applied Probability, Vol. 18, No. 2, pp. 491–520, 2008

[7] Cheridito, Patrick
Arbitrage in fractional Brownian motion models,
Finance and Stochastics, Vol. 7, No. 4, pp. 533–553, 2003

[8] Dasgupta, A. and Kallianpur, G.
Arbitrage opportunities for a class of Gladyshev processes,
Applied Mathematics and Optimization, Vol. 41, No. 3, pp. 377–385, 2000

[9] Salopek, D. M.
Tolerance to arbitrage,
Stochastic Processes and their Applications, Vol. 76, No. 2, pp. 217–230, 1998

[10] Rogers, L. C. G.
Arbitrage with fractional Brownian motion,
Mathematical Finance. An International Journal of Mathematics, Statistics and Financial Economics, Vol. 7, No. 1, pp. 95–105, 1997

[11] Teverovsky, Vadim and Taqqu, Murad S. and Willinger, Walter
Arbitrage with fractional Brownian motion,
Journal of Statistical Planning and Inference, Vol. 80, No. 1-2, pp. 211–227, 1999

[12] Willinger, Walter and Taqqu, Murad S and Teverovsky, Vadim
Stock market prices and long-range dependence,
Finance and Stochastics, Vol. 3, No. 1, pp. 1–13, 1999

[13] Poterba, James M and Summers, Lawrence H
Mean reversion in stock prices: Evidence and implication,
Journal of Financial Economics, Vol. 22, No. 1, pp. 27–59, 1988

[14] Fama, Eugene F and French, Kenneth R
Permanent and temporary components of stock prices,
Journal of Political Economy, Vol. 96, No. 2, pp. 246–273, 1988

[15] Jacobsen, Ben
Long term dependence in stock returns,
Journal of Empirical Finance, Vol. 3, No. 4, pp. 393–417, 1996

[16] Greene, Myron T and Fielitz, Bruce D
Long-term dependence in common stock returns,
Journal of Financial Economics, Vol. 4, No. 3, pp. 339–349, 1977

[17] Mandelbrot, Benoit B.
When can price be arbitraged efficiently? A limit to the validity of the random walk and martingale modelss,
The Review of Economics and Statistics, Vol. 53, pp. 225–236, 1971

[18] Guasoni, Paolo and Rásonyi, Miklós
Hedging, arbitrage and optimality with superlinear frictions,
The Annals of Applied Probability, Vol. 25, No. 4, pp. 2066–2095, 2015

[19] Lo, Andrew W
Long-Term Memory in Stock Market Price,
Econometrica, Vol. 59, No. 5, pp. 1279–1313, 1991

[20] Giraitis, Liudas and Koul, Hira L. and Surgailis, Donatas
Large sample inference for long memory processes,
Imperial College Press, London, 2012

[21] Carmona, Philippe and Coutin, Laure
Fractional Brownian motion and the Markov property,
Electronic Communications in Probability, Vol. 3, pp. 95–107, 1998

[22] Taqqu, Murad S.
Weak convergence to fractional Brownian motion and to the Rosenblatt process,
Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, Vol. 31, pp. 287–302, 1974/75

[23] Borovkov, Konstantin and Mishura, Yuliya and Novikov, Alexander and Zhitlukhin, Mikhail
Bounds for expected maxima of Gaussian processes and their discrete approximation,
Stochastics, Vol. 89, No. 1, pp. 21–37, 2017

[24] Norros, Ilkka and Valkeila, Esko and Virtamo, Jorma and others
An elementary approach to a Girsanov formula and other analytical results on fractional Brownian motions,
Bernoulli, Vol. 5, No. 4, pp. 571–587, 1999

[25] Cheridito, Patrick and others
Mixed fractional Brownian motion,
Bernoulli, Vol. 7, No. 6, pp. 913–934, 2001

[26] Mishura, Yuliya S.
Stochastic calculus for fractional Brownian motion and related processes,
Lecture Notes in Mathematics, Springer-Verlag, Berlin, Vol. 1929, 2008

[27] Mandelbrot, Benoit B and Van Ness, John W
Fractional Brownian motions, fractional noises and applications,
SIAM review, Vol. 10, No. 4, pp. 422–437, 1968

[28] Sottinen, Tommi
Fractional Brownian motion, random walks and binary market models,
Finance and Stochastics, Vol. 5, No. 3, pp. 343–355, 2001

[29] Guasoni, Paolo and Nika, Zsolt and Rasonyi, Miklos
Trading Fractional Brownian Motion,
SIAM Journal on Financial Mathematics, Vol. 10, No. 3, pp. 769–789, 2019

[30] Norros, Ilkka and Valkeila, Esko and Virtamo, Jorma and others
An elementary approach to a Girsanov formula and other analytical results on fractional Brownian motions,
Bernoulli, Vol. 5, No. 4, pp. 571–587, 1999

[31] Sun, Zhongfeng and Qin, Huizeng
Some Results on the Derivatives of the Gamma and Incomplete Gamma Function for Non-positive Integers,
International Journal of Applied Mathematics, Vol. 47, No. 3, 2017

[32] Guasoni, Paolo and Weber, Marko Hans
Nonlinear price impact and portfolio choice,
Mathematical Finance, Vol. 30, No. 2, pp. 341–376, 2020

Thank you for your attention!