The Journal of Investment Strategies (23–66) Volume 1/Number 2, Spring 2012

Gauge invariance, geometry and arbitrage

Samuel E. Vázquez
Capital Fund Management S.A., 6 Boulevard Haussmann, 75009 Paris, France;
email: samuel.vazquez@cfm.fr

Simone Farinelli
Kraus Partner Investment Solutions AG, Militärstrasse 76, 8004 Zurich,
Switzerland; email: farinelli@krauspartner.com

We identify the most general measure of arbitrage for any market model governed by Ito processes, and, on that basis, we develop dynamic arbitrage strategies. It is shown that our arbitrage measure is invariant under changes of numeraire and equivalent probability measure. Moreover, such a measure has a geometrical interpretation as a gauge connection. The connection has zero curvature if and only if there is no arbitrage. We prove an extension of the martingale pricing theorem in the case of arbitrage. In our case, the present value of any traded asset is given by the expectation of future cashflows discounted by a line integral of the gauge connection. We develop simple dynamic strategies to measure arbitrage using both simulated and real market data. We find that, within our limited data sample, the market is efficient at time horizons of one day or longer. However, we provide strong evidence for nonzero arbitrage in high-frequency intraday data. Such events seem to have a decay time of the order of one minute.

1 INTRODUCTION

The no-arbitrage principle is the cornerstone of modern financial mathematics. Put simply, an arbitrage opportunity allows an agent to make a risk-free profit with zero or negative net investment (see Delbaen and Schachermayer (2008)). Under the no-arbitrage assumption we can assign, in a complete market, a unique price to the derivative of any traded assets using the replicating portfolio method (see Musiela and Rutkowski (2007) or Cvitanic and Zapatero (2004)). The no-arbitrage principle

We would like to extend our gratitude to Eric Weinstein, Pia Malaney, Lee Smolin, Bernd Schroers, Mike Brown, Simone Severini, Jiri Hoogland, Jim Herriot and Bruce Sawhill for many discussions and ideas that influenced the results of this paper. The opinions expressed in this paper are the authors' own and do not represent those of Capital Fund Management S.A. or of Kraus Partner Investment Solutions AG. The risk control principles presented are not necessarily used by Capital Fund Management S.A. or by Kraus Partner Investment Solutions AG. This document does not provide a comprehensive description of concepts and methods used in risk control at Capital Fund Management S.A. or at Kraus Partner Investment Solutions AG.

can also be shown to be equivalent to a weaker form of economic equilibrium (see Cvitanic and Zapatero (2004)) and can therefore be seen as a form of market efficiency (see Fama (1998) and Malkiel (2003)). It is, then, not surprising that most financial and economic literature is based on the no-arbitrage assumption.

Nevertheless, the no-arbitrage principle represents a very strong assumption about market dynamics that must be tested empirically. Even when the market participants use no-arbitrage models, the ultimate price of any security that is traded in a centralized market is set by supply/demand and the complex dynamics of the order book. That is, the market sets the prices. It then makes sense to ask: how efficient are these final prices? In order to answer this question we need a measure of arbitrage.

There is a large body of empirical literature on financial arbitrage. Most of these studies focus on measuring the excess return of particular trading strategies (see, for example, Jegadeesh and Titman (1993) and Gatev *et al* (2006)). Other studies try to find violations of general no-arbitrage relations between option prices (see, for example, Ackert and Tian (1999)). However, there does not seem to be a consensus on whether the reported market “anomalies” are due to arbitrage, or simply to random fluctuations (see Malkiel (2003) and Fama (1998)). Part of the problem is that there seems to be no general measure of arbitrage that can be applied to any traded asset. One of the main goals of this paper is to define such a measure.

The second goal of this paper is to provide a geometrical interpretation of the arbitrage measure. In particular, it was speculated long ago by Ilinski (1997, 2001) and Young (1999) that arbitrage should be viewed as the “curvature” of a gauge connection, in an analogy to some physical theories. The fact that gauge theories are the natural language for describing economics was first proposed by Malaney (1996) and Weinstein (2006) in the context of the economic index problem. The need for such mathematical language can easily be seen from the fact that prices are only relational. More precisely, let $X(t) = (X_1(t), X_2(t), \dots)^\dagger$ be the price vector of all goods in the economy at time t, in some common unit (US dollars, say). Since the measuring units are arbitrary, fundamental economic laws must be invariant under the transformations:

$$X(t) \to \Lambda(t)X(t) \tag{1.1}$$

where $\Lambda(t) > 0$ is a positive stochastic factor.[1] In physics, a (local) transformation such as Equation (1.1) is known as a gauge transformation. These represent a redundancy in our description of the economic system. The laws of the economy should be gauge invariant. The need for a gauge-theoretical approach to economics was highlighted recently by Smolin (2009). The role of gauge invariance in option pricing has been studied in Hoogland and Neumann (1999a,b, 2000) and in Hoogland

[1] For example, $\Lambda(t)$ can be the euro–US dollar exchange rate.

et al (2001). For an unrelated use of differential geometric methods in (no-arbitrage) option pricing, see Labordère (2008). A recently introduced alternative definition of the gauge connection (Farinelli (2011)) can be shown to lead to an arbitrage measure equivalent to ours.

In physics, curvature is a gauge-invariant measure of the path dependency of some physical process. For example, readers familiar with electrodynamics might recall the vector potential A_μ, where $\mu = 0, 1, 2, 3$ label the space and time directions. In differential geometry for theoretical physics, A_μ is known as a gauge connection. Now consider a charged particle that is traveling along some trajectory in space-time $x^\mu(s)$, $s \in [0, 1]$. The interaction of this particle with the gauge potential is proportional to the line integral:

$$\int_\gamma A := \sum_\mu \int_0^1 A_\mu(x(s))\dot{x}^\mu(s)\,\mathrm{d}s$$

Now suppose that we make an infinitesimal change in the path of the particle $\delta x^\mu(s)$, keeping the boundary conditions fixed $\delta x^\mu(0) = \delta x^\nu(1) = 0$. The interaction changes by:

$$\delta\biggl(\int_\gamma A\biggr) = \sum_{\mu,\nu} \int_0^1 \delta x^\mu(s)\dot{x}^\nu(s)F_{\mu\nu}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is known as the curvature of A. Therefore, we see that, for zero curvature $F_{\mu\nu} = 0$, the line integral $\int_\gamma A$ is independent of the path γ. Moreover, note that the curvature is invariant under a gauge transformation of the form $A_\mu \to A_\mu + \partial_\mu \Lambda$, where Λ is any function of space-time.

We find a very similar construction in the case of mathematical finance. In particular, we show that the arbitrage curvature defined in this paper measures the path dependency of the present value of a self-financing portfolio of traded assets with fixed final payoff. The no-arbitrage principle is then equivalent to a zero-curvature condition. By analogy with the electromagnetic curvature F, we expect that any measure of arbitrage should be invariant under the gauge transformation in Equation (1.1). Moreover, the fundamental theorem of asset pricing states that the no-arbitrage principle is equivalent to the existence of a probability measure with respect to which asset prices expressed in terms of a numeraire are martingales (Musiela and Rutkowski (2007)). Therefore, we expect that any measure of arbitrage should also be invariant under a change of probability. These are, in fact, two very important properties that will characterize our arbitrage measure.

This paper takes a "macroscopic" or phenomenological approach to arbitrage. More precisely, we will study arbitrage from the point of view of general stochastic models. We do not address the question of what is causing the arbitrage. Our main goal is to identify the gauge-invariant financial observables that indicate an arbitrage opportunity. Our main assumption is that the prices of all financial instruments can be described by Ito processes. Moreover, we ignore transaction costs.[2]

The organization and main results of the paper are as follows. In Section 2 we present the class of models that we use in the rest of the paper. They are very general and include the case of stocks, bonds and commodities, and more complicated derivative products. We decompose the dynamics of these models in terms of their gauge-transformation properties with respect to Equation (1.1). We identify the gauge invariants and show that they represent an obstruction to the existence of a martingale probability measure. We conclude Section 2 with an example with three assets, and we derive a modified nonlinear Black–Scholes equation with arbitrage.

In Section 3 we give a geometrical interpretation to the gauge-invariant quantities defined in Section 2. Our main goal is to identify the stochastic gauge invariants of Section 2 with the curvature of a gauge connection. We begin with a review of the Malaney–Weinstein connection (Malaney (1996) and Weinstein (2006)), which is done in the context of differentiable economic paths. In Section 3.1 we generalize the Malaney–Weinstein construction to stochastic processes and prove an asset pricing theorem. The main result of this section is that the present value of any self-financing portfolio of traded assets is given by the conditional expectation of future cashflows, discounted by a line integral of the Malaney–Weinstein gauge connection. We show how the value functions of different portfolios replicating the same contingent claim are related to the arbitrage curvature. Readers who are interested only in the arbitrage measure and the detection techniques can skip Section 3.

In Section 4 we develop a simple algorithm to measure the arbitrage curvature using financial data. None of the results of Section 4 require an understanding of the geometry of arbitrage. We explain the main sources of error in such measurement. The algorithm is applicable to any financial instrument. In Section 5 we provide examples with financial data of stock indexes and index futures and construct dynamic arbitrage strategies. We find that, on long timescales, the market is very efficient. However, we provide strong evidence for nonzero-curvature fluctuations at short timescales of the order of one minute. We conclude in Section 6.

2 STOCHASTIC MODELS AND GAUGE INVARIANCE

Let $\mathcal{M} = \{0, 1, 2, \dots, N-1\}$ be the set of all traded securities at any point in time in the market. We will use Greek indexes $\mu, \nu, \dots$ to label members of the set $\mathcal{M}$. We denote the price of security $\mu \in \mathcal{M}$ by X_μ. Our main assumption is that the dynamics

[2] Note that, as pointed out in Shleifer and Vishny (1997), many possible arbitrage opportunities disappear once we take into account market frictions such as transaction costs. Therefore, it is important to keep in mind that, even when we measure a nontrivial curvature in the market, it does not mean that it can always be exploited in a practical trading strategy.

of all such securities is described by Ito processes of the form (see Sondermann (2006) and Shreve (2000)):

$$\mathrm{d}X_\mu := X_\mu\left(\alpha_\mu\,\mathrm{d}t + \sum_a \sigma^a_\mu\,\mathrm{d}W_a\right), \quad \forall \mu \in \mathcal{M} \tag{2.1}$$

where the W_a are standard Brownian motions such that the $W_a(t) - W_a(0)$ are independently and normally distributed random variables with:

$$\mathbb{E}[W_a(t) - W_a(0)] = 0, \qquad \mathrm{cov}[W_a(t) - W_a(0), W_b(t) - W_b(0)] = t\delta_{ab} \tag{2.2}$$

We make no assumptions about the number of Brownian terms, and hence the completeness of the market. The set of Brownian motions $\{W_a\}$ represents all the randomness in the market. Therefore, they induce a natural filtration $\mathcal{F} = (\mathcal{F}_t)_{t\geqslant 0}$ for our probability space. The coefficients α_μ and σ^a_μ can also be stochastic processes adapted[3] to the filtration $\mathcal{F}$. However, they are assumed to satisfy suitable conditions to ensure the existence of the price processes X_μ (see Lamberton and Lapeyre (2007)). This class of models is very general and includes stocks, bonds, options, etc. Moreover, the case of fat tails in the distribution of returns is also included, since this is known to be generated by stochastic volatilities σ^a_μ.

Looking back at Equation (2.1), we can see that the tangent space $\mathrm{d}X_\mu$ has a natural decomposition into the directions that contain all the randomness ($\sum_a \sigma^a_\mu\,\mathrm{d}W_a$) and those orthogonal to it. Therefore, we will make the following decomposition of the drift term in Equation (2.1):

$$\alpha_\mu = \alpha + \sum_a \beta^a \hat{\sigma}^a_\mu + \sum_{A\in\mathcal{N}} \alpha^A J^A_\mu \tag{2.3}$$

where $\mathcal{N}$ is the space spanned by basis vectors $J^A := [J^A_0, \dots, J^A_{N-1}]^\dagger$ such that:

$$\left.\begin{aligned} \sum_\mu J^A_\mu J^B_\mu &= \delta^{AB} \\ \sum_\mu J^A_\mu &= 0 \\ \sum_\mu J^A_\mu \sigma^a_\mu &= 0, \quad \forall a \end{aligned}\right\} \tag{2.4}$$

[3] In simple terms, a process p adapted to $\mathcal{F}$ means that it does not depend on future values of the Brownian motion. In other words, $p(t)$ can only depend on $\{W_a(s)\}$ up to time $s \leqslant t$.

and:

$$
\left.
\begin{aligned}
\alpha &:= \frac{1}{N}\sum_{\mu}\alpha_{\mu} \\
\alpha^{A} &:= \sum_{\mu} J_{\mu}^{A}\alpha_{\mu} \\
\hat{\sigma}_{\mu}^{a} &:= \sigma_{\mu}^{a} - \sigma^{a} \\
\sigma^{a} &:= \frac{1}{N}\sum_{\mu}\sigma_{\mu}^{a}
\end{aligned}
\right\} \tag{2.5}
$$

We will refer to $\mathcal{N}$ as the null space of the market. Note that this space is orthogonal to all the randomness in the tangent space $\mathrm{d}X_{\mu}$. However, we need to remember that $\mathcal{N}$ might be trivial. The definition of β^{a} in Equation (2.3) is not unique if the vectors $\hat{\sigma}_{\mu}^{a} := [\hat{\sigma}_{0}^{a}, \dots, \hat{\sigma}_{N-1}^{a}]^{\dagger}$ are linearly dependent. This is the case of, for example, an incomplete market with more Brownian motions than traded securities. Moreover, α^{A} is unique up to rotations in the null space. As we will see, the quantities α^{A} are the unique gauge-invariant measures of arbitrage. The two main goals of this paper are to give a geometric interpretation to the parameters α^{A}, and to set up a procedure to measure them using financial data.

Since prices are relative and only reflect an exchange rate between two products, the units used to measure X_{μ} are arbitrary. Therefore, the dynamics of the market must be invariant under a change of measuring units. In mathematical finance, this is known as a change of numeraire Musiela and Rutkowski (2007), and it can be interpreted as a gauge transformation:

$$X_{\mu}(t) \to \Lambda(t)X_{\mu}(t) \tag{2.6}$$

where Λ is a positive stochastic process that is adapted to the filtration $\mathcal{F}$. Another symmetry, which is special to the particular models of Equation (2.1), is a transformation of the probability measure. This is not really a gauge symmetry, but corresponds rather to a change of variables of the form:

$$W_{a}(t) \to W_{a}(t) + \int^{t} \delta\beta^{a}(s)\,\mathrm{d}s$$

Our next task is to study the transformation properties of the different terms in Equations (2.1) and (2.3). The following result follows.

Proposition 2.1 *Consider a change of numeraire of the form $X_{\mu} \to \Lambda X_{\mu}$, where Λ is a positive stochastic process adapted to the filtration $\mathcal{F}$ and:*

$$\mathrm{d}\Lambda := \Lambda\left(\delta\alpha\,\mathrm{d}t + \sum_{a}\delta\sigma^{a}\,\mathrm{d}W_{a}\right)$$

Then, the coefficients of the Ito processes, Equations (2.1) and (2.3), transform as:

$$\left.\begin{aligned} \alpha &\to \alpha + \delta\alpha + \sum_a \sigma^a \delta\sigma^a \\ \beta^a &\to \beta^a + \delta\sigma^a \\ \sigma^a &\to \sigma^a + \delta\sigma^a \end{aligned}\right\} \tag{2.7}$$

Finally, under a transformation of the probability measure given by the Radon–Nykodým derivative:

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}^*} = \exp\left(-\int^t \sum_a \delta\beta^a \,\mathrm{d}W_a - \frac{1}{2}\int^t \sum_a (\delta\beta^a(s))^2 \,\mathrm{d}s\right)$$

we have the mapping of standard Brownian motions:

$$W_a(t) \to W_a(t) + \int^t \delta\beta^a(s)\,\mathrm{d}s \tag{2.8}$$

and:

$$\alpha \to \alpha + \sum_a \sigma^a \delta\beta^a, \qquad \beta^a \to \beta^a + \delta\beta^a \tag{2.9}$$

In particular, it follows that $\hat{\sigma}^a_\mu$, α^A *and* J^A_μ *are invariant under such transformations.*

PROOF The result in Equation (2.7) above follows from a simple application of Ito's rule to the product $X'_\mu := \Lambda X_\mu$:

$$\begin{aligned} \mathrm{d}X'_\mu &= \mathrm{d}\Lambda X_\mu + \Lambda\,\mathrm{d}X_\mu + \mathrm{d}\langle \Lambda, X_\mu\rangle \\ &= X'_\mu\Bigg[\left(\alpha + \delta\alpha + \sum_a \sigma^a\delta\sigma^a + \sum_a(\beta^a + \delta\sigma^a)\hat{\sigma}^a_\mu + \sum_{A\in\mathcal{N}}\alpha^A J^A_\mu\right)\mathrm{d}t \\ &\qquad + \sum_a(\hat{\sigma}^a_\mu + \sigma^a + \delta\sigma^a)\,\mathrm{d}W_a\Bigg] \end{aligned} \tag{2.10}$$

where $\mathrm{d}\langle \Lambda, X_\mu\rangle = \mathrm{d}t\,\Lambda X_\mu \sum_a \delta\sigma^a\sigma^a_\mu$ is the differential of the quadratic variation. The transformation in Equation (2.9) follows from a simple differentiation of W_a in Equation (2.1):

$$\begin{aligned} \mathrm{d}X_\mu = X_\mu\Bigg[&\left(\alpha + \sum_a \sigma^a\delta\beta^a + \sum_a(\beta^a + \delta\beta^a)\hat{\sigma}^a_\mu + \sum_{A\in\mathcal{N}}\alpha^A J^A_\mu\right)\mathrm{d}t \\ &+ \sum_a(\hat{\sigma}^a_\mu + \sigma^a)\,\mathrm{d}W^*_a\Bigg] \end{aligned} \tag{2.11}$$

where we defined:

$$W_a(t) = W_a^*(t) + \int^t \delta\beta^a(s)\,\mathrm{d}s$$

Note that both α^A and J^A are unchanged by these gauge transformations. In particular, suppose that:

$$\sum_\mu J_\mu^A \sigma_\mu^a = 0$$

Then, it follows from Equation (2.4) that:

$$\sum_\mu J_\mu^A(\sigma_\mu^a + \delta\sigma^a) = \sum_\mu J_\mu^A \sigma_\mu^a = 0$$

□

So far we have taken the existence of the basis vectors J^A for granted. A constructive procedure to find such a basis, if nontrivial, is given by the following proposition.

PROPOSITION 2.2 *Let Ω be the symmetric and real $N \times N$ matrix with component:*

$$\Omega_{\mu\nu} = \sum_a \sigma_\mu^a \sigma_\nu^a, \quad \textit{where } N = \dim(\mathcal{M})$$

Moreover, define U as the matrix of all ones, eg, $U_{\mu\nu} = 1$, $\forall \mu, \nu \in \mathcal{M}$. Then, the matrix G defined as:

$$G = \Omega - \frac{1}{N}(U\Omega + \Omega U) + \frac{1}{N^2}\operatorname{Tr}(U\Omega)U \tag{2.12}$$

is gauge invariant. Let $\mathcal{N}_G$ be the null space of matrix G such that $\sum_\mu J_\mu = 0$ for any nontrivial $J \in \mathcal{N}_G$. Then $\mathcal{N}_G = \mathcal{N}$. In particular, the space $\mathcal{N}_G$ is spanned by the orthonormal zero modes of G that are orthogonal to the vector $J = (1, 1, \dots, 1)^\dagger$.

PROOF First we need to prove that the space of vectors J such that $J^\dagger \sigma^a = 0$, $\forall a$, is in one-to-one correspondence with the zero modes of Ω: $\Omega J = 0$. Obviously, if $J^\dagger \sigma^a = 0$, it follows that J is also a zero mode of Ω. To prove the converse, suppose that $\Omega J = 0$, but $J^\dagger \sigma^a = \lambda^a$, where $\lambda^a \neq 0$ for at least one value of a. Then:

$$0 = J^\dagger \Omega J = \sum_a (\lambda^a)^2$$

which can only be true if $\lambda^a = 0$, $\forall a$.

Now we turn our attention to the matrix G, defined in Equation (2.12). Using the gauge transformation $\sigma_\mu^a \to \sigma_\mu^a + \delta\sigma^a$, we can see that Ω transforms as:

$$\Omega_{\mu\nu} \to \Omega_{\mu\nu} + \sum_a \delta\sigma^a(\sigma_\mu^a + \sigma_\nu^a) + \sum_a (\delta\sigma^a)^2 \tag{2.13}$$

It is then straightforward to verify the gauge invariance of the matrix G. Next we recall that the space $\mathcal{N}$ is spanned by (nontrivial) orthonormal zero modes of Ω such that they also satisfy $\sum_\mu J_\mu = 0$. One can define a similar space $\mathcal{N}_G$ for G. It is easy to verify that any vector $J \in \mathcal{N}$ is also a vector in $\mathcal{N}_G$. On the other hand, for any vector $J' \in \mathcal{N}_G$, it follows from Equation (2.12) that $\Omega J' = 0$. Thus, we have proven that $\mathcal{N} = \mathcal{N}_G$.

It is easy to verify that $\sum_\mu G_{\mu\nu} = 0$. Therefore, the vector $J = (1, 1, \dots, 1)^\dagger$ is a particular zero mode of G. Now take any other zero mode of G, and call it J', which is orthogonal to J. It follows that $0 = J^\dagger J' = \sum_\mu J'_\mu$. Therefore, $J' \in \mathcal{N}_G = \mathcal{N}$. This completes the proof. □

So far we have talked about the full set of securities of the market. However, it is clear that the decomposition in Equation (2.3) can be done for any subset of the market. That is, suppose we observe a subset of the prices X_i, $i \in \mathcal{S} \subset \mathcal{M}$. Moreover, suppose that, within this subset, we can still find some zero modes J^A obeying:

$$\sum_i J_i^A \sigma_i^a = 0, \ \forall a, \qquad \sum_i J_i^A = 0$$

We can then easily lift these vectors to the full set $\mathcal{M}$ by taking $J_{\mathcal{M}}^A = (J^A, \mathbf{0})$. This represents a particular choice of basis in the null space $\mathcal{N}$. By observing a subsector of the market, we will only have access to some of the components of α^A. For notational convenience, in the following we will not distinguish between the full market and a subset of it.

In a next step we want to link α^A with the no-arbitrage condition. But what do we mean exactly by "no-arbitrage condition"? As a matter of fact there exist two similar conditions, the no-arbitrage condition and the no-free-lunch-with-vanishing-risk condition (NFLVR) (see Delbaen and Schachermayer (2008)), which, in discrete time, are equivalent. In continuous time NFLVR is the stronger condition and is equivalent to the existence of a martingale measure for the (discounted) asset prices (see Delbaen and Schachermayer (2008, Chapter 9.4)).

Under the NFLVR assumption (see Delbaen and Schachermayer (2008) and Hunt and Kennedy (2004)), it is always possible to find a common positive discount factor Λ and an equivalent probability measure $\mathbb{P}^* \sim \mathbb{P}$ such that the discounted prices ΛX_μ are martingales $\Lambda(t) X_\mu(t) = \mathbb{E}_t^*[\Lambda(T) X_\mu(T)]$, where $t \leqslant T$. This is known as the martingale representation theorem (see Sondermann (2006) and Shreve (2000)). In our language, this means that there is a gauge transformation mapping X_μ to $\mathbb{P}^*$-martingales. In other words, if there is no arbitrage, price processes are gauge equivalent to $\mathbb{P}^*$-martingales for some probability measure $\mathbb{P}^*$. The result of the martingale representation theorem can only be obtained if one is able to write:

$$\int_t^T \mathrm{d}(\Lambda X_\mu) := \int_t^T \gamma_\mu^a \,\mathrm{d}W_a^*$$

for some adapted process γ_μ^a. The reason is that the stochastic integral $\int_t^T \gamma_\mu^a \,\mathrm{d}W_a^*$ is a martingale:

$$\mathbb{E}_t^*\biggl[\int_t^T \gamma_\mu^a \,\mathrm{d}W_a^*\biggr] = 0$$

By Proposition 2.2, there is neither a change of probability nor a choice of a positive discount factor for which the vector $\sum_{A\in\mathcal{N}} \alpha^A J^A$ is mapped to 0 (in contrast to α and all β^a, which can indeed be made to vanish). Therefore, it is easy to see that the term $\sum_{A\in\mathcal{N}} \alpha^A J^A$ parameterizes the obstruction to the existence of a martingale probability measure for any discounted price process ΛX_μ. Its vanishing is a necessary but not sufficient condition for the existence of an equivalent martingale measure.

As α^A are gauge-invariant quantities, one expects that they should be observables. In the next section we will relate this quantity to a gauge connection and its curvature. In Section 4 we show that such a quantity can indeed be observed, and we explain simple strategies to measure it. Before concluding this section, it is instructive to study a particular example with three assets.

2.1 An example

Consider the case of three assets X_μ, $\mu = 0, 1, 2$, where X_0 is a savings account and X_1, X_2 are some other risky assets. All prices are measured in the same common units. We will assume only one Brownian motion. Therefore, the dynamics of the prices are described by:

$$\left.\begin{aligned} \mathrm{d}X_0 &= rX_0 \,\mathrm{d}t \\ \mathrm{d}X_i &= X_i[\alpha_i \,\mathrm{d}t + \sigma_i \,\mathrm{d}W], \quad i = 1, 2 \end{aligned}\right\} \tag{2.14}$$

For later convenience, we assume that the interest rate r is deterministic. In order to carry out the decomposition in Equation (2.3) we need to find a basis for the null space $\mathcal{N}$. In this case, since there is only one Brownian motion and two risky assets, there will be only one null direction. To calculate it, we start by identifying the Ω matrix:

$$\Omega = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \sigma_1^2 & \sigma_1\sigma_2 \\ 0 & \sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \tag{2.15}$$

We can now construct the G matrix using Equation (2.12). The explicit form of G is not very illuminating. The unnormalized eigenvectors of G are found to be:

$$V_1 = \begin{pmatrix} 2 - \dfrac{3\sigma_1}{\sigma_1 + \sigma_2} \\ 0 \\ 0 \end{pmatrix}, \qquad V_2 = \begin{pmatrix} -1 + \dfrac{3\sigma_1}{\sigma_1 + \sigma_2} \\ 1 \\ 0 \end{pmatrix}, \qquad V_3 = \begin{pmatrix} \dfrac{\sigma_1 + \sigma_2}{\sigma_1 - 2\sigma_2} \\ \dfrac{\sigma_2 - 2\sigma_1}{\sigma_1 - 2\sigma_2} \\ 1 \end{pmatrix} \tag{2.16}$$

where:

$$\left.\begin{aligned} GV_1 &= 0 \\ GV_2 &= 0 \\ GV_3 &= \tfrac{2}{3}(\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2)V_3 \end{aligned}\right\} \tag{2.17}$$

In order to find a basis for the null space $\mathcal{N}$ defined in Equation (2.4), we need to project V_1 or V_2 into the space orthogonal to the vector $J^0 = (1, 1, \dots, 1)^\dagger$. To do this, we define the projection matrix:

$$P_U := \tfrac{1}{3}U, \qquad P_U^2 = P_U \tag{2.18}$$

where U is the 3×3 all-ones matrix. Note that $\mathbf{1} - P_U$ projects into the space orthogonal to J^0. Our choice for the normalized null vector is then:

$$J = \frac{(\mathbf{1} - P_U)V_1}{\sqrt{[(\mathbf{1} - P_U)V_1]^\dagger(\mathbf{1} - P_U)V_1}} = \frac{1}{\sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}} \begin{pmatrix} \sigma_1 - \sigma_2 \\ \sigma_2 \\ -\sigma_1 \end{pmatrix} \tag{2.19}$$

It is easy to verify that J obeys the properties given in Equation (2.4).

We can now go back to the decomposition given in Equation (2.3). Using Equation (2.14), we find:

$$\alpha = r - \beta\hat{\sigma}_0 - \tilde{\alpha}J_0, \quad \sigma_0 = 0 \tag{2.20}$$

where $\hat{\sigma}_\mu = \sigma_\mu - \frac{1}{3}\sum_{\nu=0}^{2}\sigma_\nu$, and $\tilde{\alpha}$ is the arbitrage vector α^A, which in this case has only one component $\alpha^1 := \tilde{\alpha}$. Therefore, inserting Equation (2.20) into Equation (2.14), we can write the evolution equations as:

$$\mathrm{d}X_0 = rX_0\,\mathrm{d}t \tag{2.21}$$

$$\mathrm{d}X_1 = X_1\biggl[\biggl(r + \beta\sigma_1 + \tilde{\alpha}\frac{2\sigma_2 - \sigma_1}{\sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}}\biggr)\mathrm{d}t + \sigma_1\,\mathrm{d}W\biggr] \tag{2.22}$$

$$\mathrm{d}X_2 = X_2\biggl[\biggl(r + \beta\sigma_2 + \tilde{\alpha}\frac{\sigma_2 - 2\sigma_1}{\sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}}\biggr)\mathrm{d}t + \sigma_2\,\mathrm{d}W\biggr] \tag{2.23}$$

For $\tilde{\alpha} = 0$, Equations (2.21)–(2.23) reduce to the familiar no-arbitrage Black–Scholes dynamics. As usual, β is interpreted as the market price of risk. Note that, in this example, both risky assets are exposed to the same market risk factor W. The volatility σ_i measures the coupling to such risk. Under the no-arbitrage assumption, both assets should give the same expected return per unit of risk. This is β. However, we see that if $\tilde{\alpha} \neq 0$, X_1 and X_2 have different expected returns, even when they are exposed to the same risk. This discloses an arbitrage opportunity.

There is a very interesting consequence of Equations (2.21)–(2.23) when X_2 is any function of X_1 (eg, an option). For simplicity, consider the case where the only time dependence in $\tilde{\alpha}$ is of the form $\tilde{\alpha} = \tilde{\alpha}(t, X_1)$, where $\tilde{\alpha}(t, X_1)$ is a differentiable function of t and X_1. Moreover, the interest rate r is assumed to be deterministic. In this case we can derive a nonlinear version of the Black–Scholes equation with arbitrage. For ease of notation, let $X_1 := X$. Under our assumptions we will have that $X_2 = V(t, X)$. Then, using Ito's rule, we find:

$$\begin{aligned} \mathrm{d}V &= \partial_t V\,\mathrm{d}t + \partial_X V\,\mathrm{d}X + \tfrac{1}{2}\partial_X^2 V\,\mathrm{d}\langle X\rangle \\ &= V(\alpha_2\,\mathrm{d}t + \sigma_2\,\mathrm{d}W) \end{aligned} \tag{2.24}$$

where we identify:

$$\left.\begin{aligned} \alpha_2 &= \frac{\partial_t V}{V} + \alpha_1 X \frac{\partial_X V}{V} + \tfrac{1}{2}\sigma_1^2 X^2 \frac{\partial_X^2 V}{V} \\ \sigma_2 &= \sigma_1 X \frac{\partial_X V}{V} \end{aligned}\right\} \tag{2.25}$$

Comparing Equation (2.25) with Equation (2.23), we find that:

$$\begin{aligned} \alpha_2 &= r + \beta\sigma_2 + \tilde{\alpha}\frac{\sigma_2 - 2\sigma_1}{\sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}} \\ &= \frac{\partial_t V}{V} + \alpha_1 X \frac{\partial_X V}{V} + \tfrac{1}{2}\sigma_1^2 X^2 \frac{\partial_X^2 V}{V} \end{aligned} \tag{2.26}$$

where, from Equation (2.22):

$$\alpha_1 = r + \beta\sigma_1 + \tilde{\alpha}\frac{2\sigma_2 - \sigma_1}{\sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}} \tag{2.27}$$

Therefore, after some algebra, Equation (2.26) becomes a modified nonlinear Black–Scholes partial differential equation:

$$\begin{aligned} &\partial_t V + rX\partial_X V + \tfrac{1}{2}\sigma_1^2 X^2 \partial_X^2 V \\ &\qquad + \left(\sqrt{2}\tilde{\alpha}\left[1 + \frac{X\partial_X V}{V}\left(\frac{X\partial_X V}{V} - 1\right)\right]^{1/2} - r\right)V = 0 \end{aligned} \tag{2.28}$$

Note that, for $\tilde{\alpha} = 0$, this reduces to the familiar Black–Scholes equation. The nonlinear Black–Scholes equation is a special case of the more general pricing theorem presented in Section 3.

It is important to remember that the arbitrage parameter $\tilde{\alpha}$ in Equation (2.28) can in general depend on time and the stock price. Therefore, in principle, almost any deformation of the option price is possible. It follows that Equation (2.28) can be solved only if the arbitrage dynamics are known. For example, consider the case where we set:

$$\tilde{\alpha} := \frac{1}{2^{3/2}}(\tilde{\sigma}_1^2 - \sigma_1^2)\frac{X^2 \partial_X^2 V}{V[1 + (X\partial_X V/V)((X\partial_X V/V) - 1)]^{1/2}} \tag{2.29}$$

for some constant $\tilde{\sigma}_1^2$. Then, the option price obeys the usual Black–Scholes equation but with the "wrong" volatility:

$$\partial_t V + r(X\partial_X V - V) + \tfrac{1}{2}\tilde{\sigma}_1^2 X^2 \partial_X^2 V = 0 \tag{2.30}$$

This is a simple example of the well-known volatility arbitrage.

3 THE GAUGE CONNECTION

The application of differential geometric ideas in economics can be traced to the work of Malaney (1996) and Weinstein (2006). It was found that the solution to the apparent discrepancy among different economic growth indexes could be solved by the appropriate choice of a covariant derivative. Such a derivative has the property that a self-financing basket of goods is seen as "constant". More technically, a self-financing basket is interpreted as being "parallel transported" along a one-dimensional curve in the base manifold spanned by prices and portfolio nominals. Then there is a natural geometric index to measure the growth of such basket, which was shown to be identical to the so-called Divisa index. It is very illuminating to review this construction to gain intuition about the relation between arbitrage and curvature. In what follows, all quantities are assumed to be deterministic and differentiable. We will return to the stochastic case in the next subsection.

A covariant derivative induces a connection one-form in the base space (for the differential geometric background see Kobayashi and Nomizu (1996) and Bleecker (1981)). In Malaney (1996) and Weinstein (2006), this connection is given by:

$$A = \frac{\sum_\mu \phi_\mu \,\mathrm{d}X_\mu}{\sum_\nu \phi_\nu X_\nu} \tag{3.1}$$

where ϕ_μ are the portfolio nominals, $V = \sum_\mu \phi_\mu X_\mu$, and the base space is parameterized by the coordinates (t, ϕ_μ, X_μ). Note that under a change of numeraire

$X_\mu \to \Lambda(X)X_\mu$, the connection transforms as:

$$A \to A + \mathrm{d}\Lambda \tag{3.2}$$

This is the analog of the transformation rule of the vector potential in electrodynamics.

A self-financing portfolio can be seen as being parallel transported with the connection A as:

$$\nabla_{\dot\gamma} V = (d - A)V|_{\dot\gamma} = 0 \tag{3.3}$$

where $\nabla_{\dot\gamma}$ is the covariant derivative along the trajectory γ. The solution to this equation is simply:

$$\frac{V(T)}{V(t)} = \exp\left\{\int_\gamma A\right\} := D_\gamma \tag{3.4}$$

where γ is a particular self-financing trajectory $(s, \phi(s), X(s))$, $s \in [t, T]$, and D_γ is known as the Divisa index.

The dependence of D_γ on the choice of curve γ is parameterized by the curvature of the gauge connection, which is given by:

$$R = \mathrm{d}A = \frac{1}{(\sum_\mu \phi_\mu X_\mu)^2}\sum_{\nu,\sigma}(\phi_\nu X_\nu \,\mathrm{d}\phi_\sigma \wedge \mathrm{d}X_\sigma - \phi_\sigma X_\nu \,\mathrm{d}\phi_\nu \wedge \mathrm{d}X_\sigma) \tag{3.5}$$

Note that the curvature is invariant under a gauge transformation, as $\mathrm{d}(A + \mathrm{d}\Lambda) = \mathrm{d}A$. In the approximation where economic agents are price takers, the price trajectory $X(t)$ is given exogenously, and we are only allowed to make changes in the portfolio nominals ϕ. In other words, we can write $\mathrm{d}X_\mu = \dot X_\mu \,\mathrm{d}t$ in Equation (3.5). We can then restrict the curvature to the submanifold corresponding to the (t, ϕ_μ) coordinates. The induced curvature in this submanifold is given by:

$$\begin{aligned} R &= \frac{1}{(\sum_\mu \phi_\mu X_\mu)^2}\sum_{\nu,\sigma}\phi_\sigma X_\nu X_\sigma\left(\frac{\dot X_\nu}{X_\nu} - \frac{\dot X_\sigma}{X_\sigma}\right)\mathrm{d}\phi_\nu \wedge \mathrm{d}t \\ &:= \sum_\mu R_{\mu,t}\,\mathrm{d}\phi_\mu \wedge \mathrm{d}t \end{aligned} \tag{3.6}$$

In this case, the path dependency of the Divisa index, Equation (3.4), can be written as:

$$\delta_\gamma \log D_\gamma = \sum_\mu \int_t^T \mathrm{d}s\, R_{\mu,t}(s)\delta\phi_\mu(s) \tag{3.7}$$

where δ_γ represents a variation to the trajectory of the portfolio nominals. Therefore, we see that Equation (3.4) is independent of the path γ only if the price trajectories obey the zero-curvature condition:

$$R_{\mu,t} = 0 \implies \dot X_\mu(t) = \alpha(t)X_\mu(t), \quad \forall\mu$$

The zero-curvature condition implies that the prices of all securities evolve by the same common inflation factor.

The relation between curvature and arbitrage goes as follows. Suppose that the prices obey the zero-curvature condition given above. It follows that, for any self-financing portfolio, we have:

$$\begin{aligned} V(T) &= V(t) \exp\left\{ \int_\gamma A \right\} \\ &= V(t) \exp\left\{ \int_t^T \alpha(s)\, \mathrm{d}s \right\} \end{aligned} \tag{3.8}$$

for $T \geqslant t$. In particular, if $V(t) = 0$, it follows that $V(T) = 0$. Therefore, it is not possible to make wealth without a positive initial investment. On the other hand, suppose that the curvature is not zero. Consider two portfolio trajectories γ_1 and γ_2 such that, say, $D_{\gamma_1} > D_{\gamma_2}$ at some time $T \geqslant t$, for the same initial wealth $V_{\gamma_1}(t) = V_{\gamma_2}(t) > 0$. Now construct the difference portfolio with nominals $\phi := \phi_1 - \phi_2$ and wealth function:

$$V = V_{\gamma_1} - V_{\gamma_2} \tag{3.9}$$

Then, at time $T \geqslant t$ we have:

$$V(T) = (D_{\gamma_1} - D_{\gamma_2}) V_{\gamma_1}(t) > 0 \tag{3.10}$$

while $V(t) = 0$. In other words, we have made wealth out of nothing. In the next section we show how this construction carries over to the stochastic case.

3.1 The stochastic gauge connection

In the previous section we illustrated the relation between curvature, path dependency and arbitrage, using the Malaney–Weinstein connection. However, this construction only works for differentiable economic trajectories in the base space (ϕ, X). Nevertheless, we have found a direct analog of the Malaney–Weinstein connection for Ito processes, which we summarize in the following theorem. In order to avoid technical complications, we restrict our attention to an economy on a finite interval of time $t \in [0, T]$.

THEOREM 3.1 *Consider any self-financing portfolio $V = \sum_\mu \phi_\mu X_\mu$, so that:*

$$\mathrm{d}V = \sum_\mu \phi_\mu \, \mathrm{d}X_\mu$$

Then, if there exists a change of measure satisfying the Novikov condition, then there exists a (nonunique) equivalent probability measure $\mathbb{P}^* \sim \mathbb{P}$ under which the price *processes obey:*

$$\mathrm{d}X_\mu = X_\mu\left[\left(\alpha^* + \sum_A \alpha^A J_\mu^A\right)\mathrm{d}t + \sum_a \sigma_\mu^a\,\mathrm{d}W_a^*\right] \tag{3.11}$$

Moreover, the present value of $V(t)$ given some final payoff $V(T)$, $T \geqslant t$, is given by:

$$V_\gamma(t) = \mathbb{E}_t^*\left[V(T)\exp\left\{-\int_\gamma \Gamma\right\}\right] \tag{3.12}$$

where γ is some self-financing trajectory, and Γ is given by the expectation of the Malaney–Weinstein connection:

$$\Gamma = \mathbb{E}_t^*\left[\frac{\sum_\mu \phi_\mu\,\mathrm{d}X_\mu}{\sum_\nu \phi_\nu X_\nu}\right] = \frac{\sum_{\mu,A\in\mathcal{N}} \alpha^A J_\mu^A \phi_\mu X_\mu}{\sum_\nu \phi_\nu X_\nu}\,\mathrm{d}t + \alpha^*\,\mathrm{d}t \tag{3.13}$$

Finally, the path dependency of the present value of the portfolio, with fixed final payoff, is parameterized by:

$$\delta V_\gamma(t) = -\sum_\mu \int_t^T \mathrm{d}s\,\mathbb{E}_t^*\left[V(T)\exp\left\{-\int_\gamma \Gamma\right\}\delta\phi_\mu(s) R_{\mu,t}(s)\right] \tag{3.14}$$

where $R_{\mu,t}$ are the components of the curvature two-form defined in the reduced base space (t,ϕ):

$$\begin{aligned} R &= \mathrm{d}\Gamma \\ &= \frac{1}{(\sum_\mu \phi_\mu X_\mu)^2}\sum_{\nu,\sigma,A\in\mathcal{N}} \alpha^A X_\nu X_\sigma \phi_\sigma (J_\nu^A - J_\sigma^A)\,\mathrm{d}\phi_\nu \wedge \mathrm{d}t \\ &:= \sum_\mu R_{\mu,t}\,\mathrm{d}\phi_\mu \wedge \mathrm{d}t \end{aligned} \tag{3.15}$$

PROOF We start by writing the portfolio return as:

$$\mathrm{d}V = \sum_\mu \phi_\mu\,\mathrm{d}X_\mu := V\left(a\,\mathrm{d}t + \sum_a b^a\,\mathrm{d}W_a\right) \tag{3.16}$$

where:

$$a = \frac{\sum_\mu \alpha_\mu \phi_\mu X_\mu}{\sum_\nu \phi_\nu X_\nu}, \qquad b^a = \frac{\sum_\mu \sigma_\mu^a \phi_\mu X_\mu}{\sum_\nu \phi_\nu X_\nu} \tag{3.17}$$

Now consider the combination $V' := \Lambda V$, where we take (see Equation (2.1)):

$$\mathrm{d}\Lambda = \Lambda\left[\left(-a + \sum_a b^a \beta^a\right)\mathrm{d}t - \sum_a \beta^a\,\mathrm{d}W_a\right] \tag{3.18}$$

A simple application of Ito's rule gives:

$$dV' = V' \sum_a (b^a - \beta^a)\, dW_a \tag{3.19}$$

It is well-known that any stochastic integral of the form $\int_0^t \gamma\, dW_a$ is a martingale (see Sondermann (2006) and Shreve (2000)). Therefore, we have:

$$V(t) = \mathbb{E}_t\biggl[V(T) \exp\biggl\{ \int_t^T d\log\Lambda \biggr\} \biggr] \tag{3.20}$$

A further application of Ito's rule gives:

$$d\,\log\Lambda = -\Gamma - \tfrac{1}{2}\sum_a (\beta^a)^2\, dt - \sum_a \beta^a\, dW_a \tag{3.21}$$

where Γ is defined in Equation (3.13), with $\alpha^* := \alpha - \sum_a \beta^a \sigma^a$.

Now consider making a change of probability measure such that:

$$W_a := W_a^* - \int^t \beta^a(s)\, ds$$

It is easy to see that, under $\mathbb{P}^*$, the price processes will obey Equation (3.11) of the theorem. Moreover, the Radon–Nykodým derivative is given by:

$$\frac{d\mathbb{P}}{d\mathbb{P}^*} = \exp\biggl[-\tfrac{1}{2}\sum_a \int_t^T (\beta^a(s))^2\, ds + \sum_a \int_t^T \beta^a\, dW_a^* \biggr] \tag{3.22}$$

This Radon–Nykodým derivative is a martingale if the Novikov condition:

$$\mathbb{E}\biggl[\exp\biggl(\int_0^T \frac{1}{2}\biggl(\sum_a (\beta_a)^2 \biggr) ds \biggr) \biggr] < +\infty \tag{3.23}$$

is satisfied. Therefore, using Equations (3.21) and (3.22) in (3.20), we obtain:

$$\begin{aligned} V(t) &= \mathbb{E}_t\biggl[V(T) \exp\biggl\{ -\int_\gamma \Gamma - \tfrac{1}{2}\sum_a \int_t^T (\beta^a)^2\, dt - \sum_a \int_t^T \beta^a\, dW_a \biggr\} \biggr] \\ &= \mathbb{E}_t^*\biggl[V(T) \frac{d\mathbb{P}}{d\mathbb{P}^*} \exp\biggl\{ -\int_\gamma \Gamma + \tfrac{1}{2}\sum_a \int_t^T (\beta^a)^2\, dt - \sum_a \int_t^T \beta^a\, dW_a^* \biggr\} \biggr] \\ &= \mathbb{E}_t^*\biggl[V(T) \exp\biggl\{ -\int_\gamma \Gamma \biggr\} \biggr] \end{aligned} \tag{3.24}$$

In order to prove that Γ can be written as an expectation of the Malaney–Weinstein connection, we recall that:

$$
\begin{aligned}
\Gamma &= \sum_{\mu} \frac{\phi_\mu(t) X_\mu(t)}{\sum_\nu \phi_\nu(t) X_\nu(t)} \left(\alpha^* + \sum_{A \in \mathcal{N}} \alpha^A J_\mu^A \right) \mathrm{d}t \\
&= \lim_{\delta t \to 0} \sum_{\mu} \frac{\phi_\mu(t)}{\sum_\nu \phi_\nu(t) X_\nu(t)} \left(\frac{\mathbb{E}_t^*[X_\mu(t+\delta t)] - X_\mu(t)}{\delta t} \right) \mathrm{d}t \\
&= \lim_{\delta t \to 0} \mathbb{E}_t^* \left[\sum_{\mu} \frac{\phi_\mu(t)}{\sum_\nu \phi_\nu(t) X_\nu(t)} \left(\frac{X_\mu(t+\delta t) - X_\mu(t)}{\delta t} \right) \mathrm{d}t \right] \\
&= \mathbb{E}_t^* \left[\frac{\sum_\mu \phi_\mu \, \mathrm{d}X_\mu}{\sum_\nu \phi_\nu X_\nu} \right] \qquad (3.25)
\end{aligned}
$$

The last result of the theorem, Equation (3.14), follows simply by making a small change in the portfolio nominals, and keeping the boundary conditions on V fixed. □

Note that the curvature of Γ is zero if and only if $\alpha^A = 0$, which implies, together with the Novikov condition, the no-arbitrage condition. Moreover, the probability measure $\mathbb{P}^*$ might not be unique, as the choice of β^a in general is not. This also implies that α^* is not unique in general.

A special case of a self-financing portfolio is a portfolio containing just one base asset.

COROLLARY 3.2 *Under the same assumptions as Theorem 3.1, for all assets in the market model $\mu \in \mathcal{M}$:*

$$
X_\mu(t) = \mathbb{E}_t^* \left[X_\mu(T) \exp\left\{ -\int_t^T \left(\alpha^* + \sum_A \alpha^A J_\mu^A \right) \mathrm{d}t' \right\} \right] \qquad (3.26)
$$

In particular, under the no-arbitrage assumption $\alpha^A = 0$, we recover the classic martingale pricing theorem:

$$
X_\mu(t) = \mathbb{E}_t^* \left[X_\mu(T) \exp\left\{ -\int_t^T \alpha^* \, \mathrm{d}t' \right\} \right] \qquad (3.27)
$$

In Section 2.1 we derived a modified Black–Scholes equation for the case of three assets. Now we can use the result of Corollary 3.2 to prove a generalization of such an equation. Consider the following vector of assets:

$$
X = [X_0, X_1, \dots, X_n, \Phi_1(\boldsymbol{X}, t), \dots, \Phi_m(\boldsymbol{X}, t)]^\dagger \qquad (3.28)
$$

We will label the components of this vector by X_μ, $\mu = 0, 1, \dots, n+m$. Moreover, we assume that the Φ_i are smooth functions of the vector of underlying prices, $\boldsymbol{X} :=$

$[X_1, \dots, X_n]^\dagger$, and $\mathrm{d}X_0 = X_0 r \,\mathrm{d}t$ describes a savings account with deterministic interest rate r. The functions $\Phi_i(\boldsymbol{X}, t)$ describe a set of European-style contingent claims with final payoff $\Phi_i(\boldsymbol{X}(T), T) = f_i(\boldsymbol{X}(T))$, for some fixed $T \geqslant t$. Finally, we need to assume that α^A are either deterministic or some function of the underlying prices $\boldsymbol{X}$. These assumptions ensure that the expectation values in the right-hand side of Equation (3.26) are functions of $\boldsymbol{X}$ and t only, and so our assumption, $\Phi_i = \Phi_i(\boldsymbol{X}, t)$, is self-consistent. Under these assumptions we can prove the following corollary.

COROLLARY 3.3 (Modified Black–Scholes equation) *Under the assumptions given above, Corollary 3.2 implies that the European-style contingent claims Φ_i, $i = 1, \dots, m$, obey the nonlinear Black–Scholes equations:*

$$\partial_t \Phi_i + \sum_{j=1}^{n} \Big(\alpha^* + \sum_A \alpha^A J_j^A\Big) X_j \partial_j \Phi_i - \Big(\alpha^* + \sum_A \alpha^A J_{n+i}^A\Big) \Phi_i + \frac{1}{2} \sum_{j,k=1}^{n} \Omega_{jk} X_j X_k \partial_j \partial_k \Phi_i = 0 \quad (3.29)$$

with terminal conditions $\Phi_i(\boldsymbol{X}, T) = f_i(\boldsymbol{X})$. Moreover, the $J^A = J^A(\boldsymbol{X}, t)$ are a basis for the null space $\mathcal{N}$ of the $(n+m+1) \times (n+m+1)$ matrix Ω with components $\Omega_{\mu\nu} = \sum_a \sigma_\mu^a \sigma_\nu^a$, where the σ_i^a, $i = 1, \dots, n$, are the volatilities of the underlying securities, and we define $\sigma_0^a := 0$:

$$\sigma_{n+i}^a := \sum_{j=1}^{n} \sigma_j^a X_j \partial_j \log \Phi_i(\boldsymbol{X}, t), \quad i = 1, \dots, m \quad (3.30)$$

and:

$$\alpha^* = r - \sum_A \alpha^A J_0^A \quad (3.31)$$

PROOF Equation (3.29) of Corollary 3.3 is a simple application of the Feynman–Kac theorem to Equation (3.26) (see Shreve (2000)). In order to calculate all components of the matrix Ω, we remind the reader that the underlying prices $\boldsymbol{X}$ obey:

$$\mathrm{d}X_i = X_i \Big[\Big(\alpha^* + \sum_A \alpha^A J_i^A\Big) \mathrm{d}t + \sum_a \sigma_i^a \,\mathrm{d}W_a^*\Big], \quad i = 1, \dots, n \quad (3.32)$$

This implies that the stochastic part of $\mathrm{d}\Phi_i$ is given by:

$$\mathrm{d}\Phi_i(\boldsymbol{X}, t) = \Phi_i(\boldsymbol{X}, t) \sum_{j=1}^{n} X_j \partial_j \log \Phi_i(\boldsymbol{X}, t) \sigma_j^a \,\mathrm{d}W_a^* + \cdots \quad (3.33)$$

Therefore, the volatilities for the $X_{n+i} = \Phi_i$ securities are:

$$\sigma^a_{n+i} = \sum_{j=1}^{n} \sigma^a_j X_j \partial_j \log \Phi_i(\boldsymbol{X}, t) \tag{3.34}$$

Moreover, since X_0 is a deterministic process, it follows that $\sigma^a_0 = 0$. In order to prove Equation (3.31) of the corollary, we recall that the savings account obeys $\mathrm{d}X_0 = rX_0\,\mathrm{d}t$. This implies that $r = \alpha^* + \sum_A \alpha^A J^A_0$. This completes the proof. □

The example of Section 2.1 is a special case of Corollary 3.3, with $n = m = 1$. In this case there is only one null direction. We will use the notation $X_1 := X$, $\Phi_1 := V$, $\sigma^1_1 := \sigma_1$ and $\alpha^1 := \tilde{\alpha}$ in what follows. A choice for the basis of the null space was given in Equation (2.19), which we repeat here for the convenience of the reader:

$$J = \frac{1}{\sqrt{2}\sqrt{1 + X\partial X \log V(X\partial X \log V - 1)}} \begin{pmatrix} 1 - X\partial X \log V \\ X\partial X \log V \\ -1 \end{pmatrix} \tag{3.35}$$

It follows that Equation (3.29) becomes:

$$\begin{aligned} 0 &= \partial_t V + (\alpha^* + \tilde{\alpha} J_1) X \partial_X V - (\alpha^* + \tilde{\alpha} J_2) V + \tfrac{1}{2}\sigma_1^2 X^2 \partial_X^2 V \\ &= \partial_t V + (r + \tilde{\alpha}(J_1 - J_0)) X \partial_X V - (r + \tilde{\alpha}(J_2 - J_0)) V + \tfrac{1}{2}\sigma_1^2 X^2 \partial_X^2 V \\ &= \partial_t V + r(X\partial X V - V) + \tfrac{1}{2}\sigma_1^2 X^2 \partial_X^2 V \\ &\qquad + \sqrt{2}\tilde{\alpha} V \sqrt{1 + X\partial X \log V(X\partial X \log V - 1)} \end{aligned} \tag{3.36}$$

This is exactly what we obtained in Section 2.1 (see Equation (2.28)).

4 MEASURING ARBITRAGE CURVATURE

In this section we explain how to estimate the arbitrage parameters α^A using financial data. Given the discussion in the previous section, measuring these parameters is equivalent to measuring the "curvature" of the market. Needless to say, we can do this for a subset of all instruments only, and there are many technical difficulties, which we discuss below.

Even though α^A is a gauge invariant, it is still defined up to a rotation in the null space[4] $\mathcal{N}$. Therefore, the basic idea is to measure the rotational and gauge-invariant quantity:

$$\sum_{\mu,A} \frac{\alpha^A J^A_\mu}{X_\mu} \frac{\mathrm{d}X_\mu}{\mathrm{d}t} = \sum_A (\alpha^A)^2 \geqslant 0 \tag{4.1}$$

[4] For notational simplicity, we will still use $\mathcal{N}$ for the null space of the particular market subsector under study. However, it is important to keep in mind that this is not the null space of the full market.

where α^A in the left-hand side of this equation is expressed as:

$$\alpha^A = \sum_{\mu} \frac{J_{\mu}^A}{X_{\mu}} \frac{\mathrm{d}X_{\mu}}{\mathrm{d}t} \tag{4.2}$$

The vectors J^A must be calculated using an estimate for the quadratic variation:

$$\Omega_{\mu\nu} := \frac{\mathrm{d}\langle \log X_{\mu}, \log X_{\nu} \rangle}{\mathrm{d}t}$$

and the results of Proposition 2.2. We introduce the notation:

$$\mathcal{A}^2 := \sum_{A} (\alpha^A)^2$$

for the measurement of arbitrage curvature. A positive detection of $\mathcal{A}^2$ can be translated into a self-financing arbitrage portfolio strategy using the result of the following proposition.

PROPOSITION 4.1 (Arbitrage strategy) *Let the asset corresponding to $\mu = 0$ be the numeraire ($X_0 := 1$). If the market model satisfies the positive curvature assumption:*

$$\mathcal{A}^2 > 0 \tag{4.3}$$

then the portfolio allocation:

$$\phi_0(t) := \sum_{i=1}^{N} \int_0^t \phi_i(s)\,\mathrm{d}X_i(s) + \sum_{A} J_0^A(t)\alpha^A(t) \tag{4.4}$$

$$\phi_i(t) := \sum_{A} \frac{J_i^A(t)\alpha^A(t)}{X_i(t)}, \quad i = 1, \dots, N \tag{4.5}$$

is a self-financing arbitrage strategy delivering wealth:

$$V(t) = \int_0^t \mathcal{A}^2(s)\,\mathrm{d}s \tag{4.6}$$

PROOF First, we check that the strategy is self-financing, that is:

$$\sum_{\mu=0}^{N} (\mathrm{d}\phi_{\mu} X_{\mu} + \mathrm{d}\langle \phi_{\mu}, X_{\mu} \rangle) = 0 \tag{4.7}$$

where "d" denotes the Ito differential (see Chapter 4.1.2 in Lamberton and Lapeyre (2007)). This is proved by the following computation:

$$
\begin{aligned}
&\sum_{\mu=0}^{N}(\mathrm{d}\phi_\mu X_\mu + \mathrm{d}\langle\phi_\mu, X_\mu\rangle) \\
&\quad = \mathrm{d}\phi_0 X_0 + \mathrm{d}\langle\phi_0, X_0\rangle + \sum_{i=1}^{N}(\mathrm{d}\phi_i X_i + \mathrm{d}\langle\phi_i, X_i\rangle) \\
&\quad = \sum_A \sum_{i=1}^{N} \frac{\alpha^A J_i^A \, \mathrm{d}X_i}{X_i} + \sum_A \mathrm{d}(\alpha^A J_0^A) + \sum_{i=1}^{N} X_i \, \mathrm{d}\biggl(\sum_A \frac{\alpha^A J_i^A}{X_i}\biggr) \\
&\qquad + \sum_{i=1}^{N} \mathrm{d}\biggl\langle \sum_A \frac{\alpha^A J_i^A}{X_i}, X_i \biggr\rangle \\
&\quad = \sum_A \mathrm{d}\biggl(\alpha^A \sum_{\mu=0}^{N} J_\mu^A\biggr) \\
&\quad = 0
\end{aligned}
\tag{4.8}
$$

Since the self-financing condition is fulfilled, the portfolio value can be computed as:

$$
\begin{aligned}
V(t) &= \sum_{\mu=0}^{N} \phi_\mu(t) X_\mu(t) \\
&= \int_0^t \phi_i(s)\,\mathrm{d}X_i(s) + \sum_A \alpha^A \underbrace{\sum_{\mu=0}^{N} J_\mu^A}_{=0} \\
&= \int_0^t \sum_A \alpha^A(s) \sum_{i=1}^{N} \frac{J_i^A(s)}{X_i(s)}\,\mathrm{d}X_i(s) \\
&= \int_0^t \sum_A \alpha^A(s) \sum_{\mu=0}^{N} \frac{J_\mu^A(s)}{X_\mu(s)}\,\mathrm{d}X_\mu(s) \\
&= \int_0^t \sum_A \alpha^A(s)^2\,\mathrm{d}s \\
&= \int_0^t \mathcal{A}^2(s)\,\mathrm{d}s
\end{aligned}
\tag{4.9}
$$

Since the arbitrage curvatures are positive, we see that $V(0) = 0$ and $V(t) > 0$ for all times $t \in [0, T]$. The proof is complete. □

There is no continuous time trading in the markets, and we can only do measurements in discrete time. Moreover, our estimate of Ω will always include errors. This means that we will always have a noise term on the right-hand side of Equation (4.1). The goal of this section is to explain the basic steps used to measure arbitrage curvature, and to understand the major sources of error in such measurements. A key aspect of our algorithm is that we test directly for the gauge invariance of the arbitrage signal. This allows us to check the robustness of our estimators. We find that the gauge invariance of the arbitrage signal, as predicted by the stochastic models, is indeed obeyed with good accuracy in the real market.

4.1 Basic algorithm

In what follows, we will use a hat to signify that a variable is an estimate of some parameter. For example, $\hat{\Omega}$ is an estimate for Ω. The first problem we face is finding an estimate for the quadratic variation Ω and determining the null space $\mathcal{N}$ defined in Section 2 (if nontrivial). This is a familiar problem in volatility modeling. Since we will never observe Ω directly, it is expected that our estimate will not have any exact zero mode, but only eigenvectors with small eigenvalues. In fact, *a priori*, we do not know if the space $\mathcal{N}$ is nontrivial. We can only guess its dimension.

Let $\hat{\Omega}$ be any estimate for Ω. Then, following Proposition 2.2, we construct the matrix:

$$\hat{G} = \hat{\Omega} - \frac{1}{N}(U\hat{\Omega} + \hat{\Omega}U) + \frac{1}{N^2}\operatorname{Tr}(U\hat{\Omega})U \tag{4.10}$$

where N is the number of rows (or columns) of Ω and U is the matrix of all ones $U_{\mu\nu} = 1$, $\forall\mu,\nu$. We can then use standard algorithms to compute the eigenspace of $\hat{G}$. This will yield orthonormal eigenvectors:

$$\hat{G}\hat{J}^A = \lambda^A\hat{J}^A, \quad (\hat{J}^A)^{\dagger}\hat{J}^B = \delta^{AB}, \quad A, B = 0, 1, \dots, N-1 \tag{4.11}$$

where $\lambda^A \geqslant 0$ since $\hat{G}$ is positive semidefinite. As a matter of fact, since Ω and U commute, they have a common basis of eigenvectors and a short computation proves that G has always (at least) one zero eigenvalue and the biggest $N-1$ eigenvalues equal those of Ω, which are not negative (see Proposition 2.2). In practice, there will only be one exact zero eigenvector: $\hat{J}^0 \propto (1, 1, \dots, 1)^{\dagger}$. Summarizing, the eigenvalues of G, in increasing order of magnitude, are:

$$0 = \lambda^0 \leqslant \lambda^1 \leqslant \lambda^2 \leqslant \cdots \leqslant \lambda^{N-1} \tag{4.12}$$

It is easy to show (see Proposition 2.2) that:

$$\sum_{\mu} \hat{J}^A_{\mu} = 0 \quad \text{for } A = 1, 2, \dots, N-1 \tag{4.13}$$

Our estimate for the basis of $\mathcal{N}$ will be to chose the first k eigenvectors with the smallest eigenvalues: $\hat{J}^A$, $A = 1, \ldots, k < N - 1$. In doing this, we are assuming that $\dim(\mathcal{N}) = k$.

Once we have calculated $\hat{J}^A$, we can compute our estimate of α^A in discrete time:

$$\hat{\alpha}^A(t + \delta t) = \sum_{\mu} \frac{\hat{J}^A_\mu(t)}{\delta t X_\mu(t)} [X_\mu(t + \delta t) - X_\mu(t)] \tag{4.14}$$

Note that $\hat{J}^A(t)$ is constructed with information up to time t only. This estimate is consistent with the nonanticipating nature of Ito integrals. The time step δt is, of course, arbitrary. Our estimate for $\mathcal{A}^2$ now becomes:

$$\hat{\mathcal{A}}^2(t + \delta t) = \sum_{A=1}^{k} [\hat{\alpha}^A(t)]^2 + \sum_{A=1}^{k} \hat{\alpha}^A(t)[\hat{\alpha}^A(t + \delta t) - \hat{\alpha}^A(t)] \tag{4.15}$$

In the limit of short timescales, and if there is nontrivial arbitrage, we expect that this estimator will converge to the true signal:

$$\hat{\mathcal{A}}^2(t + \delta t) = \mathcal{A}^2(t) + \sum_{A} \alpha^A \, \mathrm{d}\alpha^A = \mathcal{A}^2(t) + \mathcal{O}(\delta t), \quad \delta t \to 0 \tag{4.16}$$

The convergence in Equation (4.16) is only valid if, in the limit $\delta t \to 0$, we have:

$$\mathbb{E}_t[\hat{\alpha}^A(t+\delta t)] - \hat{\alpha}^A(t) = \mathcal{O}(\delta t), \qquad \operatorname{cov}_t[\hat{\alpha}^A(t+\delta t), \hat{\alpha}^B(t+\delta t)] = \mathcal{O}(\delta t) \tag{4.17}$$

Therefore, we expect that, if there is nontrivial arbitrage in the market, the estimator (4.15) will give us a positive signal on average. Since the timescale is arbitrary, it is convenient to set $\delta t = 1$ henceforth.

There are several candidates for an estimator for Ω. The "right" choice of $\hat{\Omega}$ should reflect our beliefs about the true dynamics of the asset values. Here we will simply take the empirical estimator for covariance of the time series of log returns for a window of length L. More precisely, our data consists of a number of time series for the prices X_μ, $\mu = 0, \ldots, N - 1$, in certain units, say US dollars.[5] Our estimator reads:

$$\begin{aligned} \hat{\Omega}_{\mu\nu}(t) = {} & \frac{1}{L} \sum_{i=0}^{L-1} \log\left[\frac{X_\mu(t - i)}{X_\mu(t - i - 1)}\right] \log\left[\frac{X_\nu(t - i)}{X_\nu(t - i - 1)}\right] \\ & - \frac{1}{L^2} \sum_{i,j=0}^{L-1} \log\left[\frac{X_\mu(t - i)}{X_\mu(t - i - 1)}\right] \log\left[\frac{X_\nu(t - j)}{X_\nu(t - j - 1)}\right] \end{aligned} \tag{4.18}$$

For more sophisticated estimators, see Hardle *et al* (2008) and Zhang (2006). We are now in a position to summarize the most basic algorithm to detect arbitrage.

[5] We also include the US dollar itself as an asset in which we have $X_0 = 1$.

4.1.1 Algorithm

(1) Starting with the time series for X_μ, $\mu = 0, \ldots, N-1$, in an interval $[t, t-L]$, we estimate $\hat{\Omega}_{\mu\nu}(t)$ using Equation (4.18).

(2) We then calculate the $\hat{G}$ matrix using Equation (4.10), and its orthonormal eigenspace. The eigenvectors will be labeled as $\hat{J}^A$, $A = 0, 1, \ldots, N-1$, in order of increasing eigenvalues $0 = \lambda^0 \leqslant \lambda^1 \leqslant \cdots \leqslant \lambda^{N-1}$. Moreover, $\hat{J}^0 \propto (1, 1, \ldots, 1)^\dagger$.

(3) Given a guess for the dimension of the null space $k = \dim(\mathcal{N})$, we take as its basis the following eigenvectors of $\hat{G}$: $\hat{J}^A$, $A = 1, \ldots, k$.

(4) We then calculate $\hat{\alpha}^A(t+1)$ from Equation (4.14), which uses information up to time $t+1$.

(5) Roll the time window by one step, and repeat steps (1)–(4). Once we have more than one estimate for $\hat{\alpha}^A$, we can calculate our final arbitrage estimator $\hat{\mathcal{A}}^2$ from Equation (4.15).

(6) In order to explicitly check for gauge invariance, we repeat steps (1)–(5), using each asset X_μ as a numeraire. For example, if we want to use X_1 as a numeraire, we divide all elements of the time series by the corresponding element of X_1, eg, $X_\mu(s) \to X_\mu(s)/X_1(s)$, $\forall \mu$ and $s \in [t, t-L]$. Then we repeat steps (1)–(5) with the new time series. Note that this is a nontrivial transformation in the data and, in practice, we will get different estimates for $\hat{\alpha}^A$.

Before discussing the results of Algorithm 4.1.1, we need to understand what the main sources of error are in our signal. This is done in the next subsection.

4.2 Sources of error

The sources of error in our measurement of $\mathcal{A}^2$ can be divided into three groups. First, there is gauge dependence. Second, there is a gauge-invariant noise, which we will discuss below. Finally, when using high-frequency financial data, one is faced with the so-called market microstructure noise, which is partly due to the bid–ask bounce effect (see Hardle *et al* (2008)).

We begin by looking at sources of gauge dependence. Note that our construction of the estimators assumes that, under a gauge transformation, $\hat{\Omega}$ transforms like Ω (see Equation (2.13)). However, the gauge-transformation rule in the real world can be quite different, because the unknown effective dynamics could lead to gauge dependencies. We do not have an *a priori* test for this source of error. The only way to test for it is to make our calculations in different gauges and see how different the answers are. We show examples of this in the following sections.

The second source of error in our signals comes from a gauge-invariant noise term. In fact, we will see that this is the dominant noise contribution. In order to understand this noise, it is convenient to discretize the Ito integral and write our estimate for α^A as:

$$\begin{aligned}
\hat{\alpha}^A(t+1) &= \sum_{\mu} \frac{\hat{J}^A_\mu(t)}{X_\mu(t)}[X_\mu(t+1) - X_\mu(t)] \\
&= \sum_{\mu,B} \hat{J}^A_\mu(t) J^B_\mu(t)\alpha^B(t) + \sum_{\mu,a} \hat{J}^A_\mu(t)\sigma^a_\mu \beta^a(t) \\
&\quad + \sum_{\mu,a} \hat{J}^A_\mu(t)\sigma^a_\mu [W_a(t+1) - W_a(t)] \\
&:= \alpha^A_{\text{trend}}(t) + \varepsilon^A(t+1)
\end{aligned} \tag{4.19}$$

Here we have decomposed the signal in a trend:

$$\alpha^A_{\text{trend}}(t) := \sum_{\mu,B} \hat{J}^A_\mu(t) J^B_\mu(t)\alpha^B(t) + \sum_{\mu,a} \hat{J}^A_\mu(t)\sigma^a_\mu \beta^a(t) \tag{4.20}$$

and a stochastic noise term:

$$\varepsilon^A(t+1) := \sum_{\mu,a} \hat{J}^A_\mu(t)\sigma^a_\mu [W_a(t+1) - W_a(t)] \tag{4.21}$$

with $\mathbb{E}_t[\varepsilon^A(t+1)] = 0$. Since $\hat{J}^A$ is only an estimate for the real J^A, we have that $\sum_\mu \hat{J}^A \sigma^a_\mu \neq 0$ in general. Therefore, our error in the estimate of J^A will induce an extra noise term in the signal. Moreover, it will also induce some gauge dependency. To see this, note that, under a change of numeraire, we have $\sigma^a_\mu \to \sigma^a_\mu + \delta\sigma^a$ and $\beta^a \to \beta^a + \delta\sigma^a$. It is then easy to check that the trend will transform according to:

$$\alpha^A_{\text{trend}} \to \alpha^A_{\text{trend}} + \sum_{\mu,a} \hat{J}^A_\mu \sigma^a_\mu \delta\sigma^a$$

However, note that the noise term is gauge invariant. In fact, one expects the term $\sum_\mu \hat{J}^A_\mu \sigma^a_\mu$ to be quite small. Moreover, since, in Algorithm 4.1.1, the gauge transformation is of the order $\delta\sigma^a = \mathcal{O}(\sigma^a_\mu)$, we expect the gauge dependence coming from the trend to be negligible. We will see that, in real financial data, most of the signal can be accounted for by the gauge-invariant noise term.

We are interested in estimating the size of the noise contribution. For that, we compute the variance of the noise using information up to time t:

$$
\begin{aligned}
\operatorname{var}_t&\left[\sum_A \hat{\alpha}^A(t)\hat{\alpha}^A(t+1)\right]\\
&= \mathbb{E}_t\left\{\left[\sum_A \hat{\alpha}^A(t)(\hat{\alpha}^A(t+1) - \mathbb{E}_t[\hat{\alpha}^A(t+1)])\right]^2\right\}\\
&= \sum_{A,B} \hat{\alpha}^A(t)\hat{\alpha}^B(t)([\hat{J}^A(t)]^\dagger G(t)\hat{J}^B(t))\\
&= \sum_A [\hat{\alpha}^A(t)]^2\lambda^A(t) + \sum_{A,B} \hat{\alpha}^A(t)\hat{\alpha}^B(t)([\hat{J}^A(t)]^\dagger \delta G(t)\hat{J}^B(t))
\end{aligned}
\tag{4.22}
$$

where $\delta G = G - \hat{G}$. If we think that our estimate of G is good, we can neglect the $\mathcal{O}(\delta G)$ term and approximate:

$$
\operatorname{var}_t[\hat{\mathcal{A}}^2(t+1)] \approx \sum_A [\hat{\alpha}^A(t)]^2\lambda^A(t) \leqslant \sum_A [\hat{\alpha}^A(t)]^2\lambda_k(t) \tag{4.23}
$$

where we remind the reader that k is our estimate for the dimension of $\mathcal{N}$, and the eigenvalues of $\hat{G}$ have been ordered so that $\lambda^1 \leqslant \lambda^2 \leqslant \cdots \leqslant \lambda^k$.

An interesting consequence of Equation (4.23) is that we can place a fundamental bound on the size of the arbitrage curvature in order that it is detectable. We have:

$$
\mathcal{A}^2 \geqslant \sqrt{\operatorname{var}_t[\hat{\mathcal{A}}^2]} \implies \mathcal{A}^2 \geqslant \lambda_k \tag{4.24}
$$

This means that, in order to have a chance to detect arbitrage, one needs to find financial products whose time series are as correlated as possible, which implies a very small value of λ_k.

The third source of error is market microstructure noise. This effect is relevant in high-frequency data, when the size of the price movements is comparable with the bid–ask spread. In order to model this noise, it is convenient to set $X_0 := 1$ as our numeraire. The standard way of simulating this noise is to introduce an additional jump term $\eta_i(t)$ to the log prices X_i, $i = 1, \dots, N$. More precisely, the observed price is $\tilde{X}_i$ and it is given by:

$$
\log \tilde{X}_i(t) = \log X_i(t) + \eta_i(t) \tag{4.25}
$$

where X_i is the "true" Ito process, and, for simplicity, we assume that:

$$
\left.
\begin{aligned}
\mathbb{E}[X_i\eta_j] &= 0\\
\mathbb{E}[\eta_i] &= 0\\
\mathbb{E}[\eta_i\eta_j] &:= \eta^2\delta_{ij}
\end{aligned}
\right\}
\tag{4.26}
$$

Moreover, the noise terms are uncorrelated between different times. We can then show that our estimator will be contaminated by an amount:

$$\mathbb{E}[\hat{\mathcal{A}}^2] = \mathbb{E}[\mathcal{A}^2] - \gamma\frac{\eta^2}{\delta t^2} + \mathcal{O}\left(\frac{\eta^4}{\delta t^2}\right), \quad \delta t \to 0 \tag{4.27}$$

where:

$$\gamma := \dim(\mathcal{N}) - \mathbb{E}\left[\left(\sum_A J_0^A\right)^2\right]$$

It can be shown that $\gamma \geqslant 0$. Therefore, we see that the microstructure noise leads a negative contribution to our estimation of $\mathcal{A}^2$. The absolute value of such a contribution diverges as we move toward higher frequencies ($\delta t \to 0$).

One way of detecting the presence of microstructure noise is to note that:

$$\lim_{\delta t \to 0} \mathbb{E}\left[\log\left(\frac{\tilde{X}_i(t+\delta t)}{\tilde{X}_i(t)}\right)\log\left(\frac{\tilde{X}_i(t)}{\tilde{X}_i(t-\delta t)}\right)\right] = -\eta^2 < 0 \tag{4.28}$$

In other words, the microstructure noise induces a negative correlation between subsequent log returns. We find that this effect is quite pronounced for equity and futures data. However, for stock indexes, the effect seems to be negligible. This is mainly due to the fact that the microstructure noise "averages out" between all the stocks in the index.

There is an extra source of error, which is intrinsic to Algorithm 4.1.1, but only if we use a rolling window in our estimation of $\hat{J}^A$. For example, suppose that we estimate $\hat{J}^A(t)$ and then roll the window and estimate $\hat{J}^A(t+1)$. Even if the matrices $\hat{G}(t)$ and $\hat{G}(t+1)$ are near, $\hat{J}^A(t)$ and $\hat{J}^A(t+1)$ can differ by a large orthogonal transform. It can be just a sign flip, for example, since the eigenvalue equations are invariant under $\hat{J}^A \to -\hat{J}^A$. However, suppose two eigenvalues are near to each other, ie, $\lambda_1 \approx \lambda_2$. Then, any linear combination of $\hat{J}^1$ and $\hat{J}^2$ is also approximatively an eigenvector of $\hat{G}$. In physics, this is known as the problem of degenerate perturbation theory (see, for example, Sakurai (1994)). More generally, we have that:

$$\lim_{\|\hat{G}(t)-\hat{G}(t+1)\|\to 0} \hat{J}^A(t+1) = \sum_B C^{AB}\hat{J}^B(t) \tag{4.29}$$

where C is an orthogonal matrix, ie:

$$C^\dagger C = \mathbf{1} \tag{4.30}$$

The problem can be solved if we can determine C. If so, we can construct the "correct" eigenvectors:

$$\tilde{J}^A(t+1) := \sum_B (C^\dagger)^{AB}\hat{J}^B(t+1)$$

so that:

$$\lim_{\|\delta\hat{G}\|\to 0} \tilde{J}^A(t+1) = \hat{J}^A(t)$$

An approximate solution for C, denoted by $\hat{C}$, can be found by minimizing the Lagrangian:

$$\mathcal{L} = \mathrm{Tr}[(\tilde{C} - \hat{C})^\dagger(\tilde{C} - \hat{C})] + \mathrm{Tr}[\lambda(\hat{C}^\dagger\hat{C} - \mathbf{1})] \tag{4.31}$$

where the $k \times k$ real matrix $\tilde{C}$ has components:

$$\tilde{C}^{AB} := [\tilde{J}^A(t+1)]^\dagger \tilde{J}^B(t) \tag{4.32}$$

and λ is a symmetric matrix that serves as a Lagrange multiplier implementing the constraint $\hat{C}^\dagger\hat{C} = \mathbf{1}$. This is implemented in our numerical routines.

5 DYNAMIC ARBITRAGE STRATEGIES

In this section we apply Algorithm 4.1.1 to simulated financial data to construct dynamic arbitrage strategies. First, we test our arbitrage detection algorithm, in order to see whether arbitrage is indeed detectable or if it is masked by statistical noise. Then, we backtest an equity indexes strategy, which is not actually tradable, unless one can find appropriate exchange-traded funds. Here we discover arbitrage. Later, we backtest a future indexes strategy, which is directly tradable. In this case arbitrage is more difficult to find.

5.1 Arbitrage detection

To simulate the financial data we study the simple lognormal random walk model with constant coefficients (see Equation (2.1)). The solution to the stochastic differential equation (2.1) is:

$$X_\mu(t) = X_\mu(0)\exp\left[\left(\alpha_\mu - \frac{1}{2}\sum_{a=1}^{d}(\sigma_\mu^a)^2\right)t + \sum_{a=1}^{d}\sigma_\mu^a B_a(t)\right] \tag{5.1}$$

where $B(t) := [B_1(t), \dots, B_d(t)]^\dagger$ is a standard multivariate Brownian motion with:

$$\mathrm{E}[B_a(t)] = 0, \qquad \mathrm{cov}[B_a(t), B_b(t)] = t\delta_{ab} \tag{5.2}$$

for all $a, b = 1, \dots, d$. As usual, we decompose the trends as:

$$\alpha_\mu = \alpha + \sum_{a=1}^{d}\beta^a\hat{\sigma}_\mu^a + \sum_{A\in\mathcal{N}}\alpha^A J_\mu^A, \qquad \hat{\sigma}_\mu^a = \sigma_\mu^a - \frac{1}{N}\sum_{\nu=0}^{N-1}\sigma_\nu^a \tag{5.3}$$

FIGURE 1 Simulation with twenty lognormal random walks.

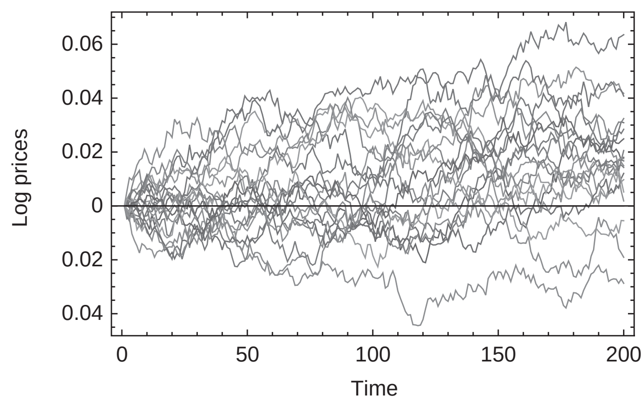

We begin with an example with $N = 21$ assets and $d = 18$ Brownian motions, which implies $k = \dim(\mathcal{N}) = 2$. We take as a first asset a bank account with a zero interest rate, and make it our numeraire. This means that we choose:

$$X_0 :\equiv 1, \qquad \sigma_0^a = 0, \qquad \alpha_0 = 0 \tag{5.4}$$

which implies:

$$\alpha = \frac{1}{N} \sum_{i=1}^{N-1} \sum_{a=1}^{d} \beta^a \sigma_i^a - \sum_{i=1}^{N-1} \sum_{A=1}^{2} \alpha^A J_i^A \tag{5.5}$$

In Figure 1 we show a particular simulation of the log prices, where we take β^a, σ_i^a, α^A from uniform random distributions in the intervals $\beta^a \in [-10^{-4}, 10^{-4}]$, $\sigma_i^a \in [-10^{-3}, 10^{-3}]$ and $\alpha^A \in [-10^{-4}, 10^{-4}]$. The simulation was generated using Mathematica. The arbitrage detection algorithm was implemented in C++. Each price was taken at a time separation of $\Delta t = 1$ (arbitrary time unit). In this particular case we calculate $\hat{\Omega}$ using the first 100 prices of the time series. In other words, we do not use a moving window. The results with the moving window are very similar.

Now suppose we assume (correctly) that we have $k = \dim(\mathcal{N}) = 2$. We then run Algorithm 4.1.1 and find the signal shown in Figure 2 on the facing page. The solid horizontal line at $\mathcal{A}^2 \approx 10^{-8}$ is the correct value of $\mathcal{A}^2$. Therefore, we see that we get an accurate estimate for the arbitrage curvature. Note that, as we discussed in the previous section, in our algorithm we compute $\hat{\mathcal{A}}^2$ using each of the different assets as a numeraire. We include error bars showing the range of values obtained using the different gauges. The results in this simulated sample are gauge invariant to such a high accuracy that the error bars cannot be appreciated.

FIGURE 2 Result for the arbitrage detection algorithm applied to the simulated data in Figure 1 on the facing page.

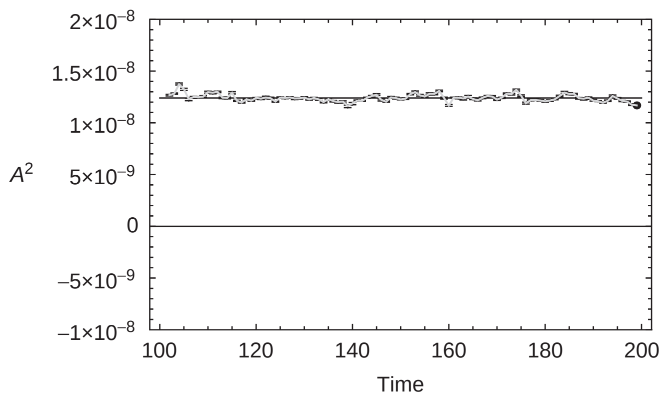

Here we assume (correctly) $k = 2$ null directions. The solid horizontal line at $\mathcal{A}^2 \sim 10^{-8}$ is the correct value of $\mathcal{A}^2$. The gray solid line is the US dollar value of the signal.

In the previous section we discussed how the main source of error in our detection technique can be related to the biggest eigenvalue λ_k of the set $\{\lambda_1, \dots, \lambda_k\}$. This led to a gauge-invariant noise term. In this simulated sample data, we find that $\lambda_k \approx 10^{-21}$, and so, using Equation (4.23), we find:

$$\sqrt{\mathrm{var}[\hat{\mathcal{A}}^2]} \approx \sqrt{10^{-8} \times 10^{-21}} \approx 10^{-15}$$

Therefore, the noise term is very small in this case. The fluctuations seen in Figure 2 are an artifact of this particular model. To understand them, we can expand Equation (5.1) as:

$$\frac{X_\mu(t+1) - X_\mu(t)}{X_\mu(t)} = \alpha_\mu + \sum_a \sigma_\mu^a B_a(1) + \varepsilon_\mu + \cdots \tag{5.6}$$

where:

$$\varepsilon_\mu = \frac{1}{2}\left(-\Omega_{\mu\mu} + \sum_{a,b} \sigma_\mu^a \sigma_\mu^b B_a(1) B_b(1)\right) \tag{5.7}$$

It is easy to show that ε_μ is gauge invariant and that:

$$\mathbb{E}[\varepsilon_\mu] = 0, \qquad \mathbb{E}[\varepsilon_\mu \varepsilon_\nu] = \tfrac{1}{2}(\Omega_{\mu\nu})^2 \tag{5.8}$$

This extra noise term, ε_μ, is the reason for the gauge-invariant fluctuations in Figure 2. The noise term vanishes if we integrate $\mathrm{d}X_\mu$ using an infinite partition of the time interval, as it is assumed in Ito integrals. Of course, this is never possible in practice.

FIGURE 3 Result for the arbitrage detection algorithm applied to the simulated data in Figure 1 on page 52.

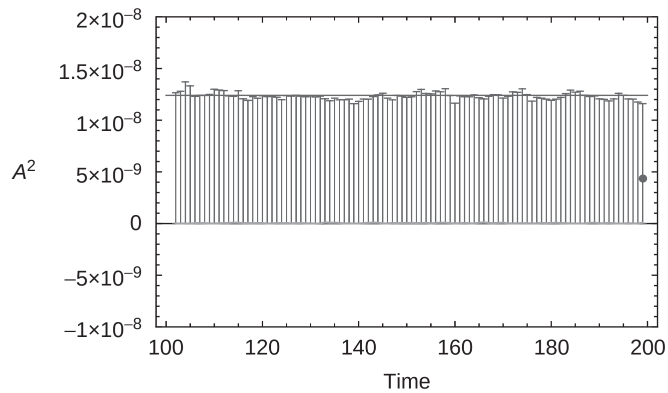

Here we assume (incorrectly) $k = 1$ null directions and use a fixed window of 100 time steps. The error bars give the range of values obtained using the different gauges. The solid dots are the mean of all results. The gray line near to the time axis is the US dollar value of the signal.

Nevertheless, we see that, in this example, the extra noise is very small compared with the arbitrage parameter $\mathcal{A}^2$. In fact, we expect this noise to be very small in general since it is of order $\mathrm{var}[\varepsilon_\mu] = O((\sigma^a_\mu)^4)$.

It is interesting to see what happens if we assume the wrong number of zero modes. For example, in Figure 3 we show what happens if we take $k = 1$. We see that we get a gauge dependent signal. Finally, in Figure 4 on the facing page we show what happens if we assume $k = 3$. In this case, the biggest eigenvalue is $\lambda_k \approx 10^{-8}$. As the figure shows, most of the fluctuations are coming from the gauge-invariant noise described in the previous section. To see this we have plotted the expected noise according to Equation (4.23):

$$\begin{aligned} \mathrm{noise}_{\pm}(t+1) &= (\mathcal{A}^2 \pm \sqrt{\mathrm{var}_t[\hat{\mathcal{A}}^2(t+1)]}) \\ &= \left(\mathcal{A}^2 \pm \sqrt{\sum_{A=1}^{k} (\hat{\alpha}^A(t))^2 \lambda^A(t)}\right) \end{aligned} \tag{5.9}$$

where $\mathcal{A}^2 \approx 10^{-8}$ is the true value of the arbitrage (which is also the mean of the signal). We see that this noise accounts for most of the fluctuations and it makes the true arbitrage signal almost undetectable. The main point we would like to make here is that the correct value of k can be estimated from the quality of the signal.

FIGURE 4 Result for the arbitrage detection algorithm applied to the simulated data in Figure 1 on page 52.

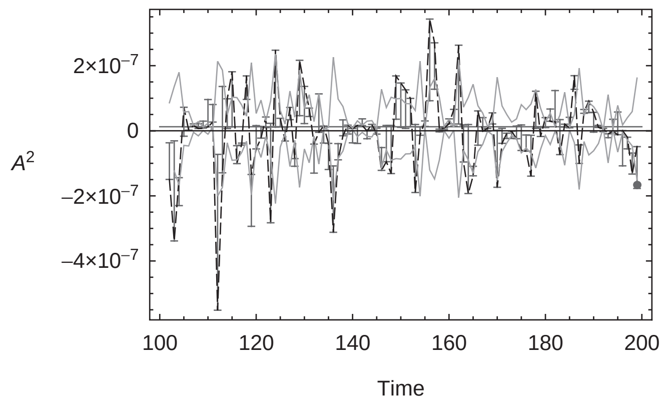

Assuming $k = 3$ null directions (incorrectly), and using a fixed window of 100 time steps. The gray lines are the noise terms noise$_{\pm}$ estimated according to Equation (5.9). The error bars give the range of values obtained using the different gauges. The solid dots are the mean of all results. The dashed line is the US dollar value of the signal.

FIGURE 5 Result for the arbitrage detection algorithm applied to the simulated data in Figure 1 on page 52, now including microstructure noise with variance $\eta^2 = 10^{-5}$.

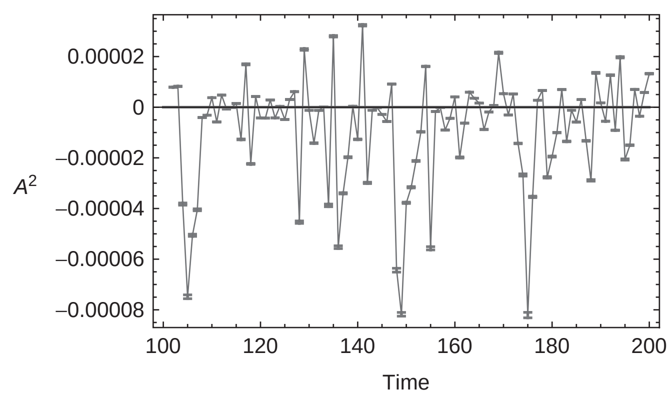

Here we assume (correctly) $k = 2$ null directions.

We can now investigate the effect of the market microstructure noise discussed in the previous section. In order to do this, we include additional white-noise terms in the price processes of Equation (5.1) as described in the previous subsection (see

FIGURE 6 Product of subsequent log returns in the presence of microstructure noise, according to Equation (5.10).

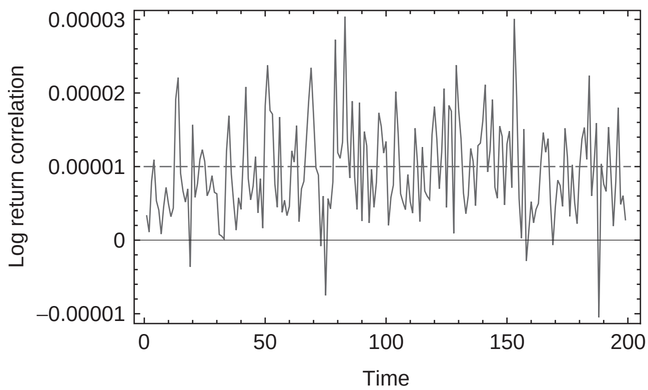

The average of this signal is equal to $\eta^2 = 10^{-5}$ (dashed horizontal line), in agreement with the theoretical prediction.

Equation (4.25)). In this particular example we choose the variance $\eta^2 = 10^{-5}$. We then apply the noise to the data of the previous example. In Figure 4 on the preceding page we show the result of the estimate of $\mathcal{A}^2$ for the contaminated data. In this case we assume (correctly) that the dimension of the null space is $k = 2$. We can clearly see how the signal is now negative on average, due to the microstructure noise. This matches the theoretical prediction in Equation (4.27). In Figure 6 we plot the product of subsequent log returns according to:

$$\hat{\eta}^2(t) := -\frac{1}{N-1}\sum_{i=1}^{N-1} \log\left(\frac{\tilde{X}_i(t+1)}{\tilde{X}_i(t)}\right) \log\left(\frac{\tilde{X}_i(t)}{\tilde{X}_i(t-1)}\right) \tag{5.10}$$

(see also Equation (4.28)), where $\tilde{X}_i$ is the contaminated price. According to Equation (4.28), we should have $\mathbb{E}[\hat{\eta}^2] = \eta^2$. This is precisely what we observe in Figure 6. In the next subsection, we will see that such signals are typical of high-frequency security prices.

5.2 Equity index strategies

Here we present some examples of our arbitrage detection algorithm applied to real financial data. We begin with a look at three major US stock indexes: the Dow Jones composite average (DJA), the NASDAQ composite index (IXIC) and the NYSE composite index (NYA). Due to their similar nature, we expect strong correlations between these indexes. Our first sample consists of daily closing prices from September 1, 2004

FIGURE 7 Arbitrage detection algorithm applied to daily closing prices of three major US indexes: DJA, IXIC and NYA.

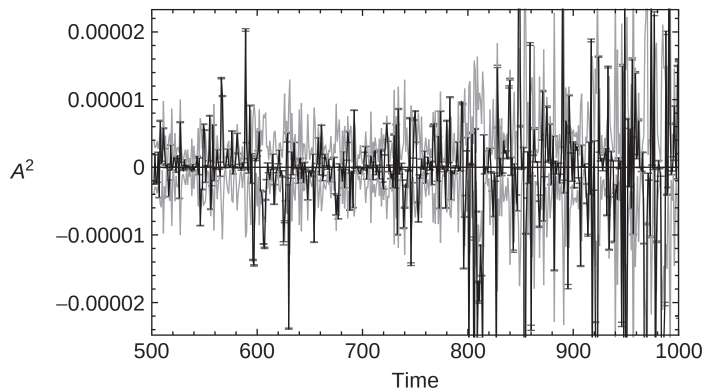

Here we show a sample of 500 data points. The gray lines are an estimate of the variance of the gauge-invariant noise using Equation (5.11). The error bars give the range of values obtained using the different gauges. The solid bars are the mean of all results. The black line is the US dollar value of the signal.

to July 16, 2009: a total of 1227 data points. The gauge-invariant matrix $\hat{G}$ has been estimated using a moving window of 500 days. We have found the following values for the eigenvalues:

$$\lambda_1 \approx 2 \times 10^{-3}, \qquad \lambda_2 \approx 5 \times 10^{-3}, \qquad \lambda_3 \approx 2 \times 10^{-2}$$

Therefore, it is reasonable to assume that the null space has only one dimension, $k = 1$. The result of the arbitrage detection algorithm is shown in Figure 7. We can see that the signal is indeed gauge invariant to a very high level of accuracy. In Figure 7 we have also included an estimate for the gauge-invariant noise term described in Section 4.2. In this case we have assumed that the average of the signal is zero (ie, no-arbitrage), and so our estimate for the expected noise is:

$$\begin{aligned} \text{noise}_{\pm}(t+1) &= \pm\sqrt{\text{var}_t[\hat{\mathcal{A}}^2(t+1)]} \\ &= \pm\sqrt{\sum_{A=1}^{k}(\hat{\alpha}^A(t))^2\lambda_A(t)} \end{aligned} \tag{5.11}$$

Looking at Figure 7, we see that the noise can explain most of the signal. Therefore, we find that our results are consistent with $\mathcal{A}^2 = 0$, and hence no arbitrage. In Figure 8 on the next page we show a histogram of the different values of $\hat{\mathcal{A}}^2$. As pointed out

FIGURE 8 Histogram of different values of $\hat{\mathscr{A}}^2$ obtained using daily market data for DJA, IXIC and NYA.

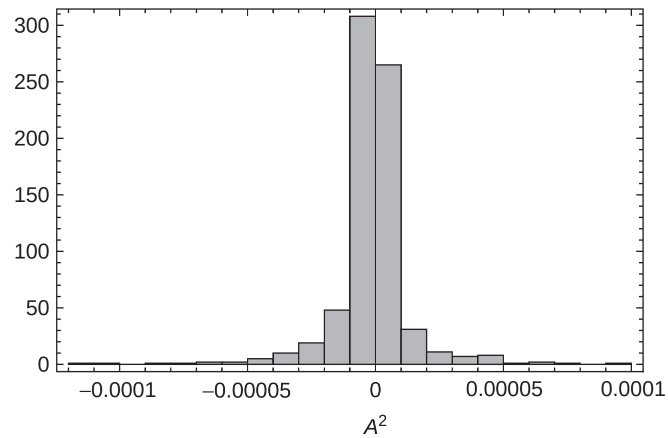

The signal-to-noise ratio is $\mathbb{E}[\hat{\mathscr{A}}^2]/\sqrt{\mathrm{var}[\hat{\mathscr{A}}^2]} \approx -0.0709$.

above, the signal is consistent with $\mathscr{A}^2 = 0$ since the signal-to-noise ratio is very low:

$$\frac{\mathbb{E}[\hat{\mathscr{A}}^2]}{\sqrt{\mathrm{var}[\hat{\mathscr{A}}^2]}} \approx -0.0709$$

It is very instructive to look at the trading strategy exploiting the arbitrage discussed in Proposition 4.1. In discrete time, the initial value of this portfolio is:

$$V(0) = \sum_{\mu} \phi_\mu(0) X_\mu(0) = 0$$

and the value at time t is simply:

$$\begin{aligned} V(t) &= \sum_{s=0}^{t-1} \sum_{A,\mu} \hat{\alpha}^A(s) \hat{J}^A_\mu(s) \frac{X_\mu(s+1) - X_\mu(s)}{X_\mu(s)} \\ &= \sum_{s=0}^{t-1} \hat{\mathscr{A}}^2(s+1) \end{aligned} \tag{5.12}$$

In Figure 9 on the facing page we show the value of this portfolio for the daily data of the three US indexes. We include the integrated profit and loss of the indexes themselves for comparison. We have multiplied the index signals by a numerical

FIGURE 9 Integrated profit and loss of the arbitrage portfolio (dashed black line) for the daily data of the US stock indexes DJA, IXIC and NYA.

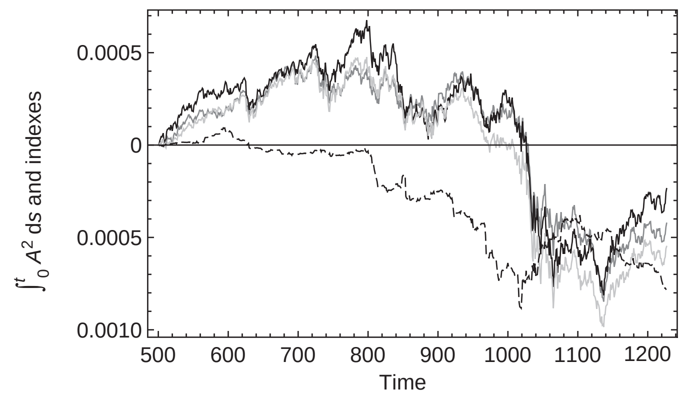

We also show the integrated profit and loss of the indexes themselves.

FIGURE 10 Arbitrage detection algorithm applied to high-frequency data for three major US indexes: DJA, IXIC and NYA.

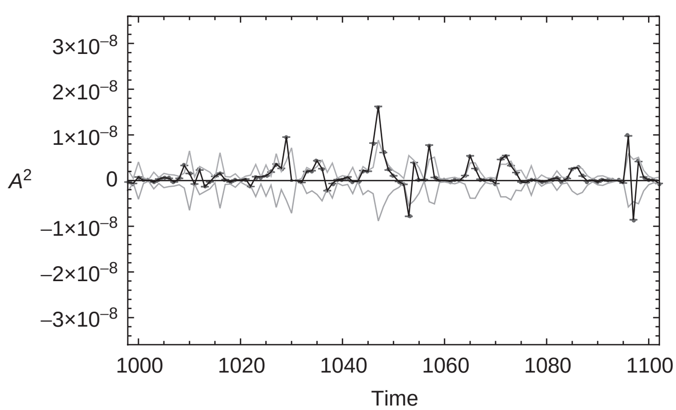

Here we show a sample of 100 data points. The gray lines are an estimate of the variance of the gauge-invariant noise using Equation (5.11). The error bars give the range of values obtained using the different gauges. The solid dots are the mean of all results. The black line is the US dollar value of the signal.

factor so that it fits in the same picture. Therefore, the overall scale on the vertical axis is irrelevant. We can see that, as expected, the performance of this portfolio is very poor for such low-frequency data.

FIGURE 11 Histogram of $\hat{\mathcal{A}}^2$ obtained using high-frequency data for DJA, IXIC and NYA.

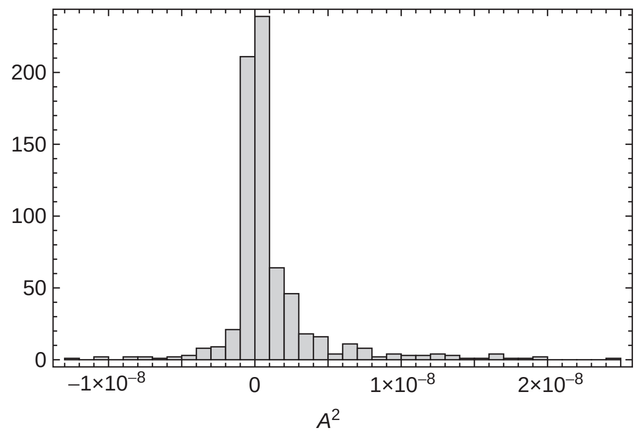

The signal-to-noise ratio is $\mathbb{E}[\hat{\mathcal{A}}^2]/\sqrt{\mathrm{var}[\hat{\mathcal{A}}^2]} \approx 0.32$.

Next we look at the same index set (DJA, IXIC, NYA), but now at short timescales. As an example, we study high-frequency data obtained on July 28, 2009. The data points are separated by 7–10 seconds. The data was collected using the "Financial-Data" package of MATHEMATICA. The gauge-invariant matrix $\hat{G}$ has been estimated using a moving window of 500 data points. We have also assumed one null direction ($k = 1$). A sample of the arbitrage detection algorithm is shown in Figure 10 on the preceding page. It is quite obvious from this figure that the signal has a very significant positive skewness. In fact, a prominent feature of the signal is a series of positive peaks. These transient events have a duration of the order of five to ten time steps, which, for this data, is about one minute. The amplitude of the peaks is quite significant compared with the noise. We argue that these peaks are precisely temporary fluctuations with $\mathcal{A}^2 \neq 0$, that is, nonzero-curvature events in the market. To show that these are not isolated events, Figure 11 shows the histogram for the full data sample. We can see significant positive skewness in the signal, compared with the daily data (see Figure 7 on page 57). In fact, we find a significant signal-to-noise ratio:

$$\frac{\mathbb{E}[\hat{\mathcal{A}}^2]}{\sqrt{\mathrm{var}[\hat{\mathcal{A}}^2]}} \approx 0.32$$

The integrated profit and loss of the arbitrage portfolio of Equation (5.12) are shown in Figure 12 on the facing page. We can see a very good performance in comparison with the daily data (see Figure 9 on the preceding page). Because of model risk, such

FIGURE 12 Integrated profit and loss of the arbitrage portfolio (dashed black line) for the high-frequency data of the US stock indexes DJA, IXIC and NYA.

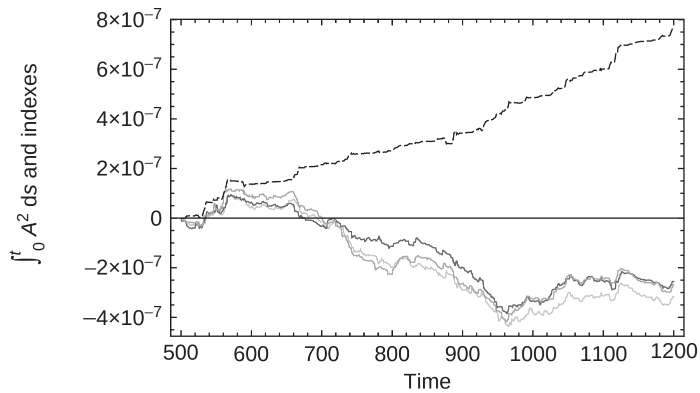

We also show the integrated profit and loss of the indexes themselves.

a portfolio can indeed have a finite probability of a loss on short timescales. However, we see that, on longer timescales (integrated signal), the probability of a loss goes to zero asymptotically as $t \to \infty$. This is an example of a statistical arbitrage as discussed in Pole (2007) and Bondarenko (2003).

We have also studied the effect of the microstructure noise on the high-frequency signal. In particular, we have computed the estimate of the noise $\hat{\eta}^2$ defined in Equation (5.10). We have found that, for this particular data sample, the contribution from such noise is very low:

$$\frac{\mathbb{E}[\hat{\eta}^2]}{\sqrt{\mathrm{var}[\hat{\eta}^2]}} \approx 0.04$$

However, if we look at traded assets such as stocks and futures, the effect becomes quite significant.

5.3 Index futures strategies

Our next data sample consists of the following set of US index futures: E-Mini S&P 500 (ESU09.CME), DJIA mini-sized (YMU09.CBT), E-Mini Nasdaq 100 (NQU09.CME) and S&P 500 index future (SPU09.CME). The data was collected on August 9, 2009, and all futures expire on September 18, 2009. We have collected prices with a frequency of seven to ten seconds separation, using the "Financial-Data" package of MATHEMATICA. These securities are highly correlated. Therefore, they are ideal for the search for the arbitrage signal. However, since these are traded

FIGURE 13 Histogram of $\hat{\mathcal{A}}^2$ obtained using high-frequency data for the index futures: ESU09.CME, YMU09.CBT, NQU09.CME, SPU09.CME.

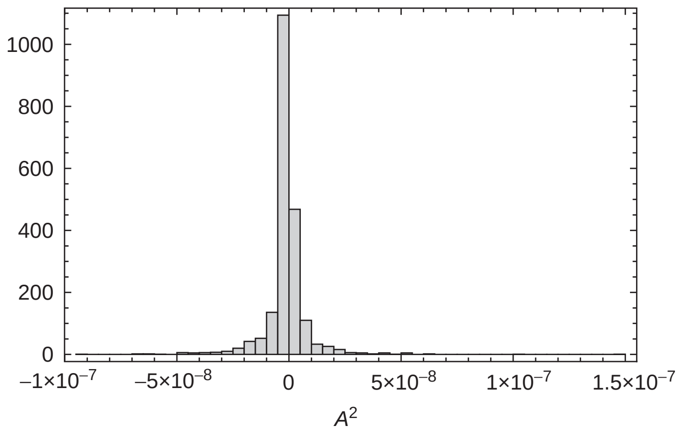

The signal-to-noise ratio is $\mathbb{E}[\hat{\mathcal{A}}^2]/\sqrt{\text{var}[\hat{\mathcal{A}}^2]} \approx -0.061$. This figure illustrates the negative effects of the market microstructure noise for the simplest trading strategy.

instruments, the effect of the bid–ask spread is more pronounced. In Figure 13 we show the histogram of the values of $\hat{\mathcal{A}}^2$ obtained by applying exactly the same algorithm as in the previous example. We get a very poor signal, which is contaminated by the microstructure noise; in fact, we get a negative mean:

$$\frac{\mathbb{E}[\hat{\mathcal{A}}^2]}{\sqrt{\text{var}[\hat{\mathcal{A}}^2]}} \approx -0.061$$

The integrated profit and loss of the simple portfolio of Equation (5.12) are shown in Figure 14 on the facing page.

We have also calculated the effect of the microstructure noise, by computing the estimate $\hat{\eta}^2$ defined in Equation (5.10). The effect for this data sample is about an order of magnitude bigger than the previous example:

$$\frac{\mathbb{E}[\hat{\eta}^2]}{\sqrt{\text{var}[\hat{\eta}^2]}} \approx 0.15$$

This noise is the main obstacle to a detection of $\mathcal{A}^2$. Nevertheless, one can devise more complicated detection methods that filter out the microstructure noise, whose effect is minimized by trading at lower frequency but using higher-frequency data to calculate the null space. As a matter of fact, when looking at the expression in Equation (4.27) describing the contamination of the arbitrage estimator by the market

FIGURE 14 Integrated profit and loss of the arbitrage portfolio (dashed black line) of Equation (5.12) for the high-frequency data of the US index futures: ESU09.CME, YMU09.CBT, NQU09.CME, SPU09.CME.

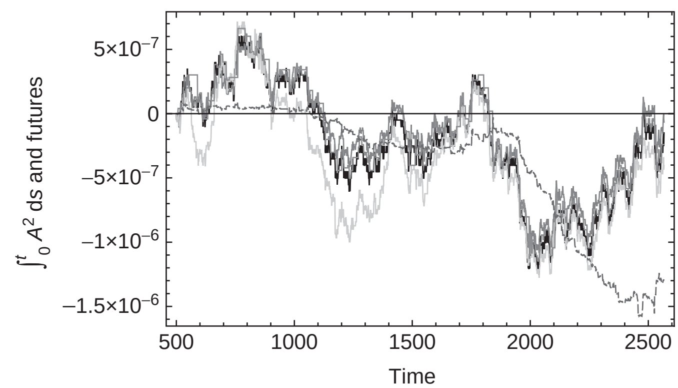

We also show the integrated profit and loss of the futures themselves. This figure illustrates the negative effects of the market microstructure noise for the simplest trading strategy.

microstructure noise, we note that, for small values of the trading period δt, the longer the period, the smaller the contamination. So, if we keep high-frequency data to compute the (approximate instantaneous) covariance matrix $\hat{\Omega}$, while increasing the trading period in such a way that the approximation relation (4.27) still holds, we can filter out the market microstructure noise and obtain a clear arbitrage signal. In Figure 15 on the next page we show the integrated profit and loss of this particular strategy.

A detailed study of similar strategies requires a precise calibration of the trading period: it must be sufficiently long that it clearly reduces the contamination but, at the same time, sufficiently short that the arbitrage contamination relation is still valid. Every asset universe requires its own calibration, which needs, as all statistical parameters utilized here do, a regular update.

6 CONCLUSIONS

In this paper we have defined a general measure of arbitrage that is invariant under changes of numeraire and equivalent probability measure. Our main assumption is that all financial instruments can be described by Ito processes. This is not a very strong assumption, as many complex financial models, including those reflecting the non-Gaussian nature of stock returns, can be modeled this way. We showed that the gauge-invariant arbitrage measure can be interpreted in terms of the curvature

FIGURE 15 Integrated profit and loss of an arbitrage portfolio (dashed black line) for low-frequency trading and high-frequency data (for null space computational purposes) of the US index futures: ESU09.CME, YMU09.CBT, NQU09.CME, SPU09.CME.

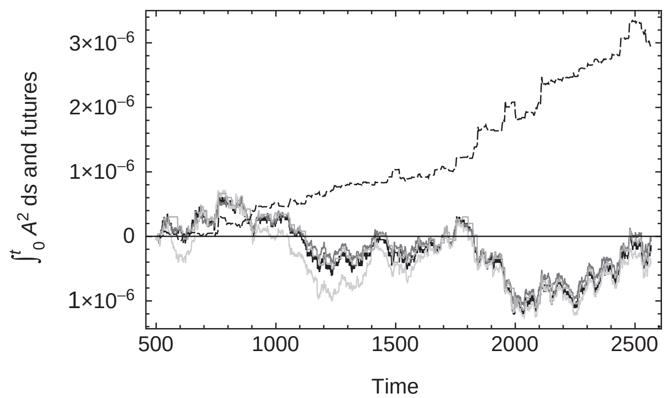

We also show the integrated profit and loss of the futures themselves. This particular strategy is designed to minimize the effects of the microstructure noise.

of the stochastic version of the Malaney–Weinstein connection (Malaney (1996) and Weinstein (2006)). The zero-curvature condition is then equivalent to the no-arbitrage principle up to the Novikov condition for the asset value dynamics. Moreover, we demonstrated a simple generalization of the classic asset pricing theorem to include arbitrage. Finally, we presented a basic algorithm to measure the market curvature using financial data. We found evidence for nonzero-curvature fluctuations in high-frequency data involving stock indexes and index futures.

From a financial perspective, we used our algorithms to exploit arbitrage systematically, and to generate profitable dynamic trading strategies. However, a distinction must be made between the case where assets are directly tradable (eg, index futures) and the case where this cannot be achieved (eg, index equities). In the tradable asset case, arbitrage is less easily detected than in the nontradable asset case and it is likely to be destroyed by the market microstructure noise. Moreover, we have completely neglected transaction costs. Even when the real asset universe (eg, exchange-traded funds) has low transaction costs, these could destroy the profitability of the dynamic arbitrage strategies depicted. The right way to tackle this problem is to develop a theory handling arbitrage and transaction costs at the same time. This will require much more empirical research, and the development of more sophisticated techniques to estimate the arbitrage curvature measure $\mathcal{A}^2$. This is left for future work. For the time being, we can still find a profitable strategy with tradable assets by optimizing the trading period to reduce the market microstructure noise and by regularly recalibrating the parameters.

From a scientific perspective, we believe that our findings represent a modest step toward an understanding of nonequilibrium market dynamics. Gauge theories provide the natural mathematical language to that aim, and arbitrage opportunities can be interpreted as a nonzero-curvature fluctuation in an economy out of equilibrium. It is interesting that most of our current economic and financial thinking relies so heavily on the assumptions of general equilibrium theory.

There has been a growing consensus that we need a better understanding of the nonequilibrium dynamics of the economy (see, for example, Farmer and Geanakoplos (2008)). In particular, we would like to understand what the relaxation timescale is for nonequilibrium fluctuations to disappear (if they do). Within the limited data sample that we have shown in this paper, the relaxation time seems to be of the order of one minute. However, this can be very different in other sectors of the market.

REFERENCES

Ackert, L. F., and Tian, Y. S. (1999). Efficiency in index options markets and trading in stock baskets. Working Paper, Federal Reserve Bank of Atlanta.

Bleecker, D. (1981). *Gauge Theory and Variational Principles*. Addison-Wesley, Boston, MA.

Bondarenko, O. (2003). Statistical arbitrage and securities prices. *Review of Financial Studies* **16**(3), 875–919.

Cvitanic, J., and Zapatero, F. (2004). *Introduction to the Economics and Mathematics of Financial Markets*. MIT Press, Cambridge, MA.

Delbaen, F., and Schachermayer, W. (2008). *The Mathematics of Arbitrage*. Springer.

Fama, E. F. (1998). Market efficiency, long-term returns and behavioral finance. *Journal of Financial Economics* **49**(3), 283–306.

Farinelli, S. (2011). Geometric arbitrage theory and market dynamics. Working Paper. URL: http://ssrn.com/abstract=1113292.

Farmer, J. D., and Geanakoplos, J. (2008). The virtues and vices of equilibrium and the future of financial economics. Working Paper. URL: http://arxiv.org/abs/0803.2996.

Gatev, E., Goetzmann, W. N., and Rouwenhorst, K. G. (2006). Pairs trading: performance of a relative-value arbitrage rule. *Review of Financial Studies* **19**(3), 797–827.

Hardle, W., Hautsch, N., and Pigorsch, U. (2008). Measuring and modeling risk using high-frequency data. Discussion Paper, Humboldt University, Berlin. URL: http://sfb649.wiwi.hu-berlin.de.

Hoogland, J., and Neumann, D. (1999a). Scale-invariance and contingent claim pricing. Working Paper. URL: http://arXiv.org:cond-mat/9906048.

Hoogland, J., and Neumann, D. (1999b). Scaling invariance in finance II: path-dependent contingent claims. Working Paper. URL: http://arXiv.org:cond-mat/9907185.

Hoogland, J., and Neumann, D. (2000). Asians and cash dividends: exploiting symmetries in pricing theory. Working Paper. URL: http://arXiv.org:cond-mat/0006133.

Hoogland, J., Neumann, D., and Vellekoop, M. (2001). Symmetries in jump-diffusion models with applications in option pricing and credit risk. Working Paper. URL: http://arXiv.org:cond-mat/0108137.

Hunt, P. J., and Kennedy, J. E. (2004). *Financial Derivatives in Theory and Practice*. Wiley, New York.

Ilinski, K. (1997). Black–Scholes equation from gauge theory of arbitrage. Preprint. URL: http://arxiv.org/abs/hep-th/9712034.

Ilinski, K. (2001). *Physics of Finance: Gauge Modelling in Non-equilibrium Pricing*. Wiley, New York.

Jegadeesh, N., and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance* **48**(1), 65–91.

Kobayashi, S., and Nomizu, K. (1996). *Foundations of Differential Geometry*, Volume I. Wiley, New York.

Labordère, P.-H. (2008). *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*. Chapman & Hall/CRC, London.

Lamberton, D., and Lapeyre, B. (2007). *Introduction to Stochastic Calculus Applied to Finance*, 2nd edn. Chapman & Hall/CRC, London.

Malaney, P. N. (1996). The index number problem: a differential geometric approach. Doctoral Thesis, Harvard University Economics Department.

Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives* **17**(1), 59–82.

Musiela, M., and Rutkowski, M. (2007). *Martingale Methods in Financial Modeling*, 2nd edn. Stochastic Modeling and Applied Probability, Volume 36. Springer.

Pole, A. (2007). *Statistical Arbitrage: Algorithmic Trading Insights and Techniques*. Wiley, New York.

Sakurai, J. J. (1994). *Modern Quantum Mechanics*. Addison-Wesley, Boston, MA.

Shleifer, A., and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance* **52**(1), 35–55.

Shreve, S. E. (2000). *Stochastic Calculus for Finance*, Volumes I and II. Springer.

Smolin, L. (2009). Time and symmetry in models of economic markets. URL: http://arxiv.org/abs/0902.4274.

Sondermann, D. (2006). *Introduction to Stochastic Calculus for Finance: A New Didactic Approach*. Lecture Notes in Economics and Mathematical Systems, Volume 579. Springer.

Van Krampen, N. G. (1981). Ito versus Stratonovich. *Journal of Statistical Physics* **24**(1), 175–187.

Weinstein, E. (2006). Gauge theory and inflation: enlarging the Wu–Yang Dictionary to a unifying Rosetta Stone for geometry in application. Paper presented at the Perimeter Institute, 2006. URL: http://pirsa.org/06050010/.

Young, K. (1999). Foreign exchange market as a lattice gauge theory. *American Journal of Physics* **67**(10), 862–868.

Zhang, L. (2006). Efficient estimation of stochastic volatility using noisy observations: a multi-scale approach. URL: http://arxiv.org/abs/math/0411397.