

On the Autocorrelation of the Stock Market*

Ian Martin 

London School of Economics and Political Science

Address correspondence to Ian Martin, Department of Finance, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK, or e-mail: i.w.martin@lse.ac.uk.

Received March 5, 2019; revised August 19, 2020; editorial decision August 21, 2020; accepted August 24, 2020

Abstract

I introduce an index of market return autocorrelation based on the prices of index options and of forward-start index options and implement it at a six-month horizon. The results suggest that the autocorrelation of the S&P 500 index was close to zero before the subprime crisis but was negative in its aftermath, attaining values around -20% to -30% . I speculate that this may reflect market perceptions about the likely reaction, via quantitative easing, of policymakers to future market moves.

Do past returns on the market forecast future returns? Is the return on the market autocorrelated? It is well known that any asset return has zero risk-neutral autocorrelation (Samuelson, 1965).¹ But true autocorrelation may diverge significantly from zero—a point first made by LeRoy (1973)—and fluctuate over time. It is not clear whether one should expect positive or negative autocorrelation; indeed, both might be present simultaneously at different horizons. The former might be attributed to the influence of return-chasing investors in the investor population, as in the models of Hong and Stein (1999) and Vayanos and Woolley (2013), or to sluggish response to information; and the latter to bid-ask bounce, to overreaction as in the model of Barberis et al. (2015), or to the response of monetary authorities to fluctuations in asset prices.²

* I thank Anthony Neuberger, Anna Cieslak, Leifu Zhang, Fabio Trojani (the editor) and two referees for their comments; Can Gao for research assistance; and Dimitri Vayanos for posing the question that prompted this article. This work was carried out with the support of the Paul Woolley Centre and the ERC (Starting Grant 639744).

1 This statement is precisely true only if interest rates are deterministic; see below.

2 This is a highly incomplete, and somewhat arbitrary, list. Many other authors have studied the properties of autocorrelation; see, for example, Roll (1984), Grossman and Miller (1988), and Campbell, Grossman, and Wang (1993).

Several authors, including Fama and French (1988), Lo and MacKinlay (1988), Poterba and Summers (1988), and Moskowitz, Ooi, and Pedersen (2012), have studied the properties of realized autocorrelation of the market return, with results that vary depending on the horizon studied and on the sample period (on the latter, see Campbell, 2018, pp. 125–127). But how can we infer the forward-looking autocorrelation perceived by sophisticated investors? One straightforward approach is simply to ask investors what they think, following Shiller (1987) and others. But the expectations reported in such surveys appear to be far from rational: for example, Greenwood and Shleifer (2014) argue that times when surveyed investors are optimistic about future returns are in fact associated with low, not high, subsequent returns.

I therefore take a different approach and ask what autocorrelation must be perceived by a rational, risk-averse investor—specifically, by an unconstrained rational investor with log utility who chooses to invest his or her wealth fully in the market. It turns out to be possible to give a precise answer to this question in terms of the prices of various types of options. The fact that the autocorrelation index is computed directly from forward-looking asset prices, rather than from historical measures, is the major innovation of the article. The price to pay is that one has to accept the log investor's perspective as being a reasonable one to adopt. Nonetheless, related approaches have proved fruitful in forecasting returns on the stock market (Martin, 2017), on individual stocks (Martin and Wagner, 2019), and on currencies (Kremens and Martin, 2019); and the approach has the obvious advantage of bringing a novel type of evidence to bear on a classic question. Moreover, as in these earlier papers, we have the benefit of not requiring statistical assumptions on the underlying process (e.g., that it is stationary or ergodic). Such assumptions are widely made in the empirical literature, but they are not uncontroversial.

The theoretical results are derived in Section 1. They show that the autocorrelation index can be calculated from the prices of European index options and of forward-start index options. The latter options are relatively exotic, but I have been able to obtain indicative price quotes from a major investment bank for a small number of days between June 2007 and December 2013. Section 2 uses these prices to calculate the autocorrelation index and compares the implied forward-looking autocorrelations that emerge to the corresponding realized autocorrelations. Section 3 concludes.

1 Measuring Autocorrelation

Today is time t ; the price of the underlying asset at time t is S_t . The goal is to measure the correlation between the return on the asset over the next period, $R_{t \rightarrow t+1}$ and the return over the following period, $R_{t+1 \rightarrow t+2}$. I assume that the underlying asset does not pay dividends, so $R_{t \rightarrow t+1} = S_{t+1}/S_t$. I write the rate at which money can be risklessly invested from time u to time v as $R_{f,u \rightarrow v}$, so $R_{f,t \rightarrow t+1}$ and $R_{f,t \rightarrow t+2}$ are one- and two-period spot rates.

I write \mathbb{E}_t^* for the risk-neutral expectation operator whose defining property is that the time- t price of a time- $(t+2)$ payoff X_{t+2} is $\frac{1}{R_{f,t \rightarrow t+2}} \mathbb{E}_t^* X_{t+2}$; and cov_t^* and corr_t^* for the corresponding risk-neutral covariance and risk-neutral correlation operators.

When seeking a measure of autocorrelation that can be computed directly from asset prices, the natural first thought is to consider *risk-neutral* autocorrelation. Unfortunately, we have the following well-known result.

Result 1. *Suppose that interest rates are deterministic. Then the return on any asset has zero risk-neutral autocorrelation: $\text{corr}_t^*(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) = 0$.*

Proof. By the defining property of the risk-neutral expectation operator, we have $\mathbb{E}_t^* R_{t \rightarrow t+1} = R_{f,t \rightarrow t+1}$ and $\mathbb{E}_{t+1}^* R_{t+1 \rightarrow t+2} = R_{f,t+1 \rightarrow t+2}$. As interest rates are deterministic, the second equality implies that $\mathbb{E}_t^* R_{t+1 \rightarrow t+2} = R_{f,t+1 \rightarrow t+2}$ by the law of iterated expectations. So we can write

$$\begin{aligned} \text{cov}_t^*(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) &= \mathbb{E}_t^*[(R_{t \rightarrow t+1} - R_{f,t \rightarrow t+1})(R_{t+1 \rightarrow t+2} - R_{f,t+1 \rightarrow t+2})] \\ &= \mathbb{E}_t^*[(R_{t \rightarrow t+1} - R_{f,t \rightarrow t+1})\mathbb{E}_{t+1}^*(R_{t+1 \rightarrow t+2} - R_{f,t+1 \rightarrow t+2})] \\ &= 0, \end{aligned}$$

using the law of iterated expectations again for the second equality. □

Although interest rates are not deterministic, they are typically extremely stable by comparison with returns on stock indices, so Result 1 rules out the autocorrelation perceived by a risk-neutral investor as a useful measure. Moreover, it is easy to adapt the proof above to show that the risk-neutral autocorrelation of *excess* returns is zero even if interest rates are stochastic.

How, then, can we define a non-trivial measure of autocorrelation? This article introduces an index that can be interpreted as the autocorrelation perceived by a rational, unconstrained investor with log utility whose wealth is fully invested in the market. The next result, which is also exploited by [Martin \(2017\)](#), provides the key to calculating this quantity.

Result 2. *Let X_T be some random variable of interest whose value becomes known at time T , and suppose that we can price a claim to $X_T R_{t \rightarrow T}$ delivered at time T . Then we can compute the expected value of X_T from the perspective of an investor with log utility whose wealth is invested in the market:*

$$\mathbb{E}_t X_T = \text{time } t \text{ price of a claim to the time } T \text{ payoff } X_T R_{t \rightarrow T}. \tag{1}$$

Proof. An investor with log utility who chooses to hold the market must perceive that the return on the market is growth-optimal. As the reciprocal of the growth-optimal return is a stochastic discount factor (SDF), the right-hand side of [Equation \(1\)](#) equals $\mathbb{E}_t \left[\frac{1}{R_{t \rightarrow T}} X_T R_{t \rightarrow T} \right]$, and the result follows. □

This result provides a general strategy for inferring the *true* expectation of the log investor from traded asset prices. If we can price the claim $X_T R_{t \rightarrow T}$, then we can infer the investor’s expectation of X_T , even if X_T is not itself a tradable payoff. In particular, Result 2 will allow us to calculate $\text{corr}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2})$.

To that end, we wish to compute

$$\text{cov}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) = \mathbb{E}_t R_{t \rightarrow t+2} - \mathbb{E}_t R_{t \rightarrow t+1} \mathbb{E}_t R_{t+1 \rightarrow t+2}. \tag{2}$$

By Result 2, $\mathbb{E}_t R_{t \rightarrow T}$ is equal to the price of a claim to the square of the return on the market, $R_{t \rightarrow T}^2$. This price can be calculated by a replication argument, as in [Martin \(2017\)](#), using the fact that

$$R_{t \rightarrow T}^2 = \left(\frac{S_T}{S_t} \right)^2 = \frac{2}{S_t^2} \int_0^\infty \max\{0, S_T - K\} dK.$$

This equation expresses the desired payoff—the squared return—as the payoff on a portfolio that holds equal quantities of calls of all strikes. Thus

$$\mathbb{E}_t R_{t \rightarrow T} = \frac{2}{S_t^2} \int_0^\infty \text{call}_{t,T}(K) dK, \quad (3)$$

and setting $T = t + 1$ and $T = t + 2$ in this expression delivers the first two expectations on the right-hand side of [Equation \(2\)](#).

It is more difficult to compute $\mathbb{E}_t R_{t+1 \rightarrow t+2}$, and doing so is the main innovation of the article. In view of Result 2, to calculate this quantity we need the time- t price of a claim to $R_{t+1 \rightarrow t+2} R_{t \rightarrow t+2}$ delivered at time $t + 2$. That is, we must price a claim to $S_{t+2}^2 / (S_t S_{t+1})$.

It will turn out that we can replicate this claim using *forward-start options*. A forward-start call option that is initiated at time t , for settlement at time $t + 2$, has the payoff

$$\max\{0, S_{t+2} - K S_{t+1} / S_t\}$$

for some fixed K . The unusual feature of a forward-start option is that its strike price, $K S_{t+1} / S_t$, is not determined until the intervening time $t + 1$. (The introduction of S_t , a known constant from the perspective of time t , is simply a convenient normalization.) In contrast, the strike price of a conventional option is determined at the initiation of the trade. I write $\text{FScall}_t(K)$ for the time- t price of the above payoff, and $\text{FSput}_t(K)$ for the price of the corresponding put payoff, $\max\{0, K S_{t+1} / S_t - S_{t+2}\}$.

If we hold a portfolio consisting of $2/S_t^2 dK$ forward-start calls for each K , the portfolio payoff is

$$\frac{2}{S_t^2} \int_0^\infty \max\{0, S_{t+2} - K S_{t+1} / S_t\} dK = \frac{S_{t+2}^2}{S_t S_{t+1}}. \quad (4)$$

Since the payoff on the portfolio of forward-start calls replicates the desired payoff, the price of the payoff $S_{t+2}^2 / (S_t S_{t+1})$ is the price of the portfolio of forward-start calls, and hence

$$\mathbb{E}_t R_{t+1 \rightarrow t+2} = \frac{2}{S_t^2} \int_0^\infty \text{FScall}_t(K) dK. \quad (5)$$

Before using [Equations \(3\)](#) and [\(5\)](#) to compute the covariance $\text{cov}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2})$, it will be convenient to rearrange them by replacing in-the-money calls and forward-start calls with out-of-the-money puts and forward-start puts. For vanilla options, we can do so by exploiting put-call parity, which in our context states that

$$\text{call}_{t,T}(K) - \text{put}_{t,T}(K) = S_t - \frac{K}{R_{f,t \rightarrow T}}.$$

The next result provides the corresponding relation for forward-start options.

Result 3 (Put-call parity for forward-start options). *Let G_t be defined by the equation $FScall_t(G_t S_t) = FSput_t(G_t S_t)$ (so G_t is observable at time t , assuming the prices of forward-start options of all strikes are available). Then*

$$FScall_t(K) - FSput_t(K) = S_t - \frac{K}{G_t}. \tag{6}$$

If interest rates are deterministic, then G_t equals the forward (gross) interest rate for investment from time $t + 1$ to $t + 2$.

Proof. The time- $(t + 2)$ payoff on a portfolio that is long a forward-start call and short a forward-start put, each with strike K , is $S_{t+2} - KS_{t+1}/S_t$. It follows that

$$FScall_t(K) - FSput_t(K) = S_t - \frac{1}{R_{f,t \rightarrow t+2}} \mathbb{E}_t^* \left[\frac{KS_{t+1}}{S_t} \right] = S_t - \lambda K,$$

where λ is the time- t price of a claim to S_{t+1}/S_t delivered at time $t + 2$. We can pin down λ by applying the equation immediately above in the case $K = G_t S_t$ to conclude that $\lambda = 1/G_t$. This gives the result of Equation (6).

If interest rates are deterministic, $\lambda = 1/R_{f,t+1 \rightarrow t+2}$. For we can replicate the payoff S_{t+1}/S_t , paid at time $t + 2$, by investing $1/R_{f,t+1 \rightarrow t+2}$ in the market from time t to $t + 1$, and then at the riskless rate from time $t + 1$ to $t + 2$. □

Martin (2017) defined the volatility index $SVIX_{t,T}^2 = \frac{1}{T-t} \text{var}_t^*(R_{t \rightarrow T}/R_{f,t \rightarrow T})$:

$$SVIX_{t,T}^2 = \frac{2}{(T-t)R_{f,t \rightarrow T}S_t^2} \left[\int_0^{S_t R_{f,t \rightarrow T}} \text{put}_{t,T}(K) dK + \int_{S_t R_{f,t \rightarrow T}}^\infty \text{call}_{t,T}(K) dK \right].$$

We can define a forward volatility index $FSVIX_t$ that is new to this article:

$$FSVIX_t^2 = \frac{2}{G_t S_t^2} \left[\int_0^{S_t G_t} FSput_t(K) dK + \int_{S_t G_t}^\infty FScall_t(K) dK \right].$$

Using the put-call parity relations to substitute out calls and forward-start calls that have low strikes (i.e., are in-the-money), and then introducing these definitions, Equations (3) and (5) can be rewritten as

$$\mathbb{E}_t R_{t \rightarrow T} = R_{f,t \rightarrow T} \left(1 + (T-t) SVIX_{t,T}^2 \right) \tag{7}$$

$$\mathbb{E}_t R_{t+1 \rightarrow t+2} = G_t (1 + FSVIX_t^2). \tag{8}$$

These definitions lead to the following characterization.

Result 4. *The forward-looking autocovariance of returns, as perceived by the log investor, can be expressed in terms of spot and forward volatility indices as*

$$\begin{aligned} \text{cov}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) &= R_{f,t \rightarrow t+2}(1 + 2\text{SVIX}_{t,t+2}^2) \\ &\quad - R_{f,t \rightarrow t+1}G_t(1 + \text{SVIX}_{t,t+1}^2)(1 + \text{FSVIX}_t^2). \end{aligned} \tag{9}$$

This expression simplifies if interest rates are deterministic:

$$\text{cov}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) = R_{f,t \rightarrow t+2}(2\text{SVIX}_{t,t+2}^2 - \text{SVIX}_{t,t+1}^2 - \text{FSVIX}_t^2 - \text{SVIX}_{t,t+1}^2 \text{FSVIX}_t^2). \tag{10}$$

Proof. Equation (9) follows on substituting Equations (7) and (8) into the definition (2) of autocovariance. If interest rates are deterministic, then $R_{f,t \rightarrow t+1}G_t = R_{f,t \rightarrow t+2}$ (because, as shown in Result 3, G_t is then equal to the forward rate from $t + 1$ to $t + 2$); Equation (10) follows. \square

Result 4 has the intuitive implication that forward-looking autocorrelation is positive if long-dated options (whose prices are embedded in $\text{SVIX}_{t,t+2}$) are sufficiently expensive relative to short-dated and forward-start options (whose prices are embedded in $\text{SVIX}_{t,t+1}$ and FSVIX_t).

The remaining task is to compute $\text{var}_t R_{t \rightarrow t+1}$ and $\text{var}_t R_{t+1 \rightarrow t+2}$. As one might by now expect, the former can be computed from vanilla options and the latter from forward-start options. We have already calculated $\mathbb{E}_t R_{t \rightarrow t+1}$ and $\mathbb{E}_t R_{t+1 \rightarrow t+2}$, so it only remains to find $\mathbb{E}_t R_{t \rightarrow t+1}^2$ and $\mathbb{E}_t R_{t+1 \rightarrow t+2}^2$. By Result 2, the first of these is equal to the time- t price of a claim to $R_{t \rightarrow t+1}^3$ paid at time $t + 1$, and since

$$\left(\frac{S_{t+1}}{S_t}\right)^3 = \frac{6}{S_t^3} \int_0^\infty K \max\{0, S_{t+1} - K\} dK,$$

the desired quantity is

$$\mathbb{E}_t R_{t \rightarrow t+1}^2 = \frac{6}{S_t^3} \int_0^\infty K \text{call}_{t,t+1}(K) dK. \tag{11}$$

The remaining term, $\mathbb{E}_t R_{t+1 \rightarrow t+2}^2$, is equal to the price of a claim to $R_{t+1 \rightarrow t+2}^2 R_{t \rightarrow t+2}$ at time $t + 2$. Since

$$R_{t+1 \rightarrow t+2}^2 R_{t \rightarrow t+2} = \frac{S_{t+2}^3}{S_t S_{t+1}^2} = \frac{6}{S_t^3} \int_0^\infty K \max\{0, S_{t+2} - K S_{t+1}/S_t\} dK,$$

we have

$$\mathbb{E}_t R_{t+1 \rightarrow t+2}^2 = \frac{6}{S_t^3} \int_0^\infty K \text{FScall}_t(K) dK. \tag{12}$$

Using the put-call parity relations to replace in-the-money calls with out-of-the-money puts, Equations (11) and (12) become

$$\mathbb{E}_t R_{t \rightarrow t+1}^2 - R_{f,t \rightarrow t+1}^2 = \frac{6}{S_t^3} \left[\int_0^{S_t R_{f,t \rightarrow t+1}} K \text{put}_{t,t+1}(K) dK + \int_{S_t R_{f,t \rightarrow t+1}}^\infty K \text{call}_{t,t+1}(K) dK \right] \tag{13}$$

and

$$\mathbb{E}_t R_{t+1 \rightarrow t+2}^2 - G_t^2 = \frac{6}{S_t^3} \left[\int_0^{S_t G_t} K \text{FSput}_t(K) dK + \int_{S_t G_t}^\infty K \text{FScall}_t(K) dK \right]. \tag{14}$$

Equations (7), (8), (13), and (14) provide the ingredients needed to calculate the autocorrelation index

$$\text{corr}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) = \frac{\mathbb{E}_t R_{t \rightarrow t+2} - \mathbb{E}_t R_{t \rightarrow t+1} \mathbb{E}_t R_{t+1 \rightarrow t+2}}{\sqrt{\text{var}_t R_{t \rightarrow t+1} \text{var}_t R_{t+1 \rightarrow t+2}}}. \tag{15}$$

1.1 The Autocorrelation Index in Homogeneous Models

In many familiar theoretical models, the autocorrelation index is exactly zero. (As we will see in the next section, this is counterfactual.) As an illustration, consider the model of Black and Scholes (1973). At time $t + 1$, a forward-start call becomes identical to a vanilla call with strike KS_{t+1}/S_t , so by the Black–Scholes formula (with volatility σ and a continuously compounded riskless rate of r), the forward-start call is worth

$$S_{t+1} \Phi \left(\frac{\log \frac{S_t}{K} + r + \frac{1}{2} \sigma^2}{\sigma} \right) - K \frac{S_{t+1}}{S_t} e^{-r} \Phi \left(\frac{\log \frac{S_t}{K} + r - \frac{1}{2} \sigma^2}{\sigma} \right)$$

at time $t + 1$. As a claim to S_{t+1} at time $t + 1$ is worth S_t at time t , the above expression implies that at time t , the forward-start call is worth the same as a one-period vanilla call: $\text{FScall}_t(K) = \text{call}_{t,t+1}(K)$. It follows by put-call parity that $\text{FSput}_t(K) = \text{put}_{t,t+1}(K)$, and hence also that $\text{FSVIX}_t^2 = \text{SVIX}_{t,t+1}^2$. The autocorrelation index therefore takes a particularly simple form: as we have $\text{SVIX}_{t,T}^2 = \frac{1}{T-t} (e^{\sigma^2(T-t)} - 1)$,

$$\text{cov}_t(R_{t \rightarrow t+1}, R_{t+1 \rightarrow t+2}) = e^{2r} (e^{2\sigma^2} - 1 - 2(e^{\sigma^2} - 1) - (e^{\sigma^2} - 1)^2) = 0.$$

That is, the autocorrelation index is zero in the Black–Scholes model. Another way to make the same point is that with constant risk aversion (through log utility) and constant volatility, the risk premium is constant, so there is no room for autocorrelation to arise through the drift term. As volatility is also constant, there is no autocorrelation at all.

More generally, let us say that a model is *homogeneous* if interest rates are constant and call prices have the property that $\text{call}_{t,T}(K) = Kg(S_t/K, T - t)$ for some function g . (Many option-pricing models have this property, including the model of Black and Scholes (1973), the jump-diffusion model of Merton (1976), the variance-gamma model of Madan, Carr, and Chang (1998), and the model of Heston (1993) among others; the local volatility framework of Dupire (1994) is an example of a setting in which the homogeneity property does *not* hold.) We have the following result.

Result 5. *In a homogeneous model, the relationship between a forward-start option and a vanilla European option is trivial. That is, we have $\text{FScall}_t(K) = \text{call}_{t,t+1}(K)$ and $\text{FSput}_t(K) = \text{put}_{t,t+1}(K)$. It follows that $\text{FSVIX}_t^2 = \text{SVIX}_{t,t+1}^2$ in homogeneous models.*

Proof. A forward-start call with strike K , initiated at time t for final settlement at time $t + 2$, has the payoff $\max\{0, S_{t+2} - KS_{t+1}/S_t\}$ at time $t + 2$. From the perspective of time $t + 1$, this is equivalent to the payoff on a vanilla call with strike KS_{t+1}/S_t . By the

homogeneity assumption, at time $t + 1$, the vanilla call (and hence also the forward-start call) is worth

$$\frac{KS_{t+1}}{S_t} g\left(\frac{S_{t+1}}{KS_{t+1}/S_t}, 1\right) = \frac{KS_{t+1}}{S_t} g(S_t/K, 1).$$

The forward-start call is therefore worth $Kg(S_t/K, 1)$ at time t . In other words, by the homogeneity property, $\text{FScall}_t(K) = \text{call}_{t,t+1}(K)$. Hence $\text{FSput}_t(K) = \text{put}_{t,t+1}(K)$, by the put-call parity relations for vanilla and forward-start options. It follows that $\text{FSVIX}_t^2 = \text{SVIX}_{t,t+1}^2$ as an immediate corollary. \square

1.2 Beyond Log Utility

Our measure of implied autocorrelation exploits asset price data alone, without reference to, say, survey forecasts, or accounting or macroeconomic data. Several recent papers, including [Martin \(2017\)](#), [Kadan and Tang \(2019\)](#), [Kremens and Martin \(2019\)](#), [Martin and Wagner \(2019\)](#), and [Schneider and Trojani \(2019\)](#), have adopted similar approaches. But we have made a stronger structural assumption on the form of the stochastic discount factor than these papers do, so it is natural to wonder whether the approach of this article can be generalized to allow for, say, power utility rather than log utility.

Unfortunately, it cannot. To be concrete, suppose we wish to compute the autocorrelation perceived by a hypothetical investor who has power utility over wealth at time $t + 2$ and who chooses to invest fully in the market. As shown by [Martin \(2017, online appendix\)](#), the “easy” terms that appear in [Equation \(2\)](#)—namely, $\mathbb{E}_t R_{t \rightarrow t+1}$ and $\mathbb{E}_t R_{t \rightarrow t+2}$ —can be calculated in this more general setting.

The difficulty lies in the term $\mathbb{E}_t R_{t \rightarrow t+2}$. To compute this quantity³ with power utility (i.e., with an SDF proportional to $R_{t \rightarrow t+2}^{-\gamma}$), we would have to replicate (and hence price) the payoff $S_{t+2}^{1+\gamma}/(S_t^\gamma S_{t+1})$. To do so by holding a portfolio of $f(K)dK$ forward-start calls for each K —where $f(K)$ is some function that we can choose freely—we would need to have

$$\int_0^\infty f(K) \max\{0, S_{t+2} - KS_{t+1}/S_t\} dK = \frac{S_{t+2}^{1+\gamma}}{S_t^\gamma S_{t+1}}.$$

In the log utility case $\gamma = 1$, [Equation \(4\)](#) shows that $f(K)$ can be taken to be a constant known at time t . More generally, dividing through by S_{t+1} , we require that

$$\int_0^\infty f(K) \max\{0, S_{t+2}/S_{t+1} - K/S_t\} dK = \frac{S_{t+2}^{1+\gamma}}{S_t^\gamma S_{t+1}^2}. \quad (16)$$

The left-hand side of [Equation \(16\)](#) is a function of S_{t+2}/S_{t+1} (as S_t is known at time t). Therefore the right-hand side must also be a function of S_{t+2}/S_{t+1} ; but this forces $\gamma = 1$. Thus, log utility is the only case in which this article’s approach works.

3 Given a random variable X and SDF M , one can compute $\mathbb{E}X$ if the payoff X/M can be priced, as $\mathbb{E}X = \mathbb{E}\left(M \frac{X}{M}\right)$, and the latter is the price of the payoff X/M .

2 Empirical Results

To show how the theoretical results of Section 1 might be applied in practice, I obtained indicative price quotes for 6-month and 12-month vanilla call and put options on the S&P 500 index, together with 6-month-into-6-month forward-start options, from a major investment bank. All prices were supplied for a range of dates—June 15, 2007; June 20, 2008; November 21, 2008; February 20, 2009; December 17, 2010; July 15, 2011; December 20, 2012; and December 20, 2013—and for at-the-money, 5% and 10% out-of-the-money strikes for puts and for calls, together with the level of S&P 500 spot and the bank's internally marked 6-month and 12-month interest rate. I also obtained daily updated prices of vanilla European call and put options on the S&P 500 index from OptionMetrics in order to plot the daily time-series shown in Figure 1.

It should be emphasized at the outset that this exercise is a first step, given the limited data I have been able to obtain. A serious empirical exploration of the theoretical results of the previous section would require better data in the form both of a longer time series and of a more extensive range of strikes at each point in time.

I calculate the autocorrelation index (15) using expressions (7), (8), (13), and (14), interpolating linearly between option prices inside the range of observed strikes. I report results using two alternative methods to extrapolate option prices outside the observed range of strikes. In the first, I assume a flat volatility smile outside the range of observed strikes, following the approach of Carr and Wu (2009): in other words, for out-of-the-money puts with moneyness below the lowest observed strike, I use the Black–Scholes implied volatility at the lowest observed strike price, and for out-of-the-money calls, I use the Black–Scholes implied volatility at the highest observed strike price. In the second, I extrapolate implied volatilities linearly outside the range of observed strikes.⁴ Lastly, I set G_t equal to the forward rate from $t + 1$ to $t + 2$.

Figure 1 shows the six-monthly autocorrelation of the S&P 500 index—that is, $\text{corr}_t(R_{t-t+6\text{mo}}, R_{t+6\text{mo}-t+12\text{mo}})$ —on a sample of dates. The solid line uses a flat volatility smile for options with strikes outside the observed range; the dashed line extrapolates implied volatility linearly outside the observed range, as described in the previous paragraph. Both methods deliver similar conclusions: autocorrelation was close to zero at the beginning of the sample period, and declined sharply following the subprime crisis.

For comparison, Figure 2 plots the realized autocorrelation of six-monthly price changes, $P_{t+6\text{mo}}/P_t$, of the S&P 500 index over various time periods. I take the end-of-month level of the S&P 500 index from CRSP, over the period January 1950 to September 2019. The sample autocorrelation depends on which start month is chosen, so in each panel, I show every possible choice. For example, the autocorrelation of January–July and July–January price changes was around 0.08 over the full sample period and around -0.2 over the most recent decade, whereas the autocorrelations of the corresponding March–September and September–March price changes were around -0.1 and -0.6 , respectively.

4 I introduce a floor at zero volatility in cases where a downward-sloping smile would lead to negative implied volatilities at very high strikes; the results are not sensitive to where the floor occurs because the associated prices of deep-out-of-the-money calls are essentially zero for any reasonable level of volatility.




Figure 1 The autocorrelation of the S&P 500, $\text{corr}_t(R_{t-t+6\text{mo}}, R_{t+6\text{mo}-t+12\text{mo}})$. The solid line imposes a flat volatility smile for strikes outside the observed range of strikes. The dashed line extrapolates implied volatility linearly outside the observed range of strikes.

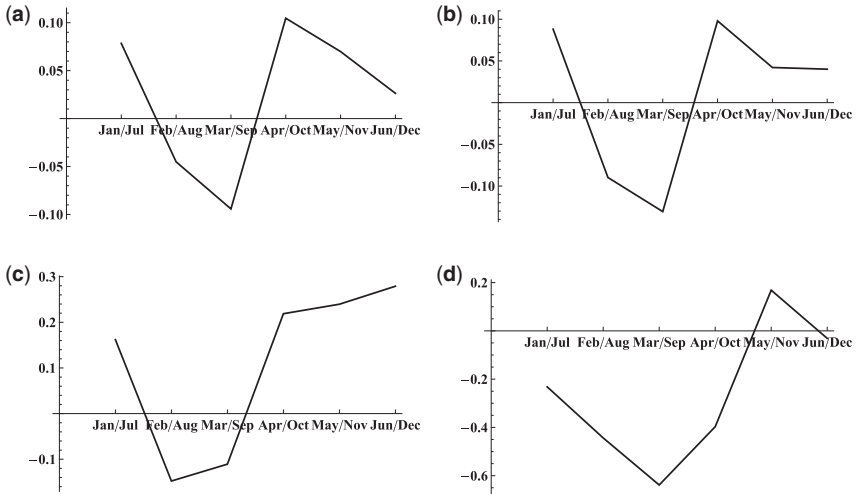


Figure 2 Realized autocorrelation in six-monthly price changes of the S&P 500 index over various time periods. (a) 1950–2019. (b) 1970–2019. (c) 1990–2019. (d) 2010–2019.

In the context of [Figure 2](#), the magnitude of the autocorrelation index shown in [Figure 1](#) appears reasonable.

Even so, it might seem, given [Figure 1](#), that strategies designed to exploit reversals should earn Sharpe ratios that are too good to be true. In order to assess this possibility, the next result shows how to use vanilla option prices to calculate the maximum attainable Sharpe ratio perceived by the log investor.

Result 6. *The maximal Sharpe ratio over the period from t to $t + n$, as perceived by the log investor, satisfies*

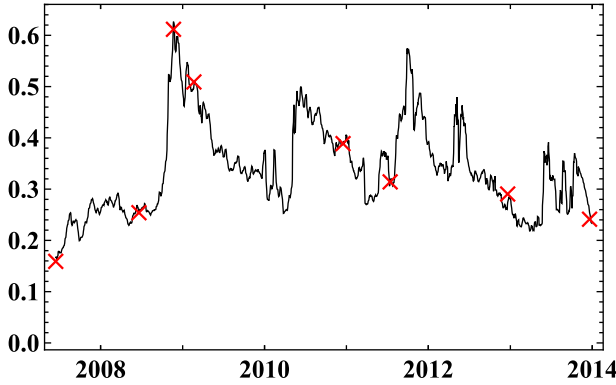


Figure 3 The maximal Sharpe ratio (one-year horizon). Crosses indicate the dates on which the autocorrelation index is computed in [Figure 1](#).

$$\max \text{ Sharpe ratio} \leq R_{f,t \rightarrow t+n} \sqrt{\int_0^\infty \frac{2S_t}{K^3} \Omega_{t,t+n}(K) dK}, \tag{17}$$

where $\Omega_{t,t+n}(K)$ is the time t price of an out-of-the-money European option with strike K expiring at time $t+n$:

$$\Omega_{t,t+n}(K) = \begin{cases} \text{put}_{t,t+n}(K) & \text{if } K \leq S_t R_{f,t \rightarrow t+n} \\ \text{call}_{t,t+n}(K) & \text{if } K > S_t R_{f,t \rightarrow t+n} \end{cases}$$

Proof. Using the result of [Hansen and Jagannathan \(1991\)](#) and the fact that $M_{t \rightarrow t+n} = 1/R_{t \rightarrow t+n}$, we have

$$\begin{aligned} \max \text{ Sharpe ratio} &\leq R_{f,t \rightarrow t+n} \sqrt{\text{var}_t \frac{1}{R_{t \rightarrow t+n}}} \\ &= \sqrt{R_{f,t \rightarrow t+n}^2 \mathbb{E}_t \frac{1}{R_{t \rightarrow t+n}^2} - 1}. \end{aligned} \tag{18}$$

By the result of [Breedon and Litzenberger \(1978\)](#), as rewritten by [Carr and Madan \(1998\)](#), the time t price of a claim to $f(S_{t+n})$ paid at time $t+n$ is

$$\mathbb{E}_t \frac{f(S_{t+n})}{R_{t \rightarrow t+n}} = \frac{f(S_t R_{f,t \rightarrow t+n})}{R_{f,t \rightarrow t+n}} + \int_0^\infty f''(K) \Omega_{t,t+n}(K) dK.$$

Setting $f(K) = S_t/K$, this implies that

$$\mathbb{E}_t \frac{1}{R_{t \rightarrow t+n}^2} = \frac{1}{R_{f,t \rightarrow t+n}^2} + \int_0^\infty \frac{2S_t}{K^3} \Omega_{t,t+n}(K) dK. \tag{19}$$

The result follows on substituting [Equation \(19\)](#) into [Equation \(18\)](#). □

Figure 3 plots the time series of the right-hand side of inequality (17) of Result 6, which provides an upper bound on the maximal Sharpe ratio at a one-year horizon. The dates on which the autocorrelation index is calculated in Figure 1 are marked with crosses. While the maximum attainable Sharpe ratio (as perceived by the log investor) spiked in late 2008, it was not implausibly high. Thus, although there are several potential ways reversal strategies might be implemented in practice, none of them has an unreasonably high Sharpe ratio from the perspective of the log investor.

3 Discussion

This article has introduced a new index of autocorrelation and constructed it at the six-month horizon using indicative prices obtained from a major investment bank on various days between mid-2007 and late 2013. Implied autocorrelation was close to zero at the beginning of the sample period but turned negative, in the range of -0.2 to -0.3 , following the subprime crisis.

Negative autocorrelation during this period may have been driven by market participants' expectations about the behavior of policymakers. The Federal Open Market Committee (FOMC) statement of September 21, 2010 contains the following paragraph,⁵ which heralded a second round of quantitative easing (QE2):

The Committee will continue to monitor the economic outlook and financial developments and is prepared to provide additional accommodation if needed to support the economic recovery and to return inflation, over time, to levels consistent with its mandate.

Based on this statement, it would have been reasonable to conclude that policy would be more expansive conditional on further declines in the market and relatively more contractionary conditional on further rises and, hence, to anticipate a decline in market autocorrelation. Indeed, Cieslak, Morse, and Vissing-Jorgensen (2018) argue that the behavior of stock returns over the "FOMC cycle" is consistent with this view (though they focus on shorter horizons and emphasize the importance of timing within the cycle). Consistent with this interpretation, the low point of the autocorrelation measure occurs in December 2010.

As the autocorrelation index depends only on asset prices, it has the great advantage of being computable, in principle, in real time. The central novel feature of the index is that it is based on the prices of forward-start index options. As shown by Hobson and Neuberger (2012), the prices of forward-start options are not tightly constrained by the prices of ostensibly closely related vanilla options. This fact is precisely what makes them interesting; nonetheless, it should be emphasized that they are exotic derivatives, with all the caveats that entails—most notably, that the forward-start option market is not nearly as liquid as the vanilla option market. A full empirical investigation of the theoretical results of this article would require considerably more data than I have been able to obtain.

5 The precise phrasing of the paragraph was discussed extensively during the meeting: see pages 78, 98, 101, 113, and 124–126 of the Transcript of the Federal Open Market Committee Meeting on September 21, 2010, which is available at <https://www.federalreserve.gov/monetarypolicy/files/FOMC20100921meeting.pdf> (last accessed on September 4, 2020). One consequence of the discussion was that the phrase "as needed" was replaced with "if needed," which was felt to emphasize the conditionality of any potential Fed action more clearly.

A further contribution of the article, however, is to point out that such options—variants of the more familiar “plain vanilla” European call and put options—have a natural economic application. It is sometimes tempting, when confronted with a cliquet, a look-back, a Napoleon, Himalayan, Bermudan, Asian, best-of, worst-of, or rainbow option, or with any other member of the bewildering menagerie of exotic derivatives, to conclude that such contracts play no more significant a role than to transfer resources between groups of quants. Precisely because there is an element of truth in this caricature, financial economists have a role to play in pointing out when some seemingly obscure derivative contract is in fact of economic interest.

References

- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer. 2015. X-CAPM: An Extrapolative Capital Asset Pricing Model. *Journal of Financial Economics* 115: 1–24.
- Black, F., and M. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81: 637–659.
- Breeden, D. T., and Litzenberger R.H. 1978. Prices of State-Contingent Claims Implicit in Option Prices. *Journal of Business* 51: 621–651.
- Campbell, J. Y. 2018. *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton, NJ: Princeton University Press.
- Campbell, J. Y., S. J. Grossman, and J. Wang. 1993. Trading Volume and Serial Correlation in Stock Returns. *The Quarterly Journal of Economics* 108: 905–939.
- Carr, P., and D. Madan. 1998. “Towards a Theory of Volatility Trading.” In R. Jarrow (ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*. London: Risk Books, pp. 417–427.
- Carr, P., and L. Wu. 2009. Variance Risk Premiums. *Review of Financial Studies* 22: 1311–1341.
- Cieslak, A., A. Morse, and A. Vissing-Jorgensen. 2018. Stock Returns over the FOMC Cycle. *Journal of Finance*, forthcoming.
- Dupire, B. 1994. Pricing with a Smile. *Risk* 7: 18–20.
- Fama, E. F., and K. R. French. 1988. Permanent and Temporary Components of Stock Prices. *Journal of Political Economy* 96: 246–273.
- Greenwood, R., and A. Shleifer. 2014. Expectations of Returns and Expected Returns. *Review of Financial Studies* 27: 714–746.
- Grossman, S. J., and M. H. Miller. 1988. Liquidity and Market Structure. *The Journal of Finance* 43: 617–633.
- Hansen, L. P., and R. Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99: 225–262.
- Heston, S. 1993. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies* 6: 327–343.
- Hobson, D., and A. Neuberger. 2012. Robust Bounds for Forward Start Options. *Mathematical Finance* 22: 31–56.
- Hong, H., and J. Stein. 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance* 54: 2143–2184.
- Kadan, O., and X. Tang. 2019. A Bound on Expected Stock Returns. *The Review of Financial Studies* 33: 1565–1617.
- Kremens, L., and I. W. R. Martin. 2019. The Quanto Theory of Exchange Rates. *American Economic Review* 109: 810–843.
- LeRoy, S. F. 1973. Risk Aversion and the Martingale Property of Stock Prices. *International Economic Review* 14: 436–446.

- Lo, A. W., and A. C. MacKinlay. 1988. Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test. *Review of Financial Studies* 1: 41–66.
- Madan, D. B., P. P. Carr, and E. C. Chang. 1998. The Variance Gamma Process and Option Pricing. *European Finance Review* 2: 79–105.
- Martin, I. W. R. 2017. What is the Expected Return on the Market? *The Quarterly Journal of Economics* 132: 367–433.
- Martin, I. W. R., and C. Wagner. 2019. What Is the Expected Return on a Stock? *The Journal of Finance* 74: 1887–1929.
- Merton, R. C. 1976. Option Pricing When Underlying Stock Price Returns Are Discontinuous. *Journal of Financial Economics* 3: 125–144.
- Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen. 2012. Time Series Momentum. *Journal of Financial Economics* 104: 228–250.
- Poterba, J. M., and L. H. Summers. 1988. Mean Reversion in Stock Prices: Evidence and Implications. *Journal of Financial Economics* 22: 27–59.
- Roll, R. 1984. A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *The Journal of Finance* 39: 1127–1139.
- Samuelson, P. A. 1965. Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review* 6: 41–49.
- Schneider, P., and F. Trojani. 2019. (Almost) Model-Free Recovery. *The Journal of Finance* 74: 323–370.
- Shiller, R. J. 1987. “Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence.” NBER Working Paper 2446.
- Vayanos, D., and P. Woolley. 2013. An Institutional Theory of Momentum and Reversal. *Review of Financial Studies* 26: 1087–1145.