

Changing Expected Returns Can Induce Spurious Serial Correlation

Kuntara Pukthuanthong, Richard Roll, and Avanidhar Subrahmanyam

January 18, 2022

Abstract

Changing expected returns can induce spurious autocorrelation in returns. We show why this happens with simple examples and investigate its prevalence in actual equity data. In a key contribution, we use shifts in *ex ante* expected return estimates from options prices, factor models, and analysts' price targets to investigate our premise. Absolute shifts in expected returns are indeed strongly and positively related to autocorrelations in the cross-section of individual stocks, as predicted by our analysis. We also show how our analysis implies spurious *cross*-autocorrelation and find supporting evidence for this phenomenon as well.

Pukthuanthong	Roll	Subrahmanyam
Department of Finance University of Missouri Columbia, MO 65211	California Institute of Technology Pasadena, CA 91125 & Emeritus, UCLA	Anderson Graduate School of Management UCLA Los Angeles, CA 90095-1481
619-807-6124	626-395-3890	310-825-5355
pukthuanthongk@missouri.edu	rroll@caltech.edu	subra@anderson.ucla.edu

Keywords: Efficient Markets, Serial Correlation, Spurious Results

JEL classifications: G10, G14

We are grateful to Michael Brennan, S.P. Kothari, Gene Fama, Andrew Lo, Ken French, and John Campbell for helpful comments.

Changing Expected Returns Can Induce Spurious Serial Correlation

Abstract

Changing expected returns can induce spurious autocorrelation in returns. We show why this happens with simple examples and investigate its prevalence in actual equity data. In a key contribution, we use shifts in *ex ante* expected return estimates from options prices, factor models, and analysts' price targets to investigate our premise. Absolute shifts in expected returns are indeed strongly and positively related to autocorrelations in the cross-section of individual stocks, as predicted by our analysis. We also show how our analysis implies spurious *cross*-autocorrelation and find supporting evidence for this phenomenon as well.

Keywords: Efficient Markets, Serial Correlation, Spurious Results

JEL classifications: G10, G14

1.1. The Phenomenon explained

For many reasons, expected returns (henceforth, ER) on assets change over time. For instance, if bonds have risk premia related positively to term, their ER decline as they mature. Equities are also subject to changing ER because lines of business evolve, firms issue more debt or retire some, they become larger, more mature and have longer and more familiar records of performance; there are probably many other driving reasons for equities.

One of the fundamental tenets of efficient markets is that returns should not be very serially dependent, i.e., autocorrelations of returns (AR, hereafter) should be rather small and insignificant. But we point out in here that this tenet can be empirically contradicted when returns are lumped together across regimes with different ER. Within each regime, when the ER is relatively constant, return AR is close to zero, yet combining different regimes and naïvely computing an unconditional AR will sometimes find it to be positive. The effect, however, is spurious in that no profit can be earned on the seemingly non-zero serial correlation.¹

To give a simple example, consider a stock that has no debt for a given decade and then borrows large amounts during the next decade.² Its debt to total asset ratio is, say, 50%, in the second decade, so it is quite a bit riskier, and its market beta is larger. Just for illustration, assume its beta is 1.0 in the first decade and then, using a simple Modigliani/Miller adjustment with no taxes, is 2.0 in the second decade. Suppose its ER in the first decade is the broad market's average, say six percent per annum or 0.5% per month. Assuming that the riskless rate is zero for ease of illustration, its ER in the second decade is twice the market's, or 1% per month.

¹ We present how our paper differs from the studies in this branch in Section I.2.

² This is not unusual among US industrial firms; see DeAngelo and Roll (2015).

However, the observed monthly return on a stock is its expected plus its unexpected return; if the latter is relatively large, any spurious AR might not be so obvious. To illustrate, the three panels of Figure 1 show simulated monthly returns for two decades that differ by doubling of the ER in the second decade, as described in the previous paragraph but with different levels of unexpected return volatility in the three panels.

Panel A illustrates the scatter diagram of returns in adjacent months when the noise is very small, 0.1% per month. The problem is readily apparent. There are two distinct (and obvious) clusters centered on means of .5% and 1%. The return AR within the first (second) decade is -0.0283 and 0.00820; neither is statistically significant. In contrast, the unconditional AR with all observations lumped into a single calculation is 0.864 and is highly significant. (The solid plotted line depicts the unconditional OLS regression of returns on first-order lagged returns while the two dashed lines depict the conditional regressions.)

Although the simple illustration in Figure 1, Panel A, deals with first-order AR, it should be obvious that second- and higher-order AR will also be spuriously overstated. This will continue until the AR lag relates a return in the second decade to a return in the first, i.e., it relates t to $t-s$ where t has one ER and $t-s$ has the other. For AR orders greater than this point, the bias will be negative, not positive.

For a slightly larger amount of volatility, a monthly standard deviation of .2%, the scatter is considerably less obvious. This is shown in Panel B which is identical in every respect to Panel A except for the variance. Even the underlying random numbers are identical.

An astute observer might notice some clustering in Panel B, but it is certainly harder to discern. The first-order return AR are exactly the same within each constant mean regime (indeed, it's easy to see that AR is unaffected by a constant ER whatever its

value.) However, the unconditional AR has fallen from 0.864 to 0.621. This might be misleading to an observer who does not notice the clustering and is eager to exploit a seeming profit.

Going further into a more realistic situation with abundant volatility, an individual stock's return standard deviation is about 30% per annum, or, with no serial correlation, about 9% per month, which is much higher than the volatilities in Panels A and B of Figure 1. Indeed, as Panel C below shows, it is really impossible to discern any clustering with such a level of volatility; the unconditional AR has fallen to -0.00656.³

Yet the clustering is still there. The underlying data are identical except for the volatility. The ER are still 0.5% and 1%, respectively in the two decades, but the noise is so overwhelming that the clustering is completely hidden. Of course, discernment would become more possible with a larger sample, but discovery can be frustratingly futile with even decades-long samples of monthly observations.

The simple examples in Figure 1 are undoubtedly too simple. There is little reason why a firm should have a constant ER in one decade and another constant in a second decade. Indeed, changes in ER must evolve for some firms at a rather steady pace while those of other firms are sudden (such as when the firm issues substantially more debt), which could occur anywhere in the record of observations. In addition to the endemic volatility of unexpected returns, such diverse experiences of various firms seem likely to render detection all that more difficult.

We should also mention that the phenomenon under study in this paper affords a new method for detecting changes in ER. If markets are efficient, serial correlation

³ The solid and dashed regression lines in Panel C cannot be distinguished from the horizontal axis.

should be small and insignificant. Consequently, under the assumption of market efficiency, return AR is intrinsic evidence that ER has changed during the sample.

Note that the argument that variations in ER can influence AR is not new, as we see in the next section. What is new in this paper is the observation that it can create spurious AR. We provide analytical verification of this observation, and, under the null of efficient markets, investigate whether the data provide evidence in favor of the existence of spurious components to AR.

Our second contribution in this paper is to link AR to shifts in *ex ante* measures of ER, obtained not only from realized returns, but also from options markets, analysts' price targets, and a standard factor model. We are not aware of earlier attempts to measure spurious AR using actual *ex ante* metrics for ER.

1.2 Literature review

The main point of this paper is that changing ER can induce spurious autocorrelations (AR), which is distinct from a spurious AR due to microstructure biases. We review the main studies on AR within this section.

Examining autocorrelation in portfolios, Fama and French (1988), Ball and Kothari (1989), and Poterba and Summers (1988) document positive AR over short horizons and negative AR over long horizon. The last study finds this evidence in both the US and 17 international stock indexes. Richardson and Stock (1989) stipulate that the joint testing of the hypothesis that portfolio autocorrelations at various lags jointly equal zero is hampered when using long-horizon data. This is due to the paucity of truly independent observations.

The evidence and explanation of the AR of individual stock returns, which is the focus of this study, is different from those of portfolio returns. While daily returns of market indices exhibit pronounced positive AR, individual stocks exhibit on average only slightly positive first-order AR. Fama (1965, 1976), French and Roll (1986), and Lo and MacKinlay (1988,1990) find that short-horizon individual security returns tend to be positively autocorrelated with no empirical evidence of significant AR for higher lags. They also show the returns of many securities are negatively autocorrelated, but larger firms' stocks tend to exhibit weak positive AR (see Chan, 1993; Sias and Starks, 1997; Chordia et al., 2005). French and Roll (1986) show that the AR(1) of the returns of the largest three equities of NYSE and Amex stocks are positive. Atchison et al (1987) report average AR close to zero.

Outside the US, Safvenblad (2000) examines the AR of individual stocks traded on the Stockholm Stock Exchange (SSE) and finds a positive AR on average with larger firms exhibiting a higher AR. In less liquid markets, positive return AR has been observed in several markets such as Austria (Huber, 1997), Finland (Berglund and Liljebloom, 1988), Israel (Ronen, 1998), and Malaysia and Singapore (Laurence, 1986).

Autocorrelation can vary with the length of time horizon. The literature has shown the time-series of short-horizon returns of individual stocks and portfolios reveal intriguing properties. Campbell (1987), French and Roll (1986), Keim and Stambaugh (1986), Conrad and Kaul (1988,1989), Lo and MacKinlay (1988) and Mech (1993) show that weekly and monthly portfolios returns are significantly and positively autocorrelated and that a positive AR is inversely related to firm size.

The underlying causes of a short-term AR could be different from those inducing a long-term AR. Ball and Kothari (1989) argue that trading frictions explain short-term AR. They argue that short intervals are characterized by noise that masks the impact of risk shifts and nonstationary ER. Studying daily returns of stocks, Brown, Harlow, and Tinic (1988) argue that short-term negative serial correlation is caused by bad

information whose effects are reversed in prices. Campbell et al. (1993) postulate that, if the main motive for trading is informationless hedging, then extreme short-term stock returns, positive or negative, will tend to be reversed when they are associated with large trading volume.

Over longer intervals, Fama and French (1988), French, Schwert, and Stambaugh (1988), and Poterba and Summers (1988) present a simple model in which aggregate returns show negative serial correlation. The model assumes that the aggregated ER is autocorrelated but mean-reverting and that revisions in these returns are independent of revisions in aggregate expected future dividends. Fama and French (1988) and Poterba and Summers (1988) argue AR is weak for the daily and weekly holding periods but stronger for long-horizon return. A large negative AR for return horizons beyond a year suggests that predictable price variation due to mean reversion accounts for large fractions of 3-5-year return variances. Fama and French (1988) posit a mean reverting process for ER. They argue that expected returns are positively correlated in the short term but mean revert in the long run, thus generating negative serial correlation over long horizons. Their paper is different from ours in two ways. First, we do not posit any particular process for ER; i.e., any change in ER will suffice. Second, they do not remark that changes in ER produce positive ARs over short horizons. In Section I.5, we show algebraically why this happens. Third, a predictable process for ER allows for profit enhancement by astute longer-term trading, while we argue that no profit is possible relative to the conditional ER in each period. Ball and Kothari (1989) provide supporting evidence of negative serial correlation over five years and argue the changing risks should be more apparent at longer horizons.

Using weekly returns of size-based portfolios over the 1962-85 period, Conrad and Kaul (1988) characterize the time pattern of ER by a stationary first-order autoregressive process. They document strong and positive first-order AR for all portfolios and even second-order AR for the smaller stock portfolios. They dismiss

microstructure as an underlying cause. However, they do not mention (1) that non-stationary ER alone is sufficient to produce the positive AR, or (2) that the observed positive AR is spurious in the sense that it does not provide a profit opportunity. In short, they present empirical evidence that ER is non-constant and that AR is positive but they do not show that the former is the cause and the latter is the effect.

Conrad and Kaul (1989) extend Conrad and Kaul (1988) by modelling the monthly ER as a decaying function of the weekly ER. They show the mean reversion in the shorter-horizon (weekly) ER explains a significant proportion of the variation in the monthly ER and that the observed return variance is explained by the variation of the ER. The impact of the latter is more pronounced for small firm portfolios. Our paper is different from these papers. Specifically, we make the point that a change in ER over time, for any reason, will produce spurious short-horizon ARs if the computation of the AR incorrectly uses a constant mean. We say the computed AR is spurious because it does not imply a profit opportunity which would be zero in an efficient market if the conditional mean is used rather than an incorrectly assumed constant mean.⁴

Campbell (1991) derives an expression for the autocovariance that depends on movement in expected returns (ER). A very similar result is in Campbell (2018, p. 138), who lists three distinct sources of autocovariance in *ex post* returns:

1. The positive covariance between dividend news and revisions in expected returns

⁴ Roll (1970) applies a similar argument to bonds and finds that much of the serial dependence in raw T-Bill returns is accounted for by changes in the level of the one-period yield.

2. The direct autocovariance of expected returns (through expected returns (ER) following an AR(1) process)

3. The negative capital loss that occurs when expected returns unexpectedly increase

They note that these three effects could conceivably be exactly offsetting but “reasonable parameter values” suggest that the third is dominant and thus actual returns are negatively autocorrelated.

We do not consider #1 or #3 because our main point is that structural (permanent) changes in ER can induce positive unconditional return autocovariance, as distinct from the conditional autocovariances induced by #1 and #3. Also, there are numerous other reasons for intertemporal movements in ER, such as permanent shifts in a firm’s risk profile or alterations in the business climate or in the firm’s policies, *inter alia*. In other words, our focus is on unconditional autocorrelation over lengthy samples associated with large and permanent changes in ER. Pastor and Stambaugh (2009) have a related approach where the expected return follows an AR(1) process but is imperfectly observable. The distinction between the conditional autocovariance and the unconditional autocovariance is relevant – but the terms that influence the conditional autocovariance also show up in a complete expression for the unconditional autocovariance. We emphasize that unconditional autocovariance has an additional effect. Pastor and Stambaugh (2012) demonstrate the distinction between the conditional and unconditional autocovariance (see also Campbell, 2018, p. 278-279).

Brennan and Wang (2010) make a different point on autocorrelations: When prices depart from fundamentals, but the departure is zero on average, expected returns will be biased due to Jensen’s inequality. They show that this bias in expected

returns is related to autocorrelations in factor model residuals. We instead make the point that changing expected returns affect return autocorrelations. Beyond the differences from earlier papers delineated above, of course, we also advance the literature by considering the relation between AR shifts and changes in actual estimates of ER from options prices and analysts price forecasts.

Cross-autocorrelations

Although our studies focus on time-series AR, a large group of studies from 1990s focuses on cross-autocorrelations (cross-AR). Lo and MacKinlay (1990) and Mech (1993) present cross-AR patterns in an attempt to investigate the sources of contrarian profits. They argue that the returns on a portfolio of small stocks is correlated with lagged returns of a large stock portfolio, but not vice versa.

The most common explanation of cross-AR is that the time series of stock prices are not sampled synchronously. Atchison et al (1987), Lo and MacKinlay (1988, 1990) and others argue that some of the cross-autocorrelations might be attributed to nonsynchronous trading problems, but to claim all of them suffer that requires markets to be unrealistically thin. Boudoukh, Richardson and Whitelaw (1994) show the cross-AR between large and small stock portfolio can be explained by the AR of small stock portfolios.

Controlling for firm size, Chordia and Swaminathan (2002) show daily and weekly returns on high volume portfolios lead returns on low volume portfolios. Nonsynchronous trading or low volume portfolio AR cannot explain this finding, but it appears that the traded returns of low volume portfolios react more slowly to information in market returns. Avramov, Chordia and Goyal (2006) show that controlling for trading volume, illiquidity causes negative cross-AR for individual stocks as price pressures caused by non-informational demands for immediacy are accommodated.

We describe the implications of our analysis for cross-AR in Section V.

1.3. Some underlying algebra for two regimes with different expected returns

In the simplest case theoretically, there are two disjoint regimes, A and B, which have different ER and (possibly) different variances. The first, Regime A, lasts for a fraction α of all observations and the subsequent Regime B follows and lasts for the complementary fraction. If ERs in the two regimes are denoted, respectively, μ_A and μ_B , the unconditional ER is $\bar{\mu} = \alpha\mu_A + (1-\alpha)\mu_B$.

An expected first-order unconditional autocovariance for regime A, computed with the unconditional ER, is $E[(R_{A,t} - \bar{\mu})(R_{A,t-1} - \bar{\mu})]$,⁵ where $R_{A,t}$ is the return during Regime A at time t and, of course, $\mu_A = E(R_{A,t})$, with analogous expressions for Regime B.⁶

This autocovariance can also be expressed as

$$E[(R_{A,t} - \mu_A + \mu_A - \bar{\mu})(R_{A,t-1} - \mu_A + \mu_A - \bar{\mu})] \quad (1.1)$$

$$= E[(R_{A,t} - \mu_A)(R_{A,t-1} - \mu_A)] + (\mu_A - \bar{\mu})[E(R_{A,t} - \mu_A) + E(R_{A,t-1} - \mu_A)] + (\mu_A - \bar{\mu})^2. \quad (1.2)$$

If markets are fully efficient, the first term in Eq. (1.2) is zero because it's a conditional autocovariance. In the second term, both expectations within the bracket are also zero. Hence the entire autocovariance is simply $(\mu_A - \bar{\mu})^2$. The analogous

⁵ Note that this is a population as opposed to a sample covariance; the latter can deviate because of sampling error. We work with population values throughout this section.

⁶ For now, we ignore the transition observation between the two regimes. It will be covered in detail subsequently.

autocovariance in Regime B is $(\mu_B - \bar{\mu})^2$. Both of these autocovariances are zero if there is no difference in ER between Regimes A and B, otherwise, they are (spuriously) positive.

The unconditional variance of returns in Regime A is

$$\begin{aligned} & E(R_{A,t} - \bar{\mu})^2 \\ &= E(R_{A,t} - \mu_A + \mu_A - \bar{\mu})^2 = \sigma_A^2 + 2(\mu_A - \bar{\mu})E(R_{A,t} - \mu_A) + (\mu_A - \bar{\mu})^2, \end{aligned} \quad (1.3)$$

where σ_A^2 denotes the conditional variance of returns during Regime A. As with the autocovariance, the second term in Eq. (1.3) is zero. Hence the unconditional variance of returns in Regime A consists of the sum of the conditional variance plus the squared difference between the conditional ER in A and the unconditional ER. Regime B is the same.

To obtain the first-order AR coefficient, ρ , we simply weight-average the terms from Regimes A and B in the numerator and denominator, while noting that the population variances of returns and lagged returns are identical. The resulting expression is

$$\rho = \frac{\alpha(\mu_A - \bar{\mu})^2 + (1-\alpha)(\mu_B - \bar{\mu})^2}{\alpha[\sigma_A^2 + (\mu_A - \bar{\mu})^2] + (1-\alpha)[\sigma_B^2 + (\mu_B - \bar{\mu})^2]} \quad (1.4)$$

Provided that the two regimes share the observations (i.e., $0 < \alpha < 1$), the AR is always positive.

The expression can be simplified because $\mu_A - \bar{\mu} = (1-\alpha)(\mu_A - \mu_B)$ while $\mu_B - \bar{\mu} = \alpha(\mu_B - \mu_A)$. Substituting and collecting terms, we obtain

$$\rho = \frac{\alpha(1-\alpha)(\mu_A - \mu_B)^2}{\alpha\sigma_A^2 + (1-\alpha)\sigma_B^2 + \alpha(1-\alpha)(\mu_A - \mu_B)^2}. \quad (1.5)$$

There are some interesting special cases. For instance, as one regime becomes much longer than the other, either α or $(1-\alpha)$ approaches zero and the AR approaches zero. Also, if the regimes are equal in length, as we assumed them to be in some of our empirical work, the AR coefficient simplifies further to

$$\rho = \frac{[(\mu_A - \mu_B) / 2]^2}{(\sigma_A^2 + \sigma_B^2) / 2 + [(\mu_A - \mu_B) / 2]^2}. \quad (1.6)$$

In any case, the coefficient is non-negative and will be biased upward if the two regimes have different ERs.

All of the above expressions for the AR coefficient reveal the influence of volatility in unexpected returns. Since the denominator contains return variances, greater volatility in return innovations implies less AR bias. The extent of the magnitudes involved is illustrated in Table 1, which is based on Eq. (1.6) with the same volatility in the two regimes.

Table 1 verifies the patterns we have already discussed by computing (1.6) with a variety of ER and volatilities. For small levels of unexpected return volatility (the first few rows of the table) and large changes in ER (the last few columns), there is substantial bias in the AR coefficient, rising to a level of more than 0.5 in an extreme case. Conversely, when volatility is substantial, there is little bias. It is still positive but would be hard to detect in a finite sample.

1.4. Autocorrelations across changes in regimes.

To this point, we have not discussed AR for an observation that occurs just as regimes shift. This cannot matter much when the sample is long and there is only one regime change. However, if regimes shift frequently, the autocovariance is affected. To see

this, a first-order unconditional autocovariance for an observation that occurs when a shift occurs from A to B is

$$E[(R_{B,t} - \bar{\mu})(R_{A,t-1} - \bar{\mu})] \quad (1.7)$$

$$= E[(R_{B,t} - \mu_B + \mu_B - \bar{\mu})(R_{A,t-1} - \mu_A + \mu_A - \bar{\mu})] \quad (1.8)$$

Reasoning along the same lines following (1.2) above, and collecting terms, the resulting autocovariance is $(\mu_B - \bar{\mu})(\mu_A - \bar{\mu}) = -\alpha(1-\alpha)(\mu_A - \mu_B)^2$, which is negative for $0 < \alpha < 1$. Hence changes of regime offset the positive AR bias discussed above. If such changes are more frequent than continuations, the AR could even be biased negatively.⁷

In general, if γ denotes the fraction of observation pairs (i.e., pairs of t and $t-1$) that are in the same regime, the overall autocovariance is $(2\gamma-1)\alpha(1-\alpha)(\mu_A - \mu_B)^2$. The fraction $(2\gamma-1)$ is a measure of regime persistence. It is close to +1 when regime shifts are infrequent while it is closer to -1 when regimes shift back and forth repeatedly and often.

Unlike the autocovariance, the unconditional variance is not affected by regime shifts. Equation (1.3) above still holds for the unconditional variance of returns in Regime A, regardless of how they are interspersed over time, and there is an identical valid expression for the unconditional variance of observations in Regime B. Hence, the general expression for the AR coefficient when a fraction $1-\gamma$ of the observations coincide with a regime shift from $t-1$ to t is, analogously to Eq. (1.5),

$$\rho = \frac{(2\gamma-1)\alpha(1-\alpha)(\mu_A - \mu_B)^2}{\alpha\sigma_A^2 + (1-\alpha)\sigma_B^2 + \alpha(1-\alpha)(\mu_A - \mu_B)^2}. \quad (1.9)$$

⁷ See Campbell (1991) for an analysis supporting negative autocorrelation.

According to (1.9), the AR coefficient in an efficient market is biased unless there are just as many reversals of regimes as continuations; in that case $\gamma = 1/2$ and $\rho = 0$ but this is probably a rather serendipitous situation.

It is interesting to speculate on whether actual regimes are more likely to persist or exhibit frequent reversals. In the case of bonds, persistence seems more likely because ER decline with shorter maturities and with lower inflation, which has historically been rather stable. For equities, driving influences such as leverage, product lines, or industry structure suggest persistence. In contrast, episodes of rapid change, such as periods around earnings announcements, could generate reversals, i.e., short-term risk and higher ER falling back to a normal level after each episode, (see Savor and Wilson (2016).) General empirical investigations of persistence and/or reversal would seem to be a useful and important topic for future research.

1.5. Multiple regimes (more than two).

A natural extension of our theory is to examine the influence of multiple regimes on the return AR coefficient. We require an expression corresponding to Eq. (1.9) with more than just regimes A and B. Assume that we know there are exactly K different possible sequential ER regimes but we do not know when any particular one of the K becomes the prevailing regime. If each can occur only once during a time series with T observations, the persistence parameter discussed previously is $\gamma = 1-(K-1)/T$. However, persistence could be larger if regimes re-occur and arrive randomly because a “shift” could simply leave the old regime in place.

Let α_t denote the fraction of observations in the t^{th} regime, $t=1,\dots,K$, and $\sum_{t=1}^K \alpha_t = 1$. Then the return AR coefficient in an efficient market is a generalization of Eq. (1.4), viz.,

$$\rho = \frac{\gamma \sum_{t=1}^K \alpha_t (\mu_t - \bar{\mu})^2 + (1-\gamma) \sum_{t=1}^{K-1} [(\mu_{t+1} - \bar{\mu})(\mu_t - \bar{\mu})] / (K-1)}{\sum_{t=1}^K \alpha_t [\sigma_t^2 + (\mu_t - \bar{\mu})^2]} \quad (1.10)$$

where μ_t is now the conditional ER in Regime t while the unconditional ER is $\bar{\mu} = \sum_{t=1}^K \alpha_t \mu_t$. Unlike our earlier development with just two regimes, the arguments of the second summation in the numerator of Eq. (1.10) are not invariably negative; indeed, if regime shifts tend to sequentially move in the same direction, most of them or perhaps all of them except one, (when μ_t moves from less than $\bar{\mu}$ to greater than $\bar{\mu}$), could be positive. Alternatively, if most regime shifts tend to be reversals, these terms could be mostly or all negative. In other words, the sign of each element in this second summation is determined by whether both μ_t and μ_{t+1} are on the same side of $\bar{\mu}$ or on opposite sides.

The first summation in the numerator of Eq. (1.10) is strictly non-negative, so the AR coefficient's sign is decreased by the prevalence of reversals, with frequency $1-\gamma$, and by whether terms in the second summation in the numerator of Eq. (1.10) tend to be negative.

Fitting Eq. (1.10) to actual data would appear to be a daunting task, though perhaps not impossible. Since the number of regimes K is an unknown, one would have to search over values of K from 1 to T while for each such value, find the best fit for the other $3K+1$ parameters (γ , $K\alpha$'s, $K\sigma$'s and $K\mu$'s.)

We leave a detailed estimation of Eq. (1.10) for future research, but note here that for large K , the second summation in the numerator of Eq (1.10) is dominated by the first. Moreover the entire fraction of Eq. (1.10) is increasing in the first summation,

which is a transformation of the volatility of ER. Hence, in our empirical work, we test for the relation between AR and the volatility of various ER proxies.

Next, we show the algebra of the volatility of the ER. For this purpose, we assume that there are multiple regimes and ER changes every month. Note that the presence of spurious serial correlation is not an artifact of the specific stochastic process followed by the mean. Even if the mean follows a standard drift with diffusion (it fluctuates with drift stochastically), our assertion still holds. Our model holds as long as there is some variation in ER; the specific process does not matter.

However, the task could be much less challenging if one is willing to make assumptions about the process that drives the temporal evolution of the ER. For example, it could change monthly but be a constant over the days within each month. When expected returns follow a general AR(1) process, they obey

$$\mu_t = c + \phi\mu_{t-1} + \sigma\varepsilon_t, \quad (1.11)$$

where the symbols are as follows:

μ_t :the expected return in period t

$\bar{\mu}$:the long-run expected return, which is, asymptotically, $c / (1 - \phi)$ ⁸

ϕ :the autoregressive parameter; $-1 < \phi < 1$ for stationarity

σ :the volatility of the AR(1) process

ε_t :a mean zero standardized IID perturbation

For simplicity, specification (1.11) ignores the chance that the expected return might become negative. This can be corrected with a square root process such as the one adopted by Cox, Ross, and Rubinstein, but only by making everything more complicated and more opaque.

⁸ This result can be confirmed by recursive substitution; i.e., $\mu_t = c + [\phi + \phi^2 + \dots + \phi^{t-1}]c + \phi^{t-1}\mu_0 + \sum_{i=0}^{t-1} \phi^i \sigma \varepsilon_{t-i}$ and so on, which delivers $E(\mu_t) = c(1 + \phi + \phi^2 + \dots) \cong c / (1 - \phi)$ because $E(\varepsilon) = 0$ and $\phi^t \mu_0 \cong 0$ for large t and $|\phi| < 1$.

It will be convenient below to rewrite equation (1.11) slightly as

$$\begin{aligned}\mu_t - \bar{\mu} &= c + \phi(\mu_{t-1} - \bar{\mu}) + (\phi - 1)\bar{\mu} + \sigma\varepsilon_t = \\ &= c + \phi(\mu_{t-1} - \bar{\mu}) + (\phi - 1)c / (1 - \phi) + \sigma\varepsilon_t = \phi(\mu_{t-1} - \bar{\mu}) + \sigma\varepsilon_t.\end{aligned}\tag{1.12}$$

Let the observed return be

$$R_t = \mu_t + \xi_t,\tag{1.13}$$

where the symbols denote:

R_t :the return in period t

ξ_t :the unexpected return in t, a mean zero IID perturbation that includes risk volatility

By assumption, $E(\xi_t \varepsilon_{t-j}) = 0 \forall j$ and $E(\xi_t \xi_{t-j}) = 0 \forall j \neq 0$.

The first order autocovariance is given by

$$\text{Cov} = E[(R_t - \bar{\mu})(R_{t-1} - \bar{\mu})] = E[(R_t - \mu_t + \mu_t - \bar{\mu})(R_{t-1} - \mu_{t-1} + \mu_{t-1} - \bar{\mu})].\tag{1.14}$$

Expanding (1.14), it becomes

$$\begin{aligned}E[(R_t - \mu_t)(R_{t-1} - \mu_{t-1})] &+ E[(R_t - \mu_t)(\mu_{t-1} - \bar{\mu})] + E[(\mu_t - \bar{\mu})(R_{t-1} - \mu_{t-1})] \\ &+ E[(\mu_t - \bar{\mu})(\mu_{t-1} - \bar{\mu})].\end{aligned}$$

Substituting, these terms become

$$E[\xi_t \xi_{t-1}] + E[\xi_t (\mu_{t-1} - \bar{\mu})] + E[(\mu_t - \bar{\mu}) \xi_{t-1}] + E\{[\phi(\mu_{t-1} - \bar{\mu}) + \sigma\varepsilon_t](\mu_{t-1} - \bar{\mu})\}.$$

By assumption, the first three of these terms are zero. The fourth term is

$$\phi E[(\mu_{t-1} - \bar{\mu})^2],\tag{1.15}$$

which is strictly positive if there is some variation in ER and ϕ is positive. The expectation in (1.15) is the conditional variance of ER as of period $t-1$. Thus, Eq. (1.15) models the (spurious) autocovariance as a function of the volatility of the expected return, σ_{μ}^2 . For most realistic applications we expect the autoregressive parameter ϕ to be positive. A constantly sign-flipping expected return is not for the most part realistic. Under $\phi > 0$ the spurious component of AR is positively related to the volatility of ER and this is the implication we test below.

II. Spurious (?) Autocorrelation in U.S. Equities.

In this section, we test our main conclusion from the previous analyses. Specifically, under the assumption that markets are efficient (so that true serial correlations are zero), we ask: Is there evidence that serial correlations do have a spurious component? As the analysis indicates, under the null of market efficiency, spurious serial correlations manifest themselves in their links with shifts in expected returns. Therefore, we consider the empirical connection between measured serial correlations and shifts in various proxies for expected returns.

II.1. Data

We collect monthly returns data from CRSP for common stocks listed on NYSE, AMEX, and NASDAQ, for the period 1966-2018. We use monthly data to mitigate the potential biases associated with nontrading and the bid-ask effect in daily data. Data with missing values are discarded. We also often use the midpoint of closing bid and ask prices to avoid these biases. We use standard methods as in Fama and French (1992) to match with Compustat in tests that use financial statements.

II.2. A first look.

We take an admittedly very rough first look at the data by computing the unconditional return AR over the entire sample of data available for each stock described in Section II.1. We then use average realized returns as proxies for expected returns and consider the link between AR and shifts in the average returns.

More specifically, we split each stock's sample in half and compute the mean return from the observations in each half. Next, we take the absolute difference between each half's mean return and rank those absolute differences across stocks, sorting them into ten deciles, where decile 1 (10) has stocks with the lowest (highest) absolute return difference.⁹ Within each decile, we compute an equally weighted average return AR. The results are shown in Figure 2.

There is clearly an upward trend in AR, albeit non-monotonic, from the lowest to the highest decile of absolute mean return differences. This clear pattern surprises us because the absolute return difference between two half-samples of each stock is a rather crude measure of changes in ER. The means are sample means, not ER, and there is no good reason why true changes in ER should be manifest in half-samples. Yet, Figure 2 seems to portray exactly the pattern we would have anticipated *a priori*.

The other surprise in Figure 2 is that AR are negative on average for every decile. The data are monthly so we did not anticipate pronounced microstructure issues, but evidently there might still be a bid/ask bounce-induced negative AR, or else there is really a tendency for stock returns to reverse themselves at a monthly frequency.

⁹ Obviously, this is a very crude measure of a change in ER. Many firms with considerable variation in ER are not captured by simple decade-long sample means, so we did not anticipate a lot of discriminatory power.

For actively traded stocks, the impact of bid/ask bounce on monthly AR can be expunged by simply eliminating the last trading day of each month. We adopt this simple cure to find the results shown in Figure 3, which corresponds to Figure 2 in every respect except for the elimination of the last trading day of every month.

The difference between Figures 2 and 3 is obvious and rather startling. The average AR in every absolute return difference decile has increased in Figure 3 and the pattern has become more monotonic.¹⁰ This suggests that AR over even monthly intervals are rather seriously contaminated by the bid/ask bounce.

To compare different groups of stocks, we next follow common practice (e.g., Fama and French, 1992) and compute separate results that exclude financial firms¹¹ and firms with negative book values. Table 2 presents t-statistics from cross-sectional regressions of unconditional AR¹² on absolute changes in mean returns computed from the first and second half of available observations for each stock. We find that the t-statistic exceeds 6.0 in each case and is highly significant (assuming there is no cross-sectional dependence in the sample AR.) This is the pattern to be anticipated if markets are efficient and there have been changes in ER for at least some firms.

Although the results in Table 2 are consistent with the phenomenon under study, we know that volatility in unexpected returns reduces the effect. We address this issue when we compute *ex ante* expected returns using various approaches in the next section.

¹⁰ Only decile 7 breaks the monotonicity

¹¹ Financial firms have 4-digit SIC codes of 6000 to 6999.

¹² To reduce the impact of bid/ask bounce, we use the mid-price of the last traded bid and ask prices

III. Analyses based on ex ante expected returns.

In Section II, we use observed mean returns as proxies for expected returns. In this section, we apply three different methodologies to estimate *ex ante* expected returns. To transform AR into a variable that is not bounded by negative and positive unity, we perform the following standard procedure:

$$\text{Transformed AR} = \text{TAR} = \frac{AR * \sqrt{T - 2}}{\sqrt{1 - AR^2}}$$

The transformed AR is not bounded, and we employ this throughout the rest of the analyses. Again, in order to avoid contamination of return serial correlations by bid-ask bounce, we compute returns from quote midpoints of the last reported bid and ask prices each day. For ER, we apply three methodologies, as detailed in the three subsections below.

III.1 Using options prices

The first approach to ex ante ER is based on Martin and Wagner (2019), who derive a formula for the ER on a stock in terms of the risk-neutral variance of the market and the stock's excess risk-neutral variance relative to that of the average stocks. Their parameters are computed from index and stock option prices. We estimate ER of stock *i* at time *t* ($E_t R_{i,t+1}$) from their Eq. (17) as follows:

$$E_t R_{i,t+1} = (R_{f,t+1} * [SVIX_t^2 + 0.5(SVIX_{i,t}^2 - \overline{SVIX_t^2})]) + R_{f,t+1}$$

where $SVIX_{i,t}^2$ is the risk-neutral variance of firm *i*, $SVIX_t^2$ is the risk-neutral variance of the index, $\overline{SVIX_t^2}$ is the average risk-neutral variance, and $R_{f,t}$ is the gross riskless rate of returns. Martin and Wagner (2019) provide the data on $SVIX_t$, $SVIX_{i,t}$, $\overline{SVIX_t^2}$,

and $R_{i,t}$.¹³ We apply their approach for several reasons. First, their ER is based on current market prices rather than historical financial characteristics. It is more parsimonious than relying on multiple characteristics. In addition, which characteristics are associated with expected returns is still inconclusive. Second, their approach does not rely only on any regression estimation, but on just the three measures of risk-neutral variance. Thus, it is less subject to bias from the estimation of factor loadings and factor construction. Third, it makes specific predictions about the relationship between ER and the three measures of risk-neutral variance whereas the relationship between characteristics and ER is less deterministic and varies with types of stocks and sample period. The caveat of this approach is the data does not start until 1996 when Option Metrics became available and the number of covered stocks is just above 800. Relying on the authors' data, our sample period ends in 2013 and is limited to firms that are constituents of S&P500.

III.2 Factor models

The second approach is factor models, which are commonly applied in academics and industry. The multi-factor arbitrage pricing theory (APT) of Ross (1976) can be expressed in two equivalent forms. The basic form expresses the total return on asset i at time t as

$$R_{i,t} = E_{i,t} + \beta_{i,1}\tilde{f}_{1,t} + \dots + \beta_{i,K}\tilde{f}_{K,t} + \tilde{\varepsilon}_{i,t} \quad (\text{III.1})$$

where the K systematic risk factors, each have a zero mean; $E_{j,t}$ is the ER on asset i at time t (it can be time-varying); and the last (firm-specific) term is an IID stochastic idiosyncratic risk component at time t for asset i . A fundamental feature of the APT is that the factor realizations are unpredictable. Any predictable component should be subsumed and incorporated by the market into expectations.

¹³ <https://onlinelibrary.wiley.com/doi/full/10.1111/jofi.12778>

If risk premiums (denoted by λ 's) are time varying, while the betas are constant, the APT's absence of arbitrage condition at time t is

$$E_{i,t} = \lambda_{0,t} + \beta_{j,1} \lambda_{1,t} + \dots + \beta_{i,k} \lambda_{k,t} \quad (\text{III.2})$$

It seems completely plausible that risk premiums are time varying; why not? For demographic and many other reasons, aggregate risk tolerances can change over time. Also, the volatilities of risk factors can change for macroeconomic reasons and actually have been observed to change in some cases (such as inflation volatility that was once considerably larger than it has been recently.) All this implies that when factor risk premiums are time varying, spurious serial correlation can be induced in securities, measured jointly across regimes.

There are various factor models. For illustration purposes, we apply 5- and 6-factor Fama-French models (FF5 and FF6). To be specific, we perform the following steps (using the FF5 factors as an example):

1. For any month j , we estimate factor loadings over the months $j-1$ to $j-60$ and save these loadings for each month j . That is, for each stock, we estimate factor loadings for FF5 ($\beta_{k,i,j}$) in every month using rolling data for 60 months where i , k and j stand for factor k , firm i , and month j .

$$R_{i,j} - R_{f,j} = \alpha + \beta_{R_m - R_f} (R_m - R_{f,j}) + \beta_{HML} HML_j + \beta_{SMB} SMB_j + \beta_{RMW} RMW_j + \beta_{CMA} CMA_j + \varepsilon_{i,j}$$

2. For each month j , we then run the Fama-MacBeth regression of monthly returns on factor loadings measured at time $j-1$ over the past 48 months and save the coefficients $\lambda_{k,i,j}$ of stock i for factor k in each month j .¹⁴

¹⁴ We follow Brennan, Chordia, and Subrahmanyam (1998) and Haugen and Baker (1996) and use the past 48 rolling months. We could instead simply estimate ER using a standard FM regression and the most recent past month but we think that approach may be noisy.

$$R_{i,j} = \hat{\alpha}_j + \hat{\lambda}_{R_m-R_f,i,j} \hat{\beta}_{R_m-R_f,i,j-1} + \hat{\lambda}_{HML,i,j} \hat{\beta}_{HML,i,j-1} + \hat{\lambda}_{SMB,i,j} \hat{\beta}_{SMB,i,j-1} + \hat{\lambda}_{RMW,i,j} \hat{\beta}_{RMW,i,j-1} + \hat{\lambda}_{CMA,i,j} \hat{\beta}_{CMA,i,j-1} + \eta_{i,j}$$

3. For the ER prediction, we use the average Fama-MacBeth coefficients and intercepts over the past 48 months (j-48 to j-1) from Step 2 and the current estimate of factor loadings from Step 1 to predict month j's ER.

$$ER_{i,j} = \frac{\sum_{j=t-48}^{t-1} \hat{\alpha}_{i,j}}{48} + \sum_k \left[\frac{\sum_{j=t-48}^{t-1} \hat{\lambda}_{k,i,j}}{48} \times \hat{\beta}_{k,i,j-1} \right]$$

III.3 Analysts' price targets

The third method is ER estimated from price targets forecasted by analysts. For this approach, we follow Engelberg, McLean, and Pontiff (2020) exactly. Specifically, we collect 12-month price targets from IBES to impute an expected return relative to the current price. The data on price targets are available monthly from the IBES starting in 1999. For each month in our sample, we use the most recent price target issued by each analyst over the last 12 months and compute the median across all such targets. The ER estimate is then simply the expected price appreciation implied by the median target price relative to the actual transaction price as of the end of the month.¹⁵

III.4 Implementation

Our basic requirement is that AR is computed with returns that are over shorter horizons than the ER. Since we compute AR from monthly returns, we split our sample into two equal halves. For each stock, we then compute the average of the

¹⁵ We ignore dividend yields for convenience.

monthly ERs within each half of the sample. We also compute the volatility of the monthly ER series for each stock.

For an ease of interpretation, we standardize an absolute change in ER (thereafter, $|\Delta ER|$) or the standard deviation of ER to have a zero mean and unit variance and then relate the transformed AR to either one of them. Table 3 shows the results of ER estimated by the Martin and Wagner (2019) approach. The $|\Delta ER|$ and volatility of ER induce AR at 10% and 1% significance levels, respectively. The economic magnitude of the change in AR is more material: A one standard deviation increase in $|\Delta ER|$ and volatility of ER increase AR by 8.38% and 24%, respectively. The impact of ER volatility on AR is almost three times stronger than that of $|\Delta ER|$. Compared to the mean of TAR of -2.88%, $|\Delta ER|$ and volatility of ER are economically and strongly impactful.

Table 4 presents the results from the ER estimated from FF5 and FF6. This estimation allows us to expand the sample to the 1982 to 2020 period and the results become stronger. Both $|\Delta ER|$ and volatility of ER estimated by FF5 and FF6 are significant at 1% and their impact is economically meaningful. A one standard deviation change in either quantity increases AR by 6.07% and 7.89%, respectively for FF5 (5.60% and 6.79%, respectively for FF6). The economic impact of volatility is about 1.4 times larger than that of $|\Delta ER|$. Similar to Table 2, the impact of $|\Delta ER|$ and volatility of ER is strong relative to the mean TAR of -9.49%.

Table 5 presents the results of ER derived from analysts' price targets and they support our conjecture as well. Both $|\Delta ER|$ and volatility of ER are significant at the 1% level. Economically, the impact of $|\Delta ER|$ and volatility of ER are lower relative to FF5 and FF6; a one standard deviation change increases AR by 1.01% and 1.64%,

respectively. The lower economic impact might be attributed to the higher imprecision of the ER generated from analysts' forecasts.¹⁶

IV. Tests with regime shifts.

In the previous sections, we examine how AR increases with $|\Delta ER|$ in two equally split sample periods. In this section, we formally identify regime shifts for individual stock returns and factors using the Quandt-Andrews test (henceforth, QA; see Quandt, 1960; Andrews, 1993; Andrews and Ploberger, 1994) and the Bai-Perron test (henceforth, BP; see Bai, 1997; Bai and Perron, 1998, 2003a and 2003b).

The QA approach tests a specified equation for one or more structural breakpoints during a time series sample by performing the single breakpoint Chow (1960) test across every two adjacent observations. Those Chow tests are then aggregated into one statistic for a test against the null hypothesis of no breakpoints anywhere in the sample. We set the QA parameters so that we infer only a single breakpoint over the entire available sample, the most significant one in terms of the largest absolute change in means. This is tantamount to applying a single Chow test.

After detecting the largest regime shift (breakpoint) using the Quandt-Andrews test, we compute the absolute difference in returns before and after the breakpoint. We regress the unconditional AR over the entire sample for the stock against this absolute difference in returns. Table 6 reports a slope coefficient t-statistic of 20.569, which supports the notion that changes in ER increase sample unconditional AR.

¹⁶ To address issues arising from cross-sectional correlation of errors, we re-estimate all regressions of Table 5 while clustering errors at the industry level, using the 48 industry-classification at Kenneth French's website (<http://tinyurl.com/mr3zpd29>). The results are virtually unchanged.

Next, we retain only the regime shifts that are significant according to the QA test and perform similar regressions with them. The t-statistic of 8.712 again supports the same inference.

Lastly, we detect significant breakpoints with the BP (1998) test. Bai (1997) and Bai and Perron (1998, 2003a) provide theoretical and computational results that extend the Quandt-Andrews framework. The ideas behind this implementation are described in Hothorn and Zeileis (2008). We find again that the absolute change in returns is significantly related to unconditional AR; regression slope t-statistic of 5.81.

V. Cross-autocorrelations.

Our basic idea also applies to cross-autocorrelation. Stocks should display spurious positive cross-autocorrelation when their expected returns are correlated and change over time. For example, two stocks with similar betas from the market model should have larger spurious positive cross-autocorrelations because their expected returns, (driven by the market's expected return and transmitted by their betas from the market model), should have stronger co-movement. Higher beta stocks have more volatile returns, both total and expected; consequently, higher beta quintiles should display more positive cross-autocorrelations unless there is no movement in the market's expected return.

The algebra of cross-autocorrelation is very similar to that presented in Section I, which derives the first-order autocorrelation for a single asset whose expected return (ER) changes once, from Regime A to Regime B. The only addendum is that we now have two assets, both of which experience a regime change at the same time. Once this is done for a single regime change, the further generalization to multiple changes, to shifts in regimes (both presented also in section I), and to non-simultaneous regime changes, are straightforward and so we do not present these analyses here.

Suppose that assets i and j both have a concurrent regime change from A to B. The conditional expected returns for i are denoted $\mu_{i,A}$ and $\mu_{i,B}$ and similarly for asset j . As before, Regime A lasts for a fraction α of all observations and Regime B lasts for the complementary fraction. Hence, the unconditional ER for i is $\bar{\mu}_i = \alpha\mu_{i,A} + (1-\alpha)\mu_{i,B}$ and similarly for j .

A first-order unconditional cross-autocovariance computed with the respective unconditional ERs, is $E[(R_{i,A,t} - \bar{\mu}_i)(R_{j,A,t-1} - \bar{\mu}_j)]$ with an analogous expression for the cross autocovariance during Regime B.¹⁷ Note also that a cross-autocovariance can be computed in two ways, with j either lagging or leading i .

The cross-autocovariance above can also be written as

$$E[(R_{i,A,t} - \mu_{i,A} + \mu_{i,A} - \bar{\mu}_i)(R_{j,A,t-1} - \mu_{j,A} + \mu_{j,A} - \bar{\mu}_j)] \quad (5.1)$$

$$\begin{aligned} &= E[(R_{i,A,t} - \mu_{i,A})(R_{j,A,t-1} - \mu_{j,A})] + (\mu_{i,A} - \bar{\mu}_i)E(R_{j,A,t-1} - \mu_{j,A}) \\ &+ (\mu_{j,A} - \bar{\mu}_j)E(R_{i,A,t} - \mu_{i,A}) + (\mu_{i,A} - \bar{\mu}_i)(\mu_{j,A} - \bar{\mu}_j). \end{aligned} \quad (5.2)$$

If markets are efficient, the first three terms in Eq. (5.2) are zero because they involve, respectively, a conditional cross-autocorrelation and two conditional expectations about conditional means. Hence, the entire cross-autocovariance is the fourth term, $(\mu_{i,A} - \bar{\mu}_i)(\mu_{j,A} - \bar{\mu}_j)$. There is an analogous expression for Regime B.

Unlike the situation for a single asset in Section I, where the analogous remaining term was a square, the sign of this remaining term in Eq. (5.2) depends on whether the two assets change expected returns in the same direction from Regime A to

¹⁷ As in section I, this is a population covariance.

Regime B. This would be the situation, for example, if both assets are driven by the market model with positive betas and the only change between regimes is in the market's expected return. However, there could be other cases where a negative cross-autocovariance might occur for two competing firms that had opposite responses to a concurrent alteration in economic circumstances. In either case, the resulting computed cross-AR does not indicate a profit opportunity.

The unconditional variance of returns for asset i in Regime A is

$$\begin{aligned} E(R_{i,A,t} - \bar{\mu}_i)^2 = \\ \sigma_{i,A}^2 + 2(\mu_{i,A} - \bar{\mu}_i)E(R_{i,A,t} - \mu_{i,A}) + (\mu_{i,A} - \bar{\mu}_i)^2, \end{aligned} \quad (5.3)$$

where $\sigma_{i,A}^2$ denotes the conditional variance of returns for asset i during Regime A.

The second term in Eq. (5.3) is zero. There are analogous expressions for asset j and Regime B.

The first-order cross-autocorrelation coefficient weight averages across regimes the remaining expression from Eq. (5.2) in the numerator and the square roots of Eq. (5.3) in the denominator, to obtain, after collecting terms and simplifying,

$$\rho = \frac{\alpha(1-\alpha)(\mu_{i,A} - \mu_{i,B})(\mu_{j,A} - \mu_{j,B})}{\{[\bar{\sigma}_i^2 + \alpha(1-\alpha)(\mu_{i,A} - \mu_{i,B})^2][\bar{\sigma}_j^2 + \alpha(1-\alpha)(\mu_{j,A} - \mu_{j,B})^2]\}^{1/2}}, \quad (5.4)$$

where the weight-averaged conditional variance is denoted $\bar{\sigma}_i^2 = \alpha\sigma_{i,A}^2 + (1-\alpha)\sigma_{i,B}^2$ for asset i and similarly for asset j . The Appendix shows that when (i) asset returns conform to the market model, (ii) all betas are positive, and (iii) only market expected returns shift across the regimes, the right-hand side of Eq. (5.4) increases in the betas of the two stocks.¹⁸ Thus, the implication is that the measured cross-AR should increase across portfolios sorted in increasing order of beta.

¹⁸ Intuitively, the numerator is proportional to the product of the betas times the squared change in the expected market return. Dividing the numerator and denominator by the product of the betas,

To assess the empirical extent of the above effect, we select stocks that have at least 30 monthly returns and compute the market model beta of each stock using all available months. We then sort stocks by their betas into quintiles and within a quintile, sort stocks by their PERMNO. We compute the cross-autocorrelations for $N/2$ sorted pairs in each quintile, where N is the number of stocks in a quintile. To avoid unnecessary sampling dependence, we use each stock only once, with its adjacent partner by the next lower PERMNO, instead of twice (with the partners both above and below it.) Then the average cross-autocorrelation is computed for the $N/2$ pairs in each quintile and transformed into a TAR between $R_{i,t}$ and $R_{j,t-1}$ and another between $R_{i,t-1}$ and $R_{j,t}$.¹⁹ As shown in Panel A of Table 7, the average TAR generally increases with beta. The result is monotonic for with CAPM betas but for $R_{i,t-1}$ with $R_{j,t}$, the average TAR is highest in Q4 (0.352) and is slightly lower in Q5 (0.330).

PERMNO might be related to firm characteristics. To guard against this possible contamination, we randomly draw, without replacement, pairs of stocks within each beta quintile and compute their cross-autocorrelation. The results are given in the Panel B of Table 7, which shows a monotonic pattern, the average cross-autocorrelation increasing with beta.

VII. Conclusions.

An absolute change in expected return, $|\Delta ER|$, or volatility in ER can induce spurious components to return autocorrelation, AR. The bias in first-order AR is generally a positive function of either the change in ER or the volatility of ER and is always a negative function of unexpected return volatility. Thus, under the null of efficient markets, measured AR should be positively related to variations in expected returns. We demonstrate this with simple examples and derive analytics that verify the

and substituting for the total variance from a market model, it can be seen that the denominator decreases in each individual beta.

¹⁹ The TAR transformation is given in the first equation of Section III.

phenomenon. We also contribute to the literature by testing for spurious AR using shifts in *actual* ex ante metrics for ER obtained from options markets, analysts' price targets, and standard factor models.

Although AR bias is mitigated by volatility in unexpected returns, we find significant evidence supporting the existence of such bias for individual US Equities. We also show that potential for spuriousness also applies to cross-autocorrelation. We observe that stocks in higher beta quintiles, whose expected returns change more with movements in the market-wide expected return, indeed exhibit more positive cross-autocorrelation, as suggested by our theory.

In future work, one could broaden our enquiry by looking into other possible measures of change in ER. These include but are not necessarily limited to: A. Slow drifts in mean returns over time; B. Events that precipitate sudden changes in ER (such as mergers, debt issuance or retirement, etc.); C. EGARCH models. In each case, we would be able to estimate the statistical significance of a change in ER and thereby sort stocks into groups that are more or less likely to have biased return AR.

In addition, it would be worthwhile to expand the sample to other asset classes, particularly bonds, which have obvious changes in ER over their lives.²⁰ The phenomenon applies to any asset class and it would be interesting to ascertain which classes are subject to it.

²⁰ Assuming, of course, that there are term premiums.

References

- Andrews, Donald W. K. and Werner Ploberger, 1994, Optimal Tests When a Nuisance Parameter is Present Only Under the Alternative *Econometrica*, 62, 1383–1414.
- Andrews, Donald W. K., 1993, Tests for Parameter Instability and Structural Change with Unknown Change Point *Econometrica*, 61, 821–856.
- Andrews, Donald W. K., 2003, Tests for Parameter Instability and Structural Change with Unknown Change Point: Corrigendum. *Econometrica* 71, 395-397.
- Atchison, Michael .D., Kurt C. Butler, and Richard R. Simonds, 1987, Non-Synchronous Security Trading and Market Index Autocorrelation. *Journal of Finance* 42, 111-118.
- Avramov, Doron, Tarun Chordia, and Amit Goyal, 2006, Liquidity and Autocorrelation in Individual Stock Returns, *Journal of Finance* 61(5), 2365-2394.
- Bai, Jushan, 1997, Estimating Multiple Breaks One at a Time, *Econometric Theory*, 13, 315–352.
- Bai, Jushan and Pierre Perron, 1998, Estimating and Testing Linear Models with Multiple Structural Changes, *Econometrica* 66, 47-78.
- Bai, Jushan and Pierre Perron, 2003a, Computation and Analysis of Multiple Structural Change Models, *Journal of Applied Econometrics* 6, 72–78.
- Bai, Jushan and Pierre Perron, 2003b, Critical Values for Multiple Structural Change Tests, *Econometrics Journal*, 18, 1–22.
- Ball, Ray, and S.P. Kothari, 1989, Nonstationary Expected returns Implications for Tests of Market Efficiency and Serial Correlation in Returns, *Journal of Financial Economics* 25, 51-74.
- Berglund, Tom, and Eva Liljeblom, 1988, Market serial correlation on a small security market: A note, *Journal of Finance* 43 (5), 1265-1274.

- Boudoukh, Jacob, Matthew P. Richardson, and Robert F. Whitelaw, 1994, A Tale of Three Schools: Insights on Autocorrelations of Short-Horizon Stock Returns, *Review of Financial Studies* 7, 539–573.
- Brennan, Michael J., Tarun Chordia, and Avanidhar Subrahmanyam, 1998, Alternative Factor Specifications Security Characteristics, and the Cross-Section of Expected Stock Returns, *Journal of Financial Economics* 49, 35-374.
- Brennan, Michael J, and Ashley W. Wang, 2010. The Mispricing Return Premium, *Review of Financial Studies* 23, 3437-3468.
- Brown, Keith, W.V. Harlow, and Seha Tinic, 1988, Risk Aversion. Uncertain Information and Market Efficiency, *Journal of Financial Economics* 22, 355-385.
- Campbell, John Y., 1987, Stock Returns and the Term Structure, *Journal of Financial Economics* 18, 373-399.
- Campbell, John Y., 1991, A Variance Decomposition for Stock Returns, *Economic Journal* 101, 157-179.
- Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton University Press, 2018.
- Chan, Kalok, 1993, Imperfect Information and Cross-Autocorrelation among Stock Prices, *Journal of Finance* 48, 1211-1230.
- Chow, Gregory C., 1960, Tests of Equality between Sets of Coefficients in Two Linear Regressions, *Econometrica* 28, 591-605.
- Chordia, Tarun, and Bhaskaran Swaminathan, 2002. Trading Volume and Cross-Autocorrelation, *Journal of Finance* 55, 913-935.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2005. Evidence on the Speed of Convergence to Market Efficiency, *Journal of Financial Economics* 76, 271-292.
- Conrad, Jennifer, and Gautam Kaul, 1988, Time Varying Expected Returns, *Journal of Business* 61, 409–425.
- Conrad, Jennifer, and Gautam Kaul, 1989, Mean Reversion in Short-Horizon Expected Returns, *Review of Financial Studies* 2, 225–240.

- Cooper, Michael, Huseyin Gulen, and Michael Schill, 2008, Asset Growth and the Cross-Section of Stock Returns, *Journal of Finance* 63, 1609-1651.
- DeAngelo, Harry, and Richard Roll, 2015, How Stable are Corporate Capital Structures? *Journal of Finance* 70, 373-418.
- Engelberg, Joseph, R. David McLean, and Jeffrey Pontiff, 2020, Analysts and Anomalies, *Journal of Accounting and Economics* 69, 101249.
- Ehsani, Sina and Juhani Linnainmaa, 2021, Factor momentum and the momentum factor, *Journal of Finance*, forthcoming.
- Fama, Eugene F., 1965, The Behavior of Stock Market Prices, *Journal of Business* 3, 34-105.
- Fama, Eugene F., 1981, Stock returns, real activity, inflation, and money. *American Economic Review* 71, 545-65.
- Fama, Eugene F., and Kenneth French, 1988, Permanent and Temporary Components of Stock Prices, *Journal of Political Economy* 96, 246-73.
- Fama, Eugene F., and Kenneth French, 1993, Common Risk Factors in Returns on Stocks and Bonds, *Journal of Financial Economics* 33, 3–56.
- French, Kenneth R. and Richard Roll, 1986, Stock Return Variances: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics* 17, 5-26.
- French, Kenneth R., G. William Schwert, and Robert Stambaugh, 1988, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-29.
- Greene, William, 2000 *Econometric Analysis*, 4th ed. Prentice-Hall International.
- Haugen, Robert, and Nardin Baker, 1996, Commonality in the Determinants of Expected Stock Returns, *Journal of Financial Economics* 41, 401-439.
- Hothorn, Torsten, and Achim Zeileis, 2008, Generalized Maximally Selected Statistics, *Biometrics* 64, 1263–1269.
- Huber, Peter, 1997, Stock Market Returns in Thin Markets: Evidence from the Vienna Stock Exchange, *Applied Financial Economics* 7, 493-498.
- Keim, Donald B. and Robert F. Stambaugh. 1986, Predicting Returns in the Stock and Bond Markets, *Journal of Financial Economics* 17, 357-390.

- Laurence, Martin M., 1986, Weak Form Efficiency in the Kuala Lumpur and Singapore Stock Markets, *Journal of Banking and Finance* 10, 431-445.
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* 1, 41–66.
- Lo, Andrew W., and A. Craig MacKinlay, 1990, When are Contrarian Profits Due to Stock Market Overreaction, *Review of Financial Studies* 3, 175–205.
- Lo, Andrew W., and A. Craig MacKinlay, 1990, An Econometric Analysis of Non-Synchronous trading. *Journal of Econometrics* 45, 181-211.
- Martin, Ian W.R., and Christian Wagner, 2019, What is the Expected Return on a Stock? *Journal of Finance* 74, 1887-1929.
- Mech, Timothy .S., 1993, Portfolio Return Autocorrelation, *Journal of Financial Economics* 34, 307-344.
- Pastor, L., and R. Stambaugh, 2009. Predictive systems: Living with imperfect predictors, *Journal of Finance* 64, 1583-1628.
- Pastor, L., and R. Stambaugh, 2012. Are stocks really less volatile in the long run, *Journal of Finance* 67, 431-478.
- Poterba, James A., and Lawrence.M. Summers, 1988. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics* 22, 27-59.
- Quandt, Richard E., 1960, Tests of the Hypothesis that a Linear Regression System Obeys Two Separate Regimes, *Journal of the American Statistical Association* 55, 324-330.
- Richardson, Matthew, and James Stock, 1989, Drawing Inferences from Statistics based on Multiyear Asset Returns, *Journal of Financial Economics* 25, 323-348.
- Roll, Richard, 1970, *The Behavior of Interest Rates: An Application of the Efficient Market Model to U.S. Treasury Bills*, Basic Books, Inc., Publishers, New York
London

- Ronen, Tavy., 1998, Trading Structure and Overnight Information: A Natural Experiment from the Tel- Aviv Stock Exchange, *Journal of Banking and Finance* 22, 489-512.
- Ross, Stephen, 1976, The Arbitrage Theory of Capital Asset Pricing, *Journal of Economic Theory* 13, 341-360.
- Safvenblad, Patrik, 2000, Trading Volume and Autocorrelation: Empirical Evidence from the Stockholm Stock Exchange, *Journal of Banking & Finance* 24, 1275-1287.
- Savor, Pavel, and Mungo Wilson, 2016, Earnings Announcements and Systematic Risk, *Journal of Finance* 71, 83-138.
- Sias, Richard W., and Laura T. Starks, 1997, Return autocorrelation and institutional investors. *Journal of Financial Economics* 46, 103-131.
- Zellner, Arnold, 1962, An Efficient Method of Estimating Seemingly Unrelated Regression Equations and Tests for Aggregation Bias, *Journal of the American Statistical Association*, 57, 348-368.

Appendix Spurious Cross-Autocorrelations and Betas

If changes in expected return are entirely driven by changes in the market's expected return transmitted through the assets' betas, which are assumed to be constant, then Eq. (5.4) can be written as

$$\rho = \frac{\alpha(1-\alpha)\beta_i\beta_j[E(R_{M,B}) - E(R_{M,A})]^2}{\{[\bar{\sigma}_i^2 + \alpha(1-\alpha)\beta_i^2[E(R_{M,B}) - E(R_{M,A})]^2][\bar{\sigma}_j^2 + \alpha(1-\alpha)\beta_j^2[E(R_{M,B}) - E(R_{M,A})]^2]\}^{1/2}} \quad (\text{A.1})$$

This is the same as

$$\rho = \frac{\beta_i\beta_j\phi}{\{[\bar{\sigma}_i^2 + \beta_i^2\phi][\bar{\sigma}_j^2 + \beta_j^2\phi]\}^{1/2}} \quad (\text{A.2})$$

where $\phi = \alpha(1-\alpha)[E(R_{M,B}) - E(R_{M,A})]^2$.

The weighted average conditional variance for asset i is

$$\bar{\sigma}_i^2 = \alpha[\beta_i^2\text{Var}(R_{M,A}) + \text{Var}(\varepsilon_{i,A})] + (1-\alpha)[\beta_i^2\text{Var}(R_{M,B}) + \text{Var}(\varepsilon_{i,B})]$$

or in simpler notation, $\bar{\sigma}_i^2 = \beta_i^2[\alpha\sigma_{M,A}^2 + (1-\alpha)\sigma_{M,B}^2] + \alpha\sigma_{\varepsilon_{i,A}}^2 + (1-\alpha)\sigma_{\varepsilon_{i,B}}^2$

where $\varepsilon_{i,A}$ is the market model idiosyncratic disturbance for asset i in Regime A and similarly for asset j.

If the market's variance and the idiosyncratic variances are the same in the two regimes, this expression simplifies to

$$\bar{\sigma}_i^2 = \beta_i^2 \sigma_M^2 + \sigma_\varepsilon^2$$

If the betas of the two assets are the same (and are constants over time) and both weight-averaged variances are also the same, Eq. (A.2) reduces to

$$\rho = \beta^2 \phi / (\beta^2 \sigma_M^2 + \sigma_\varepsilon^2 + \beta^2 \phi) = 1 / (\sigma_M^2 / \phi + \sigma_\varepsilon^2 / \beta^2 \phi + 1),$$

so the cross-AR increases with beta (assuming that idiosyncratic volatility is non-zero.)

More generally, dividing numerator and denominator of Eq. (A.2) by its numerator, it becomes

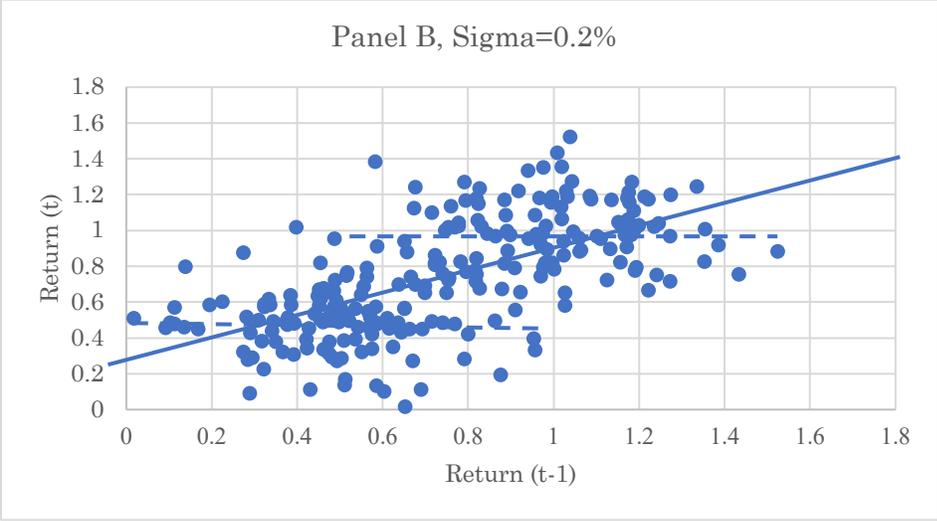
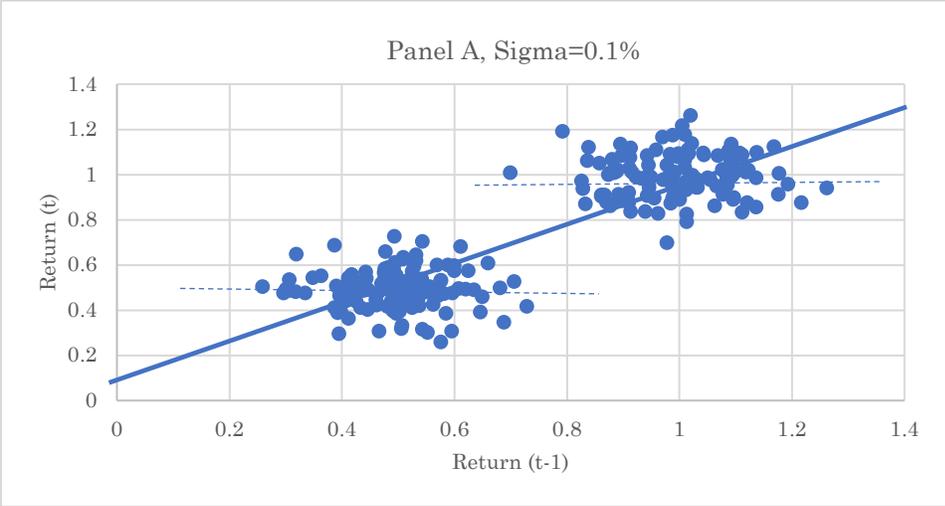
$$\rho = \frac{1}{\phi} \left\{ (\phi + \sigma_M^2)^2 + (\phi + \sigma_M^2) \left(\frac{\sigma_{\varepsilon_i}^2}{\beta_i^2} + \frac{\sigma_{\varepsilon_j}^2}{\beta_j^2} \right) + \frac{\sigma_{\varepsilon_i}^2 \sigma_{\varepsilon_j}^2}{\beta_i^2 \beta_j^2} \right\}^{-1/2}. \quad (\text{A.3})$$

A simplified version of Eq. (A.3) is

$$\rho = \phi \left[(\phi + \sigma_M^2 + \sigma_{\varepsilon_i}^2 / \beta_i^2) (\phi + \sigma_M^2 + \sigma_{\varepsilon_j}^2 / \beta_j^2) \right]^{-1/2} \quad (\text{A.4})$$

If both betas or either one increases, *ceteris paribus*, the denominator decreases and the cross-AR increases. Hence, groups of stocks ranked by beta should display higher cross-ARs in the higher betas groups. The same would be true (and would perhaps be even more accurate) for groups ranked by the ratio of beta to idiosyncratic variance.

Figure 1. Simulated scatter diagram for successive returns with two regimes that differ only in ER. In one regime, the ER is 0.5% per period and in the other regime, it is 1% per period. The three panels illustrate the impact of volatility, in standard deviation per period, which is 0.1%, 0.2% and 8.66%, respectively, in Panels A, B and C. The dashed lines show the conditional autoregression within each regime while the solid line shows the unconditional autoregression using all observations regardless of regime.



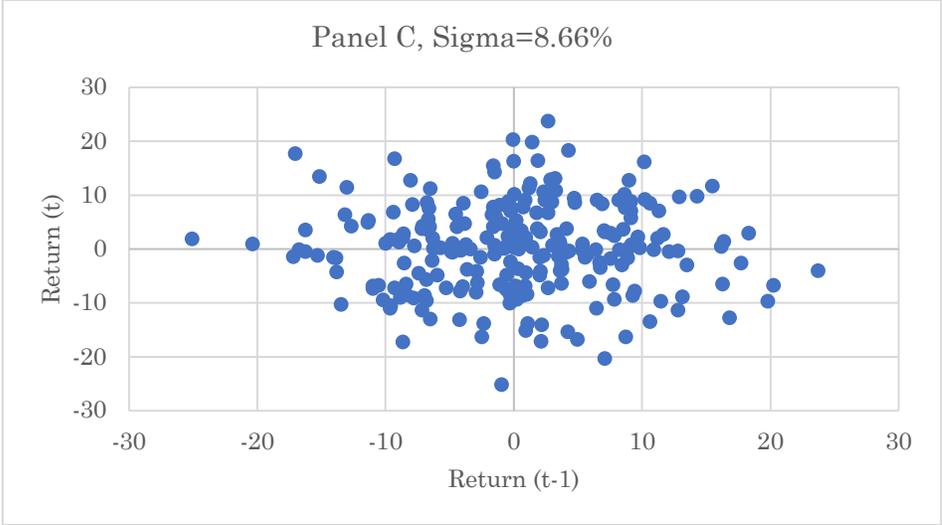


Figure 2. For each non-financial firm common stock listed on the NYSE, AMEX, and NASDAQ, with at least 2 years of presence in Compustat, the unconditional monthly return first-order AR is computed over the entire available sample. Then the stock's sample is divided in half and mean returns are computed from each half. The absolute difference between the two halves' mean returns is computed and ranked across stocks, then sorted into deciles, where decile 1 (10) has stocks with the lowest (highest) absolute mean return difference. An equally weighted average return AR is computed for each decile and is plotted here.

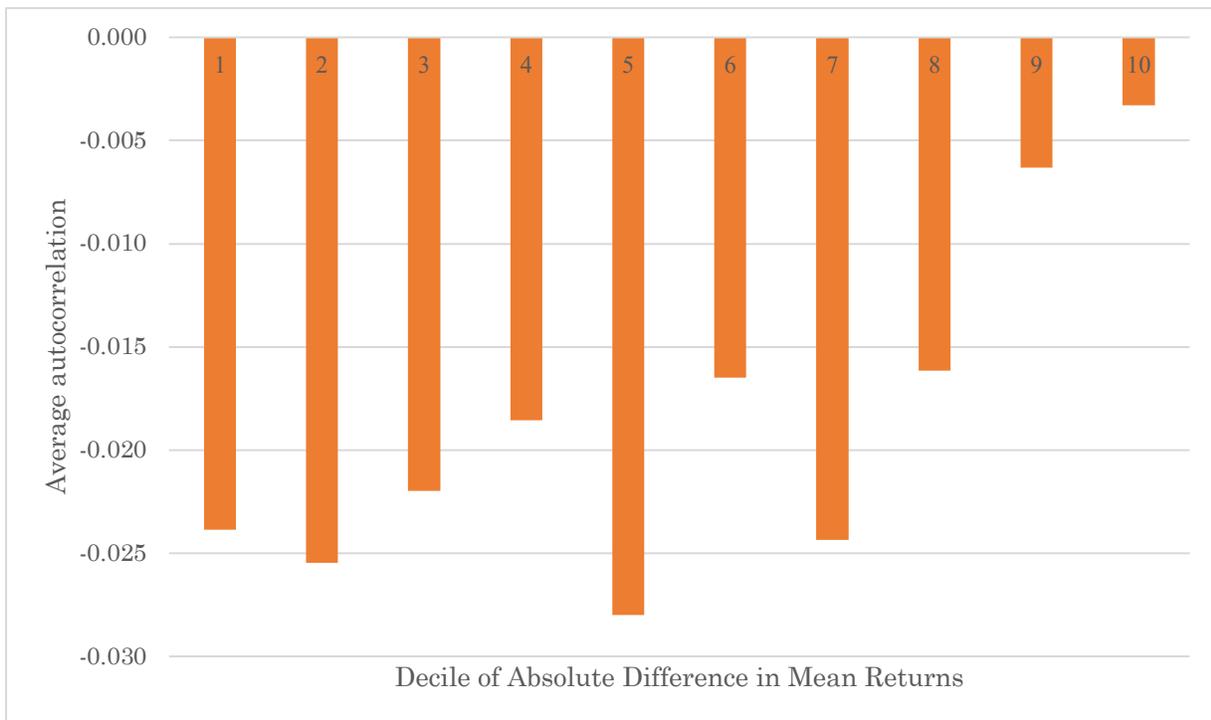


Figure 3. For each non-financial firm common stock listed on the NYSE, AMEX, and NASDAQ, with at least 2 years of presence in Compustat, the unconditional monthly return first-order AR is computed over the entire available sample while excluding the last trading day of each month. Then the stock's sample is divided in half and mean returns are computed from each half. The absolute difference between the two halves' mean returns is computed and ranked across stocks, then sorted into deciles, where decile 1 (10) has stocks with the lowest (highest) absolute mean return difference. An equally weighted average return AR is computed for each decile and is plotted here.

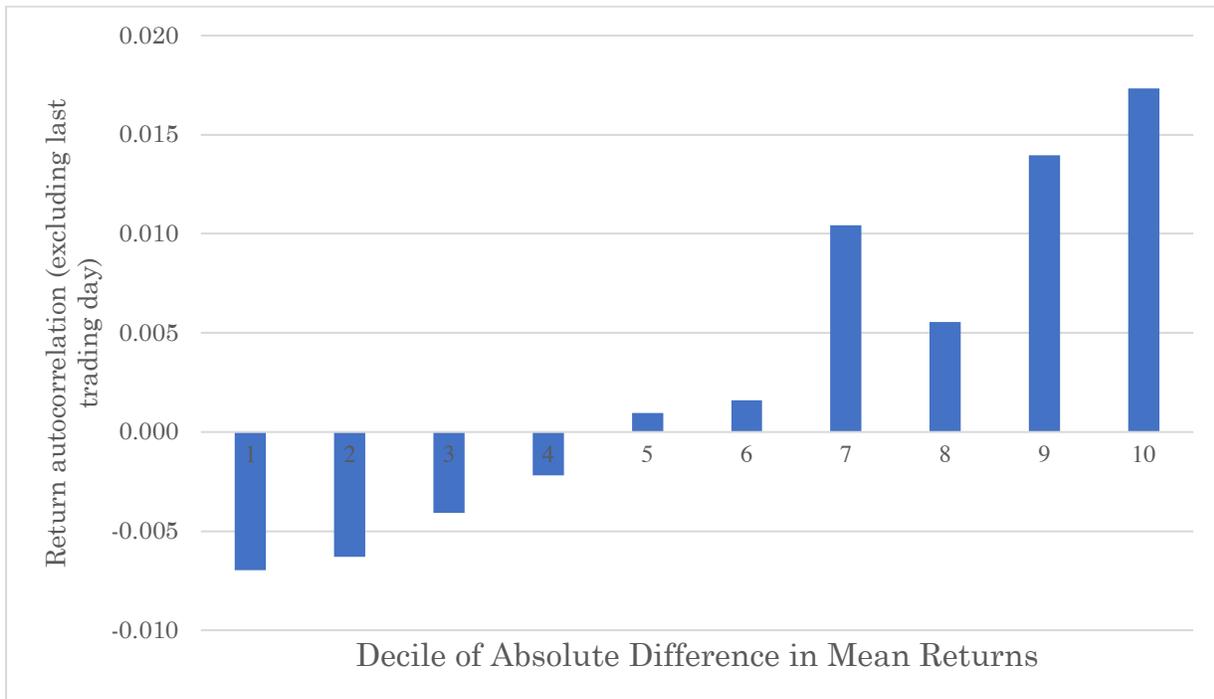


Table 1

Theoretical Spurious Return Autocorrelation Coefficients for Different Combinations of Return Volatility and Expected Return Differences in Two Regimes of Equal Length,
Monthly Observations, Annualized Means and Volatilities

This table is generated from simulated returns with two regimes that differ only in ER. Sigma is the standard deviation of returns per period in percent per annum.

Sigma	Expected Return Difference, (% per annum)						
	0	5	10	15	20	25	30
0	N/A	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
5	0	0.0383	0.1324	0.2478	0.3600	0.4581	0.5397
10	0	0.0099	0.0368	0.0761	0.1233	0.1745	0.2267
15	0	0.0044	0.0167	0.0353	0.0588	0.0859	0.1153
20	0	0.0025	0.0094	0.0202	0.0340	0.0502	0.0683
25	0	0.0016	0.0061	0.0130	0.0220	0.0327	0.0448
30	0	0.0011	0.0042	0.0091	0.0154	0.0229	0.0315
35	0	0.0008	0.0031	0.0067	0.0113	0.0170	0.0234
40	0	0.0006	0.0024	0.0051	0.0087	0.0130	0.0180

Table 2
 Bivariate Regression of Return Autocorrelation on the Absolute Change
 In Mean Observed Return between Two Half-Samples;
 All Available Firms, Non-Financials, and Always Positive Book Values

The sample of available observations for each stock is split into two halves and the explanatory variable is the absolute difference in the sample mean return in the two halves. The dependent variable is the unconditional autocorrelation with all monthly observations in each stock's entire sample expunging the last trading day of every month. The regression t-statistics are reported in the first row. N is the number of firms included in a regression. To be included, a stock must have 120 available monthly observations during 1966-2018 inclusive. ***, **, and * present 1%, 5%, and 10% significance levels, respectively.

	All Firms	Positive Book Value	Non-Financial	Non-Financial Positive Book Value
T-Statistic	8.506***	7.242***	7.572***	6.410***
N	7,679	6,262	6,300	5,001

Table 3
 Bivariate Regression of Return Autocorrelation on the Absolute Change
 In (and Volatility of) Expected Return based on Martin and Wagner (2019)

The sample is provided by Martin and Wagner (2019) consisting of 834 stocks from January 1996 to September 2013. The dependent variable is the unconditional AR, transformed based on the first equation in Section III. The sample is split into two halves and the explanatory variable is the absolute difference in the ER ($|\Delta ER|$) in the two halves or the volatility of ER. The ERs are computed based on Martin and Wagner (2019). Both $|\Delta ER|$ and volatility of ER are standardized to have a zero mean and unit variance. All stock prices are the mid-price of bid and ask closing prices to avoid bid-ask bounce. ***, **, and * present 1%, 5%, and 10% significance levels, respectively.

	Coef	t Stat	R2	Adj R ²	N
Intercept	-0.029	-0.661	0.46%	0.33%	804
$ \Delta ER $	0.084*	1.921			
	Coef	t Stat	R2	Adj R ²	
Intercept	-0.029***	-0.672	3.75%	3.63%	804
Volatility of ER	0.240***	5.593			

Table 4
 Bivariate Regression of Return Autocorrelation on the Absolute Change
 In (and Volatility of) Expected Returns estimated from 5- and 6-factor Fama-
 French Models

The sample of available observations for each stock is split into two halves and the explanatory variable is the absolute difference in the sample mean return in the two halves ($|\Delta ER|$) or the volatility of ER. The ER are estimated by 5- and 6-factors Fama French models. Both $|\Delta ER|$ and volatility of ER are standardized to have a zero mean and unit variance. The dependent variable is the unconditional AR with all monthly observations in each stock's entire sample and it is transformed based on the first equation in Section III. To be included, a stock must have 120 available monthly observations during 1967-2018 inclusive. All stock prices are the mid-price of bid and ask closing prices to avoid bid-ask bounce. ***, **, and * present 1%, 5%, and 10% significance levels, respectively. The number of observations is 14,002.

	Coeff	t-stat	N	R ²	Adj
Fama-French 5 factors model					
Intercept	-0.096***	-9.051	14,002	0.24%	0.23%
$ \Delta ER $	0.061***	5.784			
Intercept	-0.095***	-9.058	14,002	0.40%	0.39%
Volatility of ER	0.079***	7.495			
Fama-French 6 factors model					
Intercept	-0.095***	-9.049	14,002	0.20%	0.20%
$ \Delta ER $	0.056***	5.338			
Intercept	-0.095***	-9.053	14,002	0.30%	0.29%
Volatility of ER	0.068***	6.472			

Table 5
 Bivariate Regression of Return Autocorrelation on the Absolute Change
 In (and Volatility of) Expected Return Estimated from Analysts' Forecasts

The sample of available observations for each stock is split into two halves and the explanatory variable is the absolute difference in the sample mean return in the two halves ($|\Delta ER|$) or the volatility of ER. The ER are estimated by 12-month price target from analysts' forecasted data in IBES. Both $|\Delta ER|$ and volatility of ER are standardized to have a zero mean and unit variance. The dependent variable is the unconditional AR with all monthly observations in each stock's entire sample and it is transformed based on the first equation in Section III. All stock prices are measured as the mid-point of bid and ask quotes to avoid bid-ask bounce. To be included, a stock must have 120 available monthly observations during 1997-2018 inclusive. ***, **, and * present 1%, 5%, and 10% significance levels, respectively. The number of observations is 7,004.

	Estimate	t-stat	R ²	Adj R ²
Intercept	-0.057***	-31.967	0.46%	0.44%
$ \Delta ER $	0.010***	5.637		
Intercept	-0.057***	-32.08	1.21%	1.20%
Volatility of ER	0.016***	9.194		

Table 6
 Bivariate Regression of Return Autocorrelation on the Absolute Change In Mean
 Return during Periods before and after Identified Mean Return Breakpoints;
 All Available Firms and Factors

The sample of available observations for individual stocks and factors is split based on either all breakpoints or significant breakpoints identified by the Quandt-Andrews test developed by Andrews (1993, 2003), and Quandt (1960) and significant breakpoints suggested by the Bai-Perron (1998) test. Then a cross-sectional regression is computed where the explanatory variable is the absolute difference in the sample mean returns for a stock or a factor before and after the breakpoint, and the dependent variable is the unconditional AR with all monthly observations in each stock or factor's entire sample expunging the last trading day of every month. The regression slope t-statistic is reported in the first row. N is the number of firms included in a regression. To be included, a stock must have 120 available monthly observations during 1966-2018 inclusive. ***, **, and * present 1%, 5%, and 10% significance levels, respectively.

	All breakpoints from Quandt-Andrews test	Significant breakpoints from Quandt-Andrews test	Significant breakpoints from Bai-Perron test
T-Statistic	20.569***	8.712***	5.806***
N	9,186	2,481	782

Table 7
Cross-autocorrelation and Market-model Betas

This table presents cross-autocorrelation among stocks sorted by their market-model beta into quintiles. For those stocks that have at least 30 monthly returns observations, we estimate the market-model beta. We sort stocks into beta quintiles and, within a quintile, randomly draw a pair of stocks and compute their cross-autocorrelations. As a robustness check, we also sort stocks by their PERMNOs and compute cross-autocorrelations of stocks that have adjacent PERMNOs. We then take an average of the cross-autocorrelations across stock pairs within a quintile. For both methods, we do not use any stock twice and estimate cross autocorrelations between $R_{i,t}$ and $R_{j,t-1}$ (and $R_{i,t-1}$ and $R_{j,t}$) where i and j represent stocks i and j , and t is the time index, in units of a month. The autocorrelation is transformed based on the first equation in Section III.

Beta Quintile	Q1	Q2	Q3	Q4	Q5
Mean market-model beta	0.014	0.656	1.000	1.381	2.378
Pairs chosen randomly within each beta quintile (without replacement within each Beta Quintile)					
			TAR(R_{1t}, R_{2t-1})		
Mean TAR	0.071	0.111	0.181	0.210	0.244
# of unique pairs	2,011	2,015	2,011	2,016	2,006
			TAR(R_{1t-1}, R_{2t})		
Mean TAR	0.083	0.157	0.188	0.203	0.211
# of unique pairs	2,013	2,019	2,014	2,017	2,004
Pairs with adjacent PERMNOs					
			TAR(R_{1t}, R_{2t-1})		
Mean TAR	0.186	0.197	0.238	0.302	0.331
# of unique pairs	1,866	1,852	1,856	1,834	1,915
			TAR(R_{1t-1}, R_{2t})		
TAR	0.154	0.240	0.261	0.352	0.33
# of unique pairs	1,866	1,854	1,857	1,838	1,917