

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

Changing Expected Returns Can Induce Spurious Serial Correlation

Kuntara Pukthuanthong
University of Missouri

Richard Roll
California Institute of Technology



SOCIAL SCIENCE WORKING PAPER 1446

July, 2019

Changing Expected Returns Can Induce Spurious Serial Correlation

By

Kuntara Pukthuanthong and Richard Roll

July 10, 2019

Abstract

Changing expected returns can induce spurious autocorrelation in returns and possible bias in other financial models. We show why this happens with some simple examples and then investigate its prevalence in actual equity data. Assets may undergo changes in expected returns for a variety of reasons. Hence, autocorrelations computed from extended records can be subject to spurious bias. The bias might be difficult to measure because of noise but a re-sampling method (the bootstrap) discloses its rather ubiquitous presence in US equity data. Turning the phenomenon on its head, return serial correlation in an efficient market is evidence of changing expected returns. Estimated risk measures such as “beta” might also be subject to bias induced by non-stationary mean returns, but the direction of the bias is more ambiguous.

Co-Author	Affiliation	Voice	E-Mail
Pukthuanthong	University of Missouri Columbia MO 65211	619-807-6124	PukthuanthongK@ Missouri.Edu
Roll	Caltech Pasadena CA 91125	626-395-3890	RRoll@Caltech.Edu

Key words: Efficient Markets, Serial Correlation, Spurious Results

I.1. The Phenomenon Explained

For many reasons, expected returns on assets change over time. For instance, if bonds have risk premia related positively to term, their expected returns decline as they mature. Equities also are subject to changing expected return 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 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 expected returns. Within each regime, when the expected return is relatively constant, return autocorrelation is close to zero, yet combining different regimes and blindly computing an unconditional autocorrelation 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.

To give a simple example, consider a stock that has no debt for a given decade and then borrows like mad for the next decade.¹ 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 expected return 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 expected return in the second decade is twice the market's, or 1% per month.

However, the observed monthly return on a stock is its expected plus its unexpected return; if the latter is relatively large, any spurious autocorrelation 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 expected return in the second decade, as described in the previous paragraph but with different levels of noise due to unexpected return volatility in the three panels.

¹ This is not unusual among US industrial firms; see DeAngelo and Roll [2015.]

Panel A illustrates the scatter diagram of returns in adjacent months when the noise is very small, 0.1% per month. In this case, anyone who bothers to plot a scatter diagram will easily observe the problem. There are two distinct (and obvious) clusters centered on means of .5% and 1%. The return autocorrelation within the first (second) decade is -0.0283 and 0.00820; neither is statistically significant. In contrast, the unconditional autocorrelation 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 autocorrelation, it should be obvious that second- and higher-order autocorrelations will also be spuriously overstated. This will continue until the autocorrelation 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 expected return and $t-s$ has the other. For autocorrelations orders greater than this point, the bias will be negative, not positive.

For a slightly larger amount of noise, 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 noise volatility. 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 autocorrelations are exactly the same within each constant mean regime (indeed, it's easy to see that autocorrelation is unaffected by a constant expected return whatever its value.) However, the unconditional autocorrelation has fallen from 0.864 to 0.621. This might be misleading to a careless observer who fails to notice the clustering and is eager to exploit a seeming profit.

Going further into a more realistic situation with abundant noise, an individual stock often has a return standard deviation of 30% per annum, or, with no serial correlation, about $30 / \sqrt{12} \approx 8.66\%$ per month, 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 noise; the unconditional autocorrelation has fallen to -0.00656.

Yet the clustering is still there. The underlying data are identical except for the volatility. The expected returns 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 expected return in one decade and another constant in a second decade. Indeed, changes in expected return must evolve for some firms at a rather stately 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 noise, such diverse experiences of various firms seem likely to render detection all that more difficult.

One method that offers a possibility of uncovering the true pattern is the bootstrap. Re-sampling with the bootstrap essentially allows one to expand the sample size indefinitely. We employ the bootstrap here with actual data.

Finally, we should mention that the phenomenon under study in this paper affords a new method for detecting changes in expected return. If markets are efficient, serial correlation should be small and insignificant. Consequently, under the assumption of market efficiency, return autocorrelation is intrinsic evidence that expected returns have changed during the sample.

I.2. Some Underlying Theory for Two Regimes with Different Expected Returns

In the simplest case theoretically, there are two disjoint regimes, A and B, which have different expected returns 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 expected returns in the two regimes are denoted, respectively, μ_A and μ_B , the unconditional expected return is $\bar{\mu} = \alpha\mu_A + (1 - \alpha)\mu_B$.

An expected first-order unconditional autocovariance for regime A, computed with the unconditional expected return, is $E[(R_{A,t} - \bar{\mu})(R_{A,t-1} - \bar{\mu})]$,² 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.³

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 (1.2) is zero because it's the 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 autocovariance in Regime B is $(\mu_B - \bar{\mu})^2$. Both of these autocovariances are zero if there is no difference in expected return 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 (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 expected return in A and the unconditional expected return. Regime B is the same.

To obtain the first-order autocorrelation 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

² 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.

³ For now, we ignore the transition observation between the two regimes. It will be covered in detail subsequently.

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

Provided that the two regimes share the observations (i.e., $0 < \alpha < 1$), the autocorrelation 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 autocorrelation approaches zero. Also, if the regimes are equal in length, as we assumed them to be in the much of our empirical work, the autocorrelation 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 expected returns.

All of the above expressions for the autocorrelation coefficient reveal the influence of noise. Since the denominator contains return variances, noisier returns will have less autocorrelation bias. The extent of the magnitudes involved is illustrated in the following theoretical table, which is based on equation (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 expected returns and volatilities. For small levels of noise (the first few rows of the table) and large changes in expected return (the last few columns), there is substantial bias in the autocorrelation coefficient, rising to a level of more than 0.5 in an extreme case. Conversely, when noise is substantial, there is little bias. It is still positive but would be hard to detect in a finite sample.

I.3. Autocorrelations across changes in regimes.

To this point, we have not discussed autocorrelation for an observation that occurs just as regimes shift. This cannot not 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 autocorrelation bias discussed above. If such changes are more frequent than continuations, the autocorrelation could even be biased negatively.

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 autocorrelation coefficient when a fraction $1-\gamma$ of the observations coincide with a regime shift from $t-1$ to t is, analogously to (1.5),

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

According to (1.9), the autocorrelation 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 expected returns 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 expected return 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.

I.4. Multiple Regimes (More Than Just Two).

A natural extension of our theory is to examine the influence of multiple regimes on the return autocorrelation coefficient. We require an expression corresponding to (1.9) with more than just regimes A and B. Assume that we know there are exactly K different possible sequential expected return 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 autocorrelation coefficient in an efficient market is a generalization of (1.4), viz.,

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

where μ_t is now the conditional expected return in Regime t while the unconditional expected return 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 (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 (1.10) is strictly non-negative, so the autocorrelation 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 (1.11) tend to be negative.

Fitting (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 this to future research.

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

II.1. Data

We collect monthly returns data from CRSP 1966 to 2018 inclusive, for non-financial firms common stocks listed on NYSE, AMEX, and NASDAQ. We follow Fama and French (1993) and Cooper, Gulen and Schill (2009) for screening criteria. To mitigate backfilling bias, a firm must be listed in Compustat two years before it is included in the data set. That is, A small stock is defined as one with a market cap less than 20th percentile of those on the NYSE the year it was first listed. We then merge stock return data from CRSP with firm characteristics from the Compustat Daily Updates - Fundamentals Annual (file name "funda"). We exclude missing data and data having infinite values after doing the divisions described in the Appendix.

II.2. A first look.

We take an admittedly very rough first look at the data by computing the unconditional return autocorrelation over the entire sample of data available for each stock described in section II.1. We then 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.⁴ Within each decile, we compute an equally weighted average return autocorrelation. The results are shown in Figure 2.

There is clearly an upward trend in autocorrelation, 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 expected return. The means are sample means, not expected returns, and there is no good reason why true changes in expected returns 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 autocorrelations are negative on average for every decile. The data are monthly so we did not anticipate pronounced micro-structure issues, but evidently there might still be a bid/ask bounce-induced negative autocorrelation, or else there is really a tendency for stock returns to reverse themselves at a monthly frequency.

For actively traded stocks, the impact of bid/ask bounce on monthly autocorrelations 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 autocorrelation in every absolute return difference decile has increased in Figure 3 and the pattern has become more monotonic.⁵ This suggests that autocorrelations over even monthly intervals are rather seriously contaminated by the bid/ask bounce. Consequently, we will henceforth use only autocorrelations computed from monthly returns after expunging the last trading day of each month.

⁴ Obviously, this is a very crude measure of a change in expected return. Many firms with considerable variation in expected returns are not captured by simple decade-long sample means, so we did not anticipate a lot of discriminatory power.

⁵ Only decile 7 breaks the monotonicity

II.3. Variety in the Phenomenon across Groups of Stocks with Disparate Attributes.

To compare different groups of stocks, we follow common practice and eliminate financial firms⁶ and firms whose book values are not positive in some years. Data definitions and sources are in Table A-1 of the Appendix.

Table 2 presents t-statistics from cross-sectional regressions of unconditional autocorrelations⁷ on absolute changes in mean returns computed from the first and second half of available observations for each stock. There are 7,679 stocks in all and the number in each of the three sub-groups is provided in the table. The t-statistic exceeds 6.0 in each case and is highly significant (assuming there is no cross-sectional dependence in the sample autocorrelations.) This is the pattern to be anticipated if markets are efficient and there have been changes in expected returns for at least some firms.

Although the results in Table 2 are consistent with the phenomenon under study, we know that noise in unexpected returns reduces the effect. Larger firms are less volatile but they might also be less likely to experience large changes in expected return, which are offsetting effects. Hence, studying firms grouped by size might reveal which effect is dominant. Accordingly, in Table 3, firms in Table 2 are separated into larger and smaller size groups. A firm is assigned to the small size group if its market cap in its first available year was less than the 20th percentile of market caps for firms traded on NYSE (Hou, Xue, and Zhang, 2018). Table A-2 in the Appendix presents information about the distribution of sizes corresponding to each of the Panels of Table 3.⁸

All of regressions in Table 3 have significant t-statistics and large firms have larger t-statistics than small firms. Small firms are more numerous, which would increase t-statistics, *ceteris paribus*, and if smaller firms are more likely to experience changes in expected returns, their t-statistics should be correspondingly larger, *ceteris paribus*. Hence, the results in Table 3 suggest that noise is the dominant influence on autocorrelation bias relative to changes in expected returns.

⁶ Financial firms have 4-digit SIC codes of 6000 to 6999.

⁷ To reduce the impact of bid/ask bounce, we eliminated the last trading day of each month before computing the autocorrelations of monthly returns (in conformity to the findings described in section II.B.)

⁸ The distribution of Ln(Size) is in the first line of each Panel of Table A-2. The other seven lines in each panel pertain to other attributes of the firms.

Another possible way to mitigate the noise problem is to examine firm attributes that might be more related to expected returns than to unexpected returns. Based on past asset pricing literature, (Fama and French, 1993; Cooper, Gulen, and Schill, 2008; Hand et al., 2018) we decided to examine ten firm attributes: Size, Book/Market, Sales Growth, Profit Ratio, Operating Profit Ratio, Asset Growth, Age, Systematic Volatility (Beta) and two measures of the Price/Earnings Ratio (Diluted and Not Diluted.) Table A-1 in the appendix defines these attributes more precisely and gives their sources. Table A-2 provides their cross-firm distributions.

We compute a simple bivariate regression with the firm's unconditional return autocorrelation as the dependent variable and the attribute's absolute change as the explanatory variable. Table 4 presents the results.

The absolute change in size and operating profit ratio decreases autocorrelation whereas the absolute change in asset growth, beta, and P/E diluted increases autocorrelation. Comparing small and large firms, the absolute change in sales growth and operating profit ratio is significant for small firms whereas the absolute change in size and beta is significant for large firms. Asset growth and age are impactful for both types of firms. Interestingly, for small firms, the absolute change in B/M increases autocorrelation when we exclude firms with negative book value of equities. For large firms, the results are consistent across subsamples.

III. The Bootstrap.

The empirical results reported in Section II above are consistent with the phenomenon under study, but they are exposed to at least one confounding statistical issue; viz., t-statistics from cross-sectional regressions generally assume independence. To the extent that serial correlation estimates are not independent across firms, the t-statistics could be overstated. Moreover, we know that those estimates are influenced by noise, a nuisance. In an effort to overcome both problems, we now resort to the bootstrap to obtain a sharper insight into the ubiquity of the phenomenon.

We retain the same eight groups of stocks as in Table 3 and perform independent bootstraps on each group. Given the group sample sizes, ranging from just above 2,000 to more than 4,000, we re-sample (with replacement) from each group using 200 stocks at a time. For each group of 200, we compute the cross-sectional regression between the unconditional return first-order autocorrelation and the absolute change in mean return between each stock's first and second half of available observations. We repeat this 1,000 times and calculated the average of the slope across all 1,000 regressions. Since the 1,000 re-samples are chosen randomly, the results are independent, which allows a simple calculation of the standard error of the mean from the individual re-sampled regressions.⁹

Table 5 presents the results.¹⁰ For all eight groups, the mean slope is highly statistically significant, with t-statistics ranging from around 42 to around 51. Not surprisingly, large firms exhibit considerably more significance and larger coefficients than smaller firms, probably because the latter are noisier. The bootstrap, however, seems able to discern the underlying pattern that is mitigated by noise in simpler approaches.

In terms of the consequences for return serial correlation, the explanatory variable is in units of percent per month. So, for example, if a stock doubles its expected return from one half-period to the other, say 1% per month in one half period and 2% in the other, the absolute return difference would be on the order of 1%. The impact on autocorrelation is then this absolute return difference multiplied by the slope coefficient. As examples, this would amount to 0.264 for large firms and 0.232 for small firms.

IV. Bias in an estimated binary coefficient induced by non-stationary mean returns, with an application to the CAPM “Beta”

Following the approach in section I.2, we now consider the slope coefficient in a binary model when the two variables, x and y , experience simultaneous disjoint regimes with different expected

⁹ I.e., the standard error of a bootstrapped mean statistics is the cross-replication standard deviation divided by the square root of 1,000.

¹⁰ For reasons of space, we report only the mean slope t-statistic and its own t-statistic but other results are available on request from the authors.

returns. As in I.2, the regimes are labeled A and B and the unconditional expected returns are weighted averages of the respective conditional expected returns, with a fraction α in regime A and its complement in regime B; i.e., the unconditional expected returns for x and y are, respectively,

$$\begin{aligned}\bar{\mu}_x &= \alpha\mu_{A,x} + (1-\alpha)\mu_{B,x} \\ \bar{\mu}_y &= \alpha\mu_{A,y} + (1-\alpha)\mu_{B,y}\end{aligned}$$

The covariance calculated for regime A using the unconditional expected returns from above, is

$$E[(x_A - \bar{\mu}_x)(y_A - \bar{\mu}_y)] = E[(x_A - \mu_{A,x})(y_A - \mu_{A,y})] + (\mu_{A,x} - \bar{\mu}_x)(\mu_{A,y} - \bar{\mu}_y) \quad (4.1)$$

and there is an identical expression for regime B with the subscript changed.

The first term on the right side of (4.1) is the conditional covariance during regime A. If this conditional covariance does not change between regimes A and B and, for notational simplification is denoted C, then the unconditional covariance across both regimes is

$$C + \alpha(\mu_{A,x} - \bar{\mu}_x)(\mu_{A,y} - \bar{\mu}_y) + (1-\alpha)(\mu_{B,x} - \bar{\mu}_x)(\mu_{B,y} - \bar{\mu}_y). \quad (4.2)$$

Substituting for the four terms in parentheses in (4.2) and simplifying, the overall unconditional covariance becomes

$$C + \alpha(1-\alpha)(\mu_{A,x} - \mu_{B,x})(\mu_{A,y} - \mu_{B,y}) . \quad (4.3)$$

Equation (4.3) shows that the unconditional covariance is upward biased, relative to the conditional covariance, if the expected returns of x and y are displaced in the same direction between regimes A and B, (assuming, of course, that $1 < \alpha < 1$.)

Taking x as the explanatory variable in a bivariate model and following the same line of development while assuming its conditional variance is a constant V, the unconditional variance is

$$V + \alpha(1-\alpha)(\mu_{A,x} - \mu_{B,x})^2 . \quad (4.4)$$

Whenever there has been a regime change of a finite length, this variance is also upward biased (relative to the conditional variance V), so the unconditional slope coefficient, the ratio of (4.3) to (4.4) has an undetermined bias relative to the conditional coefficient, C/V .¹¹

An interesting special case involves the CAPM “Beta” coefficient, which is computed by relating x to the market return, M , and y to the return on some other asset or portfolio, j . If we employ excess returns beyond the riskless rate F , then $x = R_M - R_F$ and $y = R_j - R_F$ and the CAPM stipulates that $R_j - R_F = \beta(R_M - R_F)$ or, in the simpler notation, $y = \beta x$.

If the market’s expected returns is changes from regime A to B while the subject asset j ’s expected return changes only because of its dependence on the market through β , equation (4.3) becomes

$$C + \alpha(1 - \alpha)\beta(\mu_{A,x} - \mu_{B,x})^2 \quad (4.5)$$

while the conditional slope is $C/V = \beta$. Defining a construct $Z \equiv \alpha(1 - \alpha)(\mu_{A,x} - \mu_{B,x})^2$, the ratio of the unconditional covariance (4.5) to the unconditional variance (4.4) is

$$\frac{C + \beta Z}{V + Z} = \frac{\beta(V + Z)}{V + Z} = \beta.$$

Hence, in this particular application, the unconditional slope (β) is unbiased and equal to the conditional slope. But, there remains a caution in that excess returns must be employed in the estimation of β . If raw returns are employed, a bias could be introduced because of variation in the riskless rate. Moreover, if the subject asset happens to exhibit a time-varying “alpha”¹² its estimated β would typically be biased except in some serendipitous situation where the two biases, in the covariance and the variance, just happen to be offsetting. If any time-varying alpha moves in the same direction as the market’s expected return, it’s likely that β will be biased upward. This would explain the long-standing puzzle in asset pricing that higher β ’s have insufficiently higher mean returns.

¹¹ Note that the ratio of (4.3) to (4.4) is identical to expression (1.5) in section I for the autocorrelation coefficient after setting the conditional autocovariance, C , to zero, making the conditional variance constant, $V = \sigma_A^2 = \sigma_B^2$ and, of course, realizing that $x=y=R$. In the case of the autocorrelation coefficient, there is no ambiguity in the bias.

¹² Extra-risk expected return.

V. Future research.

In this paper, we have only studied empirically the simplest and most naïve measure of changes in expected return, the sample mean return difference between the first and last half of each stock's available observations.

In future work, we intend to broaden this enquiry by looking into other possible measure of change in expected returns. These include, but are not necessarily limited to: A. Time series slow drifts in mean returns; B. Events that precipitate sudden changes in expected returns (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 expected returns and thereby sort stocks into groups that are more or less likely to have biased return autocorrelations.

In addition, we would like to expand the sample to other asset classes, particularly bonds, which have obvious changes in expected returns over their lives.¹³ The phenomenon applies to any asset class and it would be interesting to ascertain which classes are subject to it.

VI. Conclusions.

Changes in expected returns can induce spurious bias in return autocorrelations. We demonstrate this with simple examples, derive analytics that verify the phenomenon, and find its presence in a broad sample of US equities. The bias in first-order autocorrelations is generally a positive function of the change in expected returns and is always a negative function of idiosyncratic volatility (noise.)

The autocorrelation bias is positive when the expected return change occurs across long calendar periods such as, for example, between two decades. However, the bias can be reversed and become

¹³ Assuming, of course, that there are term premiums.

negative when reversals in expected returns are frequent, such as, for example, across contiguous months.

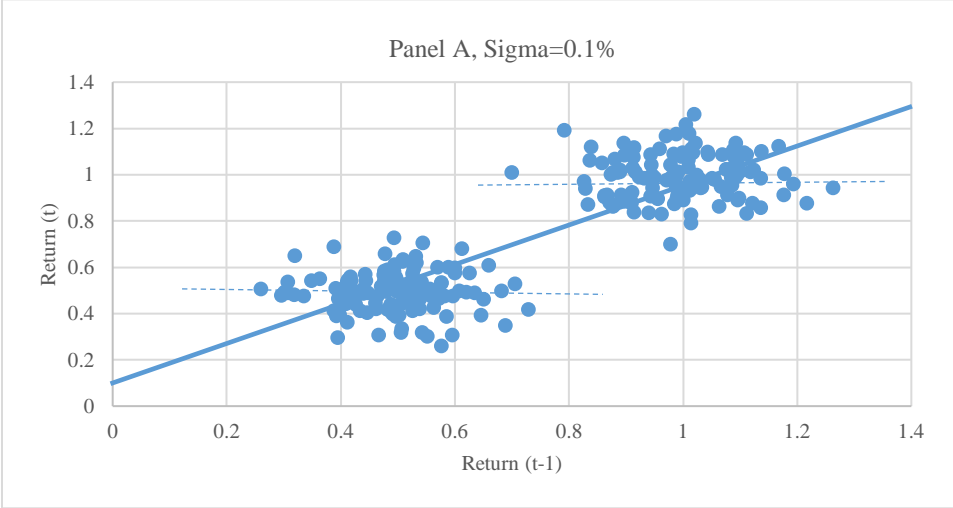
Empirically for US equities, we find that the first-order autocorrelation bias is generally positive. Moreover, larger firms tend to have more bias relative to smaller firms. This is consistent with noise being dominant over changes in expected returns (which are probably more pronounced for smaller firms.)

We also find that autocorrelation bias is related to changes in particular firm attributes that have been shown in previous literature to be related cross-sectionally to expected returns.

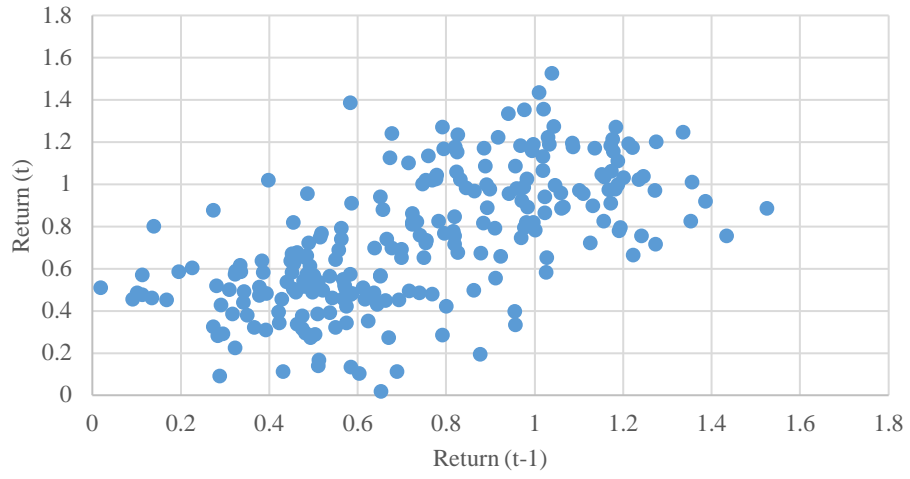
Although autocorrelation bias is mitigated by noise, we find, using the bootstrap, that it is prevalent and highly significant for US Equities.

However, the bias in autocorrelation induced by non-stationary expected returns does not necessarily extend to other financial models. For instance, the CAPM beta might be biased but the direction is rather ambiguous. Moreover, despite changes over time in expected returns, the CAPM beta is not biased at all when excess returns are used in its estimation, provided there are no extra-risk returns (“alphas),

Figure 1. Simulated scatter diagram for successive returns with two regimes that differ only in expected return. In one regime, the expected return 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.



Panel B, Sigma=0.2%



Panel C, Sigma=8.66%

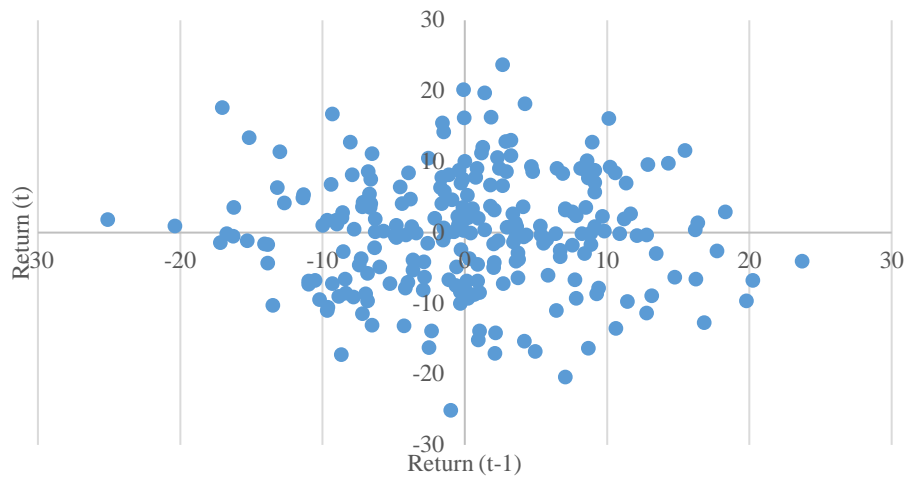


Figure 2. For each non-financial firm common stock listed on the NYSE, AMEX, and NASDAQ, with at least 2 years of presence in Compustat, the unconditional monthly return first-order autocorrelation is computed over the entire available sample. Then the stock's sample is divided in half and mean returns 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 autocorrelation is computed for each decile and is plotted here.

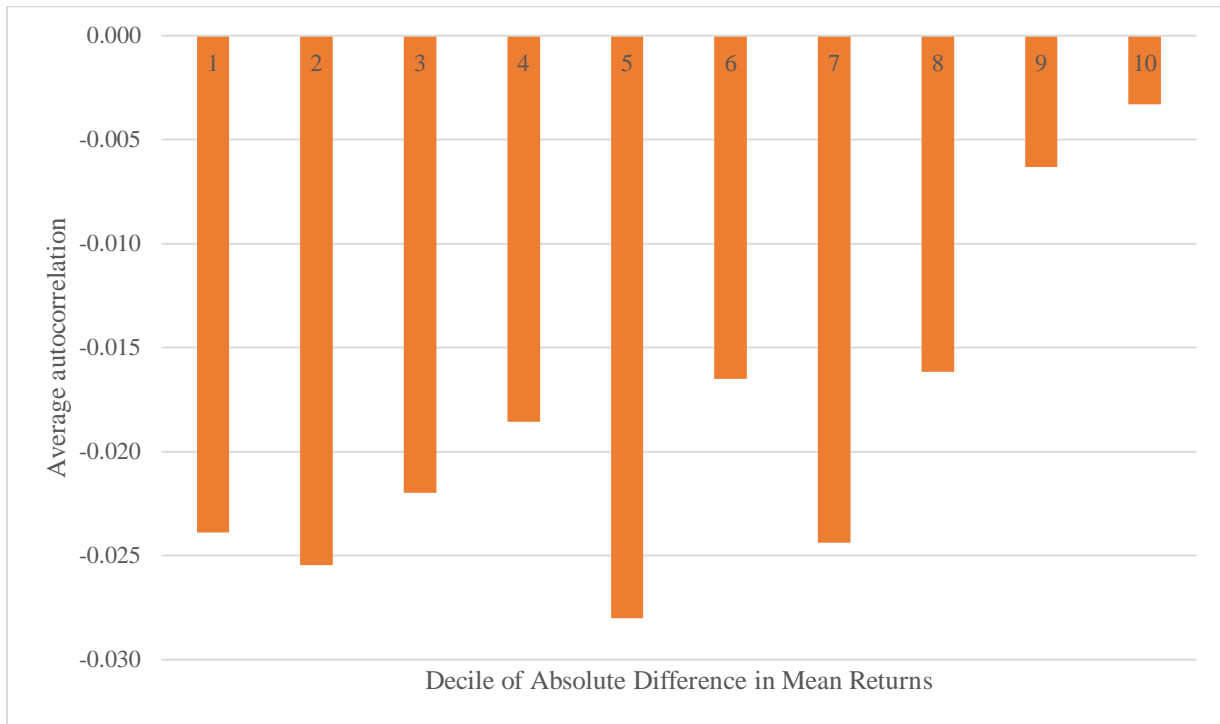


Figure 3. For each non-financial firm common stock listed on the NYSE, AMEX, and NASDAQ, with at least 2 years of presence in Compustat, the unconditional monthly return first-order autocorrelation is computed over the entire available sample while excluding the last trading day of each month. Then the 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 autocorrelation is computed for each decile and is plotted here.

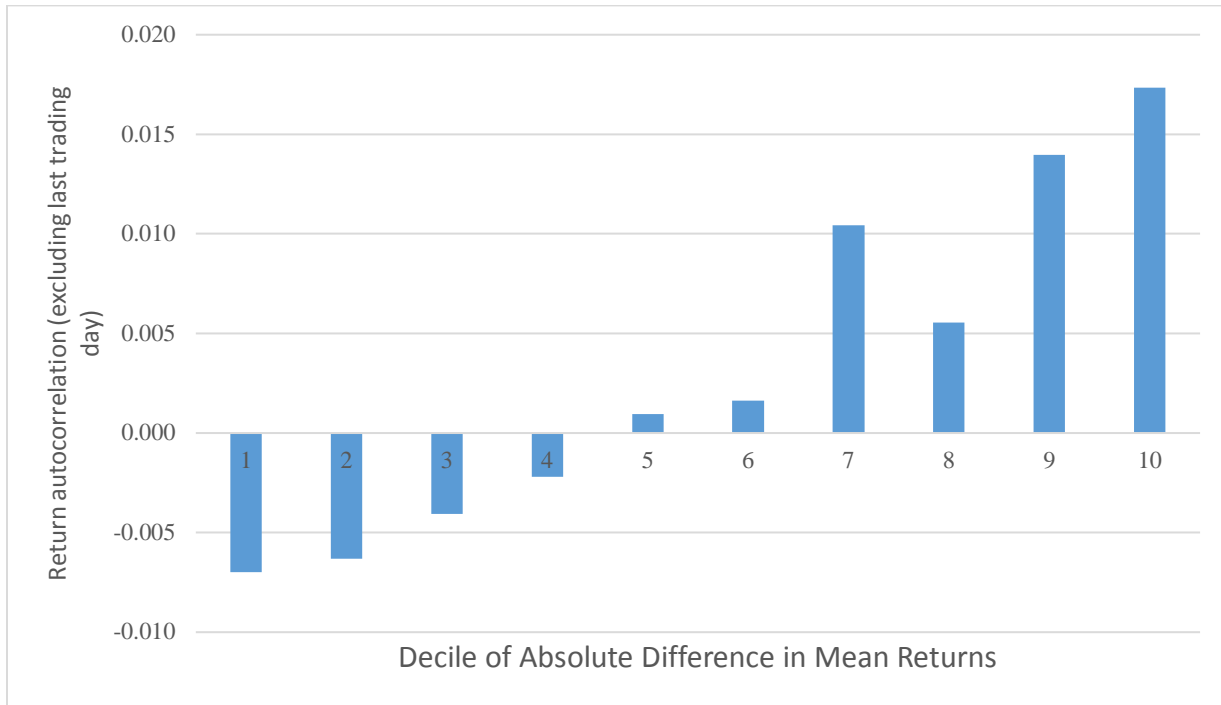


Table 1

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

Sigma ¹⁴	Expected Return Difference, (%/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 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 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 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

¹⁴ Standard Deviation of returns in percent per annum.

Table 3
 Bivariate Regression of Return Autocorrelation on the Absolute Change
 In Mean Return between Two Half-Samples;
 Small and Large Firm Sub-Samples

The sample of available observations for each stock is split into two halves and the explanatory variable is the absolute difference in the mean return in the two halves. The dependent variable is the unconditional autocorrelation with all monthly observations in each stock's entire sample, eliminating the last trading day of each month. The regression t-statistic is reported in the first line of each panel while 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
Panel A. Small Firms				
T-Statistic	5.244***	4.309***	3.826***	2.836***
N	4,805	3,870	3,909	1,906
Panel B. Large Firms				
T-Statistic	8.642***	7.590***	9.029***	8.280***
N	2,874	2,392	2,391	1,944

Table 4

Bivariate Regressions of Return Autocorrelation on the Absolute Change In Attribute between Two Half-Samples; All Firms and Small and Large Firm Sub-Samples

This table presents t-statistics from cross-sectional bivariate regressions of return autocorrelation on the absolute change in various firm attributes. The dependent variable is the monthly return autocorrelation coefficients eliminating the last trading day of each month. The independent variable is the attribute's absolute change between the first and last half of each firm's sample of observations. Panels A, B and C present, respectively, all firms, small cap stocks, and large cap stocks. Small firms are those when first listed are smaller than the 20th percentile of the market equity for NYSE stocks. 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
Panel A: All firms				
Size	-2.314**	-2.522**	-2.713***	-2.862***
B/M	1.284	3.188***	1.051	2.781***
Sales_Growth	1.806	1.109	1.783*	1.492
Profit_Ratio	-1.713	-0.226	-1.745*	-0.089
Operating_Profit_Ratio	-1.793*	-0.361	-1.847*	-0.371
Asset_Growth	2.622***	4.243**	2.485**	4.052***
Age	-2.034**	-2.629***	-2.463**	-3.069***
Beta	3.170***	3.057***	3.964***	3.759***
P/E_Diluted	1.668*	1.673*	1.660*	1.753*
P/E_Not Diluted	1.302	1.256	1.373	1.350
Panel B: Small firms				
Size	-0.107	-0.687	-0.220	-0.585
B/M	1.422	3.376***	1.550	3.123***
Sales_Growth	1.872*	0.841	1.922*	1.500
Profit_Ratio	-1.607	-0.203	-1.651*	-0.066
Operating_Profit_Ratio	-1.666*	-0.249	-1.731*	-0.259
Asset_Growth	2.167**	3.382***	2.027**	3.114***
Age	-1.806*	-1.841*	-1.341	-1.410
Beta	1.125	1.408	1.568	1.817*
P/E_Diluted	0.816	0.919	0.661	0.853
P/E_Not Diluted	0.249	0.529	0.204	0.516
Panel C: Large firms				
Size	-3.914***	-3.486***	-4.312***	-4.066***
B/M	0.540	1.234	0.262	0.881
Sales_Growth	0.290	0.849	0.116	0.910
Profit_Ratio	0.295	0.385	0.210	0.304
Operating_Profit_Ratio	-0.252	-0.193	-0.282	-0.226
Asset_Growth	3.650***	3.347***	3.666***	3.483***
Age	-2.382**	-2.961***	-3.283***	-3.878***
Beta	4.230***	3.538***	5.010***	4.221***
P/E_Diluted	1.447	1.404	1.572	1.561
P/E_Not Diluted	1.519	1.220	1.640	1.342

Table 5
 Bootstrap Bivariate Regression of Return Autocorrelation on the Absolute Change
 In Mean Return between Two Half-Samples;
 Small and Large Firm Sub-Samples

For each of eight stock groups, 200 firms at a time are randomly selected for the bootstrap. These 200 observations are split into two halves and the explanatory variable is the absolute change in mean return between the halves. The dependent variable is the return autocorrelation coefficient computed from all observations in each stock's entire sample while eliminating the last trading day of each month. Re-sampling is repeated (with replacement) 1,000 times. The table reports the mean slope (over 1,000 re-samples) and the t-statistic of the mean using the resulting standard error. N is the number of firms in each stock group. 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
Panel A. Small Firms				
Mean Slope	0.232***	0.222***	0.223***	0.178***
T-Statistic	49.047	43.800	51.112	41.519
N	4,850	3,870	3,909	3,057
Panel B. Large Firms				
Mean Slope	0.264***	0.310***	0.259***	0.290***
T-Statistic	45.416	49.638	48.604	50.867
N	2,874	2,392	2,391	1,944

References

- DeAngelo, H., and Roll, R., 2015, How Stable are Corporate Capital Structures? *Journal of Finance* 70, 373-418.
- Cooper, Michael, Gulen, Huseyin, and Schill, Michael, 2008, Asset Growth and the Cross-Section of Stock Returns, *Journal of Finance* 63(4), 1609-1651.
- Fama, Eugene, and French, Kenneth, 1993, Common risk factors in returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Hou, Kewei, Xue, Chen, Zhang, Lu, 2018, Replicating Anomalies forthcoming *Review of Financial Studies*.
- Savor, Pavel, and Mungo Wilson, 2016, Earnings Announcements and Systematic Risk, *Journal of Finance* 71, (February), 83-138.
- Scholes, M., and J. T. Williams, 1977, Estimating betas from non-synchronous trading, *Journal of Financial Economics* 5, 309-327.

Appendix

Table A-1
Data Definitions and Sources

Acronym	Calculation	Definitions of Components
Size	$\text{Ln}[\text{market_cap} = \text{csho} \times \text{prcc_f}]$	csho: Common Shares Outstanding (\$millions) prcc_f: Closing price per share at fiscal year end
B/M	ceq/market cap	ceq: Common Ordinary equity(\$millions) market cap: as in first row
Sales_Growth	$\text{Ln}(\text{sales}_t / \text{sales}_{t-1})$	sales: Sales/Turnover (Net) (\$millions) in year t
Profit_Ratio	ib/sales _t	ib: Income Before Extraordinary Items (\$millions); sales: in same year
Operating_Profit_Ratio	EBITDA/sales _t	EBITDA: Income Before Interest, Tax, Depreciation and Amortization (\$millions); sales: in same year
Asset_Growth	$\text{Ln}(\text{at}_t / \text{at}_{t-1})$	at: Assets – Total (\$millions) in year t
Age	$\text{max}(\text{fyear}) - \text{min}(\text{fyear})$	fyear: Fiscal Year min (max) first (last) year on Compustat
Beta	Beta computed based on Scholes and William (1977)	Computed from monthly observations over the entire available sample of each firm against the CRSP's value-weighted market index including distributions
P/E	prcc_f/diluted EPS	diluted EPS are earnings divided by shares outstanding plus shares of exercisable stock option grants and convertible liabilities.

Table A-2

Cross-Sectional Distribution Statistics for Firm Attributes by Groups

	N	mean	stdev	min	25%	50%	75%	max
Small Firms								
Ln(Market Cap)	4,805	3.84	1.53	-0.59	2.74	3.78	4.86	11.32
Book/Market	4,797	0.84	0.67	-19.59	0.50	0.77	1.11	6.50
Sales Growth	4,774	0.68	12.51	-2.28	0.07	0.13	0.23	670.55
Profit Ratio	4,787	-4.39	114.91	-7555.97	-0.02	0.02	0.06	39.36
Asset Growth	4,787	0.20	0.55	-0.26	0.07	0.13	0.22	24.18
Firm Age (years)	4,805	20.68	9.80	9.00	13.00	18.00	25.00	58.00
Beta	4,805	0.79	0.42	-1.28	0.47	0.76	1.07	2.33
Small Firms, Positive Book Values								
Ln(Market Cap)	3,870	3.89	1.54	-0.59	2.75	3.81	4.95	9.98
Book/Market	3,862	0.93	0.51	0.05	0.58	0.83	1.16	4.57
Sales Growth	3,843	0.28	1.17	-2.28	0.07	0.12	0.21	38.77
Profit Ratio	3,856	-1.01	12.82	-305.40	0.00	0.03	0.07	24.63
Asset Growth	3,852	0.16	0.19	-0.26	0.07	0.12	0.20	4.47
Firm Age (years)	3,870	20.50	9.84	9.00	13.00	18.00	25.00	58.00
Beta	3,870	0.76	0.42	-1.28	0.44	0.73	1.04	2.33
Small Firms, Non-Financial								
Ln(Market Cap)	3,909	3.76	1.56	-0.59	2.63	3.64	4.81	11.32
Book/Market	3,909	0.83	0.62	-7.65	0.46	0.74	1.11	4.57
Sales Growth	3,903	0.79	13.82	-2.22	0.08	0.14	0.25	670.55
Profit Ratio	3,904	-5.28	127.12	-7555.97	-0.04	0.01	0.04	24.63
Asset Growth	3,909	0.21	0.60	-0.20	0.08	0.14	0.23	24.18
Firm Age (years)	3,909	21.19	10.10	9.00	13.00	18.00	26.00	58.00
Beta	3,909	0.84	0.41	-1.28	0.54	0.83	1.12	2.33
Small Firms, Non-Financial and Positive Book Values								
Ln(Market Cap)	3,057	3.79	1.57	-0.59	2.64	3.68	4.86	9.98
Book/Market	3,057	0.92	0.54	0.07	0.54	0.81	1.18	4.57
Sales Growth	3,053	0.30	1.24	-1.00	0.08	0.13	0.22	38.77
Profit Ratio	3,054	-1.18	13.29	-293.71	-0.01	0.02	0.04	24.63
Asset Growth	3,057	0.17	0.20	-0.16	0.07	0.13	0.21	4.47
Firm Age (years)	3,057	21.10	10.24	9.00	13.00	18.00	26.00	58.00
Beta	3,057	0.82	0.41	-1.28	0.53	0.80	1.10	2.33
Large Firms								
Ln(Market Cap)	2,874	6.18	1.64	1.42	5.05	6.15	7.25	12.30
Book/Market	2,873	0.69	0.70	-22.54	0.41	0.67	0.93	5.44
Sales Growth	2,869	0.37	5.87	-3.80	0.07	0.11	0.17	255.40
Profit Ratio	2,869	-0.54	8.51	-278.77	0.01	0.04	0.08	6.79
Asset Growth	2,873	0.14	0.16	-0.18	0.07	0.11	0.17	3.31
Firm Age (years)	2,874	25.43	13.28	9.00	15.00	21.00	32.75	59.00
Beta	2,874	1.03	0.40	-0.26	0.76	1.01	1.28	2.38
Large Firms, Positive Book Values								
Ln(Market Cap)	2,392	6.22	1.61	1.42	5.08	6.17	7.25	12.30
Book/Market	2,391	0.78	0.44	0.06	0.47	0.70	0.97	5.44
Sales Growth	2,387	0.22	1.84	-3.80	0.07	0.11	0.17	83.43
Profit Ratio	2,387	-0.26	5.39	-205.03	0.02	0.05	0.09	6.79

Asset Growth	2,391	0.13	0.12	-0.17	0.07	0.11	0.16	1.86
Firm Age (years)	2,392	25.19	13.05	9.00	15.00	21.00	32.00	58.00
Beta	2,392	1.01	0.40	-0.19	0.73	0.98	1.25	2.32
Large Firms, Non-Financial								
Ln(Market Cap)	2,391	6.14	1.66	1.42	4.96	6.12	7.23	12.30
Book/Market	2,391	0.65	0.72	-22.54	0.38	0.62	0.88	5.44
Sales Growth	2,390	0.35	6.17	-0.67	0.07	0.11	0.17	255.40
Profit Ratio	2,390	-0.65	9.31	-278.77	0.01	0.04	0.07	6.79
Asset Growth	2,391	0.14	0.16	-0.18	0.07	0.11	0.17	3.31
Firm Age (years)	2,391	26.06	13.68	9.00	15.00	22.00	34.00	59.00
Beta	2,391	1.06	0.39	-0.26	0.80	1.04	1.31	2.38
Large Firms, Non-Financial and Positive Book Values								
Ln(Market Cap)	1,944	6.18	1.63	1.42	5.02	6.11	7.24	12.30
Book/Market	1,944	0.73	0.43	0.06	0.44	0.66	0.91	5.44
Sales Growth	1,943	0.16	0.38	-0.67	0.07	0.11	0.16	10.94
Profit Ratio	1,943	-0.32	5.94	-205.03	0.02	0.04	0.08	6.79
Asset Growth	1,944	0.13	0.12	-0.17	0.07	0.11	0.17	1.86
Firm Age (years)	1,944	25.83	13.48	9.00	15.00	22.00	34.00	58.00
Beta	1,944	1.03	0.39	-0.19	0.77	1.01	1.27	2.32