The Performance of Characteristic-Sorted Portfolios: Evaluating the Past and Predicting the Future

Aydoğan Alti, Travis L. Johnson, and Sheridan Titman

Discussion

Amit Goyal

23 May 2022

ABFER 9th Annual Conference

UNIL | Université de Lausanne

swiss:finance:institute

Autocorrelation in realized returns

Campbell (2018, p. 138):

1. Positive correlation between dividend news and revisions to expected returns.

2. Autocorrelation in expected returns (expected returns follow an AR(1) process).

3. Negative correlation between realized returns and expected returns (realized returns decrease when expected returns increase).

The authors focus on (2).

Data generating process

Model:

$$R_{t+1} = \mu_t + \epsilon_{t+1}$$
$$\mu_{t+1} = \mu + \lambda(\mu_t - \mu + \delta_{t+1})$$
$$\mathrm{Cov}\left(\begin{bmatrix}\epsilon_{t+1} \\ \delta_{t+1}\end{bmatrix}\right) = \begin{bmatrix}\sigma_\epsilon^2 & \rho\sigma_\epsilon\sigma_\delta \\ \rho\sigma_\epsilon\sigma_\delta & \sigma_\delta^2\end{bmatrix}.$$

Model implied moments:

$$\mathrm{Mean}(R) = \mu$$
$$\mathrm{StdDev}(R) = \sigma_r$$
$$\boxed{\begin{gathered}\mathrm{Corr}(R_{t+1}, R_t) = \gamma \\ \mathrm{Corr}\left(R_{t+j}, R_t\right) = \lambda^{j-1}\gamma.\end{gathered}}$$

Note: $\{R_t\}$ can be written as an ARMA(1,1) process.

What does that mean?

Since realized returns are not iid, inference about mean return, μ, using average realized return, $\bar{R}$, requires thinking about the entire $T \times T$ covariance matrix, Σ, of realized returns.

In particular, if realized returns are positively autocorrelated, then the variance of $\bar{R}$ (the sum of all the elements in the matrix Σ) is higher than OLS variance of $\bar{R}$ (essentially the sum of only the diagonal elements in the matrix Σ).

Correct t-statistics of the null $H_0\colon \mu = 0$ might be smaller than the OLS t-statistics.

This obviously has investment and asset pricing implications.

First reactions

This is great!

It is simple, obvious, and very important.

Dr. Watson to Sherlock Holmes “When I hear you give your reasons, the thing always appears to me to be so ridiculously simple that I could easily do it myself” (A Scandal in Bohemia).

Why did no one ever think of it before?

Actually, some people have. Poterba and Summers (1988), Conrad and Kaul (1988, 1989), and Pukthuanthong, Roll, and Subrahmanyam (2021). But they all remain focused on autocorrelation and never go on to analyze its impact on hypothesis testing for mean returns.

My discussion

One suggestion (perhaps for a follow-up paper).

One complaint (perhaps philosophical).

Suggestion (1)

The authors currently focus on implications for one portfolio. Viz., are expected returns different from zero?

However, the central insight of changing expected returns can also be used in a cross-portfolio context.

For example, in a mean-variance framework, one also needs an estimate of the covariance matrix of returns. This will also be impacted by changing expected returns.

Suggestion (2)

One will need to write down a model of changing expected returns in a cross-sectional context.

One possibility is a one-factor model with AR(1) factor and AR(1) correlated betas. Example, Hameed (1997):

$$R_{it+1} = \beta_{it} F_{t+1} + \epsilon_{it+1}$$
$$F_{t+1} = \mu_F + \phi(F_t - \mu_F) + \nu_{t+1}$$
$$\beta_{it} = \rho_i \beta_{it-1} + e_{it}$$
$$\mathrm{Corr}\left(e_{it}, e_{jt}\right) \neq 0.$$

Similar to (but not quite) a model used by Jegadeesh and Titman (1995) although they are focused on delayed price adjustment.

Suggestion (3)

Such a model micro-founds (to some extent) why expected returns are changing.

I have not worked out the math, but it should be possible to think about the 3-D covariance matrix Σ_{ij}^{t} and see the impact of induced autocorrelations and cross-correlations on the estimate of unconditional sample covariance matrix S_{ij}.

Then one can think of mv-efficient portfolio weights

$$w_{mv} \propto S_{ij}^{-1}\bar{R},$$

and do 'hypothesis testing' on w_{mv} (perhaps using Delta method in a frequentist framework or in a more involved Bayesian analysis).

Complaint (1)

Researchers do account for the possibility of autocorrelation in hypothesis testing. Usual procedure is to use Newey-West correction.

Usually, Newey-West does not change the t-statistics much. The authors find similar results. They discuss (page 17) that this is due to long-lasting autocorrelations that Newey-West misses.

How does this magic happen?

Complaint (2)

Some numbers using the Investment portfolio (quarterly over 1963-2019; not market adjusted)

	Lag H (in years)		
t-statistics	2.5	5	10
OLS	2.69	2.69	2.69
Newey-West	2.31	2.40	2.54
Model truncated to lag	2.05	1.73	1.40
Model full	1.74	1.43	1.16

Using $\gamma = 5\%$

- OLS to Newey-West: Not much action.
- Model truncated to lag to Model full: Not much action ($\Rightarrow$ higher lags not that important).
- Newey-West to model with same lags: Dramatic difference.

Complaint (3)

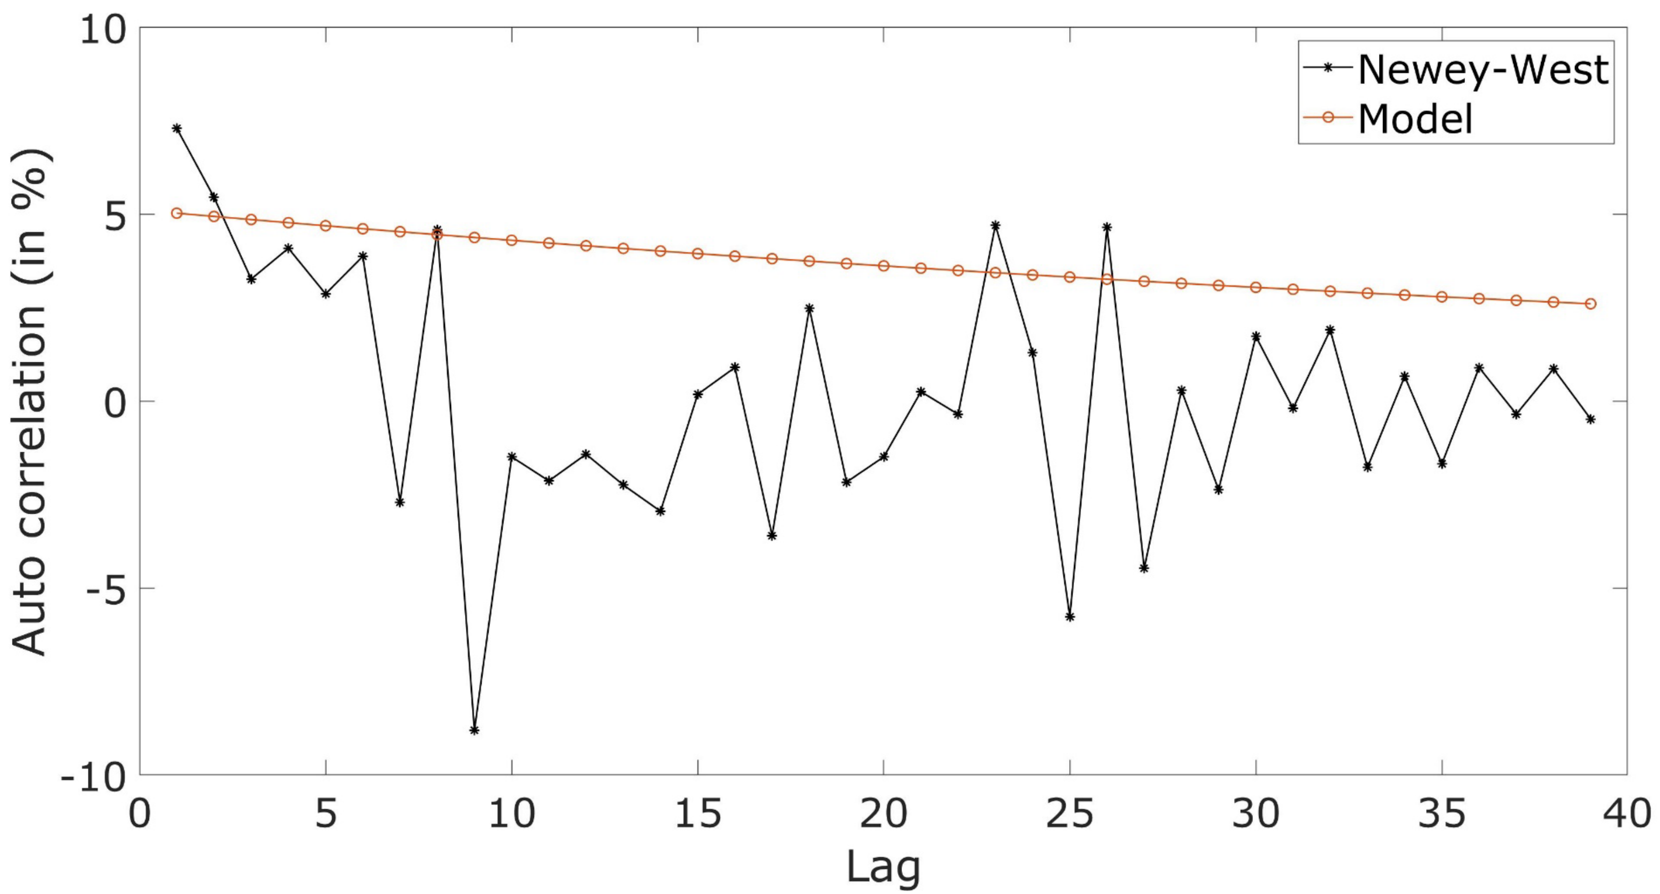

- Model has very gentle decay.
- Model has no negative autocorrelations; but they are present in the data.

Complaint (4)

Yes, each one of the autocorrelations in Newey-West may have large standard errors. But the Newey-West covariance matrix is supposed to be robust.

Which do we trust?

1. Tightly parametrized model, or

2. Non-parametric model mis-specification robust Newey-West?

Minor

- Adjusting returns from CAPM
  - Are the results due to the portfolio return or the market return?
  - What is the impact of estimation error in $\hat{\beta}$ on inferences?

- Is the unconditional Sharpe ratio?
  - $SR = \mathrm{E}(SR_t) = \mathrm{E}(\mu_t/\sigma_\epsilon) = \mu/\sigma_\epsilon$, or
  - $SR = \mathrm{E}(R)/\mathrm{SDev}(R) = \mu/\sigma_r$?

Final thoughts

Encourage everyone to read the paper