Tower Building and Stock Market Returns

Gunter Löffler*
University of Ulm

First version: April, 2010
This version: September 2010

Abstract

This papers shows that construction starts of record-breaking skyscrapers predict subsequent US stock returns. The predictive ability is significant and relatively stable. It exceeds the one of alternatives such as the prevailing historical mean or predictions based on dividend yields. One explanation for these patterns is that tower building is indicative of over-optimism; alternatively, tower building could help to identify periods of low risk aversion. The evidence indicates that the predictive content of tower building is at least partly related to overvaluation.

JEL classification: G12, G14
Key words: Predictability, Stock Market, Dividend yield, Skyscrapers

* Gunter Löffler, Institute of Finance, University of Ulm, Helmholtzstrasse 18, 89069 Ulm, Germany. Phone: ++49-731-5023597. Fax: ++49-731-5023950.
E-mail: gunter.loeffler@uni-ulm.de

1 Introduction

Ever since the story of the tower of Babel was recorded, the construction of tall towers has been associated with human hubris. From a finance theory perspective, towers are large-scale projects with uncertain future cash flows and enormous funding requirements. These observations suggest two ways why tower building might predict low future stock market returns. Either it indicates periods in which over-optimism has led to overvalued stock markets, or it helps to identify times of low risk aversion. (With low risk premia, funding costs for large-scale projects are lower, while future stock market returns are expected to be relatively low, too.) An example that illustrates the two opposite interpretations is the Chicago Spire, which has a planned height of 609 meters, exceeding existing buildings in the US as well as the planned height of the One World Trade Center. Construction of the Chicago Spire began in June 2007, a time in which (i) risk premia – as exemplified by low credit spreads – were low, and (ii) valuation levels appear to have been relatively high.

In this paper, I therefore examine whether tower building is associated with lower subsequent stock market returns. Using US data from 1871 to 2009, I document significant predictability. In the three to five years after the construction of a record-breaking new skyscraper was begun, per annum stock market returns are around 10 percentage points lower than in other years. The analysis shows that the predictive ability of tower building is significant and stable over time. This contrasts with the fragile predictive power of variables such as the dividend yield which have been studied extensively in the literature (e.g. Welch and Goyal, 2008).

The perception that tower building could be linked to economic as well as stock market performance is frequently voiced in the media.[1] Often, such articles cite the research report of Lawrence (1999). The only academic paper I am aware of is Thornton (2005), who discusses the relation between tower building, business cycles and economic crises but does not conduct a statistical analysis.

There is a large literature on predicting stock markets with dividend yields and other variables. Classical references are Campbell and Shiller (1988) and Fama and French (1988). Recent contributions include Goyal and Welch (2003), Boudoukh, Richardson, and Whitelaw (2008),

---

[1] Examples include a 2005 article in Fortune (http://money.cnn.com/magazines/fortune/fortune_archive/2005/09/05/8271392/index.htm) or a 2009 article in the Telegraph http://www.telegraph.co.uk/news/worldnews/middleeast/dubai/6934603/Burj-Dubai-The-new-pinnacle-of-vanity.html)

Campbell and Thompson (2008), Welch and Goyal (2008), and Rapach, Strauss and Zhou (2010).

The remainder of the paper is structured as follows. In Section 2, I describe the data and the methodology. Section 3 analyzes the predictive ability of regression models that include information on tower building. Section 4 concludes.

2 Data and methodology

Data on towers is obtained from the research data base of Emporis, a private information provider focusing on building-related information. The data base also contains information on planned projects and construction status. This allows to identify buildings that were started to be built but were never finished, thereby avoiding possible selection biases that might arise when only finished buildings are studied. The measure of height used is the elevation from the building's base to its highest architectural element.[2] I focus on skyscrapers, i.e. towers with residential or commercial use, and neglect telecommunications towers and other high-rise structures. Skyscrapers are often viewed as a class of their own. For example, the Sears (now Willis) tower is often described to have held the record for the world's tallest building until the building of the Petronas Towers,[3] even though it was exceeded by the CN telecommunications tower in Toronto in 1976. Furthermore, skyscrapers are more expensive to build. The Sears Tower, finished in 1974, had construction costs of USD 160 million; the CN tower, finished in 1976, cost CDN 63 million.[4] In the US, skyscrapers are so tall relative to other structures that including the latter in the analysis would not change the constructed series of record-breaking buildings. Note that in the following, I will use the words *towers* and *skyscrapers* interchangeably.

To measure tower-building activity, I identify the years in which a tower that would break the current US height record was started to be built, or in which construction was finished. (In a

---

[2] "Architectural elements include everything which is integral to the design, including sculptures, spires, screens, parapets, and decorative features." (Source: http://standards.emporis.com).

[3] E.g. http://www.willistower.com/propertyprofile.html, or emporis.com.

[4] Source: emporis.com. The 1976 average CAN/USD exchange rate was 0.9863 as reported in http://research.stlouisfed.org/fred2/data/.

sensitivity analysis, this is also done based on buildings that break the international record.) The information is recorded in the following two dummy variables:

RecordStart_t	One if construction of a tower breaking the current US record begun in year t, zero else.
RecordFinish_t:	One if construction of a tower breaking the current US record was finished in year t, zero else.

As a general measure of US building activity, I examine the number of towers exceeding a height of 100 meters. Since the number of towers above that size is trending upward in long cycles, I examine logarithms relative to their 20-year trailing average. Specifically, I construct the following variables:

100_Start_t:

$$\ln\left(\frac{1+\text{Number of } \mathit{US} \text{ towers taller than } 100 \text{ meters } \mathit{begun} \text{ in year t}}{1+\text{Average } \mathit{n}\text{umber of } \mathit{US} \text{ towers taller than } 100 \text{ meters } \mathit{begun} \text{ in t - 1, t - 20}}\right)$$

100_Finish_t:

$$\ln\left(\frac{1+\text{Number of } \mathit{US} \text{ towers taller than } 100 \text{ meters } \mathit{finished} \text{ in year t}}{1+\text{Average } \mathit{n}\text{umber of } \mathit{US} \text{ towers taller than } 100 \text{ meters } \mathit{finished} \text{ in t - 1, t - 20}}\right)$$

Table 1 lists the record-breaking buildings that enter the construction of the *Record* variables. International towers breaking the world record are also listed in the table. Note that there is no year in which two or more record-breaking towers were begun or finished. There is therefore no need to decide how multiple starts or completions should be treated in the analysis. Figure 1 shows the building activity as measured by the variables 100_Start and 100_Finish.

Stock market data is obtained from different sources. Annual US stock market returns and associated dividend information and risk-free returns from 1871 – 2009 are obtained from Robert Shiller's web site.[5] For the 1926-2009 sample, annual data on the value-weighted market portfolio and dividend information are from CRSP, made available by Michael Roberts.[6] The risk-free rate that is used for the CRSP data is the one-month Treasury bill rate taken from Ken

---

[5] http://www.econ.yale.edu/~shiller/data.htm

[6] http://finance.wharton.upenn.edu/~mrrobert/data_code.htm

French's website.[7] In the literature on stock market predictability, the dividend yield is the most widely studied and successful predictor for long-horizon stock market returns. I include its logarithm in the analysis, denoted by *dy*. Stock market returns enter the analysis as logarithmic excess returns over the risk-free rate, denoted by $r_{t,t+k}$.

In line with extant literature, predictability is analyzed through linear regressions. As the information contained in tower building activity may be reflected with a time-lag, I also include one-year lags of the tower variables. When studying the effects of tower starts, for example, I would run the following regression:

$$r_{t,t+k} = b_0 + b_1 dy_t + b_2 TowerStart_t + b_3 TowerStart_{t-1} + u_{t,t+k} \qquad (1)$$

For return horizons larger than one year (k>1), the return observations are overlapping, which induces correlations in the error terms. While OLS continues to yield consistent estimates of the coefficients *b*, standard errors are no longer reliable. Until recently, the common academic response was to use Newey and West (1987) standard errors. Ang and Bekaert (2007), however, have shown that the Newey and West procedure is sensitive to time persistence in the explanatory variables. Through simulations, Ang and Bekaert (2007) show that Hodrick (1992) standard errors are much more reliable in regressions such as (1). With k>1, I therefore use Hodrick (1992) standard errors. With one-year returns, the standard heteroscedasticity-robust estimator of White (1980) is used.

I will also explore the out-of-sample performance of predictions based on regressions such as (1). One of the metrics examined is the out-of-sample R² suggested by Campbell and Thompson (2008), which compares the mean squared error of a prediction model to the errors one would incur when using the historical mean prevailing at time t as a predictor. Let $\bar{r}_{t,t+k}$ denote the prediction derived from the historical mean prevailing at time t, and let $\hat{r}_{t,t+k}$ denote an out-of-sample regression-based prediction. The out-of-sample R² is then computed as follows:

$$R^2_{OS} = 1 - \frac{\sum_{t=m}^{T} \left(r_{t,t+k} - \hat{r}_{t,t+k}\right)^2}{\sum_{t=m}^{T} \left(r_{t,t+k} - \bar{r}_{t,t+k}\right)^2} \qquad (2)$$

---

[7] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/

where *m* is the starting year of the out-of-sample analysis. If the R^2_{OS} is positive, the prediction model outperforms the prevailing mean, which serves as a natural benchmark for evaluating predictive performance. To assess the statistical significance of the out-of-sample R², I follow Rapach, Strauss and Zhou (2010) and favor the Clark and West (2007) MSPE-adjusted statistic over the Diebold-Mariano (1995) test, which can have low power when applied to nested models.[8] To compute the MSPE-adjusted statistic, define

$$f_{t,t+k} = \left(r_{t,t+k} - \bar{r}_t\right)^2 - \left(\left(r_{t,t+k} - \hat{r}_{t,t+k}\right)^2 - \left(\bar{r}_{t,t+k} - \hat{r}_{t,t+k}\right)^2\right), \qquad (3)$$

and regress $f_{t,t+k}$ on a constant. The p-value for a one-sided test is obtained by applying the standard normal distribution to the t-statistic of the constant. For k>1, I use Newey and West standard errors with lag k.[9]

3 Tower building and stock returns in the US

3.1 In-sample analysis

I start the analysis with a regression in which future stock market returns are predicted with the logarithm of the dividend yield (dy) and the variables capturing tower construction starts. Results for prediction horizons of one, three and five years are reported in Table 2, separately for the 1871-2009 and the 1926-2009 samples. As is familiar from the literature, high dividend yields are associated with high future returns. However, the coefficients are not significant at a level of 10%, which is in line with the findings of Ang and Bekaert (2007).[10] The dummy variables capturing the construction start of record-breaking towers are consistently negative. For horizons of three and five-years, the t-statistics for the lagged construction start $RecordStart_{t-1}$ range between -2.14 and -3.35; the construction start $RecordStart_{t-1}$ is significant on the 5% level in one regression. The scaled number of building starts (with height above 100 meters) does not

---

[8] See the discussion in Rapach, Strauss and Zhou (2010). The prevailing mean is nested in regression models of type (1) because it would result from restricting coefficients other than b_0 to be zero.

[9] Note that the doubts about the reliability of Newey and West (2008) discussed in conjunction with regression (1) arise from the persistence of the predictors and therefore do not carry over to the regression that is run for the MSPE-adjusted statistic.

[10] With the Newey and West estimator, the coefficient on the dividend yield has a t-statistic of 2.46 in the five-year regression.

show a strong association with future returns. In the two cases in which coefficients are significant on the 10% level, though, they are also negative.

Coefficients of RecordStart_{t-1} are not only statistically but also economically significant. The coefficient of -.305 in the three-year regression, for example, implies that the three-year return from year 2 to year 4 after a record-breaking tower was begun is on average 30.5 percentage points lower than in years not preceded by tower-building activity. This translates into a per annum return difference of around 10 percentage points.

The predictability literature (cf. Cochrane, 2005, ch. 20) has shown that one should be careful to interpret coefficients or R^2's that increase with the horizon as evidence that forecasting ability increases with longer horizon. This observation, however, applies to persistent predictors such as the dividend yield. The *RecordStart* variables are not persistent. Their first-order autocorrelation is 0.11, while the first-order autocorrelation of the dividend yield is 0.87. It is therefore appropriate to say that the predictive power of tower starts increases with the horizon. This seems compatible with both explanations of a link between tower building and stock returns. If tower starts are indicative of overvaluation, predictive power should increase with the horizon provided that periods of overvaluation persist for more than one year; in this case, tower starts will not necessarily be observed at the end of the overvaluation period. To see why the risk premium explanation can be compatible, consider a situation where risk premia fall over an extended period. If a tower is started to be built in the middle of such a period, the short-term returns after the building start will be positive as the fall of risk premia continues to push up prices for a while.

The regressions of Table 3 examine whether the completion of towers predicts future returns, too. The answer is in the negative. None of the variables capturing the completion of towers has a significant coefficient. This conforms to expectations. Construction starts should provide a better measure of the current situation – be it overvaluation or low risk premia – than construction completions. Several record-breaking towers like the Chrysler Building or the Sears Tower have been completed despite the fact that economic conditions worsened significantly after their start.

Summarizing the results from Tables 2 and 3, it appears that there is an association of tower-building with mid-term future stock returns, and that these effects can be captured well through information about the start of record-breaking towers in the two years prior to the prediction horizon. This suggests a parsimonious model, in which tower building is represented through the variables $RecordStart_t$ and $RecordStart_{t-1}$. In the following, I will therefore concentrate on three simplified models. For purposes of comparison, I first examine the standard dividend model without tower information:

Dividend model: $$r_{t,t+k} = b_0 + b_1 dy_t + u_{t,t+k}$$

In addition, I either explain future returns with tower information only, or add the tower information to the dividend model:

Tower model: $$r_{t,t+k} = b_0 + b_1 RecordStart_t + b_2 RecordStart_{t-1} + u_{t,t+k}$$

Dividend + tower model: $$r_{t,t+k} = b_0 + b_1 dy_t + b_2 RecordStart_t + b_3 RecordStart_{t-1} + u_{t,t+k}$$

Table 4 compares the in-sample fit of these models. None of the models is able to significantly predict one-year returns. On the three-year horizon, the tower model leads to an R^2 of 22.3% over the 1926-2009 sample; the corresponding R^2 of the dividend model is 9.7%. On the five-year horizon, the two models have R^2's that are very close. Combining dividend and tower information further increases the in-sample fit, suggesting that the two variables capture different information. This is confirmed by the low correlation of the predictive variables used in the regressions. Over the 1871-2009 period, the correlation of the dividend yield variable with $RecordStart_t$ and $RecordStart_{t-1}$ is -0.047 and 0.027, respectively. I further examine whether the results are sensitive to defining *RecordStart* with US or non-US towers breaking the current world record, rather than define it with US towers breaking the US record. Results (not reported) do not change conspicuously. Over the three-year horizon, for example, the lagged tower starts continue to be significant on a 1% level in each regression. The next section will examine whether the in-sample predictability documented here carries over to out-of-sample predictability.

3.2 Out-of-sample analysis

An out-of-sample analysis mimics the situation of market participants who tried to use the information in predictive variables at a certain point τ in the past. To predict *k*-year returns starting from the end of year τ, one would run the regression (when using the dividend+tower model)

$$r_{t,t+k} = b_0 + b_1 dy_t + b_2 RecordStart_t + b_3 RecordStart_{t-1} + u_{t,t+k}, \quad t = 1,...,\tau - k, \tag{4}$$

derive coefficient estimates $\hat{b}$, and compute the prediction

$$\hat{r}_{\tau,\tau+k} = \hat{b}_0 + \hat{b}_1 dy_\tau + b_2 RecordStart_\tau + \hat{b}_3 RecordStart_{\tau-1}. \tag{5}$$

By running regressions of type (4) for each τ considered in the analysis, one obtains a series of out-of-sample predictions. Following the suggestion of Campbell and Thompson (2008), I also examine the effects of imposing a non-negativity constraint on the prediction. The motivation is that expected excess returns on an asset exposed to systematic risk should be non-negative if the average investor is risk-averse. This leads to the constrained predictor

$$\hat{r}_{\tau,\tau+k}(constrained) = \max(0, \hat{r}_{\tau,\tau+k}). \tag{6}$$

For the sake of brevity, I limit the presentation to the three-year horizon because the analysis of the previous section has shown that the in-sample predictability is strongest over the three-year horizon. The first date τ on which a prediction is made is taken to be 1915. By then, the three construction starts of 1906-1910 have fully entered the three-year return regressions.

Out-of-sample performance will suffer if coefficient estimates are unstable. It is therefore illustrative to examine how coefficients change over time. Figure 2 shows the time series of coefficients for the dividend plus tower model. The coefficient on lagged construction starts is consistently negative. It decreases following the building activity from 1926-1930, but it is largely unaffected by the subsequent construction starts.[11] The coefficient on the non-lagged $RecordStart_t$ is more volatile at the start but also quite stable afterwards. The coefficient on the dividend variable, by contrast, shows a relatively large volatility.

In accordance with prior literature, the benchmark for assessing the predictive performance of a regression model is the prevailing historical mean, computed as

$$\bar{r}_{\tau,\tau+k} = \frac{1}{\tau - k} \sum_{t=1}^{\tau-k} r_{\tau,\tau+k}. \tag{7}$$

---

[11] It should be noted that the 2006 and 2007 construction starts do only partially enter the regressions as the last three-year horizon considered is from 2007-2009.

Separately for each of the three regression models, Figure 3 shows how the sum of squared prediction errors of a given regression model compares to the sum of squared prediction errors of the historical mean. Specifically, for a given year T, the cumulative relative squared prediction errors are determined through

$$\text{Cumulative relative SSE} \;=\; \sum_{\tau=t(1915)}^{T}\left(r_{t,t+k}-\bar{r}_{t,t+k}\right)^2 - \sum_{\tau=t(1915)}^{T}\left(r_{t,t+k}-\hat{r}_{t,t+k_t}\right)^2$$

If the cumulative relative SSE is positive, the regression predictions perform better than the simple prediction based on the historical mean. As is evident from Figure 3, the in-sample performance of the tower variables carries over to the out-of-sample analysis. Including the tower variables in the predictive regressions leads to squared errors that are lower than the ones of the historical mean; squared errors are also lower than the ones arising from the dividend model, in which the dividend yield is the sole predictor. Furthermore, results do not critically depend on whether the non-negativity constraint is imposed or not. Regarding stability, it is noteworthy that the performance gains of the tower model are steadier than the ones of models that include the dividend yield.

Table 5 provides a battery of statistics on the prediction errors. Results are presented both for the entire 1915-2009 period and for two sub-periods, 1915-1945 and 1945-2009. In each case, the estimation sample starts in 1871, and the Shiller data are used throughout.

All predictors tend to underestimate returns on average, leading to a positive mean prediction error. Compared to the dividend model, however, including information on tower building reduces the bias; the bias of the tower model is not significantly different from zero. (To assess significance, the prediction errors are regressed on a constant; standard errors are based on Newey and West (1987) with a lag length of k=3.)

As is evident from the out-of-sample R^2, which is based on the root mean squared errors, models with tower information consistently outperform the historical mean. With the exception of the constrained dividend+tower model in the 1945-2009 sub-period (significance of 10%), the performance difference is significant on a level of 5% or better. Note that the out-of-sample R^2's of the tower-model are also mostly larger than the ones of the dividend model. An examination of the mean absolute error MAE shows that the superiority of the models with tower information continues to hold when moving from squared prediction errors to mean absolute errors.

The differences between the errors of unconstrained and constrained predictors are interesting because they could help to differentiate between the two possible explanations why tower building predicts returns. The constraint is motivated by an equilibrium approach. Even if risk-aversion is very low, one would not expect risk premia to be negative as the stock market is not only risky but also positively correlated with consumption risk. If tower building is indicative of overvaluation, by contrast, there is no reason for ruling out non-negative expected returns. In the tower and dividend plus tower models, errors increase if the constraint is imposed. As is evident from Figure 3, the difference is largely due to the tower building activity in the 1960s and 1970s. Using the MSPE-adjusted test to examine whether there is a difference between the unconstrained model and the corresponding constrained one yields t-statistics of 1.76 (tower model) and 1.22 (dividend plus tower model). Therefore, the unconstrained tower model outperforms the constrained one with a upper-tail significance better than 5%. The findings thus provide an indication that the predictive content of tower building is at least partly related to overvaluation. Note that in the dividend model, it is the other way round. There, imposing the constraint increase the predictive performance. Of course, one should bear in mind that it is generally difficult to distinguish between overvaluation and risk premia explanations in long-horizon returns (see Cochrane (2005) for a general discussion). Together with the low significance levels and the dependence on few events, this justifies caution when giving an interpretation in favor of overvaluation.

To check robustness of the results, I conduct several sensitivity analyses. Important results are summarized below:

1) Instead of using the Shiller data from 1871-2009, I link the Shiller data from 1871 to the end of 1926 with the CRSP data from the end of 1926 to 2009. Conclusions are not affected. For 1915-2009, the out-of-sample R^2 of the unconstrained tower model changes from 0.106 (Table 5) to 0.113.

2) Instead of using the Shiller data with an estimation sample starting in 1871, I conduct the 1945-2009 analysis with the CRSP data starting in 1926. Conclusions are not affected. The out-of-sample R^2 of the unconstrained tower model changes from 0.152 (Table 5) to 0.154.

3) Instead of using the log dividend yield as an additional predictor, I use the log 10-year price-earnings ratio computed by Robert Shiller. Over the entire sample, the price-earnings ratio leads

to a better performance than the dividend yield but conclusions are not affected. For example, the out-of-sample R^2 of the unconstrained tower model (0.106) still exceeds the one of a model including only the price-earnings ratio (0.086). Over 1945-2009, the price-earnings ratio leads to somewhat larger errors than the dividend yield.

5 Conclusion

In this paper, I have shown that construction starts of record-breaking skyscrapers in the United States predict subsequent US stock returns. The predictive ability appears to be stable. It exceeds the one of alternatives such as the prevailing historical mean or predictions based on dividend yields. Of course, one could be concerned about the fact that predictability is tied to a few tower-building episodes. Here, it is important to note that other, frequently used predictive variables such as the dividend yield are very slowly moving. Their high persistence leads to situations in which a regression is heavily affected by a few regime shifts even though it uses decades of continuous data (cf. the discussion in Cochrane, 2005, Ch. 20). Persistence also leads to problems in assessing statistical significance (Ang and Bekaert, 2007). Viewed in this light, it is not obvious that statistical tests (which by construction take the number of observations into account) should be less reliable when applied to tower-building, as compared to tests applied to dividend yields or other variables.

One explanation for the documented patterns is that tower building is indicative of over-optimism. Widespread over-optimism could lead not only to tower-building but also to overvalued stock markets. The rational asset pricing explanation is that in periods of low risk aversion, financing of large-scale projects such as record-breaking towers is easier, and expected returns are lower. It is generally difficult to disentangle rational and irrational explanations for patterns in long-run returns. In the paper, I provide indirect evidence which favors the overvaluation view. Given its weakness, it should not be over-interpreted. In any case, the results suggest that a few episodes can have a significant impact on the long-run returns earned by investors. They also show that non-financial information like tower-building may help investors to identify periods of low future expected returns.

References

Ang, A., and G. Bekaert. 2007. Return Predictability: Is It There? Review of Financial Studies 20:651–707.

Boudoukh, J., M. P. Richardson, and R. F. Whitelaw. 2008. The Myth of Long-Horizon Predictability. Review of Financial Studies 21:1577–605.

Campbell, J. Y., R. J. Shiller. 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. Review of Financial Studies 1, 195–228.

Campbell, J. Y., and S. B. Thompson. 2008. Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average? Review of Financial Studies 21:1509–31.

Clark, T. E., and K. D.West. 2007. Approximately Normal Tests for Equal Predictive Accuracy in Nested Models. Journal of Econometrics 138:291–311.

Cochrane, J. 2005. Asset Pricing. Rev. Edition. Princeton.

Diebold, F. X., and R. S. Mariano. 1995. Comparing Predictive Accuracy. Journal of Business and Economic Statistics 13:253–63.

Fama, E. F., and K. R. French. 1988. Dividend Yields and Expected Stock Returns. Journal of Financial Economics 22:3–25.

Goyal, A., and I. Welch. 2003. Predicting the Equity Premium with Dividend Ratios. Management Science 49:639–54.

Hodrick, R. J. 1992. Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement. Review of Financial Studies 5:357-386.

Lawrence, A. 1999. The Skyscraper Index: Faulty Towers! Property Report, Dresdner Kleinwort Benson Research (Cited in Thornton, 2005).

Newey, W. K., and K. D. West. 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. Econometrica 55:703–8.

Rapach, D., J. Strauss, and G. Zhou. 2010. Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy. Review of Financial Studies 3:821-62.

Thornton, M. 2005. Skyscrapers and Business Cycles. The Quarterly Journal of Austrian Economics 8, 51–74.

Welch, I., and A. Goyal. 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. Review of Financial Studies 21:1455–508.

White, H. 1980. A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, Econometrica 48, 817–38..

Table 1: List of towers that were expected to break the US or world record at the time when construction began.

Data from Emporis. The measure of height used is the elevation from the building's base to its highest architectural element.

Name	Height	Start	Finished
US buildings			
Philadelphia City Hall	167.03	1871	1901
Singer Building	186.57	1906	1908
Metropolitan Life Tower	213.36	1907	1909
Woolworth Building	241.4	1910	1913
Church Missionary Building	243.84	1926	
Chrysler Building	318.92	1928	1930
Empire State Building	381	1930	1931
One World Trade Center	417	1966	1972
Willis Tower (Sears Tower)	442.14	1970	1974
One World Trade Center	541.33	2006	
Chicago Spire	609.61	2007	
Non-US buildings			
Petronas Tower	451.9	1992	1998
Shanghai World Financial Center	492 (460 planned at start)	1997	2008
Taipei 101	509.2	1999	2004
Burj Khalifa	827.99	2004	2010

Table 2: Explaining returns on the US stock market with information related to the start of large towers

The log excess return on the US stock market from t to t+k is regressed on the log dividend yield (dy) and variables containing information about the building starts of large towers in the US. RecordStart equals 1 when a record-breaking tower was started to be built in year t. 100_Start relates the number of towers larger than 100 meters started in year t to the number of such starts in the 20 years before t. Data from 1871 – 2009 is constructed by Robert Shiller based on the S&P 500 and other series; data from 1926-2009 is for the CRSP value-weighted market portfolio. T-statistics are based on White (1980) for the one-year horizon, and on Hodrick (1992) for horizons larger than one. Stars flag significance at the 1% (***), 5% (**) or 10% (*) levels.

	1-year horizon (k=1)		3-year horizon (k=3)		5-year horizon (k=5)	
	1871-2009	1926-2009	1871-2009	1926-2009	1871-2009	1926-2009
dy	0.030	0.045	0.149	0.253	0.273	0.388
	(0.76)	(0.81)	(1.23)	(1.46)	(1.40)	(1.55)
RecordStart_t	-0.095	-0.192	-0.178	-0.514*	-0.282	-0.502**
	(-1.06)	(-1.40)	(-1.09)	(-1.78)	(-1.47)	(-2.23)
RecordStart_{t-1}	-0.006	-0.077	-0.305***	-0.658***	-0.294**	-0.526***
	(-0.08)	(-0.71)	(-2.73)	(-3.57)	(-2.14)	(-3.35)
100_Start_t	-0.025	-0.067*	-0.052	-0.108	-0.093	-0.133*
	(-0.80)	(-1.77)	(-1.01)	(-1.49)	(-1.22)	(-1.64)
100_Start_{t-1}	-0.001	0.044	0.001	0.085	0.005	0.069
	(-0.04)	(1.02)	(0.02)	(1.33)	(0.08)	(0.91)
constant	0.142	0.224	0.620	1.085*	1.098*	1.651*
	(1.09)	(1.18)	(1.59)	(1.83)	(1.75)	(1.95)
Adj. R^2	0.025	0.094	0.146	0.400	0.212	0.398
N	138	82	136	80	134	78

Table 3: Explaining returns on the US stock market with information related to the completion of large towers

The log excess return on the US stock market from t to t+k is regressed on the log dividend yield (dy) and variables containing information about the building starts of large towers in the US. RecordFinish equals 1 when a record-breaking tower was completed in year t. 100_Finish relates the number of towers larger than 100 meters completed in year t to the number of such completions in the 20 years before t. Data from 1871 – 2009 is constructed by Robert Shiller based on the S&P 500 and other series; Data from 1926-2009 is for the CRSP value-weighted market portfolio. T-statistics are based on White (1980) for the one-year horizon, and on Hodrick (1992) for horizons larger than one. Stars flag significance at the 1% (***), 5% (**) or 10% (*) levels.

	1-year horizon (k=1)		3-year horizon (k=3)		5-year horizon (k=5)	
	1871-2009	1926-2009	1871-2009	1926-2009	1871-2009	1926-2009
dy	0.054	0.109**	0.150	0.229	0.270	0.349
	(1.45)	(2.13)	(1.27)	(1.37)	(1.41)	(1.37)
RecordFinish$_t$	-0.122	-0.268	-0.095	-0.161	-0.037	0.054
	(-1.29)	(-1.52)	(-0.65)	(-0.58)	(-0.22)	(0.20)
RecordFinish$_{t-1}$	-0.030	-0.039	0.059	0.200	0.043	0.193
	(-0.33)	(-0.23)	(0.46)	(0.80)	(0.25)	(0.73)
100_Finish$_t$	-0.026	-0.068	-0.046	-0.107	-0.096	-0.147
	(-0.81)	(-1.59)	(-0.75)	(-1.38)	(-1.02)	(-1.29)
100_Finish$_{t-1}$	0.003	0.050	-0.036	0.013	-0.045	0.016
	(0.09)	(1.15)	(-0.61)	(0.18)	(-0.53)	(0.14)
constant	0.218*	0.436***	0.594	0.922	1.052*	1.442*
	(1.81)	(2.47)	(1.57)	(1.61)	(1.71)	(1.68)
Adj. R²	(1.45)	(2.13)	(1.27)	(1.37)	(1.41)	(1.37)
N	-0.122	-0.268	-0.095	-0.161	-0.037	0.054

Table 4: Parsimonious models for explaining returns on the US stock market

The log excess return on the US stock market from t to t+k is regressed on the log dividend yield (dy) and variables containing information about the building starts of large towers in the US. RecordStart equals 1 when a record-breaking tower was started to be built in year t. Data from 1871 – 2009 is constructed by Robert Shiller based on the S&P 500 and other series; Data from 1926-2009 is for the CRSP value-weighted market portfolio. T-statistics are based on White (1980) for the one-year horizon, and on Hodrick (1992) for horizons larger than one. Stars flag significance at the 1% (***), 5% (**) or 10% (*) levels. Coefficients of regression constants are not reported.

	1-year horizon (k=1)		3-year horizon (k=3)		5-year horizon (k=5)	
	1871-2009	1926-2009	1871-2009	1926-2009	1871-2009	1926-2009
Panel A: Dividend yield as predictor – Dividend Model						
Dy	0.050	0.087	0.163	0.265*	0.296	0.418*
	(1.30)	(1.59)	(1.39)	(1.69)	(1.57)	(1.73)
Adj. R²	0.006	0.022	0.039	0.097	0.086	0.186
N	138	82	136	80	134	78
Panel B: Tower building as predictor – Tower Model						
RecordStart_t	-0.125	-0.167	-0.234	-0.438	-0.351	-0.629*
	(-1.43)	(-1.38)	(-1.35)	(-1.63)	(-1.60)	(-1.92)
RecordStart_{t-1}	-0.018	-0.080	-0.303***	-0.581***	-0.292*	-0.475**
	(-0.24)	(-0.79)	(-2.66)	(-3.03)	(-1.93)	(-2.14)
Adj. R²	0.021	0.040	0.087	0.223	0.077	0.174
N	138	83	136	81	134	79
Panel C: Dividend yield and tower building as predictors – Dividend+Tower Model						
dy	0.047	0.078	0.183	0.298*	0.329*	0.469*
	(1.24)	(1.56)	(1.57)	(1.89)	(1.75)	(1.94)
RecordStart_t	-0.122	-0.215	-0.233	-0.538*	-0.378*	-0.591***
	(-1.39)	(-1.60)	(-1.35)	(-1.78)	(-1.74)	(-2.54)
RecordStart_{t-1}	-0.020	-0.076	-0.333***	-0.645***	-0.341**	-0.571***
	(-0.27)	(-0.75)	(-2.91)	(-3.50)	(-2.25)	(-3.49)
Adj. R²	0.026	0.086	0.140	0.390	0.186	0.374
N	138	82	136	80	134	78

Table 5: Analysis of prediction errors for the three-year horizon

Out-of-sample forecasts of three-year stock returns are generated with (i) a model including only the dividend yield (dividend model); (ii) a model including record tower building and lagged record tower building (tower model) and (iii) a model combining the variables from (i) and (ii), denoted as dividend+tower model Stock market data is constructed by Robert Shiller based on the S&P 500 and other series. Statistical significance of mean errors is assessed through a regression of forecast errors on a constant, using Newey and West (1987) standard errors. Statistical significance of the R^2_{os} statistic (out-of-sample R^2 relative to the historical mean) is based on the Clark and West (2007) MSPE-adjusted statistic, again computed with Newey and West (1987) standard errors. RMSE is root mean squared error, MAE mean absolute error. Stars flag significance at the 1% (***), 5% (**) or 10% (*) levels.

	Hist. mean	Dividend model		Tower model		Dividend+Tower model	
		unconstr.	constrained	unconstr.	constrained	unconstr.	constrained
Panel A: 1915-2009							
Mean	0.059	0.110**	0.098**	0.036	0.032	0.097**	0.080*
st.dev	0.344	0.329	0.327	0.328	0.331	0.311	0.319
R^2_{os}	-	0.008	0.042*	0.106***	0.089***	0.127***	0.111***
RMSE	0.347	0.345	0.339	0.328	0.331	0.324	0.327
MAE	0.263	0.271	0.264	0.243	0.247	0.252	0.253
Panel B: 1915-1945							
Mean	0.086	0.076	0.073	0.066	0.065	0.062	0.060
st.dev	0.433	0.421	0.421	0.422	0.423	0.407	0.409
R^2_{os}	-	0.062***	0.064***	0.066**	0.061**	0.132***	0.123***
RMSE	0.435	0.421	0.421	0.420	0.421	0.405	0.407
MAE	0.325	0.312	0.311	0.321	0.322	0.305	0.305
Panel C: 1945-2009							
mean	0.045	0.126***	0.110**	0.020	0.013	0.114***	0.089**
st.dev	0.285	0.268	0.264	0.265	0.271	0.245	0.259
R^2_{os}	-	-0.054	0.018	0.152***	0.120***	0.123**	0.100*
RMSE	0.286	0.294	0.284	0.264	0.269	0.268	0.272
MAE	0.227	0.245	0.236	0.200	0.205	0.219	0.221

Figure 1: US building activity in towers above a height of 100 meters.

The start of tower above 100 meters is measured by the variable 100Start$_{t,}$ defined through.

$$\ln\left(\frac{1+\mathit{Number\ of\ US\ towers\ taller\ than\ 100\ meters\ began\ in\ year\ t}}{1+\mathit{Average\ number\ of\ US\ towers\ taller\ than\ 100\ meters\ began\ in\ t-1,t-20}}\right)$$

The variable 100Finish$_{t,}$ is defined analogously for buildings finished in *t*.

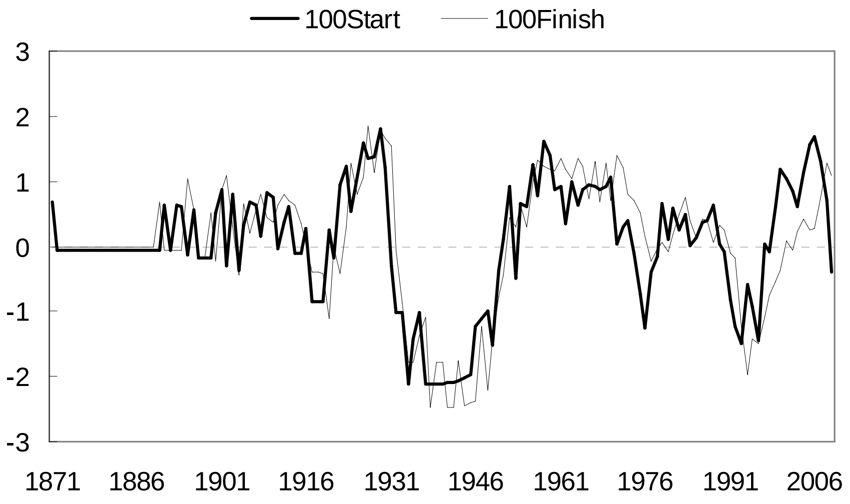

Figure 2: Recursive coefficient updates

The figure shows the coefficients estimates from the recursive regression

$$r_{t,t+k} = b_0 + b_1 dy_t + b_2 RecordStart_t + b_3 RecordStart_{t-1} + u_{t,t+k}, \qquad t = 1,...,\tau - k$$

where $r_{t,t+k}$ is the log excess return on the US stock market from year t to t+k.; dy is the log dividend yield;. RecordStart equals 1 when a record-breaking tower was started in year t. Data from 1871 – 2009 is constructed by Robert Shiller based on the S&P 500 and other series

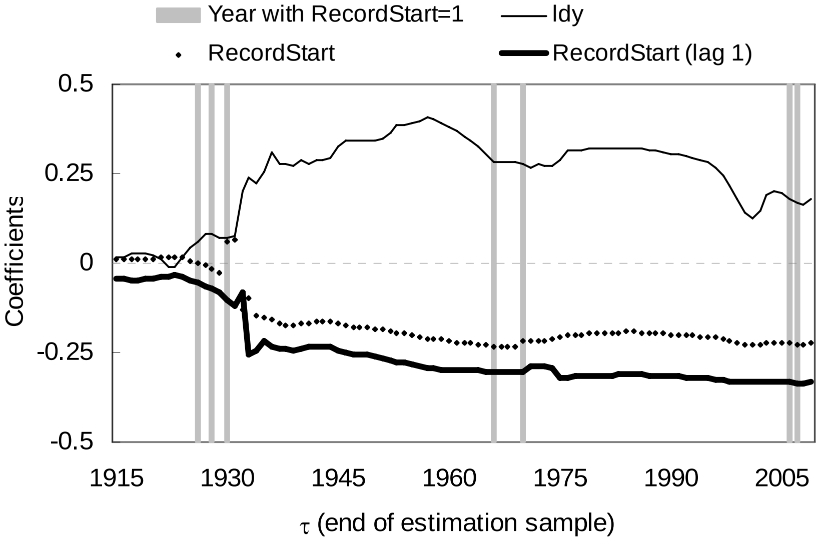

Figure 3: Cumulative out-of-sample performance relative to the prevailing mean

Based on one of the following regressions:

Dividend model: $r_{t,t+k} = b_0 + b_1 dy_t + u_{t,t+k}$

Tower model: $r_{t,t+k} = b_0 + b_1 RecordStart_t + b_2 RecordStart_{t-1} + u_{t,t+k}$

Dividend + tower model: $r_{t,t+k} = b_0 + b_1 dy_t + b_2 RecordStart_t + b_3 RecordStart_{t-1} + u_{t,t+k}$

Recursive out-of-sample predictions for the stock market return over the next three years are made with information available at time t. Variables are defined as in Figure 2. The predictions are denoted by $\hat{r}_{t,t+k}$. The figure plots the relative sum of squared predictions errors

$$\text{Cumulative relative SSE} = \sum_{\tau=t(1915)}^{T} \left(r_{t,t+k} - \bar{r}_{t,t+k} \right)^2 - \sum_{\tau=t(1915)}^{T} \left(r_{t,t+k} - \hat{r}_{t,t+k_t} \right)^2$$

where the test starts in t= 1925 or 1945, and the prevailing historical three-year mean return over years 1871 to t is denoted by $\bar{r}_{t,t+k}$. A positive cumulative relative SSE shows that the regression model performed better than the historical mean.