

SCOTTISH INSTITUTE FOR RESEARCH IN ECONOMICS



SIRE DISCUSSION PAPER

SIRE-DP-2015-82

Revisiting Shiller's excess volatility hypothesis

Dooruj Rambaccussing

UNIVERSITY OF DUNDEE

www.sire.ac.uk

Revisiting Shiller's excess volatility hypothesis¹

Abstract

One of the cornerstone of financial anomalies is that there exists money making opportunities. Shiller's excess volatility theory is re-investigated from the perspective of a trading strategy where the present value is computed using a series of simple econometric models to forecast the present value. The results show that the excess volatility may not be exploited given the data available until time t . However, when learning is introduced empirically, the simple trading strategy may offer profits, but which are likely to disappear once transaction costs are considered.

Key Words: Present Value, Excess Volatility

JEL References: G12, G14, G17

1 Introduction

The excess volatility anomaly, discovered independently by Shiller (1981) implies that stock market prices fluctuate more than present value of expected future dividends under rational expectations. Numerous reasons have been accounted for this including, nonstationary dividends (Mankiw et. al., 1985 and Campbell and Shiller, 1987) and rational learning behaviour (Lof (2014), Hommes and Zhu (2014), Timmermann, 1993)). Presence of such an anomaly would mean that profit making opportunities should be available. An example of such an application is in the work of Bulkley and Tonks (1989,1992) where the excess volatility was proven to yield above market returns in both US and UK markets.

This paper extends their strategy by considering time variation in discount rates, and dividend growth by assuming agents are econometricians, and they forecast the present value. As such, the model considers a series of econometric models ranging from the simple model to the more complicated Vector Autoregression in the Campbell-Shiller (1988) framework.

The main premise of the strategy is as follows. If index prices are higher than the present value, arbitrageurs sell the index in order to avoid losses. On the other hand, if price is lower than the present value, there will be an upward adjustment in future periods. A simple trading strategy is therefore to hold the stock index when it is underpriced and hold bonds when it is overpriced.

The main innovation of this paper builds from the fact that discount rates and dividends in Tonks and Bulkley (1989,1991) were treated as known and observed at any specific point in time. This paper seeks to test whether agents may make the most out of stock market

misalignments with present value in real-time. Agents can only forecast discount rates and dividends using some econometric models. In practice, agents will exploit this opportunity if the cost of arbitraging is less than expected returns and the price lies further away from the present value. As a result of arbitrage, price will converge to fundamentals over time as agents continue exploiting this opportunity. Economic profits based on the identification of this opportunity will only exist until price has adjusted to the fundamental value. Trading strategies should therefore have limited success over time. They will work until the market absorbs this information such that prices reflect the present value (Granger and Timmermann, 2004). This study stems from the excess volatility literature where Shiller (1981) and LeRoy and Porter (1981) document that prices tend to fluctuate much more than variance bounds of the present value. The objective of this study is to test whether a trading may be implemented in the event of this anomaly.

There has been compelling evidence that anomalies tend to disappear over time when they are identified. For instance, Horowitz (2008) finds that the size premium was considerably lower for the sample 1982–1997 as compared to the 1963–1981 sample, which coincides with the publication of the findings of Reinganum (1981) and Banz (1981). Marquering et. al. (2008) find that most of calendar anomalies have significantly weakened. Both findings are consistent with the rational learning framework in behavioural finance where agents make optimal use of the information available and adapt to new information. For a thorough understanding of the rational learning framework, see Brav (2002). The main conclusion of the rational learning framework in the presence of efficient markets is that economic profits can only be temporary.

As agents are rational learners, they will exploit any information which becomes public, which will lead to the convergence of price towards the present value. Hence a simple strategy would be to compare prices and present value and predict the direction of the movement of price in subsequent periods based on the wedge between price and predicted present value. The greater the wedge between the price and present value, the higher the probability of price adjusting towards the present value over time. The treatment of stock market bubbles is confined to the present value. It is important to note that the computation of the present value is subjective. Methodologically, five different econometric models with recursive windows are used for forecasting dividends. Recursive windows ensure that at time t , agents can use only information available until that time in order to predict future dividends and the discount rate. Once the present value is computed, the strategy involves holding the equity index when it is underpriced. In this case, a positive capital gain is expected as prices adjust towards present value. On the other hand, if prices are higher than the present value, the strategy would be to short the market index and go long on another preferably less risky asset.

2 Methodology

This paper considers two simple variants of the present value. The first case assumes that dividends do not grow over time (1). Despite the formulation being restrictive, it provides a good benchmark for alternative representations of the present value. The second model computes the present value as the discount of next period price and dividend (2).

$$PV_{1,t} = \frac{E_t(D_{t+1})}{k_{t+1}}, \quad (1)$$

$$PV_{1,t} = \frac{E_t(D_{t+1})}{k_{t+1} - g_{t+1}}. \quad (2)$$

E_t is the conditional expectations at time t , D_{t+1} is next period dividend, k_{t+1} is the discount rate and g_{t+1} is the dividend growth. It is worth noting that the problem with computing the present value is that the conditional expectations of D_{t+1} , g_{t+1} and k_{t+1} are still unknown at time t . Dividends are forecast using the following five models.

Model 1

$$D_{t+1} = \beta_0 + \beta_1 T + \varepsilon_{t+1}, \quad (3)$$

where T is the trend term. This model forecasts dividends based only on the trend term.

Model 2

$$D_{t+1} = \beta_0 + \beta_1 D_t + \varepsilon_{t+1}. \quad (4)$$

This is an Autoregressive process of order 1, implying that the best forecast of dividends is based on last year's dividend scaled by β_1 and the constant mean.

Model 3

$$Z_{t+1} = A + BZ_t + \varepsilon_{t+1}, \quad (5)$$

where $Z_{t+1} = \begin{bmatrix} P_{t+1} \\ D_{t+1} \end{bmatrix}$.

Model 3 is a Vector Autoregression of order one. It simply implies that there are feedback effects from lagged price and dividends in the dividend equation.

Model 4

$$\Delta d_{t+1} = \beta_0 + \beta_1 pd_t + \beta_2 r_t + \varepsilon_{t+1}, \quad (6)$$

Δd_{t+1} is dividend growth, pd_t the logarithm of the price dividend ratio and r_t is the realized return. It should be noted that this equation is a variant of the Campbell-Shiller log return linear form of the present value. However instead of using future return at time $t+1$ which is unknown at time t , the lagged return at t has been included. The important variable in this model is the logarithm of the price dividend ratio, which according to some empirical studies has predictive ability for dividend growth. Using forecast dividend growth, one step ahead dividend may be computed using the following: $D_{t+1} = (1 + \Delta d_{t+1}) D_t$.

Model 5:

$$Z_{t+1} = A + BZ_t + \varepsilon_{t+1}, \quad (7)$$

Where $Z_{t+1} = \begin{bmatrix} r_{t+1} \\ \Delta d_{t+1} \\ pd_t \end{bmatrix}$.

Model 5 illustrates the Campbell-Shiller identity in the context of a VAR model. In this case future returns, dividend growth and price dividend ratio have joint dynamics and help to forecast each other.

The discount rate k_{t+1} is one of the hardest variables to compute within an economic framework. Various theories and models govern the behaviour of discount rates. For instance, forward looking agents may discount the market at a premium over the average historic returns. Moreover, the term structure of interest rate models may be used to derive the

discount rate. In the context of the average rational learner, we shall assume that the discount rate is computed as the average of past realized returns. While this method is not innovative, it has some advantages. First of all, this value is always positive. Secondly, a small change in the discount rate will have huge effects on the present value. There are a variety of models which can be used to compute the discount rate, given different market conditions. For instance, it is possible to consider discount rates with forward looking premiums or which match risk attitudes. In this paper, the simplest case where agents discount the market according to the average of past realized returns is considered:

$$k_{t+1} = \sum_{j=1}^t r_{m,j}, \quad (8)$$

where $r_{m,j}$ is the market return from period $j-1$ to j .

In the presence of no model uncertainty, forecast dividends, growth and discount rates are all equal to the realized counterparts. Therefore in this setting the present value is simply:

$$PV_t = \frac{D_{t+1} + P_{t+1}}{k_{t+1} + 1} \quad (9)$$

This equation is used as a benchmark as to check for the whether markets have been underpriced.

2.1 Trading Strategy

The strategy involves comparing $PV_{i,t}$ with the current Price P_t . The strategy involves the following:

$PV_{i,t} > P_t$, hold equity index while expecting the market to go up,

$PV_{i,t} < P_t$, hold bonds while expecting the market to be bearish.

When the price is higher than the present value, it should be expected that the price should fall over time. Using this thinking, it implies that if agents hold the equity index, they will suffer from capital loss. The ideal strategy in this case would be to hold on to an asset of low risk. On the other hand, if price is lower than the present value, it should rise in subsequent periods, which will lead to an increase in the price. The theoretical return (μ_t) as a result of holding the market portfolio is therefore given by:

$$\mu_t = \frac{PV_{i,t} - P_t}{P_t}, \quad (10)$$

μ_t is the return that is expected as market adjust to the present value. μ_t can also be used as a benchmark measure to test for changing markets. If $PV_{i,t}$ is greater or less than P_t only at a margin, it may not be profitable to shift instruments in the presence of transaction costs. Therefore μ_t can act as a filter when the strategy is executed, and also provides a certainty criteria on whether to shift positions, which is ideal with individuals with different attitudes to risk. A more risk averse person, may be interested in executing the strategy when μ_t is high.

One assumption and a certain caveat of the present value model is that it assumes that reversion to the present value happens within the next year. In truth, it is known that such reversion may take years due to the presence of noise traders and speculators. However, the number of years for this reversion to happen is not uniform across time. A simplistic measure of one year of reversion is adopted. Observing μ_t over time may signal interesting behavior as it will show whether the market tends may be inflated.

3 Data and Results

Annual time series on price, one year bond yield and market returns prior from 1871 to 2013 was retrieved from Shiller's website.

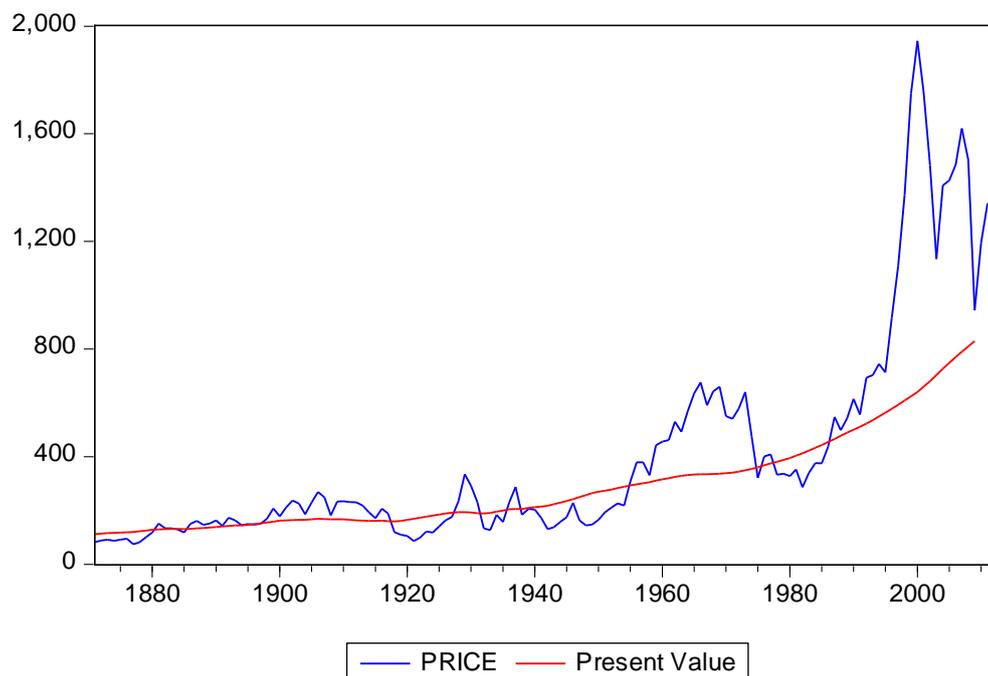


Figure 1: **Plot of present value and real S&P composite index from 1871 to 2013.** The figures have been adjusted for the real price level. The present value formulae assumes the constant interest rate.

Figure 1 illustrates the benchmark present values and the S&P index. The present value made up of the discount of one step ahead dividend and price tends to move very closely to the index. The present value based on no dividend growth appears of little relevance as the price always exceeds the present value. Underpricing and overpricing tends to be serially correlated. For instance, the stock market was underpriced from 1982 to 1986 and overpriced during the periods 2007-2008. Based on the sample size, the equity index is underpriced approximately 70 % of the time. In hindsight, the strategy postulates the correct position 60 % of the time. An interesting finding is that the the present value is lower than price when markets go down and it tends to be higher than the price when markets go up. In hindsight the strategy works well, with a terminal wealth of \$13,543 if \$1 was invested in 1927. The Buy and Hold strategy would actually yield a terminal wealth of \$214.47. The average annual return from the strategy is 12.5%. The present value in hindsight postulates the incorrect position (based on the optimal benchmark) only 10 times.

Model	PV_1		PV_2	
	W	R_t	W	R_t
1	16.69	3.69	12.80	4.06
2	12.40	3.22	12.80	4.06
3	5.58	2.22	12.80	4.06
4	9.96	2.97	12.80	4.06
5	7.68	2.61	12.80	4.06

Table 1: **Terminal Wealth and Annual Return.** The table shows the results from the strategy if returns dividends were forecast using models (1)-(5) for both variants of the present value.

The Buy and Hold strategy returns an average annual return of 8% with a terminal value of \$214.47. If the trading strategy is adopted, it can be easily be seen that none of the forecast and present value models beat the Buy and Hold. The best forecast model coupled with the one step ahead present value has an average return (4.06%) which is twice lower than the return from Buy and Hold. It is pertinent to note that when the second present value approach is adopted, all of the forecast models produce the same return. This finding simply implies that forecasting deviations do not matter when the larger part of the present value numerator is comprised of price. All ten different models show that the index is generally overpriced, thus recommending the holding of bonds throughout the sample. However, the trading strategy does not work in the strict form if the market does not fall in the next period. The theoretical return μ_t is plotted in figure 2.

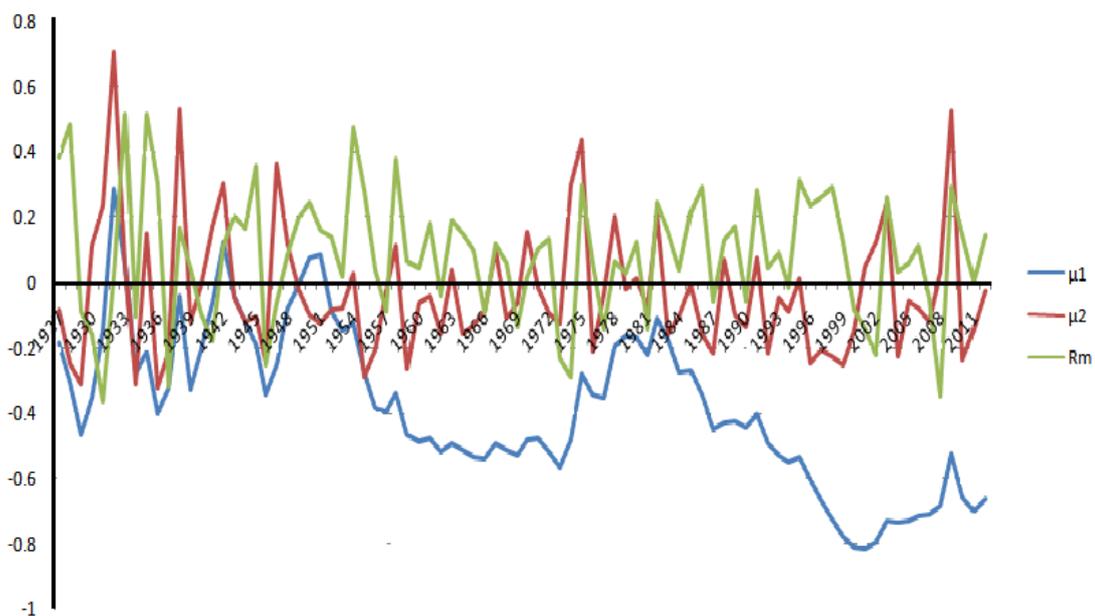


Figure 2: Plot of μ_t over time. The figure shows the plot of μ_t based on the two versions of the present value plotted over time against market returns.

Figure 2 shows the plot of μ_t when dividends are forecast using the Campbell and Shiller VAR model (equation 7). $\mu_{1,t}$ and $\mu_{2,t}$ refer to the present value models (Equation 1) and (Equation 2). Unlike $\mu_{1,t}$, which does not show any correlation with market returns, $\mu_{2,t}$ tends to display a pattern similar to the movement of stock market returns which does not show any correlation with market returns. An interesting finding from the plot is that when the present value increases, market returns increase as well. This finding has a very important behavioral implication as it means that momentum and past movements matter. However, the current trading strategy is only concerned with when μ_t crosses 0, which then posits a shift in the holding position; it does not pick up on the momentum in the present value itself.

Table 3 shows two cases in which a shift is advised if the theoretical return exceeds the constants of 0.3 and 0.6. In the case $\mu_t = 0.3$, holding the index is only advised if the theoretical return is expected to be higher than 30 %.

Model	$\mu_t = 0.3$				$\mu_t = 0.6$			
	PV_1		PV_2		PV_1		PV_2	
	W	R_t	W	R_t	W	R_t	W	R_t
1	88.06	6.05	149.36	7.73	104.97	6.92	214.48	8.24
2	61.60	5.69	206.89	8.15	193.88	7.76	214.48	8.24
3	61.60	5.69	206.89	8.15	151.59	7.52	214.48	8.24
4	61.60	5.69	206.89	8.15	151.59	7.52	214.48	8.24
5	68.78	5.86	206.89	8.15	151.59	7.52	214.48	8.24

Table 2: **Terminal Wealth and Annual Return when μ_t is 0.3 and 0.6.**

It can be seen from table 2 that if the expected return margin of greater than 0.3 is adopted, the terminal wealth increases to \$206.89. The average annual return is 8.15%, which is nearly twice as high as the return would be if the strict strategy were applied. When the higher benchmark of $\mu_t = 0.6$ is adopted, the trading strategy nests the Buy and Hold strategy. Hence, it is very interesting to note that an investor would require at least 60% of expected loss in market equity to hold bonds in an overpriced market.

3.1 μ_t based on previous positions.

In this section, An alternative measure is developed to compare with μ_t . In this case, the benchmark is based on the memory of incorrect holdings. For each model and alternative present value formula, μ_t is computed as the following:

$$\mu_t = \frac{\sum_{j=1}^{t-1} H_j}{t-1},$$

where H_j is a dummy variable which takes a value of 1 if the incorrect position was advised. This functions as a behavioral penalty point. The summation goes until $t-1$ as agents only know about their previous performances at time t . μ_t is bounded between 0 and 1, but it is heavily biased towards 1 in the initial stages when the sample size is small. Obviously, as agents learn, they get better at trading and the penalty point becomes lower. Thus, the expected return is equal to the penalty point. Table 3 illustrates the return from this scheme.

Model	PV_1		PV_2	
	W	R_t	W	R_t
1	52.30	5.75	187.32	7.94
2	100.62	6.72	187.32	7.94
3	100.24	6.86	187.32	7.94
4	128.78	7.11	187.32	7.94
5	159.28	7.47	187.32	7.94

Table 3: **Terminal Wealth and Annual Return based on Learning corrected μ_t** The table shows that that the average annual return and wealth as a result of using μ_t as a filter. The strategy is executed only if the theoretical return exceeds the filter.

Unlike results from the arbitrary case where μ_t is fixed at a high value, results from the behavioral μ_t indicate an improvement in the terminal wealth and the average rate of return. There are no significant differences created by the behavioral μ_t across the different forecasting models. However, it is very interesting to note that as μ_t becomes smaller over time, the strategy performs better.

One phenomenon revealed in 2.2 is the role of momentum. As previously discussed, there is a tendency for market returns to follow μ_1 . Therefore, the trading strategy is refined towards the following:

$$\mu_{1,t} > \mu_{1,t-1},$$

Hold the Market index

$$\mu_{1,t} < \mu_{1,t-1},$$

Hold Bonds

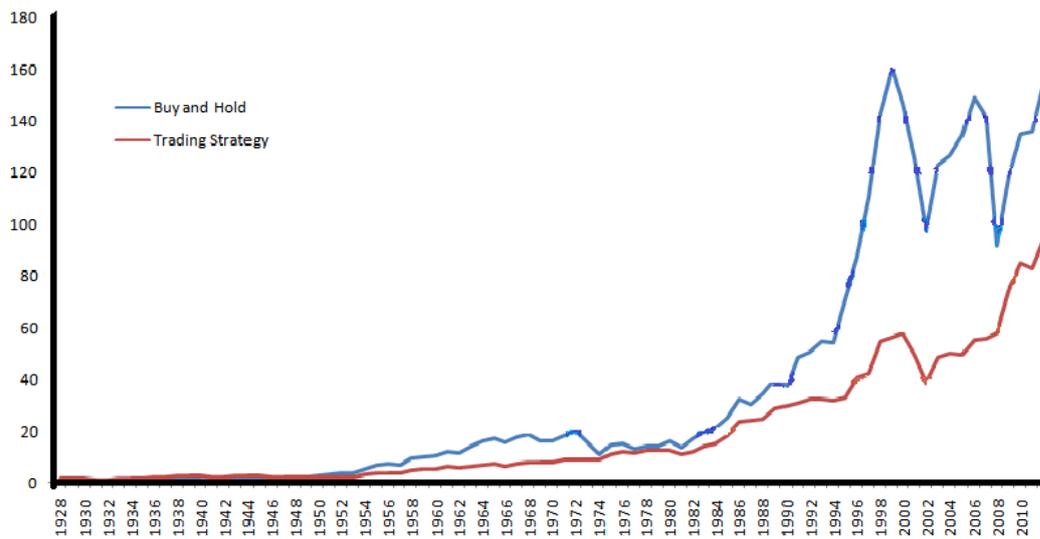


Figure 3: **Returns from the Trading strategy and Buy and Hold.** The y-axis shows \$ return, while the x-axis shows the years. The plot shows the worth of \$1 invested back in 1927.

The plot shows the cumulated returns from the trading strategy and buy and hold. Figure 3 shows the cumulated returns both Buy and Hold and the momentum strategy over the period. By taking into account previous movement of the theoretical return, the new strategy performs better than that based on the arbitraging argument. The average return from the trading strategy is 7.94 %. The trading strategy tends to perform quite well over the sample. The trading strategy postulates the correct position 59 % of the time. As shown in 3, the cumulative returns from the trading strategy are higher than those from the Buy and Hold until 1949, where they are overtaken by the return from the Buy and Hold . The period 1944–1952 appears to be the worst period for the momentum strategy, where the correct position was postulated only once. Other periods in which the strategy tends to perform poorly for a series of years are 1980–1982 and 1993–1995.

3.2 Risk Adjustment

The above computations fail to take risk into account. Although the strategy does not seem to yield higher excess returns, it may be profitable after taking risk into account . This is simply because the strategy advises holding bonds during some periods, which implies lower volatility. The periods in which bonds and equity are held should be accounted for in the

variance. A simple measure of performance to account for the periods in which bonds and equity are held is the Sweeney statistic (1988), which is given by the following:

$$X = R_s - (1 - f)R_{bh}$$

$$\sigma_x = \sigma[f(1 - f) / N]^{0.5}$$

where R_s refers to the return from the strategy and R_{bh} is the return from Buy and Hold. $1-f$ is the proportion of years that equity is held and σ measures the volatility of the stock market. The Sweeney Statistic simply involves computing the ratio $\frac{X}{\sigma_x}$. The null hypothesis of the one tail test is to see whether the returns from the strategy exceeds that of Buy and Hold at statistically significant levels. The statistic is equal to 1.86 in the case of the simple trading strategy. When memory of incorrect holdings and momentum strategies are considered the statistic is equal to -1.31 and 0.94 respectively, which does not indicate any statistical significance of higher returns from the strategy.

3.3 Decomposing Effects

The present value in real-time states that the market is more underpriced than usual. In this section, we explore the dynamics of the present value. In the following, the relationship between price and the present value is estimated. The following present value equation is estimated:

$$\ln P_t = \alpha_1 + \alpha_2 \ln(P_{t-1} + D_{t-1}) + \alpha_3 \ln(\mu_t + 1) + v_t$$

The estimated coefficients are -2.218, 1.017 and -0.765 for α_1 , α_2 and α_3 respectively. It is interesting to note that prices move inversely with regards to the discount rate. The payoff also matters for it has an elasticity close to unity. It is of interest to note that the intercept term is large, it is not statistically significant. To estimate the effects of dividend forecast errors and discount rate errors, we estimate three PROBIT Models:

$$Y_t = \alpha_1 + \alpha_2 fe_t + \alpha_3 de_t + v_t, \quad (11)$$

$$Y_t = \alpha_1 + \alpha_2 fe_{t-1} + \alpha_3 de_{t-1} + v_t, \quad (12)$$

$$Y_t = \alpha_1 + \alpha_2 fe_{t-1} + \alpha_3 de_{t-1} + \alpha_4 I_t v_t, \quad (13)$$

where $Y_t = 1$ when the position is incorrect, and 0 otherwise. fe_t refers to the ratio of the dividend forecast error to the realized dividend. de_t is the difference between the discount rate and the market returns normalized by the discount rate. I_t is a dummy variable when the holding should be in equity. The dummy variable is included to cope with asymmetries in the model. The result from the Probit regression are illustrated in table 3.3d.

	(11)	(12)	(13)
α_1	-0.934	-2.516	-1.282
	(0)	(0)	(0)
α_2	5.544	-54.25	5.906
	(0)	(0)	(0)
α_3	3.606	9.597	4.103
	(0.01)	(0.46)	(0)
α_4			0.410
			(0.36)
R^2	0.12	0.91	0.92

Table 4: **Results from PROBIT regressions.** The table shows the estimates and p-values (in brackets) from models 11, 12 and 13. $\alpha_1, \alpha_2, \alpha_3$ and α_4 show the PROBIT coefficients of the intercept term, forecast error/lagged forecast error, difference between the discount rate and market returns, and the equity position dummy variable.

In all three models (equations 11, 12 and 13), it is found that a higher difference between the discount rate and realized returns, the higher the probability of being in the wrong position. The same phenomenon applies with respect to the dividend forecasting error from the same period. However, the lagged forecast error reacts negatively with the position. The discount rate error appears to be stronger than the forecast error effect. The lagged discount rate error is also found to have a negative effect on the probability of being in the wrong asset. This phenomenon may be explained by the fact that if the discount rate is negative in the previous period, there should be some adjustment that occurs over the next period. This adjustment is updated in the discount rate at time t . When the asymmetry dummy is included, it does not change the original results by much, and it is found to be nonsignificant.

4 Conclusion

The success of these trading strategies depends on the correct estimate of the discount rate. A small change in the discount rate may lead to magnified changes to the present value and affect the decision of whether to go long on equity. It is worth noting that when the discount rate differs from realized returns, a stronger effect on the holding position is observed. The effect of forecasting errors is marginal in this case. Most importantly, this paper reiterates the importance of having a correct discount rate to compute present value.

The novelty of forecasting dividends appears futile if deemed to be independent from the discount rate. Discount rates appear to be the key indicator of market movements, as they are the main movements of the discount rate. Interesting Future studies may be suggested from the above. The trading strategy may take into account that the real present value appears to be higher when the market is rising and lower than the price when the market is falling. In terms of the trading strategy itself, a strategy may involve holding a portfolio of both equity and

bonds, where the weights are determined by a combination of previous forecasting errors and discount rate errors.

Other trading strategies include holding a stock in a particular position for more than one year until the present value is sufficiently high or low to switch positions. This ad-hoc measure of μ_t is subjective and does not lead to a higher terminal value. Another trading strategy would be to set μ_t as a weight of the forecasting error, which again does not lead to higher returns. However, as seen from the previous results, especially from the present value model (2.2), the forecasting error has little role to play in the present value. Another strategy that may be considered is an asymmetric momentum present value strategy where 3.1 is refined to consider effects on of whether μ_t is positive or negative. It may be likely that if μ_t is highly negative and then increases only marginally, the momentum strategy will posit going long on equity.

5 Bibliography

References

- Banz, R W, 1981. "The relationship between return and market value of commonstocks,"
Journal of Financial Economics, 9(1), 3 – 18.
- Brav, A & Heaton, J.B, 2002. "Competing Theories of Financial Anomalies," Review of
Financial Studies, 15(2), 575–606.
- Bulkley, G & Taylor, N, 1996. "A cross-section test of the present value model," Journal of
Empirical Finance, 2(4), 295 – 306.
- Bulkley, G & Tonks, I, 1989. "Are UK Stock Prices Excessively Volatile? Trading Rules
and Variance Bound Tests," The Economic Journal, 99(398), 1083–1098.
- Campbell, J & Shiller, R J, 1988. "The Dividend-Price Ratio and Expectations of Future
Dividends and Discount Factors," Review of Financial Studies, 1(3), 195-228
- Campbell, J & Shiller, R J, 1988. "Cointegration and Tests of the Present Value Models, "
Journal of Political Economy, 95(5), 1062-1088
- Hommes, C and Zhu, M, 2014. "Behavioural learning equilibria," Journal of Economic
Theory, 150, 778-814
- Horowitz, JL & Loughran, T & Savin, N.E, 2000. "The disappearing size effect. Research in
Economics," 54(1), 83 – 100.
- LeRoy, S F & Porter, R D, 1981. "The Present-Value Relation: Tests Based on Implied
Variance Bounds," Econometrica, 49(3), pp. 555–574.
- Lof, M 2014. " Rational Speculators, Contrarians, and Excess Volatility," Management
Science (forthcoming).
- Mankiw, N Gregory & Romer, David & Shapiro, Matthew D, 1985 "An unbiased
Reexamination of Stock Market Volatility," Journal of Finance, 40(3), 677-87

- Marquering, W, Nisser, J, Valla, T. 2006. "Disappearing anomalies: a dynamic analysis of the persistence of anomalies," *Applied Financial Economics*, 16(4), 291–302.
- Reinganum, M.R. 1981. "Misspecification of capital asset pricing. Empirical anomalies based on earnings' yields and market values," *Journal of Financial Economics*, 9(1), 19–46.
- Shiller, RJ, 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? ," *The American Economic Review*, 71(3), 3 – 18.
- Timmermann, A, Granger, CW. J, 2004. "Efficient market hypothesis and forecasting," *International Journal of Forecasting*, 20(1), 15 –27.
- Timmermann, A, 1993. "How learning in financial markets generate excess volatility and predictability in stock market prices,"*The Quarterly Journal of Economics*, 108(4), 1135-45