Does my beta look big in this?

Patrick Burns[*]

15th July 2003

Abstract

Simulations are performed which show the difficulty of actually achieving realized market neutrality. Results suggest that restrictions on the net value of the fund are particularly ineffective. A negative correlation—that is, market negativity—is proposed as a more reasonable target, both on theoretical and practical grounds. Random portfolios—portfolios that obey given constraints but are otherwise unrestricted—prove themselves to be a very effective tool to study issues such as this.

1 Introduction

How should I constrain my portfolio's beta in order to achieve market neutrality? That is the practical question to be answered.

Market neutrality is a very useful feature, and is well worth pursuing. The value of a fund to an investor is partly based on the return that it generates, and partly based on its correlation to the rest of the investor's portfolio—the lower the correlation, the more valuable it is. As [Beliossi, 2000] points out, a fund that claims market neutrality but doesn't deliver it, will evoke disappointment. Thus a methodology to produce market neutrality should be of keen interest to the fund manager.

Section 2 discusses correlation—perhaps the most natural measure of market neutrality—and its relation to beta. Experiments with various constraints aiming for market neutrality are shown in Section 3. Section 4 moves to market negativity, and Section 5 concludes.

2 Betas and Market Correlation

An asset's beta is defined as the slope coefficient in a linear regression in which market returns explain the returns of the asset. One approach to computing the beta is to manipulate elements of the variance matrix of the asset and the market. Specifically beta is equal to the covariance between the asset and the market

---

divided by the variance of the market—see, for example, [Weisberg, 1980]. We can use this definition to establish the relationship between a portfolio's beta and its correlation with the market.

A portfolio's beta is

$$\beta_P = \sum_i w_i \beta_i = \sum_i w_i \frac{\text{Cov}_{i,M}}{\text{Var}_M} \tag{1}$$

where w_i is the weight in the portfolio of asset i.

The covariance of the portfolio with the market is

$$\text{Cov}_{P,M} = \sum_i w_i \text{Cov}_{i,M} \tag{2}$$

Using these two equations and the fact that a correlation is equal to the corresponding covariance divided by the two volatilities, we get the relation that we are looking for:

$$\text{Cor}_{P,M} = \frac{\text{Vol}_M}{\text{Vol}_P} \beta_P \tag{3}$$

So the correlation between the portfolio and the market is equal to the portfolio beta times the ratio of volatilities. This seems to imply that the lower your portfolio volatility, the harder it is to be market neutral. The simulations in Section 3 suggest that the opposite is actually true.

The alert reader may notice that substituting equation 1 into equation 3 makes it look like any correlation can be achieved just by changing the scale of the weights. However, the same weights appear in the portfolio volatility, so a change in scale cancels out. In what follows we use the convention that the sum of the absolute value of the weights is 1. That is, an asset's weight is its position divided by the gross value of the portfolio.

3   Market Neutrality Simulations

A simulation study was performed to find the constraints needed to achieve a desired level of market neutrality. The variance matrices were created using a statistical factor model as described in [Burns, 2003] using one year of daily data. The betas were computed from the estimated variance matrix which was also used to compute the portfolio and market volatilities. The data are a rather random collection of 200 US stocks—some large cap, some small cap. The "market" was taken to be the S&P 500.

The realized values are for the quarter following the period over which the variance is estimated. For each combination of constraints, 1000 random portfolios were generated for each quarter using the POP Portfolio Optimization software of Burns Statistics.

In this section all of the random portfolios can have their longs and shorts interchanged and still satisfy the same constraints. Thus the distribution of

Figure 1: Absolute realized correlations by quarter.
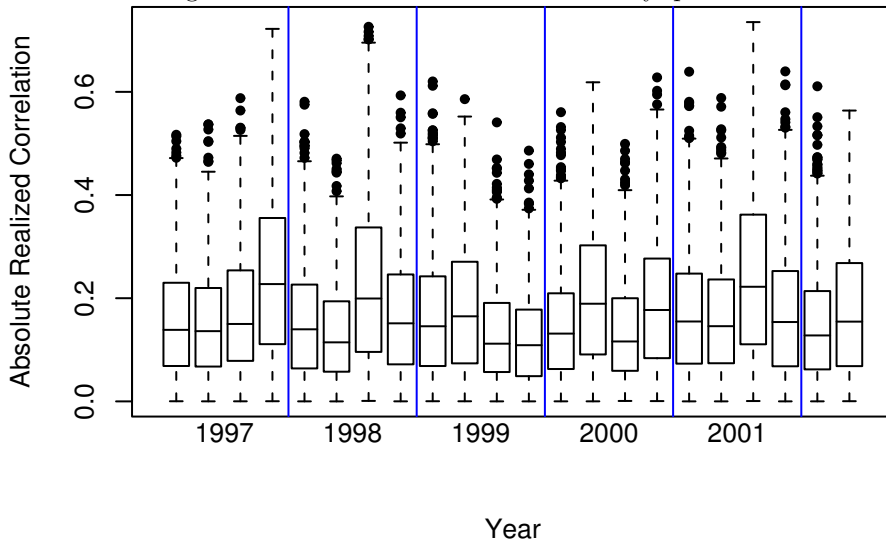
betas and correlations are symmetric about zero. In turn this implies that looking at the absolute value of the correlation is appropriate.

Figure 1 shows a boxplot of the absolute value of the realized correlation for each quarter. Each end of a box is a quartile, and the median is indicated within the box. The data are for portfolios of size 40, and constraints to not be greater than 10% were placed on each of: 1) the absolute value of the portfolio beta, 2) the volatility, and 3) the absolute value of the net relative to the gross. Clearly the correlations are not necessarily small.

Figures 2 through 4 show results in the second quarter of 2002 for portfolios of 40 stocks with volatility less than 10%, and the absolute value of the net less than 10% of the gross. Figure 2 compares the predicted and realized betas when the predicted betas are constrained to be less than 1% in absolute value. It appears that there is no relation between the predictions and the realized values. Figure 3—with the betas less restricted—shows that there is a relationship, but there is quite a bit of noise. This relationship feeds through to the correlations as shown in Figure 4, where we can see that the realized correlation can be substantially larger than the predicted correlation.

Figure 5 is similar to Figure 4 except that the absolute net value is restricted to be less than 1% of the gross instead of 10% of the gross. The realized correlations are unaffected by the increased restriction on the net. Thus if the sole purpose of the more restrictive condition on the net were to get market neutrality, then it would be counter-productive. Tightening the constraint doesn't do what is intended, but it does limit the alpha more. Generating random portfolios is an easy way to gauge if constraints are too tight or too loose.

3

Figure 2: Predicted and realized betas with the predicted beta less than 1% in absolute value.
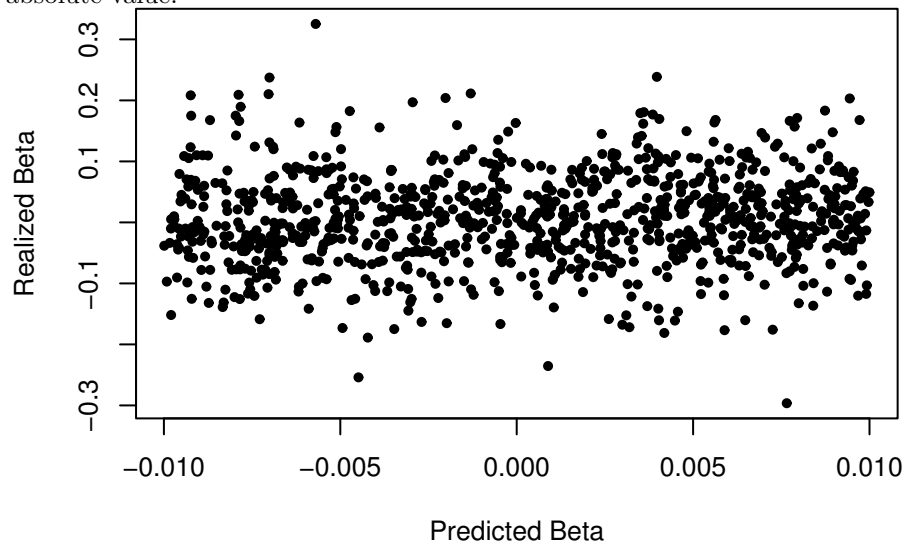

Figure 3: Predicted and realized betas with the predicted beta less than 10% in absolute value.
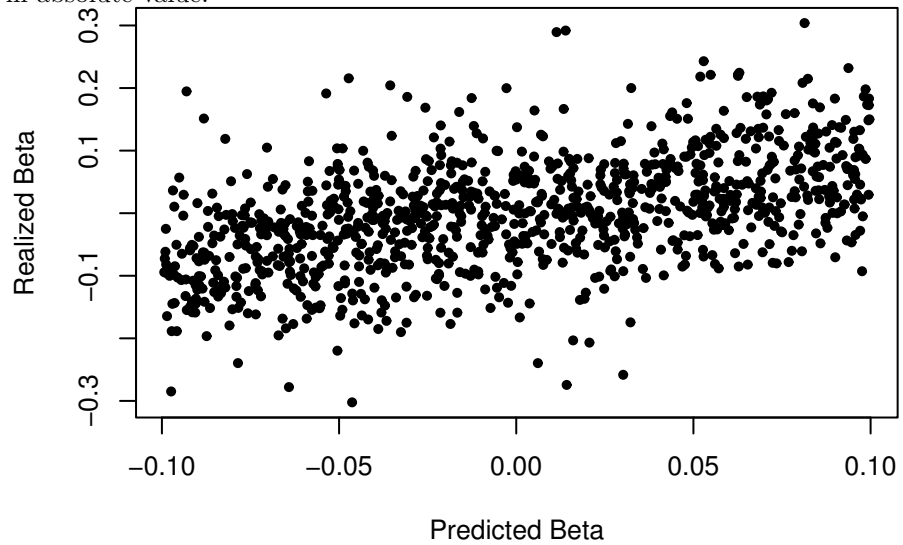

4

Figure 4: Predicted and realized correlations with the net and the predicted beta less than 10% in absolute value.
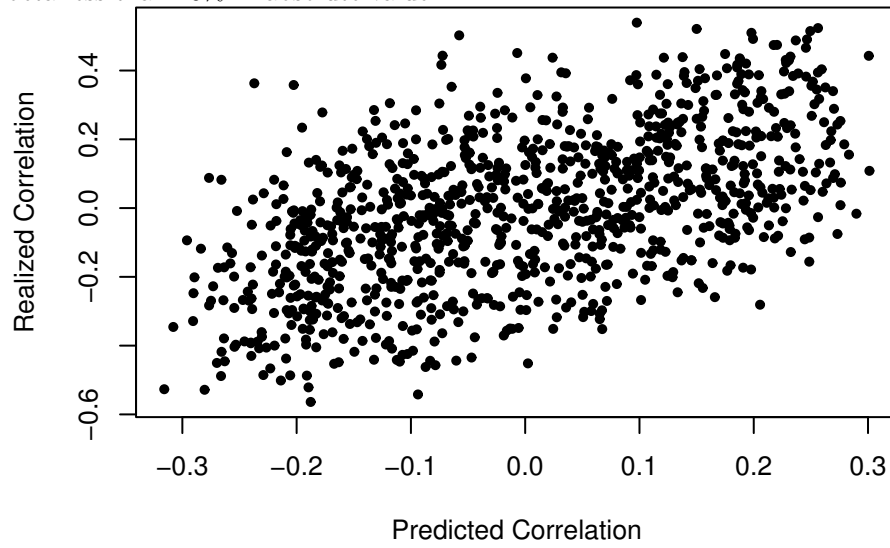

Figure 5: Predicted and realized correlations with the net less than 1% and the predicted beta less than 10%, both in absolute value.
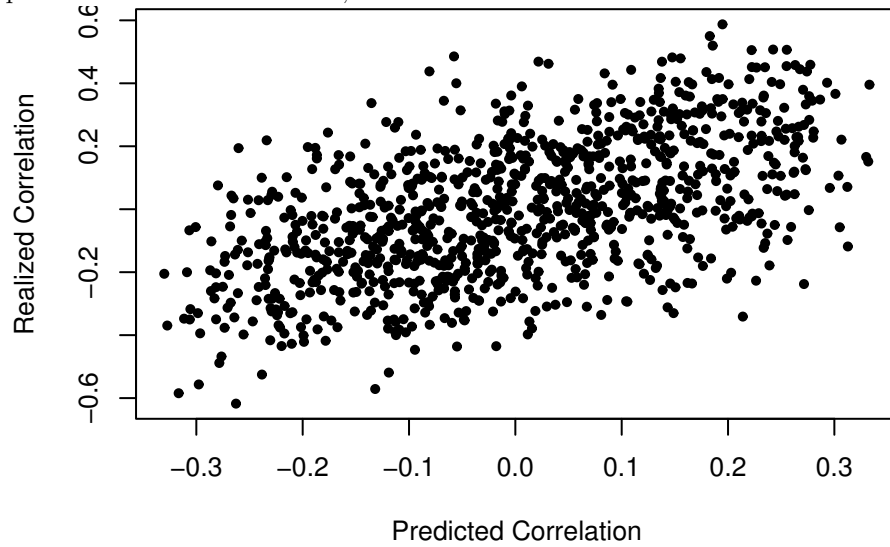
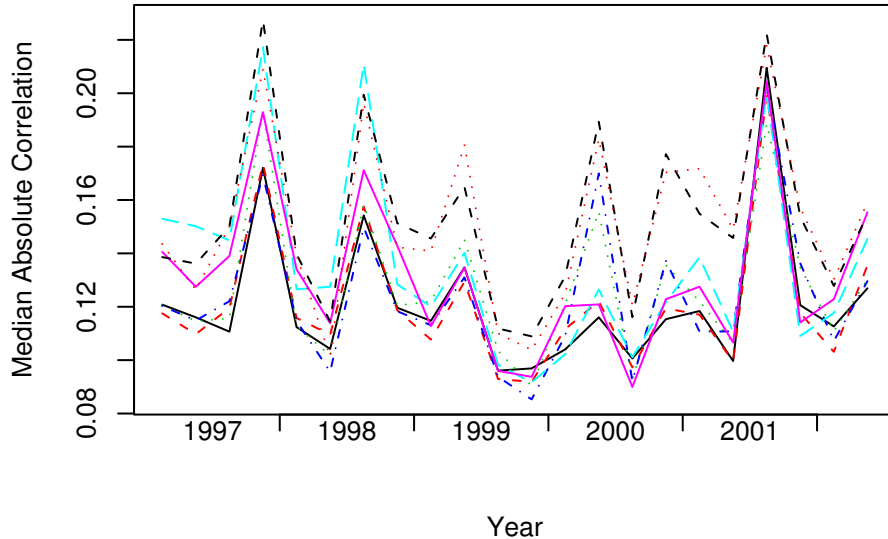Figure 6: Median of absolute correlation for eight constraint combinations.

Table 1: Mean of the median absolute correlation by constraint type.

	low value	high value
beta	12.4 (@ beta < 1%)	14.3 (@ beta < 10%)
volatility	12.7 (@ vol < 5%)	13.9 (@ vol < 10%)
net	13.2 (@ net < 1%)	13.4 (@ net < 10%)

In Figure 6 we look at the median of the absolute value of the correlation in the 1000 random portfolios for each quarter and eight constraint combinations using portfolios of size 40. The betas are constrained to be either less than 10% or less than 1% in absolute value, as is the net relative to the gross. The volatility is constrained to be either less than 10% or less than 5%. The most important factor is time—some quarters have much larger correlations than others. The levels of the constraints do have an impact, but none of the constraint combinations is dramatically superior or inferior.

We can see more clearly the effect of the various constraints by averaging over the quarters and the runs. These averages are shown in Table 1. The most significant is the change in beta from 10% to 1%, the next most important is the change in volatility from 10% to 5%. Changing the net from 10% to 1% has only a marginal effect.

A few runs using portfolios of size 100 instead of 40 were performed. The only noticeable difference is that the correlations tend to be slightly bigger with the bigger portfolios.

Figure 7: Realized correlations of random portfolios with negatively constrained betas.
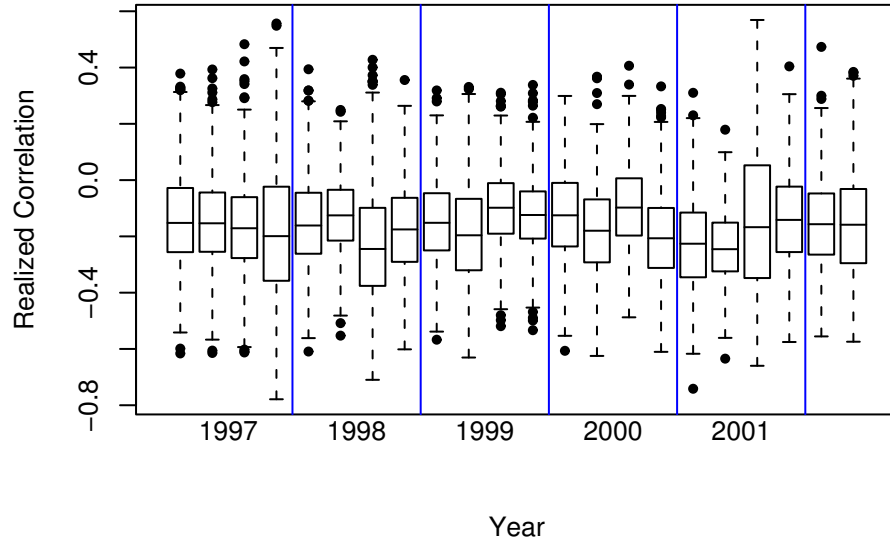

4 Market Negativity

The simulations of constraints for market neutrality show that actually achieving a small correlation with the market is not a trivial task. We also know that the only thing better than a fund with zero correlation is one with a negative correlation. Hence it makes sense to aim for a negative correlation. On average the portfolio should then be negatively correlated to the market, and have only a small chance of significant positive correlation at any point in time.

Figure 7 shows the correlations that are realized from random portfolios of size 40 with the beta constrained to be between -10% and -5%, the volatility constrained to be less than 5%, and the absolute value of the net constrained to be less than 1% of the gross. 1000 portfolios are generated for each quarter. Figure 8 shows the probability by quarter of the realized correlation being greater than 0 and of being greater than 20%. While there is still a small chance of a large positive correlation with this set of constraints, it would be highly unlikely for a large positive correlation to exist through several rebalances.

A common question is likely to be: if the fund is market negative, doesn't that mean that it will lose money on average since the market gains money on average? No. The correlation measures the relationship after the mean return has been accounted for. In theory a fund could have any combination of expected return and correlation with the market. The ideal fund would go up when the market goes up, and go up a lot when the market goes down.

Figure 8: Probability of the realized correlation being greater than 0 or greater than 20% with negatively constrained betas.
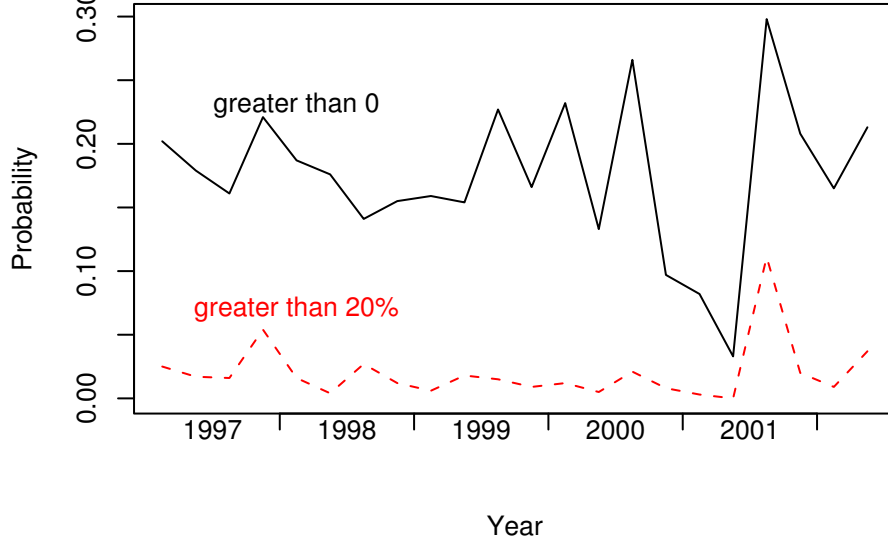
5 Discussion

While the betas used in the simulations are decidedly not optimal—they were selected for computational convenience rather than optimal prediction—even the best possible betas will engender substantial differences between the realized and the predicted correlation. Much of this difference will be observational noise. The true (but unknowable) correlation may be zero, but the 60-some days of data in a quarter can show a computed correlation that is quite different from zero. However, the investor cares about what actually happens—not some "true" state that no one ever sees or knows.

Investors also care about getting a low volatility for their entire portfolio—they do not care about market neutrality per se. In fact the perfect correlation from the investor's point of view is -1, not 0. Moreover, the measurement should be against the investor's portfolio, not the "market". With a correlation of -1 investors could have a guaranteed return, and would be more than happy if the combined return were greater than the (mythical) risk-free rate.

A correlation of -1 with a good return is a pipe dream, but aiming for a negative correlation is both feasible and sensible. The fund manager needs to find the optimal tradeoff between pushing the correlation lower and losing expected return because of the more stringent constraints.

It seems that we can answer—in broad terms at least—the original question. No lower bound on the portfolio beta should be given. The preference is for betas that are as negative as possible, so it makes no sense to put a lower bound on the

8

beta. An upper bound on the portfolio beta that is negative—perhaps between 0 and -10%—seems reasonable. This should be explored in each particular case. One means of exploration is to vary the bound on beta and see how that affects the expected return of the optimal portfolio. A complementary approach would be to generate random portfolios and see how the beta constraint changes the upper tail of the distribution of realized returns of the random portfolios.

References

[Beliossi, 2000] Beliossi, G. (2000). Take the long and short route. *Risk*, 13(11):105–107.

[Burns, 2003] Burns, P. (2003). On using statistical factor models in optimizing long-only portfolios. Working paper, Burns Statistics, http://www.burns-stat.com/.

[Weisberg, 1980] Weisberg, S. (1980). *Applied Linear Regression*. John Wiley & Sons.