NORGES BANK
INVESTMENT MANAGEMENT

01 | 2018

NO-TRADE BAND REBALANCING RULES: EXPECTED RETURNS AND TRANSACTION COSTS

DISCUSSION NOTE

In this note, we explore equity share rebalancing using no-trade band rules. We compare rules based on their implications for the behaviour of the equity share, expected returns, turnover and transaction costs.

Date 29/08/2018
ISSN 1893-966X

dn@nbim.no
www.nbim.no

SUMMARY

- We evaluate rebalancing rules that rebalance the equity share when it breaches a predefined no-trade band. These rules allow the equity share to drift over time and save transaction costs, though bands imply variation in the equity share over time that may have expected return implications. We focus on comparing alternative no-trade band rules with a continuously rebalanced benchmark, rather than revisiting the question of whether to rebalance at all.
- A wider no-trade band reduces the likelihood that the equity share is rebalanced and implies a longer horizon over which the equity share drifts in the no-trade band. This leads to larger deviations of the equity share from its strategic level and a higher equity share on average.
- We consider whether equity share drift is related to variation in expected returns over time. We use simulations to show that whether the equity share drifts positively or negatively with expected returns depends on the width of the no-trade band. This results from the balance between momentum- and reversal-style components of expected returns changing with different horizons. The ability of no-trade bands to exploit time-varying expected returns is economically small and uncertain, however, implying that no-trade bands are an ineffective strategy for capturing time-varying expected returns.
- We calibrate a transaction cost model to illustrate the extent to which transaction costs vary across rebalancing rules. Wider no-trade bands, trading partially towards the strategic level and trading more slowly all lead to cost reductions and lower turnover.
- Our analysis suggests that an investor's tolerance for a higher average equity share and higher equity share variability should be weighed against the transaction cost savings from no-trade bands. The uncertainty around the relationship between the equity share and expected returns under no-trade bands suggests that this should be a secondary consideration when comparing rules.

Introduction

An important consideration for many investors is whether and how to rebalance when differences in the relative performance of assets lead to drifts in their portfolio weights. One example of this issue is the drift in the equity share of a multi-asset portfolio with equity and fixed income components. Without rebalancing, the equity share can deviate significantly from its strategic or starting level, and as a result many investors choose to rebalance frequently.

Frequent rebalancing, however, can lead to significant trading volumes and transaction costs. An investor may therefore choose to follow a rule that reduces rebalancing events. A common approach is to establish a no-trade band around the strategic equity share, within which the equity share can vary without rebalancing occurring.

In this note, we compare a range of no-trade band rebalancing rules. The use of no-trade bands means that the equity share can drift over time, and this achieves a reduction in trading and transaction costs relative to rebalancing continuously. The drift in the equity share also implies a different distribution of the equity share over time, however, which potentially has expected return and risk implications. We assess the implications for drift in the equity share, expected returns and transaction costs for a range of no-trade band rebalancing rules. We vary the rules based on the width of the no-trade bands, and also consider alternative targets to which the equity share is returned (inner bands), and different speeds at which the equity share is rebalanced (trade size).

The starting point for our analysis is that the decision to rebalance the equity share is already in place. The case for and against rebalancing has been well-documented, where particular benefits include maintaining allocations at or near optimal portfolio weights and ensuring portfolio diversification.[1] In addition, the counter-cyclical nature of rebalancing and the discipline imposed by a rebalancing rule may be especially appropriate for long-term investors, as argued in Ang and Kjaer (2011). A large related literature focuses on the 'rebalancing premium', which is the tendency for the better diversification of rebalanced portfolios to positively skew the distribution of portfolio value at long horizons. Rather than revisiting the question of whether to rebalance or not, we focus on how and when rebalancing occurs through the comparison of alternative no-trade band rebalancing rules.

In the first part of the note, we show how no-trade band rules change the distribution of the equity share. A wider no-trade band reduces the likelihood that the equity share is rebalanced. This implies larger deviations of the equity share from its strategic level and a higher equity share on average, and implies that portfolio risk changes over time. We then consider whether the

[1] There is a large literature, often practitioner oriented, examining the case for rebalancing and comparing alternative rebalancing rules. See for example Arnott and Lovell (1993), Leland (1996), Donohue and Yip (2003), Masters (2003), Jaconetti, Kinniry and Zilbering (2010), Ilmanen and Maloney (2015), and Huss and Maloney (2017). See also NBIM Discussion Notes #1, #2, #3 and #4 (2012) for further discussion and analysis.

equity share drift within no-trade bands captures variation in expected returns over time.

One consequence of time-varying expected returns for portfolio choice is that a long-term investor should vary the equity share conditional on the level of expected returns over time (e.g. Brennan, Schwartz and Lagnado (1997) and Campbell and Viceira (1999)).[2] In this note, we consider rebalancing rules that keep the strategic equity share fixed over time. As a result, ability to time expected returns only arises through the equity share drifting positively or negatively with changes in expected returns within the no-trade region. We guide our analysis using the evidence on time-varying expected returns, but note that varying the strategic equity share would be a more direct way of exploiting time-varying expected returns.[3]

We use simulations to show that the ability of no-trade bands to exploit time-varying expected returns depends on the width of the no-trade band. We specify expected return dynamics in terms of 'reversal' and 'momentum' components, where reversal effects are much more persistent than momentum effects. We show that when momentum effects dominate, which is the case over shorter horizons, the equity share co-moves positively with expected returns within no-trade bands. A narrower band shortens the period of time the equity share drifts before rebalancing is triggered, and this helps no-trade bands to exploit expected return variation. From a practical point of view, however, these effects are economically small and statistically weak, and an investor cannot be confident that no-trade bands can meaningfully exploit time-varying expected returns.

In the second part of the note, we explore the extent to which alternative no-trade bands, inner bands and trade sizes can reduce transaction costs and turnover. The relationship between no-trade band rebalancing and transaction costs has been the subject of a number of academic studies. The focus is often on optimal portfolio choice in the presence of different types of transaction costs, where it is optimal to establish no-trade bands as a rebalancing strategy. We calibrate a transaction cost model to illustrate the extent to which costs vary across rules, using insights from this literature to aid the specification of our model, in particular for a price impact function (e.g. Frazzini, Israel and Moskowitz (2015) and Novy-Marx and Velikov (2015)). Wider no-trade bands, wider inner bands and slower trading all lead to cost reductions.

Overall, the change in the distribution of the equity share under no-trade bands should be considered alongside transaction cost savings when comparing alternative rules. A wider no-trade band saves transaction costs but leads to a higher average equity share and higher variability in the equity share. The uncertainty associated with capturing time-varying expected

[2] This implication has been derived in a partial equilibrium context, and naturally in a general equilibrium setting it is not possible for all investors to rebalance at the same time, in the same direction. Heterogeneity across investors is required to account for different rebalancing behaviour in general equilibrium.

[3] The debate over whether time-varying expected returns can be exploited in real time is ongoing, however, and the feasibility of timing the strategic equity share would need to be explored further.

returns implies that this should be a secondary consideration when comparing no-trade band rebalancing rules.

The note proceeds as follows. In Section 1, we define the features we vary in alternative no-trade band rebalancing rules, and describe the data and methodology used in our analysis. In Section 2, we explore variation in the equity share under alternative rules using bootstrapped returns. In Section 3, we consider whether no-trade rebalancing can capture time-varying expected returns. In Section 4, we specify a calibrated transaction cost model to understand the scope for reducing transactions costs using alternative rules. Section 5 concludes.

1. Definitions and Methodology

In this section, we define the main features of the no-trade rebalancing rules we consider, and outline the data and methodology used to assess alternative rules.

No-trade Band Rebalancing: Definitions and Methodology

In specifying a no-trade band rebalancing rule, the primary feature to define is the width of the no-trade band around the strategic equity share. The width of the no-trade bands defines the band outside which the equity share needs to move in order for rebalancing to be triggered. We refer to the width in terms of percentage points on either side of the strategic equity share, for example a 4 percentage point no-trade band around a 70 percent strategic level implies that the outer edges of the no-trade band are at 66 percent and 74 percent.

We also consider other features of rebalancing rules that are relevant once rebalancing has been triggered. We allow for 'inner bands' around the strategic equity share, which are the levels to rebalance to following a breach of the no-trade band. For example, a 70 percent strategic equity share with 2 percentage point inner bands implies that the equity share is rebalanced to a 72 percent (68 percent) equity share if the upper (lower) side of the no-trade band is breached.[4]

We also vary the speed at which the equity share is returned to the strategic equity share or inner band, which we refer to as 'trade size'. Trade size is defined as the fraction of the portfolio, in basis points per month, that is traded each period towards the rebalancing target. Figure 1 shows a simple illustration of these key features.

In this note, we consider symmetrical no-trade bands, inner bands and trade sizes, using the same magnitudes whether rebalancing on the upside or downside. We define no-trade bands and inner bands that are between 1 and 6 percentage points above/below the strategic equity share, and trade sizes from 25 basis points per month to a single trade ('one-step') to the strategic level. Table 1 summarises the ranges of the rule parameters.

[4] When rebalancing to an inner band, if the equity share drifts past the inner-band target though remains within the no-trade band region, no more rebalancing trades occur until the next time the no-trade band is breached.

Figure 1: No-trade band rebalancing rule definitions

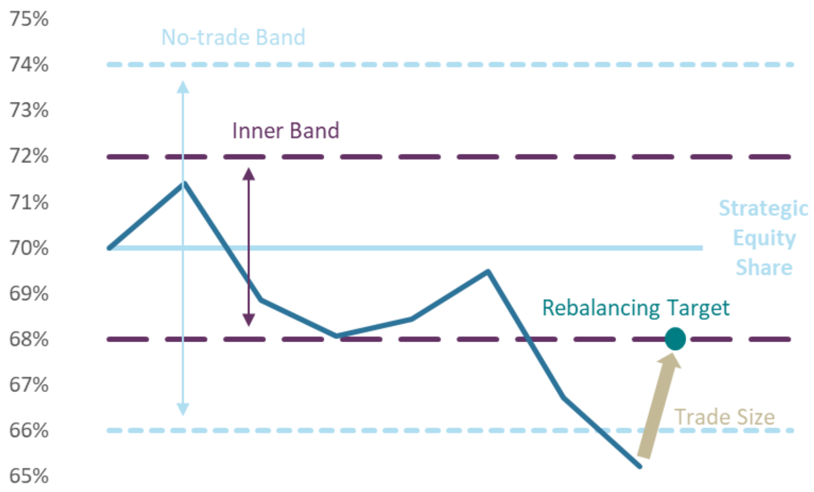

We implicitly consider zero no-trade bands when rebalancing continuously, and zero inner bands when rebalancing fully to the strategic equity share. The continuously rebalanced portfolio provides a benchmark in our analysis for comparing rules and is defined as a portfolio that rebalances the equity share to a constant 70 percent every month.[5] As discussed as in the introduction, we use continuous rebalancing as a benchmark for comparing the performance of alternative rebalancing rules, rather than focusing on the question of whether to rebalance or not.

Table 1: No-trade band rebalancing rule parameters

No-trade band (percentage points +/-)	1, 2, 3, 4, 5, 6
Inner band (percentage points +/-)	1, 2, 3, 4, 5, 6
Trade size (basis points per month)	25, 50, 100, 150, One-step

Data and Methodology

We evaluate alternative rebalancing rules for an equity and fixed income portfolio represented by aggregate US equity and fixed income returns. We use monthly returns calculated from the MSCI USA total return index and the Bloomberg Barclays US Treasury total return index over the period from January 1973 to December 2017.[6] For the majority of our analysis, we set the strategic equity share to 70 percent.

[5]While close to constant, the continuously rebalanced benchmark equity share can drift a little every other month, since we require a month for rebalancing to take place. This provides a more realistic representation of a constantly rebalanced benchmark. For all our analysis, this slight drift in the equity share is of little consequence, and all results are unaffected if using an instantaneously rebalanced/truly constant benchmark.

[6]The start of the sample period is determined by the availability of the fixed income return data.

For much of our analysis we use a bootstrapping methodology to re-sample historical returns. We block-bootstrap 1,000 50-year samples of equity and fixed income returns to which we apply the alternative rebalancing rules.[7] Resampling historical data through bootstrapping allows us to avoid making assumptions about key statistical properties of equity and fixed income returns, such as their predictability, volatility and correlation. We also consider calibrated simulation models, however, which are described in more detail later in the note.

Throughout our analysis, we estimate the effects of alternative rebalancing rules by averaging across the bootstrapped samples. We also include interval estimates around the averages to understand the statistical significance of effects, and to convey the degree of uncertainty associated with different effects.[8]

2. The Behaviour of the Equity Share under No-trade Bands

In this section, we consider the implications of no-trade band rebalancing rules for the behaviour of the equity share over time. A natural consequence of no-trade bands is that the equity share will drift within the no-trade band until rebalancing is triggered. The likelihood that the equity share breaches the no-trade band and triggers rebalancing depends on many factors. Some of these factors relate to the properties of asset returns, including the expected returns and volatilities across the equity and fixed income asset class components, and also the correlation across asset classes.

An important additional factor is the level of the strategic equity share. Figure 2 shows the average number of years between rebalancing events for different levels of the equity share, holding the no-trade band width at 4 percentage points, with no inner band and trading in one step. An increase in the strategic equity share from 60 to 70 percent with a fixed percentage point no-trade band width increases the average length of time between rebalancing events occurring. As the equity share is changed from around 50 percent, in either direction, the portfolio is increasingly dominated by one asset class and becomes more homogeneous. More extreme relative returns across the asset classes are therefore needed to trigger rebalancing. Other factors that influence the extent of deviations in the equity share, which we focus on, are concerned with the specification of the no-trade rebalancing rule.[9] A wider no-trade band, holding the strategic equity share fixed,

[7] To preserve the time series properties of returns, such as persistent variation in expected returns and volatility, we resample blocks 12 months in length. For highly persistent expected returns, the block length is likely too short and would not capture long-term reversal effects in returns, which we address by bootstrapping a VAR model later in the note. However, smaller/larger block sizes do not lead to meaningfully different results. We use a moving-block bootstrap (Kuensch (1989)), which allows for overlapping blocks and delivers better statistical properties of bootstrapped series than a non-overlapping block-bootstrap.

[8] It should be noted that this uncertainty only reflects randomness in the return-generating process and the extent to which a given rule would behave differently with alternative sample realisations. It does not capture the parameter uncertainty inherent in bootstrapping returns, simulation calibrations or transaction cost estimates.

[9] All results in the section are based on the bootstrap methodology described in Section 1. The

Figure 2: Average number of years between rebalancing events by equity share (4pp no-trade band)

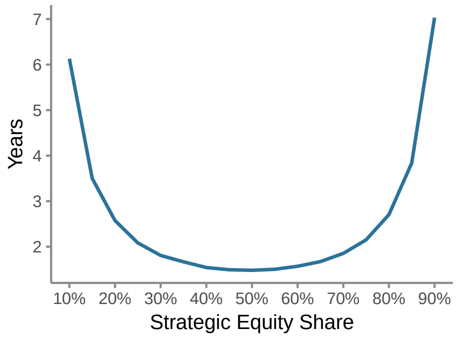

Figure 3: Average number of years between rebalancing events by no-trade band (70 percent strategic equity share)

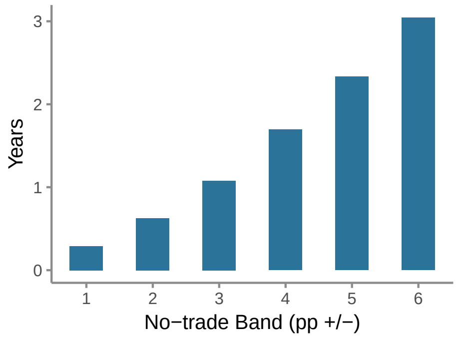

decreases the probability of rebalancing taking place. This can be seen in Figure 3, which shows the average number of years that occur between rebalancing events for alternative no-trade band widths.

One consequence of imposing no-trade bands is that the equity share is likely to be higher on average than with continuous rebalancing. Given higher average returns for equities relative to fixed income, the equity share will tend to increase over time as the equity component outperforms the fixed income component. To understand the extent to which the average equity share increases, we estimate the average equity share for alternative rebalancing rules for no-trade bands between 1 and 6 percentage points. The average deviations from a 70 percent strategic level are shown in Figure 4, in percentage points.[10] The average equity share increases from 70 percent for a 1 percentage point band to 71.2 percent for a 6 percentage point band.[11, 12]

The figure also includes the 5th and 95th percentiles of the distribution of 50-year samples to which the alternative no-trade rules are applied. The realised average equity share can vary across different samples, and the percentiles provide a sense of the significance of the average effect.[13]

In addition to differences in the average level of the equity share, varying the width of the no-trade band naturally allows additional variation in the equity share around its strategic level. We measure this effect through "tracking error": the standard deviation of relative returns for a rebalancing rule relative

results are robust to using the VAR-bootstrap methodology that we employ later in the note, which allows for possible effects of highly persistent expected returns.

[10] Naturally, this deviation could be corrected for by rebalancing towards a target level that is below the 70 percent strategic level, such that the equity share is 70 percent on average. When comparing returns for different rules, we adjust for these differences.

[11] The positive drift in the equity share also implies that the upper limit of the no-trade band is more likely to be breached than the lower limit. Based on bootstrapped samples, the probability of rebalancing down from the upper limit is approximately 60 percent, compared to 40 percent for rebalancing upwards from the lower limit.

[12] This increase in the average equity share may be over-estimated from a forward-looking perspective, to the extent that the realised equity premium in our bootstrapped samples is higher than today, evidence for which is summarised in NBIM Discussion Note #1 (2016).

[13] The intervals cannot be used to compare the significance of differences across no-trade band widths, since the co-variance needed to perform the test is not included. The widths of 1, 2 and 3pp bands are significantly different from each other at the 10 percent level. 2pp differences between band widths (e.g. 2 and 4pp bands) are significantly different at the 10 percent level for all differences except 4 and 6.

Figure 4: Average equity share deviation by no-trade band width

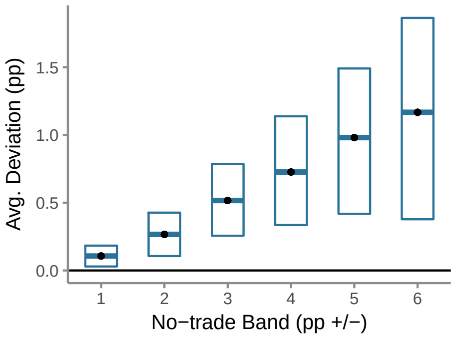

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 5: Tracking error vs. continuous rebalancing by no-trade band

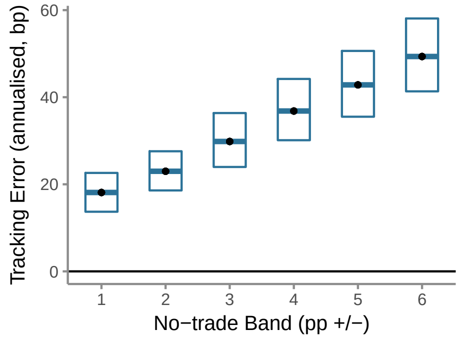

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

to continuous rebalancing. Annualised tracking error for different band widths is shown in Figure 5, where average values range from 18 basis points to 49 basis points.

Inner band size and trading speed also influence the path of the equity share within the no-trade band. To analyse the effect of alternative inner bands on the average equity share and the variability of the equity share, we hold the no-trade band fixed at 4 percentage points.[14] Figure 6 shows the average deviations from a 70 percent strategic level for different inner bands. Including an inner band of 4 percentage points increases the average equity share by around 70 basis points. It should be noted, however, that the differences across inner bands are statistically insignificant.

Figure 6: Average equity share deviation by inner band

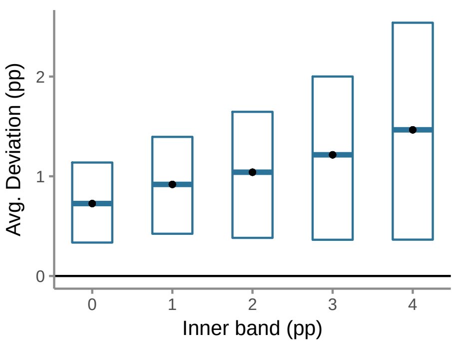

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 7: Tracking error relative to continuous rebalancing by inner band

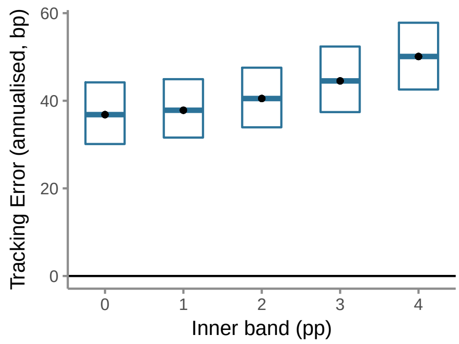

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Inner bands also lead to an increase in tracking error, where trading to points

[14] The expected impact on the average equity share of setting an inner band is somewhat ambiguous: trading to a higher point when rebalancing from an upper limit breach of the no-trade band should lead to a higher equity share on average. On the other hand, trading to a lower point when rebalancing from a lower limit should lead to a lower equity share on average.

away from the strategic equity share naturally leads to larger equity share deviations. Figure 7 shows annualised tracking error for alternative inner band widths. Average tracking error ranges from 36 basis points with no inner band to 50 when adding a 4 percentage point inner band.

We can also assess the effects of varying rebalancing trade size on the average equity share. We again fix the no-trade band at 4 percentage points and do not impose inner bands in the rebalancing rules.[15] Figure 8 shows the effect of varying trade sizes on the average equity share, and the impact is relatively small. When trading at 25 basis points per month, or trading in a single step, the average equity share varies within a range of approximately 50 to 75 basis points. Varying the trade size has a more noticeable effect on tracking error, shown in Figure 9, which decreases from 71 to 37 basis points when trading either at 25 basis points per month or in a single step, and this difference is statistically significant.

Figure 8: Average equity share deviation by trade size

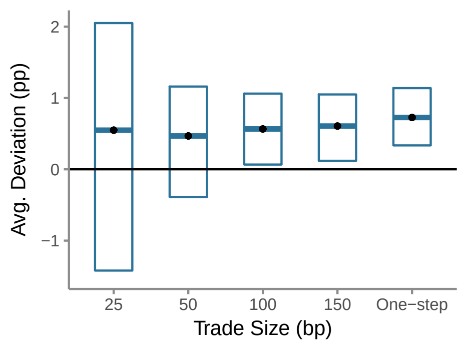

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 9: Tracking error relative to continuous rebalancing by trade size

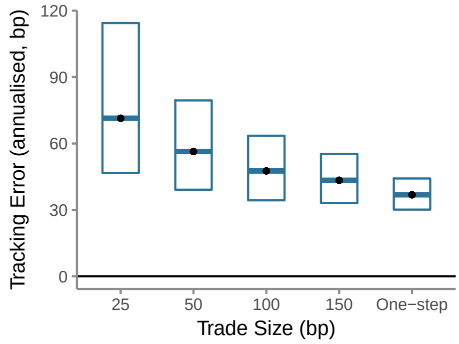

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Overall, wider no-trade bands, wider inner bands and smaller trade sizes lead to a higher equity share on average, and increased equity share variability. The increase in average equity share from no-trade band rebalancing implies that an investor is bearing additional exposure to equities. A wide no-trade band may imply a deviation from an 'strategic' or 'optimal' equity share level, and an investor may also be concerned with letting the equity share vary too much relative to this level. The variability in the equity share allowed by no-trade bands implies that the risk of the portfolio changes over time due to both time-varying market risk and time variation in the equity share.

No-trade bands may have additional expected return implications to the extent that the equity share drift within the no-trade band is related to variation in expected returns over time. In the next section, we explore the extent to which no-trade band rebalancing captures time-varying expected returns.

[15] Similar to inner bands, the net effect is somewhat ambiguous given that a slower trading speed implies a more persistent path both above and below the strategic equity share when rebalancing on the upside and downside respectively.

3. Timing Expected Returns with No-trade Bands

In this section, we explore the ability of no-trade band rules to capture variation in expected returns over time. There is an extensive literature documenting time-variation in aggregate expected equity returns.[16] For simplicity, we initially focus only on time-varying expected equity returns. There is also evidence for time variation in expected bond returns, in the volatilities of equity and bond returns, and in the correlation between equity and bond returns. At first, however, we do not include these dynamics in our simulations, in order to isolate the role of expected equity returns and to build intuition. Later in the section, we re-visit our simulation findings using a model that allows for all these additional features of asset returns, and show that our findings are unchanged.

As outlined in the previous section, the equity share drifts to varying extents depending on the rebalancing rule parameters. It is possible for the drift to have expected return implications: for example if the equity share tends to drift higher (lower) when expected returns also tend to be high (low), no-trade bands may be able to exploit variation in expected returns. We use simulation models to explore the relationship between equity share drift and expected returns under alternative rules, and ask whether alternative rules trade with or against time-varying expected returns.

Simulation Analysis: Evaluating Timing Ability

We use alternative simulations models with time-varying expected returns to generate returns and explore the timing ability of no-trade bands. We simulate monthly returns on a 70-30 percent equity and fixed income portfolio and evaluate rebalancing rule returns relative to a continuously rebalanced benchmark. This implies that we compare rebalancing strategies to a benchmark portfolio that has no link between its equity share and expected returns: by definition, a constant equity share cannot move either with or against expected returns.[17] We evaluate rebalancing strategy returns, r_t^N, based on the alpha from a regression on continuously rebalanced benchmark returns, r_t^R:

$$r_t^N = \alpha + \beta r_t^R + e_t. \tag{1}$$

The regression framework enables us to control for differences in average equity shares under alternative no-trade band rules. As outlined in Section 2, the equity share tends to drift upward within no-trade bands and is higher on average than with a continuously rebalanced benchmark. This effect is controlled for in the regression framework through a higher beta coefficient, such that a positive (negative) alpha coefficient reflects only the ability (inability) of no-trade rebalancing to exploit time-varying expected returns.[18]

[16] Cochrane (2011) provides an overview of the literature.

[17] As noted earlier, the equity share almost but not quite constant in the continuously rebalanced portfolio. The simulation results are unaffected if using a constant equity share instead.

[18] The alpha coefficient could also reflect the ability / inability to time volatility. Volatility is constant in the first two simulation models we consider. Later in this section we bootstrap a VAR model which captures time variation in the volatility of realised returns where this effect might

A positive (negative) alpha implies that the equity share co-moves positively (negatively) with expected returns under a given rebalancing rule.

A key advantage of using simulations is that expected returns can be observed directly, which would not be the case if using bootstrapped returns, where expected returns are not observed and would need to be estimated. For this reason, we choose to compare rules using simulations that are calibrated to include reasonable expected return dynamics. With simulations, we can also isolate different components of return dynamics to examine their effect on no-trade band timing performance. We proceed in this fashion, where we first simulate predictable equity returns with a single variable driving expected returns that behaves counter-cyclically relative to the market and generates a reversal effect in returns. We then extend this model to include a pro-cyclical component of expected returns which generates a momentum effect in returns.

No-trade Bands and Expected Returns: Reversal Effect

We initially analyse no-trade band performance with an equity return-generating process where expected returns are driven by a single state variable X_t:

$$r_{t+1} = r^f + X_t + \epsilon_{t+1}$$
$$X_{t+1} = \phi_0 + \phi_1 X_t + u_{t+1}.$$

r_t is the monthly equity return, r^f is the risk-free rate, X_t is the expected return and ϵ_{t+1} is the unexpected return. For simplicity, we assume no volatility in fixed income returns, and that shocks ϵ_t and u_{t+1} are homoscedastic and drawn from a multivariate normal distribution, though these assumptions do not affect our results and we later relax them. We calibrate the return-generating system in line with empirical estimates, where the parameters are calibrated to match a regression of monthly US stock returns on the dividend-price ratio over the period January 1950 to December 2016.[19, 20]

Table 2 summarises key parameters in the calibration of the system. We calibrate a small degree of predictability in returns, captured through a low R^2 from a regression of returns on the lagged predictor X_t. The process for expected returns is persistent, as captured through a high ϕ_1 coefficient. In addition, there is a strongly negative correlation between return and state variable shocks ϵ_t and u_{t+1}.

be relevant.

[19] The calibration of predictability, persistence and shock correlation is in line with the academic literature, for example Campbell (1991) and Pastor and Stambaugh (2009). Our aim is to capture the predictability of one of many possible predictors, as an illustration of the no-trade rebalancing implications. We explore additional predictors in a larger model later in this section. The risk-free rate and expected excess return are lower than has historically been the case, though our findings throughout are unchanged if using a higher level. Similarly, our results are unchanged if using international equity returns to guide our calibration.

[20] For the results that follow, a key moment in the calibration is the covariance between realised returns and expected returns, for which the shock correlation, predictability and persistence are important determinants. Our analysis is robust to moderate changes in the calibration where the covariance remains negative, consistent with the arguments in Pastor and Stambaugh (2009).

Table 2: Return system calibration

Risk-free Rate (annual)	r^f	1%
Expected Excess Equity Return (annual)	$E(r_t) - r^f$	3%
Equity Return Volatility (annual)	$\sigma(r_t)$	16%
Return Predictability (monthly)	R^2	1%
Expected Return Persistence	ϕ_1	0.99
Shock Correlation	$\rho(\epsilon_t, u_t)$	-0.90

A key parameter in the calibration is $\rho(\epsilon_t, u_t)$, the correlation between expected return and return shocks. This is strongly negative in the calibration, implying a form of counter-cyclical variation in expected returns. We refer to this as a 'reversal' component in expected returns, where negative (positive) returns are associated with positive (negative) shocks to expected returns. To the extent that no-trade band rebalancing captures this variation, the rule could be considered to exploit counter-cyclical variation in the expected returns.

It is useful to provide some context around the degree to which no-trade bands might be able to exploit time-varying expected returns, i.e. the potential magnitude of alphas of the rebalancing strategy relative to continuous rebalancing when successfully timing expected returns. To do this, we estimate alphas when varying the equity share directly in line with expected returns within our simulations.[21] We use two alternative simulation models in the section, the reversal model as outlined above, and an extended model that includes a momentum effect. Figure 10 shows the alphas from timing expected returns with the equity share for the two models. The left-most bars in Figure 10 show annualised alphas when varying the equity share between 66 and 74 percent. Varying the equity share by 4 percentage points either side of the strategic 70 percent equity share leads to alphas of 16 and 21 basis points per year (where the extended model generates higher alphas due to a higher predictive R^2 in the model).

Figure 10 also shows alphas from more aggressive equity share timing, varying the equity share by up to 50 percentage points. These improvements are more substantial, and illustrate that a significantly more aggressive change in the equity share would be needed to generate larger differences relative to a constant strategic equity share. The figure highlights an important aspect of the large literature on capturing time-varying expected returns. In portfolio choice applications, it is more common to discuss

[21] To obtain the equity share conditional on expected returns, we use a Z-score normalisation of the expected return and multiply it by 4 percentage points, which we add to the strategic 70 percent equity share. This means that when expected returns are one standard deviation above/below average, the equity share is set to 74/66 percent. The analogous methodology is used for varying the equity share by 10, 25 and 50 percentage points.

Figure 10: Timing the equity share - scope for alpha

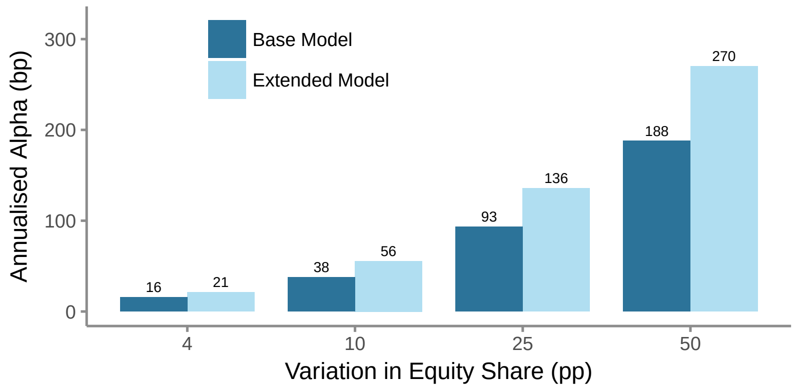

variation in optimal weights that an investor should rebalance to, which depend on the conditional level of expected returns.[22] For all the no-trade rebalancing rules we consider, the strategic equity share is held constant and in relatively narrow bands compared to those in Figure 10. The alphas from direct equity share timing should therefore be considered an upper bound to what no-trade rebalancing could achieve in terms of capturing time-varying expected returns within the simulation model.[23]

We proceed to simulate returns from the return-generating system described above, and apply no-trade rebalancing rules with different band widths (without inner bands and with one-step trade sizes) and also for a continuously rebalanced portfolio. We generate a long time series of returns, to ensure that alpha coefficients are precisely estimated, simulating 100,000 years of returns. Figure 11 shows the alphas of the no-trade rebalancing strategies from varying the band width between 1 and 6 percentage points.

Figure 11: Alpha of rebalancing strategies by no-trade band width

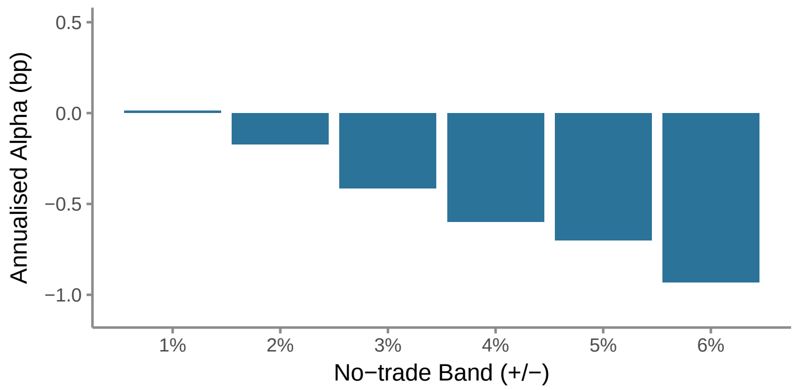

Figure 11 shows that, for any band width, no-trade bands generate a negative alpha, implying that the equity share co-moves negatively with expected returns i.e. timing in the wrong direction. Here, rebalancing trades tend to be in the correct direction given the level of expected returns: the equity share

[22] For a summary, see NBIM Discussion Note #2 (2012) "Return Predictability and Implications for Rebalancing".

[23] The upper bound may also be overstated given the in-sample normalisation of the observed expected return series. If implemented in 'real time', the annualised alphas would decrease, and even more so if a noisy proxy for expected returns were used, which would be the case in practice.

tends to be reduced (increased) when expected returns are low (high), and rebalancing is counter-cyclical in this respect. However, the equity share tends to drift in the opposite direction to expected returns, and is therefore lower than average when expected returns are high, and vice-versa. This is illustrated in a hypothetical example in Figure 12. In the shaded area, the equity share is increasing without triggering rebalancing, while expected returns are decreasing. For this period, relative to a continuous rebalanced benchmark, the equity share will be higher while expected returns are lower. This issue worsens as the band is widened, where alphas are increasingly negative, as wider bands allow the equity share to drift increasingly far in the wrong direction, and for longer periods of time.

Figure 12: Illustration of no-trade band rebalancing and expected return timing

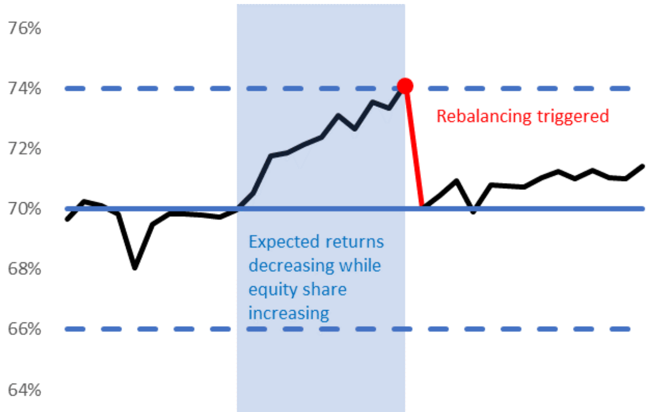

No-trade Bands and Expected Returns: Reversal vs. Momentum

Within the simulation analysis so far, a potentially important feature of returns, momentum, has not been included. The momentum effect - the tendency of strong or weak recent performance to continue - has been well documented[24] and we extend the simulation model to allow for a momentum-style effect in returns and analyse the implications for no-trade rebalancing. We do this by extending the simulation framework to include an additional variable, M_t, that drives expected returns:

$$\begin{aligned} r_{t+1} &= \mu M_t + (1-\mu)X_t + \epsilon_{t+1} \\ X_{t+1} &= \phi_0 + \phi_1 X_t + u_{t+1} \\ M_{t+1} &= \gamma_0 + \gamma_1 M_t + v_{t+1}. \end{aligned}$$

While the M_t variable does not capture momentum in terms of directly measuring past returns, a positive correlation between ϵ_t and v_t leads to the same effect. Many variables could therefore fit the 'momentum' description: variables that predict returns where the correlation between the predictive variable shocks and return shocks is greater than zero.

Based on the intuition for why no-trade bands performed poorly in the

[24] See Moskovitz, Ooi and Pedersen (2012), Neely, Rapach, Tu and Zhou (2014) and Hurst, Ooi and Pedersen (2017).

baseline model, we might expect that the addition of a momentum effect will improve their performance. The continuation of high or low returns implies that allowing the equity share to drift temporarily will allow the exploitation of momentum. It is less obvious, however, whether a momentum effect is sufficiently strong to offset the negative performance induced by the reversal effect in the previous simulations. This is partly determined through μ, which controls the relative weight on the momentum and reversal components of expected returns.

We calibrate the extended model to capture the balance between reversal and momentum effects. X_t behaves in the same way as in the previous simulation model, with the same persistence and shock correlations as earlier, and the mean and volatility of total equity returns is also kept the same as in the earlier calibration. We increase the monthly R^2 to 2 percent to reflect the inclusion of an additional predictor in the system, and calibrate the set of parameters $\mu, \rho(\epsilon_t, v_t), \rho(u_t, v_t)$ and γ_1 to capture the balance between momentum and reversal in realised returns. We do this by selecting a parameter set that attempts to match the balance between reversal and momentum captured through variance ratios over horizons up to five years.[25]

We chose the parameter set for the simulation that minimises squared errors relative to the empirical estimates of variance ratios, which gives values of $\mu = 0.70$, $\rho(\epsilon_t, v_t) = 0.30$, $\rho(u_t, v_t) = -0.20$ and $\gamma_1 = 0.20$. The correlation between the momentum state variable shock and return shocks is positive, in contrast to the reversal variable. Also, while the calibrated μ parameter implies a higher weight on the momentum than the reversal state variable, the momentum component in expected returns is much less persistent, reflected in a lower γ_1 parameter compared to ϕ_1. We simulate returns from the extended model and apply no-trade band rules to these returns. Figure 13 shows the alphas of the no-trade strategies when varying the band width between 1 and 6 percentage points, based on 100,000 years of returns from the extended model.

The alphas for the model with momentum contrast with the negative alphas of the baseline model. The alpha estimates are positive and initially increasing, but they decline as the no-trade bands widen further and fall quite sharply at 5 and 6 percentage point widths, turning negative and generating a hump-shaped profile by band width. Similar to the baseline reversal model, the magnitude of the alphas is small and far from the levels when timing the equity share directly shown earlier. In the extended model, the equity share is now able to drift with expected returns to an extent within the no-trade bands. The magnitudes of the alphas are relatively small, however, which reflects the relative inconsistency in capturing this relationship.

The hump-shaped alpha profile in Figure 13 is the result of the difference in

[25] We use US stock market returns from 1950 to 2017 to estimate variance ratios. Our chosen parameter set also generates return autocorrelations that are close to estimates from US data, for example in Koijen, Rodriguez and Sbuelz (2009). We estimate variance ratios for a large number of parameter combinations. The calibrated and empirical variance ratios are shown in Appendix A. When setting up the return-generating system, we normalise the total variance of X_t and M_t to be equal, and allow the contribution of each component to the variability in expected returns to be governed subsequently by μ and their shock correlations.

Figure 13: Alpha of no-trade band rebalancing strategies by band width – extended model

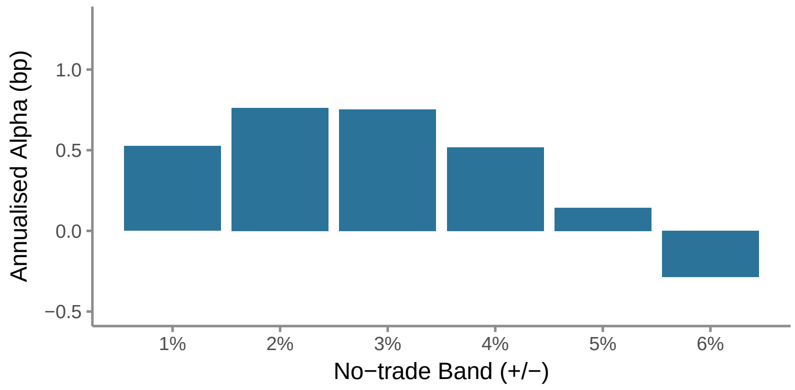

rebalancing horizons across the no-trade band widths. As discussed in Section 2, the frequency of rebalancing events decreases as the no-trade band is widened. The difference in persistence between the reversal and momentum components of expected returns means that the relative role of these two effects also varies by horizon. Figure 14 gives a depiction of these relative effects, showing the effect of a one-standard-deviation shock to the momentum and reversal components.

Figure 14: Persistence of expected return shocks - momentum vs reversal

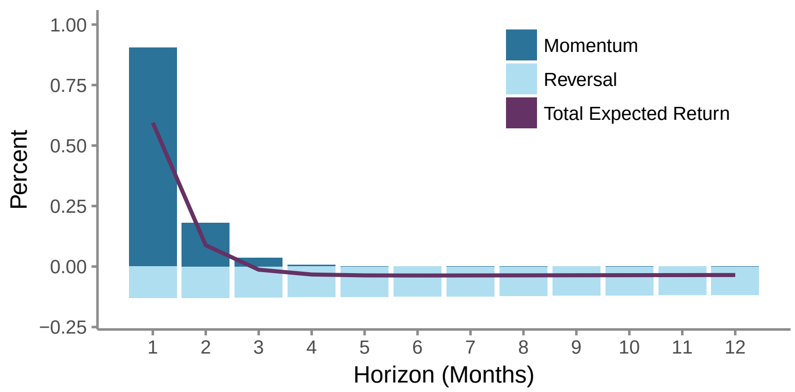

The figure illustrates how two coinciding shocks to the momentum and reversal components of expected returns affect total expected returns over longer periods. The total effect on expected returns, calculated by weighting the momentum and reversal shocks by μ and $(1-\mu)$ respectively, is expressed relative to its long-term average. The differing persistence of the two components implies that momentum tends to play a larger role at short horizons, whereas the reversal effect tends to be stronger at longer horizons. Over shorter horizons, therefore, the momentum effect dominates and pushes up total expected returns, but this effect dies out relatively quickly. The reversal shock pushes down at shorter horizons as well, but is not large enough to have a negative net effect on expected returns. As the momentum shock dies out, however, the reversal shock effect plays a larger role. As shown earlier in Figure 3, narrower bands have a shorter rebalancing horizon, implying that they have a higher exposure to momentum, while wider bands with longer horizons have higher exposure to the reversal shock effect. This horizon effect accounts for the hump shape in the extended simulation model alphas.

Alpha Distributions

We have so far outlined the importance of reversal and momentum effects for the average performance of no-trade rebalancing. This required simulation of a large number of observations in order to estimate precisely the average ability of no-trade bands to capture time-varying expected returns. It is useful also to consider the variability in this average performance for a shorter sample size, i.e. what an investor could expect to observe in any finite sample. Instead of a single long-sample simulation, we now simulate 10,000 samples of 50 years in length, using the combined reversal and momentum model, for alternative no-trade band widths. Figure 15 shows the 5th and 95th percentile intervals around the average alpha for each band width.

Figure 15: Distribution of 50-year alphas by no-trade band width – extended model

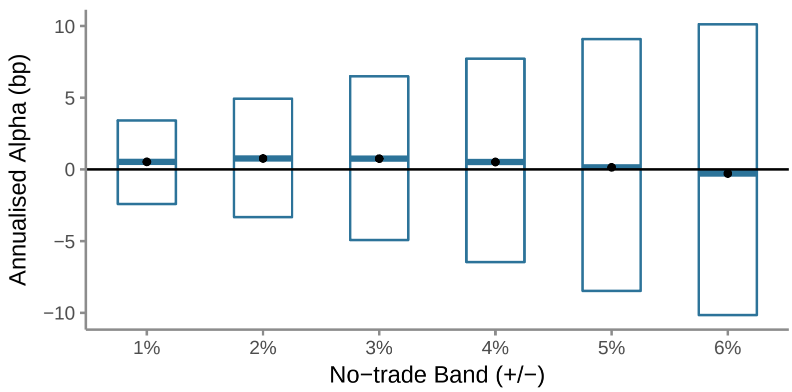

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

The alphas shown earlier in Figure 13 correspond to the mean of the alpha distributions in Figure 15. There is significant variability in the alphas, and many of the 50-year alphas are negative, for all band widths considered. This is in contrast to the effects of no-trade band rebalancing on the average equity share and tracking error shown in the previous section, which were consistently above zero. This means that while no-trade rebalancing rules might be set in such a way that they capture time-varying expected returns over the very long run, the rule will not reliably capture the positive (or negative) alphas shown earlier in Figure 13 in shorter samples. This is also the case for alternative inner band widths and trade sizes. The high degree of alpha variability implies that it is difficult to compare alternative no-trade band rules based on expected return considerations. Any differences in performance between rules that we have documented cannot be expected with a high level of confidence, even over a long period of time.

Simulation Robustness

In setting up the simulation analysis, we omitted fixed income returns from our analysis and assumed a constant co-variance matrix. To check robustness against the inclusion of fixed income returns, the possible presence of regimes, stochastic volatility or other features that are in the data but missing from our simulations, we estimate alphas using bootstrapped returns for equity and fixed income. We bootstrap a VAR system similar to Campbell, Giglio, Polk and Turley (2018) to capture persistent variation in expected returns. Bootstrapping a VAR model allows us to capture this persistence while still capturing features in the data that are not included in the simulation model.[26] The methodology is similar to the simulations, but the coefficients of the return-generating system are estimated rather than calibrated. In addition, the residuals from the estimated system are used to generate returns, as opposed to using draws from normal distributions. Details of the methodology are provided in Appendix B. Figure 16 shows the average alphas based on the bootstrapped model.[27] Similar to the simulation model with reversal and momentum components, the average alpha is positive and has a hump-shaped profile by band width. Also in line with our simulation findings, however, the estimates are very uncertain.

Figure 16: Distribution of 50-year alphas by no-trade band width – bootstrapped VAR model

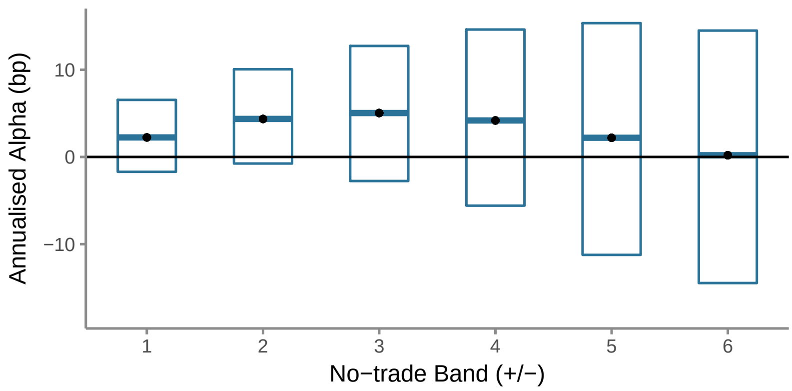

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

In addition to the VAR-based robustness, we also estimate alpha coefficients for alternative no-trade band strategies based on a simple backtest. We calculate portfolio returns of alternative no-trade band widths based on equity and fixed income returns over the sample period January 1973 to December 2017. Figure 17 shows the estimated alphas, with the intervals

[26] We choose not to use the same block-bootstrap methodology as described in Section 2. To capture long-term persistence of expected returns adequately, we would need to use a bootstrap block size that is multiple years in length, but we also want to avoid a large block size relative to the sample length.

[27] The average alphas from the bootstrapped VAR model are higher than for the extended model shown earlier. This may partly reflect the larger degree of momentum, captured through a higher short-horizon variance ratio, in the returns generated in the VAR model. In addition, it is possible that no-trade rebalancing helps to vary the equity share to take advantage of changes in volatility over time (see Moreira and Muir (2017)).

showing two-standard-error confidence intervals around the estimates, depicting their statistical significance.

Figure 17: Alpha by no-trade band width: historical backtest

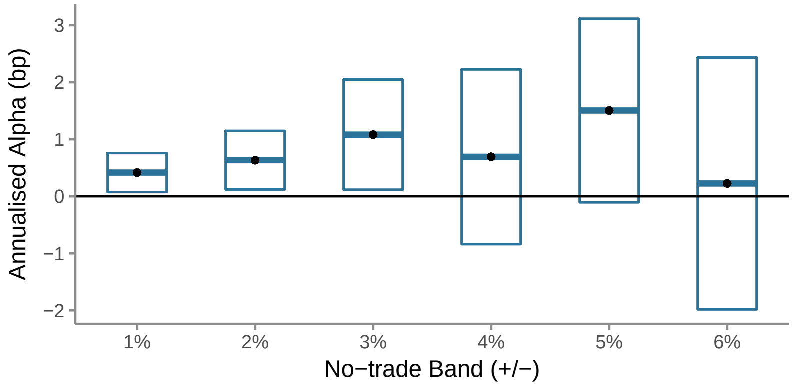

Note: Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution. Newey-West standard errors with lag length set to $T^{0.25}$

The estimates are positive for all no-trade band widths, and are of similarly small magnitudes to the simulation results. There also appears to be a hump-shaped profile to the alphas, with higher estimates for intermediate band widths. For all band widths greater than or equal to 4 percentage points, the alpha estimates are insignificantly different from zero. For band widths of 1 to 3 percentage points, the alpha estimates are significant.

In this section, we have shown that whether the equity share drifts positively or negatively with expected returns depends on the width of the no-trade band. These relationships are very uncertain, however, and for practical purposes no-trade rebalancing is likely an ineffective strategy for capturing time-varying expected returns. This implies that equity share deviations should be the main consideration against which to balance transaction cost savings from no-trade rebalancing rules. In the next section, we explore the trading cost reductions from no-trade band rules more closely.

4. Transaction Costs and No-trade Bands

In this section, we compare no-trade band rebalancing rules in terms of transaction costs and turnover. We specify and calibrate a simplified transaction cost model that allows us to illustrate the cost implications of different no-trade band widths, inner bands and trade sizes. The design of our transaction cost model is guided by the literature on trading costs in various portfolio choice applications. The earlier literature on transaction costs focuses on the implications of introducing proportional transaction costs into a standard portfolio choice problem with constant risk premia, for example in Constantinides (1979, 1986). These proportional transaction costs provide basic motivation for imposing a no-trade band around a target asset allocation.

Other studies such as Korajczyk and Sadka (2004), Frazzini, Israel and

Moskowitz (2015), Novy-Marx and Velikov (2015) and Ratcliffe, Miranda and Ang (2017) specify transaction costs with a measure of price impact incorporated into their proportional trading cost measure.[28] The specification of price impact in these papers tends to be a modified and often simplified version of price impact in the related market microstructure literature, for example Almgren (2003), Breen, Hodrick and Korajczyk (2002) and Engle, Ferstenberg and Russell (2012).

For the analysis in this section, we specify a transaction cost function that incorporates features from this literature. We inform our calibration using global data, to generate transaction costs that are more representative of the Government Pension Fund Global's global benchmarks, as opposed to the US data used so far in the note. First, we outline the transaction cost model and describe its calibration. We then proceed to illustrate the transaction cost and turnover implications of alternative rules. We focus on the specification of proportional trading costs.[29] It is important to note that the cost models we use do not incorporate or reflect any information or model calibration framework from the internal trading operations of NBIM. The model is deliberately simple and the calibrated transaction costs are conservative. Hence, the cost figures should be used for comparing rules only, not as point or level estimates of the full cost of rebalancing.

Transaction Cost Model

Naturally, transaction costs vary with the amount of trading generated by equity share rebalancing.[30] These costs consist of effective bid/ask spreads (when trading slowly at market orders) and price impact. Previous studies tend to show that the bulk of trading costs is due to the price impact of trading (see Knez and Ready (1996), Breen, Hodrick and Korajczyk (2002), Korajczyk and Sadka (2004)), which is not observable and needs to be estimated.[31] Price impact varies considerably not only across asset classes but also over time and depending on market conditions.

Our model of variable trading costs has two key drivers of price impact highlighted in the literature. The first driver is the size of the portfolio traded relative to market trading volumes, captured by the degree of market participation and measured as a fraction of average monthly volumes (AMV). The second driver is market conditions, summarised by asset class-specific 60-month rolling realised volatility (RV). We specify the following function for variable trading costs, denoted as $VC_{t,j}$, for asset class j at time t:

$$VC_{t,j} = \beta_{0,j} + \beta_{1,j}\frac{TradeSize_t}{AMV_{j,t}} + \beta_{2,j}\sqrt{\frac{TradeSize_t}{AMV_{j,t}}} + \beta_{3,j}RV_{j,t} + \beta_{4,j}\frac{TradeSize_t}{AMV_{j,t}}RV_{j,t}$$

[28] These studies are primarily concerned with evaluating the capacity of various risk factor strategies.

[29] Including calibrated fixed costs does not lead to material changes to the findings.

[30] The level of and variation in costs may be lower if jointly considering rebalancing-induced trading alongside fund inflows and outflows and cash flows from dividends, coupons etc.

[31] While the proportional trading costs (effective bid/ask spreads) are both easier to model and to estimate, they tend to be small in magnitude.

where $TradeSize_t$ is the monthly trading volume. The functional form is motivated by the empirical results presented in Frazzini, Israel and Moskowitz (2015, 2018) and Ratcliffe, Miranda and Ang (2017), where price impact is modelled as a concave function of the degree of market participation.[32]

Our calibration of variable trading costs is informed by estimates of effective bid/ask spreads using the methodology developed in Corwin and Schultz (2012). Specifically, we estimate value-weighted effective bid/ask spreads for global equities in two distinct periods.[33] Estimates from a more recent period (2015-2016), which is marked by a relatively low level of volatility, are around 60 basis points. Hence, the one-way transaction cost, which is half of the effective bid/ask spread, is close to 30 basis points. In line with this estimate, we set $\beta_{0,j}$ equal to 30 basis points.

Estimates from periods with high equity volatility, which we identify as the Lehman bankruptcy episode and the euro area sovereign debt crisis, average around 140 basis points. Hence, the variable trading costs reach 70 basis points in periods of extremely high volatility. The magnitudes of coefficients determining the sensitivity of trading costs to the degree of market participation, $\beta_{1,EQ}$ and $\beta_{2,EQ}$, are set broadly in line with the estimates in Frazzini, Israel and Moskowitz (2015, 2018).[34]

We follow a similar procedure to calibrate the variable cost function for fixed income. Informed by the estimates of effective bid/ask spreads for US Treasury and corporate bonds, we set $\beta_{0,FI}$ equal to 13 basis points and allow variable costs to reach 22 basis points in periods of high volatility. Figures 18 and 19 show the calibration of variable trading costs for equity and fixed income respectively.

We define rebalancing-induced trading volume as a fraction of the portfolio. To translate this into a degree of market participation, we need average monthly market volumes for global equities and fixed income. We estimate the average monthly trading volumes of the global equity and fixed income market using the data and the methodology described in NBIM Discussion Note #3 (2017).

Our estimate of monthly trading volumes for stocks included in the FTSE Global All Cap index is 5.4 trillion US dollars.[35] Our estimate of aggregate monthly trading volumes for bonds included in the Barclays Global Aggregate

[32] Our specification subsumes a number of notable special cases. Proportional trading costs imply setting all coefficients except for $\beta_{0,j}$ equal to zero. A quadratic trading cost function, used for instance in Garleanu and Pedersen (2013), implies a linear price impact function, which can be achieved by setting all coefficients, except $\beta_{0,j}$ and $\beta_{1,j}$, equal to zero. Our specification of variable trading cost function is simplified relative to models of price impact in the market microstructure literature. We work at the asset class rather than security-level, and our calibration of price impact is equivalent to a value-weighted average of security level price impacts. We also work with monthly as opposed to daily/transaction-level data, which implicitly assumes a constant participation rate over each month.

[33] We use the FTSE Global All Cap index as our equity universe.

[34] The estimates from the specification in Frazzini, Israel and Moskowitz (2015) do not translate directly into volatility coefficient parameters in our functional form. We set our parameters for equity and fixed income to approximate their estimated sensitivities to idiosyncratic volatility and VIX.

[35] Our estimate is slightly lower than the most recent estimate from the World Bank of 6.4 trillion US Dollars (in 2017). Part of this gap can be explained by the fact that FTSE does not include frontier markets and micro caps.

Figure 18: Calibrated variable cost function for equities

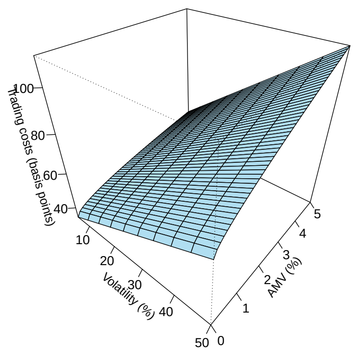

Figure 19: Calibrated variable cost function for fixed income

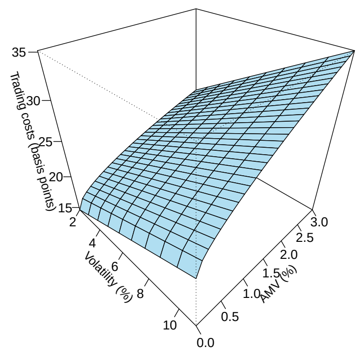

index is 15.5 trillion US dollars. In NBIM Discussion Note #3 (2017), we verify the estimates of trading volumes for US, German, UK and Japanese government bonds with aggregated data from debt management agencies and trade bodies and show that our bottom-up estimates closely match the aggregated external estimates of trading volumes.[36]

Transaction Costs of No-trade Band Rebalancing Rules

We proceed to illustrate transaction costs for alternative no-trade rebalancing rules. First, we use the transaction cost model to explore the implications of varying the strategic equity share while applying the same no-trade rebalancing rule. We illustrate the transaction cost implications of moving from a 60 to 70 per cent equity share, with a 4 percentage point no-trade band, in Figure 20. Increasing the equity share from 60 to 70 percent leads to less frequent rebalancing and lower transaction costs on average. Average annual transaction costs decrease by around 20 percent, equal to a reduction of roughly 1 basis point of the Fund's value.

Figure 20: Transaction costs from 60 vs. 70 per cent equity share

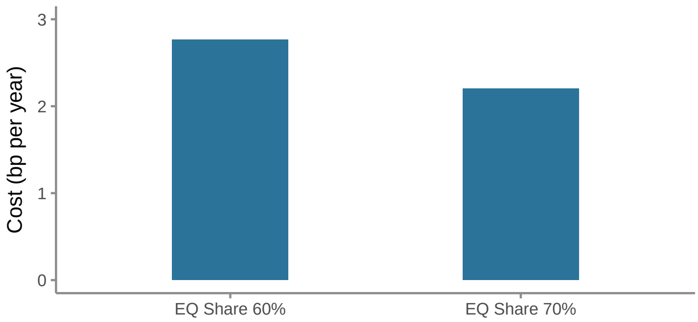

Next, we focus on the implications of varying the width of the no-trade band, without imposing an inner band, and rebalancing to the strategic equity share

[36]For comparison, SIFMA estimates monthly trading volumes of US Treasury securities at around 11 trillion US Dollars. See the table on page 37 in the SIFMA 2017 Fact Book.

in a single step. Figure 21 shows costs by no-trade band width, in basis points per year, and Figure 22 shows how the one-way portfolio turnover changes when varying the width of the no-trade band. The figures include the 5th and 95th percentiles of total costs and turnover across the 1,000 bootstrapped samples, to show the degree to which these values can vary. As a benchmark, in both figures we also include the case of a continuously rebalanced portfolio, denoted by a zero no-trade band, which provides the upper bound on total costs from rebalancing as frequently as possible. This rebalancing rule leads to transaction costs of around 4 basis points per year, on average generating turnover of around 6 percent of the portfolio per year.

Figure 21: Transaction cost distribution for alternative no-trade band widths

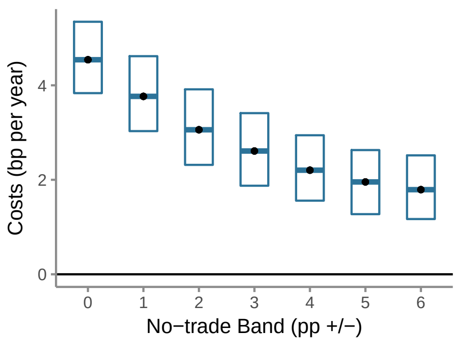

Note: No inner band, rebalancing in one step. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 22: One-way turnover distribution for alternative no-trade band widths

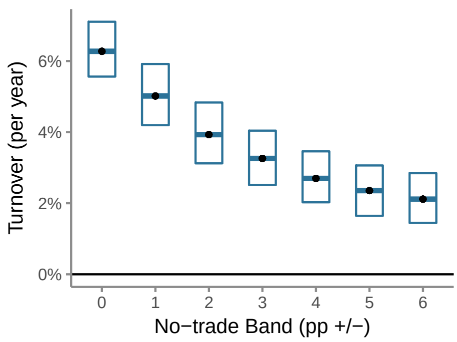

Note: No inner band, rebalancing in one step. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Imposing a no-trade band leads to less frequent rebalancing, lower turnover and lower costs. The incremental savings decrease as the no-trade bands are widened. Relative to continuous rebalancing, a cost reduction of around a 3 basis points per year can be achieved – roughly a 60 percent saving in costs - with a 6 percentage point no-trade band. The proportional reduction in turnover is similar, decreasing to around 2 percent per year for the 6 percentage point no-trade band.

Next, we explore how transaction costs vary with alternative inner band widths. Here, we fix the no-trade band width at 4 percentage points, and rebalance in a single step to the inner band. Transaction costs by inner band are shown in Figure 23, and turnover by inner band in Figure 24. Expected costs reduce by around 1 basis point by trading to the edge of the no-trade band compared to trading all the way back to the strategic equity share. In terms of turnover, trading to the no-trade band edge leads to a reduction from around 3 percent to around 1.5 percent per year.

Figure 23: Transaction cost distribution for alternative inner band widths

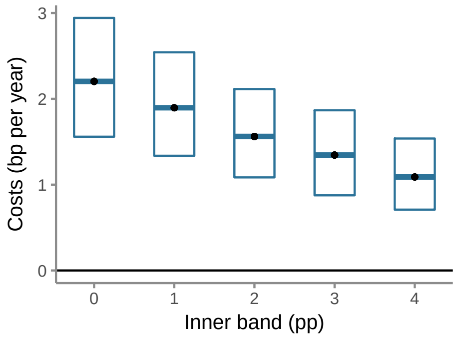

Note: 4pp no-trade band, rebalancing in one step. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 24: One-way turnover distribution for alternative inner band widths

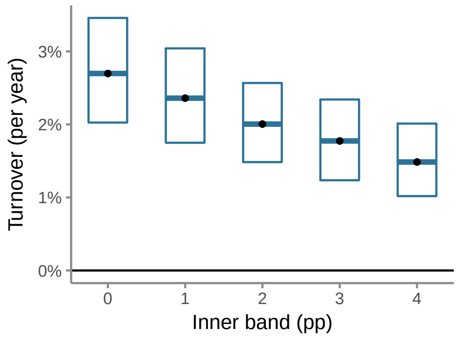

Note: 4pp no-trade band, rebalancing in one step. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Next, we illustrate the implication of varying trade sizes, assuming a 4 percentage point no-trade band width and no inner bands, shown in Figure 25. We estimate a reduction in costs of around a basis points when trading in 25 basis point monthly steps compared to a single step. In addition, average turnover per year reduces from around 2.5% trading in a single step to 1.5% when trading in 25 basis point steps.

Figure 25: Transaction cost distribution for alternative rebalancing trade sizes

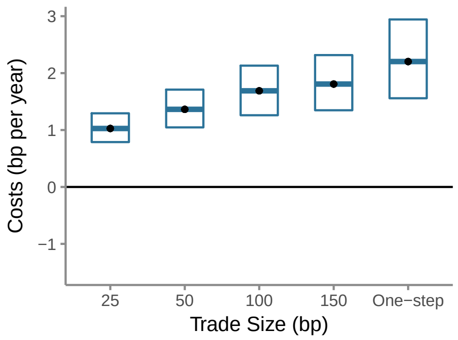

Note: 4pp no-trade band, no inner band. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

Figure 26: One-way turnover distribution for alternative rebalancing trade sizes

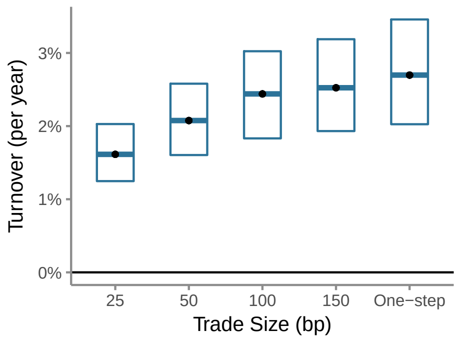

Note: 4pp no-trade band, no inner band. Dots show average of 1,000 50-year samples; intervals show 5th and 95th percentiles of distribution

It is worth noting that the incremental effects of varying trade size will to an extent depend on the width of the no-trade band. To illustrate this, Figure 27 shows how transaction costs vary for alternative trade sizes when also varying the width of the no-trade band, where the 4 percentage point no-trade band bars correspond to the values shown in Figure 25. The differences across trade sizes are larger when applying narrower no-trade bands.

In this section, we have set out a simplified transaction cost model that

Figure 27: Transaction costs by no-trade band width and trade size

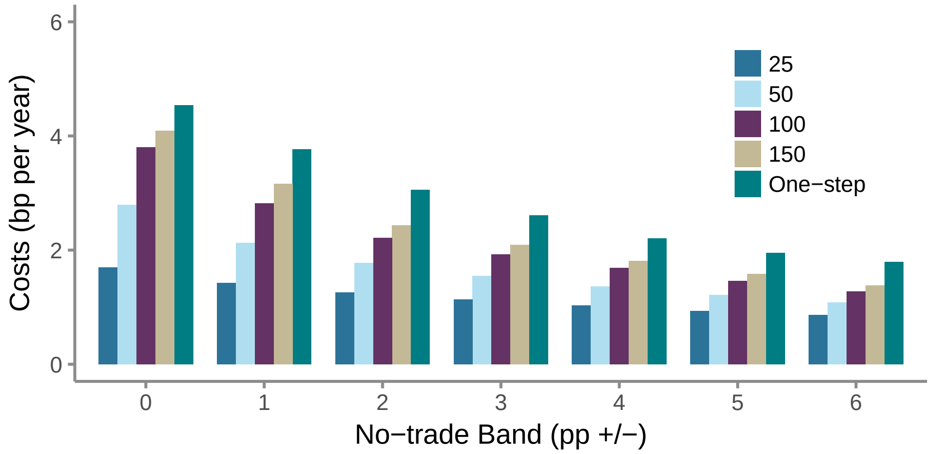

captures the price impact of trading in order to compare the transaction costs of alternative no-trade rebalancing rules. Naturally, transaction costs and turnover are always positive and a much more certain consequence of alternative rebalancing rules than expected returns. It should be noted, however, that the cost of a particular rebalancing event can vary substantially depending on market conditions, and the numbers presented in this section should not be interpreted as point estimates of transaction costs. Our calibration illustrates the degree to which wider no-trade bands, wider inner bands and trading more slowly all lead to cost reductions and lower turnover.

5. Concluding Remarks

In this note, we compare a range of no-trade band rebalancing rules. These rules can prevent rebalancing from occurring too often or too aggressively, though no-trade bands also allow the equity share to drift away from its strategic level. Using bootstrapped realised returns, simulations, and a transaction cost model, we explore the degree to which alternative no-trade rules differ in terms of equity share drift and costs.

We show that wider no-trade bands create a tendency for the equity share to drift to a higher level on average, and naturally also allow larger deviations in the equity share over time. The drift in the equity share can also relate to time-varying expected returns, though we show that no-trade bands do not reliably exploit changes in expected returns over time.

We show that there is scope for reducing transaction costs by setting wider bands, wider inner bands and trading less aggressively. Our analysis suggests that an investor's tolerance for a higher average equity share and higher equity share variability should be weighed against transaction cost savings. The uncertainty around the relationship between the equity share and expected returns under no-trade bands suggests that this should be a secondary consideration when comparing rules.

References

Almgren, R. F. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. Applied Mathematical Finance, 10(1), 1-18.

Ang, A., & Kjaer, K. (2011). Investing for the long run.

Arnott, R. D., & Lovell, R. M. (1993). Rebalancing: Why? When? How often?. The Journal of Investing, 2(1), 5-10.

Breen, W. J., Hodrick, L. S., & Korajczyk, R. A. (2002). Predicting equity liquidity. Management Science, 48(4), 470-483.

Brennan, M. J., Schwartz, E. S., & Lagnado, R. (1997). Strategic asset allocation. The Journal of Economic Dynamics and Control, 21(8-9), 1377-1403.

Campbell, J. Y., & Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. The Quarterly Journal of Economics, 114(2), 433-495.

Cochrane, J. H. (2011). Presidential address: Discount rates. The Journal of Finance, 66(4), 1047-1108.

Constantinides, G. M. (1979). Multiperiod consumption and investment behavior with convex transactions costs. Management Science, 25(11), 1127-1137.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. Journal of Political Economy, 94(4), 842-862.

Corwin, S. A., & Schultz, P. (2012). A simple way to estimate bid–ask spreads from daily high and low prices. The Journal of Finance, 67(2), 719-760.

Donohue, C., & Yip, K. (2003). Optimal portfolio rebalancing with transaction costs. The Journal of Portfolio Management, 29(4), 49-63.

Engle, R., Ferstenberg, R., & Russell, J. (2012). Measuring and modeling execution cost and risk. The Journal of Portfolio Management, 38(2), 14-28.

Frazzini, A., Israel, R., & Moskowitz, T. (2015). Trading costs of asset pricing anomalies.

Frazzini, A., Israel, R., & Moskowitz, T. (2018). Trading costs.

Hurst, B., Ooi, Y. H., & Pedersen, L. H. (2017). A century of evidence on trend-following investing.

Huss, J., & Maloney, T. (2017). Portfolio Rebalancing: Common Misconceptions. AQR research.

Ilmanen, A., & Maloney, T. (2015). Portfolio Rebalancing: Strategic Asset Allocation. AQR research.

Jaconetti, C. M., Kinniry, F. M., & Zilbering, Y. (2010). Best practices for portfolio rebalancing. Vanguard Research, July.

Knez, P. J., & Ready, M. J. (1996). Estimating the profits from trading strategies. The Review of Financial Studies, 9(4), 1121-1163.

Korajczyk, R. A., & Sadka, R. (2004). Are momentum profits robust to trading costs?. The Journal of Finance, 59(3), 1039-1082.

Kunsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. The Annals of Statistics, 1217-1241.

Leland, H. E. (1996). Optimal asset rebalancing in the presence of transactions costs. Institute of Business and Economic Research, University of California, Berkeley.

Masters, S. J. (2003). Rebalancing. The Journal of Portfolio Management, 29(3), 52-57.

Moreira, A., & Muir, T. (2017). Volatility-Managed Portfolios. The Journal of Finance, 72(4), 1611-1644.

Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. The Journal of Financial Economics, 104(2), 228-250.

NBIM Discussion Note #1 (2012). Time-varying Expected Returns and Investor Heterogeneity: Foundations for Rebalancing

NBIM Discussion Note #2 (2012). Return Predictability and Implications for Rebalancing

NBIM Discussion Note #3 (2012). Empirical Analysis of Rebalancing Strategies

NBIM Discussion Note #4 (2012). The History of Rebalancing of the Fund

NBIM Discussion Note #1 (2016). The Equity Risk Premium

NBIM Discussion Note #3 (2017). The Liquidity of a Diversified Portfolio

Neely, C. J., Rapach, D. E., Tu, J., & Zhou, G. (2014). Forecasting the equity risk premium: the role of technical indicators. Management Science, 60(7), 1772-1791.

Novy-Marx, R., & Velikov, M. (2015). A taxonomy of anomalies and their trading costs. The Review of Financial Studies, 29(1), 104-147.

Ratcliffe, R., Miranda, P., & Ang, A. (2017). Capacity of Smart Beta Strategies: A Transaction Cost Perspective.

Appendix A: Variance Ratio Calibration

Figure A-1: Variance ratios - empirical vs simulation calibration

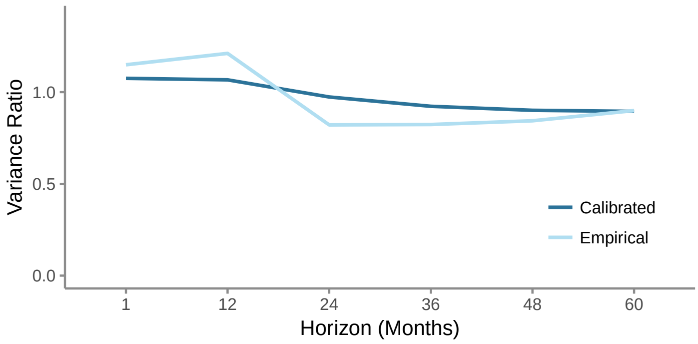

Appendix B: VAR Bootstrap Methodology

This section describes the bootstrapping methodology used in Section 3 analysing the time-varying expected return implications of no-trade bands. We bootstrap residuals of an estimated VAR model, since we want to allow for high persistence in expected returns that would require a block-bootstrap approach with a large block size relative to the size of the sample we draw from. We estimate a VAR model that shares similarities with the model in Campbell, Giglio, Polk and Turley (2018). In our state vector, we include the following variables:

- Equity Market Total Returns: Aggregate total returns on the S&P 500 equity index
- Fixed Income Total Returns: Total returns on the benchmark US 10-year government bond index
- Smoothed Price-Earnings Ratio: Real equity price divided by ten-year average of real earnings
- Term Spread: Yield difference between U.S. 10-year constant-maturity bonds and 3-month note
- Default Spread: Difference between yield on US dollar BBB- and AAA-rated corporate bonds
- Equity Volatility: 60-month rolling volatility of S&P 500 total returns
- Fixed Income Volatility: 60-month rolling volatility of US 10-year government bond index total returns

We estimate the model using a sample period from January 1950 to December 2016. We generate return series from the estimated VAR by block-bootstrapping the residual series in 12-month blocks.