They Would if They Could: Assessing the Bindingness of the Property Holding Constraints for REITs

Tobias Mühlhofer
Kelley School of Business
Indiana University *

October 2010

Abstract

This study examines the bindingness of the property holding constraints which REITs face on their portfolios (the *dealer rule*), and illustrates how these constraints hinder REITs from exploiting opportunities to time the property market and thereby generate profits by taking advantage of the market's predictability. I first simulate a set of filter-based market timing strategies, which I show outperform a buy-and-hold strategy out of sample over a wide range of levels of transaction costs. I further show that these strategies require a substantial number of transactions with holding periods that are shorter than the minimum required, and that imposing a four-year (or even the newer two-year) holding constraint significantly reduces the excess returns the strategy generates. I then analyze actual holding periods of properties in REIT portfolios and find that there seems to exist a large degree of demand for short property holding periods, and that thus the minimum holding constraint is quite binding. I proceed to model holding period as a function of current market performance, finding that REITs are more likely to hold a property beyond four years the higher the profit from the transaction, since this longer holding period is necessary in order for them to retain these capital gains. By contrast, this effect is insignificant for UPREITs, which, I argue, are not as affected by the constraint, implying that those firms which are allowed to sell properties before the required minimum time to realize capital gains often do so. The results thus give strong support to my hypothesis that REITs are severely bound by this holding constraint, and further illustrate why REITs are unable to realize short-term property appreciation profits, as shown in Mühlhofer (2008).

*Part of this research was completed while the author was visiting The University of Texas at Austin. I would like to thank John Board, Wally Boudry, Paul Cheshire, Greg Connor, Ron Donohue, Jeff Fisher, John Glascock, Kevin Habicht, Jay Hartzell, Christian Hilber, Jennifer Juergens, Chris Mayer, Francois Ortalo Magne, Henry Overman, and Sheridan Titman, as well as seminar participants at the 2008 ASSA Meetings and at Indiana University, as well as two anonymous reviewers at *Real Estate Economics* for their input. Steve Kleiman provided valuable research assistance. All other errors and omissions are my own. **E-mail: tmuhlhof@indiana.edu** Contact: Finance Department / Kelley School of Busines, Indiana University / 1309 E 10th St. / Bloomington, IN 47405 / Website: **http://tobias.muhlhofer.com**, where this paper is available for download.

1 Introduction

Equity Real Estate Investment Trusts (REITs) present a liquid and relatively cheap opportunity for smaller investors to add property-like cashflows to multi-asset portfolios. In fact, REITs own property portfolios of various sizes, manage these properties, and derive cashflows from them, and thus, fundamentally, these investment vehicles should present a more easily accessible alternative to holding a direct property portfolio. However, there exists an important hurdle placed on a REIT's portfolio management by the legislative structure under which REITs exist in the United States. This hurdle prevents REIT portfolios from closely resembling direct property portfolios, and from fully generating cashflows that are equivalent to those associated with direct property investment. In order for a property firm to be exempt from corporate income tax, and thus be classified as a REIT, among other things, its portfolio strategy is mandated to be *passive*.[1] This means that REITs face selling constraints for the properties in their portfolios: specifically, until 2008, REITs needed to hold each property for at least four years, and, once they sold, they could only sell 10% of their asset base within a given tax year. On July 30, 2008, the minimum holding period was decreased to two years, with the rest of the rule unchanged, as part of the *Housing and Economic Recovery Act of 2008*. If a REIT sells a property before the minimum holding period or sells more than 10% of its asset base, a *prohibited transaction* is said to have occurred and the REIT needs to pay a 100% gains tax on this transaction, that is, surrender all profits from the sale to the government.[2] This renders REIT managers less able to time the property market in order to exploit its predictability, and eliminates a vital component of a REIT's property cashflows, as is shown in Mühlhofer (2008).

The aim of this paper is to examine the bindingness of the minimum holding constraint, and to illustrate the way in which REITs forego possible market timing profits due to it. First, I illustrate the mechanics of how the holding period constraint affects a REIT's market timing ability, in a controlled environment, by devising and simulating a set of filter-based trading strategies with which I time the real estate market. The strategies are based on moving averages and trading bands constructed around these. Buy and sell signals are generated when the price process breaks through the moving average, or into the range delimited by the trading bands from below (buy signal) or above (sell signal).

I show that, out of sample, these strategies significantly outperform a simple buy-and-hold strategy

[1] This is known as the *dealer rule*, as the rule essentially prevents a REIT from being a dealer in property.

[2] More specifically, the four-year (or two-year), ten-percent constraint constitutes a *safe harbor rule* in this respect, meaning that any transaction that fulfills these criteria will never be classified as *prohibited*. Should a REIT wish to sell a property earlier, the firm needs to prove by some other means that the property was not 'held primarily for resale'. A firm's ability to do this will be linked to the specific circumstances of the sale; thus this part of the constraint would simply add idiosyncratic noise to this study. Furthermore, until 1999, a REIT could only derive 30% or less of its taxable income from the net gains on the sale of property, regardless of whether it was deemed to be held for resale.

under various levels of transaction costs from zero to ten percent round-trip costs; hence, these constitute profitable and, I argue, feasible market timing strategies. I then analyze the holding periods that these strategies require, and find that significant amounts of sales need to be made before the four-year or even the two-year minimum holding period. I then simulate these strategies in an environment constrained by a four-year and a two-year minimum holding period respectively, in which I ignore any sell signal that the strategy generates on a property that has been held for less than the minimum, and assess to what extent these strategies underperform their counterpart in an unconstrained environment. I find that the unconstrained trading-band strategy outperforms the buy-and-hold strategy by almost twice as much as the constrained strategy with a four-year constraint, and that this difference is statistically significant at all levels of transaction costs. With a two-year constraint, the performance shortfall of the constrained trading-band strategy is about a third, and again this difference is statistically significant at all levels of transaction costs. If timing strategies based on raw moving averages rather than trading bands are used, the results from this experiment are unaltered, with both the four-year as well as the two-year constrained strategies showing a statistically significant outperformance shortfall over the unconstrained strategy at all transaction cost levels.

This set of simulated trading results illustrates in a controlled environment some of the many ways in which a manager could possibly time the market if he were allowed to trade freely, and the loss of profits that appears once the trading constraint is imposed. While the relaxation of the minimum holding period to two years helps the situation to some extent, even with a two-year constraint there are still significant losses to profitability. It is important to note that this strategy need not necessarily be the only strategy or even the optimal strategy for a REIT manager to pursue. Merely by showing the existence of a feasible timing strategy whose ability to generate abnormal profits is significantly reduced in the presence of a minimum holding constraint, I show that this constraint reduces the opportunity set for making abnormal profits, and at the same time offer general insights about the mechanics of how this happens.

In the general finance literature, the subject of technical trading rules, and more specifically filter rules, is treated primarily as part of the literature on weak-form market efficiency and a multitude of such studies exists, for example Alexander (1961, 1964) – which seems to have pioneered the concept of filter rules, that is trading rules which *filter* small price movements to highlight larger reversions in trend – as well as Fama (1965) and Fama and Blume (1966), or more recently Blume, Easley and O'Hara (1994), Brock, Lakonishok and LeBaron (1992), LeBaron (1996). In the real estate literature, as well, many studies have documented the predictability of property markets, mainly to make a statement about their efficiency, for example Liu and

Mei (1992, 1994), Barkham and Geltner (1995), Case and Shiller (1990), Case and Quigley (1991), Gyourko and Keim (1992), Cooper, Downs and Patterson (1999) who use filter rules to construct a contrarian trading strategy in the REIT market, and especially Geltner and Mei (1995) who use technical trading strategies to illustrate that market timing profits can be made in the property market. While all these studies are loosely related to this one in that they treat technical trading strategies, they do this to make an argument about market efficiency, while I consider these in order to illustrate how REITs are hindered by the property trading constraints in generating the abnormal profits that stem from the property market's predictability. Thus, since these trading constraints and their effects seem to have been overlooked in previous literature, there is actually no direct parallel to this study.

I then analyze empirical distributions of property holding periods in a REIT's portfolio, by examining REIT property transactions data in order to directly assess the bindingness of the holding constraint. I find that there seems to be considerable demand for short property holding periods by REITs. I conduct a natural experiment in this respect, by contrasting the subsample of UPREITs (Umbrella-Partnership REITs) with that of non-UPREITs. In an UPREIT, the REIT holds shares in a limited partnership, known as the operating partnership, which then in turn holds the portfolio of properties. The partnership structure enables an UPREIT to efficiently acquire properties through Section-1031 like-kind exchanges, by allowing previous owners (known as contributing partners) to sell the property to the UPREIT by exchanging it for shares in the operating partnership's overall property portfolio. The advantage of undertaking such a transaction, is that (like all 1031 exchanges) this does not trigger a taxable event for the contributing partner, who can subsequently exchange his or her portfolio shares for cash or REIT shares, when this is most advantageous for tax purposes. This has proven to be such an efficient way of assembling a property portfolio, that a majority of REITs have adopted the UPREIT structure nowadays.

In the context of this study, an UPREIT's primary modus operandi of acquiring properties through 1031 exchanges makes this structure a fitting natural experiment. The reason for this is that because the IRS does not view a 1031 exchange as a sale, the property's basis flows from the seller (the contributing partner) to the UPREIT. This means, among other things, that the UPREIT inherits the contributing partner's holding period, from the time the property was purchased by the contributing partner. Since the private entities that ordinarily constitute contributing partners tend to hold their properties for a long time, in part due to the high transactions costs they face, this means that on most of its properties the UPREIT will have likely met the minimum holding period for the safe harbor rule within a very short time of its purchase of the property. Thus, through this mechanism, an UPREIT is virtually freed from the minimum holding periods and can easily transact under the safe harbor rule.[3] [4]

Besides this important distinction, UPREITs and non-UPREITs are operationally very similar to each other in their activities. In line with the hypothesis of a binding holding constraint, I find a large number of holding periods for UPREITs that are shorter than four years,[5] and statistically significantly more frequent trades by UPREITs than by regular REITs. This implies that those firms that can trade more often, do so. Finally, I test the specific effect that local market performance has on the conditional decision to hold a property beyond four years, given that the property was held six years or less: because of the asymmetric nature of the sanctions for violating the holding constraint (a 100% gains tax), the higher a profit a REIT makes on a particular transaction, the more important it is that a REIT hold a property beyond the minimum time in order to retain this profit.[6] As a further illustration of this point, in a falling market where the REIT sells the property at a loss, the holding constraint would have no effect, since no capital gains are made and then lost as gains tax in a prohibited transaction. On the other hand, if the constraint were not perceived as binding (for example if four years were a shorter holding time than any REIT would normally employ), market performance should have no effect on the decision to hold a property beyond four years, as in a market with irregular return patterns (whether these be predictable or not), one should not expect that a larger average annualized return be associated with any particular holding time. Throughout this part of the analysis I also examine, once again, the important distinction in this respect between regular REITs and UPREITs, and find that regular REITs are significantly more likely to hold a property beyond four years, the higher the local market return (and thus the more of a profit they are making), while UPREITs are significantly less bound by the four-year cutoff. An expanded model which accounts for both firm size and UPREIT status, in order to disentangle these two collinear effects, strengthens these results, showing UPREITs to be completely indifferent to this constraint. This result also holds when I account for possible systematic effects associated with different property types.

[3] It is also possible for a non-UPREIT to conduct a 1031 exchange. This would happen when the REIT disposes of a property. The exchange is done by selling the property and putting the proceeds into a QI (a Qualified Intermediary). Within 14 days of the sale of the old property, a new property to be acquired must be identified, and within six months, the new property must be purchased. If all these conditions are met, such a transaction is also viewed as a 1031 exchange, and thus will not be considered a prohibited transaction, regardless of previous holding period, because 1031 exchanges by definition are made on investment property, rather than dealer property. However, it seems apparent that despite this ability by non-UPREITs, an UPREIT still has a competitive advantage in this respect, as these rules are quite binding, in comparison to those governing the situation for an UPREIT, which can effectively sell the property outright whenever it deems necessary, even without an immediate replacement. This is especially useful when downsizing to cash out before a falling commercial property market, which a regular REIT cannot do through a 1031 exchange.

[4] This information is derived from an interview with Kevin Habicht, Chief Financial Officer of National Retail Properties, Inc. (NNN), a major US REIT, conducted on 9/2/2009.

[5] My data sample for this portion of the study ends in 2004, so I cannot conduct any tests with regard to the reduction in the minimum holding period.

[6] Even with respect to the safe-harbor nature of the minimum holding constraint this intuition holds: the higher a profit a firm is making from a particular transaction, the more difficult it would likely become to prove that the property was not held ‘primarily for resale’, in the specific circumstances.

While some literature which analyzes the performance of mutual fund managers in terms of market timing ability indirectly analyzes holding periods of securities in a fund's portfolio, the only other study I am aware of which analyzes trading behavior by REITs in the direct property market is Crane and Hartzell (2007), who do this in order to test for a disposition effect. However, once again, my study analyzes holding periods in order to assess the effect of the trading constraints on REIT market timing ability, and thus there should be no direct parallel to this part of the study in another line of research, as these trading constraints have been largely overlooked.

The rest of this study proceeds as follows. Section 2 presents the analysis of the filter-based timing strategy; section 3 presents the analysis of actual holding periods; section 4 concludes.

2 A Filter-Based Trading Strategy

2.1 Implementing Filter Rules in Property Portfolios

2.1.1 A Moving-Average-Based Trading-Band Filter Rule

The idea of *market timing* consists of trying to *buy low* and *sell high*, that is, in a cyclical market, to buy at a trough and sell at a peak. One of the main endeavors in devising such a strategy is therefore to find a way of observing larger-scale market turnarounds, in the midst of small-scale volatility. While many different strategies have been devised to try to do this (and some have done this with more success than others), in this study I choose to employ a timing strategy based on moving average prices, as these present a higher level of sophistication than simple return-filter based strategies, yet retain a high degree of intuitiveness.[7]

Suppose we have an asset whose price we have observed today (at time t) and at several different equally spaced intervals in the past, over a sufficiently long time window. The K-period moving average price at time t is simply defined as

$$map_{K,t} = \frac{\sum_{\tau=0}^{K-1} p_{t-\tau}}{K-1} \qquad (1)$$

Here, p_τ is the asset's price at time τ. Looking at this expression, it should be clear that in a rising market (where $p_{t-1} < p_t \; \forall \; t$), $map_{K,t} < p_t \; \forall \; K > 0$. Conversely, in a falling market (where $p_{t-1} > p_t \; \forall \; t$), $map_{K,t} > p_t \; \forall \; K > 0$. Thus when the market changes from rising to falling, or vice versa, p_t will *cut*

[7]Information on MA-based trading strategies and trading bands can be found in most technical analysis textbooks, for example Murphy (1999).

through $map_{K,t}$. Hence, this presents a possible way to observe market turnarounds, simply by computing the moving average price with every new observation and observing when $p_t = map_{K,t}$. This often occurs between two points, so that $p_{t-1} < map_{K,t-1}$ and $p_t > map_{K,t}$, identifies a trough which constitutes a *buy* signal, while $p_{t-1} > map_{K,t-1}$ and $p_t < map_{K,t}$ identifies a peak and thus a *sell* signal.

Note further that, with a monotonic price process, the size of K will determine the distance between the price and its moving average, or

$$K \quad \propto \quad |map_{K,t} - p_t| \tag{2}$$

That means the size of K can be adjusted in order to *filter* noise and trigger signals only at larger turnarounds, since a price that was monotonically increasing (decreasing) and changes to decreasing (increasing) needs to show longer consistent decreasing (increasing) tendencies, before covering the distance between the price and the moving average, if that distance is larger. Hence, the size of K becomes a filter level.

However, this implicitly creates a problem: while waiting until the asset price cuts the moving average will filter out noise reversals and signal only large reversals (depending on the size of K), it should also be clear that the actual market turnaround occurs some time before the point at which $p_t = map_{K,t}$ and that, while the price cutting the moving average is useful in filtering small shocks from large cycles, we actually miss a peak or trough if we wait until the crossover point, and thus we have sold somewhere below the peak and bought somewhere above the trough. Figure 1 illustrates this phenomenon.

One way to improve upon this is through the use of trading bands. These consist of threshold values that track alongside the moving average at a certain distance, one above and one below, as follows:

$$b_{+,t} \quad = \quad map_{K,t} + s \tag{3}$$

$$b_{-,t} \quad = \quad map_{K,t} - s \tag{4}$$

Traditionally, the distance variable s is a predetermined fixed value, and trading bands are laid out in such a way as to contain most small price movements. A larger price movement, however, will touch one of the bands, and this triggers a signal. Specifically, $p_t \leq b_{-,t}$ constitutes a *buy* signal, while $p_t \geq b_{+,t}$ constitutes a *sell* signal. It is readily apparent that if such a trading band somewhere above (below) the moving average is used to trigger a *sell* (*buy*) signal, the problem of missing the turnaround described above is somewhat alleviated. Since the two bands track a certain distance from the moving average, the filtering properties of the moving average are still retained, since a strong enough shift in the market needs to occur for the price to not just cross the moving average, but travel a certain distance from it and then return.

Bollinger (2001) makes another innovation on this trading rule, in that he devises a way of determining the value of the bandwidth parameter s. Instead of predetermining a value for s, Bollinger uses a multiple of the K-period moving standard deviation, defined as

$$s_{K,t} = \left[\frac{\sum_{\tau=0}^{K-1} \left(map_{K,t} - p_{t-\tau}\right)^2}{K-1} \right]^{\frac{1}{2}} \tag{5}$$

making the trading bands

$$B_{+,t} = map_{K,t} + \alpha s_{K,t} \tag{6}$$

$$B_{-,t} = map_{K,t} - \alpha s_{K,t} \tag{7}$$

The value of α should be determined by the trader according to the market environment.

Generally this rule uses fairly large values of α, like 1.5 or 2, with a rolling window of 20 periods over which moving averages and standard deviations are computed, to produce wide bands. I find it more appropriate for this study to use fairly narrow bands (I use a value for α of $\frac{1}{3}$), and I therefore slightly alter the trading rule. Traditionally, the price process would *touch* a band from the inside in order to trigger a timing signal, and right away retreat back inside in most cases, due to the large bandwidth used. I alter this picture in such a way that in this case the bands are more narrowly constructed, so that in a strong market upturn (downturn), $p_t > B_{+,t}$ ($p_t < B_{-,t}$). That is, in major cyclical market movement, the price is *outside* the respective trading band and breaks back inside very shortly after a turnaround: this happens, of course, some time before breaking through the moving average. Figure 2 illustrates this concept.

A buy signal is thus generated by $p_{t-1} < B_{-,t-1}$ and $p_t > B_{-,t}$ and a sell signal is generated by $p_{t-1} > B_{+,t-1}$ and $p_t < B_{+,t}$. As is apparent from Figure 2, this strategy better times turnarounds than a pure moving-average strategy. Notice also how the bandwidth changes with volatility, as is the nature of Bollinger bands.

The decision to undertake alterations to the traditional way in which this trading rule is applied stems from the low frequency and low volatility of the data for the real estate market: while this rule is usually applied to daily stock market data and a 20-day moving average, the property market data is of quarterly frequency,[8] and I use one-year (four-data-point) moving averages. Given such an environment, this modified rule seems to perform better than the traditional version. I empirically test the performance of this trading-

[8]See Section 2.2, for more information on the data set.

band strategy, both in an environment without holding period constraints, and in an environment constrained by two-year and four-year minimum holding periods. For robustness, I conduct the same test using a strategy that generates signals purely based on moving averages, as outlined at the beginning of this section.

As stated earlier, I make no claim that these be the only possible strategies with which to time the property market, nor that they be the optimal ones. The purpose of this section is simply to propose a profitable and feasible market timing strategy,[9] and to demonstrate that the minimum holding constraints significantly reduce the ability of these strategies to generate abnormal profits. This should suffice, in order to argue that the minimum holding constraints reduce the overall opportunity set available to a REIT to generate abnormal returns. At the same time, the ability to examine such a strategy in a controlled environment should offer general insights on how the holding constraint reduces the profitability of market timing strategies. While I have no way of testing whether these strategies are actually applied by REITs, the empirical distribution of holding periods by REITs shown in Table 5 does resemble that generated by these strategies, which suggests the high value of the introspection into the mechanics of how holding constraints hinder market timing profits, that can be gained from this exercise. I thus apply these filter rules to a variety of different property portfolios, testing its performance in various transaction cost environments.

2.1.2 Practical Considerations

The direct property market is characterized by poor pricing information, long transaction times, and high transaction costs, all of which may pose problems to implementing a technical trading strategy. The evidence on poor pricing information stems mainly from literature about the validity of appraisal-based data in assessing market performance.[10] In this literature, the assumption being made (at least implicitly) is generally that agents who transact in the market have some pricing information based on which they transact, and that there is an incomplete flow of information from agents in the market to appraisers, leading to appraisal-anchoring and stale-appraisal biases in an appraisal-based index. Appraisal-based indices have thus been modified through *de-smoothing* procedures such as the one used in this study in order to recreate, conceptually, a data series that is equivalent to a transactions-based index.[11] Thus, these procedures are used in order to give the academic researcher outside the market a way to infer the level of pricing information that agents within the market had at a particular time. This framework suffices for the implementability of the strategies described here, since it is agents within the market who execute these strategies, and not

[9] See the next subsection for more on the strategy's feasibility.
[10] See Clayton, Geltner and Hamilton (2001), for example.
[11] A transactions-based index is, of course, not feasible in the commercial property market at a disaggregate level, due to this market's poor liquidity.

appraisers. They will thus not implement these trades based on a publicly available index series, but upon their own price series generated with their assumed superior knowledge.

For the issue of long transaction times, I simply make the assumption that is generally made in the de-smoothing literature.[12] This assumption is that the sale price of a property is generally locked in fairly early in the transaction process, which suffices for the feasibility of the market timing strategy. A manager reacts to a signal right away, and the price is locked in early: the fact that money only changes hands several months later becomes irrelevant at this point, because the property will be sold at the price for which the signal was obtained.

Transaction costs in the property market are usually considered high; a figure of five to ten percent round-trip transaction costs is considered common. Transaction costs will, of course, be related to the size of the entity undertaking the transaction. It is conceivable, for example, that a large REIT will have an in-house legal team which draws up all sales contracts, a team of appraisers who value properties, and so forth. For such an entity, these types of expenses, which are generally considered part of the transaction costs involved in dealing with direct property, would probably fall more generally under management expenses, as these become fixed costs if the services are internalized. The minimum threshold that is not dependent on scale, is, of course, property transfer tax, which ranks from zero to about two percent, dependent on the state and city. For this reason, I run these tests with different levels of transaction cost, ranging from free transactions (which is probably never realistic, but interesting for calibration purposes) to 10% round-trip costs.[13]

A further issue that should be mentioned here is market liquidity.[14] Implicitly, I am assuming throughout this section that there will always be investors willing to take the opposite position of the manager wanting to do a trade dictated by this timing strategy. This assumption should be reasonable, in that a manager pursuing such a strategy will need to be an informed investor, in order to have access to the true pricing information discussed above; as is the case in traditional trading models, here too there should be uninformed investors in the market willing to make such a trade. Specifically, while the informed investors will be property specialists such as REIT managers, the uninformed investors could be managers of multi-asset portfolios such as pension funds, who are making a long-term, income-oriented investment and are thus not timing the

[12]See, for example Crosby and McAllister (2005)

[13]The fact that I apply transaction costs as a negative percentage return should alleviate concerns that transaction costs vary proportionally with property size, as the actual Dollar amount of the transaction cost will thus be transaction-size adjusted.

[14]Kawaguchi, Sa-Aadu and Shilling (2007), for example, show that liquidity risk is priced in property markets and suggest that this effect might be asymmetric in that supply shocks can have different effects on market liquidity than demand shocks, although their results seem to speak against such asymmetries.

market.[15] With the above assumptions this strategy should thus be feasible, and if its ability to generate abnormal profits is reduced by the four-year holding constraint, it should follow that the constraint reduces a REIT's overall opportunity set to profitably time predictable property markets.

2.2 Data and Methodology

I test the filter-based investment strategy on data from the National Property Index (NPI) compiled by the National Council of Real Estate Investment Fiduciaries (NCREIF). The entire NPI panel data set contains many levels of disaggregation, by property sector, region, state, Metropolitan Statistical Area (MSA), and all interactions of these classifications. These 1041 subindices thus proxy for property portfolios of various degrees of diversification, any combination of which could conceivably be held by a REIT. The data is quarterly and goes at its longest from the first quarter of 1978 until the second quarter of 2003. Certain subindices start later.

To overcome the problems of stale appraisals and appraisal anchoring inherent in this data, I use the de-smoothing methodology of Cho, Kawaguchi and Shilling (2003).[16]

This desmoothing model can be summarized in the following five equations:

$$
\begin{aligned}
P_t^* &= b_1 P_{t-1}^* + b_2 P_{t-4}^* + w_0 P_t && (8)\\
r_t^* &= b_1 r_{t-1}^* + b_2 r_{t-4}^* + w_0 r_t && (9)\\
r_t &= \alpha + \rho r_{t-1} + \epsilon_t, &&\\
&\quad \text{with } E(\epsilon_t) = 0,\ \sigma^2(\epsilon) = \sigma^2,\ |\rho| < 1, \text{ and } \alpha \neq 0 && (10)\\
r_t^* - \rho r_{t-1}^* &= \alpha w_0 + b_1 \left(r_{t-1}^* - \rho r_{t-2}^*\right) + b_2 \left(r_{t-4}^* - \rho r_{t-5}^*\right) + \epsilon_t', &&\\
&\quad \text{where } \epsilon_t' = w_0 \epsilon_t && (11)\\
w_0 &= 1 - b_1 - b_2 && (12)
\end{aligned}
$$

where P_t is the logarithm of the underlying (and unobservable) true property price index, P_t^* is the logarithm of the smoothed property price index, $r_t = P_t - P_{t-1}$, and $r_t^* = P_t^* - P_{t-1}^*$. Equations 8 and 9 describe

[15] Should this liquidity not exist at all, the strategy automatically becomes buy-and-hold. For any liquidity level in between perfect liquidity and perfect illiquidty, the strategy should generate at least part of its outperformance over the buy-and-hold strategy; this outperformance, in turn will be hindered by a minimum holding period, and so the constraint can only make the REIT worse off, even under partial liquidity.

[16] An alternative data source, the MIT's Transaction-Based Index (TBI), which is based on transactions rather than appraised values has recently become available. Regrettably, this data does not exist at the level of disaggregation available in the the NCREIF's dataset. Because I want to simulate a strategy which times markets on the basis of very specialized portfolios, letting these portfolios then form the basis for possible aggregation into larger portfolios, I have chosen this data over the perhaps otherwise more desirable TBI.

appraised values (returns) as a function of true values (returns) and one- and four-quarter lagged values (returns). Equation 10 describes the true property price process as long-term mean reverting, while equation 11 is a derivation of the previous two equations, and equation 12 simply constitutes a constraint on the parameters. The unsmoothed index return becomes

$$r_t = \frac{r_t^* + b_1 r_{t-1}^* + b_2 r_{t-4}^*}{w_0} \tag{13}$$

I use the same basic estimation procedure as Cho et al. (2003) to estimate equation 11, which works as follows: I insert a starting value of 0.5 for ρ to estimate b_1 and b_2. These estimates are then inserted into a version of equation 11, rearranged to yield an estimate for ρ, and this model is estimated. The new value for ρ is reinserted into equation 11 and so forth.[17] After fifty iterations of this procedure, consecutive estimates differ by less than 10^{-8}.

Since in this study I am dealing with a panel dataset, I make the following expansion to the basic estimation methodology. I first undertake the iterative estimation procedure described above to derive single values of ρ, and thus b_1, and b_2 for the entire panel dataset, and call these $b_{1,whole}$ and $b_{2,whole}$. I then proceed through the disaggregated panel subsets: for any subset j which contains 20 or more observations, I redo the iterative estimation procedure for this subset individually, and define $b_{1,j}$ and $b_{2,j}$. I then desmoothe each subset using its own values of $b_{i,j}$ if these were defined, and if not using $b_{i,whole}$. This procedure should reduce the amount of outliers produced by estimating desmoothing parameters that have little statistical power because they are based on a small number of observations.

I apply the desmoothing procedure only to the appreciation subseries of each portfolio, as the income series is not the product of an appraisal process and therefore does not suffer from appraisal bias. The desmoothed appreciation returns are then added to the income returns to generate a new series of desmoothed total index returns. Since, for the implementation of the filter rules, I require a price series, I set each portfolio subindex to 100 the quarter before the first return figure is available, and then compute series of total index levels as cumulative products of the initial level and previous returns plus one.

I test the filter-based strategies described in Section 2.1.1 against a buy-and-hold strategy, which consists of simply buying each portfolio at the time its data becomes available and holding it until the end. For the filter-based strategies, I also buy the portfolio at the beginning of its data series and hold it until the first sell signal, at which point I invest the money generated by this portfolio in 3-month Treasury bills.[18] At

[17]A coarse grid search testing starting values of ρ between 0.05 and 0.95 by increments of 0.05 always yields conversion to the same final parameter estimates, with differences of less than 10^{-8} after at most 50 iterations.

[18]While the investment rule for REITs states that the firm must consist of at least 75% real-estate assets *or government*

the next buy signal, I sell the Treasury bills and buy the portfolio again and hold it until the next sell and so forth. If the last signal for the portfolio was a buy, I sell at the end of the data series. Should the data generate two identical signals in a row, I react to the first and thereby ignore the second. For each strategy, I record total returns and then do statistical comparisons between the total returns to the two strategies, throughout the entire cross-section of available portfolios.

I further test the mean outperformance of the filter-based timing strategies over the buy-and-hold strategy under round-trip transaction costs of 1, 2, 3,..., 10 percent for each transaction. The buy-and-hold strategy would incur these transaction costs once, and the filter strategy N times. Since I am interested in relative outperformance, I waive the transaction costs for the buy-and-hold strategy and charge the transaction cost $N-1$ times for the filter strategy.

Furthermore, under transaction costs, I also test the performance of an adaptive strategy. In this strategy I also buy at the beginning of the portfolio's life, but upon obtaining the first sell signal (generated either by trading bands or the moving average alone), I test whether the portfolio made a positive profit for this transaction, net of transactions costs. If yes, I sell and invest in T-Bills, if not, I ignore the signal. Once I sell, I buy at the next buy signal and upon encountering the next sell signal I carry out the same test before implementing it. This strategy yields fewer transactions, which may be advantageous under high transactions costs, but may make a manager stay in the market through downturns, during which he could otherwise have been out of the market and cut losses. For calibration, I also implement this strategy under zero transactions costs, just testing that a profit was made whenever selling.

Finally, I test the effect of the four-year and two-year holding period restrictions on the outperformance of all types of filter-based strategies described so far. In this case, I simply ignore sell signals that occur within four or two years, respectively, of the initiation of the portfolio, or less than four or two years after a buy transaction. As with the adaptive strategy, if I ignore a sell signal, I wait for the next one before I sell, ignoring any buy signals that occur along the way.

Through these methods, I aim to illustrate the mechanics of how the minimum holding constraints affect REITs' market timing abilities. Only in this controlled environment of simulated trading strategies on actual market data is it possible to truly compare the performance of a market timing strategy with and without the minimum holding constraints.

securities and REITs could therefore fully exploit this type of strategy, due to institutional constraints, it may be the case in practice that the manager of a property portfolio may not be allowed to be invested in assets other than property (including government securities) beyond a certain period of time. Even with this constraint a manager can implement this type of strategy, as long as he is active in a sufficient number of property submarkets; in this case, assuming signals are asymmetric enough across sub-markets he will simply transfer money from one submarket into another, instead of into T-Bills.

2.3 Results and Implications

Table 1 shows the annualized cross-sectional mean outperformance (excess returns), generated by the non-adaptive and the adaptive trading-band strategy under different levels of transaction costs, both in an unconstrained environment, as well as under teh four-year minimum holding period under which REITs operated.[19] Notice first the difference in outperformance of the unconstrained non-adaptive and the adaptive strategy. Under all levels of transaction costs, the plain filter-based strategy outperforms the adaptive strategy by a considerable margin, and thus the plain strategy would constitute the optimal investment strategy in an unconstrained environment. Even under round-trip transactions costs as high as 10%, the optimal unconstrained strategy outperforms a buy-and-hold strategy by approximately 2.11 percentage points (211 basis points) per year. This result thus illustrates once again the degree to which the property market is predictable.

With the four-year holding constraint, at any level of transaction cost, the adaptive strategy outperforms the non-adaptive strategy, although the difference is probably not statistically significant. For the sake of consistency, I use the adaptive strategy for comparison with the optimal unconstrained strategy, since the point estimates of the average outperformance for this strategy are larger, and in reality a portfolio manager would likely make the same decision, when confronted with such data. Under the same rationale, I use the adaptive strategy to study the effect the new two-year minimum holding period would have. The results from the non-adaptive strategy are omitted from the table to save space.

The t-statistics at the bottom of the table compare the outperformance of the optimal strategy under the four-year and two-year constraints, respectively, to the optimum outperformance without constraints (i.e. the performance of the non-adaptive strategy). I find that, under all levels of transaction costs, the excess returns from the constrained strategy are statistically significantly lower than those from the unconstrained strategy. The four-year constraint causes a shortfall in performance between 79 and 139 basis points per year, depending on the level of transaction costs. The relevant t-statistics at the bottom of the table lie between 6.77 and 9.29, indicating that this performance shortfall is highly statistically significant at all levels of transaction costs, and that therefore the minimum holding period constraint constitutes as significant hindrance to the implementability of this market timing strategy. While the two-year minimum holding period constitutes less of a hindrance to these trading-band based market timing strategies (the shortfall in performance ranges from about 58 to 96 basis points per year), the t-statistics at the bottom of the table lie

[19] These figures are constructed by computing, for each sub-portfolio, the total performance of the timing strategy over the portfolio's life, subtracting the total performance of the buy-and-hold strategy over the portfolio's life, dividing by the length (in years) over which the portfolio is active, and then computing means and t-statistics over the cross-section of portfolios.

between 2.22 and 2.96, and therefore still indicate that the new relaxed minimum holding constraint has a significant effect on the outperformance of this market timing strategy at all levels of transaction costs.

Table 2 shows the outperformance of a strategy based on trading filters generated only by the price breaking through its moving average, as opposed to trading bands around the moving average. The point estimates of the outperformance of this strategy are actually higher by about 30 to 60 basis points, under all versions of the strategy, all levels of transaction costs, and all minimum holding period constraints than those from the trading-band strategy. While it is unclear whether the difference between these estimates is statistically significant, this question is actually not primarily relevant to the aim of this study. What we can learn from this table, however, is that the effects of the minimum holding constraints are qualitatively very similar for the moving-average only strategies as they are for the trading-band strategies. While the four-year constraint causes a slightly smaller difference in outperformance between the optimal unconstrained strategy and the optimal constrained strategy (about 70 to 110 basis points per year),[20] the difference is highly statistically significant at all levels of transaction costs, with t-statistics of 6.85 to 9.31. The effects of the two-year constraint are analogous to those of the four-year constraint. The drop in outperformance due to the constraint is 55 to 83 basis points, thus lower than that in the previous table. However, the shortfall in performance is statistically highly significant, with t-statistics of 6.03 to 7.88. Overall, as with the trading-band strategy, the relaxation of the minimum holding period to two years improves the profitability of this strategy a little, but still causes a significant performance shortfall.

It should be apparent that any outperformance by a timing strategy over a buy-and-hold strategy must necessarily be due to appreciation returns, and not income returns. Counting income returns only, the buy-and-hold strategy will necessarily outperform a market timing strategy, since on average the income generated by a property portfolio in this dataset exceeds the return on Treasury bills by 60 basis points per quarter. Hence, due to both these series' low volatilities, periods in which a property portfolio underperforms T-bills (during which a timing strategy could dictate staying out of the market) are extremely rare. Thus, on an income basis, the strategy which stays in the market the longest will generally perform best. Hence, the appreciation-driven excess returns generated by the timing strategies are actually large enough to much more than outweigh the additional income generated by the buy-and-hold strategy. The implication for this study should thus be clear: the missed returns of the constrained strategies over the unconstrained strategies are purely appreciation returns, which is also in line with the findings of Mühlhofer (2008).

I have thus demonstrated that at any level of transaction costs both the old four-year as well as the

[20] The non-adaptive strategy is the optimal one here, both in the unconstrained setting as well as for both levels of minimum holding periods.

new two-year minimum holding requirements make a REIT worse off than an unconstrained investor (at least with respect to this set of strategies, which, since they are feasible, shows in general that a REIT's opportunity set for generating abnormal profits is diminished by the constraints). To illustrate this point, Table 3 shows distributional statistics of the holding periods dictated by these strategies, for the non-adaptive trading-band strategy, the adaptive trading-band strategy (under the different transaction cost levels), as well as the non-adaptive moving-average strategy, which is the optimal version of this strategy.

Note the mass of the distribution of holding periods for the non-adaptive trading-bands strategy that is below 16 (or even 8) quarters. In fact, as is shown in Table 3, even the median of the distribution is below 16 (at 12) quarters. The first quartile actually lies at three quarters. While because of market frictions such a short holding period may not be feasible in reality, these figures still strongly show the necessity for quick trading in implementing a strategy such as this one. This strategy (and thus this distribution of holding periods) is optimal for all levels of transaction costs, so both the four-year holding constraint and the two-year holding constraint are very binding with respect to the strategy's implementability.[21] While the moving-average strategy dictates slightly longer holding periods, there is still a substantial portion of the distribution below four years or even two years, and so these results are robust for this strategy as well. REITs that would be especially affected by this constraint would tend to be large REITs, which, as was said before, may be able to trade at fairly low transaction costs, because they tend to be able to internalize many of the services required in a property trade.

It should be noted that my discussion in this section has only focused on first moments, that is, on absolute outperformance, without undertaking any risk adjustment. The reason for this is that, within each portfolio, the buy-and-hold strategy will, by construction, be more volatile than the filter strategy, because T-Bill returns are less volatile than property returns. Thus, a manager who follows the timing strategy will be invested into a strictly less volatile asset than a manager who follows the buy-and-hold strategy during those times when he is invested in T-Bills, and therefore the timing portfolio's overall volatility over the life of the portfolio will by construction be at least weakly lower than that of the buy-and-hold portfolio.

A possible objection to these results may be that REITs are set up as property companies, have ample overheads stemming from this (such as maintaining in-house acquisition and development teams), and therefore would have a difficult time justifying a withdrawal from property markets, investing primarily in treasuries. This concern can be alleviated by the fact that a cross-sectional assessment of these strategies

[21]As can be seen from these distributional statistics, even the adaptive strategy (which of course is only optimal under the four-year constraint), would theoretically require a large number of holding periods below four years, under all levels of transaction costs.

on all portfolios simultaneously reveals that at any given time, 80% or more sub-portfolios are active (i.e. the strategy is never out of more than 20% of sub-markets), except during the time from the first quarter of 1991 until the second quarter of 1994, as the commercial property market was in the process of recovering from a crisis. Therefore, it seems that in normal times a strategy such as this one does not prevent a REIT from pursuing its overall objective of being a property company, and all the overheads that are connected to this should thus be justified. In times of crisis, on the other hand, it might actually be appropriate if REITs actually sold properties and put their money in treasury securities (which they are allowed to do by law), and at the same time reduced their property-investment related overheads. Otherwise, if that is impossible for institutional reasons, the REIT could simply downsize, repurchase equity, and each investor can invest in treasuries or other securities. The fact that the money taken out of a particular market is invested in treasuries here could reflect either the REIT investing in treasuries rather than property, or the REIT downsizing (thereby also cutting overheads), and the individual investor investing in treasuries. The result would be similar.

In this section, I have presented an example of a set of trading strategies that can be used to time the direct real estate market to significantly outperform a buy-and-hold strategy, and shown that the implementation of this strategy is hindered to a considerable extent by the minimum holding period constraints which REITs face. Hence, I have shown that a REIT's opportunity set for making abnormal appreciation profits in the real estate market is diminished by the holding-period constraints it faces, and have illustrated the mechanics of how this happens.

It should be noted here again that no REIT will ever simultaneously invest in all the portfolios upon which this trading strategy is tested. In reality, a REIT will invest into a small combination of these portfolios, and the performance of each portfolio will be a random variable drawn from a performance distribution which in its mean significantly exceeds the performance distribution of the buy-and-hold strategy, and does so to a significantly lesser extent under the holding period constraints. Furthermore, in this section, I only presented one set of possible (and, I argue, feasible) trading strategies with which to time the market, and there may exist many other such strategies. In reality, while a technical strategy may be used, it would invariably also be combined with fundamentals-based signals, and thus the required amount of transactions could even increase from what is shown here, especially in the upper tail of the holding-periods distribution. In any case, however, I have shown in this section that REITs can outperform the market by pursuing an actively traded market timing strategy, and that the holding-period constraints imposed upon REITs significantly reduce their ability to do so.

3 Actual Holding Periods

3.1 The Empirical Model

After having shown the mechanics of how a REIT's profits from an active market timing strategy can be lowered by the minimum holding constraints, I now investigate whether REITs are indeed likely pursuing active strategies requiring short holding periods, by examining the empirical distribution of property holding periods that exists in the market. To directly assess the bindingness of the constraint, I then investigate whether time held is affected by the performance of the local real estate market, since the minimum holding constraint only binds when the REIT has made a profit on the transaction, as it is enforced through a 100% gains tax. This sanction becomes irrelevant if no gains were made. The total holding-period return to the local market index proxies for the actual transaction return a REIT made.[22] As discussed earlier, UPREITs have an advantage over regular REITs with respect to the holding-period constraints, because they acquire properties primarily through 1031 exchanges and their partnership structure allows them to do so efficiently. In this way, they inherit holding periods from the contributing partners, which enables them to meet the safe-harbor requirements sooner. I therefore control for a firm's UPREIT status in all empirical tests. Conceptually, for preliminary evidence, I estimate the following regression:

$$hp_{i,l} \quad = \quad \alpha + \beta_1 ret_l + \beta_2 UPREIT_i + \beta_3 ret_l \cdot UPREIT_i + \epsilon_i \tag{14}$$

Here $hp_{i,l}$ is the amount of time (in years) property i, located in area l, was owned by a specific REIT, ret_l is the total return on the local real estate market during that time, and $UPREIT_i$ is a dummy variable indicating whether the REIT that owns property i is an UPREIT. Under the null hypothesis, holding period should not be determined by any of these explanatory variables, while under the alternative of a binding holding constraint, one should observe that higher local market returns are associated with longer holding periods for regular REITs, as these firms will want to wait until the four years have elapsed, and even after that will probably not sell immediately, but at the next sell signal, similarly to the approach presented in the simulation.

Due to the binary nature of the constraint (i.e. the safe harbor rule applies exactly at four years and longer, and not before[23]), to more precisely model the specific decision to hold a property beyond four years,

[22] Purchase and sale prices for REIT portfolio properties are only available for a small subset of my data, and thus these series may suffer from a sample selection bias. The local market return I use as a proxy, on the other hand, is very readily available. As stated in Crane and Hartzell (2007), in those instances where actual prices are available, the holding period return implied by these prices is correlated with the returns implied by the local market index at about 87%, so this seems to be a good proxy.

[23] Because my data sample ends in 2004, I am unable to judge the effect of the relaxation of the minimum holding requirement

I also estimate a probit regression, as follows:

$$fourplus_{i,l} = \alpha + \beta_1 ret_l + \beta_2 UPREIT_i + \beta_3 ret_l \cdot UPREIT_i + \epsilon_i \tag{15}$$

Here, $fourplus_{i,l}$ is a binary variable which indicates whether property i was held for more or less than four years, and everything else is as defined above. Under a null hypothesis that no type of REIT is bound by the four-year constraint, one would observe that local market conditions have no effect on the decision to hold a property beyond four years for either regular REITs or UPREITs. Under the alternative hypothesis, one should observe that a regular REIT is more likely to hold a property beyond four years the higher the local market return (and hence the higher the profit made on the transaction), while an UPREIT should be indifferent to this.

3.2 Data and Methodology

For holding period data, in this section, I use the property database from SNL Datasource. For each property, I use the date it was bought by a specific REIT, the date it was sold, the identifier for the firm that had it in its portfolio, and the CBSA code for the property's location.[24] Note that, since I examine the decision to sell a property from a REIT's portfolio, I only use properties that were bought and sold within the period covered by the database (up to 2004) and not properties which were still in a REIT's portfolio at the time the database's coverage ends.[25]

I then combine this data with SNL's firm-by-firm database, from which I extract information on whether a REIT was operating as a declared UPREIT the year it sold a property.[26] Correspondingly, I define the dummy $UPREIT$ as

$$UPREIT_i = \begin{cases} 1 & \text{if a REIT was an UPREIT at time of sale} \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

from four to two years, which occurred in 2008. I therefore only mention the four-year constraint throughout this section.

[24] These Core-Based Statistical Areas (CBSAs) have taken the place of the old Metropolitan Statistical Areas (MSAs) in defining urban regions throughout the United States.

[25] SNL's database also signals whether a property was bought as part of a merger of two REITs. It seems that such properties would still underlie the four-year constraint for the acquiring firm, using the merger date as the purchase date, and thus I have included these properties in my study. However, the results are not sensitive to excluding these properties.

[26] SNL derives the UPREIT classification from REITs' 10-K filings, depending on whether a REIT declares holdings of partnership shares. For any particular firm this classification could potentially change over time; however, what is important for the bindingness of the constraint, and thus for the purpose of this study is whether or not the firm was an UPREIT at the time of sale of the property, and thus I classify firms according to this criterion.

Table 4 presents summary statistics for this data. Panel A shows the number of properties of each major type, broken down by whether the holding firm is an UPREIT. Note that there are many more properties held by UPREITs than by non-UPREITs, and that depending on property type, one form of REIT seems to dominate the other. This is especially dramatic for multi-family properties for which 1734 properties are held by UPREITs and only 138 by non-UPREITs. For health care, on the other hand, the dataset contains 229 properties held by regular REITs and 51 held by UPREITs. Note also that the dataset contains only 40 non-UPREITs, compared with 100 UPREITs. This is very much in line with the overall REIT industry, in which the UPREIT has become the dominant firm structure in recent years. Panel B of Table 4 shows summary statistics for the number of properties in the top-10 CBSAs (that is, the ten CBSAs in this dataset that contain the most properties). Note that about one third of all properties in this dataset are in a top-10 CBSA, while the other two thirds are distributed over the other 487 CBSAs. Finally, Panel C of Table 4 shows the property distribution by size of the selling firm. I classify a firm as *Large* if it is in the set of firms which constitute the upper two-thirds of total REIT industry market capitalization. Note that there are many more properties sold by large UPREITs than non-UPREITs; thus, since these two firm characteristics seem somewhat collinear, I will be accounting for both of them in an augmented version of the holding periods model. Note also that there are much fewer distinct large firms than small firms. This is indicative of the size distribution of the REIT industry in general.

As a measure of the performance of the local real estate market, I use the total return series from the MSA-level subindices of the National Property Index (NPI) from NCREIF. The appreciation components are desmoothed, as discussed in Section 2.2. Since the regional identifiers in the property dataset are the newer CBSA codes, while those in the market index dataset are the older MSA codes, I translate from one to the other, in order to match each property with its correct MSA-area subindex values. For each property, I determine the quarter it was purchased and record the level of the local real estate market subindex at the end of that quarter. An analogous procedure is used for the sale date and the current level of the local subindex.

From this data, I then construct annualized returns, as total returns over the entire holding period, divided by the holding period in years. The reason for using annualized returns comes from the assumption that the property market price process contains a drift parameter, generating a positive average annual return, equivalent to the risk premium required by investors. Using total cumulative returns that are not adjusted for holding period would thus create an endogeneity problem, as on average a longer holding period would necessarily be associated with a higher return. Using annualized returns fixes this issue. I thus define the variable $pareturn_l$ as the per annum CBSA-wide market return for area l, during the time a REIT owned property i located in CBSA l.

I further define a dummy variable $fourplus_i$ as follows:

$$fourplus_{i,l} = \begin{cases} 1 & \text{if } hp_{i,l} \geq 4 \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

Another consideration in empirically implementing the tests described above, is that, since holding periods cannot be negative, equation 14 would be empirically mis-specified as it stands, since the dependent variable has a distribution that is truncated at zero. I therefore use the natural log of holding period ($ln(hp_{i,l})$) as the dependent variable for the empirical version of this equation.

The actual models I estimate thus become

$$\begin{aligned} ln(hp_{i,l}) &= \alpha + \beta_1 pareturn_l + \beta_2 UPREIT_i \qquad (18) \\ &\quad + \beta_3 pareturn_l \cdot UPREIT_i + \epsilon_i \\ fourplus_{i,l} &= \alpha + \beta_1 pareturn_l + \beta_2 UPREIT_i \qquad (19) \\ &\quad + \beta_3 pareturn_l \cdot UPREIT_i + \epsilon_i \end{aligned}$$

As a robustness check, I further estimate a version of the probit model (equation 19) which includes firm size as well as UPREIT status. As discussed earlier, a large REIT should be able to turn properties over more quickly than a smaller one and still make a profit, as it generally faces lower transaction costs, and thus it may be advantageous to control for size in the model. Further, in the sample, larger firms have a stronger tendency to be UPREITs, and so including both size and UPREIT status might help distinguish between these two somewhat collinear firm characteristics.

This model, thus, becomes:

$$\begin{aligned} fourplus_i &= \alpha + \beta_1 pareturn_l + \beta_2 UPREIT_i \qquad (20) \\ &\quad + \beta_3 pareturn_l \cdot UPREIT_i \\ &\quad + \beta_4 Large + \beta_5 pareturn_l \cdot Large_i + \epsilon_i \end{aligned}$$

where $Large$ is a dummy variable which takes a value of 1 if a firm is part of the upper two-thirds of

REIT industry market capitalization and 0 if a firm is in the lowest third.[27] I use stock prices and shares outstanding from the Center for Research in Securities Prices (CRSP) database for the beginning of the quarter in which a property is sold, to determine firm size.

In order to observe whether localized real estate market performance affects the decision to hold a property beyond the four-year mark, I limit the data to properties held up to six years only, to better observe how decisions close to the four-year cutoff are made.[28]

3.3 Results and Implications

Table 5 shows distributional statistics for the holding periods of all properties in the SNL dataset. Notice the amount of distributional mass that lies in the low holding periods. The median holding period for all properties is exactly four years. Thus, one can definitely not say that the four-year holding constraint is irrelevant to the holding periods which REITs require: as this distributional analysis shows, the four-year holding constraint is extremely binding, and REITs seem very eager to dispose of many properties after a fairly short period of time, to such an extent that actually 50% of properties are sold before the four-year mark. Furthermore, the similarities between the empirical holding period distributions shown here, and that produced by the simulated market timing strategy should be noted. While the strategy that I implemented may not be the only way to time the market, it seems to be the case in reality that REIT managers want to engage in actively traded market timing strategies, because of the profits that these yield.

Notice next the difference in distributional statistics between regular REITs and UPREITs. As shown in Table 5, the median holding period for non-UPREITs is 5.003, while the median holding period for UPREITs is 3.836 years and the first quartiles differ by almost as much. This data clearly shows that REITs have the need to follow an active trading strategy with short holding periods, and that UPREITs do this more than non-UPREITs. Notice also that a Kolmogorov-Smirnov test strongly rejects the null hypothesis that the density functions for UPREIT holding periods and non-UPREIT holding periods are identical, in favor of the alternative that the CDF for UPREIT holding periods lies above that for non-UPREIT holding periods. This means UPREITs tend to trade more often than non-UPREITs, and this difference is strongly statistically significant.

An interesting question at this point might be why regular REITs still have such a large mass of holding periods below four years. Some of these transactions may have incurred losses and so there is no incentive

[27] This variable definition is analogous to that used in the portfolio sort in Mühlhofer (2008).

[28] Robustness tests reveal that these results are not sensitive to the selection of the cutoff, for any integer year between five and ten years.

to hold the property for the whole four years, if the firm can cut its losses by selling earlier. On some other transactions, especially where the profit is small, it may actually be advantageous to surrender profits to the government, in order to be able to leave a market before a downturn, rather than waiting until the property has actually incurred a loss, especially if market liquidity is expected to be low in the anticipated downturn. When the market has passed its trough, the REIT can buy back into this market, and thus realize higher profits at the next peak.[29] Furthermore, recall that the four-year cutoff merely constitutes a safe harbor (see Footnote 2). If a REIT can prove on its own that the property was 'not held primarily for resale', then it can sell this property before four years are over. However, as stated earlier, such a situation would be idiosyncratically motivated, and furthermore whether or not a firm manages to do this would be difficult to observe on a large cross-section.

Table 6 shows the regression results of log holding period on the annualized local market index return and the $UPREIT$ dummy. Notice, first of all, the positive and strongly significant coefficient on *pareturn*. This implies that an up-market is associated with longer holding periods: this, of course, deviates from the null hypothesis that REITs are indifferent to the four-year constraint, as under this assumption return should be unrelated to holding period length, since in a market with irregular fluctuations, there should be no relationship between total annualized returns and holding period length. Next notice the incremental effect for $pareturn \cdot UPREIT$, which is negative and significant. As the F test shows, the net effect on UPREITs is not zero, given the size of this coefficient, but, nonetheless, UPREITs show a significantly shorter holding period associated with a rising market than regular REITs. Considering that the subsample for this regression consists of only properties that were held six years or less, this combination of results fits well with the trading-restrictions hypothesis: regular REITs want to retain the profits from a rising market and thus hold properties beyond the four-year constraint. Furthermore, in most cases, just like in the simulated trading strategy, it will be optimal to wait beyond the fourth year, until a new sell signal occurs, before selling a property, which would cause holding periods that are somewhat scattered between the four-year minimum and the six-year cutoff. Hence, especially in this subset of properties being sold within six years, this result matches well with the hypothesis of a binding constraint. UPREITs, which, of course, manage to avoid the holding period constraint by acquiring properties through 1031 exchanges, show this effect much less. This should constitute some preliminary evidence for the bindingness of the constraint.

Table 7 shows the results from the probit regressions, which model the specific decision to hold a property beyond four years on the annualized local market return and firm characteristics. Once again, in Model 1, the

[29] Assuming, of course, that this peak will occur after four years or more.

coefficient for *pareturn* is positive and strongly significant, indicating that the probability that a property which is held for six years or less is held beyond the all-important four-year mark is significantly higher in a rising market. Similarly to the previous result, the coefficient for $pareturn \cdot UPREIT$ is negative and strongly significant, indicating that an UPREIT is significantly less likely to hold a property beyond four years in a rising market, which supports the trading restrictions explanation. The χ^2 statistic measuring the total net UPREIT effect shows that UPREITs are still somewhat more likely to hold a property beyond four years in a rising market than in a falling one, but comparing this result with its equivalent in Model 2 suggests that this might simply be due to an omitted variable bias generated by not including size in the model.

Controlling for size, as I do in Model 2, strengthens this result. The coefficient for *pareturn* remains largely unchanged; the absolute value of the coefficient for $pareturn \cdot UPREIT$ increases, while its standard error remains largely unchanged, thus showing an even stronger negative UPREIT effect, once size is accounted for.[30] The χ^2 test at the bottom of the table, in fact, reveals that the total coefficient for UPREITs (that is, the base effect on *pareturn* plus the incremental effect of $UPREIT \cdot pareturn$) is not distinguishable from zero. This implies that regular REITs are significantly bound by the four-year holding constraint, while UPREITs look indifferent to it, no matter how large of a profit they make from the transaction. The concept of the four-year constraint as a *safe harbor rule* also fits well into these results. If a regular REIT sells a property before four years, the firm can retain the profits from the transaction if it can prove that the property was not 'held primarily for resale'. This would be proportionally more difficult to do, the higher the profit from the sale; thus, also in this respect, the REIT would be more severely bound by the four-year constraint if it makes a larger profit from a transaction. Therefore, while in theory the four-year minimum holding period is not binding to the letter of the regulation, the spirit of the regulation forbidding active trading, together with the safe harbor rule, makes the four-year constraint de-facto binding for more profitable transactions.

Turning to size effects, *Large* and $pareturn \cdot Large$ are significant, the former negative and the latter positive. This suggests that in general large firms are more likely to sell early, but that in a rising market they are more likely to hold beyond the four-year mark than small firms. This result is most likely driven by large non-UPREIT firms: as large firms tend to be more closely followed by investment analysts than smaller firms, management would perhaps get penalized more severely for the foregone profits of selling early

[30] As remarked in Ai and Norton (2003), among others, marginal effects in probit models are non-linear, and thus in order to interpret this interaction effect correctly, the value of the cross-derivative that constitutes this effect needs to be examined over a wider range of values. In this model, the interaction effect computed in this manner is significantly negative over the entire range of predicted probabilities generated by it.

in a rising market,[31] and thus the probability of a large firm holding a property beyond four years in such a market is even greater than for a smaller firm. Neutral of market performance, however, if a large firm faces lower transactions costs, as argued earlier, its management would be more likely to actively trade its portfolio, and therefore be less likely to hold beyond four years, to cut losses, or if it is able to meet the burden of proof for transacting outside the safe harbor rules. What is more important for this investigation, however, is that the positive coefficient for $pareturn$ and the negative coefficient for $pareturn \cdot UPREIT$ are even stronger when accounting for size. Only the UPREIT structure allows a firm to overcome its four-year selling constraint, while other firm characteristics would not have this effect.

As is evident from Table 4, certain property types are held predominantly by UPREITs and a few are held predominantly by non-UPREITs. In order to rule out that the UPREIT dummy in reality proxies for a property-type effect, in Model 3 (Table 7), I include a dummy variable for each property type, as a robustness check. As can be seen from the results of this model, including these control variables has very little effect on magnitude, sign, and significance level of the remaining coefficients, and thus these results are almost unaltered by this inclusion. Hence, these results do not seem to be driven by property-type-specific phenomena for which UPREIT status merely proxies. Rather, these results seem to be driven by a wide-spread phenomenon linked to local market return and UPREIT status.

These results give strong evidence for the bindingness of the holding constraint. If a REIT is likely to make a profit on a transaction, it wants to retain these profits and therefore has no choice but to hold each property beyond the four-year mark and then sell it at the next available opportunity. UPREITs, which can dispose of properties without holding constraints, do so.

4 Conclusion

In this paper I assess the bindingness of the minimum property holding constraint faced by REITs, and illustrate how this constraint hinders REITs in their ability to time the property market to take advantage of its predictability. In order to do this, I first devise a set of filter-based trading strategies, and show that these strategies considerably outperform a simple buy-and-hold strategy over a wide range of transaction cost levels. I analyze the holding periods required by this strategy, and find that a large portion of the empirical distribution of these holding periods is shorter than either the four-year minimum, or even the new two-year minimum imposed on REITs. Further, I find that when imposing a four-year minimum holding constraint

[31] Assuming rational profit-maximizing managers such a sale might occur for liquidity reasons, for example. Otherwise there might be agency problems involved.

on this strategy it substantially underperforms the unconstrained strategy. While the performance shortfall is lower with the two-year minimum holding period, even with the new more relaxed two-year minimum the strategies statistically significantly underperform an unconstrained strategy. The differences between the returns of the constrained and unconstrained strategies are especially large for scenarios in which there are low transaction costs, thus suggesting that large REITs, which often manage to trade with very favorable costs, are especially strongly affected by the trading constraints.

While the filter-based set of strategies I use constitute only a few possible trading strategies, they allow me to show how appreciation profits from such active timing strategies are reduced by the trading constraints REITs face, and thus how these trading constraints generally reduce REITs' opportunity sets for profit making. Furthermore, most managers will combine any technical trading strategy with fundamentals-based information, possibly leading to a necessity for an even larger number of short holding periods, especially in the upper tail of the holding periods distribution. It would also seem that this illustration of how the trading constraints hinder this strategy in generating market-timing profits could be easily transferable to many other market timing strategies, and thus offers general insights about the mechanics by which the holding constraint prevents REITs from generating abnormal profits.

I then analyze actual holding periods of properties in REITs' portfolios. I find that a significant amount of the distribution of holding periods lies below or just above the four-year mark (which was the relevant minimum for the safe harbor rule throughout my sample), suggesting that REITs are eager to hold properties for a short time and that the four-year holding period is very binding. This matches with the holding periods required by the simulated trading strategy and suggests that REIT managers want to try to time the market, as, especially given the predictability of the direct property market, large profits can be made by doing this. I also find that UPREITs, which by trading efficiently through 1031 exchanges can overcome these constraints, tend to trade more often than regular REITs, which are more affected by the holding constraints.

I then analyze holding period length as determined by the current market situation and find that holding periods tend to be longer in rising markets for regular REITs and much less so for UPREITs. I also analyze the decision to hold a property for more than four years in a set of properties that were held for six years or less, and find very strongly that regular REITs are more likely to hold a property beyond four years, the more of a profit they are likely to make on the transaction, while for UPREITs this relationship is significantly weaker. These results are strengthened when firm size is accounted for in the model. In fact, with firm size included, results show that UPREITs are completely indifferent to the four year constraint, even in a rising market. Accounting for possible systematic effects associated with property types leaves these results unchanged. This, once again, lends strong support to the hypothesis that the selling constraints are binding, as UPREITs can simply sell a property whenever the manager feels the ideal time is (because they inherit the holding period of the contributing partner), while regular REITs must wait beyond the four-year mark in order to be able to retain the profits made in the rising market.

These results also fit well with the consideration that the four-year constraint merely constitutes a *safe harbor*, and that REITs can actually sell properties earlier if they can prove that the property was not 'held primarily for resale'. In fact, the more of a profit a REIT makes on a particular transaction, the more difficult it will be to prove such a claim. Thus, since the burden of proof is on the REIT, even in this expanded view, the four-year minimum holding period is effectively unavoidable for transactions in which the REIT makes a larger profit.

The results from this study should thus have some policy implications. While, in the 1960s when REITs were invented as small funds for real estate, it might have made sense to constrain them from *flipping* properties quickly, in order to safeguard the quality of the industry, in today's world in which many REITs are large vertically integrated property companies which are closely followed by analysts, it may not be necessary to have such a constraint anymore, just to guarantee quality. As I show in this study, the constraint limits REITs' abilities to make timing profits they could ordinarily make by their nature of being informed traders in a predictable market, and the constraint thus makes REITs inferior investment choices to non-REIT vehicles in this respect. While relaxing the minimum holding period to two years somewhat helps this situation, I show that this relaxed version of the safe harbor rules still constitutes a significant hindrance to a REIT's ability to realize market timing profits. While the UPREIT structure seems to allow firms to partially overcome this constraint, it may be more efficient to eliminate the constraint from legislation, as there are other factors involved in the decision on whether to set up a REIT as an UPREIT, and the necessity to create an UPREIT merely to circumvent this legislation seems inefficient at best.

References

Ai, C. and Norton, E. C.: 2003, Interaction terms in logit and probit models, *Economics Letters* **80**, 123–129.

Alexander, S.: 1961, Price movements in speculative markets: Trends or random walks, *Industrial Management Review* **2**, 7–26.

Alexander, S.: 1964, Price movements in speculative markets: Trends or random walks, *in* P. Cootner (ed.), *The Random Character of Stock Market Prices*, Cambridge, MA: Massachusetts Institute of Technology Press.

Barkham, R. and Geltner, D.: 1995, Price discovery in american and british property markets, *Real Estate Economics* **23**(1), 21–44.

Blume, L., Easley, D. and O'Hara, M.: 1994, Market statistics and technical analysis: The role of volume, *Journal of Finance* **49**, 153–181.

Bollinger, J. A.: 2001, *Bollinger on Bollinger Bands*, McGraw-Hill.

Brock, W., Lakonishok, J. and LeBaron, B.: 1992, Simple technical trading rules and the stochastic properties of stock returns, *Journal of Finance* **47**, 1731–1764.

Case, B. and Quigley, J.: 1991, Dynamics of real-estate prices, *Review of Economics and Statistics* **73**(1), 50–58.

Case, K. and Shiller, R.: 1990, Forecasting prices and excess returns in the housing market, *AREUEA Journal* **18**(3), 253–273.

Cho, H., Kawaguchi, Y. and Shilling, J. D.: 2003, Unsmoothing commercial property returns: A revision to fisher-geltner-webb's unsmoothing procedure, *Journal of Real Estate Finance and Economics* **27**(3), 393–405.

Clayton, J., Geltner, D. and Hamilton, S.: 2001, Smoothing in commercial property valuations: Evidence from individual appraisals, *Real Estate Economics* **29**(3).

Cooper, M., Downs, D. H. and Patterson, G. A.: 1999, Real estate securities and a filter-based, short-term trading strategy, *Journal of Real Estate Research* **18**(2), 313–334.

Crane, A. D. and Hartzell, J. C.: 2007, Is there a disposition effect in corporate investment decisions? evidence from real estate investment trusts, *SSRN eLibrary, http://ssrn.com/paper=1031010* .

Crosby, N. and McAllister, P.: 2005, Liquidity in commercial property markets: Deconstructing the transactions process, *Working Paper, University of Reading* .

Fama, E.: 1965, The behavior of stock markets, *Journal of Business* **38**, 34–105.

Fama, E. and Blume, M.: 1966, Filter rules and stock market trading profits, *Journal of Business* **39**, 226–241.

Geltner, D. and Mei, J.: 1995, The present value model with time-varying discount rates: Implications for commercial property valuation and investment decisions, *Journal of Real Estate Finance and Economics* **11**(2), 119–135.

Gyourko, J. and Keim, D. B.: 1992, What does the stock market tell us about real estate returns?, *AREUEA Journal* **20**, 457–485.

Kawaguchi, Y., Sa-Aadu, J. and Shilling, J. D.: 2007, Do changes in illiquidity affect investors' expectations? an analysis of the commercial property returns, *Working Paper* .

LeBaron, B.: 1996, Technical trading rule profitability and foreign exchange intervention, *Working Paper 5505, NBER* .

Liu, C. H. and Mei, J.: 1992, The predictability of returns on equity reits and their comovement with other assets, *Journal of Real Estate Finance and Economics* **5**(4), 401–418.

Liu, C. H. and Mei, J.: 1994, An analysis of real estate risk using the present value model, *Journal of Real Estate Finance and Economics* **8**, 5–20.

Mühlhofer, T.: 2008, Why do reit returns poorly reflect property returns? unrealizable appreciation gains due to trading constraints as the solution to the short-term disparity. Available at SSRN: http://ssrn.com/abstract=912807.

Murphy, J. J.: 1999, *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*, Prentice Hall Press.

R Development Core Team: 2008, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
URL: *http://www.R-project.org*

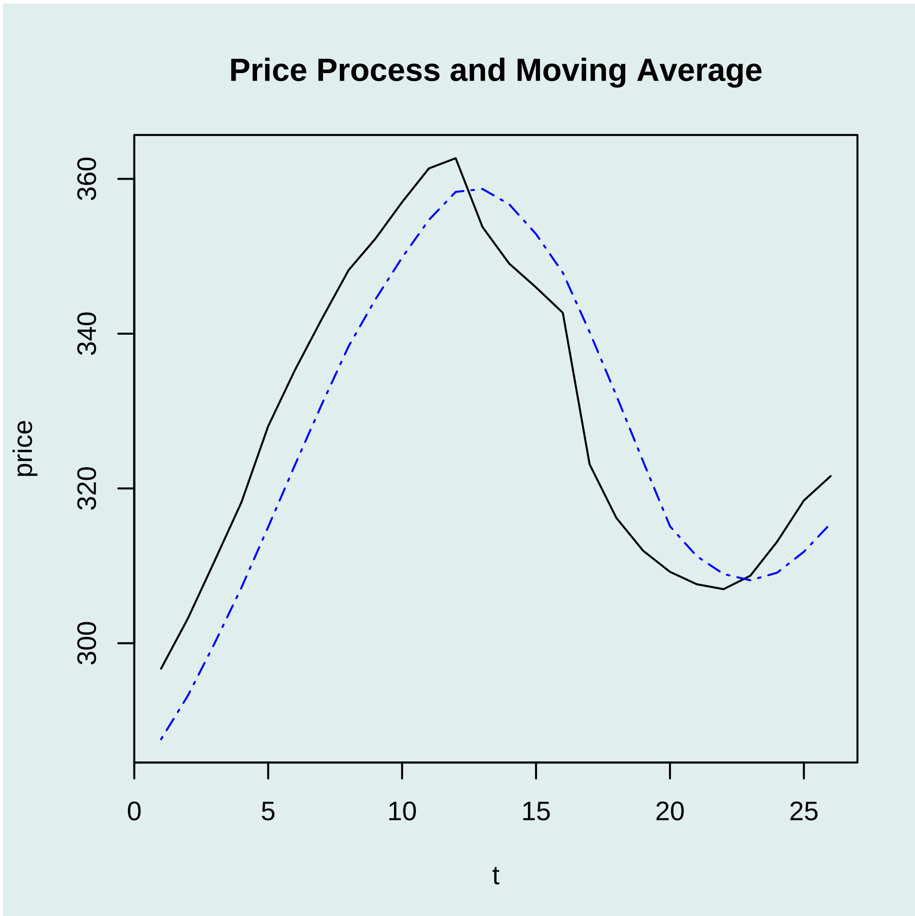

Figure 1: Illustration of the relationship between a price process and its moving average: the price series is the black solid line, while the moving average is the blue alternation of dots and dashes. Data from the NCREIF national property index.

Figure 2: Illustration of the relationship between a price process, its moving average, and a set of Bollinger bands constructed around it: the price series is the black solid line, the moving average is the blue alternation of dots and dashes, and the Bollinger bands are the red dashed lines. Data from the NCREIF national property index.

Table 1: Annualized Mean Outperformance of Trading-Band-Based Strategy over Buy-And-Hold Strategy.

This table shows the per-annum mean outperformance of the trading-band filter-based strategies over the buy-and-hold strategy, across all subportfolios of the NCREIF NPI universe and across the entire 1978-2003 time period. The adaptive strategy differs from the non-adaptive strategy in that sell signals are ignored if after transaction costs they do not lead to a positive profit for a particular round-trip transaction. Both types of strategy are tested with and without the 4-year holding constraint. Results are also reported for the optimal strategy with the 2-year holding constraint (Adaptive Strategy).

Transaction costs										
0	**1%**	**2 %**	**3%**	**4%**	**5%**	**6%**	**7%**	**8%**	**9%**	**10%**
Non-Adaptive Strategy										
0.03071***	0.02975***	0.02880***	0.02784***	0.02688***	0.02592***	0.02496***	0.02400***	0.02304***	0.02209***	0.02113***
Adaptive Strategy										
0.02306***	0.02117***	0.01972***	0.01863***	0.01771***	0.01654***	0.01592***	0.01518***	0.01504***	0.01463***	0.01484***
Non-Adaptive Strategy, 4-Year Constraint										
0.01668***	0.01630***	0.01591***	0.01552***	0.01514***	0.01476***	0.01437***	0.01399***	0.01360***	0.01322***	0.01284***
Adaptive Strategy, 4-Year Constraint										
0.01683***	0.01647***	0.01612***	0.01571***	0.01535***	0.01501***	0.01466***	0.01431***	0.01395***	0.01363***	0.01322***
Adaptive Strategy, 2-Year Constraint										
0.02111***	0.02058***	0.01994***	0.01933***	0.01881***	0.01831***	0.01769***	0.01700***	0.01637***	0.01583***	0.01530***
T-Statistics										
9.2917	9.0514	8.8094	8.5851	8.3278	8.0701	7.8188	7.5459	7.2872	7.0095	6.7677
2.9602	2.8609	2.8037	2.7362	2.6372	2.5310	2.4714	2.4486	2.3960	2.3101	2.2217

Average annualized return for buy-and-hold strategy: 0.1680

$^{\circ}$: significance level $\leq$ 10%. *: significance level $\leq$ 5%. **: significance level $\leq$ 1%. *** : significance level $\leq$ 0.1%.

The significance stars next to the outperformance values are for the hypothesis H_0: *mean outperformance of filter strategy over buy-and-hold strategy* $= 0$ against the positive one-sided alternative.

The two sets of t-statistics reported at the bottom of the table are for the hypothesis H_0: *mean outperformance of optimal unconstrained strategy (Non-Adaptive Strategy) over optimal constrained strategy (Adaptive Strategy)* $= 0$, against the positive one-sided alternative. The top line compares the unconstrained strategy to the optimal strategy with four-year constraints (Adaptive Strategy), while the bottom line compares the unconstrained strategy to the optimal strategy with two-year constraints (Adaptive Strategy).

Table 2: Annualized Mean Outperformance of Moving-Average-Based Strategy over Buy-And-Hold Strategy.

This table shows the per-annum mean outperformance of the moving-average filter-based strategies over the buy-and-hold strategy, across all subportfolios of the NCREIF NPI universe and across the entire 1978-2003 time period. The adaptive strategy differs from the non-adaptive strategy in that sell signals are ignored if after transaction costs they do not lead to a positive profit for a particular round-trip transaction. Both types of strategy are tested with and without the 4-year holding constraint. Results are also reported for the optimal strategy with the 2-year holding constraint (Non-Adaptive Strategy).

Transaction costs										
0	**1%**	**2 %**	**3%**	**4%**	**5%**	**6%**	**7%**	**8%**	**9%**	**10%**
Non-Adaptive Strategy										
0.03386***	0.03306***	0.03227***	0.03147***	0.03068***	0.02988***	0.02909	0.02829***	0.02750***	0.02670***	0.02591***
Adaptive Strategy										
0.02723***	0.02476***	0.02379***	0.02299***	0.02156***	0.02075***	0.02020***	0.01943***	0.01898***	0.01865***	0.01820***
Non-Adaptive Strategy, 4-Year Constraint										
0.02265***	0.02229***	0.02193***	0.02158***	0.02122***	0.02086***	0.02051***	0.02015***	0.01979***	0.01944***	0.01908***
Adaptive Strategy, 4-Year Constraint										
0.02158***	0.02137***	0.02105***	0.02067***	0.02031***	0.02020***	0.01983***	0.01937***	0.01903***	0.0188***	0.01849***
Non-Adaptive Strategy, 2-Year Constraint										
0.02561***	0.02508***	0.02456***	0.02403***	0.02351***	0.02299***	0.02246***	0.02194***	0.02141***	0.02089***	0.02037***
T-Statistics										
9.3144	9.0776	8.8385	8.5972	8.3539	8.1085	7.8612	7.6120	7.3610	7.1084	6.8541
7.8805	7.7009	7.5198	7.3374	7.1536	6.9686	6.7824	6.5950	6.4064	6.2168	6.0261

Average annualized return for buy-and-hold strategy: 0.1680

$^{\circ}$: significance level $\leq$ 10%. *: significance level $\leq$ 5%. **: significance level $\leq$ 1%. *** : significance level $\leq$ 0.1%.

The significance stars next to the outperformance values are for the hypothesis H_0: *mean outperformance of filter strategy over buy-and-hold strategy* $= 0$ against the positive one-sided alternative.

The two sets of t-statistics reported at the bottom of the table are for the hypothesis H_0: *mean outperformance of optimal unconstrained strategy (Non-Adaptive Strategy) over optimal constrained strategy (Non-Adaptive Strategy)* $= 0$, against the positive one-sided alternative. The top line compares the unconstrained strategy to the optimal strategy with four-year constraints (Non-Adaptive Strategy), while the bottom line compares the unconstrained strategy to the optimal strategy with two-year constraints (Non-Adaptive Strategy).

Table 3: Distributional Statistics for Holding Periods

This table shows distributional statistics for the holding periods dictated by the non-adaptive and adaptive trading-bands strategy, as well as the non-adaptive moving-average strategy (which always dominates the adaptive moving-average strategy). Holding periods for the adaptive strategy vary with the level of transaction costs, so the distributions for a variety of transaction cost levels are reported.

	1st Quartile	**Median**	**Mean**	**3rd Quartile**
Non-Adaptive Strategy, Trading Bands				
	3.00	12.00	17.00	31.00
Adaptive Strategy, Trading Bands				
Transaction Costs				
0	7.00	20.00	22.98	39.00
1%	8.00	21.00	23.85	39.00
2 %	9.00	22.00	24.45	40.00
3%	9.00	22.00	24.92	40.00
4%	9.00	23.00	25.31	40.00
5%	10.00	24.00	25.66	40.00
6%	10.00	24.00	25.87	41.00
7%	10.25	25.00	26.19	41.00
8%	11.00	25.00	26.51	42.00
9%	11.00	26.00	26.74	42.00
10%	12.00	26.00	26.91	42.00
Non-Adaptive Strategy, Moving Average				
	3.00	15.00	19.02	34.00

All figures in quarters.

Table 4: Summary Statistics for Property Dataset

Panel A: Total Number of Properties by Type and UPREIT Status of Holding Firm.

Property Type	**Non-UPREITs**	**UPREITs**	**Sum**
Multi-Family	138	1734	1872
Office	128	748	876
Industrial	170	995	1165
Shopping Center	52	480	532
Regional Mall	27	50	77
Retail: Other	191	233	424
Health Care	229	51	280
Hotel	35	238	273
Self-storage	23	21	44
Manufactured Home	0	60	60
Multi-use	2	14	16
Specialty	264	124	388
Sum	1259	4748	6007

Number of Unique Non-UPREITs: 40. Number of Unique UPREITs: 100.

Panel B: Property Distribution Throughout Top-10 Core-Based Statistical Areas (CBSAs).

CBSA Name	**Number of Properties**
New York-Northern New Jersey-Long Island, NY-NJ-PA	313
Dallas-Fort Worth-Arlington, TX	262
Los Angeles-Long Beach-Santa Ana, CA	261
Chicago-Naperville-Joliet, IL-IN-WI	252
Atlanta-Sandy Springs-Marietta, GA	219
Miami-Fort Lauderdale-Miami Beach, FL	212
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	196
Phoenix-Mesa-Scottsdale, AZ	167
Houston-Baytown-Sugar Land, TX	160
Washington-Arlington-Alexandria, DC-VA-MD-WV	130
Total in Top-10 MSAs	2172
Outside Top-10 MSAs	3835
Total	6007

Total number of unique CBSAs with at least one property covered by SNL: 497.

Table 4: Summary Statistics for Property Dataset (Continued).

Panel C: Total Number of Properties by Size and UPREIT Status of Holding Firm.

Size	**Non-UPREITs**	**UPREITs**	**Sum**
Small	1166	2503	3669
Large	93	2245	2338
Sum	1259	4748	6007

Number of Unique Small Firms: 126. Number of Unique Large Firms: 40.

Table 5: Distributional Statistics for Holding Periods of Properties owned by all REITs, Non-UPREITs, and UPREITs.

This table shows distributional statistics for holding periods of properties in REITs' portfolios, for properties owned by all REITs, by non-UPREITs only, and by UPREITs only.

	1st Quart	**Median**	**Mean**	**3rd Quart**	**N**
All REITs	2.121	4.000	5.532	6.825	8918
Non-UPREITs	2.804	5.003	6.930	9.316	2047
UPREITs	1.992	3.836	5.115	6.323	6871

Kolmogorov Smirnov test of the hypothesis that the CDF of holding periods for UPREITs is the same as that for regular REITs, against the one-sided alternative (UPREIT CDF lies above Non-UPREIT CDF): $D = 0.1974$, p-value $< 2.2 \times 10^{-16}$.

All figures in years.

Table 6: Regression Results, Ordinary Least Squares Regression, Modeling Log of Holding Period on Local Market Performance.

Dependent Variable: Natural Log of Holding Period. This table shows regression results for an ordinary least squares regression modeling the holding period of a property in a REIT's portfolio on local market performance over the time period the property was held in the portfolio, and the UPREIT status of the firm at the time the property was sold. $UPREIT$ is a dummy variable taking a value of 1 if a firm was declared an UPREIT at the time the property was sold and 0 otherwise. A test of whether the total interaction effect for UPREITs is distinguishable from 0 is included at the bottom of the table. Standard errors in parentheses.

	Model 1	Std. Error
$Intercept$	-0.0275	(0.1022)
$pareturn$	7.9262	$(0.8153)^{***}$
$UPREIT$	0.5502	$(0.1101)^{***}$
$pareturn \cdot UPREIT$	-5.3939	$(0.8784)^{***}$
N	4242	
$\overline{R^2}$	0.0602	
Model F	52.26^{***}	
F	59.09^{***}	

$^{\circ}$: significance level $\leq 10\%$. *: significance level $\leq 5\%$.
**: significance level $\leq 1\%$. *** : significance level $\leq 0.1\%$.

$pareturn$: Annualized return of the local property market index.
$UPREIT$: Dummy variable indicating whether the transacting firm is an UPREIT.

F: Test statistic for the hypothesis that $pareturn + pareturn \cdot UPREIT = 0$.

Note: only properties with holding periods of 6 years or less are included.

Table 7: Regression Results, Probit Regression Modeling the Decision to Sell a Property Before Four Years.

Dependent Variable: $fourplus$. The table shows regression results for probit regressions modeling the decision to hold a property beyond four years, on local market conditions, the selling firm's UPREIT status, and the selling firm's size. $fourplus$ is a dummy variable, which takes a value of 1 if a property was held in a REIT's portfolio for four years or longer and a value of 0 otherwise. A test of whether the total interaction effect for UPREITs is distinguishable from 0 is included for both models at the bottom of the table. Model 3 includes dummy variables for the 12 property types. Standard errors in parentheses.

	Model 1	Std. Error	Model 2	Std. Error	Model 3	Std. Error
$Intercept$	-1.2240	$(0.1827)^{***}$	-1.1980	$(0.1824)^{***}$	-0.9801	$(0.2503)^{***}$
$pareturn$	7.2280	$(1.4386)^{***}$	7.1642	$(1.4376)^{***}$	7.1643	$(1.4871)^{***}$
$UPREIT$	0.4050	$(0.1972)^{*}$	0.7625	$(0.2382)^{***}$	0.7033	$(0.2239)^{**}$
$pareturn \cdot UPREIT$	-5.5364	$(1.5575)^{***}$	-7.4349	$(1.6811)^{***}$	-7.3862	$(1.7310)^{***}$
$Large$			-0.6471	$(0.2130)^{***}$	-0.5890	$(0.1583)^{***}$
$pareturn \cdot Large$			3.2025	$(1.1684)^{**}$	2.7079	$(1.2737)^{*}$
Property-Type Dummies?	No		No		Yes	
N	6212		4242			
Model χ^2	498.91^{***}		522.00^{***}		532.8^{***}	
χ^2	8.0297^{**}		0.0916		0.0595	

$^{\circ}$: significance level $\leq 10\%$. *: significance level $\leq 5\%$.
**: significance level $\leq 1\%$. *** : significance level $\leq 0.1\%$.

$pareturn$: Annualized return of the local property market index.
$UPREIT$: Dummy variable indicating whether the transacting firm is an UPREIT.
$Large$: Dummy variable indicating whether the transacting firm is large.

χ^2: Test statistic for the hypothesis that $pareturn + pareturn \cdot UPREIT = 0$.

Note: only properties with holding periods of 6 years or less are included.