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

Tobias Mühlhofer
McCombs School of Business
University of Texas *

November 24, 2014

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. I first simulate a set of filter-based market timing strategies, which outperform a buy-and-hold strategy out of sample, and show that imposing a four-year (or even the newer two-year) holding constraint significantly reduces the excess returns the strategies generate. 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 that the trades generated by the filter strategy generally resemble actual REIT trading activity, validating the relevance of the simulation results. A direct test for the constraint reveals that REITs' propensity to hold a property beyond the minimum period increases, the higher the profit from the transaction, consistent with the asymmetric nature in which the rule is enforced. By contrast, this effect is insignificant for UPREITs, which are not as affected by the constraint. I further show that UPREITs overall achieve significantly better ex-post market timing performance than Non-UPREITs. I thus find that overall REITs are limited by the *dealer rule*.

*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: tobias.muhlhofer@gmail.com** Contact: Finance Department, 2110 Speedway, Stop B6600, Austin, TX 78712, 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. 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 regulatory 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 classified as a REIT and thus be exempt from corporate income tax, among other things, its portfolio strategy is mandated to be *passive*. This constraint is known as the *dealer rule*, as the rule essentially prevents a REIT from being a dealer in property. The rule imposes selling constraints for the properties in REITs' 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.[1] 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 also shown in Mühlhofer (2013), who argues that REIT investors do not participate in short-term property appreciation profits because of this.

For example, the privatization of Equity Office Properties (EOP), followed immediately by the

[1] 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.

sale of over 50% of the firm's portfolio presents anecdotal evidence in favor of this constraint's bindingness. Blackstone offered a substantial premium to shareholders for the share buyback, and then made a large profit from the sale of the portfolio. While many other factors were likely involved in a complex transaction such as this one, the fact remains that this could not have been accomplished by maintaining EOP's REIT status. The premia offered to shareholders, together with the subsequent profits that Blackstone made on the sale of the portfolio indicate that value was being destroyed by the REIT structure in this case.

This paper contributes to the literature by examining the bindingness of the minimum holding constraint, and illustrating the way in which REITs forgo 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 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. Subsequently, I 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 about one and a half times 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. I show that sufficient market liquidity exists during most of the sample (with only the exception of the 2008–2010 financial crisis) to make this strategy feasible. If I do not allow trading during the illiquid time of the crisis, the results are qualitatively unaltered, in that the holding constraints still impose a large and significant performance cost.

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. However, I do present evidence that the trading activity generated by this strategy is very similar to actual trading behavior of any unconstrained REITs in the market. Therefore, this suggests that the costs imposed by the holding constraints are actually borne by REIT investors in the real market. Furthermore, this part of the investigation also offers insights about the mechanics of how a minimum holding constraint hinders optimal market timing.

I then analyze empirically the trading behavior of REITs, by examining REIT property transactions data in order to directly assess the bindingness of the holding constraint to market timing. I first document that REITs indeed actively re-allocate capital across sub-markets to realize timing profits, by showing that sales in a particular sub-market are rarely followed closely by new purchases in the same market. On the topic of holding periods, I find that there seems to be considerable demand for short property holding periods by REITs, which, as stated earlier, exhibit similar trading behavior to the trading patterns generated by the simulated trading strategy.

I use a natural laboratory in this portion of the investigation, 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 gives an UPREIT the ability to easily conduct Section-1031 like-kind exchanges when making property acquisitions. This ability enables an UPREIT to mostly overcome the minimum holding period constraint, in that in a 1031 exchange the REIT essentially inherits the contributing partner's holding period.[2] 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 and even two years, and statistically significantly more frequent trades by UPREITs than by regular REITs. This implies that those firms that can trade more often, do so. I further find that a significantly higher number of UPREITs locate in markets in which my simulated trading strategy generates average holding periods of less than four years, than locate in an average market. Conversely, a significantly higher number of Non-UPREITs locate in markets in which the trading strategy generates average holding periods of more than four years, than locate in an average market. This suggests that UPREITs capitalize on their competitive advantage in turning over properties quickly, and can therefore outbid Non-UPREITs in markets that require this. Further, this helps establish another link between my simulated trading strategy and actual REIT trading behavior.

I then test the specific effect that local market performance has on the conditional decision to hold a property beyond the minimum time: 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.[3] In a falling market, on the other hand, 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. By contrast, if the constraint were not perceived as binding (for example if four or even two years were a shorter holding time

[2]The details of this are discussed in Section 3.

[3]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.

than any REIT would normally employ), market performance should have no effect on the decision to hold a property beyond the minimum, 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. I find that regular REITs are significantly more likely to hold a property beyond the minimum, the higher the local market return (and thus the more of a profit they are likely making), while UPREITs are statistically indifferent to the minimum holding periods. This result also holds when I account for possible systematic effects associated with firm size or different property types.

Finally, I document firms' ex-post success in timing markets, by computing, for each property sale, the two-year local market returns following this sale. I find that UPREITs show significantly better market timing performance than regular REITs. To rule out firm sophistication as an alternative explanation, I conduct the same test comparing small and large firms, and find that the timing performance of these two groups is statistically indistinguishable from each other.

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), Gyourko and Keim (1992), Cooper, Downs and Patterson (1999), treat this issue for publicly traded securitized real estate, while Case and Shiller (1990) and Case and Quigley (1991) treat this subject for housing markets. Brounen, Eichholtz and Ling (2007), Hartzell, Sun and Titman (2014), as well as Hochberg and Mühlhofer (2014), examine active trading performance by Real Estate portfolio managers and are therefore also related. While Brounen et al. (2007) find no relationship between the amount a manager is trading and the firm's return in public equity markets, Hartzell et al. (2014) find evidence in favor of a market timing hypothesis. Hochberg and Mühlhofer (2014) find ample evidence of value created by managers who actively re-allocate funds across sub-markets.

Most closely related to this portion of this study is Geltner and Mei (1995) who use technical trading strategies to illustrate that market timing profits can be made in the property market.[4] 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. In this study, I use the previously-documented recognition of the market's predictability and treat the subject of the trading constraints in this light, to illustrate how REITs are hindered by these abnormal profits. As such, my focus is not primarily on the fact that an active strategy outperforms the buy-and-hold strategy, but rather on the shortfall in outperformance and therefore the cost to investors induced by the trading constraints. Since these trading constraints and their effects seem to have been overlooked in previous literature, there is actually no direct parallel to this study.

For the second part of the study, 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, there are only two other studies I am aware of which analyze trading behavior by REITs in the direct property market. Crane and Hartzell (2007) do this in order to test for a disposition effect. Hochberg and Mühlhofer (2014) do this in order to gauge managerial value-added. 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. This study (together with the concurrent study by Mühlhofer (2013)) is first in the literature to systematically document the effects of this regulatory framework and the advantage of the UPREIT structure in relation to this. Mühlhofer (2013) examines the effect of the *dealer rule* and the UPREIT structure on the returns obtained by investors in a factor-modeling framework and argues that the constraint is part of the answer to the question about what drives

[4] Wheaton, Baranski and Templeton (2009), find low long-term returns for commercial property in Manhattan over the last century. However, they find substantial short-term fluctuations in these property values. My study concerns itself with capitalizing upon these.

a wedge between REIT returns and direct-property returns in the short run. The current study, on the other hand, takes a more direct approach, examining trading behavior of REITs in the direct market (both simulated and actual) and illustrating to what extent the *dealer rule* interferes with strategies for creating value through active trading in a market that should allow informed managers to do so.

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 trading behavior; Section 4 concludes.

2 A Filter-Based Trading Strategy

2.1 The Trading Strategy

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.[5] One of the advantages of this methodology over others that have been employed by past studies (such as, for example that of Geltner and Mei (1995)) is that it requires very little data and is computationally simple. As implemented in this study, the strategy I propose requires only four past quarterly price observations and no data about other markets. This enables me to make the most out the property returns data available, especially for small sub-markets. Further, as I show in Section 3, this strategy seems to approximate actual trading behavior by REITs in the property market.

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

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

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

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 through* $map_{K,t}$. This is illustrated by the relationship between the solid line and the dot-dash alternating line in Figure 1. 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.

This strategy can be improved upon, by the use of *trading bands* around the moving average. 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, this will create a better filter, as the price will cut the respective trading band sooner after a turnaround, than it will the moving average. The dashed lines in Figure 1 illustrate this concept.[6]

Bollinger (2001) determines the value of the bandwidth parameter s by using a multiple of the K-period moving standard deviation, defined as

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

making the trading bands

$$B_{+,t} = map_{K,t} + \alpha s_{K,t} \quad (6)$$

$$B_{-,t} = map_{K,t} - \alpha s_{K,t} \quad (7)$$

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

[6]It should be noted that the effectiveness of such a strategy pivots crucially on the fact that the price process for property is mean-reverting.

through the moving average. Figure 1 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 1, 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,[7] 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-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. However, as I show in Section 3, the trades generated by this strategy do seem to approximate the actual trading behavior of unconstrained REITs. The purpose of this section is then to demonstrate that the minimum holding constraints significantly reduce the ability of these strategies to generate abnormal profits, and to therefore show directly the cost imposed by this constraint. 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. I have no way of testing explicitly whether these strategies are actually applied by REITs. However, the empirical distribution of holding periods by REITs shown in Table 6 as well as the individual quarterly aggregate trade data and the data on REIT presence in CBSAs according to the transaction speed required by these markets does suggest that these strategies generate similar trading behavior to what REITs actually exhibit. This 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

[7]See Section 2.3, for more information on the data set.

different property portfolios, testing its performance in various transaction cost environments.

2.2 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.[8] 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.[9] 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 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.[10] 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

[8] See Clayton, Geltner and Hamilton (2001), for example.

[9] 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.

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

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.[11]

A further issue that should be mentioned here is market liquidity.[12] 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 market.[13] 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. In addition to these economic considerations, I also conduct an empirical examination of this issue throughout this section, by constructing a liquidity measure and comparing actual market liquidity to the number

[11] 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.

[12] Lin and Vandell (2007), as well as Kawaguchi, Sa-Aadu and Shilling (2007), for example, show that liquidity risk is priced in property markets and discuss the issues associated with this at length.

[13] Should this liquidity not exist at all, the strategy automatically becomes buy-and-hold. For any liquidity level in between perfect liquidity and perfect illiquidity, 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.

of required sale transactions generated by the trading strategies.

2.3 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 1153 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 2010. 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), which I adapt for this setting.[14]

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

$$P_t^* = b_1 P_{t-1}^* + b_2 P_{t-4}^* + w_0 P_t \quad (8)$$

$$r_t^* = b_1 r_{t-1}^* + b_2 r_{t-4}^* + w_0 r_t \quad (9)$$

$$r_t = \alpha + \rho r_{t-1} + \epsilon_t, \quad \text{with } E(\epsilon_t) = 0,\ \sigma^2(\epsilon) = \sigma^2,\ |\rho| < 1, \text{ and } \alpha \neq 0 \quad (10)$$

$$r_t^* - \rho r_{t-1}^* = \alpha w_0 + b_1 \left(r_{t-1}^* - \rho r_{t-2}^*\right) + b_2 \left(r_{t-4}^* - \rho r_{t-5}^*\right) + \epsilon_t', \quad \text{where } \epsilon_t' = w_0 \epsilon_t \quad (11)$$

$$w_0 = 1 - b_1 - b_2 \quad (12)$$

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

[14] 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 TBI.

and 9 describe 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.[15] 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 de-smoothed 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

[15] 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.

and previous returns plus one.

I test the the filter-based strategies 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.[16] At 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 annualized 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, after 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

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

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.

To directly address the issue of liquidity and examine whether problems in this respect could significantly hinder the implementation and outperformance of the active trading strategies, I estimate an empirical measure of liquidity for the dataset in each quarter. Most measures of liquidity in the commercial-property literature have been based on the methodology of Fisher, Gatzlaff, Geltner and Haurin (2003, 2004) and have thus been based primarily on trading volume within the commercial property market. However, as is discussed frequently in the general finance literature (for example Goyenko, Holden and Trzcinka (2009)), an illiquid market is characterized by trades that have a large price impact, while a liquid one is characterized by trades that do not. Therefore, while trading volume often proxies positively for liquidity, a more accurate measure would measure the price impact that such volume has. I therefore apply the Amihud Illiquidity Measure ($ILLIQ$), based on the methodology of Amihud (2002), to the commercial property market.[17]

This measure is defined as follows:

$$ILLIQ_t = \frac{|ret_t|}{volume_t} \tag{14}$$

In this notation ret_t is the return, or percent price change, of the asset during time period t and $volume_t$ is the trading volume during the same amount of time. A high price change (in either direction) with low trading volume implies a high amount of illiquidity (i.e. a low amount of liquidity), while a lower amount of price change associated with higher volume implies a low amount of illiquidity (i.e. a high amount of liquidity). For the numerator, I use the total return to the NCREIF national NPI. I define $volume_t$ as square footage transacted within the NCREIF

[17]To my knowledge, this study is the first to do this.

universe (i.e. square footage bought plus square footage sold) in quarter t divided by the total square footage of the NCREIF portfolio at the end of that quarter. Thus, $volume_t$ indicates the fraction of the NCREIF portfolio transacted in a quarter. Here, as in the original approach of Amihud (2002), the overall magnitude of this measure is not very meaningful; rather the relative fluctuations in $ILLIQ$ from quarter to quarter become important.

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.

2.4 Results and Implications

Figure 2 shows actual sales transactions from the trading-band filter strategy, compared with quarterly liquidity figures. Specifically, the top portion of the graph shows the fraction of the 1153 NCREIF sub markets in which a sell signal was generated in a particular quarter, while the bottom portion shows the market-wide Amihud Illiquidity Measure ($ILLIQ$) in the same quarter. The top portion shows relatively high amounts of sales during the late 1980s and early 1990s and during the early 2000s consistent with the market downturn experienced during those times. Further, during the crisis of 2008 to 2010, the graph shows extreme amounts of sale signals generated.

Comparing this with illiquidity, one can see that in most cases when there are relatively high amounts of sales required, illiquidity is low, which means that markets would accommodate such sales. In fact, the Spearman correlation between the fraction of markets that generate sell signals and the measure of illiquidity is -0.3318, which indicates that in times of relatively high amounts of sales required there are relatively low amounts of illiquidity (i.e. markets are fairly liquid). Thus, overall, liquidity should not be a concern for the practical implementability of the strategy suggested.

The only possible exception to this is the financial crisis of 2008 to 2010. In fact during the final quarter of 2008, the first two quarters of 2009, as well as the second quarter of 2010, I find extremely high levels of illiquidity. I therefore estimate a version of the trading strategies (both in the unconstrained and constrained setting) in which I allow no trading what so ever during these quarters, ignoring all trading signals generated during this time period. I show that even in this extreme setting of zero liquidity during this time (which was probably not the case in reality) my results fully hold.

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 the four-year and two-year minimum holding period under which REITs operated. 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.

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 7.65 percentage points (765 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 non-adaptive strategy to study the effect the new two-year minimum holding period would have, as that is the dominant strategy in this setting. The results from the adaptive strategy are omitted from the table to save space, but are available upon request.

The final two lines show the performance of the filter strategies, allowing no trading what so ever in the quarters of extreme market illiquidity of the latest financial crisis. In this setting, the outperformance is somewhat lower than without this adjustment, which illustrates that an important part of the outperformance of this strategy would have come from being out of the markets during this downturn. However, even in this setting the filter-based strategy without minimum holding period constraints significantly outperforms the buy-and-hold strategy, under all levels of transaction costs. In order to save space I have reported only the results from the optimal strategy with two-year constraints in this setting.[18] Even without allowing trading during times of extreme illiquidity, the strategy with two-year constraints significantly outperforms a buy-and-hold strategy, although the point estimates show that it does so to a lesser extent than the strategy without holding constraints, and that therefore this constraint is still costly.

The first two lines of t-statistics at the bottom of the table are results of the hypothesis test of zero outperformance of the optimal unconstrained strategy, over the optimal strategy under the four-year and two-year constraints, respectively. The third line shows results of this test with respect to the strategy without holding constraints but with no trading in illiquid quarters, over the strategy with two-year constraints and no trading in illiquid quarters. 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 170 and 227 basis points per year, depending on the level of transaction costs. The relevant t-statistics at the bottom of the table lie between 5.34 and 6.47, 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, and therefore a significant cost. 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 102 to 136 basis points per year), the t-statistics at the bottom of the table lie between 3.37 and 4.22, and therefore still indicate

[18]Since this strategy (understandably) always outperforms the strategy with four-year constraints, this constitutes the toughest hurdle with regards to showing the cost of the holding constraints in this setting, and therefore this should be the most stringent test. The results for all possible strategies in this setting are available upon request.

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. Even without allowing trades during the illiquid quarters of 2008–2010, the strategy without holding constraints outperforms the best-case strategy with holding constraints by 76 to 103 basis points, with t-statistics of 3.42 to 4.37, indicating a highly significant difference and therefore cost to investors.[19]

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 20 to 80 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 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,[20] the difference is highly statistically significant at all levels of transaction costs, with t-statistics of 5.34 to 6.47. 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 102 to 136 basis points, thus of similar size as in the previous table. Also, the shortfall in performance is statistically significant, with t-statistics of 3.37 to 4.22. The results are analogous, even when not allowing trading during the illiquid quarters of 2008–2010. 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 and therefore a significant cost to investors.

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

[19]Once again, I present the most stringent test (for the two-year constraint) in this setup in the table. The results for the four-year constraint show a larger gap and are available upon request.

[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.

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 (2013).

I have thus demonstrated that at any level of transaction costs both the old four-year as well as the 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, as I show in the next section do resemble the trading behavior of REITs). 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 always 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 two 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 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

[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.

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 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 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 two time periods; one of these goes 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. The second spans from the last quarter of 2008 until the middle of 2010, and therefore coincides with the latest financial 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 (liquidity allowing) 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's investing in treasuries rather than property, or the REIT's downsizing (thereby also cutting overheads), and the individual investor's 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. As I show in the next section, however, the overall trading activity generated by these strategies does resemble the actual trades that we see REITs make in the market. I have thus 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 them significantly reduce their ability to do so and therefore constitute a significant cost to investors.

3 Actual Holding Periods

3.1 The Empirical Setup

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. 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.

As mentioned 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.

Specifically, in an UPREIT, the REIT does not hold properties directly, but rather holds shares in a limited partnership (the operating partnership), which then in turn holds the 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 laboratory. 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.[22] [23]

Due to the binary nature of the constraint (i.e. the safe harbor rule applies exactly at four years and longer for sales made before July 30, 2008, and exactly at two years and longer for sales made after that) I estimate a probit regression which precisely models the specific decision to hold a property beyond the minimum holding period, based on local market return. This test is structured around the asymmetric nature of the enforcement of the *dealer rule*. The penalty for engaging in a *prohibited transaction* is a 100% gains tax (i.e. having to surrender all profits). This would only be

[22]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.

[23]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.

binding, if a profit was actually made, and I argue that it would be more binding, the higher the profit, as a higher profit would be more painful to give up than a lower one. Therefore, if a firm is bound by the constraint, a positive return on the property will cause the firm to want to hold the property beyond the minimum, even if this sub-optimal for market-timing purposes. If it is not bound, there should be no statistical relationship, as it would simply attempt to apply optimal market timing considerations. This would mean that some times the property would be sold earlier and some times later. Under a null hypothesis that no type of REIT is bound by the minimum holding constraint of the safe harbor rule, one would observe that local market conditions have no effect on the decision to hold a property beyond this time for either regular REITs or UPREITs, as optimal holding decisions should be unrelated to the return being made on a transaction. Under the alternative hypothesis, one should observe that a regular REIT is more likely to hold a property beyond this time, the higher the local market return (and hence the higher the profit made on the transaction), because the firm does not want to give up its capital gains. UPREITs should be indifferent to the constraint, and therefore one should see no type of statistical relationship between holding period and return. I then also test for differences in the ex-post effectiveness of realizing timing profits, between UPREITs and Non-UPREITs.

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] [25] Note that, since I examine the effect of holding constraints on 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 the end of 2009) and not properties which were still in a REIT's portfolio at the

[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] The holding period is defined as the time elapsed between purchase date and sale date. For UPREITs, properties purchased into the Operating Partnership, are still defined this way (i.e. as purchased the day they are added to the UPREIT's OP). For mergers, the property portfolio acquired through the acquisition of a REIT counts as purchased on the merger date by the acquiring firm, with respect to the dealer rule.

time the database's coverage ends.[26]

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.[27] Correspondingly, I define the dummy $UPREIT$ as

$$UPREIT_i = \begin{cases} 1 & \text{if a REIT was an UPREIT at time of sale} \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

Table 4 presents summary statistics for this data. Panel A shows the number of properties of each major type, with a distinction by whether the holding firm is an UPREIT. Note that there are many more properties held by UPREITs than by non-UPREITs and that this is the case for all property types. Note also that the dataset contains only 44 non-UPREITs, compared with 98 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 528 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 no properties sold by large 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 CBSA-level subindices of the National Property Index (NPI) from NCREIF. The appreciation

[26]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.

[27]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.

components are desmoothed, as discussed in Section 2.3. 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.[28]

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.

As the dependent variable, I define a binary variable $min.holding_{i,l}$ which shows whether a property was held beyond its minimum holding period for the safe harbor rule. This variable is equal to one for properties sold before July 30, 2008, that were held for four years or longer and for properties sold after July 30, 2008, that were held for two years or longer. The variable is equal to zero for properties whose holding period does not satisfy the minimum holding constraint.

Lastly, I define a dummy variable $crisis_i$, which is equal to one for properties sold during the year 2009 and zero otherwise. This allows some flexibility in the model, for the turbulent market condition of the last financial crisis. These conditions may have led to problems in liquidity or other types of issues which may have changed the decision-making process or the opportunity set of REIT managers.

The actual model I estimate thus becomes

$$\begin{aligned} min.holding_{i,l} &= \alpha + \beta_1 pareturn_l + \beta_2 UPREIT_i \\ &\quad + \beta_3 pareturn_l \cdot UPREIT_i + \epsilon_i \end{aligned} \tag{16}$$

[28]As mentioned earlier, direct data on actual transaction returns and intermediate cash flows is very scarcely available.

As a robustness check, I further estimate a version of this model (Equation 16) which includes firm size as well as UPREIT status together with the interaction effects of both characteristics. 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. Controlling for size also helps rule out an alternative explanation that UPREIT status only proxies for firm sophistication, which size would also do. Lastly, 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:

$$\begin{aligned} min.holding_{i,l} &= \alpha + \beta_1 pareturn_l + \beta_2 UPREIT_i \\ &\quad + \beta_3 pareturn_l \cdot UPREIT_i \\ &\quad + \beta_4 Large + \beta_5 pareturn_l \cdot Large_i + \epsilon_i \end{aligned} \tag{17}$$

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.[29] 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 (or later two-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.[30] On first glance, it might seem beneficial to add a lower-bound limitation on the holding periods of properties used for this test. However, limiting the lower end of holding periods would actually eliminate short-holding-period properties which might have been sold for a profit, and so might bias results in my favor. Therefore not applying a lower bound to this test seems like a more

[29] This variable definition is analogous to that used in the portfolio sort in Mühlhofer (2013).

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

conservative approach.

3.3 Results and Implications

I begin by giving evidence that actual REIT trading behavior reflects a pattern consistent with reallocating capital across sub-markets in order to time these, as I have assumed in the setup for the simulation in the previous section. To do this I test how often firms engage in *repeat buys* in the same CBSA. That is, I test how often a property sale by a particular firm in a particular CBSA is followed by a purchase by the same firm in the same market. If these *repeat buys* occurred often, with short time gaps between the sale and the subsequent purchase, this would indicate that firms are not pursuing a timing strategy across sub-markets. For this test, I have a total of $18,822$ transactions, counting both buys and sells. The number of *repeat buys* is 477, or about 2.5%. This indicates that firms are indeed mostly re-allocating capital across sub-markets when selling, which is consistent with REITs' engaging in timing across sectors. I also find that conditional upon engaging in a *repeat buy*, the median gap between the sale and the new purchase in the same market is 0.48 years, or a little under six months, and the mean gap is 0.94 years, or a little over 11 months. This suggests that even in the small fraction of transactions that constitute *repeat buys*, firms are not quickly swapping one property for a very similar one *next door*.

Table 5 presents further evidence on the empirical linkage between the filter-based trading strategies discussed in Section 2 and the actual trading behavior by REITs observed from SNL. For Panel A, I compute in each quarter the fraction of submarket portfolios sold by the simulated trading strategy, as well as the fraction of available properties sold by both unconstrained and constrained REITs. Sales by unconstrained firms are all sales by UPREITs, as well as sales by Non-UPREITs of properties they have held beyond four years. Sales by constrained firms are sales by Non-UPREITs of properties they have held for less than four years. The table shows correlations between the fraction of sales by the filter trading strategy and the fraction of sales by the respective firm.

The quarterly fraction of sales by unconstrained REITs shows a positive correlation of 0.3186, significant at the 5% level, with the sales generated by the filter-based trading strategy. The correlation between the sales activity of constrained firms and that generated by the simulated timing strategy is insignificant. This constitutes positive evidence that the trading strategy generated in Section 2 actually models the trading behavior of firms in the market. The property sale activity on which there are no constraints resembles that of the market timing strategy simulated earlier. This should help validate the relevance of the results previously presented, with regards to the simulation. The results suggest that the cost of the minimum holding period constraints to the simulated strategy should translate to a large extent to actual firms in the market. It is further consistent with a hypothesis of a binding constraint, that the sales activity of constrained firms does not resemble that of a market-timing strategy.

Panel B of Table 5 examines the choice by UPREITs and Non-UPREITs to locate in certain markets. If UPREITs have a competitive advantage over regular REITs in terms of profitably holding properties for a short time, they should tend to locate in markets in which the filter-based strategy generates many holding periods shorter than four years (assuming they are following a similar strategy). Conversely, one should see regular REITs retreat to markets which do not require holding periods this short.

The Panel shows, in the first line, that there are, on average 29.48 UPREITs located in a Short-Holding-Period CBSA (a CBSA for which the filter-based trading strategy generates, on average, holding periods that are shorter than four years). Across all CBSAs, there are only an average of 15.02 UPREITs per CBSA. The difference is statistically significant, at the 10% level. Consistently with the hypothesis stated above, the table also shows an average of 8.64 Non-UPREITs in each Long-Holding-Period CBSA (a CBSA for which the filter-based trading strategy generates, on average, holding periods that are longer than four years), while in the average CBSA overall there are only 2.47 Non-UPREITs. The difference is statistically significant at the 1% level. Thus, this table shows that both UPREITs and Non-UPREITs behave exactly as one would expect, to work with their competitive advantages or disadvantages, if they were pursuing trading activity similar to what is generated by the filter-based strategy simulated in the previous section. This evidence should create a strong link between the two sections of this study and thus validate the results concerning the costs to investors shown to be caused by the trading constraints with respect to the active timing strategies in the previous section.

Table 6 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 4.123 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 binding, in the sense that REITs seem to demand such short holding periods in their portfolios, and a four-year or two-year holding constraint therefore reduces their ability to engage in something that they clearly feel the need to do. In fact, REITs seem very eager to dispose of many properties after a fairly short period of time, to such an extent that actually almost half of all properties are sold before the four-year mark. With respect to a two-year holding constraint, nearly one quarter of properties is sold within this time, so even this lower minimum is relevant. 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 6, the median holding period for non-UPREITs is 4.958 years, while the median holding period for UPREITs is 4 years and the first quartiles differ by about 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 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.[31] Furthermore, recall that the four-year cutoff merely constitutes a safe harbor (see Footnote 1). 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. Further, Non-UPREITs can also help themselves through their limited ability to conduct 1031 exchanges, in order to sell some properties early. In any case, however, (recalling the results from Table 5) the trading activity of properties sold before the constraint is met does not exhibit consistent market-timing patterns.

These patterns suggest that the fraction of short holding periods by Non-UPREITs that is visible should be interpreted as evidence that these firms also see the need to transact quickly, and that they do so whenever they can. Having a minimum holding period in place thus interferes with an activity that firms seem to otherwise find necessary, and therefore this constitutes a binding constraint. The probit setup that follows now investigates this issue by taking into account whether a profit was likely made on a sales transaction and shows that, consistent with economic rationale, the constraint binds when this was the case.

Table 7 shows the results from the probit regressions, which model the specific decision to hold a property beyond the minimum holding period on the annualized local market return and firm characteristics. In an unconstrained world, annualized return should be unrelated to holding period (or a binary representation of a certain holding window, as is used here). The constraint binds if a higher annualized local return is more likely to make the firm hold a property beyond the arbitrary cutoff (as this is necessary to retain profits). Therefore it is necessary to investigate the overall effect of *pareturn* (the annualized local-market return over the life of the property holding)

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

for both types of firm, and to see whether this is statistically positive. For Non-UPREITs that effect is simply the sign and significance of the coefficient on *pareturn* itself. For UPREITs, this is given by the overall effect of local market return for this type of firms, which is the base effect of *pareturn*, plus the marginal effect of *pareturn* for UPREITs, or $pareturn + pareturn \cdot UPREIT$. This overall effect for UPREITs is tested through the χ^2 statistic at the bottom of the table. If this is insignificant, then the overall effect of *pareturn* on whether an UPREIT holds a property beyond the minimum is zero, which implies that the UPREIT is unconstrained.

In Model 1, the coefficient for *pareturn* is positive and significant at the 5% level, indicating that the overall probability that a property which is held for six years or less is held beyond the all-important minimum constraint is significantly higher in a rising market. This means that a positive profit on a property will be associated with a regular REIT's to waiting until the constraint is fulfilled, so profits do not have to be surrendered. This implies that for regular REITs the minimum holding constraint is binding. The coefficient for $UPREIT$ is insignificant. However, the fact that the UPREIT dummy is insignificant simply indicates that *holding returns constant*, UPREITs do not have a different likelihood to meet a holding period constraint than regular REITs do. However, this is not really relevant to what I am testing, as the constraint should only bind when a profit was made and not independently of *pareturn*.

The coefficient for $pareturn \cdot UPREIT$ is insignificant, although the point estimate is negative. More importantly, however, the χ^2 statistic testing the hypothesis that the total effect of $pareturn + pareturn \cdot UPREIT$ is zero, presented at the bottom of the table does not reject. This means that there is no statistical relationship between the decision of an UPREIT manager to hold a property beyond the minimum and the overall effect of local market returns. This indicates that UPREITs are indifferent to the constraint, as they do not have to fear giving up positive profits by selling early. These results are qualitatively unaltered when accounting for firm size (as in Model 2) and also for property type (as in Model 3). The result remains consistent throughout, that regular REITs are bound by the minimum holding constraint in a rising market, while UPREITs are indifferent to this, as indicated by the consistently insignificant χ^2 statistic. The small Pseudo R^2 should not be a cause for concern in this case. There are a number of other factors which, in the cross-section, determine the decision to sell a property and these regressions do not attempt to fully model this decision. Rather, there merely needs to be a statistical linkage established for regular REITs, which demonstrates that in a rising market the decision is made to hold a property beyond the minimum holding period, in order to show that the constraint binds for these firms. The lack of a linkage for UPREITs shows that for these firms the constraint is irrelevant; unrelated to market conditions, UPREITs hold some properties beyond the four-year (or two-year) mark and not others, as optimal strategies would dictate.

It should be remembered that Non-UPREITs also have a limited ability to sell before the minimum holding constraint is fulfilled and retain profits, firstly through the *facts and circumstances rule* and secondly through the limited 1031 exchanges which Non-UPREITs can make. Both will be made to the extent possible, in order to maximize the profits of the particular trade being made. If the limited 1031 exchanges, as well as the facts and circumstances method yielded a world in which this constraint did not matter, then the regression would show this. Instead, the analysis shows that, systematically, the constraint matters and that Non-UPREITs cannot satisfactorily overcome this to make them look like they are not bound. Only UPREIT status does this.

As a final result in this section, I present evidence of the ex-post market-timing performance of REITs with different characteristics. In order to do this, I compute for each property sold by a REIT the total return to the CBSA sub-market to which that property belongs, over the two years following the sale. Table 8 presents summary statistics for these figures. The mean return over two years following a property sale for all firms is 3.96%, while the median is 5.07%. While this indicates that in hindsight REITs are not timing optimally, in that they could earn more if they held properties for two more years, the cross-sectional differences among firm types should be more instructive in this case. For sales by UPREITs, the mean (median) two-year return after a sale is 3.25% (3.95%); for Non-UPREITs on the other hand, the mean (median) return over this time period is 8.17% (11.40%). The t-test for the hypothesis that the true difference in means between the two firm types is zero is strongly rejected (t-value of -10.67). It is therefore the case that UPREITs also realize significantly better ex-post market timing performance than Non-UPREITs. This is consistent with UPREITs' being less hindered by the trading constraints, and therefore being able to better follow market-timing strategies. Non-UPREITs, on the other hand operate in a constrained environment which leads them to making sub-optimal trade decisions and *leave more money on the table.*

To rule out an alternative explanation that UPREITs are simply better market timers because they possess a higher level of firm sophistication, I conduct the same test for Large versus Small firms. Here, too, Large firms should be more sophisticated than Small firms; if sophistication were the explanation for this difference, then we should see Large firms time better than Small ones. The data shows that the mean (median) two-year return after a sale for Large firms is 3.75% (3.34%), while that for Small firms is 4.00% (5.47%). The difference between the two means is not statistically significant. This result should rule out an alternative explanation of firm sophistication, and is instead consistent with the hypothesis that UPREITs, because they are unaffected by the holding constraints, realize better market timing performance.

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 minimum period and then sell it at the next available opportunity, even though this may not be optimal from a market-timing standpoint. In fact, we see that the market-timing performance of Non-UPREITs suffers, compared to that of UPREITs. UPREITs, which can dispose of properties without holding constraints, do so. Their ability to do this, then also makes them better market timers, in terms of ex-post performance.

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. 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 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 minimum the strategies statistically significantly underperform an unconstrained strategy. Thus I demonstrate that the holding constraints are potentially very costly to investors. My analysis shows that the simulated strategies that I use generate trading patterns which are similar to the actually-observed trading behavior of firms in the market. Thus, this portion of the investigation offers general insights about the mechanics by which the holding constraint prevents REITs from generating abnormal profits.

Subsequently, I 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 most of my sample), and even a significant amount of holding periods shorter than two years. This suggests that REITs are eager to hold properties for a short time when it is appropriate to do so, and that the four-year (or two-year) holding period would limit their ability to do this. This matches with the trading patterns generated by the simulated trading strategy and suggests that REIT managers want to try to time the market, as 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 model the decision to hold a property beyond the four- or two-year mark 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 does not exist. 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 minimum holding period in order to be able to retain the profits made in the rising market. Lastly, I show that UPREITs exhibit a significantly better ex-post market timing performance than Non-UPREITs which are constrained.

The results from this study should 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 is costly to investors, in that it limits REITs' abilities to make timing profits they could ordinarily make by their nature of being informed traders in a predictable market. 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 does so at the potential costs involved with such a structure.[32] From the standpoint of allowing REIT managers to capitalize on their informational advantages, it would 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 to generate value through active trading seems inefficient at best. To my knowledge, this study (together with the concurrent study by Mühlhofer (2013)) is first in the literature to systematically document the effects of this regulatory framework and this advantage of the UPREIT structure.

[32] See Han (2006).

References

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.

Amihud, Y.: 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* **5**, 31–56.

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.

Brounen, D., Eichholtz, P. and Ling, D. C.: 2007, Trading intensity and real estate performance, *Journal of Real Estate Finance and Economics* **35**, 449–474.

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.

Fisher, J., Gatzlaff, D., Geltner, D. and Haurin, D.: 2003, Controlling for the impact of variable liquidity in commercial real estate price indices, *Real Estate Economics* **31**(2), 269–303.

Fisher, J., Gatzlaff, D., Geltner, D. and Haurin, D.: 2004, An analysis of the determinants of transaction frequency of institutional commercial real estate investment property, *Real Estate Economics* **32**(2), 239–264.

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.

Goyenko, R., Holden, C. and Trzcinka, C.: 2009, Do liquidity measures measure liquidity?, *Journal of Financial Economics* **92**(2), 153–181.

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

Han, B.: 2006, Insider ownership and firm value: Evidence from real estate investment trusts, *Journal of Real Estate Finance and Economics* **32**(4), 471–493.

Hartzell, J., Sun, L. and Titman, S.: 2014, Institutional investors as monitors of corporate diversification decisions: Evidence from real estate investment trusts, *Journal of Corporate Finance* **Forthcoming**.

Hochberg, Y. V. and Mühlhofer, T.: 2014, Market timing and investment selection: Evidence from real estate investors. Working paper, available at: http://tobias.muhlhofer.com.

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* .

Lin, Z. and Vandell, K.: 2007, Illiquidity and pricing biases in the real estate market, *Real Estate Economics* **35**(3), 291–330.

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, The predictability of real estate returns and market timing, *Journal of Real Estate Finance and Economics* **8**, 115–135.

Mühlhofer, T.: 2013, Why do reit returns poorly reflect property returns? unrealizable appreciation gains due to trading constraints as the solution to the short-term disparity, *Real Estate Economics* **41**, 814–857.

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*

Wheaton, W., Baranski, M. and Templeton, C.: 2009, 100 years of commercial real estate prices in manhattan, *Real Estate Economics* **37**(1), 69–83.

Figure 1: 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.

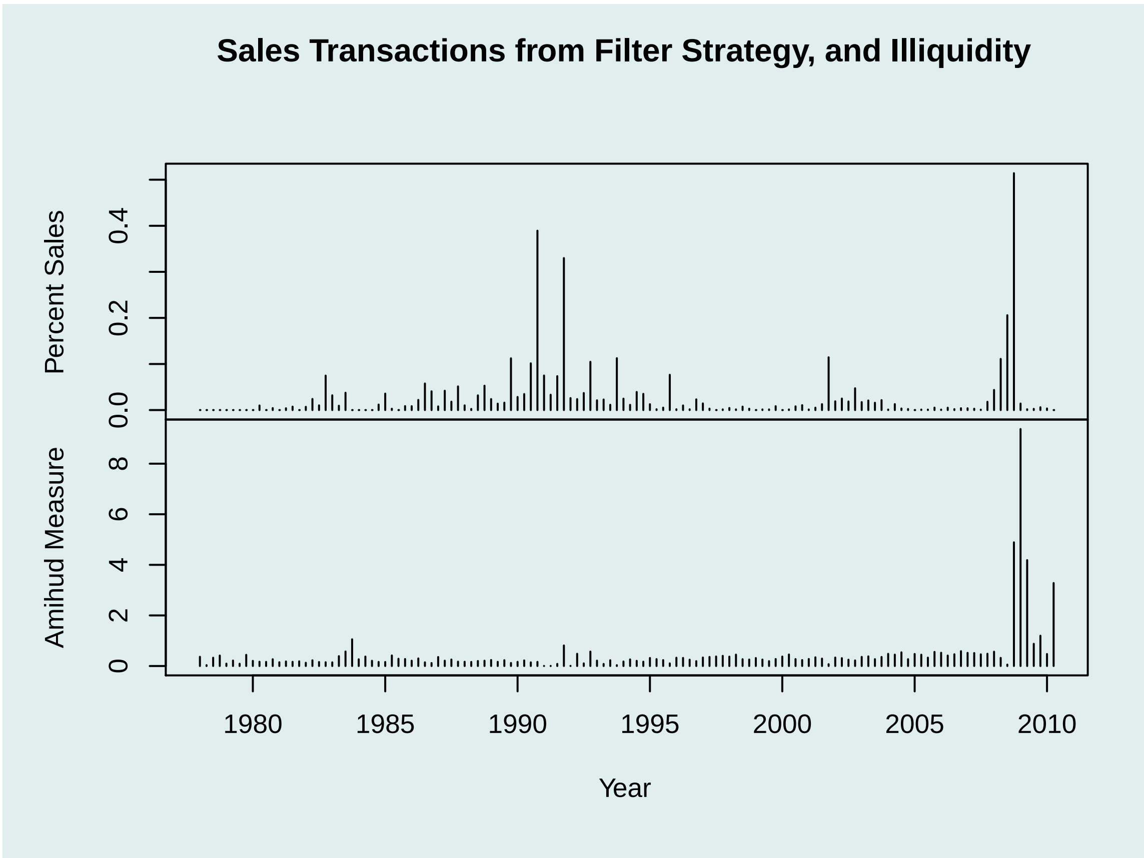

Figure 2: The top portion of this plot shows quarterly sale activity dictated by the trading-band strategy, while the bottom portion shows quarterly property market illiquidity. Sale activity is measured by the fraction of NCREIF sub-markets portfolios in which the strategy generates a *sell* signal in a given quarter. Property market illiquidity is measured by quarterly Amihud Illiquidity measure.

Spearman correlation between the two series: -0.3318.

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 1978-2010 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). Finally, the table reports outperformance of the best unconstrained strategy (Non-Adaptive) and of the best constrained strategy (Non-Adaptive, 2-Year Constraint), without allowing any trading during the 2008-2010 financial crisis when there would have been extremely low liquidity, specifically during 2008Q4 through 2009Q2, as well as 2010Q2.

Transaction costs										
0	**1%**	**2 %**	**3%**	**4%**	**5%**	**6%**	**7%**	**8%**	**9%**	**10%**
Non-Adaptive Strategy										
0.08613***	0.08517***	0.08421***	0.08324***	0.08228***	0.08132***	0.08036***	0.07940***	0.07843***	0.07747***	0.07651***
Adaptive Strategy										
0.07734***	0.07444***	0.07274***	0.07080***	0.06996***	0.06855***	0.06765***	0.0666***	0.06572***	0.06522***	0.06521***
Non-Adaptive Strategy, 4-Year Constraint										
0.06340***	0.06298***	0.06255***	0.06213***	0.06171***	0.06129***	0.06087***	0.06045***	0.06003***	0.05961***	0.05919***
Adaptive Strategy, 4-Year Constraint										
0.06361***	0.06320***	0.0628***	0.06239***	0.06191***	0.06144***	0.06106***	0.06066***	0.06027***	0.05984***	0.05944***
Non-Adaptive Strategy, 2-Year Constraint										
0.07248***	0.07186***	0.07124***	0.07062***	0.06999***	0.06937***	0.06875***	0.06813***	0.06751***	0.06688***	0.06626***
Non-Adaptive Strategy, No Trading in Illiquid Quarters										
0.05887***	0.05811***	0.05735***	0.05659***	0.05583***	0.05506***	0.05430***	0.05354***	0.05278***	0.05202***	0.05126***
Non-Adaptive Strategy, 2-Year Constraint, No Trading in Illiquid Quarters										
0.04849***	0.04801***	0.04752***	0.04704***	0.04656***	0.04608***	0.04560***	0.04512***	0.04463***	0.04415***	0.04367***
T-Statistics										
6.4749	6.3603	6.2448	6.1285	6.0358	5.9447	5.8198	5.6961	5.5705	5.4612	5.344
4.2224	4.1379	4.0532	3.9682	3.883	3.7976	3.712	3.6261	3.5401	3.4538	3.3673
4.3667	4.2729	4.1789	4.0846	3.9899	3.8949	3.7997	3.7041	3.6083	3.5121	3.4157

Average annualized return for buy-and-hold strategy: 0.1449

$^{\circ}$: 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 three 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* = 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 middle line compares the unconstrained strategy to the optimal strategy with two-year constraints (Non-Adaptive Strategy). The bottom line compares the optimal unconstrained strategy (Non-Adaptive) with no trading in illiquid quarters to the optimal strategy with two-year constraints (Non-Adaptive Strategy) with no trading in illiquid quarters.

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 1978-2010 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). Finally, the table reports outperformance of the best unconstrained strategy (Non-Adaptive) and of the best constrained strategy (Non-Adaptive, 2-Year Constraint), without allowing any trading during the 2008-2010 financial crisis when there would have been extremely low liquidity, specifically during 2008Q4 through 2009Q2, as well as 2010Q2.

Transaction costs										
0	**1%**	**2 %**	**3%**	**4%**	**5%**	**6%**	**7%**	**8%**	**9%**	**10%**
Non-Adaptive Strategy										
0.09193***	0.09115***	0.09038***	0.08960***	0.08883***	0.08805***	0.08728***	0.08651***	0.08573***	0.08496***	0.08418***
Adaptive Strategy										
0.08103***	0.07830***	0.07688***	0.07540***	0.07423***	0.07273***	0.07202***	0.07108***	0.07066***	0.07009***	0.06946***
Non-Adaptive Strategy, 4-Year Constraint										
0.07137***	0.07103***	0.07068***	0.07033***	0.06999***	0.06964***	0.06930***	0.06895***	0.06861***	0.06826***	0.06791***
Adaptive Strategy, 4-Year Constraint										
0.07026***	0.07014***	0.06982***	0.06949***	0.06909***	0.06896***	0.06864***	0.06808***	0.06777***	0.06742***	0.06711***
Non-Adaptive Strategy, 2-Year Constraint										
0.07803***	0.07751***	0.07699***	0.07647***	0.07596***	0.07544***	0.07492***	0.07440***	0.07389***	0.07337***	0.07285***
Non-Adaptive Strategy, No Trading in Illiquid Quarters										
0.05902***	0.05841***	0.05780***	0.05719***	0.05658***	0.05597***	0.05536***	0.05475***	0.05414***	0.05353***	0.05292***
Non-Adaptive Strategy, 2-Year Constraint, No Trading in Illiquid Quarters										
0.04941***	0.04901***	0.04860***	0.04820***	0.04779***	0.04738***	0.04698***	0.04657***	0.04617***	0.04576***	0.04536***
T-Statistics										
9.2532	9.1159	8.9779	8.8393	8.6999	8.56	8.4193	8.278	8.1362	7.9937	7.8506
7.6029	7.4965	7.3897	7.2825	7.1748	7.0668	6.9583	6.8495	6.7403	6.6307	6.5208
7.3371	7.2176	7.0975	6.9768	6.8557	6.734	6.6117	6.489	6.3658	6.2421	6.118

Average annualized return for buy-and-hold strategy: 0.1449

$^\circ$: 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 three 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* = 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 middle line compares the unconstrained strategy to the optimal strategy with two-year constraints (Non-Adaptive Strategy). The bottom line compares the optimal unconstrained strategy (Non-Adaptive) with no trading in illiquid quarters to the optimal strategy with two-year constraints (Non-Adaptive Strategy) with no trading in illiquid quarters.

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				
	2.00	11.00	19.45	31.00
Adaptive Strategy, Trading Bands				
Transaction Costs				
0	6.00	17.00	23.58	39.00
1%	6.00	18.00	24.18	40.00
2 %	6.00	19.00	24.85	41.00
3%	7.0	19.5	25.4	41.0
4%	7.00	20.00	25.75	42.00
5%	7.00	20.00	26.04	42.00
6%	8.00	21.00	26.37	43.00
7%	8.00	21.00	26.62	43.00
8%	8.00	21.00	26.86	44.00
9%	8.00	22.00	27.09	44.00
10%	9.00	22.00	27.23	44.00
Non-Adaptive Strategy, Moving Average				
	3.00	15.00	21.41	36.00

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**
Health Care	45	106	151
Hotel	0	4	4
Industrial	517	2384	2901
Multi-Family	213	1611	1824
Multi-use	2	23	25
Office	292	1734	2026
Regional Mall	18	110	128
Residential	0	3	3
Retail: Other	146	563	709
Self-storage	13	72	85
Shopping Center	74	1253	1327
Specialty	26	312	338
Sum	1346	8175	9521
Average Property Size (sqf)	130,076	148,131	
Average Property Age upon Sale (Years)	22.7	19.7	

Number of Unique Non-UPREITs: 44. Number of Unique UPREITs: 98.

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

CBSA Name	**Number of Properties**
Chicago-Naperville-Joliet, IL-IN-WI (Metro)	621
Atlanta-Sandy Springs-Marietta, GA (Metro)	502
Dallas-Fort Worth-Arlington, TX (Metro)	439
Los Angeles-Long Beach-Santa Ana, CA (Metro)	423
New York-Northern New Jersey-Long Island, NY-NJ-PA (Metro)	420
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD (Metro)	296
Houston-Sugar Land-Baytown, TX (Metro)	283
Washington-Arlington-Alexandria, DC-VA-MD-WV (Metro)	269
Miami-Fort Lauderdale-Pompano Beach-Homestead, FL (Metro)	255
Phoenix-Mesa-Scottsdale, AZ (Metro)	228
Total in Top-10 CBSAs	3736
Outside Top-10 CBSAs	5785
Total	9521

Total number of unique CBSAs with at least one property covered by SNL: 538.

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	1346	6539	7885
Large	0	1636	1636
Sum	1346	8175	9521

Number of Unique Small Firms: 133. Number of Unique Large Firms: 15.

Table 5: The Relationship Between Simulated Trades and Actual Trading Activity by REITs

This table shows evidence for the relationship between the trading activity generated by the trading-band filter strategy and the actual behavior of REITs in the property market. Panel A shows correlations between the percentage of the portfolio (percentage of sub-markets) sold in each quarter in the simulated filter strategy and the percentage of the portfolio (percentage of properties) sold by firms which underlie the holding constraint and firms which do not. Panel B shows average numbers of firms per CBSA. The top line compares the number of UPREITs per Short-Holding-Period CBSA (CBSAs for which the filter strategy generates an average holding period of less than four years) with the average number of UPREITs in all CBSAs. The bottom line compares the number of Non-UPREITs per Long-Holding-Period CBSA (CBSAs for which the filter strategy generates an average holding period of four years or more) with the average number of Non-UPREITs in all CBSAs.

Panel A: Correlations of percent sales by quarter, between simulated filter-based trades and actual trades.

	Unconstrained Firms	**Constrained Firms**
Correlation of Percent Sales By Quarter	0.3186*	0.1013
t-statistic	(2.49)	(0.72)

° : significance level ≤ 10%. *: significance level ≤ 5%. **: significance level ≤ 1%. *** : significance level ≤ 0.1%.

The t-statistic is for the hypothesis test that the true correlation is zero.

Unconstrained Firms: All UPREITs, and Non UPREITs which have held properties for more than four years.

Constrained Firms: Non UPREITs which have held properties for less than four years.

Note: these figures are for sales made up to 2008, as the minimum holding period changed thereafter.

Panel B: Firms per CBSA, in Long-Holding-Period and Short-Holding-Period CBSAs.

Subsample Type	**In Subsample**	**In Whole Sample**	**T-statistic for Difference**
UPREITs, Short-Holding-Period CBSAs	29.48	15.02	1.74°
Non-UPREITs, Long-Holding-Period CBSAs	8.64	2.47	3.35**

° : significance level ≤ 10%. *: significance level ≤ 5%. **: significance level ≤ 1%. *** : significance level ≤ 0.1%.

The t-statistic is for the hypothesis test that the true difference in means is zero.

Short-Holding-Period CBSAs: CBSAs for which the filter-based strategy generates an average holding period of four years or less.

Long-Holding-Period CBSAs: CBSAs for which the filter-based strategy generates an average holding period of four years or more.

Table 6: 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	2.070	4.123	4.878	7.140	9521
Non-UPREITs	2.899	4.958	5.453	7.915	1346
UPREITs	1.993	4.000	4.783	7.021	8175

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): $D = 0.1147$, p-value $< 1.13 \times 10^{-13}$.

All figures in years.

Table 7: Regression Results, Probit Regression Modeling the Decision to Sell a Property Before the Minimum Holding Period.

Dependent Variable: *min.holding*. The table shows regression results for probit regressions modeling the decision to hold a property beyond the minimum holding constraint (four years for sales occurring in or before July 30, 2008 and two years for sales occurring after this date), on local market conditions, the selling firm's UPREIT status, and the selling firm's size. *min.holding* is a dummy variable, which takes a value of 1 if a property was held in a REIT's portfolio for the time of the minimum constraint 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 all models at the bottom of the table. Model 3 includes dummy variables for the 12 property types. Z-statistics are provided. Disturbances are clustered by firm.

	Model 1	z-value	Model 2	z-value	Model 3	z-value
$(Intercept)$	−0.7939	−3.34***	−0.7919	−3.35***	−1.2982	−4.02***
$crisis$	2.2154	6.55***	2.1924	6.13***	2.2966	6.6***
$pareturn$	3.0930	2.13*	3.0790	2.14*	2.9105	2.02*
$UPREIT$	0.1093	0.37	0.1472	0.45	0.1219	0.35
$Large$			−0.1455	−0.48	−0.1958	−0.71
$pareturn \cdot UPREIT$	−2.3603	−1.33	−2.4328	−1.2	−2.0724	−1.02
$pareturn \cdot Large$			−0.0049	0	−0.1550	−0.1
Property-Type Dummies?	No		No		Yes	
N	4691		4691		4691	
Pseudo R^2	0.0231		0.0244		0.0335	
Model χ^2	76.19***		80.58***		909.44***	
$\chi^2, H_o : pareturn + pareturn \cdot UPREIT = 0$	0.49		0.19		0.34	

° : significance level ≤ 10%. *: significance level ≤ 5%.
: significance level ≤ 1%. * : significance level ≤ 0.1%.

$crisis$: A dummy variable equal to one for sales that occurred during 2009 and zero otherwise.
$pareturn$: 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.

$\chi^2, H_o : pareturn + pareturn \cdot UPREIT = 0$: 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 8: Distributional Statistics for Local Market Returns, Over 2 Years following a Sale

This table shows distributional statistics for local market returns over the two years following a property sale. These are measured as returns to the NCREIF NPI portfolio to the CBSA in which the sold property is located. The table shows these statistics for all REITs, UPREITs, Non-UPREITs, as well as Large and Small firms, without distinguishing by UPREIT status. The bottom of the table presents t-statstics testing the hypothesis that the true difference in mean returns between two groups of firms is zero, against the two-sided alternative.

	1st Quartile	**Median**	**Mean**	**3rd Quartile**
All Firms	−0.0489	0.0507	0.0396	0.1550
UPREITs	−0.0556	0.0395	0.0325	0.1548
Non-UPREITs	0.0027	0.1140	0.0817	0.1590
Large Firms	−0.0367	0.0334	0.0375	0.1637
Small Firms	−0.0490	0.0547	0.0400	0.1521
t-statistic UPREITs, Non-UPREITs: -10.67^{***}				
t-statistic Large, Small: -0.46				

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