Exploiting Property Characteristics in Commercial Real Estate Portfolio Allocation

Alberto Plazzi, Walter Torous, and Rossen Valkanov

Alberto Plazzi
is an assistant professor of finance and a Swiss Finance Institute Junior Chair at the University of Lugano, in Switzerland.
alberto.plazzi@usi.ch

Walter Torous
is the Lee and Seymour Graff Professor of Finance at the UCLA Anderson School of Management in Los Angeles, LA.
walter.n.torous@anderson.ucla.edu

Rossen Valkanov
is a professor of finance at the UCSD Rady School of Management in San Diego, CA.
rvalkanov@ucsd.edu

Commercial real estate is an important asset class. Current estimates put the value of investment-grade commercial properties in the U.S. at approximately $3 trillion. Direct investment in commercial real estate by pension funds is significant and is expected to increase in the upcoming years. Despite its growing importance, many questions still remain unexplored in the management of commercial real estate portfolios. For example, how should investors allocate their wealth across different commercial properties? How do the risk–return profiles of property types—apartments, industrial properties, offices, and retail properties—differ from one another? How should investors alter the composition of their commercial real estate portfolios to take advantage of movements in expected returns arising from changing underlying macroeconomic conditions?

Answers to these questions have been hampered by a number of factors including both data inadequacies and methodological difficulties. First, previous research on commercial property portfolio management has relied predominantly on aggregate property indices and so cannot provide insights into portfolio allocation at a disaggregated level. Second, commercial property returns are typically based on appraisal values. Because appraisals are updated infrequently, they tend to lag market values and render the resultant return series excessively smooth. As a result, the moments of these smoothed returns will systematically differ from the moments of the true market returns thereby potentially resulting in a misallocation in commercial real estate portfolios. Finally, recent evidence suggests that property-specific characteristics are related to the moments of commercial property returns. For example, Plazzi, Torous, and Valkanov [2010] provided empirical evidence consistent with a property's cap rate being informative about its subsequent returns.[1] Incorporating property-specific characteristics has the potential to improve the performance of commercial real estate portfolios. Unfortunately, the traditional mean-variance approach would require explicitly modeling the expected returns, variances, and covariances of all properties as functions of these characteristics. This task becomes computationally burdensome as the number of properties in the portfolio increases.

In this article, we apply recent advances in portfolio management (Brandt, Santa-Clara, and Valkanov [2009]) to efficiently incorporate the information contained in property-specific conditioning variables to the allocation of commercial real estate portfolios. In particular, we investigate whether a property's cap rate and other property-specific characteristics provide information that improves property portfolio performance. We do so by parameterizing the portfolio weight of each individual property as a function of its specific characteristics. The fact that a single function of characteristics applies to all properties over time significantly reduces the computational requirements of portfolio management. The coefficients of this portfolio policy are estimated by maximizing the average utility of a typical real estate investor. The Brandt, Santa-Clara, and Valkanov [2009] approach also allows us to easily impose non-negative weights on the property holdings. Unlike common stocks, it is at present difficult to take short positions in commercial properties because corresponding derivatives markets are either nonexistent or are very illiquid.

To estimate the optimal portfolio policy function, we rely on a large cross-sectional database of information on individual properties compiled by NCREIF (National Council of Real Estate Investment Fiduciaries). NCREIF assets are institutional-grade commercial properties managed by investment fiduciaries on behalf of tax-exempt pension funds. The current valuation of NCREIF properties is approximately $240 billion.[2] The disaggregated NCREIF data allow us to construct "pseudo-market" prices of NCREIF properties using the hedonic-type model of Fisher, Geltner, and Pollakowski [2007]. Intuitively, we rely on the fitted relation between available transaction prices and their corresponding lagged appraisal values and other variables to predict the prices at which properties that were only appraised would have sold for. Total returns based on these prices appear to mimic the time-series behavior of market-based returns. For example, the first-order autocorrelation coefficients of these returns are much lower in absolute value than those of returns based on appraisals.

Along with a property's cap rate, we rely on a number of additional building-specific conditioning variables to characterize a property's portfolio weight. These include a property's size measured by its appraisal value as well as a property's vacancy rate. Size plays an important role in the return performance of common stocks as well as properties (Pai and Geltner [2007]). Vacancy rates are related to residential property returns (Wheaton [1990]) and we explore their importance for commercial real estate returns. In constructing the optimal portfolios, we also consider whether a property is located in a larger and more liquid commercial property market—New York, Washington DC, San Francisco, Los Angeles, Chicago, and Boston.

Our empirical results are consistent with these property-specific conditioning variables playing an economically as well as statistically significant role in commercial real estate portfolio allocation. In particular, we find that the optimal portfolio places more weight, relative to the benchmark appraisal-weighted NCREIF portfolio, on properties having high cap rates, low vacancy rates, and on larger buildings as measured by their appraisal values. The size effect tends to be more significant in the larger and more liquid commercial real estate markets. The nature of the optimal portfolio also varies across property types. For example, the portfolio allocation of apartments and retail properties can be improved by relying more on cap rates, while vacancy rates appear to be more informative for the allocation of office portfolios. Optimal portfolios also vary over economic expansions versus economic recessions. For example, in recessions, optimal portfolios are aggressively tilted toward larger properties. Importantly, our findings continue to hold in the more realistic context of restricting the portfolio weights to be rebalanced annually.

The particular conditioning variables that we choose are simply by way of illustrating our general approach to modeling the effects of property characteristics on forming commercial real estate portfolios. Practitioners and other researchers can test the importance of other characteristics by using the methodology we present here. In addition, this article does not address at least one issue that may limit the practical implementation of our proposed portfolio policies. In practice, a portfolio policy should incorporate the transaction costs necessarily incurred in buying and selling properties when rebalancing portfolios. Nevertheless, our results point toward actual commercial real estate investors being able to improve the risk-adjusted performance of their portfolios by explicitly taking into account property characteristics such as cap rates, vacancy rates, and appraisal values.

The plan of this article is as follows. We begin by discussing the data relied upon in our empirical analysis. The two-step procedure used to estimate the predicted prices of all individual properties in the NCREIF database is detailed. The property characteristics capturing variation in commercial real estate's opportunity set are also introduced. We next briefly review the Brandt, Santa-Clara, and Valkanov [2009] methodology, emphasizing how it can be adapted to incorporate changing underlying economic conditions known to be important to the performance of commercial real estate and to impose a no-short-sales constraint, which is necessary when dealing with commercial properties. We then discuss our results both for a portfolio of all NCREIF properties as well as building type–specific portfolios. We conclude with a summary of our results.

DATA

We rely on the disaggregated information compiled by NCREIF. This information includes, among other items, a property's location and type, its net operating income and any capital expenditures, as well as the property's price. If a property is sold during a quarter, NCREIF records the net price (net of transaction costs) at which the property sold. Otherwise, NCREIF reports an appraisal value calculated either by an in-house expert (internal appraisal) or by an independent appraiser (external appraisal). Our sample begins in 1984 Q2 and ends at 2009 Q1.

The first step in our analysis is to construct a total return series for each individual property in the NCREIF database. This task is complicated by the fact that NCREIF properties, like other properties, do not transact frequently. For example, out of a total of 173,307 price observations in our NCREIF sample, only 4,875—about 3%—represent actual sales transactions. These appraisal values, however, are subject to a temporal lag bias as they tend to lag in time true contemporaneous market values. This, in turn, will smooth periodic returns based on appraisal values. Consequently, the moments of these smoothed returns will systematically differ from the moments of the true market returns. For example, the volatility of returns based on appraisals will be biased downward. Similarly, estimates of the correlations between returns and any conditioning variable will also be biased.[3] As a result, by relying on appraisal returns, the joint structure of returns would be improperly measured, thus potentially misallocating commercial real estate in a portfolio context.

Ideally, a portfolio analysis of commercial real estate would make use of transaction prices for a large number of properties over a sufficiently long period of time. To approximate this ideal, we rely on the hedonic regression model of Fisher, Geltner, and Pollakowski [2007] to *estimate* market-based prices of every individual property followed by NCREIF.

We do so by using a two-stage procedure. In the first stage, all transactions in the NCREIF database are used to estimate a hedonic price model in which corresponding transaction prices are regressed against properties' lagged appraisal values as well as several dummy variables controlling for time, property type, and location.[4] The key insight here is that while an appraisal value may represent a noisy estimate of a property's true market value, it serves as a valuable hedonic summarizing a building's characteristics, which are either observable, such as its size, or are unobservable, such as its quality. The estimated coefficients from this regression are then used in a second stage to construct predicted, or pseudo-market, prices based on the appraisal values and other characteristics of those properties that did not transact in a given quarter.[5]

A series of filters are subsequently applied to ensure that our results are not driven by outliers, which may reflect data entry and other errors.[6] The resultant pooled data serves as the basis of our estimation efforts.

Fisher, Geltner, and Pollakowski [2007] followed a similar procedure in constructing the TBI (Transaction-Based Index), an aggregate price index of commercial real estate based on the NCREIF data. However, to construct the TBI, they applied their first-stage estimates to a representative property mirroring the average characteristics of an NCREIF property. We, instead, determine the predicted prices for each individual property in the NCREIF database. Following Fisher, Geltner, and Pollakowski [2007], in order to reduce the estimation error in the first stage, the predicted prices are estimated using the pooled sample of all property prices during the 1984 Q2–2009 Q1 period, and using the individual property-type samples during the 1994 Q2–2009 Q1 subperiod.

Using these predicted prices, denoted by $P_{i,t}$, we calculate the log (total) return for property i in quarter $t + 1$ as

$$r_{i,t+1} = \ln\left(\frac{P_{i,t+1} + NOI_{i,t+1} - CAPX_{i,t+1}}{P_{i,t}}\right) \tag{1}$$

where $NOI_{i,t+1}$ denotes the building's net operating income (income minus operating expenses) earned during the period $[t;\ t + 1]$ while $CAPX_{i,t+1}$ represents corresponding capital expenditures.[7]

JPM-PLAZZI.indd 3 8/4/11 9:49:13 PM

Given these predicted individual property returns, our empirical methodology investigates whether a portfolio allocation across commercial real estate can be improved by relying on conditioning variables. In order to be relevant, these variables must capture variation, either cross-sectional or time-series, in commercial real estate's investment opportunity set. Guided by economic theory, evidence from previous studies, as well as data availability, we select the following conditioning variables:

Cap rate. A property's capitalization rate is calculated as the ratio between its net operating income ($NOI_{i,t}$) and its predicted value, $cap_{i,t} = NOI_{i,t}/P_{i,t}$. A cap rate corresponds to a stock's dividend–price ratio. There is reliable evidence in the finance literature consistent with a stock's dividend–price ratio predicting subsequent stock returns. Similarly, Plazzi, Torous, and Valkanov [2010] documented that a property's cap rate predicts subsequent property returns.

Size. We also include size as measured by a property's appraised value. The size effect is a prominent feature of common stock performance with small stocks, as measured by their market capitalization, earning a sizable return premium. However, Pai and Geltner [2007] investigated the size effect in the context of the institutional commercial real estate market and found that larger, as opposed to smaller properties, earn a return premium.

Vacancy rate. Vacancy rates are the final property characteristic that we rely on. Vacancy rates proxy for the supply versus demand relation prevailing in commercial real estate markets. As such, vacancy rates may capture changes in the commercial real estate opportunity set and thus subsequent property expected returns. Wheaton [1990] provided a theoretical argument for why vacancies and residential real estate values are negatively related, at least, contemporaneously. Empirically, Frew and Jud [1988] found vacancy rates to be a key factor in the determination of commercial office rents, while Smith [1974] directly linked local geographic and economic conditions to vacancy rates.

Summary Statistics

Summary statistics of the predicted returns, conditioning variables, and other related series are provided in Exhibit 1.

Exhibit 1

Summary Statistics

Panel A tabulates summary statistics for the value-weighted commercial property return series; the cap rate, vacancy rate, and size using the value-weighted average of all properties; and the CFNAI, an index of economic activity. Panel B tabulates their time-series correlation coefficients.

Panel A

	all	apt	ind	off	rtl	cap	vac	size	CFNAI
Mean	0.0257	0.0256	0.0327	0.0275	0.0284	0.0184	0.2677	17.671	–0.0204
Std	0.0665	0.0378	0.0648	0.0360	0.0395	0.0024	0.2741	0.3694	0.6413
AR(1)	–0.3156	0.1434	–0.2933	0.2143	–0.1005	0.8501	0.9831	0.9543	0.7925
Csd	0.0644	0.0515	0.0612	0.0715	0.0609	0.0088	0.2455	1.0866	–

Panel B

	all	apt	ind	off	rtl	cap	vac	size	CFNAI
all	1								
apt	0.288	1							
ind	0.656	0.075	1						
off	0.575	0.315	0.330	1					
rtl	0.715	–0.060	0.528	0.264	1				
cap	0.296	0.314	0.111	0.196	0.052	1			
vac	–0.002	0.115	–0.029	–0.109	0.041	–0.107	1		
size	–0.167	–0.537	–0.148	–0.387	–0.121	–0.395	–0.761	1	
CFNAI	0.205	0.487	0.218	0.406	0.174	0.251	0.246	–0.545	1

Given the returns of individual properties, we construct a number of representative property portfolios. To do so, we aggregate our individual property returns on a value-weighted basis, where an individual property's portfolio weight is proportional to its appraised value. The return to this market capitalization portfolio of all NCREIF properties is denoted by *r*. We also consider portfolios based on a particular property type—apartment (apt), industrial properties (ind), office properties (off), and retail (rtl)—and denote the corresponding portfolio returns by r_{apt}, r_{ind}, r_{off}, and r_{rtl}, respectively.

Means and standard deviations, both time-series and cross-sectional, as well as first-order autocorrelation coefficients of the various return series are tabulated in Panel A of Exhibit 1. Average returns are comparable across property types though industrial property returns are, on average, slightly higher and more variable. Office property return have the highest cross-sectional variability. Notice that the autocorrelation of the return series is quite low (in absolute value) across the property types. This is consistent with the time-series behavior of market-based returns.

We also construct appraisal-weighted aggregate series of the conditioning variables. All of the conditioning variables are persistent over time. Vacancy rates and size are particularly sticky. To the extent that these variables provide valuable information about the moments of property returns, it appears that this information does not quickly dissipate.

We are also interested in understanding how commercial real estate investment opportunities vary with economic conditions. To measure time variation in underlying economic conditions, we rely on the Chicago Fed National Activity Index (CFNAI), a monthly coincident indicator of broad-based economic activity originated by Stock and Watson [1999].[8] A positive value of CFNAI corresponds to a macroeconomic expansion while a negative CFNAI value coincides with a macroeconomic contraction. As can be seen from Exhibit 1, CFNAI is also persistent.

We now turn our attention to the time-series correlations among these series. For consistency with our subsequent analysis, we calculate correlations between the return series and the *lagged* conditioning variables. Panel B of Exhibit 1 shows that apartment returns are virtually uncorrelated with industrial and retail property returns, while industrial and retail property returns are themselves highly correlated. All returns are negatively correlated with the size of the corresponding property measured by its appraised value, while all returns and cap rates are positively correlated. We see very little correlation between vacancy rates and returns. Finally, a positive correlation prevails between returns and our measure of aggregate economic activity.

METHODOLOGY

We parameterize portfolio weights directly as a function of property characteristics. For each quarter *t*, there are a large number N_t of commercial properties in the NCREIF property universe. For each property, we have its return $r_{i,t+1}$ as well as a set of corresponding property characteristics collected in a *k*-dimensional vector $x_{i,t}$. An investor's problem is to allocate a portfolio across the N_t properties using the information contained in $x_{i,t}$.[9]

The fraction of wealth invested in property *i* at time *t* is denoted by $w_{i,t}$. The investor chooses portfolio weights to maximize the conditional expected utility of the resultant property portfolio return $r_{p,t+1}$,

$$\max_{\{w_{i,t}\}_{i=1}^{N_t}} E_t\left[u\left(r_{p,t+1}\right)\right] \qquad (2)$$

where $r_{p,t+1} = \sum_{i=1}^{N_{t+1}} w_{i,t} r_{i,t+1}$ and *u* denotes a pre-specified utility function.

Basic Case

In the basic case, we parameterize portfolio weights as a linear function of the property characteristics,

$$w_{i,t} = \bar{w}_{i,t} + \frac{1}{N_t}\theta' x_{i,t} \qquad (3)$$

where $\bar{w}_{i,t}$ are benchmark weights and θ is a k-dimensional vector of parameters to be estimated.

The vector $x_{i,t}$ is normalized to have a zero cross-sectional mean and a variance of unity at each time *t*. This implies that the second term in Equation (3) sums to zero across properties at time *t* and thus captures deviations from the benchmark weights. This allows the portfolio allocation to be tilted either toward or away from $\bar{w}_{i,t}$. The tilting is done according to the conditioning information contained in $x_{i,t}$.

We choose $\bar{w}_{i,t}$ to be market capitalization weights. That is, our benchmark is the portfolio of all NCREIF properties, which invests in each available property in proportion to its current appraised value. However, any other benchmark can beused. For example, we can assess whether conditioning information can improve a portfolio manager's commercial real estate allocation by choosing $\bar{w}_{i,t}$ to be the manager's current portfolio weights.

Given estimates of θ, the weights associated with the optimal portfolio policy are fully observable. It is important to note that θ is constant across time and across properties. Rebalancing of the portfolio away from the benchmark $\bar{w}_{i,t}$ is only due to the characteristics $x_{i,t}$ differing across properties and across time.

We maximize the sample analogue of Equation (2) with respect to the unknown parameters θ to estimate the optimal portfolio weights,

$$\max_{\theta} \frac{1}{T}\sum_{t=1}^{T} u\left(\sum_{i=1}^{N_{t+1}}\left(\bar{w}_{i,t}+\frac{1}{N_t}\theta' x_{i,t}\right)r_{i,t+1}\right) \tag{4}$$

This problem is relatively simple to optimize. A large and changing number of properties across time as well as a large number of conditioning variables can be easily accommodated.

Macroeconomic Variation

In the basic case, the coefficients of the portfolio policy θ are constant through time. This implies that the relation between property characteristics and the distribution of property returns is time invariant. However, this simplifying assumption may not be realistic in the case of commercial real estate whose performance is closely related to business cycle movements in the underlying economy.

To allow for the possible time variation in the coefficients of the portfolio policy, we follow Brandt, Santa-Clara, and Valkanov [2009] and explicitly model the coefficients as functions of a business cycle variable z_t. In this case, the portfolio policy stated in Equation (3) is extended as

$$w_{i,t} = \bar{w}_{i,t} + \frac{1}{N_t}\theta'\left(z_t \otimes x_{i,t}\right) \tag{5}$$

where ⊗ denotes the Kronecker product of two vectors. Now the impact of property characteristics on the property portfolio weights varies with the realization of the business cycle variable z_t. In our application, business cycle variation in economic conditions will be captured by the coincident indicator CFNAI.

No-Short-Sales Constraint

The portfolio policies considered to this point do not constrain the portfolio weights. As a consequence, for some properties the weights may turn out to be negative and require the properties to be shorted. While shorting is common in the case of stocks, it is not feasible for properties. Therefore, we need to modify a portfolio policy to directly impose the constraint that portfolio weights are non-negative.

There are a number of alternative ways of imposing non-negative portfolio weights. We directly impose the following non-negative weights in the portfolio optimization:

$$w_{i,t}^{+} = \frac{\max\left[0, w_{i,t}\right]}{\sum_{i=1}^{N_{t+1}} \max\left[0, w_{i,t}\right]} \tag{6}$$

This approach is straightforward to implement and interpret. However, unlike the previous two specifications, it is nonlinear in the conditioning variables.

RESULTS

Exhibit 2 displays the optimal portfolio results for our basic specification, Equation (3), and imposes the non-negativity constraint, Equation (6), when estimated using all NCREIF properties. Here and throughout, we assume that the investor has power utility with a coefficient of relative risk aversion equal to five.[10] All portfolios are assumed to be rebalanced every quarter. The results for the benchmark NCREIF market-cap weights are displayed in column I, while in subsequent columns, we sequentially add the cap rate and the other posited conditioning variables. As in standard regression analysis, a coefficient can be interpreted as the marginal effect of a particular variable $x_{i,t}$ on the optimal portfolio policy function. Moreover, since the conditioning variables are standardized, the coefficients are directly comparable both within as well as across specifications with a higher absolute value indicating a variable having a greater effect in the portfolio policy.

Exhibit 2
Portfolio Allocation: All Property Types

Optimal portfolio policy coefficients with non-negative weights estimated for all properties for alternative specifications assuming quarterly rebalancing.

	I	II	III	IV	V	VI
θ_{cap}		10.612	5.450	3.952	11.753	9.817
Std.Err.		1.134	1.292	3429	5.886	7.312
θ_{vac}			–9.169	–11.660	–13.977	–9.083
Std.Err.			1.222	2.995	4.716	3.711
θ_{size}				12.035	6.827	9.098
Std.Err.				3.417	3.216	4.483
θ_{top6}					–1.564	3.326
Std.Err.					0.981	2.120
$\theta_{top6 \times size}$					8.291	6.711
Std.Err.					2.712	2.733
θ_{East}						5.592
Std.Err.						9.069
θ_{West}						6.241
Std.Err.						10.995
$\theta_{MidWest}$						1.408
Std.Err.						9.149
θ_{South}						5.471
Std.Err.						10.416
LR p-value	–	0.000	0.000	0.000	0.000	0.000
$\lvert\omega_i\rvert \times 100$	0.068	0.068	0.068	0.068	0.068	0.068
max $\omega_i \times 100$	1.550	0.582	0.831	2.218	2.362	2.477
$\sum I(\omega_i = 0)/\mathrm{N}_t$	0.000	0.055	0.262	0.415	0.360	0.484
min N_t	707	707	707	707	707	707
max N_t	4518	4518	4518	4518	4518	4518
$\bar{r}$	0.103	0.115	0.124	0.129	0.130	0.130
σ	0.133	0.131	0.135	0.137	0.137	0.138
$CE(r)$	0.58	0.071	0.079	0.082	0.083	0.082
$SR(r)$	0.472	0.567	0.628	0.652	0.659	0.650
α	–	0.003	0.005	0.006	0.006	0.006
β	–	0.986	1.006	1.027	1.027	1.032
$\sigma(\epsilon)$	–	0.010	0.015	0.012	0.011	0.013

From Exhibit 2 we see that, all else being equal, the optimal portfolio places more weight, relative to the market-cap weights, on properties with high cap rates, low vacancy rates, and large buildings as measured by their appraised value (Columns II–IV). Being located in a larger and more liquid property market, either New York, Washington DC, San Francisco, Los Angeles, Chicago, or Boston and captured by the indicator variable *top*6, is by itself statistically insignificant. But the size effect is significantly enhanced within such markets (Column V). We also include location dummies in the final specification (Column VI) and see that none are statistically significant. This implies that the portfolio of all NCREIF properties weighted by their appraised values is well diversified by location across the country. The benefits of holding the optimal portfolio are evident in the Sharpe ratio's increase from less than 0.5 in the case of the benchmark portfolio to nearly 0.7 for

EXHIBIT 3
Portfolio Allocation by Property Type

Optimal portfolio policy coefficients with non-negative weights separately estimated using data for apartments (apt), industrial properties (ind), offices (off), and retail properties (rtl) assuming quarterly rebalancing.

	apt		ind		off		rtl	
θ_{cap}	–	19.746	–	4.879	–	13.149	–	10.780
Std.Err.	–	4.179	–	5.156	–	8.151	–	4.791
θ_{vac}	–	–1.212	–	–4.970	–	–9.770	–	–6.482
Std.Err.	–	1.224	–	3.087	–	5.689	–	4.274
θ_{size}	–	0.831	–	4.903	–	–0.214	–	9.199
Std.Err.	–	0.448	–	2.759	–	1.753	–	4.327
$\theta_{top6 \times size}$	–	2.842	–	5.759	–	5.398	–	–3.706
Std.Err.	–	0.708	–	2.355	–	2.467	–	1.712
LR p-value	–	0.000	–	0.000	–	0.000	–	0.000
$\lvert\omega_i\rvert \times 100$	0.272	0.272	0.154	0.154	0.210	0.210	0.291	0.291
max $\omega_i \times 100$	2.493	3.540	3.650	4.820	2.853	2.527	3.955	4.220
$\sum I(\omega_i = 0)/N_t$	–	0.272	–	0.464	–	0.259	–	0.305
min N_t	203	203	313	313	292	292	167	167
max N_t	921	921	1585	1585	1204	1204	778	778
$\bar{r}$	0.103	0.152	0.131	0.155	0.110	0.141	0.114	0.136
σ	0.076	0.082	0.130	0.131	0.072	0.074	0.079	0.083
$CE(r)$	0.088	0.136	0.094	0.117	0.097	0.127	0.100	0.122
$SR(r)$	0.828	1.375	0.700	0.877	0.969	1.359	0.933	1.167
α	–	0.011	–	0.006	–	0.007	–	0.005
β	–	1.052	–	1.007	–	1.019	–	1.015
$\sigma(\epsilon)$	–	0.019	–	0.01	–	0.012	–	0.019

the optimal portfolio (Column V). Importantly, with a relative risk aversion of five, we see a large gain in the certainty equivalent return[11] as a result of investing in the optimal portfolio relative to holding the benchmark.

Property characteristics are also important in explaining deviations of optimal portfolio weights from market-capitalization weights when investment is restricted solely to a particular property type. These results are tabulated in Exhibit 3 along with the corresponding property type–specific benchmark market-cap portfolios. For all property types, the optimal portfolios are tilted toward high cap-rate properties. This cap-rate effect is strongest for apartments and retail properties. In addition, the optimal office and industrial property portfolios are very sensitive to vacancy rates with relatively more wealth invested in low-vacancy-rate buildings. From Exhibit 3 we also see that the optimal apartment and industrial property portfolios are tilted toward large properties especially in the larger and more-liquid property markets. By contrast, this size effect is present only for office properties located in the larger and more-liquid property markets, but is diminished for retail properties located in these property markets. The benefits of holding the optimal property-specific portfolios are evident across all property types, but especially for apartments and offices. In the case of apartments, investing in the optimal portfolio relative to holding the benchmark increases the Sharpe ratio from approximately 0.8 to approximately 1.4. Similarly, investing in the optimal office portfolio relative to holding its benchmark increases the Sharpe ratio from approximately 1 to approximately 1.4. These gains translate into substantial increases in the annualized certainty-equivalent returns when compared to the benchmark portfolios, ranging from 2.2% for retail properties to as high as 4.8% for apartments.

EXHIBIT 4

Portfolio Allocation: Interaction with the CFNAI, All Property Types, and By Property Type

Optimal portfolio policy coefficients with non-negative weights estimated when the conditioning variables are interacted with a dummy variable, which equals one if the CFNAI is greater than zero (expansion), and zero otherwise. The corresponding coefficients are denoted by + and –, respectively. Quarterly rebalancing is assumed.

	all		apt		ind		off		rtl	
θ_{cap+}	–	8.040	–	26.805	–	10.777	–	12.494	–	25.750
Std.Err.	–	5.417	–	7.369	–	9.284	–	10.365	–	33.465
θ_{cap-}	–	3.379	–	23.313	–	2.609	–	4.156	–	4.507
Std.Err.	–	8.206	–	8.732	–	6.808	–	10.213	–	36.660
θ_{vac+}	–	–14.768	–	–1.086	–	–6.659	–	–8.301	–	–13.688
Std.Err.	–	5.767	–	1.921	–	5.620	–	6.011	–	20.288
θ_{vac-}	–	–3.715	–	–0.674	–	–4.985	–	–5.088	–	–10.617
Std.Err.	–	1953	–	1.917	–	6.379	–	4.457	–	23.974
θ_{size+}	–	3.007	–	1.368	–	7.280	–	–4.110	–	4.902
Std.Err.	–	2.447	–	0.705	–	5.208	–	2.542	–	8.529
θ_{size-}	–	9.676	–	1.983	–	–1.016	–	4.613	–	32.664
Std.Err.	–	4.668	–	1.320	–	3.682	–	4.968	–	51.003
$\theta_{(top6 \times size)+}$	–	3.693	–	0.599	–	0.083	–	1.469	–	–6.730
Std.Err.	–	1.898	–	0.729	–	1.217	–	2.644	–	12.751
$\theta_{(top6 \times size)-}$	–	7.529	–	5.064	–	12.393	–	8.623	–	6.401
Std.Err.	–	3.367	–	1.871	–	8.315	–	5.789	–	11.109
LR p-value	–	0.000	–	0.000	–	0.000	–	0.000	–	0.000
$\lvert\omega_i\rvert \times 100$	0.068	0.068	0.272	0.272	0.154	0.154	0.210	0.210	0.291	0.291
$\max \omega_i \times 100$	1.550	2.545	2.493	3.597	3.650	4.655	2.853	2.744	3.955	4.457
$\sum I(\omega_i = 0)/N_t$	–	0.491	–	0.283	–	0.427	–	0.382	–	0.427
$\min N_t$	707	707	203	203	313	313	292	292	167	167
$\max N_t$	4518	4518	921	921	313	1585	1204	1204	778	778
$\bar{r}$	0.103	0.134	0.103	0.156	0.131	0.159	0.110	0.149	0.114	0.143
σ	0.133	0.139	0.076	0.084	0.130	0.132	0.072	0.081	0.079	0.088
$CE(r)$	0.058	0.086	0.088	0.139	0.094	0.121	0.097	0.133	0.100	0.127
$SR(r)$	0.472	0.674	0.828	1.376	0.700	0.903	0.969	1.338	0.933	1.175
α	–	0.007	–	0.011	–	0.007	–	0.009	–	0.007
β	–	1.038	–	1.082	–	1.011	–	1.039	–	1.022
$\sigma(\in)$	–	0.020	–	0.021	–	0.012	–	0.031	–	0.034

The differences in the estimated θ coefficients across property types evident in Exhibit 3 are consistent with apartments, industrial properties, offices, and retail properties being characterized by different risk–return profiles. For example, cap rates do not play a statistically significant role in the optimal industrial property portfolio, and vacancy rates are not statistically significant in the optimal apartment portfolio.

Exhibit 4 allows the effects of the property characteristics on the optimal portfolio weights, overall and property type–specific, to vary with the realization of the coincident indicator CFNAI. By estimating the specification in Equation (5), this proxy for the business cycle variable z_t allows us to investigate how an investor can time the composition of optimal property portfolios.

Turning our attention first to the results for all NCREIF properties, we see in Exhibit 4 that in expansions, $CFNAI > 0$, the optimal portfolio is tilted more in the direction of both high-cap-rate and low-vacancy-rate properties. In recessions, $CFNAI < 0$, the optimal portfolio is tilted more in the direction of large properties, especially those located in the larger and more-liquid property markets captured by the variable *top*6. This is a counter-cyclical investment policy that in recessions directs investors to seek out more valuable properties with corresponding stronger cash flows.

EXHIBIT 5

Portfolio Allocation: Appraisal-Based Returns with Quarterly and Annual Rebalancing, All Property Types

Optimal portfolio policy coefficients with non-negative weights estimated on all property types using either quarterly or annual non-overlapping total returns based on appraisal values.

	Quarterly		Annual	
θ_{cap}	–	6.382	–	5.262
Std.Err.	–	5.288	–	3.525
θ_{vac}	–	–12.534	–	–21.039
Std.Err.	–	5.418	–	8.701
θ_{size}	–	–0.163	–	–1.650
Std.Err.	–	1.293	–	0.436
$\theta_{top6 \times size}$	–	6.887	–	7.209
Std.Err.	–	2.537	–	2.764
LRT p-value	–	0.000	–	0.000
$\lvert\omega_i\rvert \times 100$	0.072	0.072	0.084	0.084
max $\omega_i \times 100$	1.607	1.982	1.864	1.742
$\sum I(\omega_i \leq 0)/N_t$	0.000	0.347	–	0.368
min N_t	694	694	635	635
max N_t	4300	4300	2495	2495
$\bar{r}$	0.087	0.104	0.089	0.104
σr	0.045	0.046	0.064	0.065
$CE(r)$	0.081	0.099	0.078	0.093
$SR(r)$	1.035	1.385	0.762	0.983
α	–	0.005	–	0.017
β	–	0.993	–	0.974
$\sigma(\in)$	–	0.012	–	0.019

From Exhibit 4, we can see that the tilt toward high-cap-rate properties in expansions is especially evident in industrial, office, and retail properties. The tilt toward low-vacancy-rate properties in expansions is optimal for all property types. The tilt toward larger properties in recessions is due primarily to office and retail properties. Interestingly, the optimal office portfolio tilts toward smaller properties in expansions. Also, regardless of the property type, optimal portfolio holdings of properties located in the larger and more-liquid property markets are seen to tilt toward large properties in recessions.

To this point, we have considered only portfolio policies with quarterly rebalancing. In Exhibit 5, we report the results of two additional, but more practical, implementations of the proposed methodology to commercial real estate. In particular, we consider a quarterly rebalancing strategy based on appraisal values as opposed to predicted prices. The results, displayed in the first column of the exhibit, are not substantially different from those in the previous exhibits with the exception of the coefficient on the cap rate, which is slightly lower. Also, we consider an annual rebalancing strategy in which the returns are computed without overlap from the appraisal values. This is the most restrictive, but also the most realistic, specification that we consider. The estimates of resultant portfolio policy, presented in the second column of Exhibit 5, are again not substantially different from our previous findings. In fact, the vacancy rate and property size in the large and more-liquid property markets seem to be even more important characteristics than previously reported.

CONCLUSION

A commercial property is characterized by more than simply the fact that it is an office building or an industrial warehouse located in the U.S. Northeast or the U.S. South. Yet heretofore commercial real estate portfolio analytics have relied primarily on property type and property location when allocating investment across commercial properties. Since other property characteristics, for example, property cap rates, are related to the moments of property returns, we apply the parametric portfolio allocation approach of Brandt, Santa-Clara, and Valkanov [2009] to efficiently incorporate this property-specific information into commercial real estate portfolio management. Not surprisingly, taking such information into account significantly improves the risk-adjusted performance of commercial real estate portfolios relative to property portfolios that are well diversified across property types and property locations.

When considering the universe of all NCREIF properties, we find that the optimal portfolio weights are tilted more toward high-cap-rate, low-vacancy-rate, and high-appraisal-value properties when compared to a benchmark portfolio that holds these properties in proportion to their appraisal values. These portfolio policies, however, are shown to vary with prevailing economic conditions. For example, in recessions, optimal portfolios are aggressively tilted toward high appraisal value properties reflecting these properties' stronger cash flows.

The methodology presented in this article can be used by practitioners and other researchers to test the importance of other property characteristics in commercial real estate portfolio allocation. While our results are encouraging, they should be extended so as to improve the practical implementation of our optimal portfolio policies. For example, it is important to investigate the sensitivity of these results to including transaction costs and other market frictions. Also, while all of this article's results are in sample, it would be interesting to investigate whether the optimal strategy estimated in sample would yield equally impressive results out of sample. We leave investigating these and other issues to future research.

ENDNOTES

We gratefully acknowledge RERI for their financial support and NCREIF for making their data available to us. We thank David Geltner for sharing his TBI code and Holly Horrigan for her excellent research assistance. We would also like to thank Joseph Pagliari, James Shilling, and an anonymous referee for their valuable comments on an earlier draft, and Martha Peyton for her encouragement. All remaining errors, however, are our own.

[1]Similarly, Campbell et al. [2009] showed that the rent-price ratio reliably forecasts residential property returns.

[2]NCREIF aggregates the confidential information contributed by its members and provides indices based on aggregate data, such as the quarterly NCREIF Property Index (NPI), for use by the real estate industry.

[3]See Geltner [1991] and [1993] for an analysis of the effects of the appraisal procedure on aggregate indices and individual properties.

[4]Property type–specific results, when reported, are based solely on data specific to the particular property type. Also, because there are fewer than 50 transactions per quarter in the NCREIF database until the mid-1990s, we follow Fisher, Geltner, and Pollakowski [2007] and use only data beginning in 1994 Q2 when estimating property-specific returns to minimize their estimation error.

[5]The methodology also accounts for transaction sample selection bias in the first stage using a Heckman [1979] two-step approach. Moreover, a Bayesian noise filtering technique is applied to reduce the effect of noise in the quarterly series due to the limited number of transactions. See Fisher, Geltner, and Pollakowski [2007] for further details on this estimation procedure.

[6]First, we drop observations for which the ratio between *NOI* and current price exceeds 0.2 in absolute value. This is done to eliminate cases where the *NOI* is too large in negative or positive terms relative to the price of the building. We also drop returns larger than 0.8 or smaller than −0.4. These represent less than 0.5% of the overall number of observations.

[7]When a property transacts, the predicted price during that quarter is set equal to the transaction price.

[8]This index is based on the first principal component of 85 economic activity series and is constructed to have an average value of zero and a standard deviation of one. Because economic activity tends to grow at a trend, an index reading of zero corresponds to the economy growing at trend.

[9]The decision of how much to allocate to commercial real estate relative to other asset classes, such as stocks and bonds, is assumed to have already been made.

[10]Defined as $u(W) = \frac{W^{1-\gamma}}{1-\gamma}$, where γ is the coefficient of relative risk aversion. It should be noted that the investor will be reluctant to have the optimal portfolio deviate from the benchmark portfolio for sufficiently large relative risk aversion γ.

[11]The certainty equivalent return is the fixed known return an investor is indifferent in receiving as compared to the uncertain return generated by the portfolio.

REFERENCES

Brandt, M., P. Santa-Clara, and R. Valkanov. "Parametric Portfolio Policies: Exploiting Characteristics in the Cross Section of Equity Returns." *Review of Financial Studies*, 22 (2009), pp. 3411-3447.

Campbell, D., M. Davis, J. Gallin, and R. Martin. "What Moves Housing Markets: A Variance Decomposition of the Price-Rent Ratio." *Journal of Urban Economics*, 66 (2009), pp. 75-89.

Fisher, J., D. Geltner, and H. Pollakowski. "A Quarterly Transactions-based Index of Institutional Real Estate Investment Performance and Movements in Supply and Demand." *Journal of Real Estate Finance and Economics*, 34 (2007), pp. 5-33.

Frew, J., and D.G. Jud. "The Vacancy Rate and Rent Levels in the Commercial Office Market." *Journal of Real Estate Research*, 3 (1988), pp. 1-8.

Geltner, D. "Smoothing in Appraisal-Based Returns." *Journal of Real Estate Finance and Economics*, 4 (1991), pp. 327-345.

——. "Temporal Aggregation in Real Estate Return Indices." *AREUEA Journal*, 21 (1993), pp. 141-166.

Heckman, J. "Sample Selection Bias as a Specification Error." *Econometrica*, 47 (1979), pp. 153-161.

Pai, A., and D. Geltner. "Stocks Are from Mars, Real Estate Is from Venus." *The Journal of Portfolio Management*, 33 (2007), pp. 134-144.

Plazzi, A., W. Torous, and R. Valkanov "Expected Returns and Expected Growth in Rents of Commercial Real Estate." *Review of Financial Studies*, 23 (2010), pp. 3469–3519.

Smith, L.B. "A Note on the Price Adjustment Mechanism for Rental Housing." *American Economic Review*, 64 (1974), pp. 478–481.

Wheaton, W.C. "Vacancy, Search, and Prices in a Housing Market Matching Model." *Journal of Political Economy*, 98 (1990), pp. 1270–1292.