The Term Structure of Lease Rates*

Ernesto Aldana, Andrea Chegut, Jacob S. Sagi

Fall 2025

Abstract

In financial markets, forward contracts reflect market perception of future price dynamics. Less transparent markets, like those for real assets that can be leased, lack such tools, often hindered by both data availability and heterogeneity in the underlying asset. We introduce a methodology for unbundling information in leases, allowing us to filter a dynamic term structure of forward lease rates for a median-quality asset. We implement the methodology using a panel of NYC and Boston office leases between 2005 and 2017. The imputed term structure reflects different expectations by users and owners about future rental market conditions. Beyond shedding new light on rental market dynamics, our model can be used to quantify risk and reward for leasing strategies. To illustrate this, we examine the financial viability of a nascent "long"-"short" space market strategy commonly used by coworking providers in the last business cycle.

JEL R33, G12, G17

Keywords Commercial Real Estate, Contract Term, Term Structure

*Ernesto Aldana, Clemson University. Andrea Chegut, Massachusetts Institute of Technology. Jacob S. Sagi University of North Carolina at Chapel Hill.

Acknowledgements: We dedicate this paper to the memory of Dr. Andrea Chegut, our coauthor, whose untimely passing away deeply affected all who knew her. We are grateful to Compstak and JLL for providing data. We are also grateful for comments from Brent Ambrose, and for seminar feedback at the Real Estate Research Institute, the University of North Carolina, and the Massachusetts Institute of Technology.

1 Introduction

Many firms lease rather than purchase highly illiquid capital assets (Eisfeldt and Rampini, 2009; Rampini and Viswanathan, 2013). Examples include commercial real estate (CRE), airplanes, ships, manufacturing equipment, infrastructure, and medical equipment. Whereas purchasing capital outright reflects the full useful life of an asset, a lease reflects the value of the asset over the term of the lease. In principle, leases with different maturities for identical capital assets reflect a term structure of the "forward curve" — essentially, the risk-adjusted future value curve — of the asset's stream of benefits. In turn, the forward curve conveys information about the asset's market or industry. In cases where the industry has macro significance, forward curves may even provide information about the macroeconomy. The petroleum industry, for instance, furnishes an example where information in forward markets may have macroeconomic significance. Of course, the term structure of oil forwards is directly observed from traded contracts rather than, say, leases on oil fields. This, however, precisely underlies the need that our paper attempts to address: Leases on assets that provide non-homogenous benefits may, and should, contain important information. How might one unpack this information?

We present a methodology for estimating the forward term structure of a flow-variable, which we view as a commodity, from leases of capital stock that produces the commodity. Importantly, the commodity may be of heterogeneous quality. Although the idea is quite general, we explore it in the specific case of office leases where it is possible to obtain rich data of lease terms.[1] Within our context, the commodity is the occupancy of office space, and the commodity is not uniform because the quality of the space will vary within and across buildings. Our methodology allows us to characterize cyclical dynamics of the term structure of "the forward price of space", much in the way that one might refer to the term structure of any other commodity, currency, or interest rate. While well-illustrated in the context of real estate, the concept extends to any real asset market where right-of-use contracts are prevalent.

[1] According to Savills' "Around the World in Dollars and Cents (2016)", the aggregate value of CRE that could potentially attract institutional investors is commensurate in size with the bond market. Leases in CRE determine CRE cash flow and are therefore fundamental to the operation and valuation of CRE assets. Importantly for our purposes, the collection of newly executed lease agreement at any given time represents, among other things, the market's assessment of the current and anticipated price of space (per unit time).

To understand our proposed methodology, first consider a real estate market example in which space for rent is homogeneous. A lease is simply a bundled collection of forward commitments.[2] Because the bundling of financial securities is linear in payoffs, by observing enough leases and accounting for observation error, one can back out a term structure of forward rates. Leased spaces vary by quality in practice, however, violating the homogeneity assumption. To account for this, we assume that, on average, any two equal-maturity leases are equivalent up to a latent "quality deflator". In other words, after adjusting for a filtered quality factor, we assume that two leases with the same maturity only differ by an observation error that depends on unobserved transaction details but is independent of quality considerations. Our main contribution is to simultaneously solve the quality adjustment *and* the portfolio unbundling problem to arrive at estimated *quality-adjusted* forward lease rate dynamics. Doing so requires filtering both the market term structure and each lease's quality adjustment factor. The result is a term structure of lease rates corresponding to a median quality asset in that market.

We demonstrate the methodology using data on office lease contracts for New York City (NYC) over the period 2005-2016, and for Boston between 2006-2017. Specifically, we employ a Kalman filter to back out a vector autoregressive process governing a spot, 5-year, and 10-year key lease rates for median-quality assets.[3] The spot rate should be interpreted as the month-to-month price per square foot for an average transaction in median-quality space. Correspondingly, the 5-year (resp. 10-year) lease rate as determined at date t should be interpreted as the price at which a tenant could expect to lock-in one square foot of median-quality space for one month starting at date $t + 5$ (resp $t + 10$). We control for lease stipulations such as tenant improvement allotments, free rent concessions, and rent escalations. Because of insufficient data, we do not control for tenant credit or lease renewal options, but we do address this potential shortcoming in a discussion section.

When applied to the office leasing market in NYC and Boston, our methodology

[2] To see this, note that the lessee commits to occupy space at some time in the future and pay a pre-specified price for the right to such occupancy. Moreover, leases normally refer to multiple periods (months or quarters) and multiple rent payment amounts, which can be seen as a sequence, i.e., a bundle, of forward commitments.

[3] Consistent with the literature on the term structure of interest rates, each point on the term structure is assumed to be a linear interpolation of its nearest key rates. This assumption manageably reduces the size of the state variable space to three.

leads to several novel and interesting insights. First, we document that over the observation period the term structure is typically upward sloping in both cities. Second, the curvature of the term structure is typically negative for NYC, but it is positive for Boston. The shape of the term structure can potentially be explained by the tension between the anticipated rental growth rates versus the depreciation/obsolescence of the space over time. Seen this way, a concave term structure in a given market may signify market anticipation of decelerating rental growth or accelerated depreciation (e.g., from insufficient future capital expenditure).

Another novel insight from our analysis is that average lease dynamics across the key rates in each city are dominated by a single fundamental shock. This contrasts with existing term-structure literature that tends to find different independent shocks corresponding to short-term, persistent, and growth-rate influences (e.g. Chiang, Hughen and Sagi, 2015). Moreover, in contrast with common single-factor commodity models (e.g., CIR or Ornstein-Uhlenbeck processes) where the greatest variation is found in the spot rate, our estimates suggest that, unconditionally, long-dated space forwards are more volatile than shorter-dated space forwards. This is consistent with a market in which landlords and tenants face large adjustment costs to short-term supply or needs for space, so that changes in the rental market are more pronounced in longer-dated transactions.

We illustrate the utility of our approach by applying the estimated model to analyzing a short-term leasing strategy common in certain real asset markets. In the office market context, this strategy corresponds to the so-called co-working model. Specifically, we evaluate the ex-ante risk-reward attributes of a leasing strategy that is "long" a long-dated lease on space and "short" a sequence of short-term leases.[4] Based on our estimation, long-dated space forward prices are typically higher than spot rents. It is therefore not surprising that, in our sample period, a purely spatial co-working strategy is generally *ex-ante* unprofitable in both Boston and NYC, even assuming full occupancy of the short-term rental. To make the strategy consistently profitable on a risk-return basis, one or a combination of the following must be true: The long-term lease base rent must be below market (e.g., bottom quartile or tercile), the short-term leased space must be more intensively used (e.g., 20% increased user

[4]Coworking companies are known to employ multiple revenue strategies to support their business model. The "long-short" space strategy is just one of them.

density), or income must be derived from other sources (e.g., tenant services). Overall, our analysis makes clear that the co-working "long-short" space strategy, while potentially lucrative under the right conditions, is by itself not sufficiently financially viable across all market dynamics to justify a co-working business; i.e., additional revenue and risk mitigation is required from other sources. From an underwriting perspective, our work suggests that it may be sensible for co-working spaces to be treated more like hotels than office properties, but this has not been the prevailing practice (Chegut and Langen, 2019). Noting that our sample period predates the spectacular demise of WeWork, our analysis can be said to point to the substantial risks that accompanied the famous firm's strategy.

1.1 Literature

Earlier lease valuation models can be found in Miller and Upton (1976), McConnell and Schallheim (1983), and Schallheim and McConnell (1985). Miller and Upton (1976) provides a model where the firm decides optimally between leasing or buying and asset, whereas McConnell and Schallheim (1983) and Schallheim and McConnell (1985) analyze the valuation of leases with different embedded options. More recent literature, focusing on real estate, builds on the model presented in Grenadier (1995), in which lease rates result from simultaneous equilibria in the leasing market and the underlying asset market, and where competing developers behave optimally. In the model, leases are contingent claims on building values and are determined in the real estate market by conditions such as the number of competing developers, expectations about future demand for space, and current construction activity.[5]

With the possible exception of Gavazza (2011), who studies airplane lease terms, empirical work in finance and economics tends to focus on real estate leases where data is relatively more abundant. Existing empirical studies typically treat the lease term is a regressor in a model where the current lease rate is the dependent variable. Many, though not all, such studies also tend to find a lease term structure with positive slope (e.g., Wheaton and Torto, 1994).

Agarwal et al. (2011) present a model in which the term structure of leases is

[5]This model is extended to the case of non-perfect competition in the developer market in Grenadier (2005). Other related work includes Clapham and Gunnelin (2003), Ambrose, Hendershott and Kłosek (2002), and Ambrose and Yildirim (2008).

determined endogenously by the tenant's capital structure and space market conditions. The value of the service flow follows, as in Stanton and Wallace (2009), a geometric Brownian motion (GBM). The focus, however, is on quantifying the impact of credit risk on the part of lessee. Comparative statics of the model show that an increase in credit risk is linked to a steeper (positive) term structure slope. More recently, Brueckner and Rosenthal (2024) confirm this prediction. Their results suggest that the term structure has a positive slope and that, in going from less to more risky lessees, the term structure steepens by roughly 5%. Because we do not have data on lessee credit risk, we cannot control for it. Instead, credit risk heterogeneity is captured in our empirical specification as part of our estimate of the latent lease "quality" distribution. More recently, Takahashi and Yoshida (2025) estimate spot rental rates using CompStak data, also relying on an underlying GBM process and no-arbitrage assumptions to pin down the term structure.

Our main contribution is in estimating a dynamic model that is flexible enough to capture different components of the term structure and their joint dynamics, despite the presence of quality heterogeneity. In particular, and unlike standard models in the literature, we do not assume that the spot lease rate follows a GBM.[6] We also argue that standard "no-arbitrage" constraints on the lease rate term structure are unlikely to apply in the presence of realistic physical and market frictions. Thus, using a flexible dynamic model specification has the potential to better capture term structure attributes that would otherwise not be apparent if one imposes either GBM or no-arbitrage constraints.

2 Methodology

Our analysis relies on the assumption that the present value of the stream of lease contract *cash flow*, which need not replicate the pattern of usage, equals the present value of *contract usage*.[7] Detailed lease data typically specify periodic cash flow, including concessions, along with the corresponding periods of occupancy. Our assumption allows us to view each lease as financially equivalent to a bundle of forward

[6] Sagi (2021) demonstrates how search frictions, which are arguably prevalent in the leasing market, can lead to dramatic departures from GBM.

[7] One simple example in which contractual cash flow does not follow the same pattern as contractual usage is when a lessee pays upfront for multi-period usage.

contracts on space occupancy, with the present value of the bundle (and therefore the lease) given by the discounted sum of forward prices. Given enough heterogeneity in lease maturities one can effectively "unwind" a set of lease-equivalent bundles to arrive at the constituent forward contract prices. To make the idea concrete, we provide a simple example.

Example Consider three distinct but simultaneously signed leases for three distinct but identical assets. The first has a three-period term and commits the lessee to a constant contract rent of 5.0 paid in each of the next three periods. The second lease has a two-period term corresponding to a rent of 4.5 paid in each of the next two periods. The third lease, also with a two-period term, starts in period two and pays a rent of 4.0 and 7.0 in periods two and three, respectively. For the sake of this example, interest rates are constant at 0% per period. Further assuming away counterparty risk, and given that the three leased assets are identical, leads to three equations each of which sets the present value of the cash flow from a lease to the present value of its forward commitment. Specifically, denoting a current forward commitment to providing space in period i as F_i yields,

$$\begin{aligned} 5+5+5 &= F_1+F_2+F_3, \\ 4.5+4.5+0 &= F_1+F_2, \\ 0+4+7 &= F_2+F_3. \end{aligned}$$

The (unique) solution to the set of equations above is $F_1 = 4, F_2 = 5$ and $F_3 = 6$. □

To account for non-zero interest rates, one need only discount each cash flow and forward term by its term rate (i.e., multiply each by the price of a corresponding \$1 zero-coupon bond). Moreover, in practice, heterogeneity in the quality of space means that the set of linear equations linking the present value of rents and forward space commitments is only approximate. We address differences in quality by assuming that each lease in our sample is drawn from a distribution of quality, which we assume to be Gaussian and constant over time. In other words, we assume that the *quality-adjusted* present value of a lease equals the discounted sum of forward prices.

Two complicating factors that we do not address head-on in this paper are the potential presence of valuable (to the lessee) renewal options and lease default options. The presence of these options will bias imputed forward rates up, and the impact may systematically differ across horizons. This means that our imputed forward rates should be interpreted to be gross of average renewal option and systematic credit spreads.[8]

Another challenge to implementing the general idea outlined above is the fact that leases for real capital assets tend to cluster around the industry-standard maturities. For instance, in the office data we explore later, leases tend to cluster around maturities of of 5-, 7-, 10-, and 15-years, and there are very few short-term leases. Thus, there is unlikely to be enough heterogeneity in lease terms to identify all points along the term structure. To deal with this, following the literature on the term structure of interest rates, we assume that all forward lease rates at a given time can be derived from a small set of *key forward rates* via linear interpolation. In our application, we consider three key forward rates: the spot price of occupancy, corresponding to contracting for the next month of space; the five-year forward rate, corresponding to locking in today one month's occupancy starting in five years; and the ten-year forward rate, similarly defined.[9] For instance, under the linear interpolation assumption, the forward price of one month's occupancy starting in 2.5 years from now is a 50:50 weighting of the spot and the 5-year rates.

2.1 Details of the Fundamental Observation Equation

For any lease executed at some calendar month t, we can compute the sequence of periodic cash flow commitments from the initial commencement of lessee use to the final period. These includes concessions, such as periods with free rent, as well as rent escalations. We note that some lease concessions entail negative cash flow.[10] Denote the

[8]In principle, if one were able to control for the credit worthiness of each tenant one would be able to back out a term-dependent credit spread. We do not have such data. Likewise, among the leases we employ in our estimation, the vast majority of leases lack any information regarding renewal options, and in the few leases with observable data, the information is too vague to interpret quantitatively.

[9]We restrict attention to three key rates, despite the small amount of clustering around 7- and 15-year leases because the increase in parameters that we estimate in our structural state-space model is quadratic in the number of key rates.

[10]This is often the case with commercial real estate leases featuring tenant improvements. These are capital investments made by the lessor and used to build the space to suit the tenant.

periodic cash flow sequence for a lease i executed at period t as $(c_{i,t,0}, c_{i,t,1}, \ldots, c_{i,t,T_i})$, where the last contract period ends at $t + T_i$.

Denote as $d_{t,\tau}$ the discount factor at period t for a risk-free obligation due at month $t + \tau$. This factor is also the price of a \$1 zero-coupon bond maturing at $t + \tau$. Next, denote as $F_{t,\tau}$ the forward price, determined at month t, of locking in a commitment to one period of asset use at period $t + \tau$. As mentioned earlier, we assume that $F_{t,\tau}$ is determined by linear interpolation/extrapolation from a vector of key rates, F_t^K. Assume, for instance, monthly periods with spot, 5-year, and 10-year key rates. Then $F_t^K = (F_{t,0}, F_{t,60}, F_{t,120})'$. Moreover, the forward price of usage commencing one month after contract execution would be given by $F_{t,1} = \frac{59}{60}F_{t,0} + \frac{1}{60}F_{t,60}$. We then denote as v_τ the vector of linear interpolation/extrapolation coefficients applied to the vector of key rates to generate $F_{t,\tau}$. In the example just given, $v_1 = (\frac{59}{60}, \frac{1}{60}, 0)$.

If all leases corresponded to the same asset quality, we would proceed as follows. For a given lease i at month t we define $PV_{i,t}$ to be the sum of lease contract cash flows, $(c_{i,t,0}, c_{i,t,1}, \ldots, c_{i,t,T})$, discounted to the present using the ZCB rates:

$$PV_{i,t} = \sum_{\tau=0}^{T_i} d_{t,\tau} c_{i,t,\tau}, \tag{1}$$

Let $t + \tau_{i,c}$ be the usage commencement date. The corresponding expression for discounted forward claims on the same space is:

$$\begin{aligned} \sum_{\tau=\tau_{i,c}}^{T_i} d_{t,\tau} F_{t,\tau} &= \sum_{\tau=\tau_{i,c}}^{T_i} d_{t,\tau} v_\tau \cdot F_t^K \\ &= w'_{i,t} \cdot F_t^K, \end{aligned} \tag{2}$$

where the jth component of the vector $w'_{i,t}$ is equal to $\sum_{\tau=\tau_{i,c}}^{T_i} d_{t,\tau}(v_\tau)_j$, and where $(v_\tau)_j$ is the jth component of v_τ. Note that the expression in Eq. (2) is linear in the lease key rates.

Our main hypothesis is that there is an underlying prevailing term structure of lease rates in a given lease market and that it is a determinant of lease prices. Lease prices are also affected by the quality of the space, which is heterogeneous. We incorporate this variation in quality into our framework by assuming the existence of a lease-specific quality adjustment factor, g_i, that can be applied to the lease's present valued. We further assume, without loss of generality, that each lease's quality adjustment factor is drawn from a distribution with unit mean and variance σ_ε^2, so that the average quality lease value is unadjusted.[11] In other words, we are asserting that

$$g_i PV_{i,t} = w'_{i,t} \cdot F_t^K + \text{mean-zero independent observation error}, \tag{3}$$

For each lease, $PV_{i,t}$ on the left side is observable in our data, as is the vector of weights, $w_{i,t}$. This is the empirical analogue of the simple example provided earlier, with the addition of the quality adjustment factor and observation noise. One problem with estimating this equation is that longer leases will have larger $PV_{i,t}$'s and correspondingly larger $w_{t,i}$ components (i.e., a short-maturity lease is a smaller bundle of forward obligations than a long-maturity lease). In fact, the $PV_{i,t}$'s will roughly scale with the lease term, suggesting that deviation from average pricing will also scale with lease maturity. To control for the expected heteroskedasticity in the error term, we normalize both the left and right side of the equation above by $w'_{t,i} \cdot \mathbf{1}$ (the sum of components of $w_{t,i}$). To that end, we define the *normalized* discounted lease cash flow as,

$$nPV_{i,t} = \frac{PV_{i,t}}{w'_{t,i} \cdot \mathbf{1}}. \tag{4}$$

and restate our observation equation as

$$g_i \; nPV_{i,t} = \frac{w'_{i,t} \cdot F_t^K}{w'_{t,i} \cdot \mathbf{1}} + u_{i,t}, \tag{5}$$

where the $u_{i,t}$'s are assumed to be identically and independently distributed (iid) across leases (i.e., there is no systematic deviation from average pricing for any subset of leases). One can further decompose the left side of Equation (5) by expressing the quality deflator in terms of deviations around its unit mean, $g_i = 1 - \varepsilon_i$, where ε_i denotes a mean-zero iid *hedonic noise* component, assumed to be independent of

[11] The mean of the quality adjustment factor's distribution defines the level of the key rates for the average quality lease.

the $u_{i,t}$'s. Rearranging terms yields

$$nPV_{i,t} = \frac{w'_{i,t} \cdot F_t^K}{w'_{t,i} \cdot \mathbf{1}} + \varepsilon_i nPV_{i,t} + u_{i,t}. \tag{6}$$

We pause to note that $nPV_{i,t}$ can be interpreted as what is often denoted in industry as an *effective rent*. To see why that is, consider an average quality lease (i.e., $\varepsilon_i = 0$) and assume no observation error (i.e., $u_{i,t} = 0$). If the term-structure were flat, then $F_t^K = f_t\mathbf{1}$, where f_t is the rental rate at period t across all horizons. Then (6) reduces to $nPV_{i,t} = f_t$—the flat prevailing market rental rate. More generally, $nPV_{i,t}$ is a weighted average of rent paid from lease commencement to maturity. Expressed in this manner, equation (6) identifies the relationship between the effective rent $nPV_{i,t}$ and a weighted average of lease key rates.

Let $\overline{nPV_t}$ denote a cross-sectional average over $nPV_{i,t}$. Our final assumption is that, fixing the weight vector across leases (i.e., $w_{i,t} = w_t$) as well as the key rate vector, averaging over the cross-section of leases yields $\overline{nPV_t} = \frac{w'_t \cdot F_t^K}{w'_t \cdot \mathbf{1}}$. Under this assumption, the term $\varepsilon_i nPV_{i,t}$ has zero mean. Because ε_i and $u_{i,t}$ are independent, it follows that the variance of the error term is given by $\sigma_\varepsilon^2 nPV_{i,t}^2 + \sigma_u^2$. This allows us to interpret equation (6) as an observation equation in which the quality distribution variance component is separately identifiable from the observation noise component.

2.2 Filtering the key rates and quality distribution

Equation (6) links the key rates to observable quantities. Because the structure of the residual variance is known, one may be tempted to estimate the key rates each period using Generalized Least Squares. For realistic lease panels, however, the resulting key rate estimates are far too noisy. Moreover, independent period-by-period estimates of the key rates ignore any time-series correlation structure which can help better pin down the estimates.

We now describe our approach to the estimation of the term structure, which relies on the Kalman filter. We model the (unobserved) key rate dynamics as a VAR(1). This is done both for ease of estimation and also because, in our limited sample period, one cannot reject the hypothesis that effective rents are mean-reverting.[12]

[12] Mean reversion is unlikely to hold in a sample period spanning several decades where inflation alone will lead to secular growth in rents. We expect that over a longer time period, the dynamics

The assumed evolution is:

$$F_{t+1}^K = \bar{F} + \rho F_t^K + \epsilon_{t+1}, \tag{7}$$

where $\epsilon_{t+1} \sim N(0, Q)$ is uncorrelated with F_t^K and Q is some positive semi-definite matrix, $\bar{F}$ is a constant vector, and ρ is a square AR(1) matrix with real eigenvalues in $[0, 1)$.[13]

If, in addition to $\bar{F}, \rho$ and Q, the variances of the u_t's and ε_i's in the lease observation equations are known, then these can be used in conjunction with the Kalman filter to back out the key rates as long as one is also willing to assume that $u_{i,t}$ and ε_i are Gaussian.[14] To do that, one can assume that one period before the first in the dataset, the key rates are drawn from the prior distribution given by the unconditional mean and variance of the process for F_t.[15]

The above procedure allows one to impute the time series of key rates assuming that the various model parameters are known. To estimate these, we search for the parameter set that maximizes the Gaussian likelihood of observed data. Details are available in most graduate texts on econometric time-series techniques.

We now turn to applying the outlined methodology to a panel of office leases in New York and Boston.

3 Data

We use proprietary data on individual lease contracts corresponding to commercial office properties in New York City and Boston. For each city, we have two data sources: JLL and CompStak.[16] The JLL data comes from the company's internal records on office lease transactions and is used essentially as a "check" but also to

we identify will be complemented by a slow-moving growth factor.

[13]The constraint on the eigenvalues ensures that the process is mean-reverting and non-oscillating.

[14]We view Gaussian ε_i to be a crude approximation, as such an assumption is not consistent with the earlier assumption that $E[\varepsilon_i nPV_{i,t}] = 0$, which is required to identify σ_ε^2 and σ_u^2.

[15]The unconditional mean of F_t is given by $(I - \rho)^{-1}\bar{F}$. Its associated unconditional variance, V, is given by the solution to the linear matrix-valued equation, $V = Q + \rho V \rho'$. Once again, we caution that the unconditional mean and variance we estimate reflect only the sample period we observe. Over a significantly longer, or inflationary, time period we would expect that the time series of forward rates would not be mean-stationary.

[16]CompStak is a commercial real estate data company that provides information on comparable lease transactions, or comps, for over 50 United States cities. Jones Lang LaSalle Incorporated (JLL) is a provider of commercial real estate services with offices in 80 countries.

enhance the sample size, and hence statistical power, of the CompStak data. In the case of CompStak, data for each reported lease contract is collected directly from real estate brokers or other entities involved with the transaction. In exchange for the lease contract information, a reporting entity receives other brokers' lease information or "comps". CompStak staff validate each newly entered lease transaction for consistency and plausibility. This data collection process may alleviate, at least partially, concerns related to sample selection bias, data misreporting and measurement error.[17]

Lease transaction data includes details on the characteristics of the lease contract, the property, and tenants. Reported lease contract terms comprise the transaction date, the commencement date, the term of the lease, the type (e.g., net or gross), the brokers that were involved in the transaction, whether the lease has renewal options, the size of the space, the floors it occupies within the property, the rent schedule, and any concessions to the tenant. The rent schedule field reports the monthly rent per square foot over the life of the contract (including rent escalations), while the concessions field reports the number of months of free rent and/or tenant improvements (TI) offered to the tenant. The property characteristics available in our data set are the property type (e.g., office, retail, etc.), the address, a quality designation ("Class"), the size and age of the building, as well as the number of stories. The data set also includes the name of the landlord and the tenant, as well as the industry of the latter. In our analysis, we use a subset of these features for commercial office spaces in Manhattan and Boston.

Importantly, we have information on the type of lease. Gross and full-service leases are essentially equivalent and are all-inclusive (i.e., the tenants are only responsible for the associated quoted rents). Modified gross and net (including "NN" and "NNN") leases place some or all of the burden of operating and maintenance expenses on the tenant. Because we do not have access to information on property expenses, we restrict our analysis to gross and full-service leases, which constitute the largest category of transactions in both data sets.

Table 1 reports the number, by year, of leases in our final sample. To minimize the impact on the estimation of inconsistencies in reporting and record creation across both data sources, we drop quarters in which fewer than 30 contracts were

[17] CompStak outlines their verification approach on their website.

Table 1: **Gross office leases by year of contract execution and data source**. The table shows the distribution of lease executions across time and data source in our final sample for NYC and Boston.

	NYC			Boston		
	Compstak	JLL	Total	Compstak	JLL	Total
2005	125	1	126	-	-	-
2006	329	3	332	166	-	166
2007	376	226	602	234	23	257
2008	356	204	560	330	86	416
2009	326	195	521	411	142	553
2010	481	262	743	402	191	593
2011	410	484	894	402	210	612
2012	83	892	975	414	259	673
2013	325	549	874	179	253	432
2014	293	662	955	241	261	502
2015	-	741	741	195	305	500
2016	78	260	338	89	289	378
2017	-	-	-	114	43	157
Total	3182	4479	7661	3177	2062	5239

originated in the two data sets, combined. Furthermore, we exclude leases with outlier effective rents (outside the 0.5% and 99.5% quantiles of the combined sample for each city). For the case of New York City, this procedure yields a final sample of 7,661 transactions executed between the second quarter of 2005 and the second quarter of 2016. In the case of Boston, the final sample consists of 5,239 observations corresponding to 39 quarters between 2006 and 2017.[18]

Table 2 reports on key characteristics for the leases in our final sample. Appendix A.1 describes how the data sets are combined and duplicate leases are identified. Lease term is measured in years and is on average 8.7 years in NYC, whereas the mean lease term in Boston is almost 3 years shorter. Time to commencement measures the number of months between the execution date and the commencement date, where it takes NYC leases 1.2 months to commence, it takes 4.3 months for Boston.[19] Time

[18] The third quarter of 2017 is missing from the Boston sample because both data sources contain less than 30 transactions corresponding to office gross leases in this quarter.

[19] In Appendix A.1 we show that the average time to commencement for NYC is driven down by the information from JLL, in which the commencement date often coincides with the transaction date, yielding a time to commencement of 0 months (Table A.1).

Table 2: **Gross office leases by year of contract execution and data source**. The table shows summary statistics by city for the final sample used in the estimation. All USD figures are per square foot.

	Mean	S.D.	p1	p25	p50	p75	p99
NYC							
Lease term (years)	8.70	4.04	1.08	5.08	10.00	10.50	20.00
Time to commencement (months)	1.21	4.67	0.00	0.00	0.00	1.00	19.00
Time to expiration (years)	8.80	4.10	1.17	5.25	10.00	10.58	20.28
Starting rent (USD per month)	4.99	2.01	2.17	3.58	4.58	5.92	11.37
Average rent (USD per month)	4.69	1.95	1.97	3.33	4.27	5.58	11.11
Average rent increase (USD per year)	0.03	0.05	0.00	0.00	0.03	0.05	0.12
Number of rent bumps	0.89	0.83	0.00	0.00	1.00	1.00	3.00
Average bump duration (months)	54.30	19.71	12.00	44.00	57.33	60.00	120.00
Tenant improvements (USD)	31.51	29.56	0.00	0.00	25.00	60.00	100.00
Free rent (months)	4.98	3.98	0.00	2.00	4.00	7.00	15.00
Boston							
Lease term (years)	5.82	2.67	1.00	4.17	5.08	7.00	13.00
Time to commencement (months)	4.26	5.94	0.00	2.00	3.00	5.00	28.00
Time to expiration (years)	6.17	2.78	1.17	4.71	5.50	7.42	14.52
Starting rent (USD per month)	2.70	1.04	0.98	1.92	2.51	3.33	5.72
Average rent (USD per month)	2.52	1.01	0.95	1.75	2.36	3.13	5.49
Average rent increase (USD per year)	0.05	0.08	0.00	0.03	0.05	0.07	0.22
Number of rent bumps	3.47	2.50	0.00	2.00	4.00	5.00	10.00
Average bump duration (months)	20.13	17.95	9.00	11.67	12.00	20.50	95.62
Tenant improvements (USD)	24.40	21.39	0.00	5.00	20.00	40.00	80.00
Free rent (months)	2.62	2.76	0.00	0.00	2.00	4.00	12.00

to expiration is the number of years between the execution date and the expiration of the lease. Similar to lease term, it takes about 9 years for a lease to expire on NYC and a little over 6 years in Boston, on average.

Starting and average rent are measured in USD per square foot per month. Average rent is the mean of all monthly payments implied by the lease, taking into account months of free rent and tenant improvements (TI).[20] Average rent increase represents the average yearly rent hike in USD, i.e., the difference between rent due at the end of the lease and starting rent, divided by lease term. The average rent increase for Boston space is almost twice that of NYC space, \$0.05 to \$0.03, respectively.

The number of rent bumps represents the number of times the rent is updated during the leases, excluding the free rent period. According to our data, anticipated rent updates are more prevalent in Boston than in NYC, where most leases stipulate one rent bump during the life of the contract. In contrast, Boston lease contracts include an average of more than three rent updates. Consequently, the average rent bump duration, i.e. the length of the period between two different rent payments, is substantially higher in NYC than in Boston, 54 vs. 20 months, respectively.

Tenant concessions like TI are measured in USD per square foot. The TI for NYC space is about \$7.00 more per square foot than for space in Boston. Free rent, expressed in months, is almost twice as high in NYC as in Boston, 4.98 and 2.62 months, respectively. Importantly, TIs are substantial when given, amounting to nearly a full year's rent, while the median free rent concession amounts to four months or one third of yearly rent for the case of NYC and one sixth of a year's rent in Boston.

To sum up, office leases in NYC are unconditionally longer and pricier than in Boston. On the other hand, Boston leases incorporate relatively more price updates throughout the life of the contract than in the case of NYC. In the following section, we describe how we use the information in office lease contracts to estimate the underlying term structure of the price of office space.

Using equation (6), Figure 1 plots average monthly effective rent per square foot for NYC and Boston leases in our final sample. Effective rents in Boston and NYC

[20] Out of the 7,661 NYC leases in our sample, the TI field is missing from 841, and 144 leases do not have free rent information. In the case of Boston, among the 5,239 leases, 870 are missing the TI field, and 1,388 leases have no information under free rent. In all of these cases, we assume that the TI and/or the number of free rent months are zero.

follow similar patterns, albeit with stark differences in level and variance, with NYC experiencing both higher and more volatile effective rents. In the run-up to the GFC, both cities observed rental rate increases of roughly 50%. Rents then dropped during the crisis and started on a slow but clear recovery from around 2011 until the end of the sample.[21]

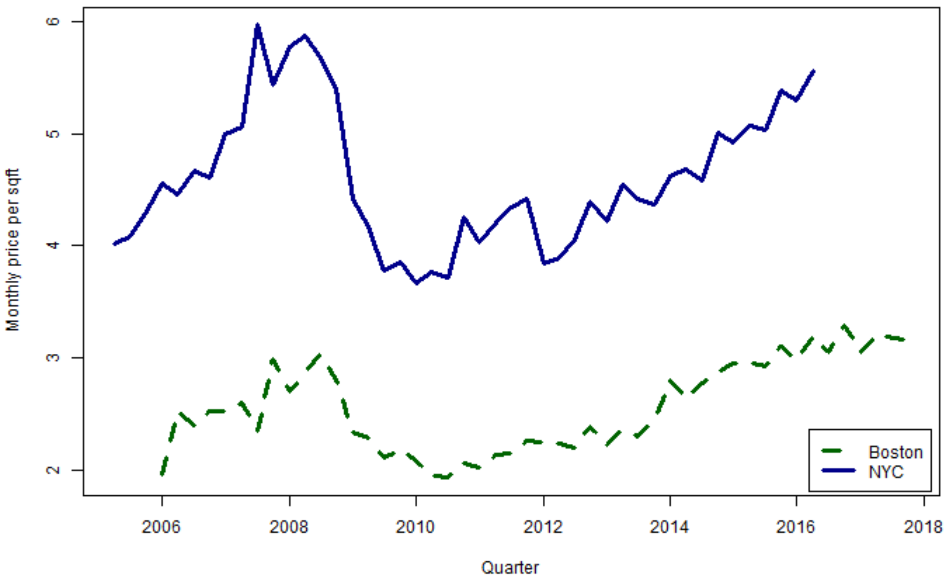

Figure 1: **Monthly effective rent per square foot**. The plot shows the average monthly effective rental rate calculated in Equation (5) for the Boston (dashed green line) and NYC (solid blue line) office markets. The rate is monthly USD per square foot.

4 Empirical estimation

Following the methodology outlined in Section 2, we estimate the dynamics of key rates for each of NYC and Boston. For each city, we separately model the key rates using a 3-component vector corresponding to a spot, 5-year and 10-year forward rates. We bucket leases by quarter to estimate average quarterly key rates and conduct separate estimation processes for each city.

[21]According to CommercialCafe.com, NYC (Boston) rents stood at \$5.75 (\$3.84) pSqFt per month as of 1Q2025.

In estimating the VAR(1) parameters of each model, we further subject the parameter space to a set of constraints that guarantee impulse-response functions that are monotonic and share the same sign in response to fundamental shocks. This is consistent with two hypotheses about the leasing market. First, an increase in one key rate should not cause a decrease in another key rate (i.e., at the market level, occupancy at different points in time are not substitutes). Second, at least to first order, the leasing market is informationally efficient in the sense that a shock to one key rate should not impact another key rate with delay. Further justification and an outline of the econometric methodology of imposing such constraints is provided in Aldana and Sagi (2025).[22] This approach permits us to simultaneously estimate both the model parameters and a "best guess" at the time series of realized key rates (referred to as the "smoothed" time-series).

Because the model features many parameters, our search for the maximum likelihood parameter set is done by first generating a pseudo-uniform (Sobol) grid of 5,000 parameter sets. We then evaluate the constraints on each of these points and discard all parameter sets that do not satisfy the constraints, along with those that lead to singularities in the Kalman filter matrices. This procedure yields 300 *valid* parameter sets for the NYC sample, and 1,466 for the case of Boston. We use each of these points to initialize a separate search for a local optimum and keep converged searches (about two thirds). Next, we rank converged searches according to likelihood and keep points that lie reasonably close to the maximum likelihood estimate.Our procedure yields 147 local maxima for NYC and 286 for Boston.

The resulting points, while similar in terms of likelihood, cover a wide range of model parameters with notable differences with respect to key rate volatility. This raises concerns about the relying exclusively on the global maximum likelihood estimate. To address this, we employ Bayesian model averaging (see, for example, Fragoso, Bertoli and Louzada, 2018) as follows. We first calculate for each convergent point estimate a "combined likelihood" score: The product of the Kalman filter likelihood with the likelihood from fitting to the unconditional moments of the effective rent time series in each city.[23] Then, we construct a probabilistic mixture of

[22] We note that constraints on impulse response functions are common in the macroeconomics literature on the identification of vector autoregression (see Uhlig, 2005; Kilian and Lütkepohl, 2017). More details on our constrained maximum likelihood estimation, as well as the specific constraints on the matrices ρ and Q are discussed in Appendix A.2.

[23] We describe the calculation of our combined likelihood score in Appendix A.3.

Table 3: **Observation error term variance estimates**. The table shows the blended estimates for the variance of the terms in the observation equation error, $\sigma_\varepsilon^2 nPV_{t,i}^2 + \sigma_u^2$.The parameters of each individual local maximum were combined following the procedure described in Appendix A.3. The first row displays the variance and standard deviation of the idiosyncratic error term $u_{t,i}$. The second row corresponds to the *hedonic noise component* ε_i.

NYC		Boston	
Var	S.D.	Var	S.D.
Idiosyncratic error term $u_{t,i}$			
0.2112	0.4596	0.2055	0.4533
Quality dispersion ε_i			
0.0974	0.3120	0.0769	0.2773

the converged point estimates using their combined likelihood scores as weights. We refer to these as *blended* estimates.

Table 3 displays the blended estimates for the variance of the observation error and the variance of the quality dispersion, $u_{t,i}$ and ε_i, respectively.[24] The point estimates indicate for the two office rental markets are similar, with NYC exhibiting about 10% higher quality heterogeneity. The measurement standard error is close to \$0.50, which represents roughly 13% of the effective rent in NYC but closer to 20% in Boston. This may reflect a less efficient price-setting environment in Boston (i.e., Boston may exhibit more search frictions). The quality dispersion is roughly 0.3, meaning that the quality adjustment factor (i.e., g_i) has a 95% confidence interval of roughly $[0.4, 1.6]$.

Table 4 reports the blended unconditional mean and variance components of the key rate vector, F_t^K, as implied by the estimated state equation parameters. The figures in the table convey a sense of the average term structure over our sample period. In NYC, because the unconditional 5-year key rate is larger than both the spot and the 10-year key rate, we observe an increasing yet concave "typical" term structure. In the case of Boston, the term structure is monotonically increasing and convex, suggesting structural differences between the two office markets.

An eigenvalue decomposition of the unconditional variance matrices reveals that, in each of NYC and Boston, there is one dominant independent shock contributing

[24]Blended state equation parameter estimates are shown in Appendix Table A.3.

Table 4: **Unconditional mean and variance of key rates**. The table shows the blended estimates for the unconditional mean and variance of the key rate vector, F_t^K. For each individual local maximum, the unconditional mean of F_t^K is given by $(I-\rho)^{-1}\bar{F}$. The unconditional variance, V, is given by the solution to the linear matrix-valued equation $V = Q + \rho V \rho'$. These estimates were then combined following the procedure described in Appendix A.3

	NYC			Boston		
	Spot	5yr	10yr	Spot	5yr	10yr
Mean						
	3.5219	3.9876	3.7784	2.0258	2.2165	3.0778
Variance						
	0.3644	0.4102	0.4358	0.1881	0.1286	0.1813
	0.4102	0.5416	0.5485	0.1286	0.1060	0.1422
	0.4358	0.5485	0.5841	0.1813	0.1422	0.3086

to around 90% of the unconditional key rate variance.[25] It is this dominant factor in each city which is associated with the higher 10-year diagonal key rate variance (the (3,3) component in the variances of Table 4). In other words, the unconditional variance of the spot contract is substantially smaller than that of the 10-year contract, and all three variance components are largely driven by a single factor. This is unusual for commodities exhibiting a mean-reverting price. In the standard term structure literature, whether dealing with commodities, currencies, or interest rates, long-horizon forward rates tend to vary less than the spot (the so-called "Samuelson Effect"). It appears that lease forward rates exhibit the opposite behavior. This might arise because there is greater scope for adjusting the supply and demand for long-horizon space commitments. For instance, users of space may be relatively inflexible concerning their current needs, while suppliers of space cannot quickly increase their "inventory" of space and may be disinclined to keep short-term space vacant for fear of foregoing rents. In other words, supply and demand for short-term occupancy are relatively fixed in the short run. Because more adjustment is possible for future perceived needs and availability, the market for locking-in long-dated space may exhibit greater variability in response to changing economic conditions.

In comparing forward claims on space with other commodities, it is important to

[25] Specifically, the largest eigenvalue of the unconditional variance matrix accounts for 96% of the variance in NYC and 88% in the case of Boston.

note that it is not possible to "store" a spot claim to space in order to take advantage of large swings in the forward price of space.[26] In other words, there is less scope for arbitrage in the space market than with other consumable commodities.

To examine how an economic shock differentially impacts the key rates and eventually "decays", we undertake an impulse-response exercise. In it, we consider a one standard deviation disturbance to the dominant fundamental shock underlying the dynamical system for the key rates. The shock is applied to the key rates at their unconditional mean. We plot the results for both cities in Figure 2. In both cases, the 10-year forward rate reacts strongly to the initial shock (see the estimates for the Q matrix of disturbances documented in Table A.4 of the Appendix). In NYC, however, all three rates initially react with roughly similar magnitudes, possibly indicating that the dominant type of shock in NYC acts much like a shift of the entire term structure. In Boston, by contrast, the dominant shock, by far, is most pronounced in the 10-year rate. This may potentially reflect more rigidity in the way shorter-term demand for and supply of space shift in Boston. After the initial shock, all key rates decay back to their unconditional means, but much more slowly in NYC.

From the estimated model parameters, we can use the Kalman filter to produce optimal estimates for the time-series of key rates. There are two ways to estimate these time series. One can produce an estimate of key rates each quarter that depends only on information available up to that quarter. This is termed a "filtered" estimate. Alternatively, each quarter, one can employ all observed information (including observations from subsequent quarters) to impute the key rates. This is termed a "smoothed" estimate. We opt to report the latter because our goal is to provide our best possible estimate for the unobserved key rates.

Figure 3 plots blended smoothed estimates of the key rates over our sample period, for both NYC (solid lines) and Boston (dotted lines). Several features stand out. Firstly, there is substantial time-variation in the level of the key rates in both cities. Second, one can see that the 10-year lease rates are generally more volatile than the shorter maturity rates, as suggested by the impulse-response exercise of Figure 2. This behavior in the case of NYC leads to a near-overlap of the spot and 10yr rate

[26] Unlike oil, which can be stored if not consumed, a spot claim to space cannot be "stored". In that sense, space is much like electricity. It is produced by the asset (i.e., the property) but must be either consumed or forgone.

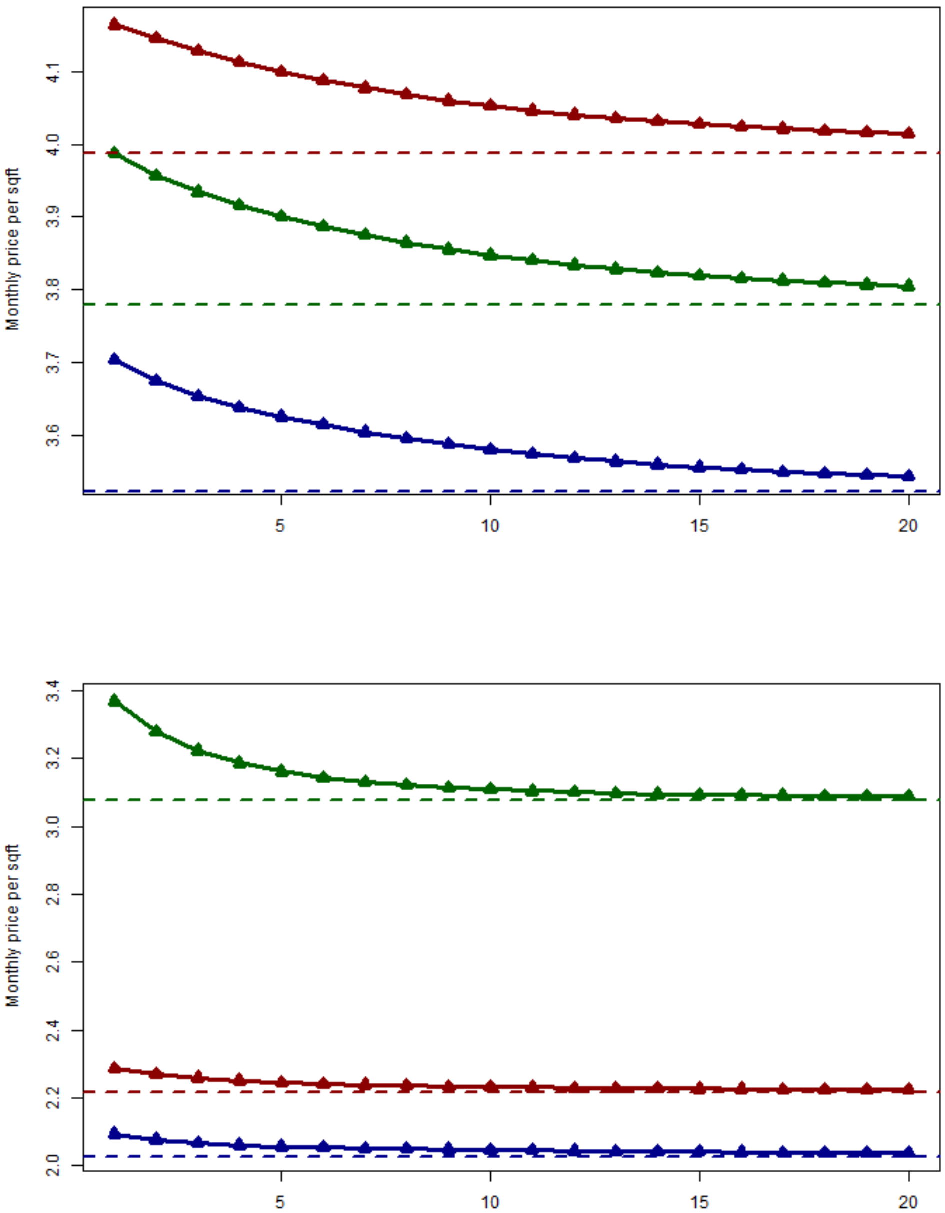

Figure 2: **Impulse-response analysis**. The figure shows the estimated response of office lease key rates to a one-standard-deviation shock to the dominant fundamental factor underlying the dynamical system. The blue lines corresponds to the spot rate, while the red and green lines denote the 5yr and 10yr rates, respectively. In both NYC (upper panel) and Boston (lower panel), the shock impacts the 10-yr strongly. Dashed horizontal lines represent blended estimates for the unconditional mean.

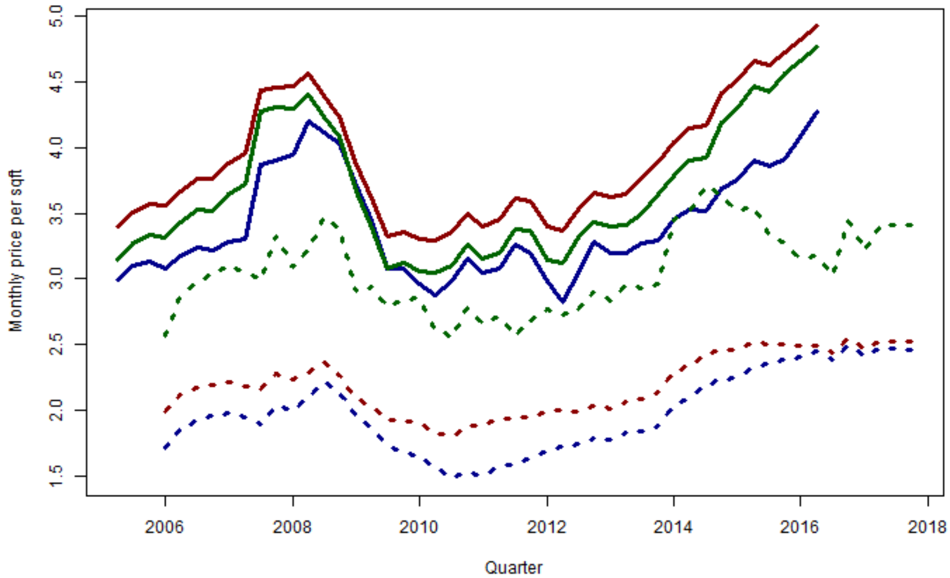

Figure 3: **Term structure of lease rates**. The figure depicts blended estimates for the key forward and spot rates for NYC (solid lines) and Boston (dotted lines) office properties, as given by the Kalman smoother. The blue lines corresponds to the spot rate, while the red and green lines denote the 5yr and 10yr rates, respectively.

in the periods immediately following the GFC.

Our smoothed key rates allow us to explore the shape of the term structure of office lease rates. Figures 4 and 5 plot the blended estimates for the slope and curvature, respectively, of the term structures implied by the key rates. The slope is calculated as $\frac{1}{10}(F_{t,120} - F_{t,0})$ while the curvature is $(F_{t,120} - 2F_{t,60} + F_{t,0})$. The top (bottom) plot in each figure corresponds to the NYC (Boston) term structure. Because the key rates are estimated rather than directly observed, we additionally plot two-standard-deviation confidence intervals for the slope and curvature.[27]

One trait that the plots highlight is that, although the NYC term structure is typically upward sloping with negative curvature, the slope substantially shrank during 2009, on the heels of the Great Financial Crisis. Boston, on the other hand, displays relatively more stable measures of the shape of its term structure, exhibiting an upward-sloping and convex term structure throughout the sample period. Another

[27]The Kalman smoother generates a covariance matrix for the estimation errors of the key rates at each quarter. These covariances matrices are then blended following the process described in Appendix A.3 to calculate the confidence intervals.

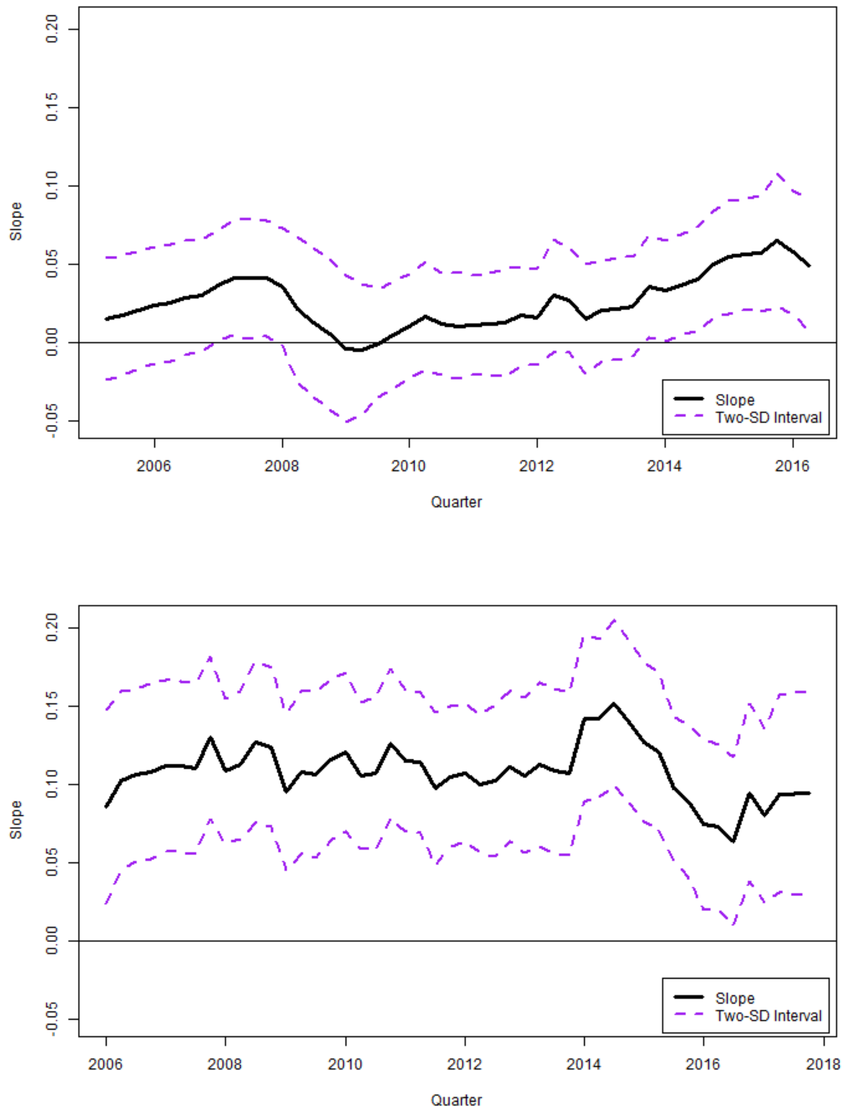

Figure 4: **Slope of the term structure of lease rates**. The figure shows blended estimates for the slope of the term structure measured as $\frac{1}{10}(F_{10} - F_0)$, one tenth of the difference between the 10-year forward and the spot rates for NYC (top) and Boston (bottom) office properties, as given by the Kalman smoother. The dashed lines represent two-standard-deviation confidence intervals.

takeaway is that the slope is much more pronounced in Boston than in NYC. In the next subsection we discuss how factors like growth expectations and physical depreciation/obsolescence might contribute to the shape of the term structure.

Our results so far speak to the price of space occupancy for the average property in terms of quality, i.e., the case in which the *hedonic noise* component, ε_i, is

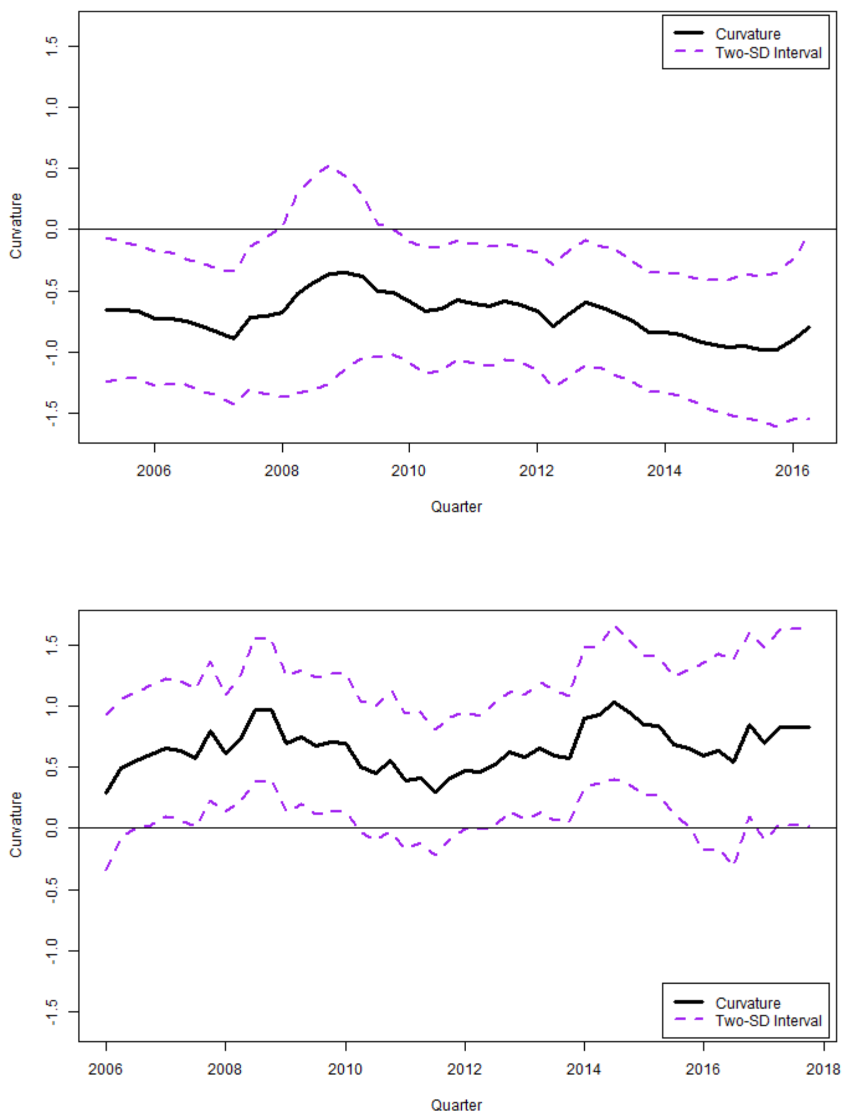

Figure 5: **Curvature of the term structure of lease rates**. The figure shows blended estimates for the curvature of the term structure measured as $F_0 + F_{10} - 2F_5$, the sum of the spot and the 10-year minus twice the 5-year forward rates for NYC (top) and Boston (bottom) office properties, as given by the Kalman smoother. The dashed lines represent two-standard-deviation confidence intervals.

zero.[28] However, given our estimates for the key rates and the variance of the quality adjustment factor, we can infer the behavior of rental rates across the quality distribution. Specifically, by using our estimate for the standard deviation of the quality adjustment factor, σ_ε, we can vary g_i to obtain average rental rates for office

[28]This corresponds to a quality adjustment factor g_i equal to 1.

space whose quality is, for example, within two standard deviations from the median quality. We plot these intervals for the 10-year key rate in Figure 6. In the figure, the solid white lines depict the 10-year rate at median building quality and the green solid lines above (below) bound a two-standard-deviation interval.

Consistent with our estimates of quality dispersion (Table 3), there are marked discrepancies in the range of long-term office rental prices between Boston and NYC. Figure 6 shows that most office properties in Boston have a 10-year lease key rate between \$2 and \$8 per square foot, throughout our sample period. The range for New York, on the other hand, is much wider: While the 10-year price for low-quality office space was around \$2 per square foot throughout the sample period, high-quality office space reached approximately \$12 per square foot both before the GFC and at the end of our sample period. The difference in ranges could be explained by, among other things, a greater degree of diversity in the physical attributes of NYC office buildings (e.g., there are more trophy properties in NYC).

4.1 Discussion

It is legitimate to ask whether our estimates of the key rate time series, and therefore the term structure of lease rates, are grounded in substance rather than driven by spurious variables for which we've failed to control. Moreover, it is equally legitimate to ask what kind of economic forces might lead to the imputed shapes of the term structure in our data. We begin by addressing the latter question, and then turn to discussing possible unobserved influences.

4.1.1 Comparison with Other Commodities

To gain an appreciation for what may drive lease forward rates, it seems natural to compare them with other types of forward contracts. The most important economic force restricting the term structure of forward contracts is the possibility of arbitrage. The price specified by a τ-years forward contract on a dividend paying asset with current value S_t must be given by

$$F_{t,\tau}d_{t,\tau} = S_t - pv_t[Div(t, t+\tau)] + pv_t[f(t, t+\tau)],$$

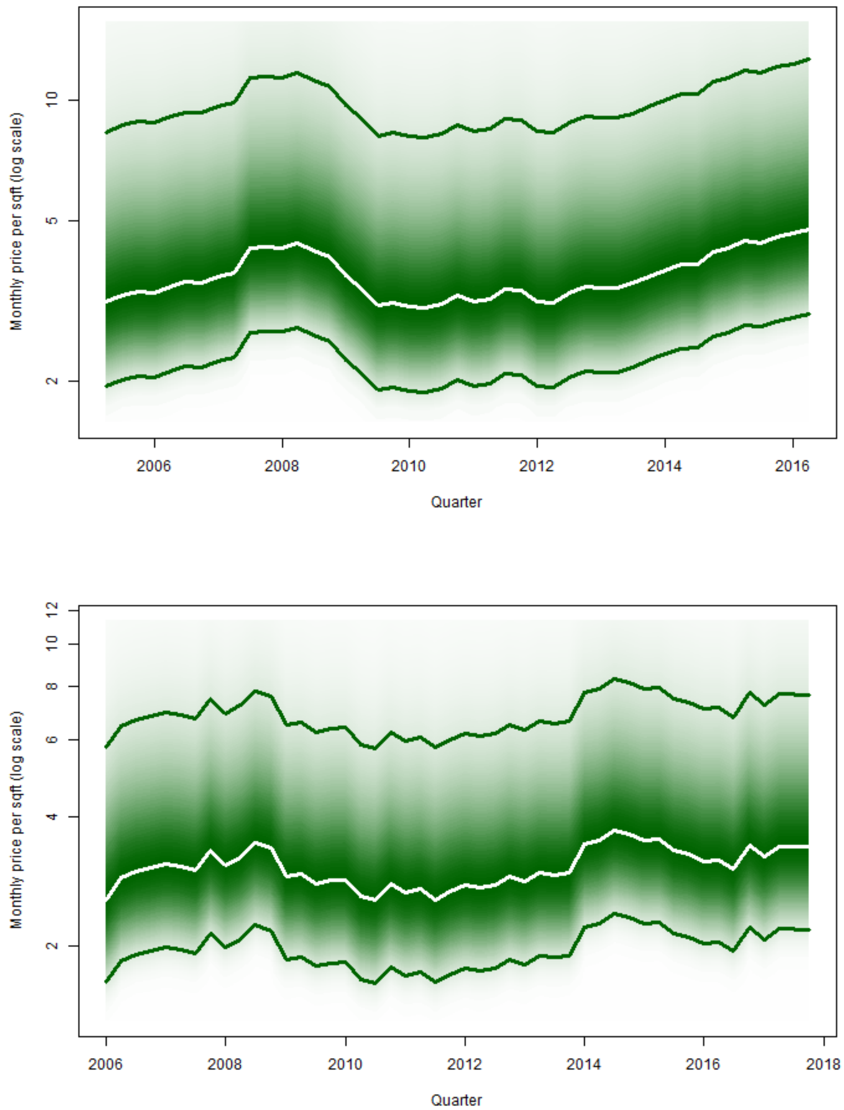

Figure 6: **Quality dispersion around the 10yr key rate**. The figure depicts how the 10yr key rate varies continuously along the distribution of quality for NYC (top) and Boston (bottom) office properties. The solid white line corresponds to the median quality ($g_i = 1$), while the solid green lines denote two-standard-deviation intervals. The shaded area illustrates the quality density.

where $d_{t,\tau}$ is the price of a \$1 zero coupon bond maturing at $t+\tau$, $pv_t[Div(t, t+\tau)]$ is the present value of all benefits ("dividends") derived from ownership of the asset between t and $t+\tau$, and $pv_t[f(t, t+\tau)]$ is the present value of a market friction component, $f(t, t+\tau)$, that must be in the interval $[-bc(t,\tau), sc(t,\tau)]$. Here, $bc(t,\tau)$ is the cost to borrow (or short-sell) the asset between t and $t+\tau$, and $sc(t,\tau)$ is the cost to store the asset between t and $t + \tau$. In plain language, absent market frictions, the forward price is the spot, less the capitalized value of cash flow forgone before taking delivery of the asset. As most textbooks describe it, deviations from this pricing in a frictionless market leads to an arbitrage opportunity from cash-and-carry or reverse cash-and-carry strategies. In the former, when the forward price is too high, one borrows money risk-free, purchases and stores the asset while simultaneously selling it forward, invests the dividend cash flow until delivery, and upon delivery uses the forward sale price to pay off the debt. In a reverse cash-and-carry, if the forward price is too low, the asset is sold short while simultaneously purchasing the asset forward, the proceeds are invested in a risk-free bond and used to pay off the dividend commitments to the asset holder, and the remainder is used to pay the forward price and settle the asset loan.

The friction working against a (reverse) cash-and-carry strategy is (short-selling) storage costs. The important point we wish to make is that the frictions associated with executing a cash-and-carry strategy (or its reverse) in the space market are prohibitively large, meaning that arguments based on arbitrage are unlikely to bound forward prices in a meaningful way. To see this, contrast the forward market on space with the forward market on, say, copper. Firstly, the commodity underlying a forward lease component is not homogeneous across time. For example, the 10-year forward component of a lease corresponds to space in a property that is ten years older than the underlying space for a spot component. The quality of copper, on the other hand, is essentially time-invariant. Second, and relatedly, the copper promised in a ten-year forward commitment can come from anywhere so long as its purity exceeds certain standards; by contrast, it would be generally prohibitively expensive for a tenant to agree to accept the promised space in the ten-year forward component of a lease from anywhere other than the same source (i.e., building) as the spot commitment of space.[29] Finally, a forward component of a lease delivers occupancy to be consumed at a set period of time, whereas a forward commitment to copper delivers a good that can be consumed anytime once it is acquired (subject to storage costs). In other words, the benefits from the delivery of a lease forward cannot be "stored", and this restricts cash-and-carry forms of arbitrage. In this regard, lease forwards

[29] Correspondingly, individual forward lease commitments cannot be costlessly stripped from a lease agreement. I.e., few office tenants can move their operations for only a single month while their originally leased space is let to another entity for only that month.

in the space market resemble those in the electricity market where storage costs are prohibitive and thus a claim delivered today must be consumed today or foregone. Correspondingly, reverse cash-and-carry is not possible because the cost to borrow spot space is the full spot rental rate whereas the cost to borrow a non-ephemeral commodity (like copper) corresponds more closely to the cost of deferring its use.[30]

To summarize, inherent inhomogeneities and the inability to meaningfully store the benefits that a lease delivers obviate typical no-arbitrage constraints on lease forwards that might otherwise apply to standard commodities. It is worth emphasizing that the noted features of a space market hold for other forms of leased capital (e.g., airplanes, ships, infrastructure, etc.).

4.1.2 Factors contributing to the shape of the term structure

Our estimates suggest a term structure that is often upward sloping and concave in the case of NYC, and convex in the Boston office market. Below, we list some economic influences that might explain the imputed shapes.

Anticipated price increases Inflation and anticipated increases in demand for space relative to supply will place upward pressure on forward lease rates. If the market anticipates rental growth rates that decelerate from an initial high level, then this can lead to an upward sloping and concave term structure. Alternatively, a convex term structure could arise if market participants anticipate demand for office space to accelerate relative to its supply over the medium (or long) term.

Quality deterioration and obsolescence For a given degree of quality, the leases we "unbundle" are only guaranteed to have that quality when the lease is executed. Properties deteriorate and become obsolete, and it's considered received wisdom in the commercial real estate industry that, absent intensive capital investment, most properties will drift down the quality spectrum as newer or more updated space becomes locally available. This means that the 10-year forward lease rates we impute

[30] The analogy with electricity would have us compare the property itself to the generation plant. One may be able to store or borrow the plant, as one would with a building, but one cannot (at this point) meaningfully store a month's production of electricity from a large plant, and to borrow one month's production one would have to pay its full price because its use could not be deferred.

represent space that, on average, has depreciated relative to the spot market.

In other words, one would expect deterioration and obsolescence to contribute towards a downward sloping term structure. Moreover, it should be clear that accelerating (decelerating) depreciation and obsolescence would contribute negatively (positively) to the convexity of the term structure. Because physical depreciation and obsolescence may accelerate or decelerate at different points in a building life cycle, the overall market pattern may depend on the vintage of the overall stock.

Credit losses. Typical forward agreements mitigate counterparty risk using margin accounts. Although leases may incorporate some kind of escrowed funds (e.g., damage deposits), these may not be enough to offset losses from lease defaults. We are unable to control for tenant credit losses in the imputation of forward lease rates, meaning that they may incorporate some credit spread, though arguably less than corporate bonds.[31] With high-quality bonds, one also typically observes default rates that increase with term — this is because a well-underwritten investment-grade issuance is less likely to default in its first year than in subsequent years, and because credit is mean-reverting.

To compensate for possible default, the corresponding forward lease rate would have to increase.[32] In turn, and absent a correlation between lease maturity and tenant risk, considered below, the default dynamics described above should contribute positively to both the slope and convexity of the term structure.

Lease renewal options. Like the option to default on their lease commitment, a renewal option can only be exercised by the tenant. Although renewal can reduce lease commissions paid by landlords, one may generally expect that their value is greater to the tenant, meaning that their value should be expressed through higher rents. Because the "cost" of the renewal option must be "paid" throughout the lease term, the existence of renewal options could, all else equal, lead to higher rents in short-term leases than in long-term leases, i.e., it would bias down the imputed slope and convexity of the term structure.[33]

[31]Leased assets can arguably be recovered and more efficiently redeployed by owners than is the case with secured assets and lenders.

[32]This is confirmed in Brueckner and Rosenthal (2024).

[33]Lease renewal options tend to contain a market rent clause, meaning that the renewal lease rate would be market determined.

Assortative matching (maturity selection bias) Our methodology relies on the assumption that differences in term across tenants and space (for same-class leases) is statistical noise. If there are systematic differences in the quality of tenants or space across lease terms, then this will lead to systematic distortions in our imputed term structure of lease forwards.

Examples of such systematic differences include the possibility that greater default risk is associated with mid-maturity leases. This could happen, for instance, if landlords will only sign long-term leases with the highest-credit tenants. In such a case, long-maturity leases will have lower effective rents as would shorter-term leases for reasons mentioned earlier. If risk concentrates in mid-maturity leases, then this would contribute to negative convexity in the term structure. Although we cannot rule this out, we do note that, while NYC's properties display negative convexity, the curvature of Boston's office term structure is positive throughout the period of analysis. It is unclear why risk concentration would be more pronounced in mid-term NYC leases, relative to contracts in Boston. Likewise, if renewal options are more prevalent in NYC mid-term leases then this too would enhance the negative convexity of the term structure. Here too, it is not clear how such an explanation accounts for the difference across markets. Overall, to investigate these possibilities further, we would need to explore a complementary dataset that included details on renewal options.

By summarizing the influences identified above, one can attempt to account for the observed term structure of lease rates over our sample period in the absence of selection bias. Specifically, anticipated growth will contribute to an upward sloping term structure. Natural depreciation and obsolescence will pull the term structure the other way. The term structure we observe in NYC can be rationalized by dynamics in rental growth expectations that are themselves concave (i.e., greater growth rates in the short/mid- than in the long-term), and in the long-run are further slowed by increasing expected quality depreciation. In the case of Boston, the term structure is convex: both the slope and curvature are positive throughout the sample period. The pattern is consistent with stable high and positive growth expectations over the long term, with the corresponding compounding leading to convexity. The latter could also be supported by a decelerating pattern of physical depreciation.

5 Application

By estimating a dynamic model of the term structure of leases, we are in a unique position to address questions concerning the risk and reward associated with spatial market strategies. For example, the eventful corporate story of WeWork has focused attention on the viability of short-term leasing of office space. Co-working companies, such as WeWork, provide high-quality office space and amenities/services on a short-term basis where the contracts can range from a day to a year. One element of the business strategy consists of obtaining long-term rights to a space, and then offering it to many short-term users. Profitability and risk can arise separately from what we refer to as the spatial and service components of the business strategy. The spatial component consists of benefiting from the difference between short-term and long-term lease rates and/or the intensification of short-term usage of leased space. The latter consists of increasing the number of users of the space per unit of area. Moreover, the space may be more intensively used and charged a premium for shorter duration contracts, e.g., a liquidity premium for shorter term contracts to the tenant.[34] The service component comes from profit margins created by the provision of amenities (e.g., internet usage, food/beverage, document printing, storage, furniture, collective buying of insurance and financial services, and even atmosphere).

The success of a co-working business relies on both the spatial and the service components. Although we are not in a position to assess the risk and reward from the service component, our model permits us to analyze the spatial component of a co-working business. In particular, we take the position that the more the business model depends on the service component to be profitable, the more it resembles service-intensive real estate investments (e.g., hotels) rather than traditional office properties. Indeed, our analysis in this section suggests that, absent non-trivial intensification or skill in obtaining long-term claims to high-quality space at below market prices, the spatial component is a drag on profits and a source of substantial financial risk.

Our analysis is based on estimating, in each quarter of our sample period, the profits and risk from paying to lease a median-quality space for 10 years, and financing this by turning around and leasing the space to users on a short term (quarterly)

[34] Chegut and Langen (2019) found from a sample of 151 coworking providers in NYC that charges for a desk per month increase for daily and monthly contracts relative to quarterly contracts.

basis. Because our model allows us to both forecast the dynamics of lease rates *and* calculate the standard error around that forecast, we can calculate both the expected profit and the profitability standard deviation from this strategy. Moreover, we can vary assumptions about average vacancy, intensification, and skill required to secure

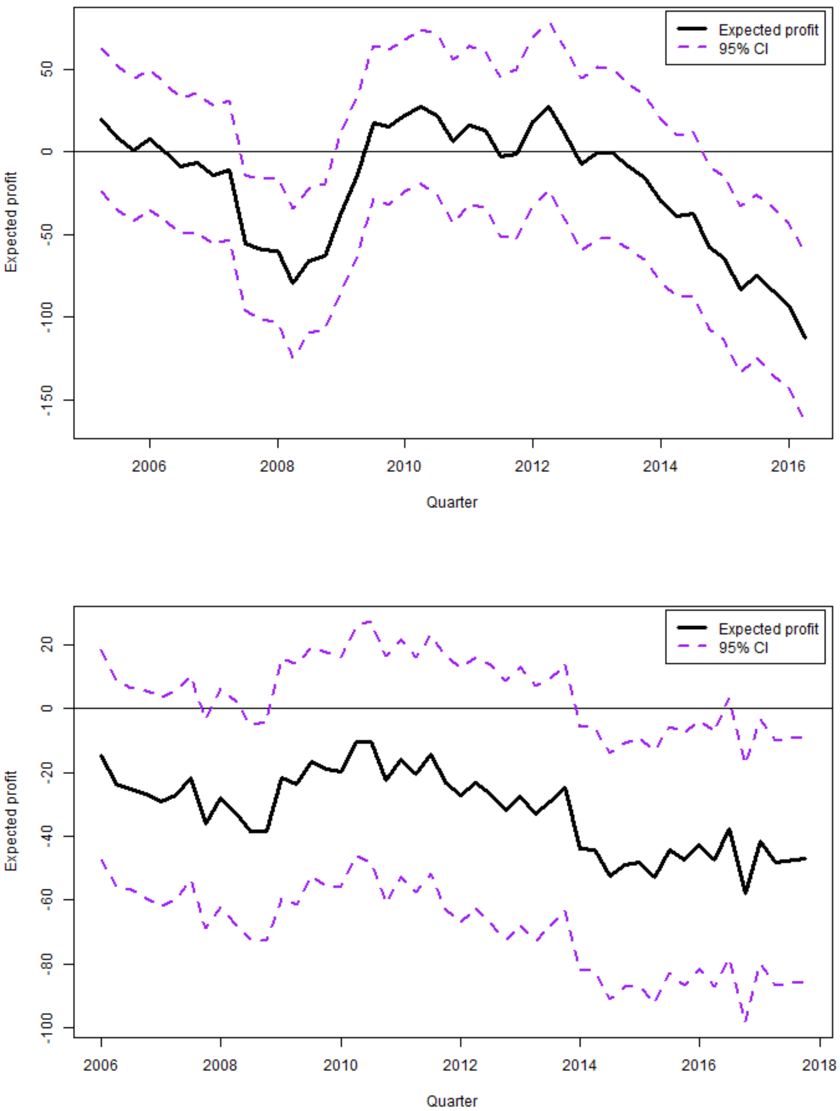

Figure 7: **Present value of profit pSqFt from long-short space strategy assuming 100% occupancy**. The figure depicts the time series of expected profits for a strategy in which a 10-year lease is financed by leasing the space on a rolling quarterly basis. Estimated model-derived profits from NYC (Boston) office properties are depicted in the top (bottom) panel. The strategy assumes 100% occupancy (there is a short-term occupant in every quarter over the 10 years). The profits are in dollars per square foot and discounted to the present using the contemporaneous zero-coupon yield curve. Dashed lines denote 95% confidence intervals.

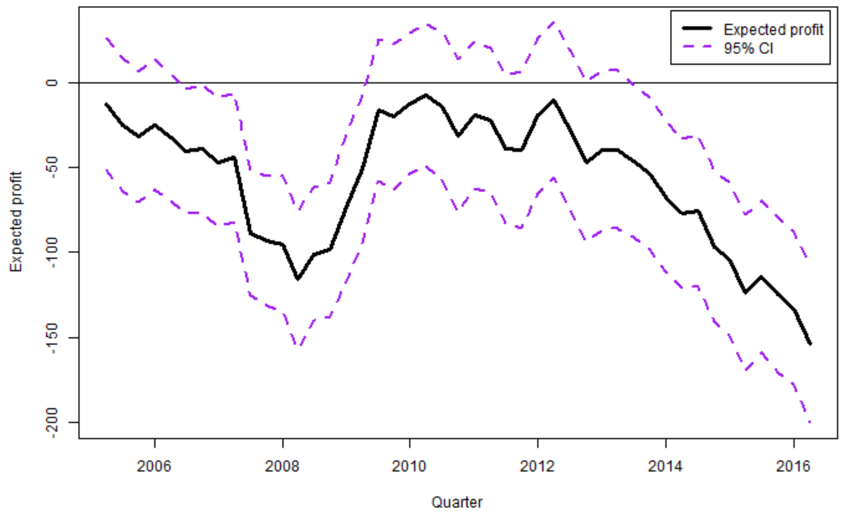

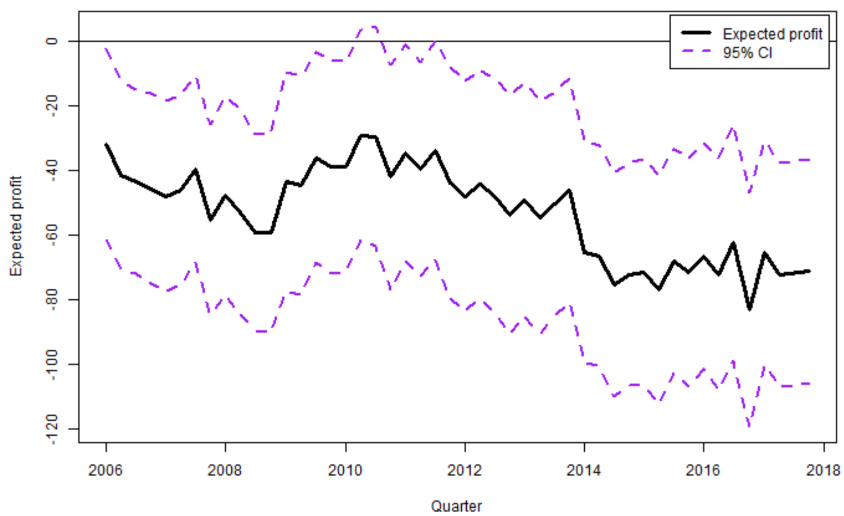

Figure 8: **Present value of profit pSqFt from long-short space strategy assuming 90% occupancy**. The figure depicts the time series of expected profits for a strategy in which a 10-year lease is financed by leasing the space on a rolling quarterly basis. Estimated model-derived profits from NYC (Boston) office properties are depicted in the top (bottom) panel. The profits are in dollars per square foot and discounted to the present using the contemporaneous zero-coupon yield curve. Dashed lines denote 95% confidence intervals.

Appendix A.5 details how we calculate the expected profit and associated standard deviation for the strategy.[35] Figure 7 depicts the results. To fix ideas, consider executing the strategy in 2010Q1 in the NYC market. We first calculate the *filtered*

[35]For the sake of simplicity, we use the parameters estimated by maximizing the Kalman filter's likelihood. A comparison between the 10-year key rates corresponding to the combined estimates and the maximum likelihood estimation can be found in Figure A1.

key rates in 2010Q1 from the Kalman filter (using only information dating from 2010Q1 or earlier). These roughly correspond to the smoothed 2010Q1 key rates seen in Figure 3. The lease forward key rates can then be used to price an average 10-year lease (i.e., derive the average rent that would have been paid on a 10-year lease in 2010Q1). Next, we use the 2010Q1 filtered key rates to forecast quarterly lease rates for the next 40 quarters, together with their Gaussian forecast error. The difference between the forecasted quarterly lease rates and the fixed 10-year rental payment is then discounted to the present using 2010Q1 zero coupon bond prices. This is a measure of the (discounted) expected profitability of the strategy, as might have been assessed in 2010Q1. The forecast errors are then used to derive a 95% confidence interval for this quantity. The procedure is then repeated for 2010Q2, and so on.[36]

Figure 7 indicates that during the sample window, the strategy was not expected to be profitable in NYC office properties without the charge of liquidity or service premiums. This is because the lease term structure is mostly upward sloping. According to the point estimates, the only times when the strategy would have been profitable was a brief period in 2005-2006 and during the years immediately following the GFC. Even in those two cases, however, expected profits are statistically indistinguishable from zero, as shown by the confidence intervals. The bottom panel of Figure 7 shows that the performance of the strategy when applied to Boston office properties is comparably poor: The point estimates for the expected profits are negative throughout the sample period. In calculating the profitability of the long-short space strategy outlined above, we assumed that the space was successfully leased out on a short-term basis in each of the 40 quarters that the 10-year lease was in place. Under a more reasonable 90% occupancy assumption, there are no quarters in which the point profit estimates are positive for either NYC or Boston. In other words, the strategy is unprofitable at any point in time during our sample period assuming a 10% vacancy rate. This is depicted in Figure 8.

To provide a sense of the riskiness involved, we calculate an *ex ante* "Sharpe Ratio" for the various strategies. For a strategy originated in a given quarter, this is done by dividing the annualized present value of strategy profits by the corresponding

[36]Because quality in our model corresponds to a multiplicative adjustment, the profitability results using for space that is different from median quality would be a multiple of the results depicted in Figure 7.

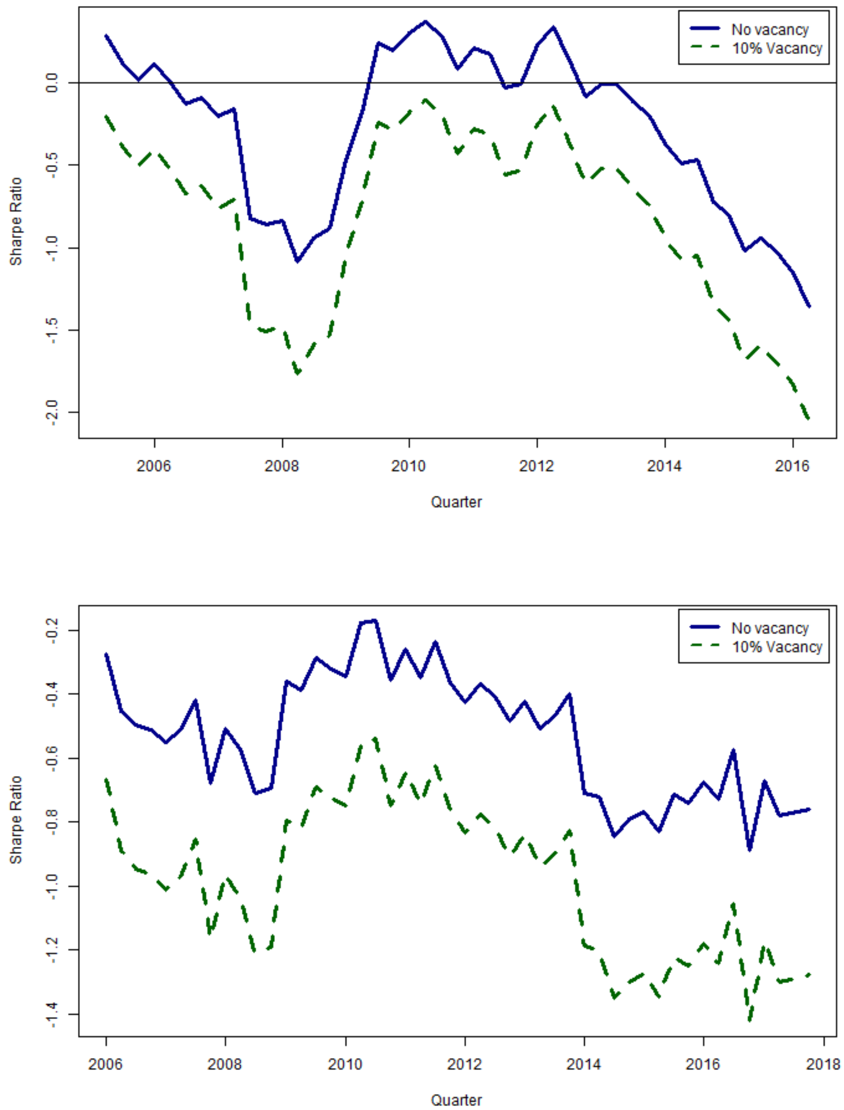

Figure 9: **Sharpe ratios of long-short space strategy**. The figure depicts the time series of ex-ante Sharpe ratios for a strategy in which a 10-year lease is financed by leasing the space on a rolling quarterly basis. Estimated model-derived ratios for NYC (Boston) office properties are depicted in the top (bottom) panel. The Sharpe ratio is calculated by dividing the annualized expected strategy profits in a given quarter by the corresponding annualized standard deviation of the profits.

annualized standard deviation of those discounted and aggregated profits. Details are in Appendix A.5. We show the resulting time series in Figure 9. For comparison, consider that typical equity market risk corresponds to annualized Sharpe Ratios of roughly 50%. What the plots show is that even when the long-short space strategy is profitable (in the case of NYC and at full occupancy), it does not provide sufficiently high compensation for the risk entailed. By and large, our takeaway is that without additional advantages, the long-short space strategy yields negative profits at substantial risk.[37]

5.1 Turning a profit in a long-short space strategy

The preceding analysis of risk and reward associated with the spatial component of a co-working business model assumes that space is acquired and re-leased at average market rates. Table 3, which documents estimated dispersion in observed effective rents, indicates that lease rates can vary widely from average market rates. This suggests that a talented (or lucky) negotiator may be able to obtain a below-market rate on the 10-year lease.[38] Likewise, our analysis assumes that the use intensity of the space obtained with the 10-year lease is equal to the use intensity of the short-term renter. Because co-working businesses may be able to configure their space to accommodate a higher density of users per square foot and (or) charge a liquidity premium for short-term use of space, they may be able to generate more revenues through increased intensity. However, should profit margins from services or liquidity premiums decline, the strategy may become a liability for the landlord during unexpected times of distress in the market.

To understand the role that negotiating or"deal scouting" skills may play in the profitability of co-working, we focus on the maximum long-term office rent that would make the spatial strategy sustainable. For each quarter in our sample period, we shift the 10-year lease rate until the long-short space strategy breaks even (zero profits) at 90% occupancy. We then continue to shift the 10-year rent until the strategy achieves a 50% Sharpe ratio (also at 90% occupancy). Next, we incorporate the observation error to calculate the percentile in the distribution of median-quality leases to which these thresholds correspond. The results are depicted in Figure 10. In the case of NYC, the 10-year lease rate would have to be in the bottom tercile of the lease distribution (holding quality constant) for the strategy to be profitable throughout most of our sample. That said, the *skill* required to find a sufficiently low long-term lease rate was substantially lower during the years following the GFC, when, at times, rates slightly below the median allowed the strategy to break even.

[37]We note that, under our model assumptions, the Sharpe Ratio we calculate is invariant to the quality of the space.

[38]The results in Antunes Batista da Silva, Liu and Hutchison (2025) suggest that co-working companies, indeed, are generally able to sing long-term leases at lower rates than other tenants.

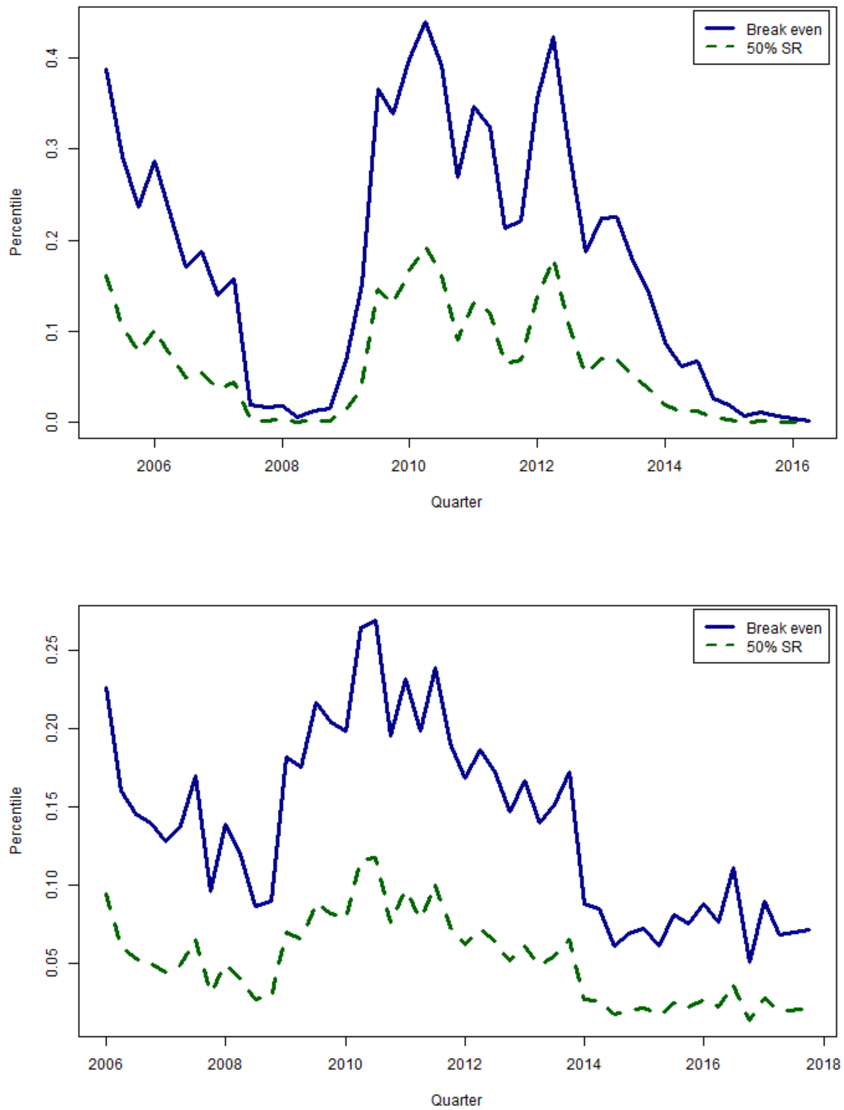

Figure 10: **Percentage of profitable 10-year leases in a long-short strategy**. The figure displays an estimate of the proportion of NYC (top) and Boston (bottom) office leases that would allow an investor to achieve different levels of profitability. For each quarter in our analysis, we calculate the the 10-year lease rate that allows the strategy to break-even (blue, solid line) or attain a 50% Sharpe ratio (green, dashed line), and show its rank in the conditional distribution of 10-year leases given that year's observation error variance (Table 3). Occupancy of 90% is assumed.

In contrast, less than 5% of long-term leases would have allowed the strategy to break even between 2007 and 2008. More importantly, typically less than 10% of available 10-year leases would permit a competitive investment Sharpe Ratio.

The bottom panel of Figure 10 shows that co-working companies in Boston may struggle even more to find 10-year rental rates that lead to positive profits for their spatial strategy. The maximum 10-year break-even lease rate for the strategy in Boston lies below the bottom quartile for all but two quarters in our sample period. Competitive risk-return tradeoff would have been even harder to achieve.

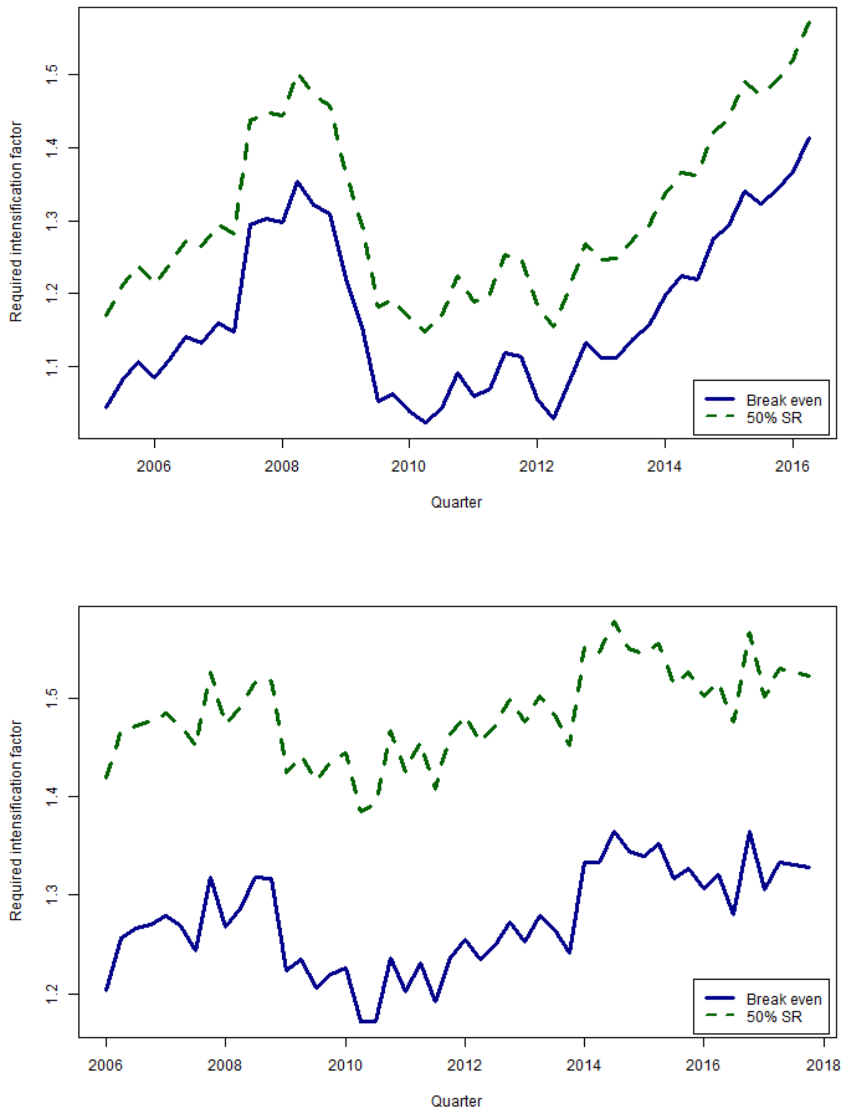

Figure 11: **Required intensification factor for profitable long-short strategy**. The figure displays the factor by which short term lease rates must be multiplied for the strategy to achieve different levels of profitability (zero-profit or 50% Sharpe ratio). The top graph corresponds to NYC office leases, while the bottom one depicts the required intensification factor for Boston leases. Occupancy of 90% is assumed.

Figure 11 depicts the threshold of intensification factor needed to achieve zero profitability or a 50% annualized Sharpe ratio at 90% occupancy. The intensification factor is used to multiply the short-term rents from the strategy. For instance, an intensification factor of 1.20 corresponds to a 20% increase in the short-term rents achieved throughout the 10 years the strategy is in play. In NYC, except for the periods during which the strategy is profitable according to the point estimates (2005-2006 and immediately after the GFC), the break-even intensification factor was above 10% during most of the sample period. To achieve a risk-adjusted compensation comparable to the stock market (50% Sharpe Ratio), however, NYC co-working firms would have needed to obtain above 20% higher revenue per square foot, with the intensification factor reaching 50% during the crisis and by the end of our sample. Consistent with the systematically low profitability of the spatial strategy in Boston, the required intensification factors in that city are higher than in NYC. To break even, a co-working firm would have needed to raise revenue per square foot by over 20% during the whole period of analysis. Achieving a 50% Sharpe Ratio, on the other hand, would have entailed an intensification factor of at least 1.4, and consistently around 1.5 during the final quarters in our sample.

In practice, skill at obtaining space at cheaper prices and intensification of use may be combined. Overall, our analysis allows us to quantify metrics that would translate into a profitable strategy as well as those that would translate into acceptable risk-adjusted returns. We re-emphasize that our focus here is on the spatial component of the co-working strategy. To the extent that services and liquidity are key to co-working profitability, its risk profile might better resemble that of service-intensive real estate (e.g., hotels).

6 Conclusion

In the large market for leasing real assets, rental income is fundamental to understanding valuation and return dynamics. However, relatively little work has been done to explore how the pricing of newly originated leases with different maturities evolves over time. We develop a methodology for unbundling leases into constituent forward contracts and imputing the dynamics of the subsequent term structure, all while allowing for asset quality heterogeneity. We then demonstrate the methodology using data on NYC and Boston office lease transactions from Compstak and JLL, from 2005 to 2017. We are able to characterize the dynamics of the term structure in each city and then use the estimated model to shed light on the viability of co-working as a business strategy.

Lease rates imputed from our model, especially for longer-maturity forward rates, could be informative for real asset markets in the same way that commodity forward rates are informative about the macroeconomy. In addition, because our methodology estimates a dynamic model, it is possible to use it to quantitatively study and understand optimal vacancy (i.e., inventory) strategy and optimal lease portfolio construction. Finally, by characterizing the distinct features of a term structure of lease rates, we inform the economic models that seek to study such markets.

References

Agarwal, Sumit, Brent W. Ambrose, Hongming Huang, and Yildiray Yildirim. 2011. "The term structure of lease rates with endogenous default triggers and tenant capital structure: Theory and evidence." *Journal of Financial and Quantitative Analysis*, 46(2): 553–584.

Aldana, Ernesto, and Jacob S. Sagi. 2025. "On constraining the sign and decay behavior of impulse response functions in vector autoregressions." *Working paper.*

Ambrose, Brent W., and Yildiray Yildirim. 2008. "Credit risk and the term structure of lease rates: a reduced form approach." *The Journal of Real Estate Finance and Economics*, 37(3): 281–298.

Ambrose, Brent W., Patric H. Hendershott, and Małgorzata M. Kłosek. 2002. "Pricing upward-only adjusting leases." *The Journal of Real Estate Finance and Economics*, 25(1): 33–49.

Antunes Batista da Silva, Fernanda, Nan Liu, and Norman Hutchison. 2025. "Office term structures and the tenancy of flexible workspace providers." *Journal of Property Research*, 1–25.

Brueckner, Jan K, and Stuart S Rosenthal. 2024. "Tenant Riskiness, Contract Length, and the Term Structure of Commercial Leases."

Chegut, Andrea, and Mike Langen. 2019. "The financial impacts of coworking: rental prices and market dynamics in the commercial office market." *Available at SSRN 3481142*.

Chiang, I-Hsuan Ethan, W Keener Hughen, and Jacob S Sagi. 2015. "Estimating oil risk factors using information from equity and derivatives markets." *The Journal of Finance*, 70(2): 769–804.

Clapham, Eric, and Åke Gunnelin. 2003. "Rental expectations and the term structure of lease rates." *Real Estate Economics*, 31(4): 647–670.

Eisfeldt, Andrea L, and Adriano A Rampini. 2009. "Leasing, ability to repossess, and debt capacity." *The Review of Financial Studies*, 22(4): 1621–1657.

Fragoso, Tiago M, Wesley Bertoli, and Francisco Louzada. 2018. "Bayesian model averaging: A systematic review and conceptual classification." *International Statistical Review*, 86(1): 1–28.

Gavazza, Alessandro. 2011. "Leasing and secondary markets: Theory and evidence from commercial aircraft." *Journal of Political Economy*, 119(2): 325–377.

Grenadier, Steven R. 1995. "Valuing lease contracts A real-options approach." *Journal of Financial Economics*, 38(3): 297 – 331.

Grenadier, Steven R. 2005. "An equilibrium analysis of real estate leases." *The Journal of Business*, 78(4): 1173–1214.

Kilian, Lutz, and Helmut Lütkepohl. 2017. *Structural vector autoregressive analysis.* Cambridge University Press.

McConnell, John J., and James S. Schallheim. 1983. "Valuation of asset leasing contracts." *Journal of Financial Economics*, 12(2): 237–261.

Miller, Merton H., and Charles W. Upton. 1976. "Leasing, buying, and the cost of capital services." *The Journal of Finance*, 31(3): 761–786.

Rampini, Adriano A, and Sridhar Viswanathan. 2013. "Collateral and capital structure." *Journal of Financial Economics*, 109(2): 466–492.

Sagi, Jacob S. 2021. "Asset-level risk and return in real estate investments." *The Review of Financial Studies*, 34(8): 3647–3694.

Schallheim, James S., and John J. McConnell. 1985. "A model for the determination of "fair" premiums on lease cancellation insurance policies." *The Journal of Finance*, 40(5): 1439–1457.

Stanton, Richard, and Nancy Wallace. 2009. "An empirical test of a contingent claims lease valuation model." *Journal of Real Estate Research*, 31(1): 1–26.

Takahashi, Masashi, and Jiro Yoshida. 2025. "Measuring Expectations in Real Estate Markets from Long-Term Lease Contracts." *Available at SSRN 5733262.*

Uhlig, Harald. 2005. "What are the effects of monetary policy on output? Results from an agnostic identification procedure." *Journal of Monetary Economics*, 52(2): 381–419.

Wheaton, William C., and Raymond G. Torto. 1994. "Office rent indices and their behavior over time." *Journal of urban Economics*, 35(2): 121–139.

A Appendices

A.1 Data Appendix

Below we provide a description of each individual data set and describe how we combined the information from both sources to obtain our working samples for both New York City and Boston.

Our raw data sets contain records of individual leases in New York City and Boston. Table 1 reports the number, by year, of leases for which we observe the full rent schedule. In the CompStak data set, the rent schedule is directly recorded for most leases. For a fraction of the leases with missing information in the Rent Schedule field, we compute it by combining the information from Starting Rent and Rent Bump USD/Rent Bump Percent. The rent schedule in the JLL data includes the sequence of prices that will be charged throughout the life of the lease but does not specify the lengths of the periods during which each of these prices will be charged. For simplicity, we assume periods of equal length, given by the ratio of lease term to number of periods in the rent schedule. From each data source, we remove all leases executed in quarters in which less than 30 contracts were originated. Tables A.1 and A.2 provide summary statistics by city for JLL and Compstak, respectively.

We use these data points to calculate effective rents and determine the periods of occupancy. Starting from the rent schedule, we construct, for each lease, a series of monthly cash flows, assuming that payments take place on the first day of each occupancy month.[39]. We adjust these cash flows according to the concession information available for each lease. First, we account for free rent months by setting the

[39]These cash flows correspond to the $c_{i,t,\tau}$ in Equation 1. For example, if occupancy starts on 2/13/2008, we assume rent payments take place on 2/13, 3/13, 4/13, and so on.

Table A.1: **Gross office leases by year of contract execution (JLL).** The table shows summary statistics by city for the leases sourced from JLL in the final sample used in the estimation.

	Mean	S.D.	p1	p25	p50	p75	p99
NYC							
Lease term (years)	8.54	4.17	1.00	5.08	10.00	10.50	20.00
Time to commencement (months)	0.40	4.25	0.00	0.00	0.00	0.00	14.00
Time to expiration (years)	8.57	4.22	1.00	5.08	10.00	10.50	20.00
Starting rent (USD)	4.95	2.05	2.08	3.50	4.58	5.83	12.50
Average rent (USD)	4.65	1.99	1.87	3.30	4.23	5.50	12.07
Average rent increase (USD per year)	0.03	0.06	0.00	0.00	0.03	0.05	0.12
Number of rent bumps	0.86	0.85	0.00	0.00	1.00	1.00	3.00
Average bump duration (months)	53.76	20.15	12.00	42.00	57.50	60.00	120.00
Tenant improvements (USD)	33.31	30.58	0.00	0.00	30.00	60.00	100.00
Free rent (months)	5.00	4.00	0.00	2.00	4.00	7.00	16.00
Boston							
Lease term (years)	6.14	2.65	1.72	5.00	5.25	7.25	13.75
Time to commencement (months)	4.37	6.80	0.00	2.00	3.00	5.00	33.39
Time to expiration (years)	6.50	2.77	2.00	5.08	5.67	7.58	15.20
Starting rent (USD)	2.80	1.03	1.08	2.00	2.67	3.46	5.67
Average rent (USD)	2.60	1.01	1.01	1.82	2.47	3.25	5.48
Average rent increase (USD per year)	0.06	0.09	0.00	0.04	0.06	0.07	0.20
Number of rent bumps	4.07	2.35	1.00	2.00	4.00	5.00	10.00
Average bump duration (months)	16.11	10.74	9.00	11.80	12.00	13.00	60.50
Tenant improvements (USD)	26.94	21.94	0.00	10.00	25.00	40.00	90.00
Free rent (months)	3.69	2.63	1.00	2.00	3.00	5.00	12.92

Table A.2: **Gross office leases by year of contract execution (CompStak)**. The table shows summary statistics by city for the leases sourced from CompStak in the final sample used in the estimation.

	Mean	S.D.	p1	p25	p50	p75	p99
NYC							
Lease term (years)	8.93	3.84	1.95	5.17	10.00	10.50	20.00
Time to commencement (months)	2.36	4.99	0.00	0.00	1.00	3.00	26.00
Time to expiration (years)	9.13	3.90	2.00	5.33	10.00	10.75	20.42
Starting rent (USD)	5.05	1.94	2.25	3.63	4.58	6.06	10.83
Average rent (USD)	4.76	1.90	2.13	3.36	4.33	5.66	10.63
Average rent increase (USD per year)	0.04	0.05	0.00	0.00	0.04	0.05	0.13
Number of rent bumps	0.92	0.81	0.00	0.00	1.00	1.00	3.00
Average bump duration (months)	55.07	19.05	12.00	46.33	57.00	60.00	120.00
Tenant improvements (USD)	29.31	28.12	0.00	0.00	25.00	50.00	100.00
Free rent (months)	4.96	3.94	0.00	2.00	4.00	7.00	15.00
Boston							
Lease term (years)	5.61	2.67	1.00	4.00	5.00	7.00	13.00
Time to commencement (months)	4.19	5.31	0.00	2.00	3.00	5.00	26.00
Time to expiration (years)	5.96	2.76	1.08	4.17	5.33	7.25	13.62
Starting rent (USD)	2.64	1.05	0.95	1.88	2.42	3.22	5.75
Average rent (USD)	2.47	1.00	0.92	1.72	2.30	3.00	5.51
Average rent increase (USD per year)	0.05	0.07	-0.01	0.02	0.05	0.07	0.23
Number of rent bumps	3.07	2.51	0.00	1.00	3.00	4.00	10.00
Average bump duration (months)	22.74	20.96	9.00	11.67	12.00	26.67	118.00
Tenant improvements (USD)	22.71	20.86	0.00	4.35	20.00	35.00	75.00
Free rent (months)	2.14	2.68	0.00	0.00	1.00	3.00	12.00

rent payment to zero in every month within the free month period.[40] In the case of tenant improvements, we assume that the full amount is given to the tenant on the first day of occupancy and thus subtract the TI amount from the first month's cash flow. [41] We calculate lease present values by discounting these cash flows using factors constructed with the continuously compounded Zero Coupon Bond from OptionMetrics. Finally, we follow the procedure described in Section 3 to calculate effective rents ($nPV_{t,i}$).

A.1.1 Combining Data Sources

We combine the records from both data sources to maximize the amount of information available for the estimation. To control for the possibility of duplicate records, i.e., leases that appear in both data sets, we identify the building and floors in which each rented space is located. For NYC, we use the lease address to obtain their Building Identification Number.[42] Since official building identification numbers are not available in the case of Boston, we use directly use each contract's address to generate internal building IDs. We then use our building identifiers to find duplicate leases across data sources. We consider a lease as duplicate if there is a contract pertaining to the same building and the same exact set of floors in both data sets.[43] Since the transaction date is more precisely recorded in the CompStak data,[44] we keep all CompStak records and drop duplicate leases from the JLL data sets before joining. Summary statistics for the complete sample are shown in Table 2.

A.2 Constrained Maximum Likelihood

For each city, we simultaneously back out the time series of key rates and estimate the unknown parameters in the state and observation equations. The unknown

[40] We make the assumption that the free rent period takes place without interruptions and starts on the first day of occupancy.

[41] We generally don't have the precise date or manner in which tenant improvement (TI) amounts are disbursed.

[42] The Building Identification Number (BIN) is assigned to every building in NYC by the City's Department of Planning Geosupport Services. See, for example, Chapter VI of the Geosupport System's User Programming Guide, available at: https://www.nyc.gov/assets/planning/download/pdf/data-maps/open-data/upg.pdf.

[43] Defining duplicate leases as those corresponding to the same building during the same quarter yielded similar summary stats and estimation results.

[44] See, for example, the Time to commencement summary statistics for NYC in Table A.1.

parameters in the state equation (equation 7) are the three elements of the $\bar{F}$ vector, the nine elements of the 3x3 ρ matrix, and the six distinct elements of the 3x3 symmetric matrix Q, the variance-covariance matrix of ϵ_{t+1}. From the observation equation, we estimate the variance parameters σ_ε^2 and σ_u^2.

Our estimation consists of maximizing the Kalman filter likelihood, subject to the following set of constraints.

A.2.1 Constraints on the eigenvalues

. To ensure that the autoregressive process of the key rates is non-oscillating and mean-reverting, we restrict the eigenvalues of ρ to the real interval $[0, 1]$. We similarly restrict the Q matrix to be positive semi-definite and focus on parameter regions in which the volatility of the key rates is not implausibly high from an economic point of view. We achieve this by constraining the eigenvalues of Q to lie within the range $[0, .13]$. To implement these restrictions on the eigenvalues, we start from a Schur decomposition of both matrices and impose box constraints on the diagonal elements of the upper-triangular matrix of the decomposition.[45]

A.2.2 Other constraints on ρ and Q

. We restrict the shape of the autoregressive matrix ρ to make the behavior of key rates compatible with the economics of the market for space. In particular, we impose restrictions that guarantee that 1) facing a fundamental shock, all key rates react in the same direction, and 2) the impact of the fundamental shock on each key rate monotonically decays over time. Computationally, we impose the following conditions:

1. The pairwise inner products of the rows of Σ are positive, where Σ is the *square root* matrix of Q, i.e., $\Sigma\Sigma' = Q$.

2. $-V(\Lambda^n \ln \Lambda)V^{-1}Q$ is non-negative, where V is the matrix whose columns are the eigenvectors of ρ, and Λ and is the diagonal matrix of eigenvalues of ρ

3. $V(\Lambda^n \ln \Lambda)V^{-1}Q(V')^{-1}(\Lambda^n \ln \Lambda)V'$ is non-negative.

[45] A Schur decomposition yields an upper-triangular matrix and a rotation matrix. The eigenvalues lie on the diagonal of the upper-triangular matrix.

A.3 Blended estimates

We consider two criteria for ranking each valid optimum: the Kalman filter likelihood and a moment-matching score. Our moment-matching score measures unconditional time series fit with the three unconditional moments of the effective rent: Mean, standard deviation, and first-order autocorrelation. For each valid optimum, we simulate 20,000 effective-rent (nPV) quarterly time series using the estimated parameters and the $\omega_{t,i}$ in the original sample.[46] We calculate the mean $n\hat{P}V$, standard deviation $\hat{\sigma}_{nPV}$, and first-order autocorrelation $\hat{\rho}_{nPV}$ of each simulated series and calculate the average and standard deviation of these moments across all simulations. We use the averages and standard deviations to obtain our moment-matching score, MMS:

$$MMS \equiv \left(\frac{Mean(n\hat{P}V) - n\bar{P}V}{\sigma(n\hat{P}V)}\right)^2 + \left(\frac{Mean(\hat{\sigma}_{nPV}) - \sigma_{nPV}}{\sigma(\hat{\sigma}_{nPV})}\right)^2 + \left(\frac{Mean(\hat{\rho}_{nPV}) - \rho_{nPV}}{\sigma(\hat{\rho}_{nPV})}\right)^2,$$

which is asymptotically distributed as χ^2_3.

We then calculate a *combined* score as the sum of the log likelihood and the log of the p-value of the moment-matching score:

$$\text{Combined score} = \ell + \log(\text{p-val}(MMS)).$$

We normalize the lowest *Combined score* to 0 and present most of our results as a weighted average of all valid local optima, with weights given by

$$\omega_i = \frac{\exp(\text{Combined score}_i)}{\sum_n \exp(\text{Combined score}_n)}$$

The resulting blended estimates for the state equation parameters are shown in table A.3. Figure A1 compares the estimated 10-year key rate using blended estimates (blue solid line) and maximum likelihood estimation (gray dotted line).

For calculations that involve variances, like the standard deviation of the slope and curvature of the term structure, we compute the weighted average of the second moment, and obtain the variance as the difference between the weighted second moment and the square of the weighted first moment:

[46] See below for a description of the simulation process.

Table A.3: **Blended parameter estimates**. The table shows the blended estimates for the parameters in the state equation (Equation 7). In each individual estimation, the eigenvalues of ρ were restricted to $[0, 1]$, while the eigenvalues of Q were constrained to be non-negative and below 0.13. Local optima were weighted and combined following the procedure described in Appendix A.3.

	NYC			Boston	
NYC.1	NYC.2	NYC.3	Boston.1	Boston.2	Boston.3
$\bar{F}$					
0.3108	0.2907	0.2037	-0.4521	-0.0782	0.0507
ρ					
0.4605	0.2585	0.1494	0.0684	1.315	-0.1898
-0.0884	0.8619	0.1557	-0.5477	1.6424	-0.0807
-0.1638	0.2659	0.821	-0.6086	1.1492	0.5533
Q					
0.0446	0.0274	0.0325	0.009	0.006	0.018
0.0274	0.036	0.0376	0.006	0.0087	0.0192
0.0325	0.0376	0.0489	0.018	0.0192	0.0875
Variance estimates					
σ_ε^2		σ_u^2	σ_ε^2		σ_u^2
0.0974		0.2112	0.0769		0.2055

Table A.4: **Eigenvalues of Q**. The table displays the combined estimates for the eigenvalues of the Q matrix.

	Eigenvalues of Q		
NYC	0.1102	0.0170	0.0024
Boston	0.0987	0.0053	0.0012

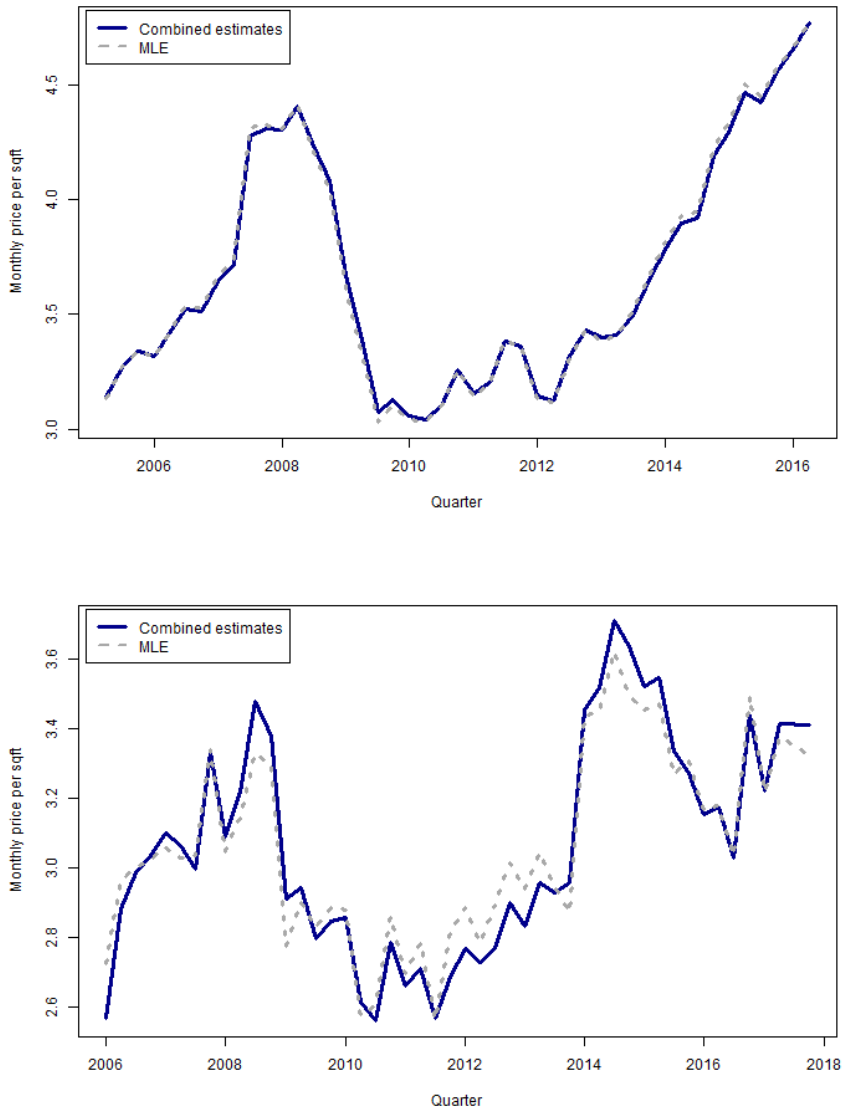

Figure A1: **10-year key rate estimated by two different methods**. The figure shows the smoothed estimate for the 10-year key rate obtained via blended estimates (blue solid line) and maximum likelihood estimator (gray dotted line) for NYC (top panel) and Boston (bottom panel).

$$Var_{combined}(X) = \sum_n \{\omega_n(Var_n(X) + E_n[X]^2)\} - \left(\sum_n \omega_n E_n[X]\right)^2.$$

A.3.1 Simulation Process

The simulation process consists of the following steps, repeated 10,000 times:

1. Use estimated error-variance parameters to draw from the error distributions of the state equation, ϵ, and the observation equation, $\tilde{u}$ and $\tilde{\varepsilon}$. The number of draws from the state-equation error term is given by the number of quarters in the original sample (NYC or Boston), whereas the number of draws from the observation-equation terms corresponds to the number of observations in the original sample.

2. Use estimated Kalman-filter parameters and simulated shocks to calculate the time series of key rates, $\hat{F}_t$.

3. Use simulated key rates, drawn observation-equation errors $\tilde{u}$ and $\tilde{\varepsilon}$ and the $w_{t,i}$ in the original sample to calculate the simulated $n\hat{P}V$:

$$n\hat{P}V_{t,i} = \frac{w'_{t,i}\hat{F}_t + \tilde{u}_{t,i}}{1 - \tilde{\varepsilon}_i}.$$

   Let $\max(nPV_{t,i})$ correspond to the observed highest effective rent in the sample, and let t^* and w^* be the corresponding origination date and key-rate weight vector. To avoid extreme outliers, we cap $\tilde{\varepsilon}$ at $\varepsilon_{max} = 1 - \frac{w^{*\prime} F_{t*}}{\max(nPV_{t,i})}$, which is the value of ε for the highest effective rent if one assumes that the observation noise for that lease is set to its average value (i.e., zero).

4. Compute the average, $n\hat{P}V_t$, by quarter to obtain the simulated-effective rent time series.

5. Calculate the mean, standard deviation, and first-order autocorrelation of the simulated effective-rent time series: $n\hat{P}V$, $\hat{\sigma}_{nPV}$, $\hat{\rho}_{nPV}$

6. Repeat steps 2-5 inverting the signs of the simulated error terms, i.e., $-\epsilon$, $-\tilde{u}$, and $-\tilde{\varepsilon}$.

A.4 Model Dynamics

we document helpful results on model dynamics. The key rate model itself takes the form:

$$F_{t+1} = \rho F_t + \bar{F} + \Sigma\tilde{\varepsilon}_{t+1}.$$

where F_t is the vector of key rates and ε_s is a three-vector of independent and normal Gaussian random variables. Setting $F_U \equiv (I-\rho)^{-1}\bar{F}$, one can rewrite the dynamics of the key rates as follows

$$F_{t+1} - F_U = \rho(F_t - F_U) + \Sigma\tilde{\varepsilon}_{t+1}, \tag{8}$$

which describes the demeaned dynamics of the key rates.

The distribution of future key rates is obtained by iterating equation (8):

$$F_{t+\tau} - F_U = \rho^\tau(F_t - F_U) + \sum_{j=1}^{\tau} \rho^{j-1}\Sigma\tilde{\varepsilon}_{t+\tau-j+1}. \tag{9}$$

From this, one can immediately deduce the Gaussian distribution of $F_{t+\tau}$:

$$F_{t+\tau} \sim F_U + N\left(\rho^\tau(F_t - F_U),\ \sum_{j=1}^{\tau} \rho^{j-1}\Sigma\Sigma'(\rho')^{j-1}\right). \tag{10}$$

A.5 Distribution of long-short portfolio profits

To simplify matters, we assume that each rental period corresponds to a quarter (i.e., rents are paid quarterly), but rents are quoted in monthly terms. We do this to match the estimated model dynamics to the lease contract. Consider a strategy where, beginning at date t, a short position is undertaken in a long-term lease with maturity of T years, together with a rolling long position in a one-quarter lease every quarter from date t until one quarter prior to T. The question at hand: What is the expected distribution of cash flow?

To address this, first consider the fixed (short) leg of the transaction (the long-term lease) and denote the model-implied monthly lease rate at date t for a T-year lease as $\ell_t(T)$. The forward equivalence relation implies that

$$3\sum_{m=1}^{4\times T} d_{j,t}v_j' \cdot F_t = 3\ell_t(T)\sum_{m=1}^{4\times T} d_{j,t},$$

where $d_{j,t}$ is the price of a j-quarter strip bond at date t and v_j is the vector of weights applied to the key-rate vector, F_t to arrive at a j-month forward price (e.g., $v_1' = (0.05, 0.95, 0)$ and $v_{30}' = (0.5, 0.5, 0)$). The factor of 3 corresponds to the fact that the lease rate is quoted as monthly but each lease period is a quarter. Thus,

$$\ell_t(T) = \frac{\sum_{m=1}^{4\times T} d_{j,t}v_j' \cdot F_t}{\sum_{m=1}^{4\times T} d_{j,t}}. \tag{11}$$

Each quarter starting at date t the strategy rolls over a new one-quarter lease. The profit at date $t+\tau$ is therefore $3(v_1' \cdot F_{t+\tau} - \ell_t(T))$. To assess the overall strategy profitability and its riskiness, one has to settle on a convention for accumulating them over time. This could be done by discounting the strategy cash flow to date t, capitalizing it forward to date $t+T$, or even using a simple sum. Each approach has its benefits and shortcomings. Consider for now the sum of strategy profits, discounted to date t:

$$P(t,T) = 3\sum_{\tau=1}^{4T} d_{\tau,t}\Big(v_1' \cdot F_{t+\tau-1} - \ell_t(T)\Big). \tag{12}$$

Correspondingly, from equation (9),

$$E_t[P(t,T)] = 3\sum_{\tau=1}^{4T} d_{\tau,t}\Big(v_1' \cdot \big(F_U + \rho^\tau(F_t - F_U)\big) - \ell_t(T)\Big). \tag{13}$$

To calculate the variance, first observe that

$$\begin{aligned}
F_{t+1} - E_t[F_{t+1}] &= \Sigma\tilde{\varepsilon}_{t+1} \\
F_{t+2} - E_t[F_{t+2}] &= \rho\Sigma\tilde{\varepsilon}_{t+1} + \Sigma\tilde{\varepsilon}_{t+2} \\
F_{t+3} - E_t[F_{t+3}] &= \rho^2\Sigma\tilde{\varepsilon}_{t+1} + \Sigma\rho\tilde{\varepsilon}_{t+2} + \Sigma\tilde{\varepsilon}_{t+3} \\
&\vdots \\
F_{t+S} - E_t[F_{t+S}] &= \rho^{S-1}\Sigma\tilde{\varepsilon}_{t+1} + \Sigma\rho^{S-2}\tilde{\varepsilon}_{t+2} + \rho^{S-3}\tilde{\varepsilon}_{t+3} + \ldots + \Sigma\tilde{\varepsilon}_{t+S} \\
F_{t+\tau} - E_t[F_{t+\tau}] &= \sum_{\tau'=1}^{\tau} \rho^{\tau-\tau'}\Sigma\varepsilon_{t+\tau'}
\end{aligned}$$

From this telescoping pattern, one can rewrite any sum over demeaned forwards as a sum over distinct shocks:

$$\sum_{\tau=1}^{S} w'_\tau \cdot \big(F_{t+\tau} - E_t[F_{t+\tau}]\big) = \sum_{\tau=1}^{S}\Big(\sum_{\tau'=\tau}^{S} w'_{\tau'} \cdot \rho^{\tau'-\tau}\Big)\Sigma\tilde{\varepsilon}_{t+\tau}.$$

Because the shocks are orthogonal, one can calculate the variance as:

$$\mathrm{VAR}_t\Big[\sum_{\tau=1}^{S} w'_\tau \cdot \big(F_{t+\tau} - E_t[F_{t+\tau}]\big)\Big] = \sum_{\tau=1}^{S}\Big(\sum_{\tau'=\tau}^{S} w'_{\tau'} \cdot \rho^{\tau'-\tau}\Big)\Sigma\Sigma'\Big(\sum_{\tau'=\tau}^{S} w'_{\tau'} \cdot \rho^{\tau'-\tau}\Big)'$$

Returning to the calculation of the profit variance using equation (9), identify w_τ with $3d_{\tau,t}v_1$ and set $S = 4T$ to yield

$$\mathrm{VAR}_t[P(t,T)] = 9\sum_{\tau=1}^{4T}\Big(\sum_{\tau'=\tau}^{4T} d_{\tau',t}v'_1 \cdot \rho^{\tau'-\tau}\Big)\Sigma\Sigma'\Big(\sum_{\tau'=\tau}^{4T} d_{\tau',t}v'_1 \cdot \rho^{\tau'-\tau}\Big)'. \qquad (14)$$

NOTE: To calculate these statistics for the forward capitalized profits we would replace $d_{\tau,t}$ with $\frac{d_{\tau,t}}{d_{4T,t}}$. If the aggregated strategy profits are measured as a simple sum, we'd replace $d_{\tau,t}$ with 1.