THE LINNEMAN LETTER

Volume 20, Issue 3 | Published since 2001 | Fall 2020

The following piece was co-authored with Matt Larriva, CFA of FCP®, a privately-held national real estate investment company.

If Interest Rates Determine Cap Rates, Where Is the Evidence?

The positive relationship between location (location, location) and value is the best-known relationship in real estate, but the relationship between interest rates and cap rates is a close second. The supposedly positive connection is reiterated by brokers (CBRE), assumed by industry groups (NAREIT), posited by education sites (Investopedia), and examined by academic journals (Briefings in Real Estate Finance). At first blush, the presumption of a positive relationship seems reasonable. After all, real estate has a bond-like component in its perpetual income stream and the discount rate should be closely related to the interest rate. When the discount rate falls, value should rise, establishing a tight relationship between interest rates and real rates. Or, if that argument does not convince you, consider that the weighted average cost of capital presumably decreases proportionally with rates. If the WACC is used as a discount rate, then it will move roughly with rates (holding all else equal) and should also create a tight connection between cap rates and interest rates. And should those two relationships fail to satisfy you, then there is always the argument that decreased borrowing costs will force funds into the market and create a demand-pull inflation on pricing. With all these pricing mechanisms at work, the positive relationship should be a foregone conclusion. But if real rates drive cap rates, then why is the empirical relationship so weak? And how could we have the same cap rates in grossly different real rate environments? For example, how did we have cap rates of 5.7% in January 2007, when the 10-year Treasury yield was at 4.7%, and cap rates 50 bps higher in July 2012 when the 10-year Treasury yield was 300 bps lower?

Background. Of course, all patterns have aberrations, and one contrary observation does not a proof make. But looking at the graph of the 10-year Treasury yield versus Green Street's major sector cap rates hardly shows a tight positive relationship (Figure 1). The same is true of real rates and all property cap rates (Figure 2).

At first blush, an R-squared of 0.68 might suggest that the relationship is adequate. But then again, that same level of correlation exists in the relationship between pool-drownings and Nicholas Cage films. And both relationships pale in comparison to the 0.95 R-square of cheese consumption and death-by-bed-sheet-entanglement (Figures 3 and 4).

If Statistics 101 teaches us anything, it is that simple *correlation does not mean causation*. And if Advanced Statistics taught us anything, it was that grad school is very expensive when you are paying your own way. We also remember something about the perils of randomly correlated time series. On that note, the scatterplot of interest rates against cap rates warrants investigation (Figure 5).

This plot appears to reflect the common wisdom that interest rates move with cap rates. But when we separate the interest rate and cap rate data by time period, a very different picture emerges. Specifically,

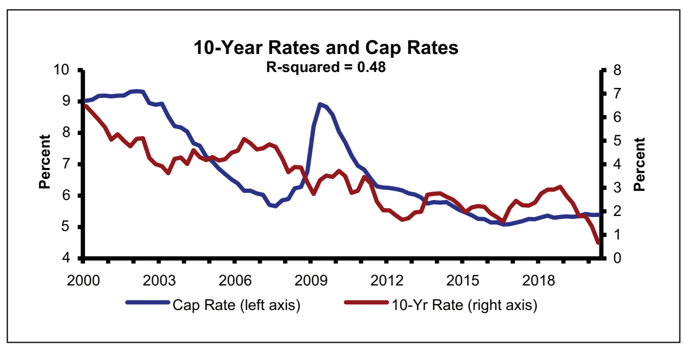

figure 1

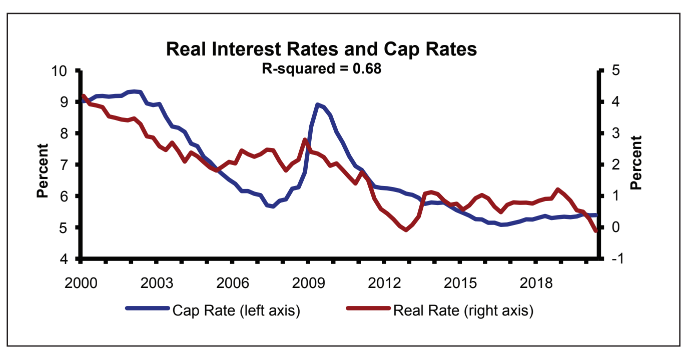

figure 2

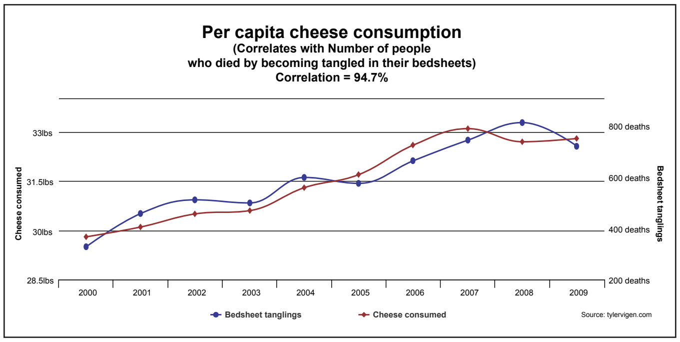

figure 3

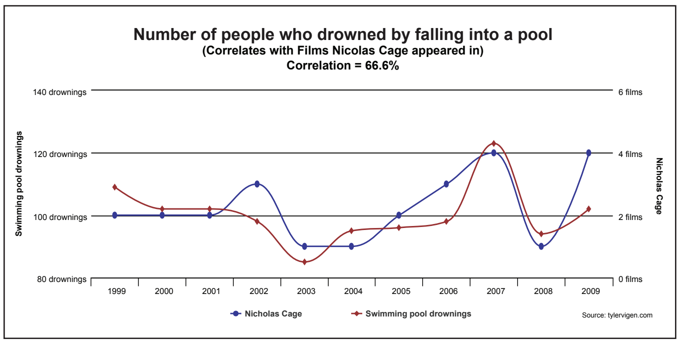

figure 4

as seen in Figure 6, three very separate patterns exist over time, each with negative (rather than the expected positive) relationships between 10-year yields and cap rates. The same phenomenon exists between cap rates and real interest rates (Figures 7 and 8).

This is known in statistics as a lurking variable problem, or Simpson's paradox: a false pattern that appears when distinct relationships are comingled. Specifically, the time period of the observation tells us much more about cap rates than does than the interest rate. That is, one can estimate the cap rate with a much higher accuracy by simply knowing what time period

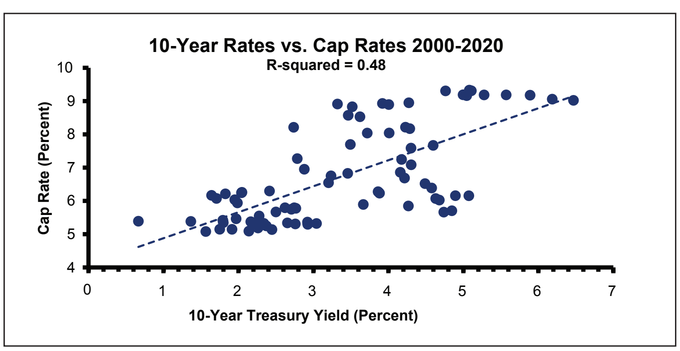

figure 5

one was trying to estimate. Knowing what the cap rate was in 2000, one could estimate with near-perfect accuracy what the cap rate would be in 2002 without knowledge of the 10-year Treasury yield. In fact, we find that both nominal and real rates, versus cap rates, show spurious patterns over time.

Statistically, one can often obviate the problems of spurious correlations by looking at the *changes* in the values, instead of the values. This exercise, with its 0.02 R-squared (Figure 9), reveals that cap rates are not driven by either real or nominal interest rates. Something else is determining cap rates.

A stronger relationship. If not interest rates, what determines cap rates? To be sure, the inability of real rates to explain cap rates has been explored by a number of researchers over the past decade. Some authors find that the relationship between interest rates is weak but noteworthy; others cite interest rates as important only in certain circumstances. In this search for a more comprehensive understanding of the factors that drive cap rates, Linneman (2015), and Chervachidze & Wheaton (2013) independently arrive at a similar metric: the flow of funds into commercial real estate.

Linneman's thought experiment is helpful: if you knew for certain that three times as much capital was competing for the same real estate a year from now, what would happen to real estate values? His answer is that they would be roughly three times higher, irrespective of interest rates. Linneman then finds that the flow of commercial mortgage funds has by far the greatest empirical impact on cap rates. His statistical analysis uses the 10-year Treasury yield, outstanding multifamily and commercial mortgages as a percent of GDP, and the unemployment rate as potential explanatory variables. With these, he established a model that tightly forecasted cap rates for a variety of property types and allowed him to make the prescient statement in 2015:

"…our bet is on the flood of liquidity, which could easily increase by more than 25%, keeping cap rates low even as other fundamentals exert upward pressure. So worrying about interest rates increasing appears not to be worth the effort. Instead, take advantage of the era of abnormally low rates by locking in debt financing for as long as possible, and watch the flow of mortgage funds as the key driver of changing cap rates."

In today's environment, where both real rates and nominal interest rates are uniquely low for as far as one

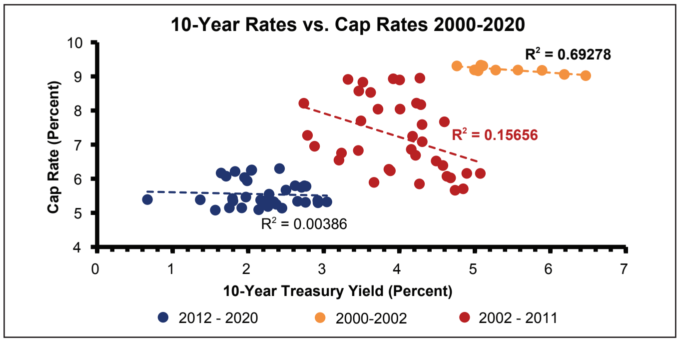

figure 6

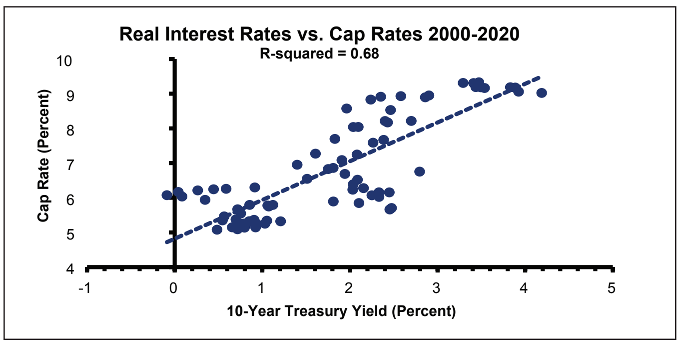

figure 7

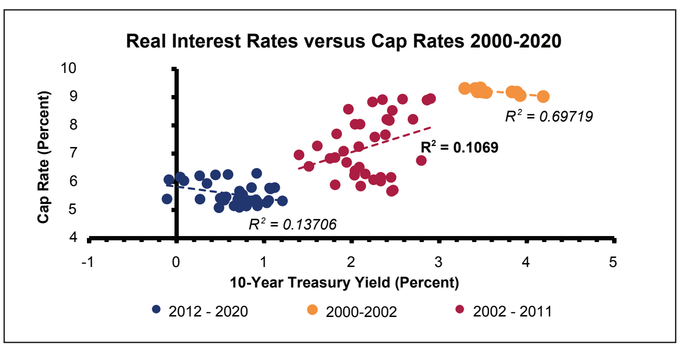

figure 8

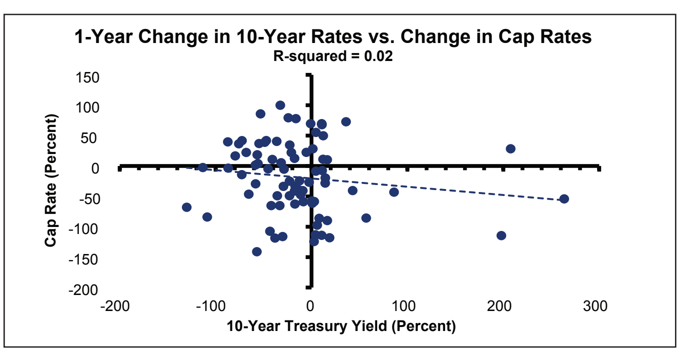

figure 9

can see, this result warrants revisiting. The aim here is to re-examine his fund flow model for robustness and ask if fund flows are simply another lurking variable. We utilize more cutting-edge statistical techniques than Linneman's original work, use better cap rate data (Green Street's transactional cap rate series for both apartments and offices) and examine how the model fares not just as a description of cap rates but also as a predictor of cap rates.

That last point is important. There is a distinction between a model which describes versus one that can predict. While we can *describe* the events of the past world conflicts—Antietam through D-Day—we struggle to forecast when and where such events will happen next. It makes the study of the events no less important, as we better understand market dynamics, but understanding World War II does not help one predict Operation Desert Storm. So too, it is possible that the fund flows explain cap rates but do not predict cap rates over the next year. We address both challenges.

Modeling Cap Rates as Functions of Different Variables. Linneman originally used the 10-year Treasury yield, the flow of mortgage funds relative to GDP, and unemployment rates to explain cap rates. We also use past values of the multifamily and office cap rates, fund flows (mortgage debt outstanding as a ratio of GDP), and U.S. unemployment rates.

This choice of variables is founded in economic theory, with cap rates determined by past cap rates, current supply and demand dynamics, and risk. We use one variable to capture each component. We regressed cap rates on earlier cap rates, the unemployment rate (to capture risk), and mortgage debt as a portion of GDP to capture the flow of funds.

To model these variables, we use a more sophisticated statistical model, which allows for multivariate time series analysis and addresses an array of knotty statistical issues.

For further reading on the specifics of our model, see our technical paper. The salient point is that statistics has a test of causality called *Granger Causality* which asks, "Am I better able to forecast cap rates if I know what funds flows are?" We find the answer is clearly yes, confirming Linneman's original experiment and empirical result.

Results: Does It Work? Does the model work descriptively? We examined the efficacy in forecasting one period ahead and conclude that the model tracks cap rates quite nicely and with very low errors. Note that this is the result of building a model on all the data available (2005-2020) and then using that model to forecast the series. While our study focused on the multifamily and office sectors, we believe (consistent with Linneman's original work) this analysis extends to other sectors.

The fit is excellent, though the model indicates a bigger jump in cap rates during the Financial Crisis than actually occurred. This is most likely because "extend and pretend" lender forbearance limited market discovery pricing. But also note that interest rates plunged during the Financial Crisis, while cap rates soared, contrary to the supposed positive correlation.

To explore the predictive power of our model, that is its ability to forecast cap rates for periods beyond the data, we estimated a new model at each period using only the data *prior* to that period. In this way, we test the model's predictive ability *out-of-sample*.

While the measures of fit decrease, they do so only marginally, and overall, the fit is good. This speaks to the efficacy of the model in capturing the dynamics of cap rates.

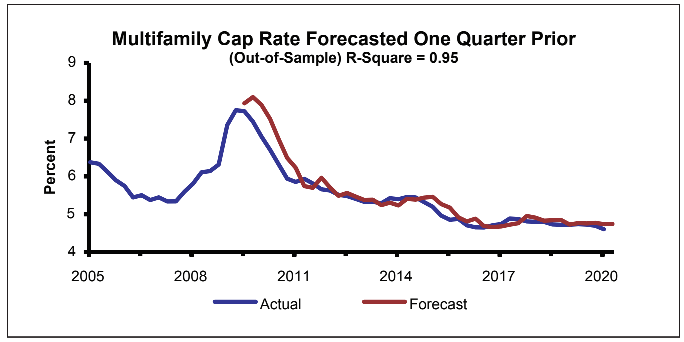

figure 12

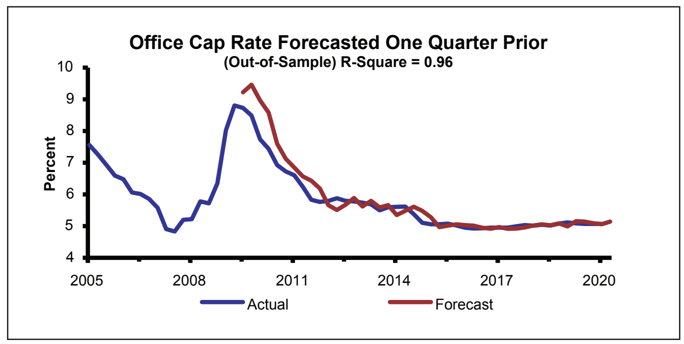

figure 13

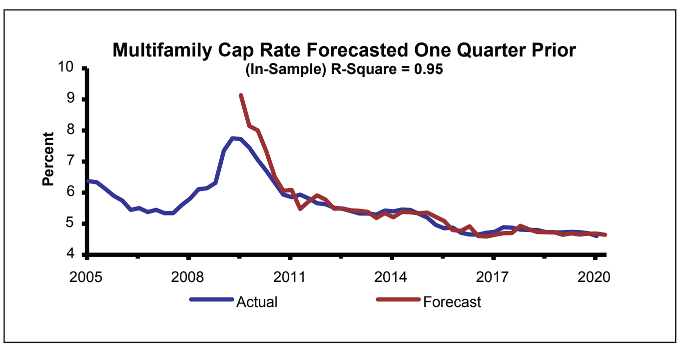

figure 10

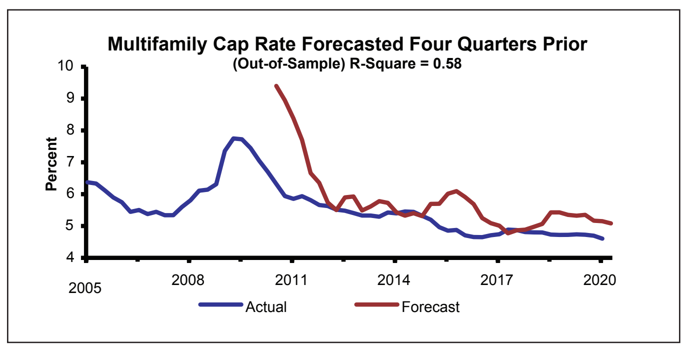

figure 14

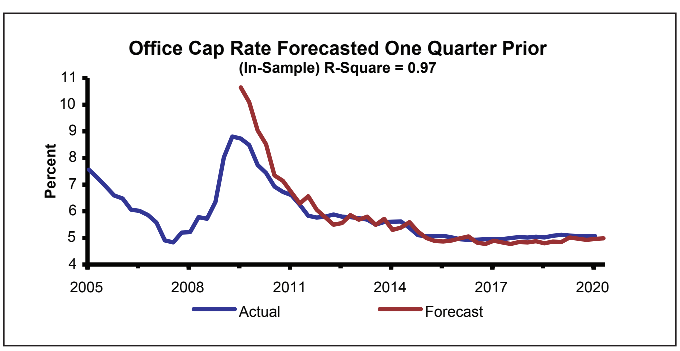

figure 11

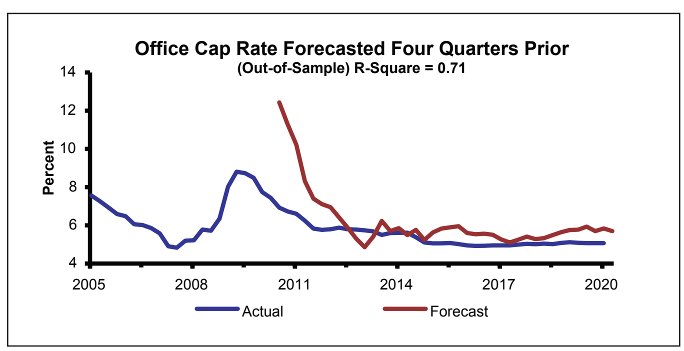

figure 15

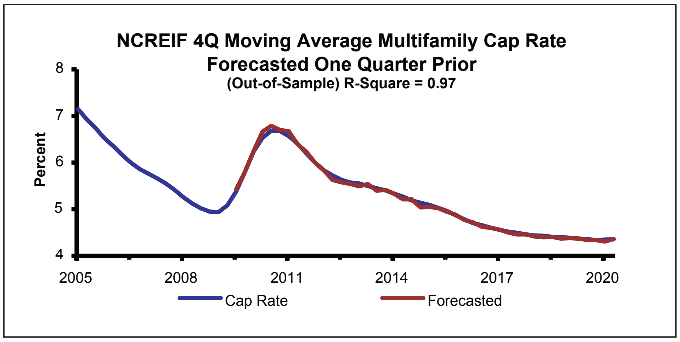

figure 16

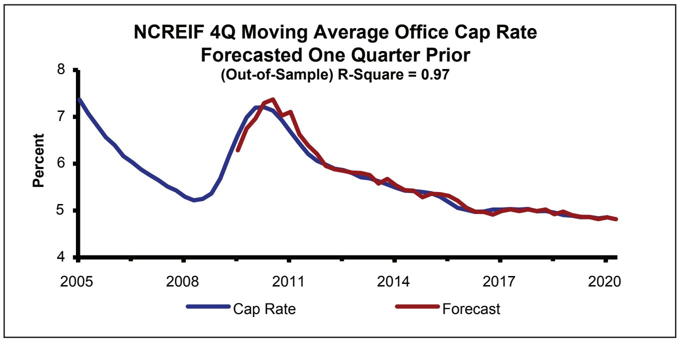

figure 17

As to the more challenging test of evaluating the model's ability to forecast cap rates one year in the future, we evaluate the four-quarter ahead cap rate forecast and find that the forecasts are highly volatile during the Financial Crisis (again due to forbearance limiting price discovery) but are less noisy and more accurate outside in other periods. It is reasonable to ask, what is the efficacy of a model that has an R-squared of 0.6-0.7? Interestingly, this R-squared is similar to that of the original real rates versus cap rates graph, on which many practitioners have relied. The difference is, our fund flow model (above) is as relatively accurate *predictively* as the real rates model is *descriptively*. Furthermore, the fund flow model is strong statistically, validated out-of-sample, and is based on variables proven to statistically cause cap rates.

How helpful is it to have a four-quarter ahead cap rate forecast? One-year cap rate forecasts can aid fund managers with optionality as to when to exit, looking for additional return in a space where the difference between top and middle quartile performance is mere percentage points. This result also helps the individual buyers and sellers looking to purchase or sell. Such sellers have optionality as to time and would be served greatly by knowing which direction pricing would head in the next year.

It is also worth noting that the accuracy achieved above is the highest of any published research to date and is based on transactional data. Most previous works have focused on NAREIT appraisal data and have not extended past the one-quarter prediction framework.

Evaluating our model on this benchmark, we see very high R-squared values.

What Does It Mean? Empirically, the in-sample and out-of-sample forecasts are quite robust, but examining why is tantamount. So how big are the impacts we find? The model's coefficients produce the following sensitivities:

Cap Rate Response to a 100-bp increase in:	Mortgage Debt as a % of GDP	Unemployment
Multifamily Cap Rates	-22	1
Office Cap Rates	-65	3

figure 18

We find that a change in the unemployment rate from 5% to 4% lowers cap rates by a negligible one to three basis points. Thus, even the 600-bp increase in the unemployment rated during the Financial Crisis only raised cap rates by 6-18 bps, and the inverse as unemployment fell. This is not really economically significant though it is statistically precise.

More importantly, we find that when mortgage debt grows 100 bps faster (slower) than GDP, cap rates fall (rise) by 22 and 65 basis points for multifamily and office properties respectively. If debt grows by 10%, relative to GDP, cap rates stand to compress by 220-650 bps. This is a dramatic impact.

So we clearly find that an increase in mortgage debt as a percent of GDP drives down cap rates, and an increase in unemployment slightly drives up cap rates. And this stands to reason, as these two variables provide insight to the risk side and the demand side of pricing, through unemployment and mortgage debt, respectively.

In sum, we confirm Linneman's earlier finding that the connection between both multifamily and office cap rates and interest rates is weak, while the connection with flow of funds is the powerful driving force. Given that, we encourage investors to look to the flow of mortgages relative to GDP (specifically its change) as an indication of where cap rates should go in the near term and perhaps the longer term.

Our model finds that a spike in unemployment is very weakly negative for real estate valuation in the short term, but in the longer-term, the view on rates has not changed, as the flow of funds itself has been stable the past five years, with all real estate mortgage debt at 75% of GDP.

As monetary infusions spike, rates dive, and equity valuations move upwards, there is value to having a model which suggests a single variable of focus. As of 1Q 2020 there is an increase in the amount of mortgage debt as a percent of GDP. Granted, this is in large part due to the compression of GDP rather than the expansion of mortgage debt, but the model as stated accounts for this. And while unemployment is certainly wide of normal, the net impacts of these factors suggests stable-to-decreasing cap rates for the near-term.

Over the next year, we expect multifamily cap rates to drop 10 basis points while the office sector drops 20 basis points. There are, of course, ways that this dynamic can be muted. Two that come to mind are surprise inflation and cloudy price discovery via forbearance. The former may cause an exodus from real estate into higher-yielding asset types, while the latter may unhinge pricing from supply and demand dynamics. In all cases, we stand with George Box, who said, "All models are approximations. Assumptions, whether implied or clearly stated, are never exactly true. All models are wrong, but some models are useful. So the question you need to ask is not 'Is the model true?' (it never is) but 'Is the model good enough for this particular application?'" As interest rates approach zero, we submit the funds flow metric as a *useful* model in the current application.

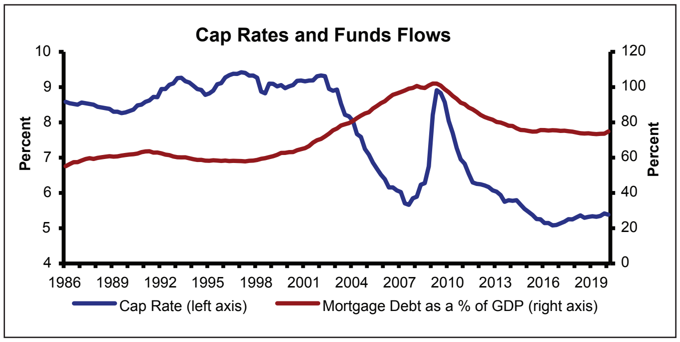

figure 19

About Dr. Peter Linneman

Dr. Linneman, who holds both Masters and Doctorate degrees in economics from the University of Chicago, is the Principal of Linneman Associates. For nearly four decades, he has provided strategic and financial advice to leading corporations. Through Linneman Associates, he provides strategic and M&A analysis, market studies, and feasibility analysis to a number of leading U.S. and international companies. In addition, he serves as an advisor to and a board member of several public and private firms.

Dr. Linneman is the author of the leading real estate finance textbook, *Real Estate Finance and Investments: Risks and Opportunities*, now in its fifth edition. His teaching and research focuses on real estate and investment strategies, mergers and acquisitions, and international markets. He has published over 100 articles during his career. He is widely recognized as one of the leading strategic thinkers in the real estate industry.

He also served as the Albert Sussman Professor of Real Estate, Finance, and Business and Public Policy at the Wharton School of Business at the University of Pennsylvania until his retirement in 2011. A member of Wharton's faculty since 1979, he served as the founding chairman of Wharton's Real Estate Department and the Director of Wharton's Zell-Lurie Real Estate Center for 13 years. He is the founding co-editor of *The Wharton Real Estate Review*.

About Matt Larriva

Matt Larriva is Vice President for Research and Data Analytics at FCP® where he leverages data science and quantitative methods to enhance asset selection, management, and disposition across real estate investments. His active research focuses on multivariate timeseries analysis as it applies to cap rate forecasting (Vector Error Correction Methods as models of US Real Estate Markets (working)), and random forest methods as they apply to submarket rent-growth in unstructured data (Semi-Greedy Construction of Oblique-Split Decision Trees (2019)).

Prior to joining FCP, Matt worked at Green Street Advisors as the head of the US and UK Quant teams, managing the data efforts in the sell-side real estate research firm. While at Green Street, he focused on REIT trading strategies, data productization, and analytic enhancements. Matt completed his undergraduate degree in economics at the Wharton School of the University of Pennsylvania. He holds a Masters in applied statistics from UCLA and is a CFA charter holder.

FCP is a privately held real estate investment company that has invested in or financed more than $6 billion in assets since its founding in 1999. For further information, visit www.fcpdc.com.