Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Construction-at-Risk: Financial conditions and downside risk in residential construction

Myeongwon Lee
Sungkyunkwan University
Doojin Ryu
sharpjin@skku.edu

Sungkyunkwan University

Research Article

Keywords: Construction-at-risk, Downside risk, Financial conditions, Housing supply, Quantile forecasting, Residential construction

Posted Date: September 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9870498/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Construction-at-Risk: Financial conditions and downside risk in residential construction

Contribution of this study

This study contributes to housing and built-environment research by introducing a construction-at-risk framework for monitoring downside risk in residential construction quantities. Existing housing-risk studies mainly focus on house prices, housing returns, mortgage credit, or aggregate residential investment. These outcomes are important, but they do not fully capture stress in the production side of the housing market. House prices are equilibrium outcomes shaped by both demand and supply, and they may remain stable even when the construction pipeline weakens or when approved projects fail to move into actual construction. This study shifts the focus from price-side housing risk to quantity-side construction risk and asks whether adverse movements in residential construction can be forecast before they become visible in aggregate housing investment or house-price dynamics. A central contribution is the distinction between building permits and housing starts as stage-specific measures of the residential supply process. Building permits capture the planning-stage pipeline, where projects enter the approved supply process. Housing starts capture the execution stage, where planned projects become actual construction activity. This distinction is important for built-environment research because housing supply is produced through a staged and frictional process rather than through immediate adjustment. A decline in permits may indicate deterioration in the future construction pipeline, whereas a decline in starts may signal execution-stage stress arising from financing constraints, builder balance sheets, labor availability, material procurement, or project-level frictions. By modeling PERMIT and HOUST separately, this study provides a framework for distinguishing pipeline-stage weakness from execution-stage stress.

This study also contributes to the macro-financial housing literature by showing that broad financial conditions contain useful early-warning information about lower-tail construction outcomes. Using the Chicago Fed National Financial Conditions Index, the Adjusted National Financial Conditions Index, and the official Risk, Credit, and Leverage subindexes, the study evaluates whether financial-condition information improves recursive out-of-sample quantile forecasts relative to a historical lower-tail benchmark. The results show that broad financial conditions reduce lower-tail forecast losses for both building permits and housing starts. The gains are concentrated at the 20th percentile and become stronger during high-stress periods. This pattern supports the construction-at-risk interpretation: financial conditions are especially informative about adverse construction outcomes, not merely about the central tendency of housing activity.

A further contribution is the comparison with standard mortgage and interest-rate benchmarks. Mortgage rates, Treasury yields, term spreads, and mortgage spreads are natural predictors of housing activity because they capture borrowing costs and the interest-rate environment. However, the evidence shows that these rate-based benchmarks do not reproduce the lower-tail forecasting gains generated by broader financial-condition measures. This result indicates that construction downside risk is not captured by mortgage-market tightness or rate-

channel information alone. Instead, broader macro-financial conditions, especially the joint information contained in Risk, Credit, and Leverage, provide incremental predictive content for monitoring stress in residential construction.

Methodologically, this study adapts the lower-tail forecasting logic of growth-at-risk research to the housing supply process. Rather than forecasting average construction growth, the study constructs a recursive construction-risk target by removing own-history predictability from PERMIT and HOUST growth and evaluating the lower tail of the resulting cumulative construction innovations. The recursive out-of-sample design ensures that each forecast uses only information available at the forecast origin. The study also uses recursively estimated Risk-Credit-Leverage principal components to summarize common and compositional information in the official Chicago Fed financial-condition subindexes without look-ahead bias. This design provides a transparent and implementable monitoring framework rather than a purely in-sample explanatory exercise.

Our findings have direct policy and practical relevance. For policymakers, housing planners, lenders, and construction-market monitors, the framework provides an early-warning tool for tracking downside risk in the residential supply process. PERMIT-based lower-tail forecasts can identify deterioration in the future supply pipeline, while HOUST-based lower-tail forecasts can identify stress in the execution of planned construction. This distinction is valuable because housing-market weakness may first appear in construction quantities before it is fully reflected in prices, aggregate investment, or employment. We offer a practical monitoring approach for identifying stage-specific vulnerabilities in the residential built environment.

Our main contribution is to show that downside risk in residential construction can be measured, forecast, and monitored as a distinct housing-supply risk. By combining stage-specific construction quantities with broad financial-condition information and recursive lower-tail forecasting, the construction-at-risk framework extends housing-risk analysis beyond prices and average activity. It provides evidence that macro-financial conditions help identify the lower-tail deterioration of residential construction and that this predictive information is most useful when housing supply is viewed as a staged production process.

Highlights

·Construction-at-risk tracks lower-tail risk in residential construction quantities.
·Building permits and housing starts separate pipeline and execution-stage risk.
·Broad financial conditions improve lower-tail forecasts beyond rate benchmarks.
·Forecast gains are concentrated at the 20^{th} percentile and in high-stress periods.

Abstract

We introduce a construction-at-risk framework for monitoring downside risk in U.S. residential construction. Building permits and housing starts are used as quantity-side measures of the housing supply process: PERMIT captures the planning-stage pipeline, while HOUST captures the execution stage in which planned projects become actual construction activity. The analysis examines whether broad financial conditions provide early-warning information about adverse movements in these construction outcomes. Chicago Fed financial-condition measures reduce

lower-tail forecast losses for both PERMIT and HOUST. The gains are concentrated at the 20th percentile, are not reproduced by mortgage-rate benchmarks, and are larger during periods of elevated financial stress. The results suggest that lower-tail construction forecasts provide a useful monitoring tool for housing supply risk and for distinguishing pipeline-stage weakness from execution-stage stress in the staged production of the residential built environment.

Keywords: Construction-at-risk; Downside risk; Financial conditions; Housing supply; Quantile forecasting; Residential construction
JEL Classification: C53, E32, E44, G10, R31

1. Introduction

Housing is a central link between the macroeconomy and financial markets. It is the largest asset for many households, a key source of collateral in mortgage lending, and an important component of bank balance sheets. When the housing cycle weakens, the effects can extend beyond house prices to consumption, residential investment, construction employment, regional lending, and financial-sector risk. Residential investment and housing activity have long been viewed as informative components of business-cycle conditions (Davis and Heathcote, 2005; Leamer, 2015).

Most housing-risk research focuses on house prices. This focus is natural because house prices are closely tied to household wealth, collateral values, and mortgage default risk. House prices, however, are equilibrium outcomes. They reflect both demand and supply conditions. A rising price does not necessarily imply low housing risk, and a stable price does not rule out construction-side stress. When land availability, regulation, labor, material costs, or financing conditions restrict supply responses, prices and construction quantities can adjust differently. This means that house-price movements alone do not fully capture housing-market adjustment, motivating a separate focus on downside risk in residential construction quantities. This construction-side perspective is especially relevant for housing and built-environment research because residential supply is produced through a staged process rather than through immediate price adjustment. Aastveit, Albuquerque, and Anundsen (2023) show that U.S. housing supply elasticities have declined and that building permits have become less responsive as supply constraints have become more important. Murray (2020) further shows that planned or approved supply need not translate immediately into actual production when developers face timing incentives and project-level frictions.

We study two measures of residential construction: building permits and housing starts. We use PERMIT to capture the planning-stage pipeline of residential construction and HOUST to capture the execution stage in which planned or approved projects become actual construction activity. Because residential construction is a dynamic investment process with adjustment costs and timing frictions (Topel and Rosen, 1988), permits and starts may share the same housing cycle without reflecting the same decision margin. A fall in permits may point to a weakening future pipeline, whereas a fall in starts may reflect execution-stage disruptions related to financing, labor, or material procurement, or lenders' willingness to bear construction risk. Studies have developed forecasting models and predictive frameworks for housing prices, housing returns, and housing investment (Cañizares Martínez, de Bondt, and Gieseck, 2023;

Chae, Cho, Kim, and Ryu, 2025). Prior work reinforces the importance of housing-market predictability, but it mainly focuses on prices, returns, or aggregate investment. The missing link is the lower-tail forecasting of construction quantities themselves. Existing studies on housing, growth-at-risk, and financial conditions address different aspects of this question but do not provide an out-of-sample (OOS) framework for monitoring the lower tail of residential construction quantities. We address this gap by developing a construction-at-risk framework for PERMIT and HOUST.

Our main question is whether broad financial conditions improve lower-tail forecasts of residential construction activity relative to a historical benchmark. We use the Chicago Fed National Financial Conditions Index (NFCI), the Adjusted National Financial Conditions Index (ANFCI), and the Risk, Credit, and Leverage subindexes. NFCI and ANFCI measure broad financial conditions. Risk, Credit, and Leverage are official subindexes within the Chicago Fed NFCI framework. The aim is to examine which types of financial condition information are useful for forecasting downside risk in the construction sector. We compare broad financial-condition measures with standard mortgage and interest-rate benchmarks. This comparison asks whether the lower-tail predictive content of financial conditions merely reflects familiar rate-channel information or whether broader financial conditions contain additional information for the downside risk of construction.

We construct the forecasting object by transforming PERMIT and HOUST into monthly growth rates, removing the component predicted by each series' own past growth, and cumulating the resulting AR residuals over the forecast horizon. We then estimate recursive OOS quantile forecasts and evaluate whether financial-condition information reduces quantile loss relative to a recursive historical lower-tail benchmark. The recursive design ensures that each forecast uses only information available at the forecast origin. We report a one-month-lag predictor specification as a conservative check on timing. We construct Risk-Credit-Leverage (RCL) principal components as a compact, recursively estimated summary of common and compositional financial-stress information. The results show that broad financial conditions reduce lower-tail forecast losses for both PERMIT and HOUST, with gains concentrated at the 20^{th} percentile and stronger performance during periods of elevated financial stress. Overall, the construction-at-risk framework shifts the forecasting object from house prices and average housing activity to the lower tail of residential construction quantities and links macro-financial conditions to stage-specific downside risk in the residential supply process. We make four contributions. First, we study construction-side downside risk using PERMIT and HOUST as stage-specific measures of quantity. Second, we show that the predictive content of financial conditions is concentrated in the lower tail rather than in the center or upper part of the construction distribution. This pattern supports a construction-at-risk interpretation: financial conditions are especially informative about adverse construction outcomes, not merely average construction activity. Third, we compare broad financial-condition measures with standard mortgage and interest-rate benchmarks to examine whether NFCI-type information merely reflects familiar rate channels. Fourth, we use the official Risk, Credit, and Leverage subindexes of the Chicago Fed NFCI as a parsimonious forecasting device for summarizing common and compositional financial-stress information. These contributions define a lower-tail forecasting approach to residential construction quantities. This forecasting perspective is useful because PERMIT and HOUST are closer to specific decision margins in the housing

supply process than house prices are. A lower-tail forecast for PERMIT can signal deterioration in the future construction pipeline, while a lower-tail forecast for HOUST can signal weakness in the execution of planned activity. The framework provides a monitoring perspective for tracking downside risk in residential construction before it is fully reflected in aggregate investment or house-price movements.

The rest of this study is organized as follows. Section 2 reviews the related literature and research background. Section 3 describes the data and methodology. Section 4 reports the empirical results. Section 5 concludes.

2. Related Literature and Research Background

This study builds on three strands of research. The first examines housing activity and the business cycle. Residential investment and housing starts are among the most cyclical components of real activity and have been used as informative indicators of macroeconomic conditions (Davis and Heathcote, 2005; Leamer, 2015). Building permits also contain forward-looking information about real activity. Strauss (2013) shows that building permits forecast state-level job and income growth and often outperform standard indicators such as housing prices, interest rates, and leading indicator indexes. Forecasting work emphasizes the difficulty of modeling housing investment. Cañizares Martínez, de Bondt, and Gieseck (2023) document substantial model uncertainty in housing investment forecasting and cross-country differences in the empirical determinants of housing investment. Their findings motivate the use of construction activity as a housing-supply forecasting object. Our study differs in that it focuses on lower-tail risk in construction quantities rather than on point forecasts of aggregate housing investment or average housing activity.

The second strand links house prices, housing supply, and quantity adjustment. House prices are important because they affect household wealth, collateral constraints, mortgage default risk, and local credit conditions. Housing-market forecasting studies continue to focus on house prices, housing price indices, and housing returns. Chae, Cho, Kim, and Ryu (2025) construct a Gangnam ripple index and examine its predictive power for Korean housing market returns. These studies are closely related in their focus on housing-market predictability, but we shift the forecasting object from prices and returns to construction quantities. Price movements alone do not fully describe housing-market adjustment. Evidence from regulated housing markets shows that conventional house-price modeling can be difficult (Tu, de Haan, and Boelhouwer, 2017). Supply elasticity, land constraints, regulation, and permitting frictions shape both price and quantity responses (Glaeser, Gyourko, and Saiz, 2008; Saiz, 2010; Paciorek, 2013). Housing investment involves adjustment costs and timing frictions, which make construction activity a dynamic process rather than an immediate supply response (Topel and Rosen, 1988). More broadly, built-environment development can involve institutional and stakeholder-related frictions, as shown in transaction-cost analyses of urban renewal decision-making (Zhuang, Qian, Visscher, and Elsinga, 2020). Expectations also matter in housing-market dynamics. Hong and Ryu (2023) develop a model of house price dynamics with expectations, and Park, Hong, and Ryu (2024) examine heterogeneous expectations in the housing market using an agent-based framework. Strauss (2013) links building permits more directly to consumer expectations and future economic activity, supporting the use of PERMIT as a forward-looking measure for the construction pipeline. Permits are closer to the approval and

pipeline margin, while starts are closer to the execution margin. Housing-supply research supports this stage-specific interpretation. Aastveit, Albuquerque, and Anundsen (2023) show that building permits can respond weakly when supply constraints become more important, while Murray (2020) shows that planned or zoned supply does not automatically become new housing production. These findings support the use of PERMIT and HOUST as separate outcomes for monitoring downside risk in the residential supply process.

The third strand connects financial conditions to real activity and downside risk. Financial frictions can amplify real fluctuations through borrowers' balance sheets and external finance premia (Bernanke, Gertler, and Gilchrist, 1999), and financial intermediation frictions can affect real activity through intermediaries' balance sheets and credit supply (Gertler and Kiyotaki, 2010). Credit spreads and excess bond premia contain information about future economic activity (Gilchrist and Zakrajšek, 2012). Kim, Cho, and Ryu (2025) link global risk aversion to future real economic activity, further supporting the view that financial variables can contain forward-looking information about subsequent real outcomes. The same macro-financial logic motivates our test of whether broad financial conditions contain information about downside risk on the construction side.

A related macro-financial literature focuses on lower-tail forecasting. Growth-at-risk research estimates the conditional distribution of future economic activity and shows that financial conditions are closely related to downside growth risk (Adrian, Boyarchenko, and Giannone, 2019). Adams, Adrian, Boyarchenko, and Giannone (2021) construct risks around consensus forecasts of GDP growth, unemployment, and inflation and show that tight financial conditions forecast downside growth risk, upside unemployment risk, and greater inflation uncertainty. Macro-finance work measures time-varying macroeconomic tail risk using a long-run international panel and links these risks to asset prices and future returns (Marfè and Pénasse, 2024). Systemic risk research shows that financial risk measures can contain OOS predictive information for the lower tail of future macroeconomic outcomes (Giglio, Kelly, and Pruitt, 2016). International evidence also indicates that combinations of systemic risk measures can improve macroeconomic forecasts and that the strength of these predictive relationships can vary across quantiles of systemic risk indicators (Pan, Zhou, and Dong, 2025). We apply this lower-tail forecasting perspective to residential construction. The forecasting object is not aggregate GDP growth, industrial production, financial returns, or macroeconomic disaster risk, but the lower tail of construction quantities measured by building permits and housing starts.

The gap lies at the intersection of housing research, growth-at-risk analysis, and financial-conditions forecasting. Housing research examines prices, investment, starts, permits, expectations, and supply constraints. Growth-at-risk analysis examines the lower tail of aggregate macroeconomic outcomes, while financial-conditions forecasting studies broad real-activity and downside-risk predictability. Yet no existing framework provides an out-of-sample design for monitoring the lower tail of residential construction quantities, even though construction is the stage at which housing-market stress becomes real activity. We address this gap by developing a construction-at-risk framework for PERMIT and HOUST.

3. Data and Methodology

We use monthly U.S. residential construction and financial-condition data. The construction series are PERMIT and HOUST. PERMIT is used as the planning-stage construction measure,

and HOUST is used as the execution-stage construction measure. The main measures of financial conditions are NFCI, ANFCI, and the Risk, Credit, and Leverage subindexes within the Chicago Fed NFCI framework. These measures allow financial stress to be measured more broadly than by a single mortgage rate or Treasury rate variable. Financial conditions indices are designed to combine information from multiple segments of the financial system and to track the link between financial markets and real activity (Brave and Butters, 2011). NFCI covers money markets, debt and equity markets, the banking system, and the shadow banking system. ANFCI adjusts financial conditions for movements associated with current economic activity and inflation. Risk, Credit, and Leverage provide the official decomposition of the NFCI framework. Risk reflects volatility and funding risk conditions; Credit reflects credit conditions; and Leverage reflects debt and equity measures. We use these subindexes as the main macro-financial predictor set within the official NFCI taxonomy. The baseline predictor set is intentionally limited to financial-condition variables and rate-based benchmarks to keep the forecasting exercise focused on the incremental predictive content of macro-financial conditions.

We include standard mortgage and interest-rate measures: the 30-year fixed mortgage rate, the 10-year Treasury yield, the term spread, defined as the 10-year Treasury yield minus the 3-month Treasury bill rate, and the mortgage spread. We use these variables as benchmark predictors. They do not directly measure construction supply; instead, they proxy for the demand and rate environment in which construction decisions are made. Mortgage rates capture homebuyer borrowing costs, while long-term Treasury yields and term spreads summarize the broader interest-rate environment. Term spreads have also been widely used as predictors of future business-cycle conditions (Estrella and Hardouvelis, 1991; Estrella and Mishkin, 1998). The mortgage spread, defined as the 30-year fixed mortgage rate minus the 10-year Treasury yield, provides a simple measure of mortgage-market tightness. These benchmarks allow us to ask whether broad financial-condition measures add information beyond standard mortgage and interest-rate conditions. All predictors are aligned at the monthly frequency. In the baseline forecasting convention, the predictor vector dated month t is used to forecast the cumulative construction outcome over months $t+1$ through $t+h$. This convention treats financial-condition variables as information available at the forecast-origin month. Because exact within-month release timing can vary across data sources, we report a conservative robustness check in which all financial predictors are lagged by one month. Under this lagged specification, only information dated $t-1$ or earlier is used to forecast the outcome over $t+1$ through $t+h$. Table 1 reports summary statistics and sample coverage for the processed monthly variables. Table 1 shows that HOUST growth is more volatile than PERMIT growth in the processed monthly sample. This pattern is consistent with the view that actual starts can fluctuate more sharply than the pipeline of building permits. This descriptive difference does not imply that HOUST is more sensitive to financial conditions. The empirical question is whether financial-condition measures reduce lower-tail forecast losses for each construction series relative to the historical benchmark. Figure 1 plots PERMIT and HOUST over the sample period. The two series move with the broad housing cycle, but their timing and amplitude are not identical. This visual difference supports the empirical design that treats PERMIT as a planning-stage construction measure and HOUST as an execution-stage construction measure.

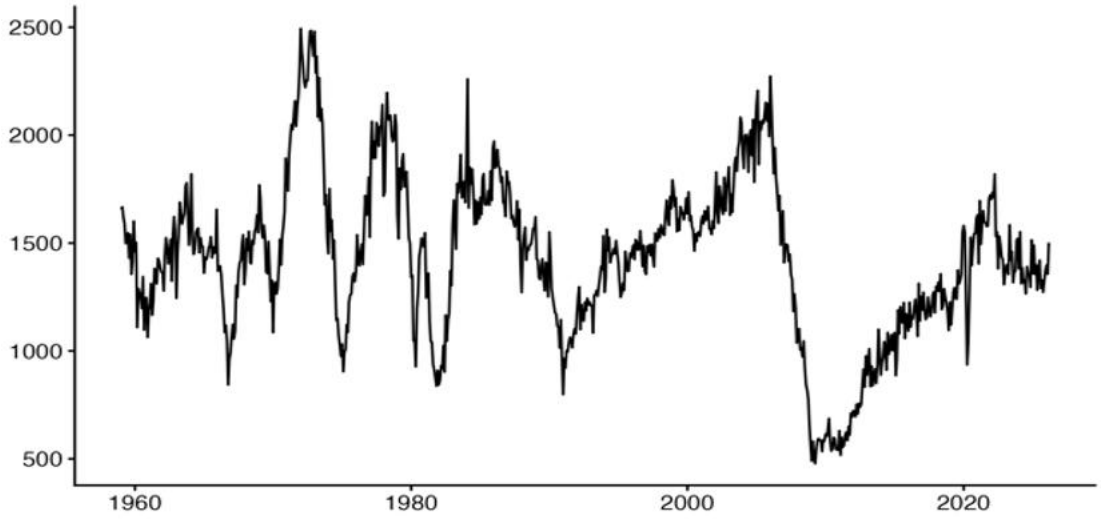
Table 1. Summary statistics and sample coverage

	N	Start	End	Mean	SD	Min	P25	Med	P75	Max
PERMIT growth	794	1960-02	2026-03	0.029	5.957	-27.41	-3.03	0.11	3.54	29.21
HOUST growth	806	1959-02	2026-03	-0.012	8.217	-30.67	-4.72	0.21	4.90	25.67
NFCI	664	1971-01	2026-04	0.000	0.999	-1.09	-0.59	-0.39	0.16	5.15
ANFCI	664	1971-01	2026-04	0.001	0.998	-1.32	-0.58	-0.38	0.21	5.15
Risk	664	1971-01	2026-04	0.000	0.999	-1.15	-0.62	-0.39	0.21	5.25
Credit	664	1971-01	2026-04	0.000	1.001	-3.14	-0.25	-0.06	0.19	3.36
Leverage	664	1971-01	2026-04	0.000	1.000	-1.98	-0.57	-0.14	0.29	3.90
Mortgage rate	661	1971-04	2026-04	7.688	3.189	2.68	5.45	7.21	9.22	18.45
10Y Treasury	772	1962-01	2026-04	5.815	2.936	0.62	3.89	5.44	7.53	15.32
Term spread	533	1982-01	2026-05	1.538	1.244	-1.74	0.55	1.58	2.55	4.15
Mortgage spread	661	1971-04	2026-04	1.760	0.517	0.53	1.43	1.67	1.98	4.85

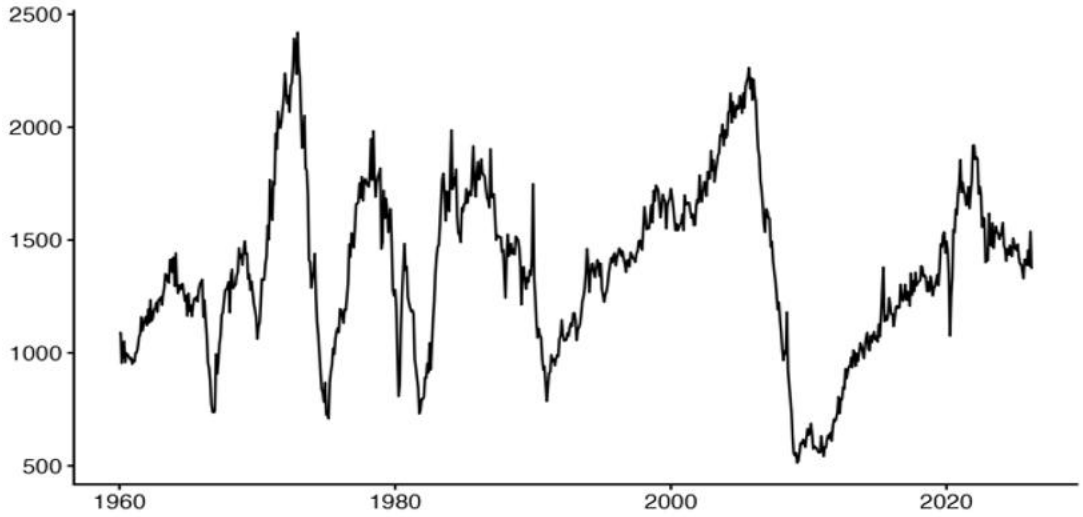
Notes. This table shows summary statistics for the monthly variables used in the empirical analysis. *N*, *Start*, *End*, *Mean*, *SD*, *Min*, *P25*, *Med*, *P75*, and *Max* denote the number of non-missing observations, first available month, last available month, mean, standard deviation, minimum, 25th percentile, median, 75th percentile, and maximum, respectively. *PERMIT growth* and *HOUST growth* are measured in percent, financial-condition variables are reported in standardized index units, and mortgage-rate variables are measured in percentage points.

Figure 1. Residential construction activity

Panel A. HOUST



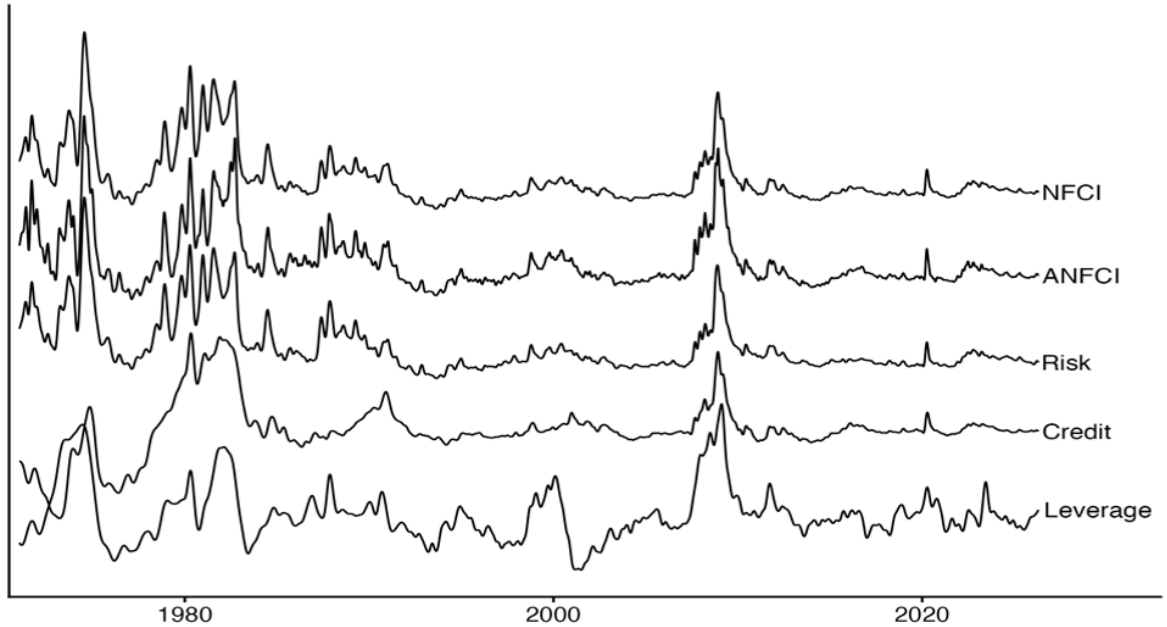
Panel B. PERMIT



Notes. This figure plots monthly residential construction activity. Panel A reports new privately owned housing units started (HOUST). Panel B reports new private housing units authorized by building permits (PERMIT). The *x*-axis reports calendar time.

Figure 1 indicates that PERMIT and HOUST should not be treated as a single construction indicator. PERMIT is closer to the entry of projects into the approved pipeline, while HOUST is closer to actual construction activity. This distinction is useful for examining whether financial conditions contain different information for pipeline-stage and execution-stage downside risk. Figure 2 plots the Chicago Fed financial-condition measures used in the forecasting analysis. NFCI and ANFCI measure broad financial conditions, while Risk, Credit, and Leverage provide the official subindex decomposition within the NFCI framework.

Figure 2. Chicago Fed financial-condition variables



Notes. This figure plots monthly Chicago Fed financial-condition variables. *NFCI*, *ANFCI*, *Risk*, *Credit*, and *Leverage* are standardized and vertically shifted for visual comparison. The *x*-axis reports calendar time. The right-hand labels indicate the corresponding series.

Figure 2 shows that the Chicago Fed financial-condition measures often move together during stress periods, while the subindexes remain distinct. This pattern motivates two forecasting exercises. The first evaluates each subindex separately. The second uses RCL principal components to capture the common movement and relative composition of Risk, Credit, and Leverage. This follows the broader idea that overlapping financial risk information can be represented using dimension-reduction methods in lower-tail forecasting applications, while our implementation is limited to the three official NFCI subindexes (Giglio, Kelly, and Pruitt, 2016). We next construct the forecasting object. For each residential construction series Y_t, where Y_t is either $PERMIT_t$ or $HOUST_t$, we estimate the following autoregressive (AR) model for monthly log growth:

$$g_t^Y = \alpha^Y + \sum_{j=1}^{p_Y} \phi_j^Y g_{t-j}^Y + \varepsilon_t^Y, \quad \text{where } g_t^Y = 100(\log Y_t - \log Y_{t-1}). \tag{1}$$

Here, g_t^Y denotes the monthly log growth rate of the corresponding residential construction series, a^Y is the intercept, p_Y is the lag order, ϕ_j^Y is the AR coefficient at lag *j*, and ε_t^Y is the AR-residual. The residual ε_t^Y is not interpreted as a structural housing or financial shock. It is a reduced-form AR residual obtained by removing the component of construction growth explained by the target's own past values. The AR lag order is selected recursively within the available training sample using the Akaike information criterion, with a maximum lag length of six months. Thus, neither the lag order nor the AR coefficients are selected using future observations. The minimum initial training window is 120 monthly observations. This recursive target construction is used to reduce the role of own-history predictability in PERMIT and HOUST before evaluating the incremental forecasting content of financial conditions. This target construction follows the lower-tail forecasting perspective of Giglio, Kelly, and Pruitt (2016), who evaluate whether financial-risk information predicts future macroeconomic downside outcomes; we adapt this idea to residential construction activity. Let *h* denote the forecast horizon, measured in months. The *h*-month construction-risk target is defined as the cumulative AR residual:

$$y_{t+h}^{Y,h} = \sum_{j=1}^{h} \varepsilon_{t+j}^Y, \tag{2}$$

where $y_{t+h}^{Y,h}$ denotes the cumulative AR residual from *t*+1 through *t*+*h*. Appendix B reports the corresponding results using alternative construction-outcome definitions based on raw monthly growth and macro-adjusted outcomes. The conditional τ-quantile of the cumulative target at forecast origin *t* is defined as:

$$q_{t+h|t,\tau}^{Y,h} = Q_\tau\left(y_{t+h}^{Y,h} \middle| \mathcal{F}_t\right), \tag{3}$$

where $\mathcal{F}_t$ is the information set available at forecast origin *t*, and τ is the target quantile. The baseline uses h=3 and τ=0.20. The 20th percentile is chosen to capture moderate downside construction risk while retaining enough lower-tail observations for recursive OOS evaluation. More extreme quantiles would be closer to severe crash risk, but would leave fewer effective tail observations and increase estimation uncertainty in the recursive OOS setting. Appendix

A reports tail-event counts for alternative lower-tail quantiles and coverage diagnostics for the baseline 20th-percentile forecasts. Appendix A shows that the model-implied baseline 20th-percentile forecasts produce empirical hit rates reasonably close to the nominal level, whereas more extreme 10th-percentile targets generate fewer effective tail observations and less stable coverage. The empirical design uses recursive OOS quantile forecasting. At each forecast origin *t*, the estimation sample includes only past observations whose *h*-month target has already been realized:

$$T_t = \{s: s + h \leq t\}, \tag{4}$$

where T_t is the training sample available at forecast origin *t*. Operationally, an observation dated *s* is included in the training sample at forecast origin *t* only if $s + h \leq t$. This restriction prevents the estimator from using target realizations that would not yet be observed when the forecast is made. For example, when *h*=3, the target attached to month s uses residuals from *s*+1 through *s*+3, and it can be included in the training sample only from month *s*+3 onward. We estimate the conditional quantile of the cumulative construction-risk target recursively at each forecast origin. The recursive quantile regression yields model-specific coefficient estimates $\hat{\alpha}_{\tau,t}^{Y,m}$ and $\hat{\beta}_{\tau,t}^{Y,m}$. The loss function is described as follows:

$$\rho_\tau(u) = u(\tau - 1\{u < 0\}), \tag{5}$$

where $1\{\cdot\}$ is an indicator function. The model-implied conditional quantile forecast is:

$$\hat{q}_{(t+h|t,\tau)}^{Y,h,m} = \hat{\alpha}_{\tau,t}^{Y,m} + (X_t^m)'\hat{\beta}_{\tau,t}^{Y,m}, \tag{6}$$

where the left-hand side denotes the conditional τ-quantile forecast generated by model *m*. The coefficient estimates are obtained from the recursive quantile regression using the training sample T_t. The benchmark forecast is a recursive unconditional lower-tail forecast based on the historical empirical distribution of the target. At each forecast origin *t*, the benchmark is computed as the empirical τ-quantile of realized cumulative targets available in the training sample T_t. Because the target variable is already constructed from recursively estimated AR residuals, the historical benchmark should be interpreted as an unconditional lower-tail benchmark for construction innovations after removing own-history predictability. Financial condition models are evaluated for their incremental ability to forecast the lower tail of construction outcomes not explained by the target's own lagged growth. Appendix B further reports robustness checks using raw growth outcomes and macro-adjusted construction outcomes, showing that the main findings are not an artifact of the AR-residual target definition. OOS quantile loss of model *m* is defined as:

$$L_t^{Y,m} = \rho_\tau\left(y_{t+h}^{Y,h} - \hat{q}_{t+h|t,\tau}^{Y,h,m}\right), \tag{7}$$

where $L_t^{Y,m}$ is the OOS quantile loss of model *m* for target *Y*. The loss for the historical benchmark, $L_t^{Y,hist}$ is defined analogously using the recursive historical forecast. The main

performance measure is the percentage reduction in OOS quantile loss relative to the historical benchmark:

$$qR^2_{Y,m}(\%) = 100\left[1 - \frac{\sum_{t\in O} L_t^{Y,m}}{\sum_{t\in O} L_t^{Y,hist}}\right], \tag{8}$$

where *O* denotes the OOS evaluation period. To test whether model *m* improves upon the historical benchmark, the loss differential is defined as:

$$d_t^{Y,m} = L_t^{Y,hist} - L_t^{Y,m}, \tag{9}$$

where $d_t^{Y,m} > 0$ indicates that model *m* produces a lower quantile loss than the historical benchmark at forecast origin *t*. The one-sided forecast-improvement test is:

$$H_0: E\left[d_t^{Y,m}\right] \leq 0, \tag{10}$$

where the null hypothesis states that model *m* does not reduce OOS quantile loss relative to the historical benchmark. This loss-differential comparison follows Diebold and Mariano (1995) and West (1996), with the loss function defined by the quantile loss. Let $\overline{d^{Y,m}}$ be the OOS sample average loss differential. The test statistic is:

$$t_{NW}^{Y,m} = \overline{d^{Y,m}} / SE_{NW}\left(\overline{d^{Y,m}}\right), \tag{11}$$

where $SE_{NW}\left(\overline{d^{Y,m}}\right)$ is the Newey-West heteroskedasticity- and autocorrelation-consistent (HAC) standard error of the average loss differential. This adjustment is needed here because the target in Equation (3) is an *h*-month cumulative AR residual. When *h*>1, adjacent forecast targets share some monthly residual components. For example, with *h*=3, the target evaluated at *t+3* contains residuals from *t+1*, *t+2*, and t+3, whereas the next target evaluated at *t+4* contains residuals from *t+2*, *t+3*, and *t+4*. Because the two targets share the residual components at *t+2* and *t+3*, the resulting loss differential can be serially correlated. The HAC truncation lag is set equal to the forecast horizon *h*, with a minimum lag of one month. Thus, the baseline *h*=3 specification uses a HAC lag of three, while the horizon-robustness exercises use horizon-specific HAC lags. This choice is intended to account for the mechanical overlap in the *h*-month cumulative targets, and the block-bootstrap *p*-values reported below provide an additional finite-sample check. The bootstrap *p*-values are computed using a moving-block bootstrap applied to the loss-differential series. We use 499 bootstrap replications, a block length of six months, and a fixed random seed for reproducibility. The bootstrap is implemented under the null by recentering the loss differential around its sample mean before resampling blocks. The reported bootstrap *p*-value is one-sided and corresponds to the fraction of bootstrap average loss differentials that are at least as large as the observed average loss differential. The baseline predictors include the Chicago Fed NFCI, ANFCI, and the Risk, Credit, and Leverage subindexes.

To summarize common and compositional information parsimoniously, we construct recursive RCL principal components. The RCL-PC specification is not a high-dimensional

factor model; it is a low-dimensional macro-financial forecasting specification based on the official NFCI taxonomy. We use only the three official NFCI subindexes: Risk, Credit, and Leverage. The purpose of PCA is not to reduce a large set of predictors, but to derive low-dimensional forecasting variables from overlapping macro-financial information within the Risk, Credit, and Leverage structure. The recursive loading diagnostics support this use of the components. PC1 loads positively on Risk, Credit, and Leverage throughout the evaluation period, so we use PC1 as a common RCL component. PC2 is orthogonal to PC1 and mainly contrasts Credit with Risk and Leverage. We use PC2 as a relative-composition component within the RCL subindexes. This component is economically useful because the downside risk of construction may depend not only on the overall level of financial stress but also on its composition. A configuration in which Credit moves differently from Risk and Leverage may indicate whether construction weakness is more associated with credit availability, broader risk sentiment, or balance-sheet pressure. We interpret PC2 as a compositional forecasting variable rather than a structural channel. On average, PC1 explains about 65% of the variation in the three subindexes, while PC1 and PC2 together explain about 89%. Appendix C reports the full-sample loadings, explained-variance shares, and recursive loading summaries used to support this construction. The PCA is implemented recursively to avoid look-ahead bias. At each forecast origin, the mean, standard deviation, covariance matrix, and loading vectors are computed using only the training sample available at that time. Full-sample standardization and full-sample PCA are not used, so the RCL component scores do not incorporate information from future observations.

$$x_t = (Risk_t, Credit_t, Leverage_t)', \tag{12}$$

where x_t collects the Risk, Credit, and Leverage subindexes. At each forecast origin *t*, the vector is standardized using only the training sample:

$$z_{s,t} = D_t^{-1}(x_s - \hat{\mu}_t), \quad s \in T_t, \tag{13}$$

where $\hat{\mu}_t$ is the training-sample mean vector and D_t is the diagonal matrix of training-sample standard deviations. The second subscript in $z_{s,t}$ indicates that observation *s* is standardized using the training sample available at forecast origin *t*. Let C_t be the covariance matrix of the standardized vector in the training sample. The principal-component loading vectors solve:

$$C_t v_{k,t} = \lambda_{k,t} v_{k,t}, \tag{14}$$

where $v_{k,t}$ is the *k*-th loading vector and $\lambda_{k,t}$ is the corresponding eigenvalue. The RCL principal-component score at forecast origin *t* is:

$$F_{k,t} = v_{k,t}' z_{t,t}, \tag{15}$$

where $F_{k,t}$ is the *k*-th RCL principal-component score. The main RCL specification uses the first two RCL principal-component scores, $F_{1,t}$ and $F_{2,t}$:

$$Q_{\tau}\left(y_{t+h}^{Y,h} \mid F_{1,t}, F_{2,t}\right) = \alpha_{\tau}^{Y} + \beta_{1,\tau}^{Y} F_{1,t} + \beta_{2,\tau}^{Y} F_{2,t}. \quad (16)$$

This specification uses the first two RCL components as forecasting variables in the recursive quantile regression. The first component captures the common movement across the Risk, Credit, and Leverage subindexes, while the second component captures the leading remaining variation within the same subindex structure. The components serve as low-dimensional forecasting variables that summarize common and compositional information within the RCL subindexes.

4. Empirical Results

This section shows the OOS forecasting results. We begin with the baseline specification, which evaluates whether financial-condition measures reduce lower-tail forecast losses for PERMIT and HOUST relative to the historical lower-tail benchmark. We then examine whether the gains are concentrated in the lower tail, whether they are reproduced by standard mortgage and interest-rate benchmarks, and whether the results are sensitive to the forecast horizon, timing convention, or major crisis episodes. We report high-stress-window diagnostics to assess whether forecast gains are stronger when financial stress is elevated. Table 2 reports the baseline results for h=3 and τ=0.20. For the baseline models reported in Table 2, the OOS evaluation period covers forecast origins from January 1990 to December 2025, yielding 432 monthly forecast origins. The table compares individual NFCI-related measures, the mortgage-spread benchmark, and the RCL-PC specifications. Positive values indicate that the model improves on the benchmark.

Table 2. Baseline OOS quantile forecasting results

	Model	qR² (%)	NW-p	Boot-p	N
PERMIT	NFCI	6.47**	0.016	0.024	432
PERMIT	ANFCI	5.93**	0.016	0.024	432
PERMIT	Risk	5.82**	0.030	0.050	432
PERMIT	Credit	-0.24	0.594	0.449	432
PERMIT	Leverage	6.33*	0.062	0.080	432
PERMIT	Mortgage spread	-0.82	0.796	0.745	432
PERMIT	RCL PC1	7.70**	0.029	0.046	432
PERMIT	RCL PC1+PC2	8.34**	0.022	0.036	432
HOUST	NFCI	5.92**	0.023	0.046	432
HOUST	ANFCI	6.46**	0.015	0.036	432
HOUST	Risk	5.61**	0.032	0.052	432
HOUST	Credit	-0.03	0.515	0.411	432
HOUST	Leverage	7.33**	0.026	0.044	432
HOUST	Mortgage spread	-1.25	0.999	0.998	432

HOUST	RCL PC1	7.85**	0.014	0.030	432
HOUST	RCL PC1+PC2	8.65**	0.011	0.030	432

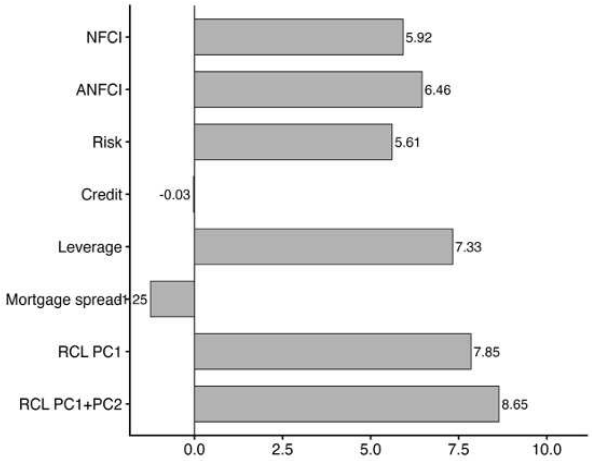
Notes. This table presents baseline OOS quantile forecasting results. The baseline uses h=3 and τ=0.20. *Model* denotes the target and forecasting specification. qR^2 *(%)* is measured relative to the historical lower-tail benchmark. *NW-p* and *Boot-p* denote Newey-West HAC and moving-block bootstrap *p*-values, respectively. *N* denotes the number of OOS forecast origins. * and ** indicate significance at the 10% and 5% levels.

Table 2 shows that NFCI, ANFCI, Risk, Leverage, and the RCL-PC specifications reduce lower-tail forecast losses for both PERMIT and HOUST. Risk and Leverage deliver more consistent improvements than Credit alone does, while the mortgage-spread benchmark yields limited gains. This does not imply that credit conditions or mortgage-market information are unimportant. Credit spreads are known to contain information about future economic activity (Gilchrist and Zakrajšek, 2012). In this setting, the Credit subindex alone does not capture the full set of financial-condition information relevant for construction-side downside risk; credit-related information appears more useful when interpreted jointly with risk and leverage conditions. The strongest baseline performance comes from the RCL PC1+PC2 specification, with loss reductions of 8.3% for PERMIT and 8.7% for HOUST, consistent with the loading diagnostics discussed above. PC1 captures common movement across Risk, Credit, and Leverage, while PC2 primarily reflects the relative configuration of Credit versus the Risk-Leverage side of the RCL structure. The contribution of PC2 suggests that the composition of financial stress contains forecasting information beyond the common stress level summarized by PC1. Lower-tail construction forecasts benefit from a parsimonious summary of both common RCL variation and the leading residual contrast across the RCL subindexes. This interpretation is statistical and predictive: PC1 and PC2 summarize variation in the official subindexes, but they are not used to identify separate structural channels. Because the baseline exercise compares multiple forecasting specifications, Appendix D reports Benjamini-Hochberg adjusted *p*-values for the baseline Newey-West tests. The adjustment weakens some individual-model significance levels, but the RCL specifications remain among the strongest performers for both PERMIT and HOUST. Economically, these results imply that broad financial conditions contain information about downside construction risk beyond the target's own history and beyond a recursive historical benchmark. The finding is relevant for macro-financial monitoring because a deterioration in construction activity can be transmitted to residential investment, construction employment, local credit conditions, and regional real activity. The RCL PC1+PC2 and Leverage gains are slightly stronger for HOUST than for PERMIT. This is consistent with the execution-stage interpretation of housing starts: financial conditions may be especially informative when planned projects enter actual construction, a stage at which financing availability, builder balance sheets, and project-level frictions are more likely to matter. At the same time, the positive PERMIT results indicate that financial conditions also contain information about the future construction pipeline. The evidence points to stage-specific predictive content within the residential construction process, rather than a single undifferentiated construction indicator. Figures 3 and 4 provide visual summaries of the baseline evidence. Figure 3 reports the baseline loss reductions across models, while Figure 4

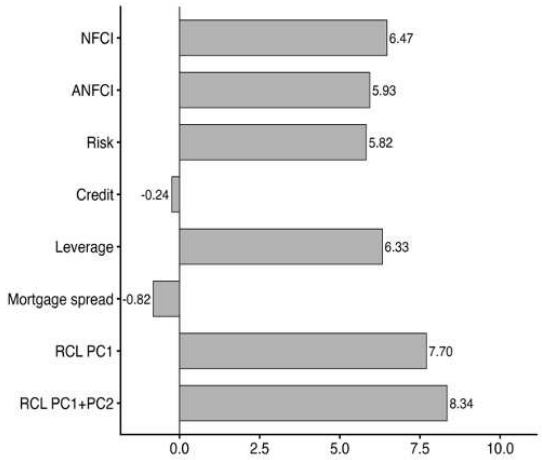
plots the realized cumulative AR-residual target, the conditional 20th-percentile forecast, and the historical 20th-percentile benchmark for the main RCL PC1+PC2 specification.

Figure 3. Baseline loss reduction

Panel A. HOUST



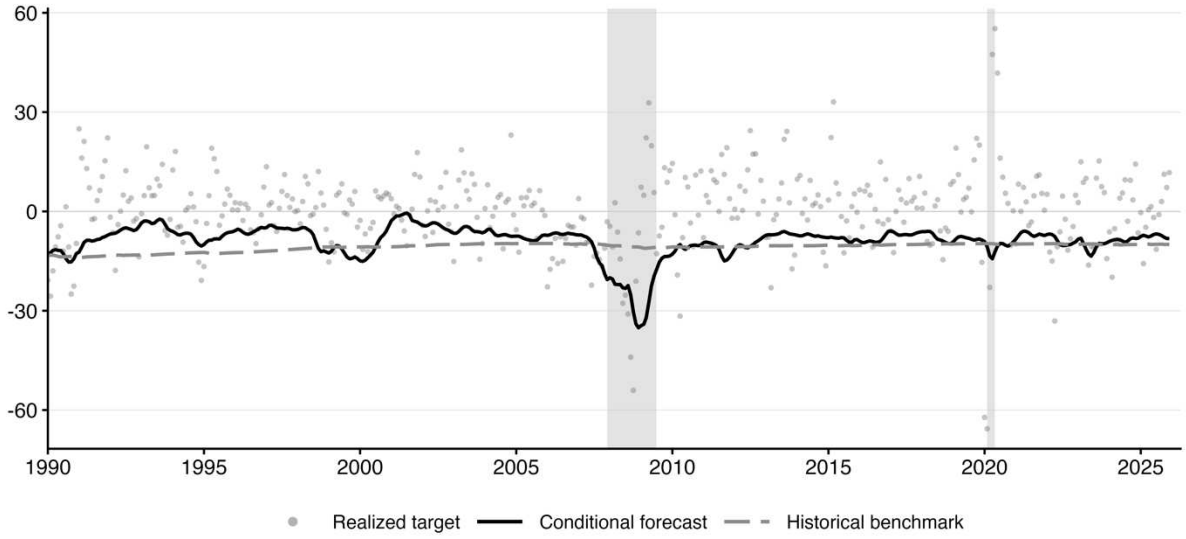
Panel B. PERMIT



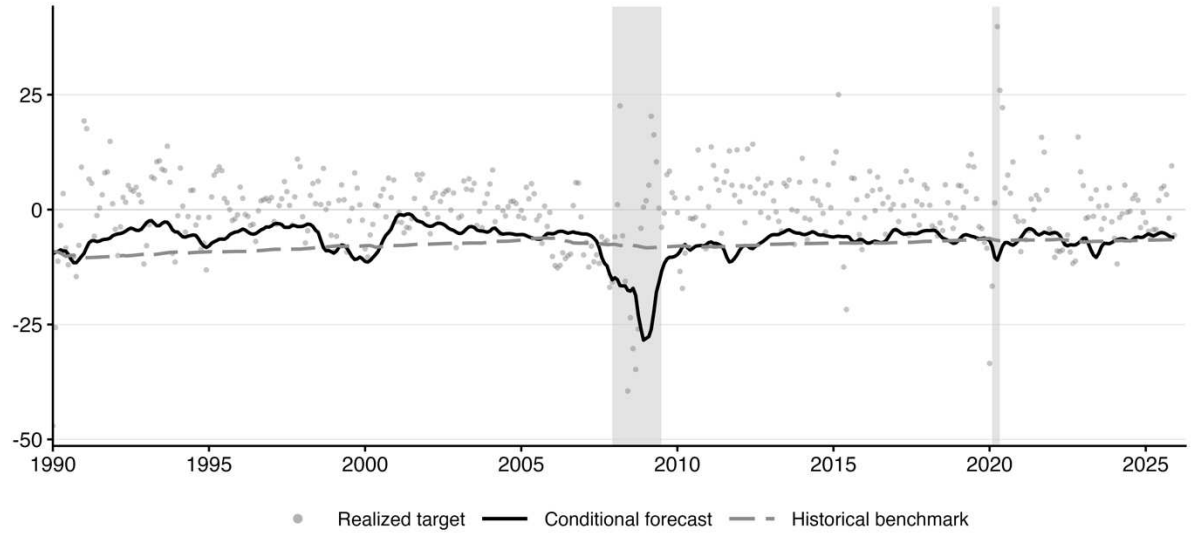
Notes. This figure presents baseline loss reductions by target and model. Panel A reports *HOUST*, and Panel B reports *PERMIT*. The baseline uses h=3 and τ=0.20. The x-axis reports loss reductions in percent, and the y-axis reports forecasting models.

Figure 4. Lower-tail forecast paths

Panel A. HOUST



Panel B. PERMIT



Notes. This figure plots realized cumulative AR-residual targets and lower-tail forecasts for the RCL PC1+PC2 specification. Panel A reports *HOUST*, and Panel B reports *PERMIT*. The *x*-axis reports calendar time. The *y*-axis reports the realized target and the corresponding 20th-percentile forecasts. The dots denote realized cumulative AR-residual targets. The solid line denotes the conditional 20th-percentile forecast, and the dashed line denotes the historical 20th-percentile benchmark. Shaded areas indicate the GFC and COVID windows.

Figure 3 provides the visual counterpart to Table 2. Figure 4 shows that the RCL PC1+PC2 forecast varies over time rather than tracking the historical benchmark by a constant shift. The historical benchmark moves gradually because it relies solely on the recursive empirical distribution of past realized outcomes. The conditional forecast responds more sharply when financial conditions deteriorate. Around the GFC window, the conditional lower-tail forecast moves downward and remains below the historical benchmark, especially for HOUST. This pattern illustrates how broad financial conditions are associated with a more adverse conditional lower-tail forecast during periods of financial stress. The COVID window shows a different pattern: realized outcomes are volatile, but the forecast movement is less similar to that of the GFC episode, consistent with the view that pandemic-period construction disruptions also reflected nonfinancial shocks such as public health restrictions and supply bottlenecks. The figure is illustrative rather than a separate test, but it helps explain when the loss reductions in Table 2 arise. Table 3 asks whether the predictive content of financial conditions is concentrated in the lower tail. If financial-condition information were useful for forecasting the entire conditional distribution of PERMIT and HOUST, similar loss reductions would also appear at the median or upper quantile. If the gains are concentrated at the 20th percentile, the evidence is more consistent with the construction-at-risk framing. Table 3 compares quantile-specific forecast performance across the main financial-condition and benchmark models at the 20th, 50th, and 75th percentiles.

Table 3. Quantile comparison
Panel A. PERMIT

Model	20th	50th	75th

PERMIT	NFCI	6.47**	0.22	-0.72
PERMIT	ANFCI	5.93**	0.07	-1.74
PERMIT	Risk	5.82**	0.15	-0.79
PERMIT	Credit	-0.24	-0.15	-1.01
PERMIT	Leverage	6.33*	0.02	-2.45
PERMIT	Mortgage spread	-0.82	-1.89	-3.75
PERMIT	RCL PC1	7.70**	0.19	-1.28
PERMIT	RCL PC1+PC2	8.34**	-0.22	-0.95
Panel B. HOUST				
	Model	20th	50th	75th
HOUST	NFCI	5.92**	0.88	-0.21
HOUST	ANFCI	6.46**	0.40	-0.06
HOUST	Risk	5.61**	0.67	-0.29
HOUST	Credit	-0.03	-0.31	-0.34
HOUST	Leverage	7.33**	0.66	-0.41
HOUST	Mortgage spread	-1.25	-1.12	-1.66
HOUST	RCL PC1	7.85**	1.13	-0.02
HOUST	RCL PC1+PC2	8.65**	0.87	-1.75

Notes. This table presents quantile-comparison results across the main financial-condition and benchmark models. Panel A reports *PERMIT*, and Panel B reports *HOUST*. Columns labeled *20th*, *50th*, and *75th* report loss reduction at the corresponding percentiles. Loss reduction is measured relative to the historical quantile benchmark at the same quantile. * and ** indicate significance at the 10% and 5% levels, respectively.

Table 3 shows that the forecast gains are concentrated at the lower quantile across the main financial-condition specifications. For PERMIT, NFCI, ANFCI, Risk, Leverage, RCL PC1, and RCL PC1+PC2 all produce positive lower-tail loss reductions, whereas the median and upper-quantile results are close to zero or negative. For HOUST, the same pattern is even clearer: NFCI, ANFCI, Risk, Leverage, RCL PC1, and RCL PC1+PC2 improve the 20th-percentile forecasts, while the 50th- and 75th-percentile gains are small or negative. Credit alone and the mortgage-spread benchmark do not deliver comparable gains in the lower tail. The pattern indicates that the predictive content of financial conditions is concentrated in downside construction outcomes rather than in the center or upper part of the construction distribution. From a housing-supply and built-environment perspective, this lower-tail concentration matters because monitoring concerns are asymmetric: the relevant signal is not the central path of construction activity, but the risk of adverse construction outcomes. The quantile-specific results support the interpretation of construction-at-risk as a supply-side monitoring framework rather than a general housing-activity forecasting exercise. The next comparison asks whether this lower-tail predictive content is distinct from standard mortgage-market information. Table 4 directly compares selected financial-condition models with the mortgage-spread benchmark.

Table 4. Direct loss comparison against the mortgage-spread benchmark

	Model	Benchmark	Loss reduction (%)	NW-p	Boot-p	N
PERMIT	ANFCI	Mortgage spread	6.80**	0.025	0.038	432
PERMIT	NFCI	Mortgage spread	7.33**	0.024	0.030	432
PERMIT	Leverage	Mortgage spread	7.19*	0.062	0.092	432
PERMIT	RCL PC1	Mortgage spread	8.55**	0.034	0.056	432
PERMIT	RCL PC1+PC2	Mortgage spread	9.19**	0.027	0.050	432
HOUST	ANFCI	Mortgage spread	7.79***	0.007	0.018	432
HOUST	NFCI	Mortgage spread	7.26**	0.011	0.026	432
HOUST	Leverage	Mortgage spread	8.65**	0.013	0.028	432
HOUST	RCL PC1	Mortgage spread	9.16***	0.006	0.018	432
HOUST	RCL PC1+PC2	Mortgage spread	9.95***	0.005	0.018	432

Notes. This table presents direct loss comparisons against the mortgage-spread benchmark. *Model* denotes the target and financial-condition specification. *Benchmark* denotes the comparison model. *Loss reduction (%)* is measured relative to the mortgage-spread model. *NW-p* and *Boot-p* denote Newey-West HAC and moving-block bootstrap *p*-values, respectively. *N* denotes the number of OOS forecast origins. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 4 compares selected financial-condition specifications directly with the mortgage-spread benchmark. All reported financial-condition specifications produce lower OOS quantile loss than the mortgage-spread model for both PERMIT and HOUST. The RCL PC1+PC2 specification reduces loss by 9.19% for PERMIT and 9.95% for HOUST relative to the mortgage-spread benchmark. NFCI and ANFCI also deliver positive and statistically significant reductions in losses for both construction series. These results do not imply that mortgage-market information is unimportant. Mortgage rates are closely related to borrowing costs and housing affordability, and the mortgage spread provides a simple measure of mortgage-market tightness. The comparison highlights a narrower point: in this lower-tail construction-forecasting exercise, mortgage spread alone does not reproduce the predictive information contained in broader financial-condition measures. NFCI, ANFCI, and the RCL specifications incorporate information from risk, funding, credit, and leverage conditions that are not captured by the mortgage spread alone. The mortgage-spread comparison is intentionally narrow because it directly captures the wedge between mortgage rates and long-term Treasury yields. Appendix E extends this direct comparison to alternative rate-based benchmarks, including the mortgage rate level, the 10-year Treasury yield, and term spread measures. The RCL PC1+PC2 specification continues to yield positive loss reductions relative to these benchmarks, although the statistical strength is greater for HOUST than for PERMIT. These additional rate-benchmark checks support the interpretation that the main results are not simply a relabeling of standard interest-rate information. Table 5 examines whether the baseline results depend on the three-month horizon. Although h=3 is the baseline, financial-

condition information may be reflected in construction downside risk over shorter or longer horizons. We keep τ=0.20 fixed and compare loss reductions at 1-, 3-, 6-, 9-, and 12-month forecast horizons.

Table 5. Horizon robustness

	Model	*h*=1	*h*=3	*h*=6	*h*=9	*h*=12
PERMIT	NFCI	5.41***	6.47**	7.24***	8.92**	4.37
PERMIT	ANFCI	3.87**	5.93**	5.06**	7.49*	3.87
PERMIT	Leverage	3.90**	6.33*	7.97	11.12*	8.33
PERMIT	RCL PC1+PC2	5.51***	8.34**	10.43**	13.36**	11.39**
PERMIT	Mortgage spread	-1.30	-0.82	-2.78	-2.47	-3.65
HOUST	NFCI	2.89*	5.92**	7.83**	7.56*	5.43
HOUST	ANFCI	3.08*	6.46**	7.60**	4.84*	3.15
HOUST	Leverage	2.40*	7.33**	10.05**	9.83*	9.74*
HOUST	RCL PC1+PC2	3.34**	8.65**	12.17**	8.29*	11.13**
HOUST	Mortgage spread	0.09	-1.25	-4.08	-3.80	-4.25

Notes. This table presents horizon-robustness results at τ=0.20. *Model* denotes the target and forecasting specification. Columns report loss reduction for forecast horizons *h*=1,3,6,9, and 12 months, where *h* denotes the forecast horizon in months. Entries report qR^2 (%), defined as the percentage reduction in OOS quantile loss relative to the recursive historical lower-tail benchmark. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 5 shows that the RCL PC1+PC2 gains extend beyond the baseline three-month horizon. For both PERMIT and HOUST, RCL PC1+PC2 yields positive loss reductions across all reported horizons. NFCI and ANFCI show positive performance over several horizons, while Leverage remains relatively strong, especially for HOUST. The mortgage spread is close to zero or negative for most horizons. These results suggest that financial-condition information is not merely capturing short-lived monthly noise. It remains useful across several forecast horizons when predicting downside movements in residential construction. At the same time, the table should not be read as evidence of a structural transmission lag. A larger loss reduction at a particular horizon does not identify the timing of a causal effect. The main message is narrower: the baseline result is not driven solely by the choice of *h*=3. The longer-horizon results should be interpreted as evidence of robustness rather than as independent tests, because cumulative targets overlap more strongly as h increases. For this reason, the inference uses horizon-specific Newey-West lag lengths, and the table is used to assess whether the sign and magnitude of the baseline result are stable across horizons rather than to select an optimal forecasting horizon. Table 6 reports timing and crisis-exclusion checks for the main RCL PC1+PC2 specification. Panel A uses lagged financial predictors. This check addresses the possibility that same-month financial variables contain information from within the forecast month that would not be available at the beginning of the month. Panel B removes the global financial crisis (GFC) period, the COVID/post-COVID period, and both periods from the OOS evaluation sample. The purpose is to examine whether the baseline RCL result is mechanically driven by the same-month information or by a single major crisis episode.

Table 6. Robustness checks for the main RCL specification

	Target	Model	qR^2 (%)	NW-p	Boot-p	N
Panel A. Lagged predictor timing check						
Lagged predictor	PERMIT	RCL PC1+PC2	6.48**	0.042	0.062	432
Lagged predictor	HOUST	RCL PC1+PC2	6.72**	0.026	0.036	432
Panel B. Crisis-exclusion diagnostics						
Ex-COVID	PERMIT	RCL PC1+PC2	8.54**	0.021	0.036	429
Ex-COVID	HOUST	RCL PC1+PC2	8.85**	0.012	0.034	429
Ex-GFC	PERMIT	RCL PC1+PC2	6.89***	0.002	<0.001	413
Ex-GFC	HOUST	RCL PC1+PC2	7.52***	<0.001	<0.001	413
Ex-GFC/COVID	PERMIT	RCL PC1+PC2	7.08***	0.001	<0.001	410
Ex-GFC/COVID	HOUST	RCL PC1+PC2	7.69***	0.001	0.002	410

Notes. This table presents robustness checks for the main RCL PC1+PC2 specification. Panel A reports the *Lagged-predictor timing check*. Panel B reports *Crisis-exclusion diagnostics*. *Ex-* denotes that the indicated window is excluded from the OOS evaluation sample. Rows report robustness checks by target. *qR^2 (%)* is measured relative to the historical lower-tail benchmark. *NW-p* and *Boot-p* denote Newey-West HAC and moving-block bootstrap *p*-values, respectively. *N* denotes the number of OOS forecast origins. ** and *** indicate significance at the 5% and 1% levels, respectively.

Table 6 reports two focused checks for the RCL PC1+PC2 specification. This specification is used here because it delivers the strongest baseline performance in Table 2 and provides a compact representation of the Risk, Credit, and Leverage subindexes. The purpose is not to make RCL PC1+PC2 the only relevant model, but to examine whether the strongest financial-condition result is sensitive to timing conventions or major crisis episodes. Panel A shows that the loss reductions remain positive when financial predictors are lagged by one month. This check is deliberately conservative. It rules out the possibility that the baseline results are driven solely by financial-condition observations that become available only after the forecast-origin month begins. The smaller but still positive loss reductions indicate that the predictive content is not eliminated when the information set is shifted back by one month. This reduces concern that the baseline result relies on information from the same-month financial variables that may not be available at the forecast origin. Panel B shows that the result remains positive after excluding the GFC period, the COVID/post-COVID period, and both periods. The evidence does not appear to be mechanically driven by a single crisis episode. These crisis-exclusion exercises should not be interpreted as evidence that all stress episodes generate identical forecast performance; they show only that the main RCL result is not mechanically concentrated in a pre-specified crisis window and survives the main timing and episode-based concerns. We next examine whether forecast performance varies with the level of financial stress. Table 7 evaluates already generated OOS forecasts within high-stress windows defined by NFCI, ANFCI, Credit, Leverage, and the RCL average. The models are not re-estimated within these windows. The table asks a narrower question: whether forecast losses are lower during months when financial stress is elevated.

Table 7. High-stress window diagnostics

Stress-window	N	HOUST	PERMIT
NFCI top 20%	88	10.06	7.30
NFCI top 10%	44	13.08	10.53
ANFCI top 10%	44	14.54	10.74
Credit top 10%	44	13.43	13.73
Leverage top 10%	44	15.46	9.44
RCL average top 10%	44	14.77	9.93

Notes. This table presents high-stress window diagnostics for the main RCL PC1+PC2 forecast. Rows show *Stress-window* definitions. *N* denotes the number of months in each stress window. *HOUST* and *PERMIT* report loss reduction (%) relative to the historical lower-tail benchmark. Forecasts are estimated over the full OOS evaluation sample and evaluated within each stress window.

Table 7 reports high-stress-window diagnostics for the already generated RCL PC1+PC2 forecasts. The models are not re-estimated within these windows, and the high-stress windows are not used for model selection. The purpose is narrower: to examine whether realized OOS sample loss reductions are larger in months with elevated financial stress. The loss reductions exceed the full-sample baseline in most high-stress windows, especially for HOUST. The larger high-stress-window gains for HOUST reinforce the stage-specific interpretation. Financial stress appears particularly informative when the forecasting object is actual construction execution, although PERMIT shows positive gains as a pipeline-stage measure. This pattern is consistent with the idea that execution-stage construction activity is more exposed to financing availability, balance-sheet pressure, and project-level frictions during stress periods. For monitoring purposes, this implies that HOUST may provide a particularly useful execution-stage signal during periods of elevated financial stress, whereas PERMIT remains useful for tracking pipeline-stage weakness. This diagnostic should be interpreted as a stress-window forecast evaluation rather than a causal test. The top-10% windows contain only 44 observations, and they are defined using the financial-stress indicators themselves. The evidence is suggestive and complements, rather than replaces, the full-sample recursive forecasting results.

5. Conclusion

We develop a construction-at-risk framework for monitoring downside risk in U.S. residential construction. Building permits and housing starts are used as stage-specific quantity measures: PERMIT captures the planning-stage pipeline, while HOUST captures the execution stage of construction activity. This distinction allows us to study lower-tail risk within the housing supply process rather than price-side housing risk alone. The empirical results show that broad financial-condition measures reduce lower-tail forecast losses for both PERMIT and HOUST. NFCI, ANFCI, Risk, and Leverage generally produce positive loss reductions, and specifications that use RCL information jointly deliver the most stable performance. Across the main financial-condition specifications, the forecast gains are concentrated at the 20th percentile and are small or negative at the median and upper quantile, indicating that financial

conditions are informative primarily about downside construction outcomes rather than general movements in the construction distribution. The gains are not reproduced by the mortgage-spread benchmark and are larger during periods of elevated financial stress. PERMIT gains indicate that financial conditions contain information about weakening in the future construction pipeline. The stronger HOUST gains, especially for Leverage, RCL PC1+PC2, and high-stress windows, suggest that financial conditions are particularly informative when planned projects move into actual construction. This pattern is consistent with an execution-stage interpretation in which financing availability, builder balance sheets, and project-level frictions are more likely to matter for actual construction activity. These findings suggest that residential construction downside risk is a useful macro-financial and housing-supply monitoring indicator that captures construction weakness not fully visible in house prices, average construction growth, or aggregate residential investment. The evidence should be interpreted as predictive rather than causal: the analysis shows that financial-condition information reduces OOS quantile loss, not that the RCL specification identifies a structural financing channel. For housing and built-environment monitoring, the results imply that permits and starts should be tracked not only as contemporaneous housing indicators but also as stage-specific downside-risk measures. PERMIT-based lower-tail forecasts can indicate weakness in the future supply pipeline, whereas HOUST-based lower-tail forecasts can indicate stress in the execution of planned construction activity. Future work can extend this framework by incorporating a full release-calendar design, regional construction data, construction-lending standards, developer balance-sheet information, and local housing-supply constraints.

References

Aastveit, K.A., Albuquerque, B., & Anundsen, A.K. (2023). Changing supply elasticities and regional housing booms. *Journal of Money, Credit and Banking*, *55*(7), 1749–1783. https://doi.org/10.1111/jmcb.13009

Adams, P.A., Adrian, T., Boyarchenko, N., & Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, *37*(3), 1173–1191. https://doi.org/10.1016/j.ijforecast.2021.01.003

Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable growth. *American Economic Review*, *109*(4), 1263–1289. https://doi.org/10.1257/aer.20161923

Bernanke, B.S., Gertler, M., & Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics*, *1*(Part C), 1341–1393. https://doi.org/10.1016/s1574-0048(99)10034-x

Brave, S.A., & Butters, R.A. (2011). Monitoring financial stability: A financial conditions index approach. *Economic Perspectives*, *35*(1), 22–43.

Cañizares Martínez, C., de Bondt, G.J., & Gieseck, A. (2023). Forecasting housing investment. *Journal of Forecasting*, *42*(3), 543–565. https://doi.org/10.1002/for.2946

Chae, J., Cho, H., Kim, H., & Ryu, D. (2025). Ripple effect and predictive power of housing market returns: Evidence from the Gangnam area. *Journal of Housing and the Built Environment*, *40*(3), 1189–1224. https://doi.org/10.1007/s10901-025-10201-5

Davis, M.A., & Heathcote, J. (2005). Housing and the business cycle. *International Economic Review*, *46*(3), 751–784. https://doi.org/10.1111/j.1468-2354.2005.00345.x

Diebold, F.X., & Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Estrella, A., & Hardouvelis, G.A. (1991). The term structure as a predictor of real economic activity. *Journal of Finance*, *46*(2), 555–576. https://doi.org/10.1111/j.1540-6261.1991.tb02674.x

Estrella, A., & Mishkin, F.S. (1998). Predicting U.S. recessions: Financial variables as leading indicators. *Review of Economics and Statistics*, *80*(1), 45–61. https://doi.org/10.1162/003465398557320

Gertler, M., & Kiyotaki, N. (2010). Financial intermediation and credit policy in business cycle analysis. *Handbook of Monetary Economics*, *3*, 547–599. https://doi.org/10.1016/b978-0-444-53238-1.00011-9

Giglio, S., Kelly, B., & Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics*, *119*(3), 457–471. https://doi.org/10.1016/j.jfineco.2016.01.010

Gilchrist, S., & Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, *102*(4), 1692–1720. https://doi.org/10.1257/aer.102.4.1692

Glaeser, E.L., Gyourko, J., & Saiz, A. (2008). Housing supply and housing bubbles. *Journal of Urban Economics*, *64*(2), 198–217. https://doi.org/10.1016/j.jue.2008.07.007

Hong, J., & Ryu, D. (2023). Expectations and the housing market: A model of house price dynamics. *Bulletin of Economic Research*, *75*(4), 1242–1266. https://doi.org/10.1111/boer.12408

Kim, J., Cho, H., & Ryu, D. (2025). Global risk aversion: Driving force of future real economic activity. *Journal of Forecasting*, *44*(2), 706–729. https://doi.org/10.1002/for.3203

Leamer, E.E. (2015). Housing really is the business cycle: What survives the lessons of 2008–09? *Journal of Money, Credit and Banking*, *47*(S1), 43–50. https://doi.org/10.1111/jmcb.12189

Marfè, R., & Pénasse, J. (2024). Measuring macroeconomic tail risk. *Journal of Financial Economics*, *156*, 103838. https://doi.org/10.1016/j.jfineco.2024.103838

Murray, C.K. (2020). Time is money: How landbanking constrains housing supply. *Journal of Housing Economics*, *49*, 101708. https://doi.org/10.1016/j.jhe.2020.101708

Paciorek, A. (2013). Supply constraints and housing market dynamics. *Journal of Urban Economics*, *77*, 11–26. https://doi.org/10.1016/j.jue.2013.04.001

Pan, Z., Zhou, Y., & Dong, Q. (2025). Systemic risk and macroeconomic activity: International evidence. *Journal of Management Science and Engineering*, *10*(4), 507–533. https://doi.org/10.1016/j.jmse.2025.09.001

Park, D., Hong, J., & Ryu, D. (2024). Heterogeneous expectations in the housing market: A sugarscape agent-based model. *Journal of Housing and the Built Environment*, *39*(3), 1465–1489. https://doi.org/10.1007/s10901-024-10116-7

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, *125*(3), 1253–1296. https://doi.org/10.1162/qjec.2010.125.3.1253

Strauss, J. (2013). Does Housing Drive state-level job growth? Building permits and consumer expectations forecast a state's economic activity. *Journal of Urban Economics*, *73*(1), 77–93. https://doi.org/10.1016/j.jue.2012.07.005

Topel, R., & Rosen, S. (1988). Housing investment in the United States. *Journal of Political Economy*, *96*(4), 718–740. https://doi.org/10.1086/261560

Tu, Q., de Haan, J., & Boelhouwer, P. (2017). The mismatch between conventional house price modeling and regulated markets: Insights from the Netherlands. *Journal of Housing and the Built Environment*, *32*(3), 599–619. https://doi.org/10.1007/s10901-016-9529-y

West, K.D. (1996). Asymptotic inference about predictive ability. *Econometrica*, *64*(5), 1067–1084. https://doi.org/10.2307/2171956

Zhuang, T., Qian, Q.K., Visscher, H.J., & Elsinga, M.G. (2020). An analysis of urban renewal decision-making in China from the perspective of transaction costs theory: The case of Chongqing. *Journal of Housing and the Built Environment*, *35*(4), 1177–1199. https://doi.org/10.1007/s10901-020-09733-9

Appendix

Appendix A. Baseline quantile-choice and coverage diagnostics

Panel A. Tail-event counts by quantile

	h	τ	OOS N	Exp.Hits	Act.Hits	Hit rate	GFC	COVID	Non-crisis
PERMIT	3	0.10	432	43.2	21	4.9	7	1	13
PERMIT	3	0.20	432	86.4	64	14.8	7	1	56
HOUST	3	0.10	432	43.2	22	5.1	6	2	14
HOUST	3	0.20	432	86.4	67	15.5	8	2	57

Panel B. Forecast coverage at the baseline quantile

	Model	Exp.rate	Act.rate	Diff	Hits	N
HOUST	ANFCI	20.0	20.4	0.4	88	432
HOUST	NFCI	20.0	22.0	2.0	95	432
HOUST	RCL PC1+PC2	20.0	19.2	-0.8	83	432
PERMIT	ANFCI	20.0	18.3	-1.7	79	432
PERMIT	NFCI	20.0	19.2	-0.8	83	432
PERMIT	RCL PC1+PC2	20.0	18.1	-1.9	78	432

Notes. Panel A reports tail-event counts for alternative lower-tail quantiles at the baseline horizon. τ denotes the quantile. *Exp.Hits* equals $\tau \times$ *OOS N*. *Act.Hits* count realized cumulative construction outcomes below the recursive historical quantile benchmark. *Hit rate* is *Act.Hits*/*OOS N*. *GFC*, *COVID*, and Non-crisis report the timing of actual hits. Panel B reports coverage diagnostics for selected conditional 20th-percentile forecasts. *Exp.rate* is the nominal quantile level, *Act.rate* is the empirical hit rate, *Diff* is *Act.rate minus Exp.rate*, *Hits* is the number of observations below the model-implied conditional quantile, and *N* is the number of OOS observations.

Appendix B. Alternative construction-outcome definitions

Panel A. Raw monthly growth outcomes

	Model	qR² (%)	NW-p	Boot-p	N
PERMIT	NFCI	7.52**	0.015	0.028	432

PERMIT	ANFCI	8.56**	0.015	0.030	432
PERMIT	Leverage	7.06	0.106	0.142	432
PERMIT	Mortgage spread	-1.95	0.996	0.992	432
PERMIT	RCL PC1+PC2	9.93**	0.030	0.046	432
HOUST	NFCI	4.04**	0.046	0.060	432
HOUST	ANFCI	4.05**	0.046	0.052	432
HOUST	Leverage	6.75**	0.020	0.038	432
HOUST	Mortgage spread	-0.76	0.999	0.998	432
HOUST	RCL PC1+PC2	7.17**	0.018	0.032	432

Panel B. Macro-adjusted construction outcomes

	Model	qR² (%)	NW-p	Boot-p	N
HOUST	ANFCI	6.75**	0.036	0.050	428
HOUST	Leverage	6.06	0.103	0.134	428
HOUST	Mortgage spread	0.05	0.467	0.455	428
HOUST	NFCI	4.74*	0.081	0.114	428
HOUST	RCL PC1+PC2	7.14*	0.054	0.084	428
PERMIT	ANFCI	5.25**	0.013	0.032	428
PERMIT	Leverage	4.37	0.168	0.196	428
PERMIT	Mortgage spread	-0.76	0.917	0.904	428
PERMIT	NFCI	4.39***	0.010	0.020	428
PERMIT	RCL PC1+PC2	4.92	0.127	0.142	428

Notes. This table reports robustness results using alternative definitions of construction-outcome. Panel A uses raw monthly growth outcomes, and Panel B uses macro-adjusted outcomes. *qR² (%)* is measured relative to the corresponding historical lower-tail benchmark. *NW-p* and *Boot-p* denote Newey-West HAC and moving-block bootstrap *p*-values, respectively. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Appendix C. RCL principal-component diagnostics

Panel A. Full-sample loadings

Subindex	PC1	PC2	PC3
Risk	0.6230	0.2777	0.7313
Credit	0.4950	-0.8638	-0.0936
Leverage	0.6057	0.4203	-0.6756

Panel B. Full-sample explained variance

Component	Var.Share	Cum.Share
PC1	0.6504	0.6504

PC2	0.2318	0.8822
PC3	0.1178	1.0000

Panel C. Recursive loading summary

	Mean	SD	P25	Med	P75
Risk-PC1	0.6330	0.0086	0.6250	0.6353	0.6387
Credit-PC1	0.4801	0.0158	0.4667	0.4902	0.4943
Leverage-PC1	0.6070	0.0050	0.6050	0.6057	0.6111
Risk-PC2	0.2411	0.0292	0.2135	0.2356	0.2668
Credit-PC2	-0.8670	0.0087	-0.8741	-0.8642	-0.8631
Leverage-PC2	0.4346	0.0198	0.4251	0.4302	0.4338
Var (PC1)	0.6546	0.0085	0.6507	0.6566	0.6611
Var (PC1–PC2)	0.8944	0.0085	0.8878	0.8917	0.9010

Notes. Panel A reports full-sample PCA loadings for the *Risk*, *Credit*, and *Leverage* subindexes. Panel B reports explained and cumulative variance shares. *Var.Share* and *Cum.Share* denote component-specific and cumulative explained-variance shares. Panel C summarizes recursive PCA loadings across forecast origins. *Risk-PC1*, *Credit-PC1*, and *Leverage-PC1* denote the loadings of each subindex on PC1; *Risk-PC2*, *Credit-PC2*, and *Leverage-PC2* are defined analogously for PC2. *Var(PC1)* and *Var(PC1–PC2)* denote the variance explained by PC1 and by PC1–PC2 jointly. *SD* denotes standard deviation, *Med* denotes median, and P25 and P75 denote the 25th and 75th percentiles.

Appendix D. Baseline results with BH-adjusted *p*-values

	Model	qR² (%)	NW-p	Boot-p	BH-p
PERMIT	NFCI	6.47**	0.016	0.024	0.095
PERMIT	ANFCI	5.93**	0.016	0.024	0.095
PERMIT	Risk	5.82**	0.030	0.050	0.095
PERMIT	Credit	-0.24	0.594	0.449	0.709
PERMIT	Leverage	6.33*	0.062	0.080	0.124
PERMIT	Mortgage rate	3.18	0.143	0.150	0.229
PERMIT	10Y Treasury	2.51	0.232	0.273	0.338
PERMIT	Term spread	-0.64	0.621	0.611	0.709
PERMIT	Term spread (monthly)	-0.64	0.619	0.607	0.709
PERMIT	Mortgage spread	-0.82	0.796	0.745	0.849
PERMIT	RCL PC1	7.70**	0.029	0.046	0.095
PERMIT	RCL PC1+PC2	8.34**	0.022	0.036	0.095
HOUST	NFCI	5.92**	0.023	0.046	0.069
HOUST	ANFCI	6.46**	0.015	0.036	0.069

HOUST	Risk	5.61**	0.032	0.052	0.069
HOUST	Credit	-0.03	0.515	0.411	0.687
HOUST	Leverage	7.33**	0.026	0.044	0.069
HOUST	Mortgage rate	1.74	0.179	0.188	0.287
HOUST	10Y Treasury	1.40	0.273	0.273	0.398
HOUST	Term spread	-0.62	0.706	0.725	0.807
HOUST	Term spread (monthly)	-0.61	0.701	0.717	0.807
HOUST	Mortgage spread	-1.25	0.999	0.998	0.999
HOUST	RCL PC1	7.85**	0.014	0.026	0.069
HOUST	RCL PC1+PC2	8.65**	0.011	0.030	0.069

Notes. This table reports baseline OOS quantile-forecast results with Benjamini-Hochberg-adjusted *p*-values. The baseline uses $h = 3$ and $\tau = 0.20$. qR^2 *(%)* denotes the percentage reduction in OOS quantile loss relative to the recursive historical lower-tail benchmark. *NW-p* is the one-sided Newey-West HAC *p*-value, *Boot-p* is the one-sided moving-block bootstrap *p*-value, and *BH-p* is the Benjamini-Hochberg adjusted *p*-value computed within each target-horizon-quantile baseline comparison. * and ** denote significance at the 10% and 5% levels based on unadjusted HAC *p*-values.

Appendix E. Direct comparison against alternative rate benchmarks

	Model	Benchmark	Loss reduction (%)	NW-p	Boot-p
PERMIT	RCL PC1+PC2	Mortgage spread	9.19**	0.027	0.050
PERMIT	RCL PC1+PC2	Mortgage rate	5.44	0.144	0.164
PERMIT	RCL PC1+PC2	10Y Treasury	5.40	0.174	0.206
PERMIT	RCL PC1+PC2	Term spread	5.28	0.176	0.176
PERMIT	RCL PC1+PC2	Term spread (monthly)	5.28	0.177	0.178
HOUST	RCL PC1+PC2	Mortgage spread	9.95***	0.005	0.018
HOUST	RCL PC1+PC2	Mortgage rate	7.21**	0.044	0.076
HOUST	RCL PC1+PC2	10Y Treasury	7.75**	0.040	0.076
HOUST	RCL PC1+PC2	Term spread	6.64*	0.081	0.098
HOUST	RCL PC1+PC2	Term spread (monthly)	6.63*	0.081	0.098

Notes. This table reports direct loss comparisons between the RCL PC1+PC2 specification and alternative rate-based benchmark models. The baseline uses $h = 3$ and $\tau = 0.20$. *Loss reduction (%)* is measured relative to the benchmark model listed in each row. *NW-p* is the one-sided Newey-West HAC *p*-value, and *Boot-p* is the one-sided moving-block bootstrap *p*-value. *, **, and *** denote significance at the 10%, 5%, and 1% levels based on unadjusted HAC *p*-values.