A Work Project, presented as part of the requirements for the Award of a Master’s degree in Economics from the Nova School of Business and Economics.

Measuring Downside Exposure in Portugal’s Real Estate Market: A House Price-at-Risk Framework

Tiago Alberto Reis Campos

Work project carried out under the supervision of:

Prof. Paulo M. M. Rodrigues

20/12/2023

Abstract

This work project introduces a House Price-at-Risk (HaR) model tailored to capture the dynamics of the Portuguese housing market. Through quantile regressions, it is demonstrated that a model incorporating the quarterly house price growth, a financial stress index, and residential gross fixed capital formation proves to be a reliable predictor of the market's trajectory one-year ahead, particularly for lower percentiles. This real-time monitoring of downside risk is crucial to prevent financial crises, given the significant role typically played by the property market. Still, even though overvaluation has been reaching record levels, our findings indicate that downside exposure remains limited.

Keywords: House Price-at-Risk, Quantile Regression, Portugal, Housing Market, Macroprudential Analysis

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

I. Introduction

The housing market plays a pivotal role in both the social and economic aspects of a society. It is a fundamental component of a nation's economy and is closely intertwined with the daily lives of its citizens. During the last decades, the real estate market accounted for between 20% to 25% of the European Union's GDP, consistently maintaining its substantial contribution as one of the largest components of the overall economy over time. In 2017, the sector represented approximately 20% of Portugal's GDP, falling below the union's average. However, it had reached 25% in 2001 and had been on a declining trend since then (Rodrigues, 2021).

Additionally, the property market wields a substantial impact on an economy due to the intricate connection between house price dynamics, the overall macroeconomic environment, and financial stability. This interdependence is crucial for understanding the broader economic landscape and the potential ripple effects of housing market fluctuations. For instance, in 2017, residential investment accounted for around 7.5% of Portugal's GDP, 250 basis points below the European average, but its significance lies in the strong and discernible correlation, often characterized by lead-lag dynamics, between itself and broader macroeconomic trends (Rodrigues, 2021).

Furthermore, given that housing constitutes a significant share of a household's total net worth, fluctuations in its prices can alter consumption patterns via wealth effects. Effectively, economic theory states that households' wealth is a key driver of aggregate demand (see e.g., Friedman, 1957, Modigliani and Brumberg, 1954, and Lourenço and Rodrigues, 2017). In addition, real estate serves as valuable collateral for loans. Consequently, when house prices decline, the resulting decrease in homeowners' net worth affects their borrowing capacity, subsequently impacting consumer spending. This phenomenon is commonly referred to as the 'financial accelerator', where pro-cyclical movements in borrowers' balance sheets amplify and propagate business cycles (Bernanke, 1999).[1] Claessens, Kose, and Terrones (2012) have demonstrated that recessions tend to be more severe and last longer when there are rapid and substantial declines in house prices. Their research reveals that over two-thirds of approximately 50 systemic banking crises in recent decades were preceded by the occurrence of boom-bust patterns in property prices[2].

After the financial crisis, the concurrent surge in housing prices across numerous countries has sparked concerns regarding the possible repercussions of significant and widespread declines in the near future. Indeed, the rising real estate prices, coupled with deteriorating fundamentals such as income and household debt, have likely caused housing prices to deviate from their long-term equilibrium. Effectively, in several Euro Area countries, the assessments of house price overvaluation in the fourth quarter of 2022 surpassed the 10% mark (ECB, 2023). Moreover, in Portugal, it was shown that this divergence became more pronounced since 2017 (Rodrigues, 2021). More recently, this overvaluation measure has reached record levels as this work project will explore later.

Nevertheless, more recently, mortgage loan origination and house prices have both decelerated significantly amid higher interest rates (ECB, 2023). The swiftness and extent of the shift in real estate markets will have a significant impact on the potential stress imposed on the financial system, particularly if the correction takes on a chaotic and discrete nature. Hence, as housing bubbles[3] are usually followed by substantial price declines, it becomes crucial to assess the extent of potential downside risks, their implications for financial stability, and the measures that can be taken to prevent these risks from materializing.

Until now, academic literature concerning house prices has predominantly focused on the

---

[1] A similar pattern was noted in the lead-up to the global financial crisis (Alter, Feng, and Valckx, 2018).

[2] Furthermore, Deghi et al. (2020) show that large declines in house prices forecast future risks to economic growth.

[3] According to Stiglitz (1990), a bubble is said to occur when "fundamental" factors fail justify the current price. Thus, the price is high today only because investors believe it will be higher tomorrow.

examination of average real estate prices, through OLS regressions, and disparities from what is considered the equilibrium, using quantiles regression (QR) techniques. Concretely, QR, originally introduced by Koenker and Bassett (1978), is a non-parametric approach employed to estimate the conditional quantiles of the variable of interest. This technique provides a more comprehensive depiction of the data, enabling us to assess the impact of a covariate across the entire distribution of the response variable, rather than solely focusing on its conditional mean. Machado and Sousa (2006) proposed a framework for characterizing asset price exuberance using quantile regression. The benefit of employing this methodology lies in its capacity to assess the extremes of the empirical distribution of the series, thus enabling the identification of periods in which prices diverged from their macroeconomic drivers. Afterwards, this method has been applied to property prices, namely in the works of Lourenço and Rodrigues (2015) for Euro Area countries, and Rodrigues (2021) for Portugal.

Contemporaneous research is applying the Growth-at-Risk approach pioneered by Adrian et al. (2019) to the real estate market, often referred to as House Price-at-Risk (HaR). This technique is designed to detect the buildup of downside risk within the housing market, emphasizing the notion that it is substantial and sudden declines in property prices that pose the greatest threat to financial stability[4]. It involves predicting extreme outcomes in the lower end of the conditional distribution of house prices, typically at the 5th or 10th percentile, to proactively identify whether these exorbitant realizations will correspond to significant declines in the property market.

Indeed, several oversight organizations, including the International Monetary Fund (Deghi et al., 2020) and the European Central Bank (ECB, 2021), have created their unique HaR

[4] Effectively, Deghi et al. (2020) demonstrate that although higher debt-to-GDP ratios do not influence the median growth of house prices, they do considerably raise the likelihood of housing crises.

measures[5]. This underscores the significance of this approach in the broader macroeconomic context. Additionally, the Bank of Spain has established a HaR model aimed at specifically assessing downside risk within the Spanish housing market (Gergely and Rodríguez-Moreno, 2022).

Considering these inquiries and the most recent advancements, this work project intends to make a valuable addition to the existing literature by implementing the HaR methodology to the Portuguese real estate market. The objective is to achieve a more precise and real-time evaluation of the risk associated with extreme events. Furthermore, historical data will be leveraged to gain deeper insights into past macroeconomic and financial crises, such as the Portuguese sovereign debt crisis, with a particular focus on disentangling the effects of each explanatory variable.

II. Data

Looking at Figure 1, it becomes evident that the financial crisis left a substantial imprint on the Portuguese housing market. In fact, during the period from 2008 to 2014, there was a noteworthy drop of more than 15% in real house prices. Subsequently, price levels not only rebounded but even surpassed previous historical peaks, experiencing a surge of approximately 45% since 2014. Nonetheless, indications of a slowdown in property price growth have surfaced in the Euro Area, as reported by the ECB (2023), primarily due to tightened financial conditions.

[5] These models use a financial conditions index, real GDP growth, credit growth, an overvaluation measure (which is computed through the comparison between the observed property prices and its fundamentals), mortgage rates, population, among others, as control variables.

Figure 1 – Real House Prices 2015=100

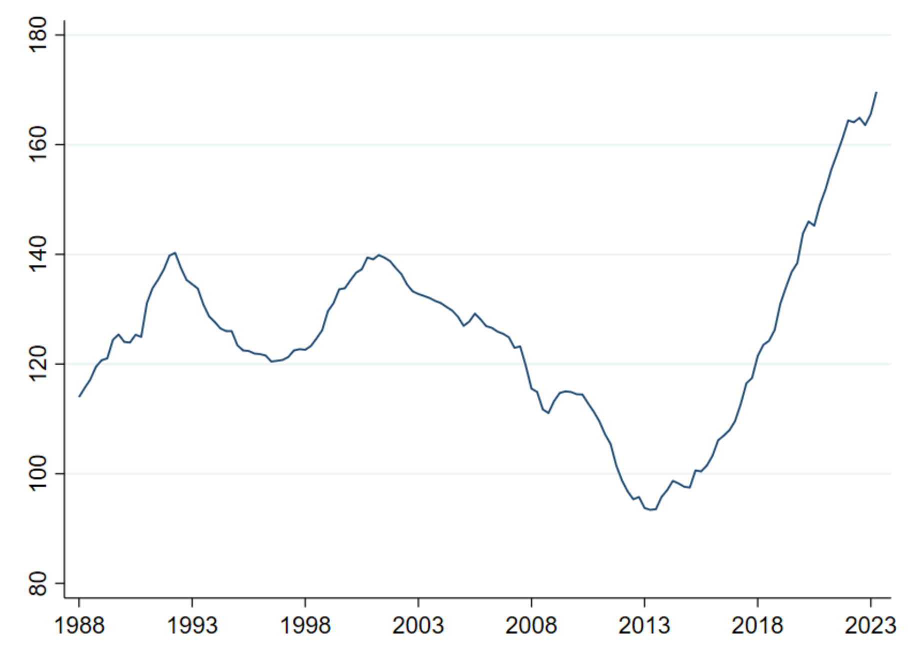

Furthermore, Figure A1 in the Appendix illustrates the distribution of one-year changes in house prices. The median stood at 0.91%, implying that most quarters display positive changes, while the 10th percentile, at -5.2%, highlights that substantial downturns are statistically significant.

When examining the primary factors responsible for these fluctuations, this analysis considers both demand and supply factors. On the demand side, one natural choice is disposable income per capita. An increase in income makes properties more accessible to economic agents, stimulating demand and consequently driving prices upward. Additionally, real mortgage rates play a pivotal role as they directly affect borrowing costs. To address demographic factors that influence housing demand, it is also include the labor force in our analysis. Conversely, on the supply side, residential gross fixed capital formation is employed as a proxy for the supply of housing[6].

Furthermore, while most studies often center their attention on the central tendencies of house prices, this analysis specifically highlights the underlying factors that exert a more pronounced impact on the lower percentiles. Thus, following the approach outlined in Gergely and

[6] The findings of OECD (2018), based on a cross-country panel analysis encompassing 20 OECD nations, including Portugal, empirically confirm the theoretical relationships mentioned.

Rodríguez-Moreno (2022) and Deghi et al. (2020), this work project will include not only the amount of credit extended to housing but also a financial stress index and an overvaluation measure in our analysis since it is shown that they both significantly influence fat tail risks in the real estate market.

Tables A1 and A2 in the Appendix provide additional details on the aforementioned variables. In addition, Figure A3 illustrates the correlations between each of them and the annualized real house price growth rate one year ahead. Notably, for instance, the financial index (FCI) exhibits a negative correlation with the 10^{th} percentile, while disposable income (Yd) displays a positive correlation across the various quantiles of our variable of interest. It should be noted that most variables are used with a natural logarithm and first-differences transformation to account for non-stationarity.

III. Methodology

As previously noted, for a formal description of the conditional connection between future housing price growth and current financial and economic factors, we rely on quantile regressions. It enables us to discern the influence of alterations in a group of conditioning variables on the structure of a dependent variable's distribution at different quantile levels. The results obtained through this approach offer valuable insights for the characterization of future probability distributions associated with the variable of interest and, consequently, the assessment of downside risks.

The quantile regression model is formally defined as follows:

$$Q_{y_{t+h}}(\tau|X_t) = X_t'\beta_\tau \tag{1}$$

where the left-hand side corresponds to the conditional quantile of the response variable, τ represents the quantile, y_{t+h} is the h-quarter changes in house prices following the forecast origin at time t, while X_t denotes the vector of explanatory variables, and β_τ is the quantile specific vector of coefficients. Effectively, this coefficient vector can be estimated consistently by minimizing the quantile-weighted absolute value of errors:

$$\hat{\beta}_{\tau,h} = arg_{\beta_\tau} min \sum_{t=1}^{T} \tau \cdot \mathbb{1}_{(y_{t+h} \geq X_t' \beta_\tau)} | y_{t+h} - X_t' \beta_\tau | + (1-\tau) \cdot \mathbb{1}_{(y_{t+h} < X_t' \beta_\tau)} | y_{t+h} - X_t' \beta_\tau | \tag{2}$$

where $\mathbb{1}$ is an indicator function signaling whether the estimated errors are positive or negative. More specifically, this analysis employed a QR model, using the variables mentioned earlier[7]. These variables have been identified in previous literature, both theoretically and empirically, as factors expected to influence house prices.

The process of model selection was guided by a rigorous evaluation of their out-of-sample performance. A total of twenty-one unique models were examined, created by pairing the seven predictor variables in combinations of two. Additionally, each model includes the current quarterly growth in real house prices. The forecast target dates range from 2003Q4+h to 2023Q2 for h-quarter-ahead predictions, providing 78-h out-of-sample periods for thorough evaluation of forecast accuracy.

Indeed, a model that excels only within the confines of the data used for its estimation may fall short in predicting future outcomes beyond the sample. Yet, the ability to do so is crucial for providing forecasts and early warnings, enabling policymakers and market participants to prepare for timely responses. Therefore, the model's out-of-sample accuracy was assessed using a quantile-based R^2 measure (Koenker and Machado, 1999):

$$quantile\ R^2 = 1 - \frac{\sum_{t=1}^{T} \tau \cdot \mathbb{1}_{(y_{t+h} \geq \hat{y}_{t+h})} | y_{t+h} - \hat{y}_{t+h} | + (1-\tau) \cdot \mathbb{1}_{(y_{t+h} < \hat{y}_{t+h})} | y_{t+h} - \hat{y}_{t+h} |}{\sum_{t=1}^{T} \tau \cdot \mathbb{1}_{(y_{t+h} \geq \bar{y})} | y_{t+h} - \bar{y} | + (1-\tau) \cdot \mathbb{1}_{(y_{t+h} < \bar{y})} | y_{t+h} - \bar{y} |} \tag{3}$$

7 See Data section.

After finding the best model, this work project conducted additional validation tests. Initially, predictions derived from both in- and out-of-sample perspectives were compared, specifically focusing on the 10th percentile. This examination aims to further assess the model's efficacy in predicting outcomes beyond its training data.

To conclude, the predictions of the optimal model were employed to compute the empirical cumulative distribution of the Probability Integral Transform (PIT) in order to evaluate their density. Specifically, the PIT measures the percentage of predictions that are below any given percentile for each h-step ahead forecast. In other words, the PIT for a given probability density function $\hat{\phi}_{t+h}$ corresponds to the cumulative density distribution of the function evaluated at Y_{t+h}:

$$z_{t+h} = \int_{-\infty}^{Y_{t+h}} \hat{\phi}_{t+h}\,(\varepsilon|\Theta)d\varepsilon \tag{4}$$

where Θ is the information set at time t and parameters of the function $\hat{\phi}_{t+h}$ are re-estimated at each t = I, …, T using expanding windows of observations. If the conditional density is correctly specified, then this probability integral transformed series should be i.i.d. U[0,1] (Diebold et al., 1998).

To account for sample uncertainty, this analysis employed the Kolmogorov-Smirnov (K-S) test. The K-S statistic takes a non-parametric approach to quantify the distance between the empirical distribution function of the sample and the cumulative distribution function of the reference distribution:

$$D_n = sup_x|F_n(x) - F(x)| \tag{5}$$

where sup_x is the largest value of the set of distances. Hence, the statistic takes the largest absolute difference between the two cumulative distribution functions, empirical and theoretical, across all x values.

Afterwards, the best model was used to evaluate the historical performance of the HaR model, using the 10th, 50th and 90th percentile as well as the 4-quarter horizon, where especial attention was placed on the lower quantiles as they serve as indicators of potential downside risk. Hence, this work project disentangled the effects of each predictor on the forecasts over time of the 10th quantile. In addition, a probability distribution function was generated for the predictions for the second quarter of 2023, encompassing quantiles ranging from 1 to 99. In this context, a skew-normal distribution was fitted to the predictions to accommodate the observed asymmetry in the data.

The skew-normal distribution was introduced by Azzalini (1985) and is defined by the following expression:

$$f(y;\, \lambda, \delta, \alpha) = \frac{2\theta(z)\Phi(\alpha z)}{\delta}\,, \quad y \in \mathbb{R}\ (\alpha, \lambda \in \mathbb{R}, \delta \in \mathbb{R}^{+}) \tag{6}$$

where $z = \frac{y-\lambda}{\delta}$, θ and Φ denote the pdf and the cumulative distribution function, respectively, λ is a location parameter, δ a scale parameter, and α a shape parameter.

Subsequently, this project excluded the quarterly growth of real house prices as an explanatory variable to conduct further robustness tests. This strategy allows us to assess whether our two remaining predictors still uphold their standing as the optimal model. In addition, it will verify if the earlier findings remain applicable even when omitting the current price levels, which could have introduced bias into our assessment of downside risk.

IV. Results

Model Selection

As mentioned above, the process of model selection involved employing an out-of-sample accuracy measure, specifically the quantile R^2. The outcomes are detailed in Table A3 in the Appendix, revealing that the optimal model across different quantiles is the one incorporating the financial conditions index (FCI) and the gross fixed capital formation (GFCF). Although certain models may exhibit marginal superiority in specific quantiles, the optimal one systematically ranks among the top performers for each quantile, distinguishing itself as the sole model to achieve this level of consistency. Figure 2 explores the quantile R^2 of each decile using the best model. Moreover, through a comprehensive analysis of the interplay between each optimal variable and all others, we can affirm that the financial conditions index proves to be a reliable predictor for the 10th quantile, while gross fixed capital formation excels in higher ones. Hence, combining these two yields the most effective solution across the wider spectrum of percentiles.

Figure 2 – Quantile R^2 Results of the Optimal Model

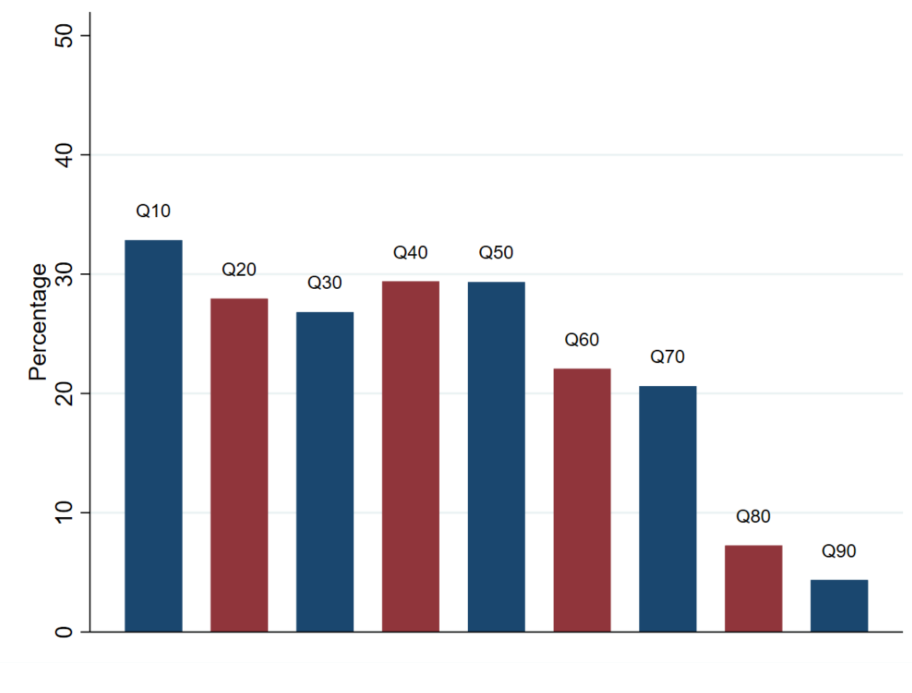

Upon reviewing the results, it becomes apparent that this model exhibits a solid performance, especially in lower quantiles, exceeding the base model's (which only employs a constant) performance by over 30%. This aligns with the primary objective of the project, as it aims to study downside exposure. Also, the model effectively captures the variance of our variable of interest up to the 70th percentile and, while its performance decreases beyond the 80th, it still outperforms the base model and surpasses the majority of other models in the selection poll.

In addition, going back to the results of Table A3, it is noteworthy that the degree of overvaluation surprisingly fails to effectively predict lower quantiles, challenging both theoretical expectations and empirical observations in other countries (Gergely and Rodríguez-Moreno, 2022). Figure 3 shows that the most significant housing market crises in Portugal were likely not driven by this measure of overvaluation, as the values were not exceptionally high during those periods. In fact, those downturns were primarily propelled by financial upheavals, as will be further explored later. However, it is worth noting that overvaluation tends to be lower when house prices experience faster growth one year ahead, providing an explanation for why this measure proves to be a more effective predictor for the 90th percentile.

Figure 3 – Overvaluation over Time

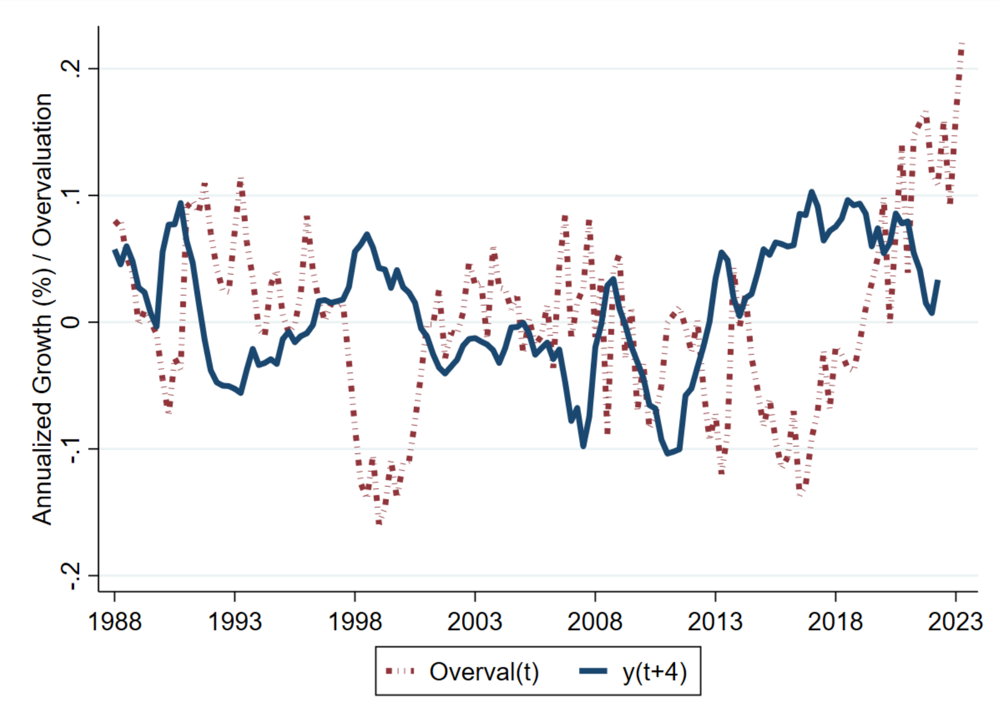

Model Validation

Before moving forward to explore the results of the optimal model, it is relevant to run further validation tests. Figure 4 compares the in- and out-of-sample predictions of the 10th percentile using the best model. Effectively, it is shown that both values are very similar during the period under analysis, suggesting that the model is performing well not only on the data it was trained on, but also on new, out-of-sample data.

Figure 4 – In- and Out-of-Sample Predictions for Q10

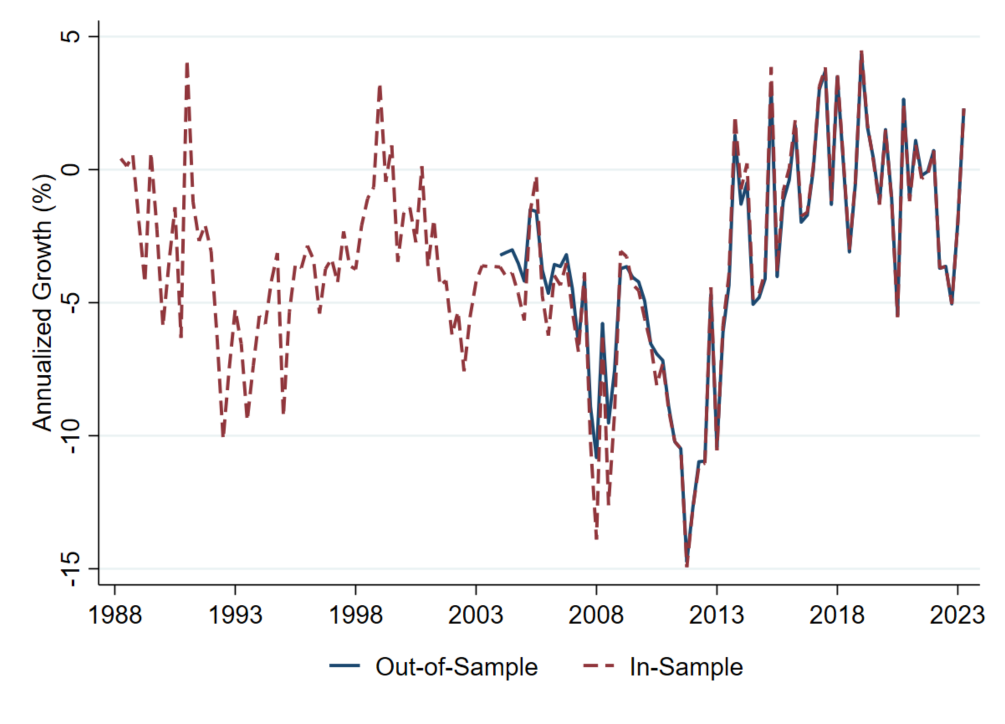

Additionally, Figure 5 employs the out-of-sample predictions, generated by the best model for each decile, to construct the empirical cumulative distribution function of the PIT. Naturally, as explored in the methodology section, the model's accuracy is higher when its empirical cumulative distribution function closely aligns with the 45-degree line, representing the cumulative distribution of a uniform distribution. The critical values were computed through the Kolmogorov-Smirnov test, as in Deghi et al. (2020) and Adrian et al. (2019). The results show that we cannot reject the null hypothesis that the sample is drawn from the reference distribution, implying that the model is well calibrated.

Figure 5 – PIT for Annualized Real House Price Growth Predictions

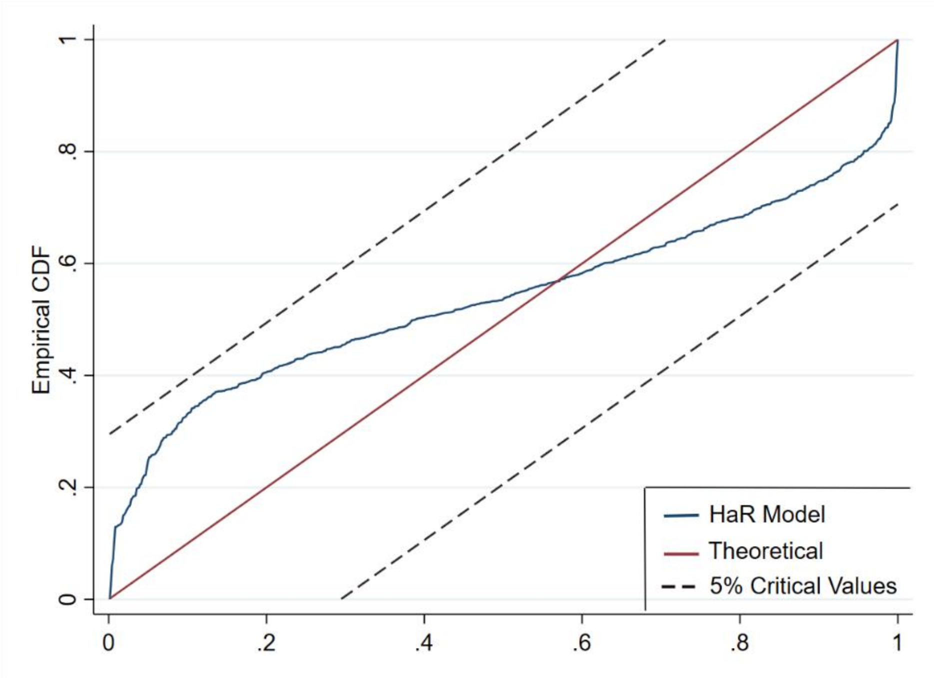

Predictions

Having completed the validation process for our model, the analysis can now move on. The optimal model is employed in Figure 6, illustrating the historical progression of house price growth alongside the forecasted percentiles (10^{th}, 50^{th}, and 90^{th}), estimated on a recursive basis. To ensure comparability, the dates were synchronized and, hence, for each period 't' it is presented the actual change in house price between 't' and 't+h' (realization) and the forecasted quantiles at the corresponding time 't'. Specifically, this analysis uses 'h'=4 to depict the one-year ahead predictions. Moreover, it encompasses data up to the second quarter of 2023, the latest available information, used for forecasting the corresponding period in 2024.

Figure 6 - Historical Realization and HaR, 4-Quarter Window

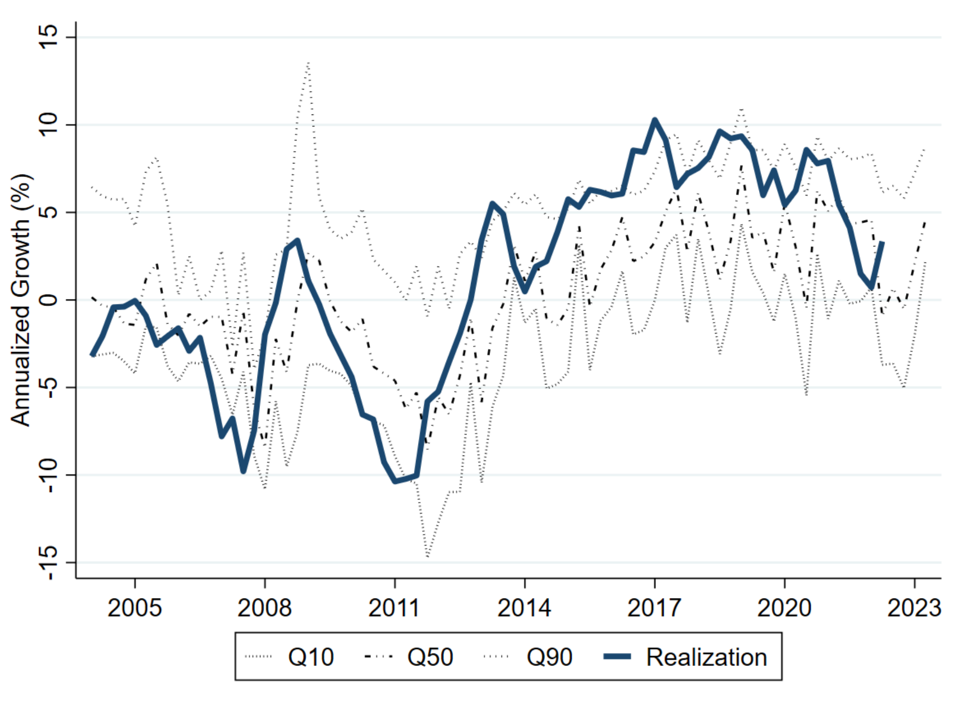

During the tumultuous 2008 financial crisis, particularly impactful in Portugal, realizations more frequently gravitated towards the lower forecasted quantile, a trend that persisted until 2011. Effectively, actual outcomes clung persistently to lower percentiles during that timeframe, and despite initial signs of recovery in prices following the initial shock, they subsequently plummeted once again. It is widely known that the aftermath of the financial crisis endured for an extended period, particularly in Southern European countries such as Portugal. In fact, the initial housing and financial shock not only prolonged the challenges but also exacerbated them, contributing to a bigger issue due to the mounting public debt in these nations, which lead to the need for international financial assistance and the implementation of austerity measures.

Now, the focus will turn to the 10^{th} percentile. Figure 7 unravels the impacts of each previously chosen optimal factor, omitting the constant, on the projected 10^{th} percentile over time. This analysis aims to provide a clearer understanding, elucidating, for instance, the primary drivers contributing to the buildup of downside risk during the financial crisis.

Figure 7 - Effects of Optimal Factors on HaR

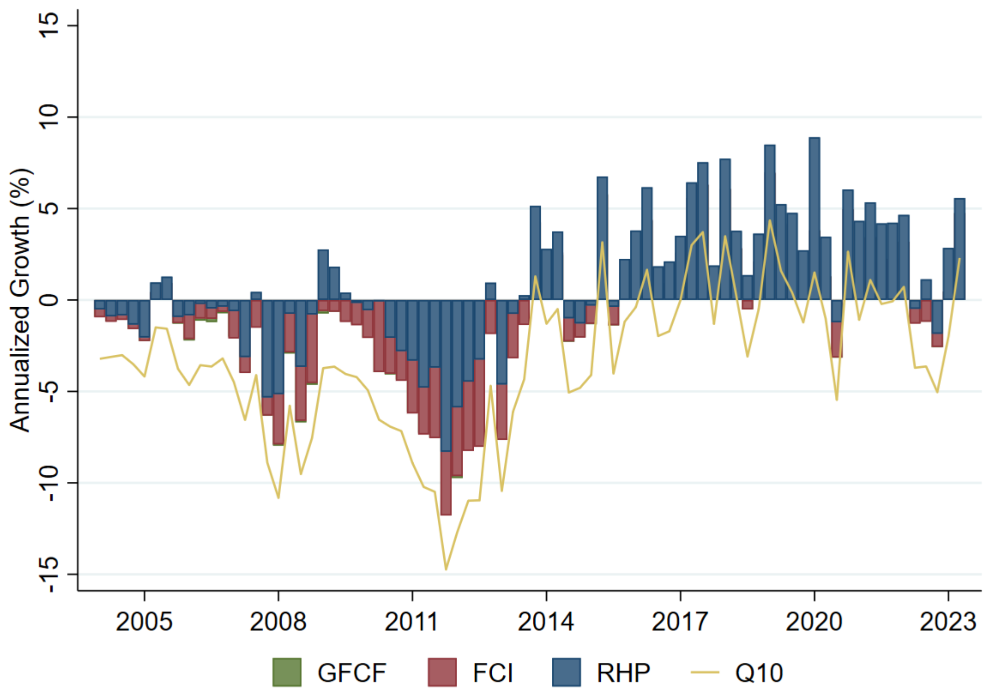

Examining Figure 7 allows for the inference that the initial escalation of downside risk before 2008 was primarily propelled by a deceleration in real house prices. In fact, research suggests that during that period there was a widespread practice of purchasing houses with the optimistic anticipation of future price growth, the so-called turkey problem or naïve empiricism (Gerardi et al., 2009). However, during the initial deceleration of prices, individuals engaged in this practice may be caught unprepared, leading to rapid real estate sales and a fast increase in supply as they adjust expectations.

Moreover, housing prices indirectly encapsulate the comprehensive dynamics of the real estate market, thereby already encompassing the impacts of various factors, such as the mortgage rate, which, although pertinent, are not explicitly incorporated as variables in the model. Also, as discussed earlier, housing constitutes a significant portion of a household's total wealth. Thus, a decline in house prices diminishes household wealth, potentially impacting demand for such assets in the future.

Subsequently, between the initial and second periods marked by significant downturns, real house prices appeared to stabilize, aligning with improved model forecasts for the 10th percentiles in the future. However, despite this stabilization, financial conditions remained unfavorable, as evidenced by the enduring significance of the "FCI" effect during that timeframe. Afterwards, prices once again entered a downward spiral, leading to the consideration that the model might have exhibited an overreaction to the easing decline in real house prices. On the other hand, it suggests that financial turmoil do indeed have a notable influence on the downside risk within the Portuguese real estate market.

Now, upon closer examination of the forecasts presented in Figure 8, employing values of 'h' ranging from 1 to 4 to encompass all available information, it becomes evident that the predictions are favorable across all depicted periods and percentiles. Even so, it is important to delve further into the analysis, particularly for the one-year forecast, as elaborated below.

Figure 8 – Predictions h-Quarters Ahead

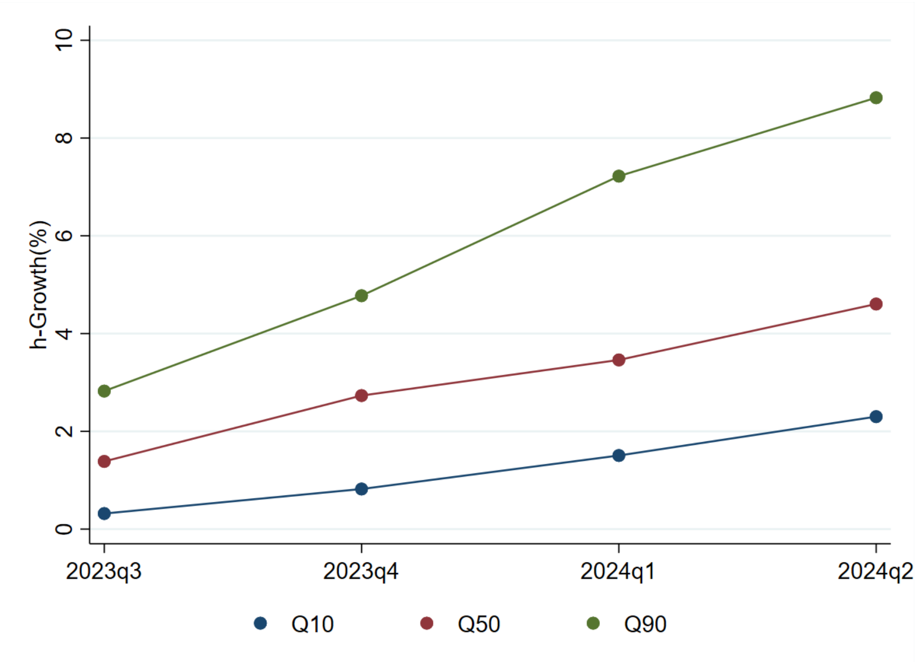

Effectively, the final objective of this House Prices-at-Risk (HaR) approach has always been real-time risk monitoring. To fulfill this goal, a meticulous analysis of the one-year ahead forecast using the most recent data becomes imperative. Figure 9, on the left, visually represents the forecasted density function and its 10th quantile. In the right panel it is shown the decomposition of the HaR model at that 10th percentile into each explanatory variable of the model (as well as the constant).

Figure 9 – House Price-at-Risk in 2023Q2

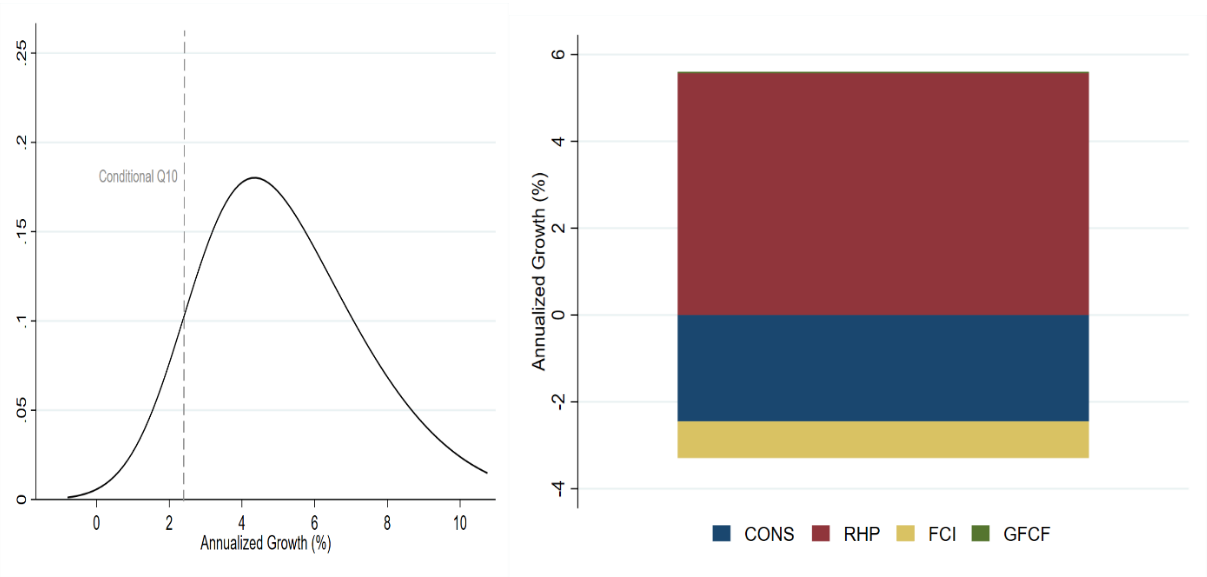

The annualized real house prices growth unconditional 10th percentile, spanning the entire analysis period, stands at -5.2%, while the most up to date conditional Q10 is 2.3%. Thus, the conditional Q10 is not only positive but also well surpasses the unconditional measure, indicating that downside risk in the market remains subdued. Additionally, examining the right panel reveals that the current predicted Q10 is primarily influenced by the quarterly growth of real house prices (RHP), which captures the latest dynamics of the market, including the recent upswing in inflation.

Robustness Tests

Still, as noticed before, there are instances in which the quarterly growth of house prices might potentially understate our downside risk measure, especially during periods when the "fundamentals" show signs of instability, but housing prices remain resilient. To explore this, this work project implements a second model with identical explanatory factors, except for

“RHP”. In that context, if the lower quantiles remain subdued, it serves as a positive indication that the fundamentals are doing well, in line with the behavior of the market. Figure 10 explores the results obtained through the recursive implementation of the alternative model.

Figure 10- Historical Realization and HaR of Alternative Model, 4-Quarter Window

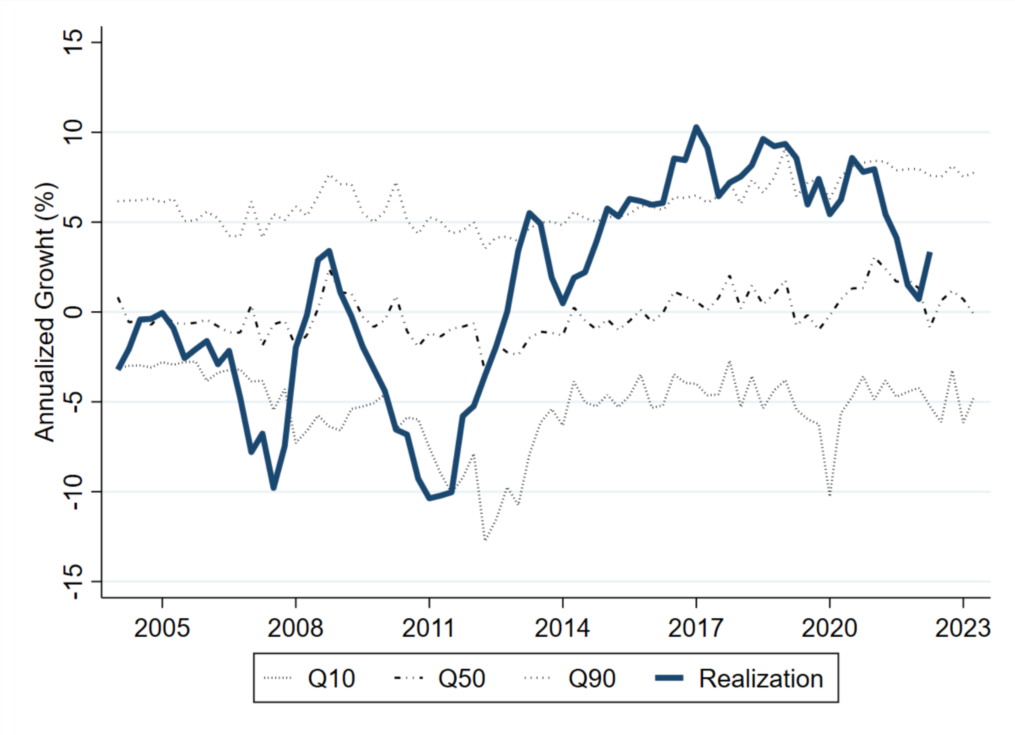

Clearly, this model exhibits greater stability compared to the previous one, as it does not react to fluctuations in house prices, which can undergo substantial variations over time. However, upon reviewing the out-of-sample accuracy measurements presented in Table A4 of the Appendix, it is possible to infer that models without “RHP” perform notably worse than their counterparts. Once more, this is logical as quarterly real house prices encapsulate, both directly and indirectly, information that underlies the pricing system.

Yet, this model performs significantly better than the base one. The forecasts for the 10th percentile are lower during the financial crisis, indicating that the model accurately captured, albeit with a slight initial delay, the contemporaneous challenges within the housing market. Moreover, it did not exhibit an exaggerated response to the moderating prices in between the two most substantial declines, in contrast to the behavior of the first model.

Finally, an analysis of the second quarter of 2023 suggests that the current downside risk is currently contained, aligning with the findings of the previous model. Figure 11 delves deeper into the most up to date forecasts of the model. In addition, Figure A5 in the Appendix explores further the predictions for each h-quarter horizon.

Figure 11 - House Price-at-Risk in 2023Q2 Alternative Model

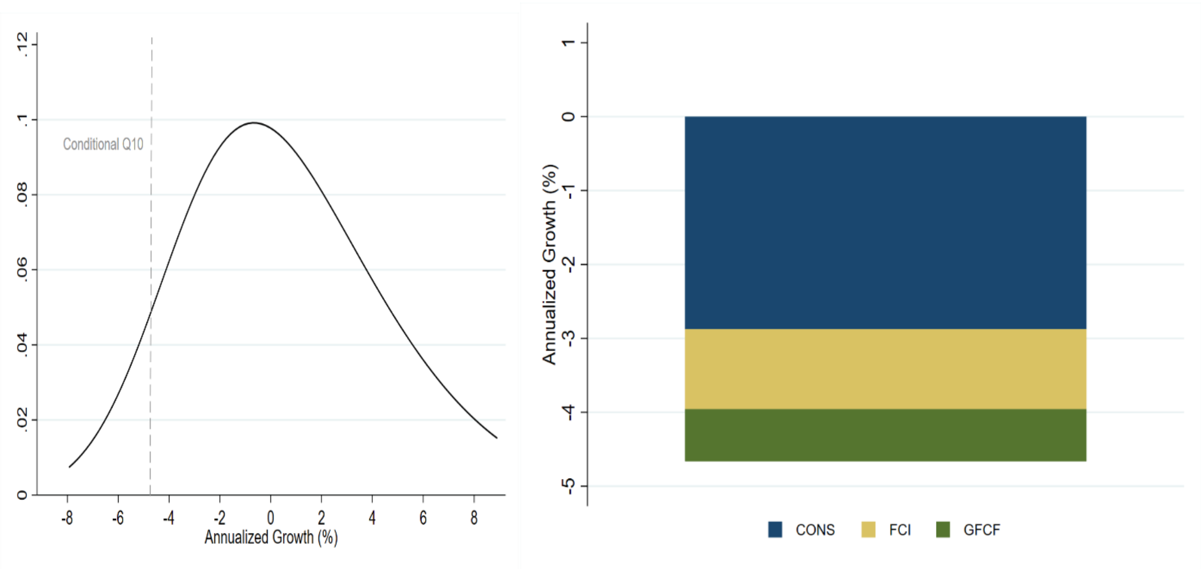

On the left side, the probability density function for 2023Q2 exhibits, as anticipated, a higher variance compared to the one derived from the optimal model since it has a worse performance. Even so, the unconditional Q10 is -5.2%, as mentioned before, while the conditional one is only -4.7%, affirming that potential losses are currently limited. On the other hand, the right panel indicates that the effects of the financial conditions index and gross fixed capital formation on the 10th percentile remain muted, signifying that the fundamentals of the housing market remain stable, despite challenges such as the Covid-19 pandemic, elevated inflation, and signs of a recession in Europe.

V. Final Remarks

In Portugal, data on housing prices is accessible only from 1988 onwards, yielding a total of 142 data points. This is evidently insufficient to encompass the richness and complexity of such macroeconomic variables. The challenge is further amplified for the HaR approach employed in this work project, which focuses on lower percentiles that, by definition, possess even fewer observations. Moreover, it presents obvious challenges for long-term forecasting.

To address these limitations, the emphasis of this analysis has been on the 10^{th} quantile, deviating from the more conventional 5^{th} quantile in the literature. In addition, this project confines the forecasting horizon to four quarters. However, these adjustments introduce additional challenges. For instance, the impact of certain variables on lower percentiles may experience a considerable lag, which is not accounted for in the one-year ahead model. Academic research indicates that monetary policies demonstrate a substantial transmission lag, especially in developed economies, ranging from 25 to 50 months (Havranek and Rusnak, 2013).

Moreover, setbacks arise due to the influence of exogenous variables, some of which could be difficult to measure, resulting in unexplained variance of our variable of interest. For instance, in the Portuguese context, tourism and foreign investment have been identified as significant factors contributing to the current upward pressure on prices. However, as our model does not integrate these variables, and considering their relatively recent impact on the real estate market, it is deduced that overvaluation is reaching record levels even though this might not be the case. Nonetheless, there has been a prevailing belief that the housing market is indeed in a bubble that would burst anytime soon and, effectively, the data suggests that the pricing system has been exhibiting exuberant behavior. Yet, this does not necessarily imply an imminent sharp decline in prices to realign with its "fundamentals". The stability of the financial conditions index, a historically significant factor influencing lower quantiles of house price growth, suggests the possibility of a soft landing. In fact, prices have been slowing down in a controlled way. Even so, it remains crucial to diligently monitor downside risk, given the unprecedented levels of overvaluation, and, therefore, past data may not be suited to properly forecast lower quantiles.

Despite all the challenges, this work project has endeavored to do the best with the available resources to create the Portugal's first House Prices-at-Risk model, which has the potential to offer a better understanding of actual risk within the real estate market. An idea that is based on science and data but, evidently, needs to be complemented with a more holistic view on the economy, given the discussed limitations of the model and broader issues associated with empiricism. As David Hume, a devoted empiricist, stated: "just because something has happened regularly in the past, does not mean it will happen the same way in the future". Still, we believe it is more likely.

References

Adrian, Tobias, Federico Grinberg, Nellie Liang, and Sheheryar Malik. 2022. "The Term Structure of Growth-at-Risk." American Economic Journal: Macroeconomics 14 (3): 283–323.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone. 2019. "Vulnerable Growth." The American Economic Review 109 (4): 1263–89.

Azzalini, Adelchi. 1985. "A Class of Distributions Which Includes the Normal Ones." Scandinavian Journal of Statistics 12 (2): 171–78.

Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework". Handbook of Macroeconomics. Vol. 1. Elsevier.

Claessens, Stijn, M. Ayhan Köse, and Marco E. Terrones. 2012. "How Do Business and Financial Cycles Interact?" Journal of International Economics 87 (1): 178–90.

Deghi, Andrea, Mitsuru Katagiri, Sohaib Shahid, and Nico Valckx. 2020. "Predicting Downside Risks to House Prices and Macro-Financial Stability." IMF Working Paper 20 (11).

Diebold, Francis X., Todd A. Gunther, and Anthony S. Tay. 1998. "Evaluating Density Forecasts with Applications to Financial Risk Management." International Economic Review 39 (4): 863.

ECB. 2021. Financial Stability Review, November.

ECB. 2023. Financial Stability Review, May.

Geng, Nan. 2018. “Fundamental Drivers of House Prices in Advanced Economies.” IMF Working Paper 18 (164): 1.

Ganics, Gergely, and María Rodríguez-Moreno. 2022. “A House Price-at-Risk Model to Monitor the Downside Risk for the Spanish Housing Market.” Documentos De Trabajo.

Gerardi, Kristopher, Andreas Lehnert, Shane M. Sherland, and Paul S. Willen. 2009. “Making Sense of the Subprime Crisis.” *Brookings Papers on Economic Activity* 2008 (2): 69–159.

Havráne, Tomáš, and Marek Rusnák. 2012. “Transmission Lags of Monetary Policy: A Meta-Analysis.” RePEc: Research Papers in Economics, October.

Koenker, Roger, and José a. F. Machado. 1999. “Goodness of Fit and Related Inference Processes for Quantile Regression.” Journal of the American Statistical Association 94 (448): 1296–1310.

Koenker, Roger, and Gilbert W. Bassett. 1978. “Regression Quantiles.” Econometrica 46 (1): 33.

Lourenço, Rita, and Paulo M. M. Rodrigues. 2015. “House Prices: Bubbles, Exuberance or Something Else? Evidence from Euro Area Countries.” Working Papers 2015. Banco de Portugal Economic and Research Department.

Machado, José Ferreira, and Jorge Pedro Sousa. 2006. “Identifying Asset Price Booms and Busts with Quantile Regressions.” RePEc: Research Papers in Economics, January.

Marchenko, Yulia, and Marc G. Genton. 2010. “A Suite of Commands for Fitting the Skew-Normal and Skew-t Models.” The Stata Journal 10 (4): 507–39.

Rodrigues, Paulo M. M., and Rita Fradique Lourenço. 2017. “House Prices in Portugal - What Happened since the Crisis?” Banco De Portugal, January.

Rodrigues, Paulo M. M., ed. 2022. The Real Estate Market in Portugal: Prices, Rents, Tourism and Accessibility. Estudos da Fundação.

Stiglitz, Joseph E. 1990. “Symposium on Bubbles.” Journal of Economic Perspectives 4 (2): 13–18.

Appendix

Figure A1 – Frequency Distribution of House Price Growth

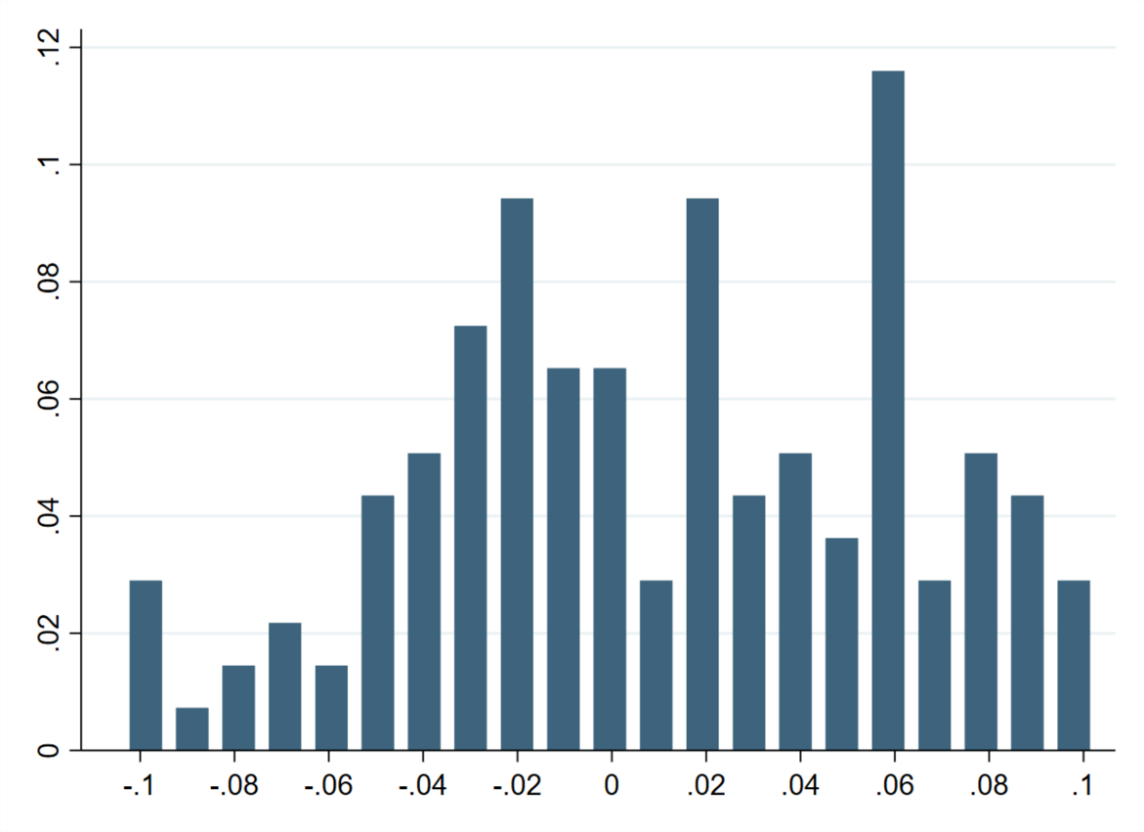

Figure A2 – Correlations Between Potential Determinants and House Price Growth

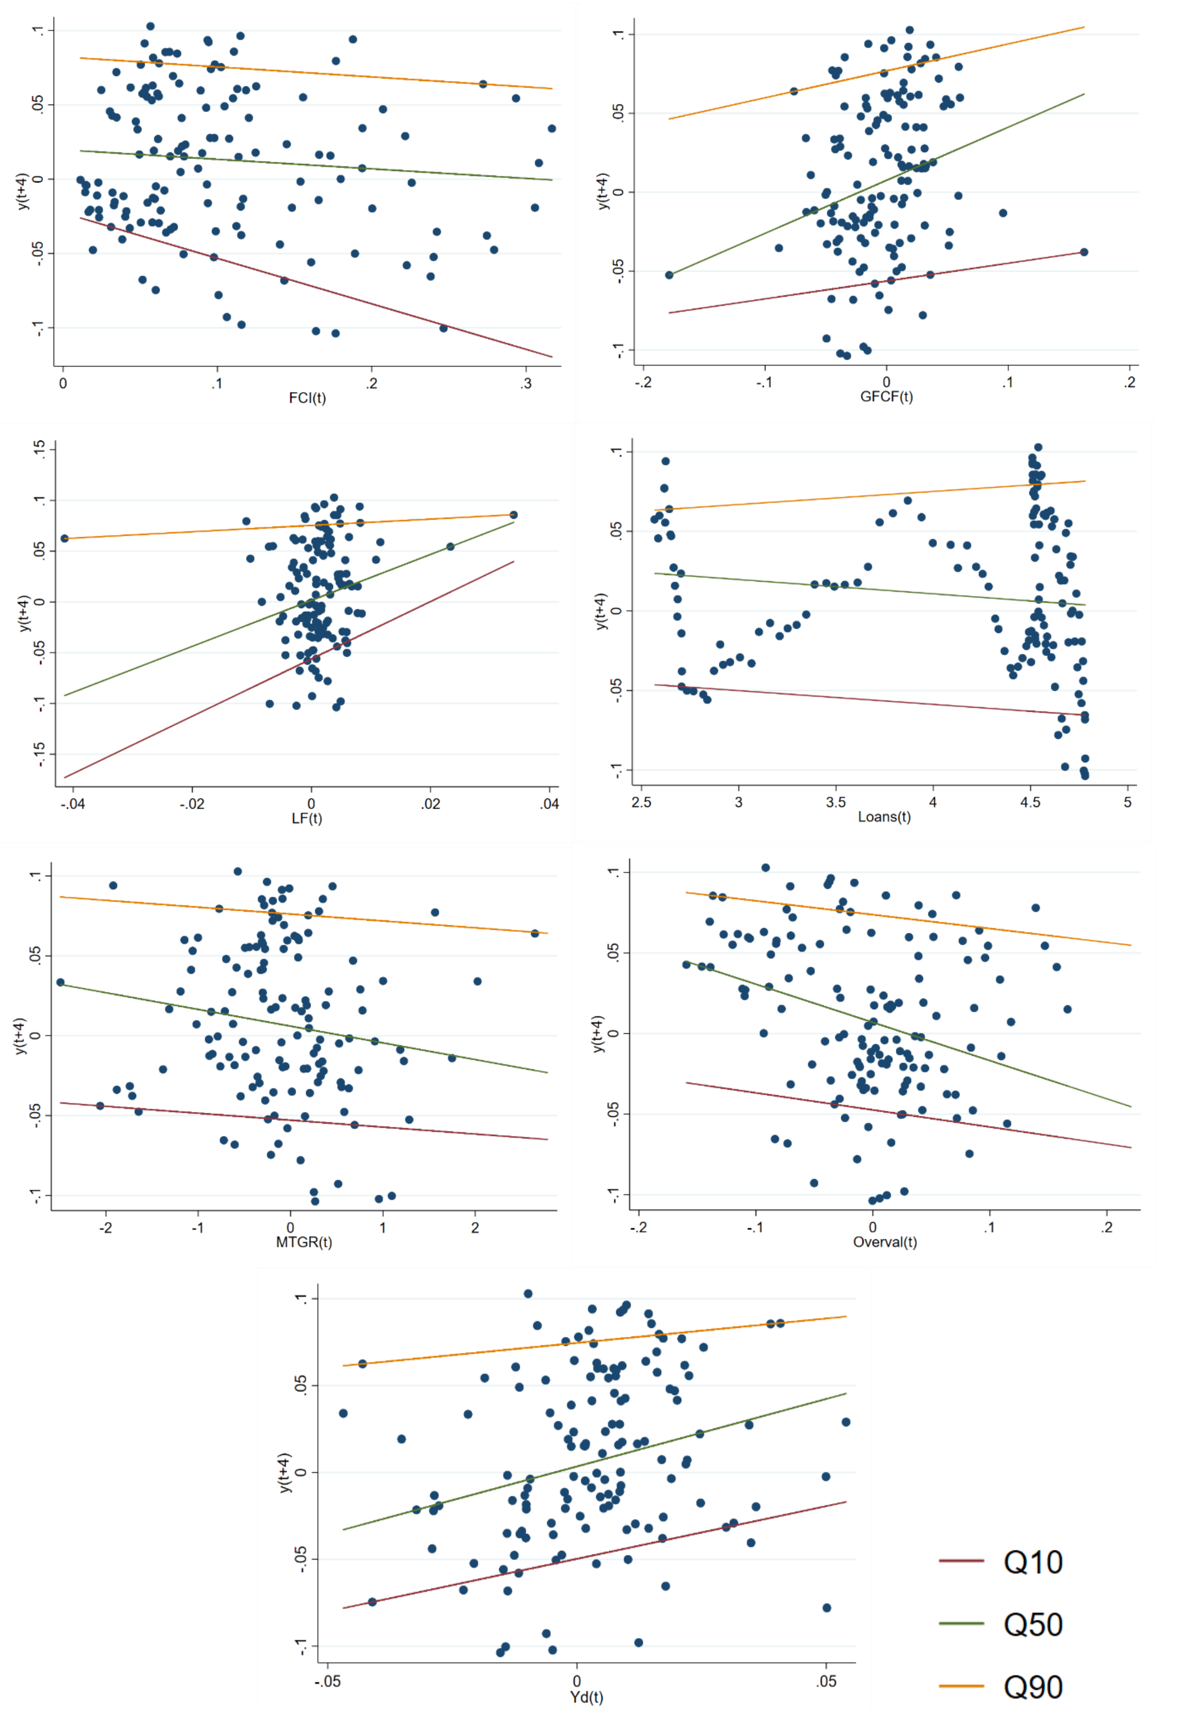

Figure A3 – Impact of Optimal Factors on House Price Growth

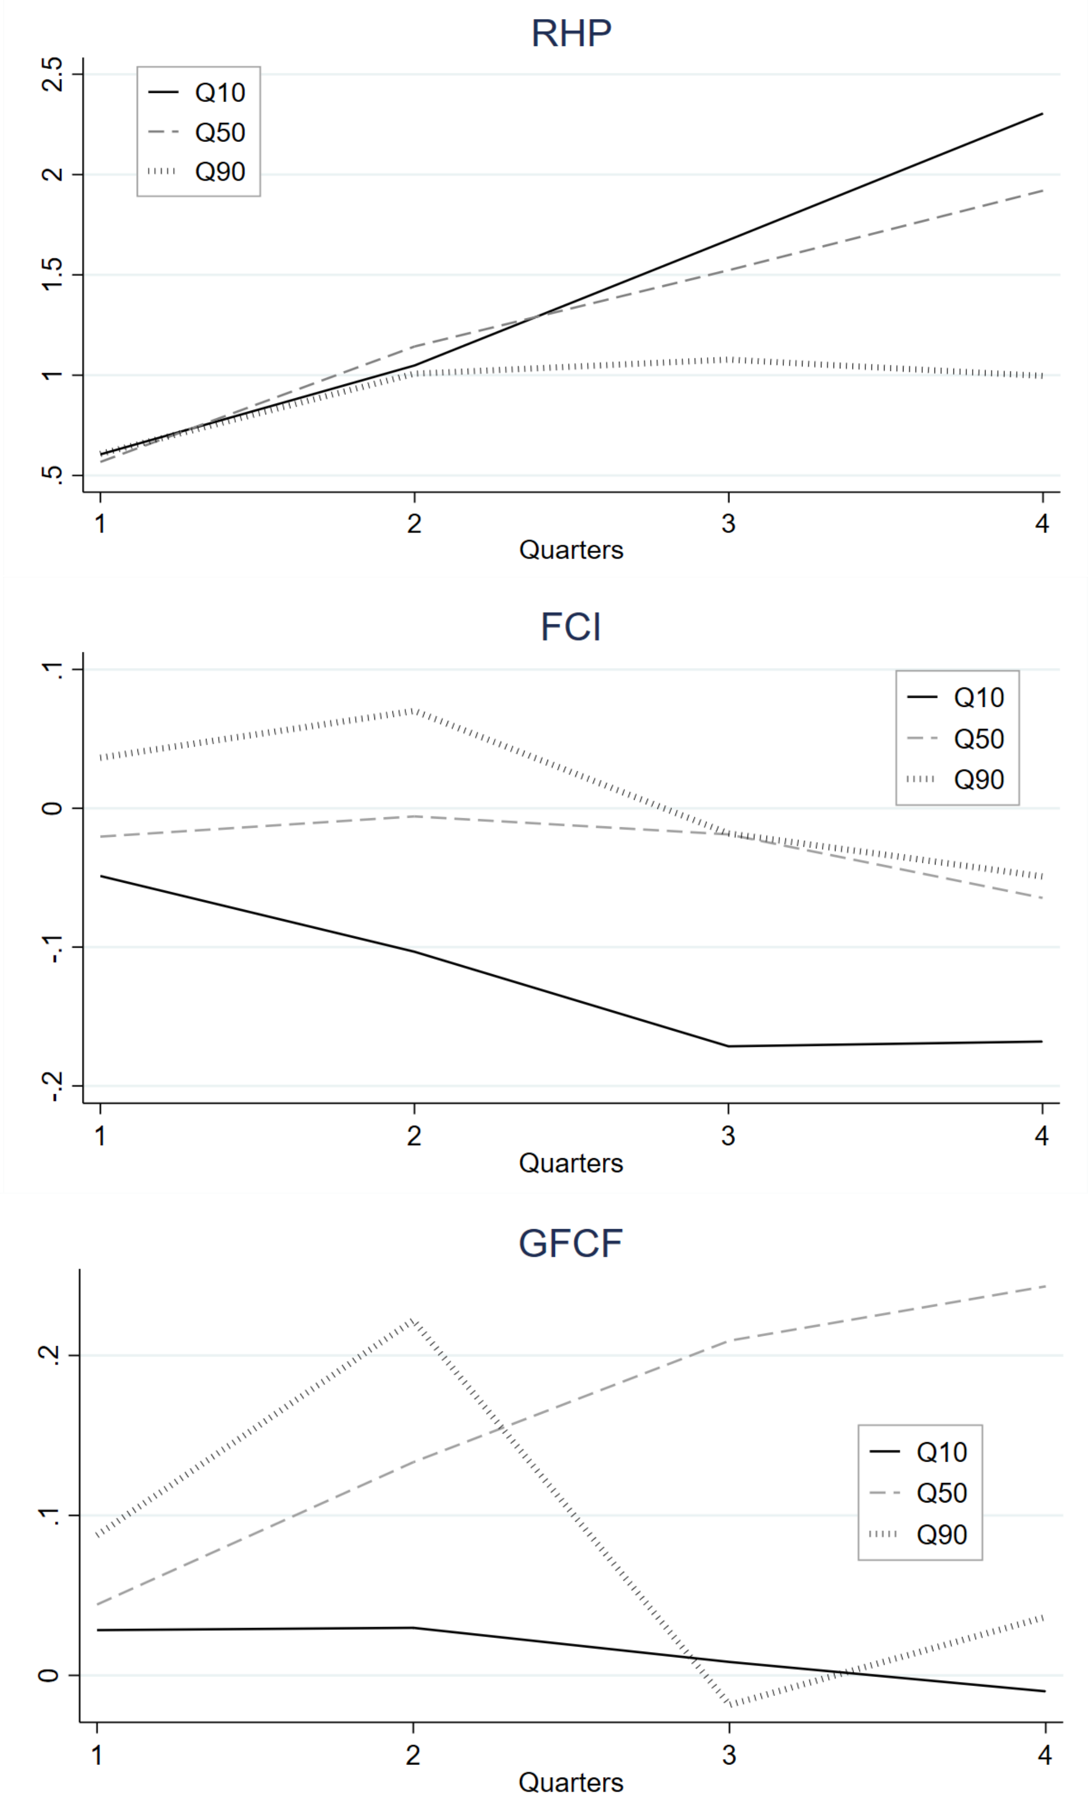

Figure A4 – Effects of Optimal Factors on Alternative Model's HaR

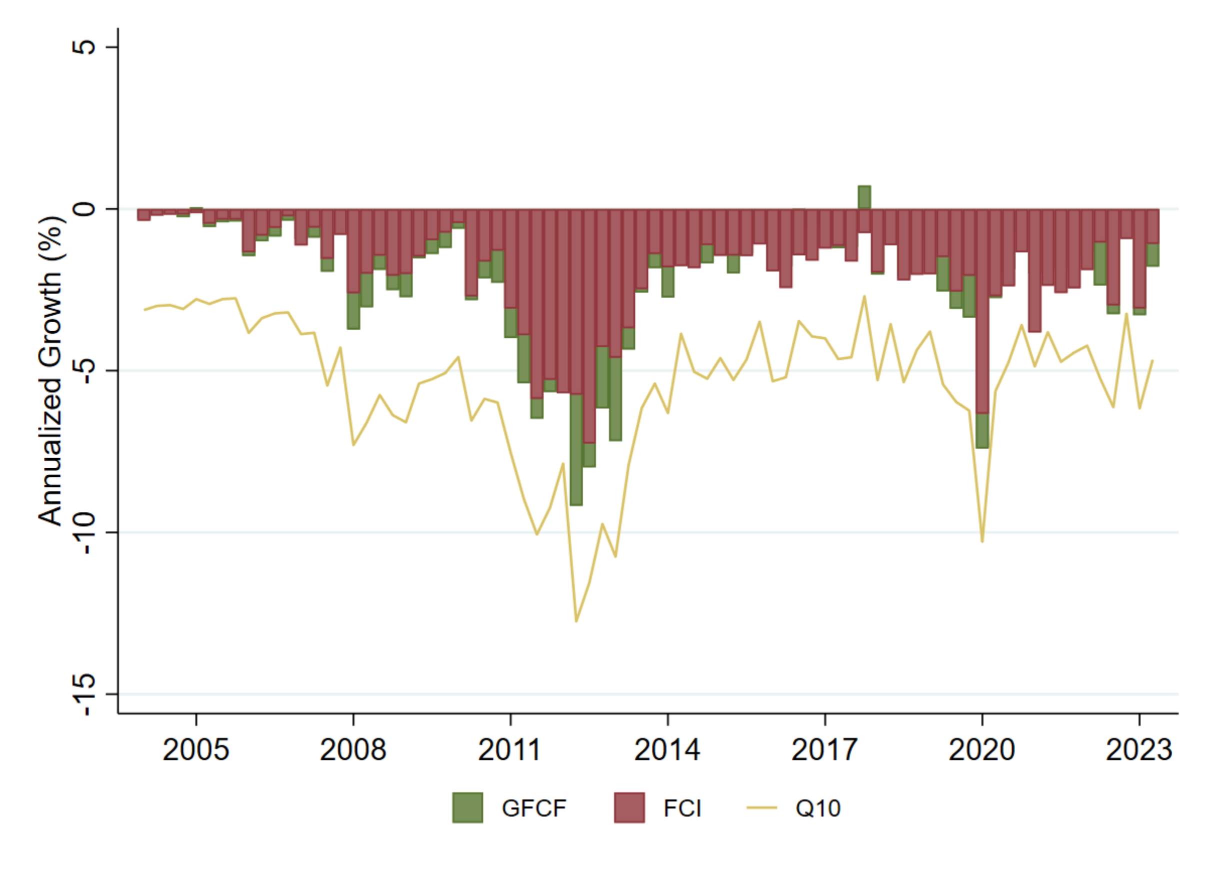

Figure A5 – Predictions h-Quarters Ahead Alternative Model

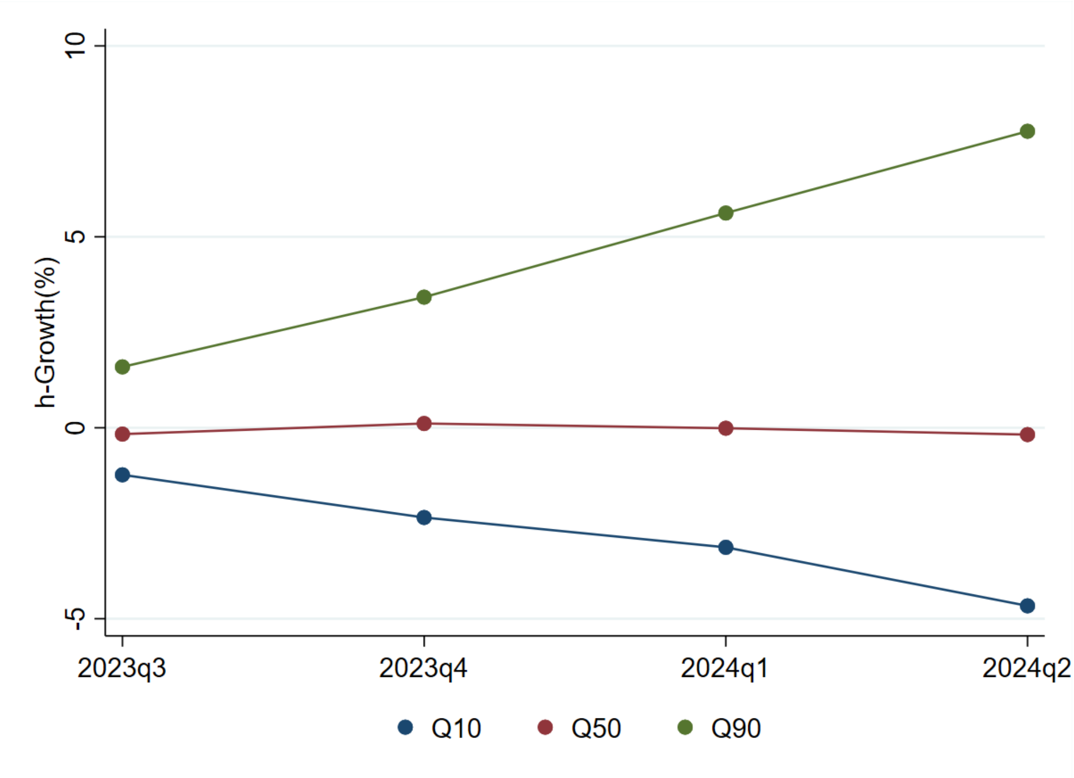

Table A1 – Summary of Variables

Variable	**Description**	**Source**
Real House Prices (RHP)	Quarterly change in residential real estate prices adjusted for inflation.	INE/Confidencial
Index of Financial Stress (FCI)	Stress index identifying periods of high financial stress in various markets.	ECB
Gross Fixed Capital Formation (GFCF)	Quarter-on-quarter changes in gross fixed capital formation within the housing market.	INE/BdP
Housing Loans (Loans)	Natural logarithm of real property loans.	BdP/ECB
Mortgage Rate (MTGR)	Quarterly difference in financing costs adjusted for inflation.	ECB
Labour Force (LF)	Quarter-on-quarter changes in the economically active population.	BdP
Disposable Income (Yd)	Quarter-on-quarter changes in real per capita disposable income.	BdP
Overvaluation (Overval)	A measure derived by comparing housing prices with the underlying market fundamentals.	Author's calculation

Notes: INE refers to the Portuguese National Institute of Statistics, BdP is the Bank of Portugal, while ECB stands for the European Central Bank.

Table A2 – Descriptive Statistics

	Units	# Obs.	Mean	Median	P10	P90	SD
RHP	Percent	141	0.28	0.11	-1.59	2.29	1.54
FCI	-	141	0.102	0.078	0.032	0.207	0.071
GFCF	Percent	141	-0.42	-0.27	-4.47	3.59	3.79
Loans	-	141	4.09	4.51	2.71	4.73	0.76
MTGR	-	141	-0.11	-0.1	-1.02	0.74	0.8
LF	Percent	141	0.13	0.11	-0.3	0.6	0.62
Yd	Percent	141	0.34	0.4	-1.54	2.25	0.18
Overval	-	141	0	0	-.093	.091	.072

Table A3 – Out-of-Sample Accuracy Measure for Different Specifications

Q10	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.328	0.372	0.345	0.339	0.366	0.160
GFCF			0.156	0.214	0.255	0.221	0.082
Loans				0.163	0.242	0.243	0.077
MTGR					0.217	0.228	0.085
LF						0.233	0.058
Yd							0.041
Overval							

Q50	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.293	0.147	0.267	0.233	0.245	0.280
GFCF			0.302	0.279	0.276	0.302	0.320
Loans				0.266	0.233	0.261	0.245
MTGR					0.240	0.270	0.312
LF						0.272	0.288
Yd							0.294
Overval							

Q90	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.044	-1.573	-0.162	-0.566	-0.390	0.131
GFCF			-0.910	-0.111	-0.211	-0.281	0.129
Loans				-0.867	-1.233	-0.754	-1.518
MTGR					-0.323	-0.276	0.106
LF						-0.594	-0.126
Yd							0.175
Overval							

Table A4 – Out-of-Sample Accuracy Measure for Different Specifications of Alternative Model

Q10	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.091	0.075	0.079	0.103	0.055	0.013
GFCF			-0.062	0.037	0.076	0.088	-0.089
Loans				-0.101	-0.055	-0.059	-0.092
MTGR					0.053	0.088	-0.183
LF						0.030	-0.164
Yd							-0.108
Overval							

Q50	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.043	-0.083	-0.002	-0.027	-0.039	0.001
GFCF			0.028	0.051	-0.004	-0.047	0.026
Loans				-0.063	-0.076	-0.048	-0.056
MTGR					-0.047	-0.049	0.000
LF						-0.040	-0.001
Yd							-0.025
Overval							

Q90	FCI	GFCF	Loans	MTGR	LF	Yd	Overval
FCI		0.003	-0.783	-0.113	-0.082	-0.592	-0.030
GFCF			-1.169	-0.038	-0.118	-0.565	-0.152
Loans				-0.918	-0.619	-1.059	-2.154
MTGR					-0.079	-0.586	-0.082
LF						-0.704	-0.192
Yd							-0.590
Overval							