Global Housing Returns, Discount Rates, and the Emergence of the Safe Asset, 1465-2024*

Paul Schmelzing†

First version: May 14, 2025. This version: August 27, 2025

Abstract

This paper reconstructs global housing returns over centuries and presents new stylized facts, by combining archival data with a novel backcasting machine learning approach. Contrary to consensus, housing valuations have been highly dynamic over the long-horizon – and patterns over recent decades align with deeper trends of rising (excess) returns and prices. Housing lends itself elegantly to a reconstruction of plausible ranges of discount rates over time: I show that discount rates exhibit a clear secular downward trend, consistent with rising housing valuations. The counterpart is a rising "safety premium", the emergence of which I attempt to pinpoint in time.

JEL Codes: G12, G19, N01.

---

*I thank Cornel Zwierlein for sharing German housing statistics, Thomas Wenderoth for Nuremberg rental data, and Walter Bauernfeind for discussing details of the TOPO N database project, Sébastien Pouget, William Goetzmann and David Le Bris for sharing discount rate data, and Allison Shertzer for sharing U.S. data in advance, as well as Kenneth Rogoff, John Cochrane, Andrey Malenko, Christophe Spaenjers, Barry Eichengreen, Allison Shertzer, Vincent Bogousslavsky, and Paul Goldsmith-Pinkham (discussant) for comments, as well as participants at the 2025 SFS Cavalcade, Hoover Institution EPWG, and IMF Fiscal Affairs division seminars – and I thank Dan Smith, Guoqi Chen and Ruiyang Wang for outstanding research assistance.

†Boston College and Hoover Institution, Stanford. Contact: schmelzi@bc.edu, schmelzing@stanford.edu.

1 Introduction

How have very long-run asset prices evolved, and did recent decades feature any key inflection points in the context of centuries of data? What is the backdrop to the "low expected returns" environment characterizing the early 21st century? This paper makes the case that a critical amount of empirical evidence has accumulated over recent years to reach better answers to these questions. In particular, this paper will address these questions from the perspective of real estate: since at least the Renaissance, real estate – in the form of structures and land – consistently represented the largest financial asset by value, a status it continues to hold right into the present. Together with sovereign bonds, real estate accounted for the majority of both modern and early modern wealth and investors even centuries ago adjusted portfolios depending on the relative yield spreads between these two assets.[1]

But despite important recent forays into the historical performance of real estate – much of it motivated by the generational "boom and bust" of housing in the 20th and 21st centuries – its comprehensive long-horizon properties remain widely understudied, certainly compared to progress made for other asset classes. Yet this long-horizon could be essential, since we also better appreciate that long-horizon analyses can potentially qualify very basic insights into the financial and econometric properties of an asset class and the backdrop to the elevated valuations of the early 21st century documented for key assets (Rogoff et al., 2024). For instance, the canonical paper establishing the benchmark methodology of modern house price indices and its basic econometric properties – Case and Shiller (1989) – only covered the horizon of 1970-1986, and expanding this methodology consistently into the past has always run into major source and measurement problems. Related obstacles make even recent ambitious "global" housing overviews – and propositions of specific "turning points" in the 20th century – likely sensitive to cleaner estimations and much longer data.[2]

More generally, housing assets play a central role in understanding the broader contours of the very long-horizon asset universe, not just for their absolute size, but given that their valuations are known to embed very long-horizon cash flow expectations (Giglio et al., 2015), from which discount rate estimates can be derived. Indeed, despite great interest in the finance literature over recent years in discount rate variation, its secular empirical properties have surprisingly not received much attention (Cochrane, 2011), including the very basic question on the existence of long-horizon time trends and structural breaks in plausible long-horizon approximation.

The first contribution of this paper is therefore to make key empirical progress, and offer new very long-horizon real estate time series that replicate modern methodologies. Property-level repeat sales sources are introduced that allow the reconstruction of house price indices consistent

[1] On capital stock and financial asset composition estimates over time see Goldsmith (1985), and more recent data in Schmelzing (2026). On investors diversifying their portfolios dynamically between residential real estate and sovereign bonds in 18th century markets, see Korevaar (2021).

[2] See Knoll et al. (2017) for the proposition of a "hockey stick" trajectory of global house prices, inflecting only in the 1960s, with the authors using data that starts in the late 1800s.

with modern benchmark methodologies (Case-Shiller repeat sales index). In addition, I introduce primary data for a number of covariates of house prices and returns over the very long run, including a new annual series for (real) German mortgage rates, and building cost indices for key advanced economies. Armed with these and other long-run time series, the empirical environment lends itself almost ideally to a novel machine learning (ML) exercise, which will comprise the second key contribution of this paper. Nonparametric machine learning approaches are already widely used for forward-looking predictive approaches in the finance literature – but their potential for backward reconstructions ("backcasting") of financial time series has so far not been realized. Machine learning trained on modern data can achieve very high out-of-sample predictive power and help more generally to reconstruct financial trends over past centuries: I demonstrate this on the basis of house price reconstructions for the U.S., the U.K., France, and the Netherlands – yielding new multi-century "global" price and return indices.

Having established new price and return series through two very different independent methodologies (archival source and ML), we can establish the commonalities across them, and newly assess consensus views in the real estate literature – which is the third contribution in this paper. Real estate does indeed show forecastability over the very long-run: as I show, such forecastability holds for a variety of real estate return measures, including plausible adjustments for housing costs, and also excess returns. Ex post real estate total returns are increasing over the very long run, and both primary source indices and those reconstructed via machine learning agree on key qualifications of current literature: the very long history of housing markets and valuations suggest that existing propositions of a "hockey stick" in global housing markets that only inflected in the mid-20th century appear untenable: this paper shows that house prices and housing returns were highly dynamic prior to the 20th century, and on many measures the elevated house valuations of the 21st century appear to merely extend deeper multi-century trends.

These results emphasize long-run continuities behind contemporary asset pricing – but what drives such long-horizon patterns housing returns and valuations? Leading current frameworks put discount rate variation at the center – but the long-run context to discount rate variation remains obscure. Housing markets, however, lend themselves elegantly to a first comprehensive analysis of discount rate trends, which this paper then undertakes.

To address these questions, in the final part of the paper, I therefore reconstruct various plausible definitions of discount rates over time. This includes a new primary-sourced multi-century data set for mortgage rates – arguably the most relevant discount rate for housing markets. This new data – together with the fact that rent growth is a stationary variable – allows me to compare implied present values of housing with actual realized prices, over the multi-century horizon. The data echoes the continuities of the current low discount rate environment – suggesting that discount rates are secularly falling, and therefore constitute a plausible prime driver of rising housing valuations – as opposed to explanations that center around rising rent growth expectations. Importantly, comparing implied present values in housing with realized prices does not confirm the existence of a "housing boom" in recent decades. Relative to the trend fall in discount rates, instead, housing in advanced economies tended to be valued at a discount.

Across plausible definitions, the fall in discount rates was flatter than the fall in sovereign rates: I identify an important cross-over with sovereign rates that occurred around the late 1600s, and show that the counterpart of falling discount rates are secularly rising "safety premia" for sovereign assets.

Regarding the paper outline – after discussing literature, part 3 will be concerned with the construction of the new German multi-century primary source real house price index, including cost estimates, with section 4 then presenting machine learning reconstructions for global series, on the basis of long-run historical covariates (such as building cost indices). Section 5 focuses on German and global total returns, their decomposition, and their respective trends. Afterwards, we will present comprehensive econometric analyses in section 6, where I formally try to establish structural breaks, stationarity, and adjustment speeds in real estate over the long-run – before turning, in section 7, to the holistic view of housing and the reconstruction of discount rates – arguing that rising housing valuations are plausibly related to secularly falling discount rates – partly on the basis of new multi-century reconstructions of mortgage rates from archives. Here, I also assess the spread between rental and sovereign yields over time, and structural breaks. Since circa the late 1600s, the spread is generally positive and moderately rising, thus pointing towards the "emergence" of sovereign assets as assets commanding a premium in markets then. Part 8 concludes.

2 Literature

The analysis of long-run housing market dynamics has been reinvigorated over recent years, though not only significant empirical drawbacks, but also key controversies remain over both basic trends in real estate itself – as well as its performance relative to other asset classes.

In a comprehensive paper, Knoll et al. (2017) reconstructed house price series for 14 advanced economies over the years 1870-2012, mostly using selected urban price data points for earlier half of their sample. The authors find that real house prices for residential properties in advanced economies were stagnant in the pre-1945 era, and afterwards began a long-run acceleration from the second half of the 20th century. Specifically, the authors concluded that:

- "From the last quarter of the nineteenth to the mid-twentieth century, house prices in most industrial economies were largely constant in real (CPI-deflated) terms. By the 1960s they were, on average, not much higher than they were on the eve of World War I. They have been on a long and pronounced ascent since then, giving rise to a hockey-stick pattern of house prices in the long run."

To reach this proposition, Knoll et al. (2017) undertake a widely used structural break test for the German and 13 other real house price indices and identify the years around 1964 as the key inflection point from which house prices began their secular "hockey stick" acceleration.

This German inflection point (1964) is close or even identical to the majority of other advanced economies the authors test.[3]

In a subsequent extension of the work, Jorda et al. (2019) contextualize long-run real estate asset performance in the universe of asset returns and suggested that real total housing returns outperformed equities over the long-run, even when adjusting for the volatility of such returns.

Methodologically, the approach in Knoll et al. (2017) and Jorda et al. (2019) has been challenged by Chambers et al. (2021) on the basis of new U.K. data over 1901-1983, and by Eichholtz et al. (2021) for Dutch and French 1800-1970 data. Both studies introduce new primary housing data and suggest that once rental income and expenditure components are properly integrated into long-run residential housing return data, housing in fact performed much worse than suggested in Jorda et al. (2019), and in fact fails to outperform equities when adjusted for volatility.

In a recent study, Lyons et al. (2024) reconstruct U.S. house price and total return series over 1890-2006, using a hedonic index approach covering 30 cities on an annual basis using newspaper sources: their data revises several stylized facts in the Case-Shiller index, and proposes that the muted volatility for housing as an asset class suggested by Jorda et al. (2019) could be the result of the aggregation method on the national level: however, the authors also find stagnant real house price dynamics prior to the mid-20th century, and –like all related recent longer-run literature – exclude any econometric analyses.

Methodologically, machine learning advances have overwhelmingly been used to project *forward* price and return variables in finance and economics literature, for instance Gu et al. (2020). One methodological innovation of this paper is to use machine learning to backcast financial variables, and to cross-validate any general results via new primary (archival) sources. The literature on such financial backcasting is very young – though similar approaches have been influential for climate-related reconstructions. Most recently in economics, a new interdisciplinary study exists in Koch et al. (2024), who attempt to "backcast" per capita GDP and improve on the classic Maddison income data set. A currently unpublished study discussing a similar approach for historical England/U.K. wage reconstructions exists in Paker et al. (2025).

As for the drivers of housing boom and bust cycles, there continue to exist alternative emphases, many of which are not mutually exclusive. Demographic factors and personal beliefs, for instance, have recently been invoked (Landvoigt, 2017; Korevaar and Francke, 2024). But a substantial majority of studies stresses the role of credit markets, and here especially the credit supply side, prominently Mian and Sufi (2009). The fall in U.S. mortgage rates has been identified as a key variable capturing an array of lending constraints, the weakening of which promoted rising credit supply in the wake of the 2008 crisis (Justiniano et al., 2019). It is therefore particularly intriguing to consider the very long-horizon trend in housing cost of capital – and the existing gap in such empirics motivates my introduction of multi-century German mortgage data in this paper.

[3]Eight out of the 14 advanced economy real house price indices record the key structural break between 1952-1964 on their basis; the authors use a standard Bai and Perron (2003) break test, allowing for up to three breaks.

The specific econometric literature that does exist and is concerned with forecastability of house prices (returns) deals virtually exclusively with post-1945 U.S. data: relevant contributions here include Campbell et al. (2009), who study U.S. data over 1975-2007, and the canonical Case and Shiller (1989), who found forecastability in U.S. house price data over 1970-1986. An overview of recent real estate-specific econometric literature is provided in Ghysels et al. (2013).

Relevant for the final, discount rate section, recent long-horizon equity research has associated historical time-variation in dividend yields as evidence of variation in "expected returns" (Golez and Koudjis, 2018), with a specific contribution in LeBris et al. (2019) applying a present value approach for the long run, which enables them to directly compare implied and realized historical equity valuations, an approach I will extend and apply to housing. More generally, Giglio et al. (2015) have made the case that residential housing valuations embed long horizon expectations of discount rates, with a substantial share of price differentials in U.K. and Singaporean contemporary markets being explained by cashflow expectations in the distant future – partly on this basis, we are motivated to reconstruct long-run past implied discount rates in housing: here it is particularly striking that an extensive finance literature operates with the canonical assumption of "time varying discount rates" – but that the understanding of the *secular* time-variation is still highly understudied Cochrane (2011). In the context of this literature, section 7 of this paper will undertake attempts to shed some critical empirical light on plausible past ranges of discount rates, using both past and contemporary propositions on how to approximate "discount rates", including the classic Campbell and Shiller (1988).

Finally, Korevaar (2021) demonstrates that investors in 18th century Holland actively shifted portfolio exposure from sovereign bonds to residential real estate conditional on expected returns and perceptions of asset "safety": low sovereign yields in this environment boosted housing total returns – and this suggests that a direct comparison of rental yields and sovereign yields, together with their relative valuations, could better pinpoint the emergence of sovereigns as "safe asset providers", the historical performance of which has attracted much recent interest in finance literature, e.g. Lustig et al. (2023) or Coppola et al. (2024).

3 Constructing a primary source long-run repeat-sales house price index (RHPI), Germany 1465-2024

Our basis for the construction of the new German housing price and total return indices are city level individual property records, known in German as *Häuserbücher* or *Hauschronik*. These are comprehensive compilations of cadaster and other official municipal records that allow a full chronological reconstruction of property level ownership, transactions, repeat sales prices, personal information on the occupants, and (occasionally) rental details. In effect, the data allows a construction of an index methodologically mirroring the widely used Case-Shiller repeat sales index for the U.S., which is still regarded as the most robust long-run approach controlling for heterogeneous housing characteristics.

Major progress has been made in recent years in the digitization of methodologically sufficient housing records over very long periods of time. German archives have spearheaded some of these initiatives, which – combined with new multi-century granular data for related asset classes – makes this country especially suited for new efforts at a holistic and robust general asset market reconstruction. Figure 1 displays an example, an entry from the new "Topo N" project for the city of Nuremberg (one of the key early modern financial hubs of Europe), which as of 2025 has systematically linked close to 4,000 property level records via unit-identifiers, and merged this cadaster information with ownership histories, sales prices, and other notable events related to the property.[4] Such advancements mirror those in other cities and allow repeat-sales observations for thousands of properties over centuries for consistent individual housing units – for more than a dozen key German cities from the late medieval period.

These new sources have several advantages: perhaps most importantly, their level of detail on the property level allow a full repeat-sales index construction, in line with benchmark contemporary indices: such a methodology is considered the "gold standard" among various alternative approaches, since it allows a better control for quality changes in housing, a distinct concern for real estate as an asset class in existing literature. Second, as these records merge official with notary information, they contain *both* tax assessment values and market values of the property, with the latter being much preferred as a measure of actual price dynamics, though much literature (including the Case-Shiller index) has had to rely on assessment or self-reported values for earlier periods thus far, which contain problematic biases.[5] Third, the records are unusually consistent and complete: the average length for individual properties included in the index that is continuously documented stands at 201.5 years, with a maximum of 407 years, and more than one-third with over 300 years of continuous records. Typically, property related information is quite abundant in this source type, with a significant number of property level records reporting details such as building year of the property, underlying land values, incidences of significant material damages, and quality changes (significant renovations, additions, or mergers with other properties).

In the construction of the index, I focus solely on actual market transaction prices – ignoring tax assessment values of properties, which are often given, as these are noisy given variation in property tax rates, irregular assessment, and underreporting of actual property values by owners.

For the headline index, and all headline return figures, I value-weight observations across all properties for the headline indices, which rebalances index weights based on property values every single year (using interpolations for years where individual properties are not on the market). In the Appendix, I also show alternatively an arithmetically weighted version, which suggests higher (and upwards sloping) real house price growth prior to the Thirty Years War, but otherwise is closely comparable regarding secular trends, econometric properties, and inflection points.[6]

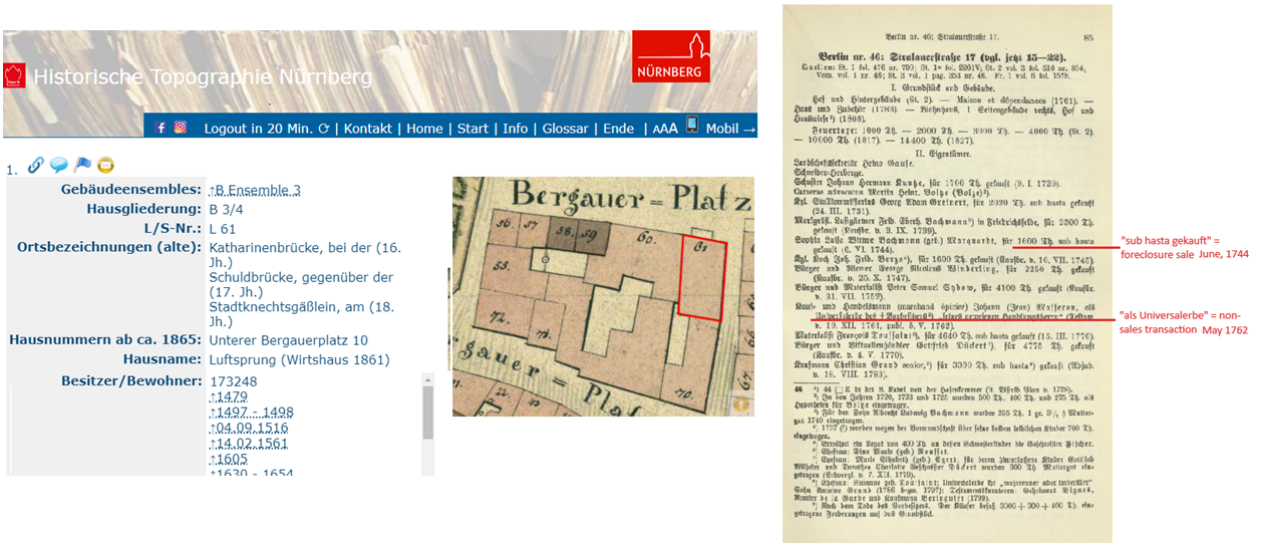
[4]See methodological details in Razum et al. (2023).

[5]Not least, non-sales transactions (inheritances, transfers, or property swaps). The Case-Shiller index for U.S. pre-depression home values uses self-reported survey data points via – these drawbacks have been critiqued in White (2009) and elsewhere.

[6]See further discussion and data in Appendix section 4 and Figure A.4.

Figure 1: Primary source exhibits: Nuremberg TOPO N project, *Unterer Bergauerplatz 10*; and Berlin "Häuserbuch", *Stalauerstrasse 17*.

[1] **[2]**



Notes: The picture displays example entries from the primary sources I use to construct the new German house price index. On the left handside ([1]), a sample entry from the Nuremberg *Häuserbuch* ("Topo N" Project), for the address of *Unterer Bergauerplatz No. 10*. Consistent cadaster units are traced here in their ownership history, sales transactions, sales prices, and other notable events associated with the house. In this case, the first ownership entry (category *Besitzer / Bewohner*) is recorded in the year 1479. Entry via: https://online-service2.nuernberg.de/stadtarchiv/zeig.FAU?sid=574BA1843dm=6ind=2ipos=4220332f34. On the right handside ([2]), a sample page from the Berlin *Häuserbuch* that covers property level transaction histories spanning the 16th-20th centuries, published via Lüdicke (1933). Here the property recorded is *Stalauerstrasse 17*. To give a sense of the information covered (beyond traditional cash sales), the red callouts highlight a foreclosure sale in June 1744 ("sub hasta"), and below details on the property being passed on as inheritence in May 1762.

3.1 Summary Statistics

The German residential house sample covers eleven cities, spanning a heterogenous geographical and demographic range in the country over time. Three cities are in Northern Germany, eight in Southern Germany, and the size of cities ranges from 1,300 (Stengberg, as of 1500) to 538,000 (Munich, as of 1905). All the key financial hubs of Germany/the Holy Roman Empire over the early modern and modern period are part of the sample, including Nuremberg, Frankfurt, Munich, and Berlin. While such ex post successful cities may feature some degree of "superstar city" bias, the inclusion of smaller and relatively stagnant towns balances the aggregate sample regarding these possible factors.[7]

In total, I rely on more than 1,400 early modern repeat sales at this point. The underlying properties are exclusively single or multi-family residential objects, typically in the "old" parts

[7]It has been argued that "Superstar cities" may exhibit faster house price growth, biasing indices upwards; ex post, my city sample includes both "successful" and "unsuccessful" cities with very stable or stagnating population growth (e.g. Stengberg).

of the respective cities. A significant share of the properties in the sample continues to exist as of 2024, based on a cursory survey – even though house numbers and other designations have repeatedly changed.[8] The houses in our sample have a median value of 594 Rheinflorin (Rfl., or *Rheingulden*) in 1465, a figure that has risen to 2,025 Rfl by the year 1650, 4,537 Rfl. by the year 1750, followed by 18,131 Rfl. by the year 1850, and finally 33,753 Rfl. at the eve of World War One (1910). The small town of Stengberg records the lowest median house prices over the sample period at 224 Rfl., while Berlin records the highest median house price, at 14,392 Rfl.

The 202 properties between them account for more than 44,000 property-years, with records starting as early as the 14th century for Nuremberg, Frankfurt, and Erfurt. The final property sale observation in our sample is for the year 1943, in Nuremberg. Of course, a sample of several hundred properties must still be considered "small" by the standards of the German housing stock, even in the early modern period. After all, more than 4,000 individual housing units existed in the early modern city of Nuremberg alone. Yet, our sample is designed to be representative of the German residential housing stock given its geographic composition – and is easily scalable based on the new sources identified.

Tables 1.1-1.2 summarize the new German primary data, by period and across variables.

3.2 A new repeat-sales house price index

Existing repeat sale indices typically use CPI data to adjust nominal house price series to real bases. Shiller (2015, chapter 3), for one, uses year-to-year U.S. CPI index changes to deflate respective nominal home price changes and create a U.S. "real home price index" over 1890-2014. Similarly, Knoll et al. (2017) deflate nominal house prices with a 5-year lagged realized CPI (that is, the average over t-5 to t-1, to deflate nominal price at t=0), to obtain real house price indices. I broadly follow the latter approach, using lagged inflation, and also test in the Appendix (section 4) alternative inflation data.

Such repeat-sales indices are regarded as "state of the art" indices, and not least they are widely considered the best methodology to control for quality improvements in homes, though alternative "hedonic indices" have been used recently for quality-adjusted U.S. series (Lyons et al., 2024). Yet, while they can capture many of the quality-related changes in housing as a wider asset class – for instance, they are by design well equipped to deal with the land component of the home price, as they keep it constant at the property level – even repeat-sales indices are far from *perfectly* stripping out pure quality-related improvements. We note here that we follow existing approaches in incorporating quality-related adjustments to the best extent possible, and dedicate an extensive subsection in the appendix discussing quality-related changes in our sample, together with better estimates on their effect on secular prices. For instance, the long-run suggests that the median German homeowner only possesses a 30-50% larger home in terms of living space, a

[8] House numbers in any case are an innovation of the 18th century. In the pre-1700 era of our sample, the properties are typically designated by their owner, or by lengthy descriptions of their location relative to prominent buildings in the city (the church, city hall, or wells).

Table 1.1: Summary Statistics for German primary data, repeat sales, 1/2

	# properties	# repeat-sales	Avg sales gain	Med sales gain
1465-1565	57	134	111.7	23.6
1565-1665	111	312	114.1	25.8
1665-1765	145	399	120.6	29.6
1765-1865	154	652	92.5	34.1
1865-1943	63	167	96.1	43.2
Total				
1465-1943	202	1,473	104.7	33.3

Notes: The table reports the sample composition of the repeat-sales observations for residential real estate in Germany. Avg and Med sales gain measure the nominal change in price at the time of sale over the respective most recent previous sales price of the property, expressed in cumulative percentage terms (properties arithmetically weighted).

fact underscoring that such quality improvements are without a doubt underlying any long-run dynamics in the asset class.

Figure 2 now displays our new long-run German house price index (RHPI), here in real terms (log), over the full period 1465-2020 (1465=100).

The index displays a sharp visible inflection during the Thirty Years War, when structural inflation rates surge, a significant demographic shock takes place, and much existing housing stock is physically destroyed.[9] Regardless of the specific inflation basis chosen, a key commonality of the various iterations of indices reconstructed is the sharp acceleration of the index values from the late 18th century. In real terms, on our benchmark basis, the index value rises by a factor of 2.2 in real terms until the terminal values in 2020.

This suggests that real house prices – at least in one key geography where such returns can be measured granularly over time – did *not* begin to accelerate only from the mid-20th century, as key literature proposes (Shiller, 2015; Knoll et al., 2017). Certainly, a long-run "hockey stick" pattern is difficult to maintain – though our data does suggest that there appears to have been a temporary correction from the early 1900s, perhaps even a correction that is only dominated by the Thirty Years War crash three hundred years earlier.

While our benchmark approach in this section – to value-weight individual transaction observations – is consistent with contemporary state of the art indices (notably the Case-Shiller index for the U.S.), key sub-periods and absolute index levels are certainly sensitive to particular sales weightings and also the specific inflation adjustment approach: though none is inconsistent with the broader conclusions reached here (especially regarding the pre-1945 real price dynamics), the Appendix

[9]The financial analysis of the Thirty Years' War is still underdeveloped. A discussion of the initial 1618-23 economic shock that is showing up in our new index exists via Kindleberger (1991). A thorough attempt to measure property and demographic losses during the war exists via von Hippel (1978): his figures suggest demographic and physical capital losses exceeding either of the 20th century World Wars.

Table 1.2: Summary Statistics for German primary data, repeat sales, 2/2

	gross nom TR	gross real TR	rental yd	real cap gain	hp
1465-1565	7.19	6.10	7.29	-0.35	23.8
1565-1665	5.71	3.48	4.73	-1.28	33.2
1665-1765	5.96	5.52	5.80	0.53	42.6
1765-1865	7.84	7.12	7.54	1.26	29.5
1865-1943	7.79	6.17	6.76	1.45	22.2
1943-2020	10.52	7.80	6.76	1.45	–
Total					
1465-2020	7.37	5.92	5.51	0.42	–

Notes: The table reports performance measures for the repeat-sales observations for residential real estate in Germany. All variables except "hp" are reported in percentage terms per annum. "hp" denotes the "holding period" (properties arithmetically weighted) and is reported in years.

reports a number of variations of the German benchmark index presented here.[10]

3.3 Renting and owner-occupation

According to Wenderoth (2025), even lower middle and lower income households typically lived in their own single family home by the late 15th century – with around 30% of taxable citizens in German urban areas living in rented houses or apartments. This figure rises to around 60% by the year 1800, and the development is relatively even across German cities.[11]

Overall, therefore, rental income constituted a clear investment motivation for buyers of real estate assets over the entire observation horizon, and should be an integral part of total return figures, and to assess actual income risk.

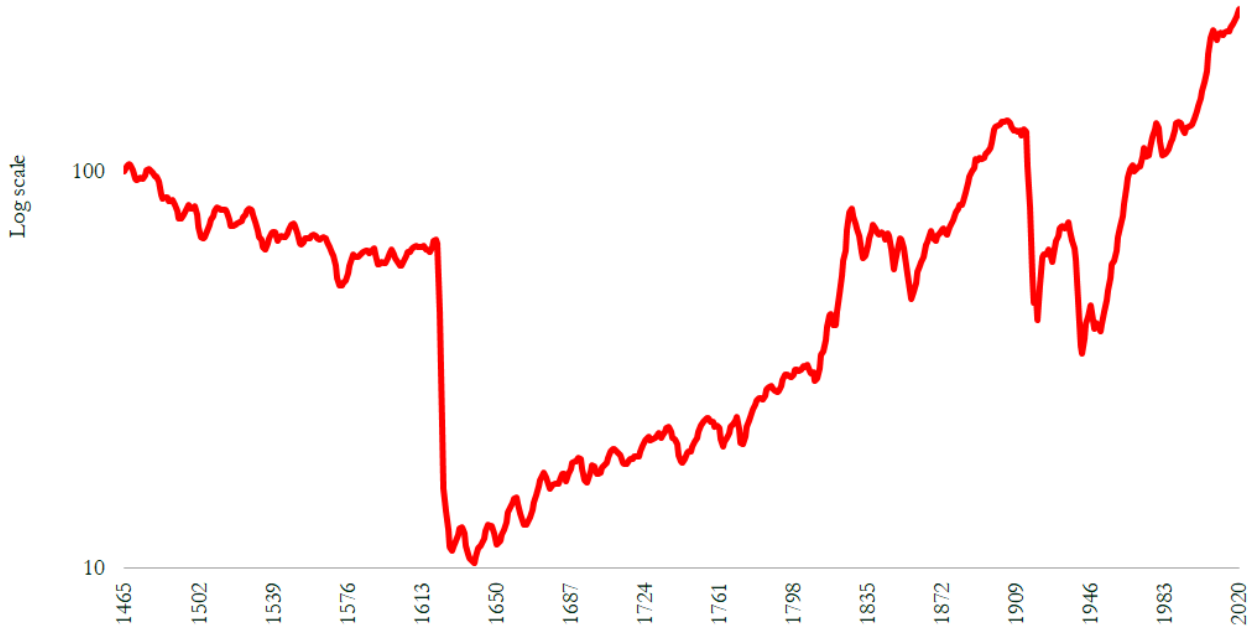
As has been noted in related papers, historical rental data in primary sources is generally more elusive, hampered in part by the fact that rental contracts were privately negotiated and did not have to be registered with notaries or municipal authorities (unlike sales), in part by a convoluted nomenclature in early modern contracts.[12] All this complicates the construction of price and rental series featuring the *identical* housing properties: key recent literature has navigated this challenge in different ways, with some constructing separate housing and rental series (featuring different properties in each of the two samples) which are then merged to construct a rental yield

[10]Recall that multiple iterations of particular weightings and inflation adjustments are to be found in the Appendix, particularly sections 4 and 6.

[11]The Wenderoth data supersedes similar but patchier data in Dirlmeier (1978), I thank the author for sharing his data, which focuses on Southern German cities and which also suggests that generally the share of renters rises with city size in the early modern period.

[12]Housing contracts – besides mentioning the assessment or sales price – often refer to *Guelten*, or *Zins* ("interest") when denoting rental payments, with such terminology adding to confusion.

Figure 2: Fused German real house price index (RHPI, log), 1465-2020.



Notes: The Figure displays a German real house price index running over 1465-2020 (1465=100, log scale), value-weighting German cities. Here we splice the new data introduced in this paper for 1465-1910 (using the "Schmelzing" inflation adjustment basis), with the existing German house price data in Knoll et al. (2017) over 1911-2020 (as updated in Jorda et al. (2019), and denoted "JST"). The gaps that JST record for selected German interwar years are filled by taking the arithmetic average of the real house price change for Denmark, France, Switzerland, and Belgium.

series (e.g. Lyons et al. (2024)). However, our *Häuserbücher* contain a critical amount of detail, at least for multiple key cities, which allows a direct linking of rental and price series – and several hundreds of additional rental-only data points have been identified in related primary and secondary sources – for instance new series in Wenderoth (2025) based on municipal rent contracts – which allow the replication of existing approaches that construct a separate rental index series that are subsequently merged with price series to derive rental yields (all rental observations are summarized in Appendix table A.2.2).

3.4 Costs and net total returns

Thus far, we have not comprehensively measured the cost component of residential real estate for our long-run sample. The existing literature mainly works with general assumptions, with nuanced unit-level data being scarce. Generally, relevant costs include taxes on the municipal and state level, maintenance costs, depreciation, and costs resulting from vacancy. For Germany, the level of detail in the new sources allows dynamic, though not unit level, estimates which we will

incorporate into a construction for net returns.[13]

3.4.1 Taxes

Beginning with the tax side, the information is relatively comprehensive. Some of the German *Häuserbücher* contain occasional unit-level tax information, most of which concur with our existing knowledge of aggregate city level rates.[14] For instance, among larger cities in our sample, we know that for Munich, an annual imposition of one Penny per Pound is levied from the mid-15th century, where it remains for centuries – a rate equivalent to 0.41% of the (self-assessed) value of the house. The city of Frankfurt charged 5/6 of one percent as real estate sales tax from the year 1613 onwards; the city of Augsburg charged 1/4 of one percent of the value of real estate and land in wealth taxes; the city of Leipzig charged 1 basis point of the (self-assessed) value of housing in irregular wealth tax impositions from 1481 onwards. It is to be stressed that these impositions occurred in irregular intervals – often connected to specific external threats for which cash was quickly needed and liquidity was otherwise constrained.[15]

Afterwards, property tax rates are rather stable from the 15th century, with evidence of a moderate upwards trend. For Berlin, Voigt (1901) reports tax rates of 0.3% of the property value in the 17th and 18th centuries, which rise to 1.25% by the first half of the 19th century. Afterwards, Müller (1881) provides annual rate of 1.9% of the market value for rented residential housing in Berlin during the 1870s and 1880s. With much corresponding data for other cities in our sample in broad agreement, Eberstadt (1903) provides rates for representative Hessian cities, over 1890-1901, arriving at a lower 0.6% of property value for rental properties. Against this backdrop, in the subsequent net return calculations, I operate with a constant 0.15% annual tax rate relative to the market house price prior to 1800, and increase this figure to 0.5% from 1800. Slight variations within the range of property rates provided in the literature, or the adoption of a more dynamic rate over time, do not alter any of our basic findings.[16]

3.4.2 Non-tax costs

What are the resulting time series estimates for long-run total gross costs, as a percentage of value, over our German sample?

[13] Most subsequent discussions will continue to focus on gross housing total returns (nominal and real), but the appendix shows that all main results discussed in the main paper (structural breaks, stationarity) hold for net bases. See also the visualization via Appendix Figure A.1 and further discussion there.

[14] For instance, in Nuremberg in 1500, we learn via TOPO N, B 6/5 that the house worth 190 Rfl was taxed at an annual 0.5 Rfl, that is, at 0.3%.

[15] For a systematic overview of tax rates, see Isenmann (2012). The Munich rates are via Solleder (1938, 208ff.). Further city level details in Mayr (1931) for Augsburg; Frankfurt data via Braeuer (1915).

[16] Wenderoth (2025, 91) documents occasional impositions of taxes on rental income: for instance, Hamburg from the year 1627 imposed a 6.25% levy on rental income. Since these taxes are not documented more widely, I ignore them for the net yield calculations.

It would be methodologically ideal of course to match property level price, vacancy, and rent data with matching property level expenditures – however, such information is sporadic. For some cities in our sample, interesting internal cadaster calculations exist. For instance, the Berlin cadaster authorities during the 1870s apparently calculated with a 2% total gross expenditure figure relative to the current sales value of a rental residential property for total costs, including maintenance, labor, and fire insurance, but not adjusting for vacancy rates (Müller, 1881). A related official publication, also for Berlin, estimates an average vacancy rate for residential housing over the period 1880-1910 of 3% (Ascher, 1917). While providing some general reality checks, however, in line with existing literature, we will have to abstract from property and matched city level data and assume plausible aggregate figures. We will also have to abstract from some relevant city level provisions that are relevant on the margin, but cannot be generalized: for instance, Voigt (1901) reports that for Berlin in the early modern age, renovation costs are partially tax-exempt (up to 15% of total renovation costs) – lacking more specifics about the scope and duration of such provisions, we will disregard such details.

Given the less robust data situation across countries, recent long-run total return calculations thus have assumed time series based on general assumptions, rather than asset-level data. An example is Eichholtz et al. (2021), who estimate non-tax costs of 32% of rental yields as expenditure basis for Paris and Amsterdam, closely in line with the estimates in Chambers et al. (2021), who have actual property level data on expenditures for their U.K. sample. These estimates include maintenance, vacancy, insurance, and labor costs.

The Appendix (Table A.3) systematically compiles all expenditure share estimates in existing literature (and component breakdowns) that form the basis for these net return adjustments. Proceeding from here, I will follow the majority of finance literature and mainly concentrate on the gross return figures – though I stress that all main results hold for the net return constructions, and results for German net return variations are reported alongside in all subsequent tables.[17]

4 Reconstructing Global House Prices – A Machine-Learning Approach

In a vibrant new literature, it has recently been argued that nonparametric machine learning (ML) approaches can clearly outperform traditional methods in empirical asset pricing applications: specifically, such "supervised learning" models that assume no functional form of the underlying training data have been applied to a broader range of financial forecasting debates, including U.S. equity prices (Gu et al., 2020), bond risk premia (Bianchi et al., 2021), as well as broader corporate finance and portfolio choice (Duarte et al., 2024).[18]

[17]See further Appendix section 2 and Table A.3. Recall, via Figure A.1, that we observe that the general secular "U-shape" of German returns is maintained – what would upset such an identity is a systematic rise in expenditure shares over time, for which there is at this point limited evidence, however.

[18]For relevant surveys of finance and economics applications and their foundations, including the subset of "supervised learning" approaches used here, see Masini et al. (2023) and Kelly and Xiu (2023).

While the majority of financial applications are currently concerned with optimizing the prediction of future returns on the basis of a limited set of past realized data points, the financial historian faces a simpler problem: estimating *past realized returns (prices)* on the basis of a limited set of past realized known covariates (which are strictly already available for end of the t period). In other words, we are operating in an easier environment because instead of a look-ahead forecast, we are interested in a backcast target variable reconstruction – an environment where important covariates of the target variable for the given period are already known. In this sense, the approach can also be thought of as a simple computational optimization of a nonparametric regression exercise – with the advantage that the supervised ML environment can optimize the nonparametric environment for a large set of covariates, and determine the ideal statistical structure. However, it is important to note that this means the model setup as used here does not inform about actual *forward-looking* (that is, beyond the cutoff of the training data set, this being the year 2020) predictive conditions.[19]

Importantly, while advanced machine learning models have already been prominently used for similar time series reconstructions in related fields – notably long-run climate trend and demographic reconstructions – a relevant application to finance has been fully absent, even though very recent published and unpublished studies demonstrate the potential for key economic time series applications – including advancements over the classic long-run Maddison per capita GDP series, as well as benchmark wage series.[20]

While each methodology faces its own set of challenges, machine learning approaches could offer a powerful cross-check to, and extensions of, our primary source housing reconstruction approach that concentrated on Germany (as well as cross-checks to previous early modern price reconstructions in the literature, per the appendix). The situation is particularly promising given the new availability – or available reconstructions – of large sets of (thus far unexploited) known covariates for key advanced economies over centuries.

4.1 Training the model on 1845-2020 housing data

Specifically, we train several nonparametric machine-learning models on both modern nominal and real housing capital gains as their target variables. A plethora of available models exists, of course all within a highly dynamic field – our focus will be on a subset of latest generation models optimized for multi horizon tabular fore- and backcasting applications.

Among the models we test are a Time-series Dense Encoder model (TiDE) (Das et al., 2024), a class of long-horizon forecasting models considered to be computationally more efficient than others, but more simplistic; a Seq2Seq+ model (Keneshloo et al., 2020), which utilizes deep reinforcement learning approaches; and a Temporal Fusion Transformer (TFT) model (Lim et al., 2021); both the Seq2Seq+ and TFT models incorporate Long Short-Term Memory (LSTM), a deep learning feature

[19]For this approach, all forward-looking covariate data would not yet be available at the time of prediction, time t.

[20]Here see Koch et al. (2024) for backcasting per capita GDP data, and Paker et al. (2025) for an unpublished demonstration for U.K. wage data; somewhat related are Bochow et al. (2025) and Borup et al. (2023).

shown to be particularly potent at asset pricing applications such as predicting bond risk premia, for instance (Cong et al., 2021; Chen et al., 2024). A fourth model tested is a Random Forest (RF) model, tuned with conventional parameters – a type of ML approach that has been shown to perform impressively in forward-looking financial prediction environments (Gu et al., 2020).

The advantage of these models is not just their optimization for long-horizon fore- and backcasting exercises – but their strictly nonparametric foundation: all models assume no particular functional form of the data, and batch predictions resulting from such training therefore contain no particular biases on the econometric features of time series properties that we are interested in.[21]

The first step consists of the training stage, to train the machine-learning model, with annual residential house price change (in both nominal and real terms) as its target variable. We train the model on housing capital gains on the basis of the most recent modern and early modern research datasets where a sufficient sample of additional covariates exists. Training and validation splits are undertaken in line with related previous finance literature: specifically, I use a 80:10:10 training:validation:test split.[22] The r-squared reported refers to the performance of the training data set on the validation and test data set – hence is an out-of-sample (OOS) measure. By design, the r-squared for backcasting environments will be higher than in classic forward-looking prediction environments, and ML approaches achieve elevated OOS more generally, as emphasized by Mullainathan and Spiess (2017) and others.

The covariates importance across our specific ML models is detailed via Figure 3. Specifically, the thirteen covariates include the government deficit-ratio using Dincecco (2013), a financial crisis dummy using Metrick and Schmelzing (2025), population growth, 3-year lagged population growth, nominal interest rates, change in nominal interest rates, real interest rates, inflation, 7-year lagged inflation, change in real per capita income, a FX peg dummy, and national construction cost indices spanning 1871-2020 that are sourced from secondary literature and/or respective statistical agencies (detailed in Appendix section 4.2).

The source of the training dataset target values (year-on-year housing capital gains, 1871-2020) is primarily the widely used "JST dataset" (Jorda et al., 2017) and sources therein – with the exception of U.S. data for which we utilize the superseding data in Lyons et al. (2024).[23] For nominal housing capital gains, the annual values are directly observable in "JST". To obtain observations for real year-on-year housing capital gains, we adjust the JST nominal capital gains with year-on-year CPI inflation rates in the same JST data set.[24]

[21]Where such LSTM models struggle are environments with more meaningful data gaps in the training set – our environment uses the (comprehensive) JST environment where such data gaps are virtually absent (even though, as we argue for the case of Germany, the estimation procedures for given periods can be challenged).

[22]For instance, Bianchi et al. (2021) use a training:validation/test dataset split of 85:15. The backcasting p.c. GDP model in Koch et al. (2024) uses 80:20.

[23]One further addition is the inclusion of UK data over 1845-1870 on the basis of the Millennium Dataset (Dimsdale and Thomas, 2016), which is also reported in conceptually consistent annual residential house price change.

[24]Per above, lagged inflation rate adjustments are used in Knoll et al. (2017), though various other approaches use year-on-year adjustments.

Table 2: Machine-learning model performance. Target: DM annual house price change

	$r^2(OOS)$	MAE	forecast horizon	context window
Target: nominal y-o-y price change				
Seq2Seq+, 110-36	**.544**	**.032**	36	110
Temporal Fusion Transformer (TFT)	.464	.033	36	115
Time Series Dense Encoder (TiDE)	.184	.039	32	110
Random Forest (RF)	.348	.050	20	80
Target: real y-o-y price change				
Seq2Seq+, RHP, 115-36	.572	.035	36	115
Temporal Fusion Transformer (TFT)	.428	.040	36	115
Time Series Dense Encoder (TiDE)	.169	.040	36	115
Random Forest (RF)	-.125	.059	20	80

Notes: The table reports training results from several machine-learning (ML) approaches, all with annual nominal (real) residential real estate price change as their target variable. The training data set consists of house price changes in addition to thirteen co-variant variables, annually over 1870-2020, primarily using "JST" (Jorda et al., 2017), except for Germany (excluded up to 1962 to prevent leakage) and the U.S. (where Lyons et al. (2024) is used). The training:validation:test split of the data in all models follows typical tunings in literature, and except for the Random Forest (RF) model is 80:10:10. For the Random Forest model, the forecast horizon / context window variables refer to the training:test data set split of 80:20.

Using this set of time series for training, I report the model training performance in Table 2. The first three rows report nominal capital gain performance. Here we observe that a *Seq2seq+* model with specific tuning parameters performs best and achieves an OOS r-squared over the full training dataset of 0.544 when optimized for a 36-year forecast horizon, and a 115-year context window. The following three rows report results for real capital gains as the target variable. Here, we obtain an even better OOS r-squared of 0.572 when using *Seq2Seq+*, though it is weaker in terms of MAE. While I consider the nominal capital gains basis the benchmark basis in the main part of the paper, reconstructions for the real capital gains basis are reported in Appendix section 4.2, on both GW and AW global weightings.

It is important to note in this context that such r-squared (out of sample sample) may be expected to be even higher – however, we have to strictly limit ourselves to the sample of covariates that are available to a sufficient degree for the early modern period as well. The TiDe and the TFT models record slightly lower r-squared, up to 0.464 for nominal changes – with all models achieving a relatively tight range of MAE performance between .032 and .05.

4.2 Variable importance

Via Figure 3, we provide the variable importance, by model, for the machine-learning methodologies. This exercise is analogous to determining variable importance in a classic regression analysis. The importance refers to the training dataset that utilizes 1870-2020 data, across the

four machine-learning models – Seq2Seq+, TFT, RF, and TiDE. The former (Seq2Seq+, nominal cap gains, bolded in Table 2) constitutes the "baseline" machine-learning model with the highest realized r-squared (for this the variable importance is called out in the Figure).

We generally note that variable importance can vary substantially across the four models, but that a handful of variables retain high important across approaches – including the nominal (long-maturity) sovereign interest rate, building costs, and population growth rates (lagged and unlagged). The finding that building costs are among the most important variable across all models (in fact constituting the top variable in three out of the five models) is particularly interesting given the notion in recent literature that they only play a minor role in explaining price dynamics (Knoll et al., 2017): on this, I compare approaches in more detail in the Appendix, and the greater relevance of building costs does not appear to be a result of my particular weighting approaches.[25] Beyond building costs, however, the identification of population growth and interest rates squares well with existing notions: finally, perhaps notable is the relevance of the deficit ratio and financial crises for our benchmark model, since they are not regularly considered in traditional regression settings.

4.3 New global RHPI series

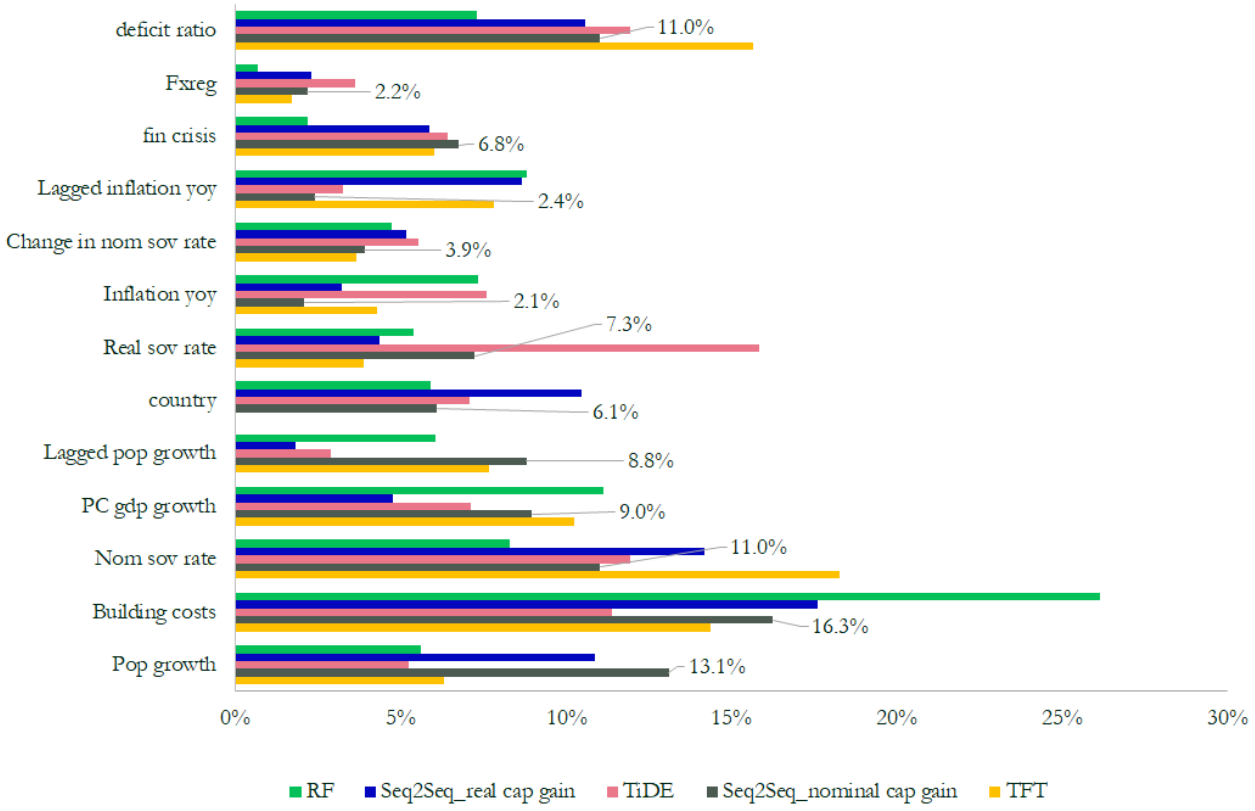
In the next step, I use the best-performing training model to generate a predictive backcasting house price growth series spanning 1502-1870: for Germany, the backcasted series runs to 1962, for the U.S. it runs from 1786 to 1889, and for the U.K. it runs from 1560 to 1844. This best-performing model for both nominal and real capital gains is the Seq2Seq+ model with a forecasting horizon of 36 years, and context windows of 110 and 115 years. TFT models perform only marginally worse across metrics, with the TiDE models being meaningfully lower than either two alternatives.

In Figure 4, I display the benchmark Seq2Seq+ machine-learning reconstruction for nominal house price changes – with predicted nominal values then being adjusted by existing lagged inflation data to obtain real values. This is done for five key advanced economies since the early 16th century. All are indexed to 1800=100, and utilize the full set of sixteen covariates for the 1502-1869 period, to obtain the target value predictions. For these pre-1869 variables and data points I mainly rely on published GDP and population time series in the Maddison database (Bolt and van Zanden, 2020), as well as financial time series in Dincecco (2013), Rogoff et al. (2024), Schmelzing (2026), and Metrick and Schmelzing (2025). A key exception is that I hand collect and/or reconstruct building cost indices for the pre-1870 era across the five key economies, which represents a variable with high variable importance for the post-1870 era across models (Figure 3). The Appendix also displays "global" reconstructions on the other model bases, including for real housing capital gains as the direct target variable: all key structural results emphasized here hold for these alternative variations, though they at times differ on absolute house price levels.[26]

[25] See Appendix section 4.2. Knoll et al. (2017) prefer to focus on "output cost" indices that also include profit margins, while I mainly focus on indices that capture materials costs and labor costs, but not profit margins.

[26] This includes the second best-performing model, the Seq2Seq+ real capital gains basis, which is exhibiting consistent

Figure 3: Variable importance, by ML model.



Notes: Variable importance for the top-13 most influential variables in each model. Variable importance within each model is normalized to sum to one, excluding "time". Called out with importance data labels is our benchmark Seq2Seq+ model targeting nominal capital gains, in grey bars (compare Table 2).

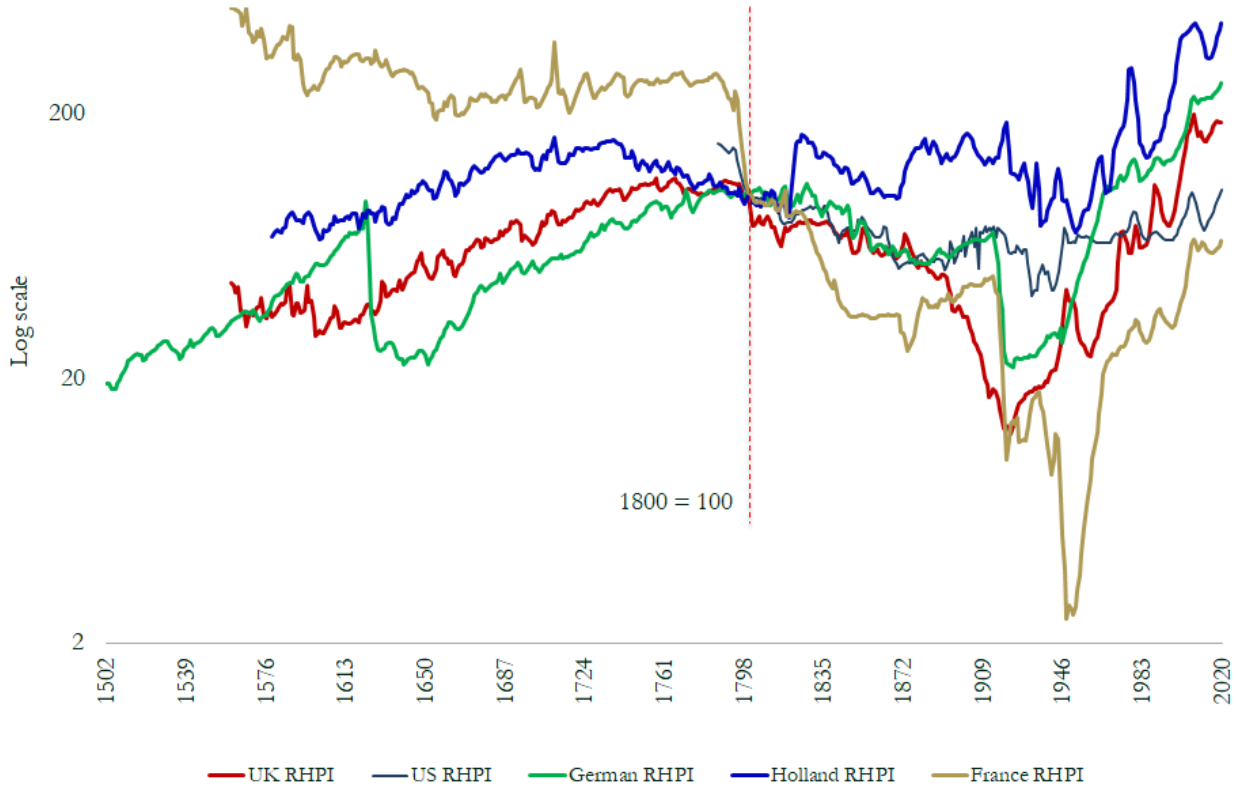
What secular patterns do we observe on the country-level? First, the period of circa 1500-1650 is marked by moderate but consistent declines in real house prices, interrupted by a sharp volatility around the outbreak of the Thirty Years' War. After the conclusion of hostilities, house prices in the Holy Roman Empire rebound secularly, a turnaround also visible in French and Dutch data, with both of them being traditionally considered key beneficiaries of the Treaty of Westphalia (1648). France's own secular appreciation cycle comes to a dramatic reversal with the French Revolution. For most countries, real house prices through the 19th century appear in fact to remain below pre-Revolution peaks. It is during the first quarter of the 20th century that French, German, U.S., and British real house prices bottom out, with the Netherlands being a key exception.

Generally, therefore, country-level inflections appear to reflect major political-financial breaks well documented in related literature – with a variety of anecdotal evidence in support of the housing-specific inflections[27] – and subsequent econometric tests (section 6 below) being able to

general contours, including the major inflection point in the late 18th century, see discussion and Figure in Appendix section 4.1.

[27]For instance Ambrose et al. (2013), who documented a "market implosion" in Amsterdam real house prices between 1781 and 1814. Our new Netherlands RHPI peaks in 1782 and bottoms in 1819. Note that the Ambrose et al. (2013) data (like all other pre-1870) data is not part of the ML training data set, so our Dutch RHPI represents an independent

Figure 4: Seq2Seq+ Machine-Learning real house price index (LHS), 5 countries, 1800=100, 1502-2020.



Notes: The Figure displays real house price indices based on the Seq2Seq machine-learning model, for five advanced economies, all based to 1800=100 (log scale). U.S. data is ML-backcasted over 1786-1889, U.K. data over 1560-1844, German data over 1502-1962, French data over 1503-1869, and Holland data over 1578-1869: we then switch to the respective existing price series.

validate such initial narrative evidence.

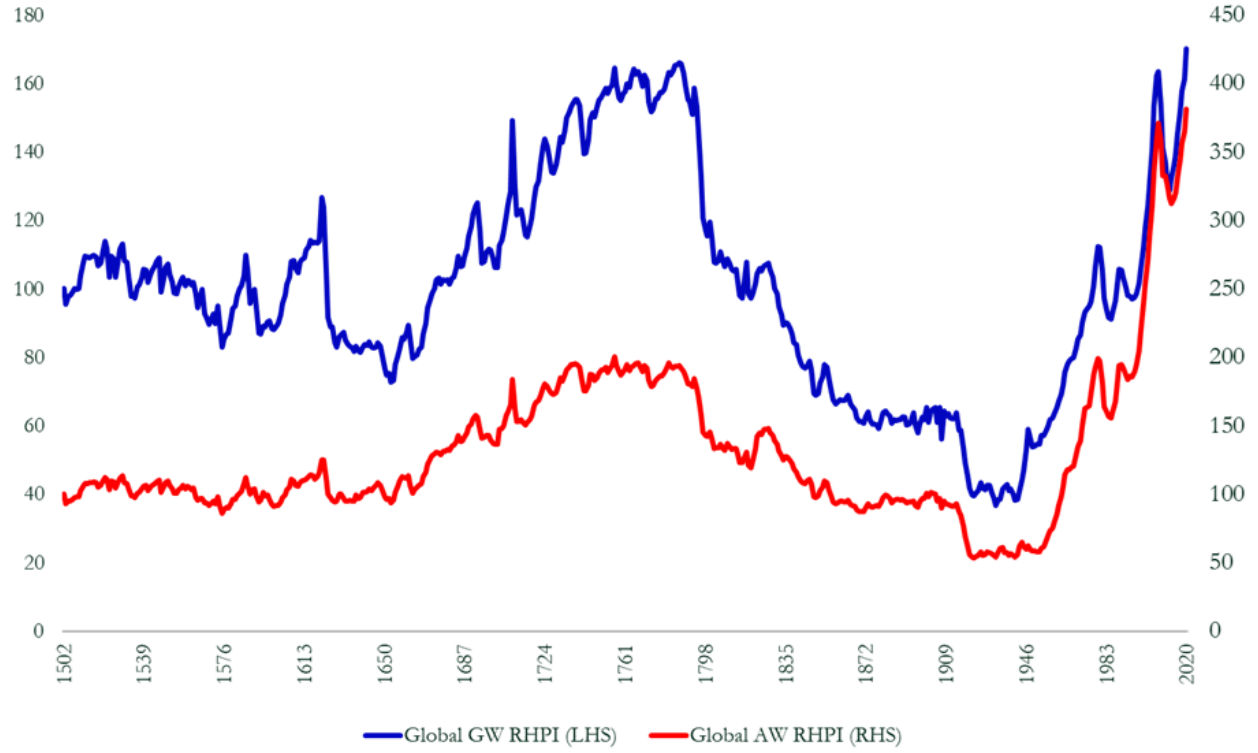
Next, Figure 5 allows a holistic analysis of the "long cycles" in global house prices over centuries, using aggregation methods recently undertaken for related financial variables[28]. A GDP-weighted aggregation (GW) and an arithmetically weighted (AW) aggregation of country-level series are presented here, covering a substantial share of "advanced economy" GDP over centuries. Particularly on the GDP-weighted basis, the contextualization of "recent", 20th century performance of housing markets, is highly revealing.

Most notably, perhaps, our reconstructions suggest that on a GDP-weighted basis only most recently, as of the early 2000s, did global real house prices surpass all-time peaks previously reached in the immediate years prior to the French Revolution. All-time lows, on the other hand, were reached on both aggregation bases in the interwar period of the 20th century, with the

validation.

[28]See Rogoff et al. (2024), where an identical weighting approach is presented for global interest rate time series, on the basis of annual-level aggregate GDP series.

Figure 5: Global Real House Price Index (RHPI), Seq2Seq+ bases, 1502=100.



Notes: The Figure displays Global Real House Price Indices, on the GDP-weighted (GW, LHS) bases and the arithmetically-weighted (AW, RHS), indexed to 1502=100. The sample spans France, Germany, U.K., U.S., and the Netherlands per Figure 4. Post-1870 values (post-1962 for Germany) are based on "JST" data points – and on Lyons et al. (2024) for U.S. data points over 1890-2006.

two world wars dragging real house prices down to all earlier local bear markets, including the post-Thirty Years War bottom. This latter observation also holds for arithmetically weighted (AW) real house prices – though here we note that the surpassing of all-time peaks in real house prices occurs earlier, specifically in the late 1990s.[29] The Appendix contains robustness exercises for these global ML indices, including the full reconstructions for using real capital gains as the target variable – on this basis, too, only most recently new all-time real price peaks were reached.[30]

[29]On the AW basis, a local peak is reached in the year 1756 (index level 200.4), with this level being surpassed in 1998 (index level 204.7). On the GW basis, the local peak occurs in 1786, at the eve of the French Revolution (index level 165.9), and is first surpassed again in 2020 (index level 170.1).

[30]See Appendix section 4 and visualized via Figure A.6: per this Figure, only on the real capital gain target GW basis, 2020 price levels remain below those during the late 1700s.

5 Nominal and real total ex post returns, and their decomposition, Germany and "Global"

5.1 Decomposition: rental yields, real estate capital gains, and the cost of capital

Recent debates have made much of the precise decomposition of residential real estate total returns: while Knoll et al. (2017) have suggested a decisive role for rental returns – one that plays the main role in accounting for superior returns of housing over equities – Chambers et al. (2021) have questioned this decomposition and posited much lower rental yields for residential estate over time.[31] Our new data may illuminate therefore the precise contribution of different return elements, and assess whether recent declines in rental yields – and corresponding large housing capital gains – represent historical outliers, or should rather be expected to continue.

Indeed, via Table 3 that decomposes returns on the Global bases, a key observation stands out: historically, the majority of residential real estate total returns was accounted for by rental yields. Indeed, over extended periods of time in the early modern era, real capital gains are *negative* for residential housing, most severely over the subperiod 1760-1860. Rental yields, in comparison, appear to be far less volatile – over the course of centuries fluctuating in a relatively tight range of 5.1 – 6.9%. These basic observations are echoed in the primary source derived data on Germany, into which I will delve somewhat deeper next.

5.2 Total ex post nominal returns

The separate consideration of house price changes, rental yields, and expenditure components now allows the fusion into a total gross return series for German residential real estate since 1466.

We follow the total return definition in Jorda et al. (2019), specifically defining housing total return R as:

$$R_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}} + Y_{i,t} \tag{1}$$

with P denoting the nominal price change year-on-year for housing index in country i, at time t, and Y representing the rental yield.

What are the main takeaways from aggregates and the decomposition? We observe that the series over four centuries displays a "U-shape", initially at levels around 12%, but declining trendwise, and reaching all-time lows around the mid- and late 1500s, at levels just over 4%. From then onwards, total gross returns recover secularly, and return to double-digit levels by the mid-1700s. From then onwards, returns appear quite stable, trending around 10% for the remainder of the sample, despite occasional sharp spikes.

[31] Specifically, Chambers et al. (2021) find rental returns that are just three-quarters the Knoll et al. (2017) estimate (3.0% p.a. vs 4.0% p.a.), but they also revise downwards real capital gains contributions for U.K. data over 1901-1983 (-0.7% p.a. vs 0.7% p.a.); rental growth rates are close to zero.

5.3 Total ex post real returns, and excess returns (EX)

Inflation has evolved dynamically in the early modern and modern periods, and therefore real total returns might represent a more accurate measure of actual asset performance. We once more follow Jorda et al. (2019), by defining housing real total returns r. Let π denote the seven-year lagged CPI price change as measured by Rogoff et al. (2024) for the respective country. Then r:

$$r_{i,t} = \frac{1 + R_{i,t}}{1 + \pi_{i,t}} - 1 \tag{2}$$

Given multi-year holding periods in residential real estate, the lagged indices appear more suitable than year-on-year inflation adjustments, and also makes results more consistent with existing long-run real asset price series including sovereign yields.[32] For our German primary source series, we can use the Schmelzing (2026) inflation basis – and when that is done, this results in a real total return figure over the entire period of 5.9% p.a., of which 5.5% are rental yield returns, and just 0.4% p.a. result from capital gains.[33]

These German return patterns are mirrored on the ML-generated global bases. Table 3 further summarizes these measures of global returns, and their decomposition, by period. The availability of sovereign asset data allows us to next turn to housing "excess returns" – a key measure of interest in the literature, not least in the Case-Shiller canon (Case and Shiller, 1990). Recent debates have focused predominantly on *levels* of excess returns, rather than secular trends, and we will attempt to shed more light on the latter aspect in particular.

My excess return definition follows previous literature – where it is standard to calculate housing returns over long-maturity bond yields including (Campbell et al., 2009; Eichholtz et al., 2021) – specifically non-leveraged gross excess returns EX for year t are:

$$EX_{i,t} = R_{i,t} - i_{i,t} \tag{3}$$

With i representing the nominal long-maturity interest rate for the respective country (global), sourced per Rogoff et al. (2024).[34]

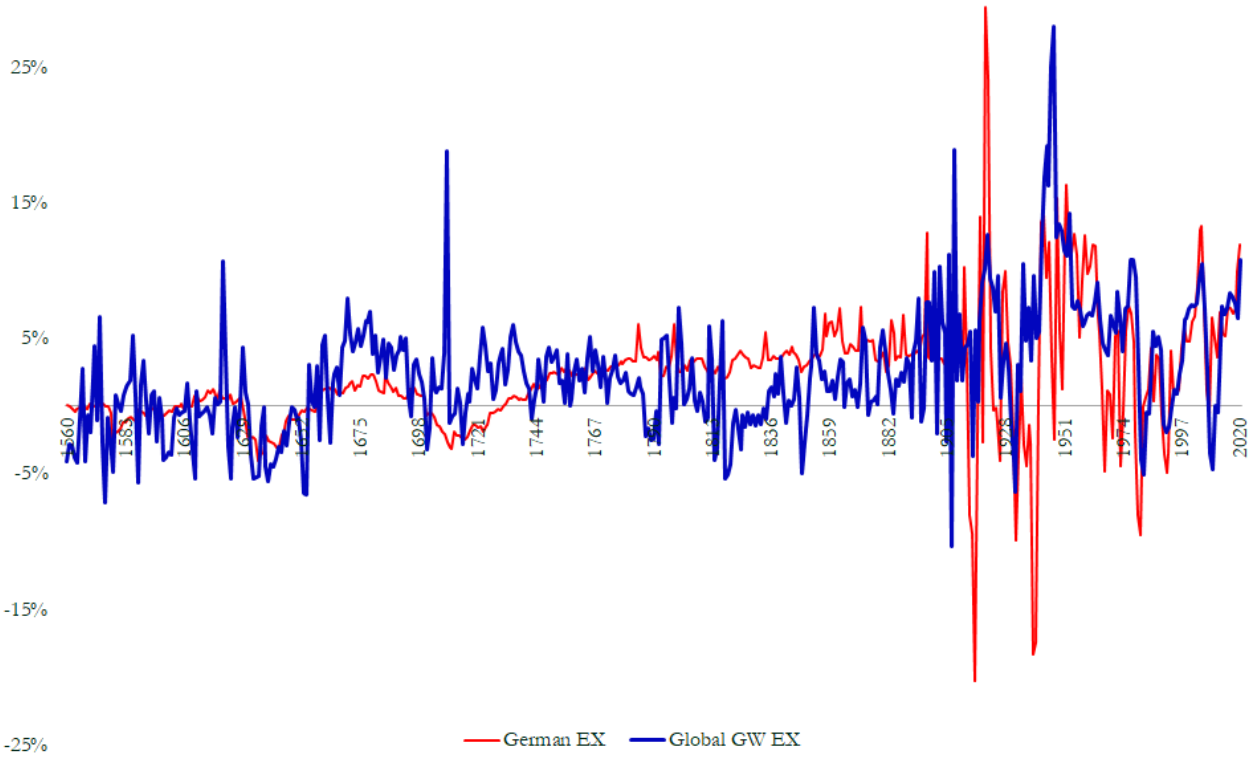
On this basis, Figure 6 displays global residential real estate excess returns (gross) for the German and the Global GW sample, spanning 1560-2020. It is visually highly suggestive that the series are secularly upwards trending, while it is also suggesting that volatility has increased on both bases. While increasingly positive on average, excess returns were frequently negative: specifically, 29% of country years in the German sample, and 30.1% of the Global GW sample are negative excess return years.

[32] For instance, Schmelzing (2026) and Rogoff et al. (2024) use the same seven-year lagged adjustments.

[33] Appendix Figure A.1 also displays the *net* total return series for Germany, spanning 1466-1866, incorporating our cost and tax estimates (as per section 3.4). It should be noted that this net real return figure does not yet include depreciation and vacancy cost estimates.

[34] Note that in Eichholtz et al. (2021), the authors calculate excess returns on a net housing return basis, but subtract

Figure 6: Excess Returns (EX), Residential Real Estate, German and Global GW basis, 1560-2020.



Notes: The Figure displays the German and Global GW excess return series using the residential real estate gross total return definition per formula (2), in the latter case weighting the five constituent countries by their rolling GDP-share ("GW"). From gross ex post residential real estate returns, the sovereign German and Global GW bond rates respectively (also pre-tax) are subtracted using identical country weights, with bond rates via Rogoff et al. (2024).

5.4 Summary and secular trends

What these new time series confirm across different methodologies is that housing generated lower but meaningfully positive average real returns even in the early modern period – despite the fact that negative (excess) returns are also frequent. There is little evidence of "stagnant" early modern housing markets, on either price or total (excess) return bases – in other words the "left tail" of the proposed "hockey stick" is anything but flat. In most recent decades, the return composition has undergone dramatic shifts, with capital gains now accounting for an unprecedented share of total returns in advanced economies, even outstripping the period between 1660-1760, which was characterized by the recovery and re-building efforts after the destructive Thirty Years' War on the European Continent.

What is also evident is that the volatility of housing returns increased meaningfully by the time of the 20th century. We do not investigate volatility and sharpe ratios in more detail here, but the finding of volatility rising in housing of the 20th century is generally in line with existing notions,

gross nominal long-maturity yields. However, bonds were taxed in the early modern period, which might distort such a calculation somewhat. Our gross housing return basis for the excess return calculation is consistently pre-tax.

Table 3: Summary Statistics for Global data, GW (AW) basis, 1560-2020.

	nom TR	real TR	real capital gains	rental yd
1560-1660	7.3 (10.0)	5.9 (8.3)	-0.1 (-0.1)	6.0 (8.2)
1660-1760	8.0 (7.7)	7.6 (7.3)	0.7 (0.6)	6.9 (6.7)
1760-1860	6.3 (6.5)	4.7 (5.2)	-0.7 (-0.7)	5.4 (5.9)
1860-1960	9.5 (9.4)	6.9 (6.6)	0.1 (0.0)	6.8 (6.6)
1960-2020	10.6 (11.2)	6.8 (7.5)	1.7 (2.7)	5.1 (4.8)
Total				
1560-2020	8.1 (8.8)	6.3 (6.9)	0.2 (0.3)	6.1 (6.6)

Notes: The table reports the sample composition of the ML-derived observations for global residential real estate. All figures for GW weighting basis and in per annum terms, with figures in brackets reporting the alternative AW weighting basis. "TR" = total (ex post) returns. "rental yd" = rental yield.

e.g. Ambrose et al. (2013).

6 Econometrics: regressions and structural breaks

We turn to econometrics, an investigation of "turning points" in housing markets over time, and the question whether our visual impression of a long-run upwards trend in housing total returns can be confirmed through formal tests.

We recall that in their seminal paper, Case and Shiller (1989) analyzed U.S. city-level real house price index changes and excess returns, and rejected random walks in both variables. In representative fashion, the authors use regression analyses on post-1945 U.S. residential real estate data, a horizon virtually all subsequent econometric work focused on, and typically using regression analyses (for a summary see Ghysels et al. (2013)). Basic regression exercises on our German data hint at the relevance of (real) interest rates and demographics over the long run, with the variable importance data (Figure 3) confirming such notions in more detail.[35] None of the existing studies to my knowledge identified secular *trends* in returns, however, and I am unaware of any consistent econometric exercises being applied to sufficiently long-run data: existing research emphasizes that only at such horizons, the full statistical power of conventional econometric tests can be leveraged.[36] At the very least, therefore, one would expect that our data can confirm existing

[35]Proceeding in the "traditional way" for our new long-run German data, basic OLS regressions do support the idea that real interest rates and demographics may be important factors also over the very long-run: see tables A.7.1-A.7.2 in Appendix, and brief discussion. In terms of return drivers, the exercise broadly confirms that higher population growth appears to have spurred demand for additional housing, and the rebound in population growth from the late 17th century thus provides a plausible backdrop to the supra-secular rise in real prices and returns. Here, the finding of significance for real sovereign interest rates is intriguing, and already represents a first hint at the possible long-horizon importance of discount rate factors for returns.

[36]See Rogoff et al. (2024) for interest rates: here, the true positive share for the ADF-GLS test falls off when samples of "only" 100 years of annual observations are used (to 68.8%), and drops even lower for 50 years (29.8%).

consensus on a much more general level.

Where we would expect potentially more substantial revisions are on the debates of "structural breaks" or "inflection points" in global housing: the long horizon could reveal the sensitivities embedded in existing short horizon literature, including the proposition that the mid-20th century marks a decisive break in the asset class.

6.1 Stationarity and forecastability

We begin with ADF-GLS tests (Elliott et al., 1996), a standard approach to confirm or reject forecastability (unit roots) and the existence of stationary time trends. Via Table 4 we can, first, clearly confirm forecastability for all German total return (gross and net) as well as real price series. The stationarity holds for variations of the test without a time trend, always at the 1% significance level, per the Appendix: there I also report a variety of ADF-GLS robustness checks, all country level results not shown in the main table, and also total return decompositions.[37]

Next, we observe that the confirmation of trend stationarity extends to our non-German, non-primary source price and return time series – including clear 1% significance for our benchmark "global" constructions across AW and GW methodologies: once more, these results hold even when not assuming a time trend, and they are also confirmed in a decomposition exercise that tests rent growth separately (reported via Appendix table A.15), a finding that suggests assuming a degree of income growth forecastability and that I will exploit further in section 7.

6.2 Bai-Perron

Next, we test the long-run series via the standard structural break test of Bai and Perron (1998). The Bai-Perron methodology tests time series for up to five structural breaks, has been widely used in the financial and economic literature, and has in particular been employed to identify "turning points" in the canonical housing literature. Specifically, Knoll et al. (2017) use Bai and Perron (1998) as their basis to suggest the 1960s as the inflection point in advanced economy housing markets.

The test is however sensitive to the sample length used, and meaningfully improved sample length in particular have been shown to revise canonical structural break narratives in recent macro-finance literature (Rogoff et al., 2024).

I report the results from our long-run series in Table 5, which employs a standard set-up of the test.[38] I test both nominal return series introduced in section 5 (gross and net bases), as well as the benchmark real housing price indices. For the former pair, we observe that several breaks are revealed: in particular, both gross and net returns share the years 1537 and 1640 as break dates.

[37] See Tables A.7-A.8, along with discussion there. There, I test variations without a time trend, and raw year-on-year (as opposed to log) change variations. In the Appendix, I also detail ADF-GLS results for log rent growth (see table A.15), which equally confirm clear trend stationarity.

[38] As is conventional in the literature, the Bai-Perron is run with a maximum of five structural breaks allowed.

Table 4: ADF-GLS, German primary series, ML series, and components			
	no of lags	t statistic	optimal lag
Germany primary source, 1465-1910 (2020)			
Real gross housing log return, 1466-1910	3	-3.013	
	2	-2.979	Seq, SIC, MAIC
	1	-4.563	
Real net housing log return, 1466-1910	3	-4.854	
	2	-4.718	Seq, SIC, MAIC
	1	-7.091	
RHPI log change, 1503-2020	3	-10.202	Seq, SIC
	2	-9.331	MAIC
	1	-11.863	
ML generated, real price y-o-y log change, 1560-2020			
U.K. RHPI log change	3	-8.897	MAIC
	2	-9.948	Seq
	1	-12.561	SIC
U.S. RHPI log change	3	-8.166	
	2	-8.992	Seq, SIC
	1	-8.699	MAIC
Global GW RHPI log change	3	-8.378	Seq, MAIC
	2	-10.103	
	1	-11.695	SIC
Global AW RHPI log change	3	-4.4537	Seq, SIC, MAIC
	2	-5.314	
	1	-6.945	
Total gross real returns, log, 1560-2020			
Global GW log	3	-5.077	Seq, MAIC
	2	-5.887	SIC
	1	-9.248	
Global AW log	3	-3.220	Seq, MAIC
	2	-3.604	SIC
	1	-5.6461	

Notes: The table reports the ADF-GLS test statistic for several choices of the number of lags k (with a maximum of three lags). The regression includes a constant. The test assumes a time trend (see appendix for ADF-GLS results without time trend). The critical values at the 1, 5, and 10 percent significance levels are the following: -3.48 (1 percent); -2.87 (5 percent); -2.59 (10 percent). "Optimal lag" indicates the optimal number of lags according to the sequential procedure ("Seq"), the Bayesian Information Criterion (SIC), or the Modified Information Criterion (MAIC). The test rejects when the test statistic is negative and larger (in absolute value) than the critical value.

The gross return series shows an additional break for the year 1742. Meanwhile, both global real price indices (GW and AW) record no breaks for the full sample length – in particular no break in the 20th century is revealed.

6.3 Half-lives and adjustment speeds

A variety of asset classes have been analyzed with regards to their adjustment speeds – but curiously, real estate as the largest asset class has not been one of them. We can try to fill this gap, and now ask – how long does it take for real estate returns to mean-revert? The finding of trend stationarity over time only assumes economic relevance if its adjustment speed is sufficiently low. Below, via Table 6, I report results of a conventional half-live test to measure the adjustment speeds for real estate total returns – using the approach recently detailed for sovereign bond yields in Rogoff et al. (2024). We observe half-lives generally around 1-3 years for our "global" total return series across different weighting methods – and also observe that these figures on most bases appear to be moderately rising over time. As such, these adjustment speeds can be regarded (by any conventional definition) as economically meaningful, including for the early modern sub-periods.

How do these real estate adjustment speeds compare to other asset classes for which data exists? Compared to sovereign real rates – where half-lives have recently been found to average between two and eight years over the long-run – real estate returns thus show slightly faster adjustment speeds (Rogoff et al., 2024). Both asset classes show a moderate upwards trend over time in these adjustment speeds. For modern U.S. stock prices, Balvers et al. (2000) found adjustment speeds (half-lives) of 3.1-3.5 years.[39]

7 Discount rates, present valuations, and the 17th century safe asset inflection

Recent propositions that a meaningful amount of contemporary house price variation captures very long-run cash-flow expectations is particularly intriguing (Giglio et al., 2015; Bracke et al., 2018) – as we are lacking even a simplistic reconstruction of relevant discount rate in *past realized* housing prices. Progress on reconstructing past data would contextualize current quantitative estimates.[40]

More broadly, a closer analysis of discount rates embedded in housing helps us distinguish better

[39]The Appendix reports further half-life results for primary source series, see table A.12: the time trend towards longer half-lives is absent in this series, though the absolute adjustment speeds are closely comparable to results in table 6.

[40]Giglio et al. (2015) use U.K. freehold and leasehold pricing discrepancies to back out pricing related to very-long run cashflow realizations, using data over recent years. Their estimate of embedded long-run forward discount rates stands at 2.6% in real terms (a 1% risk free rate, plus a 1.6% risk premium).

Table 5: Structural break tests, Bai-Perron basis		
	no of breaks	break years
Primary source series, Germany, 1465-1910/2020		
Real gross total return yoy	3	1537, 1640, 1742
Real net total return yoy	2	1537, 1640
Real price index change (AW)	0	–
Real price index change (VW)	3	1518, 1573, 1878
Real price index, 1465-2020, fused S-JST	0	–
Mortgage rates, nominal, 1465-2020	1	1685
Mortgage rates, real, 1465-2020	2	1666 1799
Machine-Learning real house price index (RHPI) series, 1560-2020		
Global GW RHPI	0	–
Global AW RHPI	0	–
U.K. RHPI	3	1713, 1882, 1966
U.S. RHPI, 1786-2020	0	–
Netherlands RHPI, 1578-2020	3	1768, 1921, 1975
Germany RHPI	0	–
Total gross return (TR) series, 1560-2020		
Global GW real TR	0	–
Global GW nominal TR	0	–
Global AW real TR	3	1713, 1882, 1966
Global AW nominal TR	0	–
Germany real TR	3	1768, 1921, 1975
Germany nominal TR	0	–
Rental-Sovereign yield spread, 1560-2020		
Global GW rental-sovereign spread	4	1628 1659 1966 1996
Global AW rental-sovereign spread	4	1849 1930 1965 1996

Notes: The table reports the results of the sequential Bai and Perron's test (Bai and Perron, 1998). The test is implemented in Matlab using the Matlab function "pbreak" from Pierre Perron's website. The test is applied to linearly detrended series using a trimming parameter of 5 percent and a heteroskedasticity and autocorrelation consistent variance estimator. The significance level is 5 percent, and the maximum number of allowed break points is 5 for all series. When the test rejects the hypothesis of no breaks, the estimated number of breaks and the break dates are reported in the second and third columns, respectively. "AW" denotes the arithmetically weighted index basis, "VW" the value-weighted basis, and "GW" the GDP-weighted bases. The series "Real price index, 1465-2020, fused S-JST" splices the newly introduced German primary series from this paper over 1465-1910 with the "JST" data series for Germany over 1911-2020: corresponding to the series shown in Figure 2.

Table 6: Half-Lives of Total Returns, Full Sample and Subsamples				
	α	Confid. Interv. α	h	Confid. Interv. h
GW real total returns				
1560–2020	0.50	(0.41; 0.60)	0.77	(0.69; 0.88)
1750–2020	0.43	(0.34; 0.56)	0.71	(0.63; 0.83)
1914–2020	0.62	(0.48; 0.78)	2.13	(0.96; 3.03)
GW nominal total returns				
1578–2020	0.72	(0.63; 0.79)	0.79	(0.71; 0.93)
1750–2020	0.74	(0.63; 0.83)	2.02	(0.69; 2.34)
1914–2020	0.81	(0.67; 0.94)	3.28	(1.79; 9.08)
AW real total returns				
1560–2020	0.28	(0.49; 0.67)	1.07	(0.90; 1.30)
1750–2020	0.62	(0.52; 0.73)	1.24	(0.91; 1.74)
1914–2020	0.74	(0.60; 0.91)	1.82	(1.24; 4.73)
AW nominal total returns				
1560–2020	0.84	(0.76; .90)	0.99	(0.87; 1.45)
1750–2020	0.83	(0.72; 0.92)	1.14	(0.86; 2.07)
1914–2020	0.87	(0.70; 1)	1.60	(1; 1000)

Notes: This table reports median unbiased estimates and 90 percent confidence intervals of α based on Hansen's (1999) grid-bootstrap as well as median unbiased estimates and 90 percent confidence intervals of the half-life (h) based on Steinsson (2008). The regression is $y_t = \mu_0 + \mu_1 t + \alpha y_{t-1} + \sum_{j=1}^{p} \gamma_j \Delta y_{t-j} + \varepsilon_t$, where α is the largest root.

the commonplace notion of "housing boom and busts", with recent literature strongly associating rising real house prices over recent decades with a "global house price boom" (Knoll et al., 2017). In the finance literature, however, it has been preferable to employ present value approaches to distinguish, say, discount rate effects from true departures from fundamentals.

Indeed, present value approaches have been used extensively to assess drivers of return contributors in housing and other asset classes. This includes select long-run asset pricing contributions, including for equities Golez and Koudjis (2018) and LeBris et al. (2019), and select housing markets (Campbell et al., 2009; Ambrose et al., 2013).

All such studies follow a variation of the classic Campbell and Shiller (1988) decomposition, as adopting the Gordon growth model, with the present value (PV) of housing estimated by

$$PV_t = \frac{\Sigma_{CF,t+1+j}}{(1+DR)^(1+t+j)-1} \tag{4}$$

When it comes to housing, stationarity of rent growth is typically assumed (Ambrose et al., 2013; Giglio et al., 2015) – the equivalent in the equity literature being forecastability of dividend streams, e.g. per Koijen and van Nieuwenburgh (2011); LeBris et al. (2019). Stationarity of rent growth is indeed econometrically confirmed for our specific series used, and it is also plausible against widespread evidence of early modern rent growth stickiness.[41]

Analogous to exercises assessing long-horizon historical equity valuations (LeBris et al., 2019), this fact allows us to proceed with a present value model under plausible assumptions, on the basis of equation (4), and compare the global GW house price to the implied PV at time t. The spread between actual and implied PV values then allows us to detect long-run mispricing in housing markets, under the assumptions provided.

The key required input that remains is the discount rate (DR). The discount rate is typically not observable on the aggregate (market) level, and in both contemporary and long-run finance literature There are various conceptual propositions to define DR. In the long-run space, definitions range from VAR estimates of real interest rates (Ambrose et al., 2013) to using legal interest rate ceilings (LeBris et al., 2019). Given the range of plausible definitions, and to avoid strong presumptions, I will proceed holistically, and consider a total of seven definitions that can be reconstructed over the long-run we are interested in – some of which are properly reconstructed for the first time here, others directly use given time series from recent secondary literature.[42] A key variation of housing discount rates is the mortgage rate, which I will briefly focus on separately (section 7.1), introducing a first multi-century dataset. Afterwards, I proceed with a holistic DR approach that allows us to compare realized housing prices to implied PVs, and more robustly

[41] See Appendix Table A.14 for ADF-GLS results on all rent growth series, including from independent secondary sources, confirming typically 1% stationarity significance. Historical evidence on rent stickiness comes, among many others, from Wenderoth (2025), who documents that in the city of Nuremberg, early modern rent contracts were completely unchanged over a 140-year period.

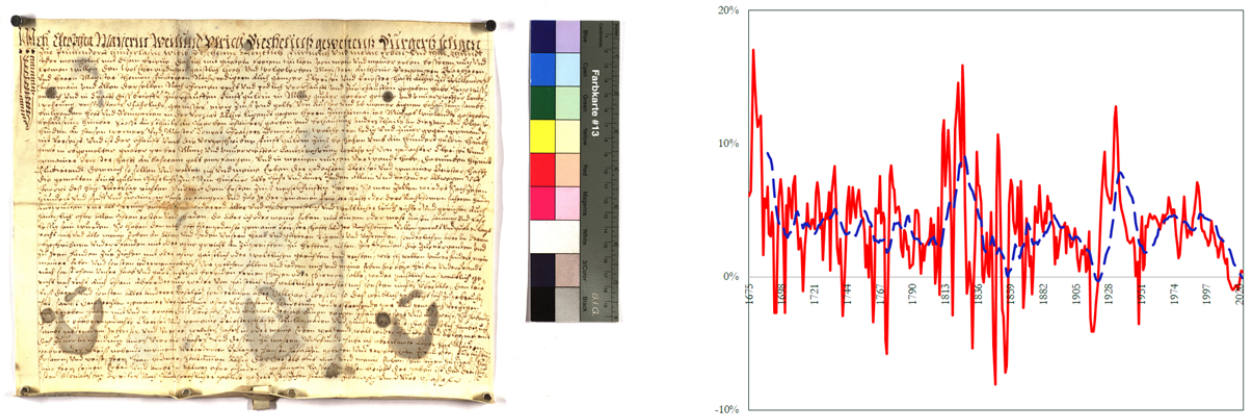
[42] Some, such as Campbell and Shiller (1988)'s proposition to use nominal per capita nondurable consumption growth rates can be measured in past data, but are frequently negative, even when smoothed. I will not consider this measure therefore for conceptual reasons, and thus follow existing literature to assume that the DR needs to be at least zero.

trace secular housing "booms" and "busts" over time.

7.1 Mortgage rates, 1311-2022

Existing literature treats mortgage rates as perhaps the most relevant discount rate to consider for housing markets (Campbell and Cocco, 2015; Justiniano et al., 2019), and indeed the variable importance (Figure 3) underscored the relevance of credit factors. But our long-horizon understanding and data for this variable remains obscure – typically mortgage rate dynamics are considered for the most recent decades at best, and it has not been established whether there are any relevant long-horizon trends in this variable.[43] In Schmelzing (2026), a more detailed discussion introduces a new long horizon mortgage interest rate series, using consistent definitions and maturities, for German secured mortgages. Figure 7 displays a sample contract from this new mortgage interest rate series, which is primarily based on newly digitized archival sources for the pre-1870 era – with the Appendix (section 3.1) providing some further details.

Figure 7: Representative German mortgage contract, dated April 30, 1624 (left handside); and German real mortgage rates, 1685-2022 (right handside).



Notes: The left handside of the Figure image displays a mortgage contract dated April 30, 1624. It is part of the mortgage rate sample over 1311-2024 in Schmelzing (2026). The contracting parties in this case are Cleopha Mayerin zu Pfullendorf on the debtor side, and the monastery of Pfullendorf on the creditor side. The principal amount is stated as 100 Rheinflorin (Rfl.), and the interest stated as 5 Rfl. p.a. Contract held in the archives of the Archbishopric Freiburg, Erzbischöfliches Archiv Freiburg, Urkundensammlung Haid UH095, in: Monasterium.net, URL </mom/DE-EAF/Haid/UH095/charter>, accessed November 11, 2024; on the right handside, displayed are inflation-adjusted German (secured) mortgage rates, using the progressively lagged seven-year inflation rate for Germany in Rogoff et al. (2024) to adjust.

Overall – visible in the snapshot via Figure 7, right handside – what the evidence from credit markets appears to show is that the "recent" fall in mortgage rates that has been put at the core

[43]There are of course various country-level investigations of past housing and mortgage dynamics, but none integrated mortgage rates comprehensively into long-run consistent time series also fusing contemporary data or analyzed time series properties: see for example Hoffman et al. (2000) for France, Clark (2010) for the U.K., and Homer and Sylla (2005) who cover U.S. mortgage rates since 1900.

of explanations of the 2000s housing boom and its subsequent crash has in fact deep historical origins. That is, real and nominal mortgage rates are a downward trending financial variable, with clear trend stationarity confirmed via ADF-GLS tests (Appendix table A.9) – and also echoing our related structural break results in not revealing strong evidence of inflection points during the 1980s or 1990s (Table 5).

The mortgage data adds further to independent evidence that discount rates should be modeled as a secularly downward trending variable. In practical terms, this observation adds weight to the idea that rising valuations in housing may be driven by secular discount rate effects, rather than rising rent growth expectations. To be clear, the idea that credit channels critically influence house price variations and potential "booms and busts" – as for instance in Justiniano et al. (2019) or Greenwald and Guren (2025) – is explicitly supported by the new data: however, existing research clearly needs to abstract from the idea of a key "inflection point" in housing credit in or around 1980, and instead attempt to understand the longer-run continuities, where prices, valuations, and credit clearly interact secularly.

7.2 Discount rates, 1560-2020

Having introduced a new housing discount rate approximation, we can consider this variable amid a broader range of discount rate propositions, and proceed in the reconstruction of implied present values in housing – with the aim of comparing such data to realized prices. The range of DR approximations over time captures a plausible range of risk premia realized in actual financial markets across advanced economy assets. In addition to mortgage rates, I consider an additional six variables to approximate discount rates (all related to conceptual propositions made in relevant literature):

- **Variation 1**: Following LeBris et al. (2019) and using a constant 5% discount rate assumption, given historical legal (private) interest rate ceilings, and extending this flat rate into the present.
- **Variation 2**: Following LeBris et al. (2019) in their "implied discount rate" reconstruction for the French Bazacle company, using 1560-1927.[44]
- **Variation 3**: Following the new reconstructions in this present paper and using nominal mortgage rates for Germany. This series is a new archival-based time series (for pre-1870 data points), and discussed further via section 7.1.
- **Variation 4**: Following Rogoff et al. (2024) and using global GW sovereign interest rates. This choice is partly motivated by Campbell and Shiller (1988)'s model 3 construction using Treasury rates. As related literature emphasizes, it should be remembered that this series is not necessarily "risk free", especially in the early modern period.

[44]The Bazacle company is nationalized in 1946, and therefore the final years of the time series display negative implied discount rates, which the authors exclude.

- **Variation 5**: Following Campbell and Shiller (1988)'s "Version 3" and using the square of annual equity real total returns. Here, empirically, I use Golez and Koudjis (2018) equity data for the leading economy, which they record annually over 1629-2015.[45]

- **Variation 6**: Following Campbell and Shiller (1988)'s "Version 4" and using the nominal aggregate nondurable per capita consumption growth. Historically, I take here the German per capita food consumption growth as measured annually by Pfister (2022) over 1501-1855.

- **Variation 7**: Following the observations in Giglio et al. (2015), I construct a discount rate series that assumes a 1.6% risk premium in residential housing over the risk-free rate. The dynamic risk-free rate definition is here taken from Golez and Koudjis (2018). Over the period 1980-2020, I linearly interpolate the series ("GMS-GK") to terminate at the 2.6% (real) value proposed by Giglio et al. (2015).

This seven-pronged reconstruction allows us conceptual variation, and for all following benchmark exercises I will proceed with tracing the median estimate of the discount rate across the approaches on an annual basis. I plot this median discount rate via Figure 8, on the GW and AW bases. The Appendix contains further results for alternative variations, with broadly identical results: not least, for all but one of the DR variations, the full-sample records an average realized price discount relative to implied PVs.[46]

I then take this median DR and calculate the implied present value, on an annual basis. We can compare the resulting implied PV to the realized actual price, and calculate the difference. Figure 9 charts this difference over time, with positive values implying that realized values are above implied PVs – therefore housing markets are "overvalued" relative to fundamental value; conversely, negative values suggest housing markets at this time are "undervalued" relative to the implied discounted cashflows (rental income). The Figure visualizes that under reasonable assumptions for discount rates (DR), here approximated by the "median estimate" of the seven approaches in Table 6, actual housing prices began their long-run valuation history at a premium relative to discount rates: prior to the Thirty Years' War, indeed, the premium stands at just over 3% to implied values (Table 6, row 9). Then, however, the signs flip, and from the mid-17th century, implied present values in housing consistently stand *above* realized values: in other words, despite the acceleration in real house price growth recorded previously, these price gains apparently fell short of implied gains as suggested by lower discount rates.

What stands out further is the gradual drift towards "neutral" valuation levels over ca. 1700-1800, at the latter date housing valuations secularly reverse, thus echoing the evidence from price side (we recall the inflection around the year 1800 in Figure 5). And not least, spikes towards positive

[45] Golez and Koudjis (2018) define the variables for the "leading economy" (Holland to 1812, U.K. to 1870, then the U.S. to 2015) over time, including the "risk-free rate" but emphasize that there remains a default risk component in this interest rate data.

[46] See in particular Appendix table A.14: here, I record all premium/discount series, by discount rate (DR) definition, and by sub-period. For two variations (mortgage rates and GW sovereign rates), the period of 1980-2020 records positive realized premia over implied PVs.

Table 6: Summary for Discount Rate estimates, and implied present value (PV).

	1560-1650	1650-1820	1820-1914	1914-1980	1980-2020	**Full P**
(1) – Bris et al. 1, IRC	5.00	5.00	5.00	5.00	5.00	5.00
(2) – Bris et al. 2, Bazacle	4.99	1.11	3.92	–	–	2.88
(3) – Mortgage rates	6.04	4.94	4.25	6.66	4.85	5.25
(4) – GW Sov. rates	8.29	5.50	4.71	5.14	5.55	5.82
(5) – Campbell-Shiller v3	0.12	0.14	-0.45	–	–	0.01
(6) – Campbell-Shiller v4	2.64	1.85	2.15	4.88	3.12	2.61
(7) – GMS-GK	4.95	4.86	5.03	4.95	4.13	4.86
(8) – Median DR est.	6.00	4.57	4.19	5.39	4.54	4.81
(9) Actual house price, relative to PV, discount (-) premium (+)	.033	-.180	-.141	-.157	-.123	-.177

Notes: The table reports discount rate estimates (rows 1-7, in %). The final row reports the implied discount of the actual Global GW house price value for each individual period relative to the present value model that uses the median DR estimate, as derived from the sample of DR estimates. For all sources of the DR estimates, see the text. "Full P" denotes average for entire period, 1560-2020.

premia occur around the 1920s and – notably – during the high inflation era of the 1970s and 1980s.[47] But perhaps most strikingly, following the brief spike towards positive premia during the inflation era, valuations once more revert and fall back into deeply negative territory over ca. 1990-2020.

In other words, our present value approach reveals that it is consistent with evidence to suggest that rising real house prices since the 1980s are primarily a reaction to falling discount rates – and have, if anything, *undershot* the decline in discount rates since then. There is from this perspective little evidence of a notable "global housing boom", the presumption of which has become a commonplace in the literature.[48] I emphasize that – though I have focused on the median DR estimate – all general contours discussed are consistent across the all individual DRs introduced. Sharply divergent secular trends in valuation trends would have to fall outside the bounds of the DR range recorded in Table 6.[49]

7.3 Discussion – historical context and plausibility

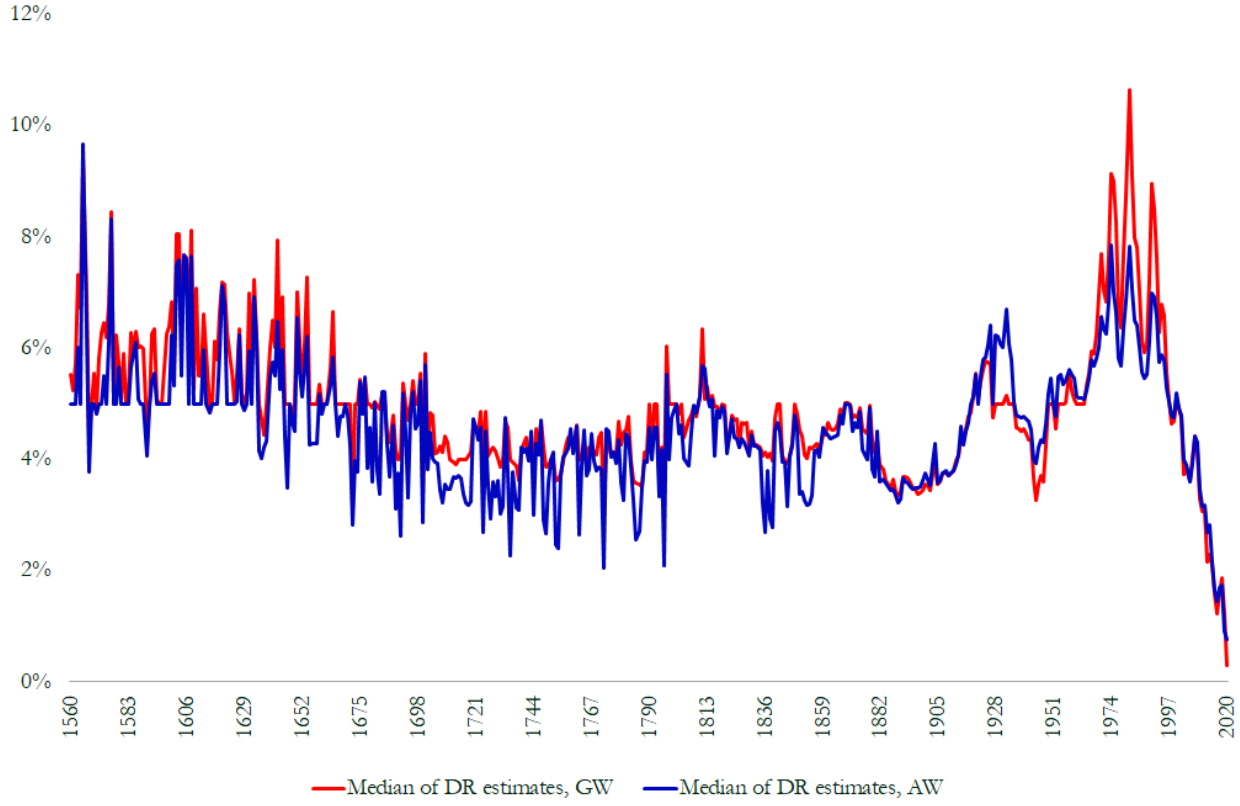
The econometric and empirical evidence raised intriguing hypotheses, and the fact that on the price and return sides our primary source approach for Germany led to similar conclusions as one driven by machine learning techniques is encouraging with regards to the actual historical

[47] I do not explore further whether there is any systematic relationship of the valuation variables to inflation.

[48] A notable exception to the "boom literature" are Himmelberg et al. (2005).

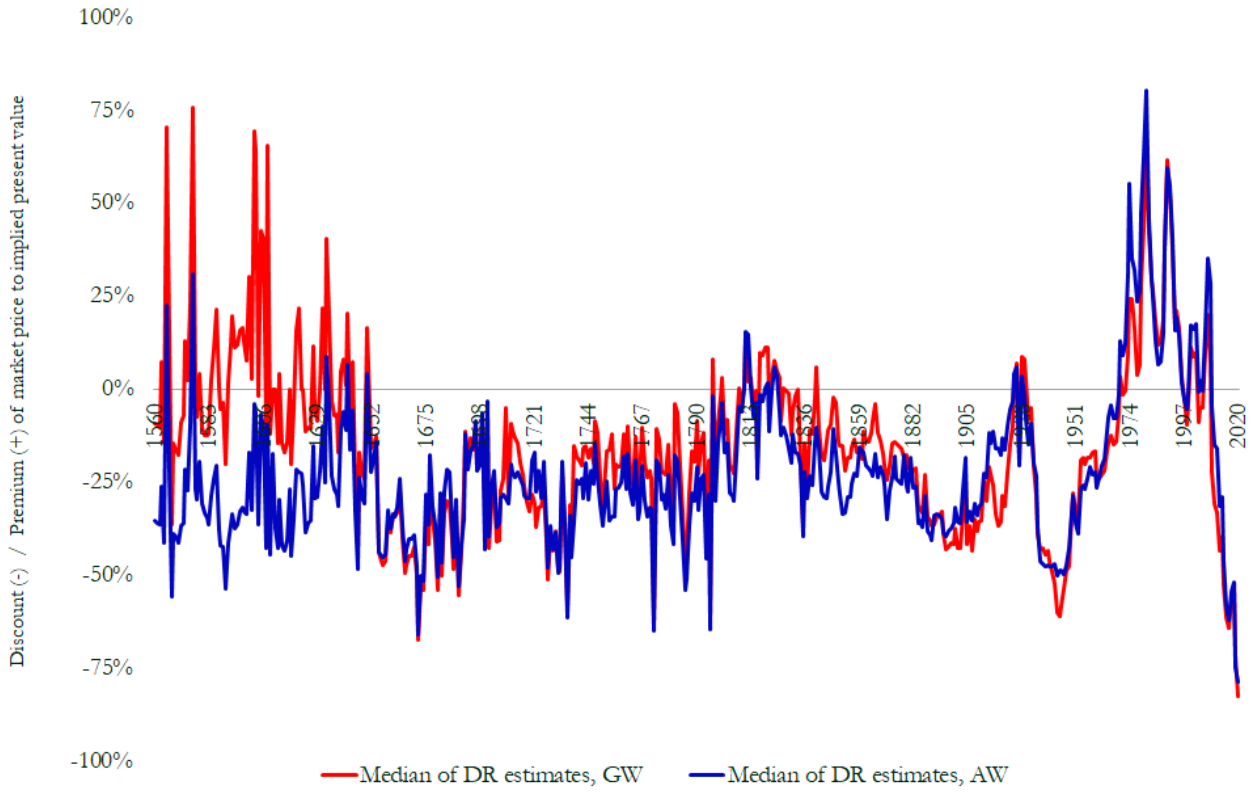
[49] For this, see appendix Table A.14 which records all individual DR implied valuations, by period.

Figure 8: Median Discount Rate (DR) estimate, GW and AW basis, 1560-2020.



Notes: The plot depicts the median estimated Discount Rate (DR) across the seven constructions discussed in the text and Table 6. Two weightings are presented, the GDP-weighted version (GW), and arithmetically weighted version (AW).

Figure 9: Premium (+) or Discount (-) of actual housing prices relative to implied PV, 1560-2020.



Notes: The plot depicts premium (+) or discount (-) of actual Global GW and AW house prices to the implied present value (PV) as calculated via Equation (4), and using the respective median estimate of the Discount Rate (DR) as presented in Figure 8.

plausibility of results.

Here I seek to dive deeper into the context, and cross-check the main data results against recent macro-finance debates, and also historical-political evidence, in which the trends and inflections described should have left traces, however anecdotal such an exercise must (by design) remain.

Generally, the long-horizon results put doubts on the idea that a relevant housing inflection point occurred in the 20th century post-war era, as suggested on the basis of prices by Knoll et al. (2017), or on the basis of the cost of capital (Justiniano et al., 2019). The more striking emphasis is that housing market dynamics are characterized by a significant degree of long-run continuity.

We begin with the price side, where the new headline exercises summarized via Figure 3 and Table 5, instead suggested a broader global inflection point in housing around the late 18th century (ca. 1780). Clearly if such a chronology is plausible, it should have left a variety of imprints.

Indeed, we note that Ambrose et al. (2013) previously posited a "market implosion" in Amsterdam real house prices between 1781 and 1814: their price series are not part of our training or prediction data sets and therefore independently align well with our corresponding price and present value inflections via Figures 5 and 9; similarly, for France, Hoffman et al. (2000) document a well-aligned epic credit boom prevailing throughout ca. 1720-1780 – followed by a unprecedented crash that leads to subdued credit volumes well into the 19th century. And in Germany, covering residential housing markets in Berlin under Frederick William (r. 1713-40) and Frederick the Great (r. 1740-86), Voigt (1901, 76) reports that:

- "The completely new observation of house speculation, a rapid speculative change of ownership of the land, a shortage of rental units and a sharp increase in rental prices, must have made a fundamental impression on the Berliners of the time; time and again, the local sources and writers keep coming back to this puzzling phenomenon..."[50]

By 1765, specific edits were passed in Prussia to reign in on house speculation, though apparently with limited effect, at least until the death of Frederick the Great.[51] These documentary echoes must remain anecdotal, but they are entirely consistent with our empirical reconstructions, and back up the general idea of a major housing boom during the 1700s, followed by a late-century bust – even though this does not contradict the idea that relative to implied present values, the second half of the 18th century recorded continued though shrinking value discounts.

Though notably not for the 2000s, the valuation data *does* align well with other existing notions of a "boom" in housing during the interwar period, including in the U.S. (White, 2009). In addition, it has been largely forgotten how a dedicated literature at the time attempted to make sense of the 1970s and 1980s "housing boom": Mankiw and Weil (1989), notably, analyzed the U.S. "housing

[50]Translated from the original: "Die vollständig neue Erscheinung einer Häuserspekulation, eines raschen spekulativen Besitzwechsels der Grundstücke, einer Wohnungsnot und einer schellen Steigerung der Mietpreise muß auf die damaligen Berliner einer ungeheuren Eindruck genmacht haben; immer wieder kommen die Berliner Lokalschriftsteller auf das merkwürdige Phänomen zurück..."

[51]Prussian edict against "Housing Excesses" of April 15, 1765, via Voigt (1901, 77).

boom" of that era and found a demand peak fuelled by demographic factors occurring in 1980 – precisely the year when our PV-implied housing premia peak in Figure 9.

Fundamentally, underlying this new chronology is the notion that real estate was regarded as a (in relative terms) "safe asset" in the pre-1600 era. While a detailed examination must be conferred to separate research, it is worth emphasizing here that the general proposition of real estate (land and its structures) representing the safest investment in the early modern (pre-1600) age, has for long been a commonplace in the historical literature, though evading any chronological precision as to pricing and cross-overs. As a representative case, we can take McFarlane's analysis of John Fastolf's sprawling portfolio, the famed English mercenary and baron of the 15th century. Regarding the relative safety of investments at the time, we learn that:

- "If a man was anxious to make a wise disposition of his fortune, land had at least one clear advantage. From the negligence or dishonesty of executors there was no sure refuge, but land was less easily misappropriated than cash. It did not evaporate as movables could during the nonage of the heir; it was better protected in the courts; whether its financial yield was high or low, it was altogether safer."[52]

In this sense, our new chronology, including the empirics suggesting a valuation premium of real estate in the pre-1600 era, with the asset being regarded as the prime "safe asset" aligns well with long-standing narrative evidence in the historical literature. But there has been very little synthesis with finance literature: the renewed interest in historical "safe asset providers" and asset pricing conditions in the latter debates simply assumes that sovereigns are providers of the "safe asset" by circa 1700 – for instance via Lustig et al. (2023) and Coppola et al. (2024).

With our new data we can bridge these two literatures at least to a small extent – and now pinpoint chronologically, and with recourse to appropriate (though not exhaustive) data, the historical "cross-over" that saw the emergence of sovereigns as such safe entities, at least as proxied by asset prices from two of the most relevant markets.

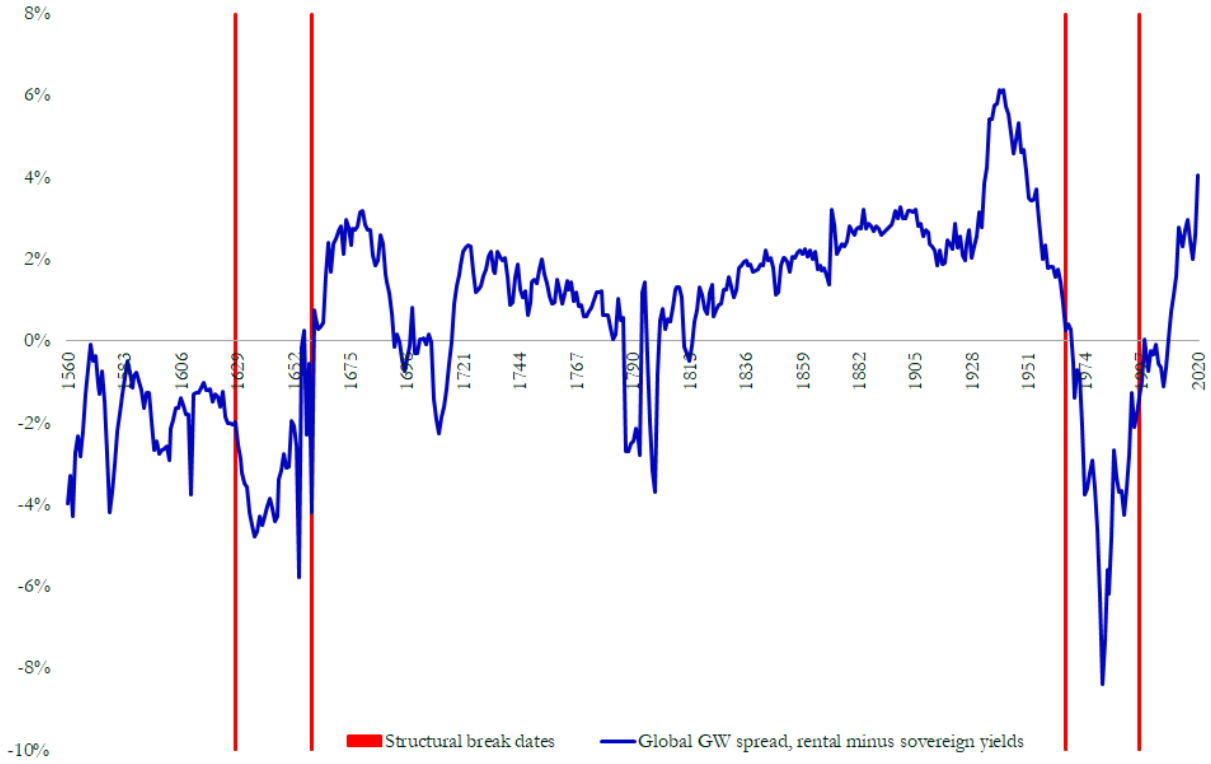
7.4 Rental yields, 1560-2020

The counterpart of discount rates are risk premia. I close this section by visualizing risk premia via Figure 10, which displays the GW basis spread between housing rental yields and sovereign GW yields, over the full sample period, 1560-2020. Red vertical lines indicate the dates of structural breaks as suggested by the Bai-Perron test on this spread, via Table 5: these years are 1628, 1659, 1966, and 1996. Just like we observed falling discount rates, it is plausible to observe rising risk premia in housing – or, conversely, posit rising "safety premia" in sovereign assets.[53] The counterpart of positive housing valuation premia in the initial decades of the sample (ca. 1560-1640

[52] McFarlane (1957, 109).

[53] The linear trend of the GW "safety premium" (defined as the slope of the blue line in Figure 10): +0.62 basis points per annum.

Figure 10: Global GW rental minus sovereign yield, and structural breaks, 1560-2020.



Notes: The plot depicts the spread between Global GW rental yields and sovereign nominal rates. Red vertical lines represent GW structural breaks as reported by the sequential Bai-Perron test per Table 5: the years are 1628, 1659, 1966, and 1996.

in Figure 9) are the negative "safety premia" of the sovereign assets in the decades prior to the Thirty Years' War in Figure 10.

Analogous to the discount rate perspective, what the comparative fusion in the two Figures and exercises conveys, hence, is that several alternative data constructions appear to confirm that a fundamental re-alignment between sovereign and housing assets occurred in the mid-17th century, when realized housing valuations begin to structurally trade at discounts to implied present values – and rental yields structurally rise above sovereign yields for the same set of five countries, when drawing on alternative discount rate variations, and under alternative weighting schemes.

In other words, we can here measure that it was "just before" the early 1700s that sovereign assets apparently started commanding a safety premium – one that we can then take for granted by the 1700s, and one that is measurable on cross-country levels beyond the "leading market" at the time, traditionally associated with England/the U.K.

Of course, a wide variety of alternative approximations of "safety premia" have been proposed, for instance measures that use corporate bond spreads (Krishnamurthy and Vissing-Jorgensen, 2012). However, the relevance of reconstructing rental yield-sovereign yield basis as in this section consists in the fact that these two markets combined comprised by all accounts the two largest (by value) and most liquid investable markets in continuous time, across which early modern

investors shuffled exposure – and are therefore representative.[54]

Therefore, the new data and relative asset performance allows not just the chronological pinpointing of the "emergence" of sovereign yields as an asset commanding a valuation premium relative to the next most important component of the asset universe; crucially, it also establishes that this valuation premium commanded by the sovereign sector is secularly *increasing*. This is a new stylized fact that – perhaps surprisingly – has not been explicitly documented thus far, despite heightened interest in the time variation of discount rates (risk premia) in the asset pricing literature of recent years (Cochrane, 2011).

8 Conclusion

Returning to our initial challenges – what, then, can be learned from reconstructing multi century asset pricing data? Real estate consistently represented the biggest asset class over time, and embedded in its valuations are particular long-horizon asset pricing expectations: yet relative to other assets, our very long-run understanding lags behind. This paper argued that recent leaps in both primary source availability and machine learning models now allow us to seriously advance. In the first contribution of this paper, I took advantage of German primary data innovations, which now allow a granular "bottom up" reconstruction of house prices and total returns: using *Häuserbücher* and newly digitized data, we can construct a new multi-century repeat sales index. Combining such data with machine learning models in a second step – for which I utilize long-run covariates reconstructions, including sourcing new building cost indices –, I provided the first housing price and total (ex post and excess) return series for housing markets covering a substantial share of global aggregate GDP over centuries.

And indeed, on the comprehensive new data I assessed the truly "long cycles" in global housing and sought to contextualize "recent" secular trends in real estate valuations. Importantly, both "bottom up" and new "top down" (ML) approaches – approaches that piece together past realized data in a "backcasting" approach – agree that the idea of a "hockey stick" evolution of global house prices (and returns) appears untenable. Sharp time variation in housing valuations has taken place well before the mid-20th century – and if anything the drastic market inflections around the years 1648 and 1800 appear to be important precedents deserving future study.

Leveraging the much-improved statistical power from long-horizon time series, I was able to confirm stationarity, relatively few structural breaks, and economically relevant adjustment speeds in housing variables. These results confirm some existing stylized facts on a much broader level (forecastability and, say, the principal relevance of credit and demographic factors), but importantly demonstrate that in real estate, decisive multi-century trends appear to underpin contemporary valuation "peaks" as well as return trends in recent data. On a broader level, this paper argued

[54] Goldsmith (1985) is the classic empirical source for capital stock decompositions in the early modern period. Korevaar (2021)'s data for Amsterdam during 1700-50 confirms combined portfolio shares above 75% for sovereign bonds plus real estate.

that we want to take these long-horizon trends in asset pricing seriously, not least because it has shown to significantly enhance related econometric asset pricing analyses.

However, time-variation in valuations, and the observation of secular run-ups in real house prices – over the 1700s, for instance – does not necessarily imply the existence of corresponding "booms" and "bubbles". Indeed, to properly qualify the cyclical and secular patterns in the new price and valuation data, I leaned on recent present value approaches in finance literature which emphasize the role of discount rates.

Indeed, housing lends itself elegantly to a first comprehensive reconstruction of "discount rates" over the long-run, across a plausible historical range of risk premia – not least in the context of recent literature that emphasizes the strong role of long-horizon discount rates on housing valuations. Amid a broad range of plausible approaches, I reconstructed a set of seven relevant conceptual approaches over the long-run, including a new primary source mortgage rate series. A key new stylized fact that is revealed in this paper is that the discount rate should be modeled – by all accounts available, and across a plausible (though not exhaustive) range of past realized risk premia – as a secularly downward trending variable. This fact – perhaps surprisingly – has not been explicitly documented thus far, despite heightened interest in the time variation of discount rates (risk premia) in the asset pricing literature over recent years (Cochrane, 2011).

Utilizing these new discount rate series, I compared implied present values to realized housing prices over time, and found that it results in a nuanced picture of secular "booms and busts" in housing. Importantly, housing appears to have been valued at a premium to implied present values prior to the 1600s, but then entered a sustained era of being valued at a discount: in other words, under conventional assumptions that abstract not least from the fact that we cannot cleanly measure expectations, the rise in real house prices since circa the 1700s appears not to have outstripped the fall in plausible discount rates.

For contemporary purposes, the approach revealed that housing indeed traded at valuation premia during the interwar period and the high inflation phase of the 1970s and 1980s, as consistent with some existing research – but since then, the rise in (real) house valuations does not show signs of a "boom" dynamic. This observation goes against the commonplace proposition of a "housing boom" in the late 20th century across advanced economies – whether the proposed channel is one centering on the price of land (Knoll et al., 2017), or a fall in mortgage rates (Justiniano et al., 2019).

Mechanically, the counterpart of falling discount rates are rising risk premia. Abstracting from particular assumptions in the PV approach, I closed with new raw long-horizon data that confirms that, initially, in the early modern period sovereign interest rates recorded positive yield spreads over rental yields for the same set of key countries, but then dropped more aggressively with time, with an important cross-over occurring around the mid-17th century. Mechanically, this results in a "safety premium" (as defined in this paper) emerging for sovereign assets at this point, with most finance literature simply assuming this measure to be positive by 1700. I argued that the new data also suggests that this safety premium exhibits clear time trends and is secularly increasing for advanced economies, a new stylized fact relevant not least for the vibrant recent finance literature.

References

Adams, Donald R. 1975. Construction Industry in the early Nineteenth Century. *Journal of Economic History* 35(4): 794–816.

Allen, Robert C. 1988. The Price of Freehold Land and the Interest Rate in the Seventeenth and Eighteenth Centuries. *Economic History Review* 41(1): 33–50.

Allen, Robert C. 2001. The great divergence in European wages and prices from the Middle Ages to the First World War. *Explorations in Economic History* 38: 411–447.

Ambrose, Brent W., Piet Eichholtz, and Thies Lindenthal. 2013. House Prices and Fundamentals: 355 Years of Evidence. *Journal of Money, Credit, and Banking* 45(2): 477–491.

Arnold, Wilhelm. 1861. *Zur Geschichte des Eigentums in den deutschen Staedten. Mit Urkunden*. H. Georg.

Ascher, Siegfried. 1917. *Die Wohnungsmieten in Berlin von 1880-1910. Eine Statistische Untersuchung.* C. Heymann.

Bai, Jushan, and Pierre Perron. 1998. Estimating and Testing Linear Models with Multiple Structural Breaks. *Econometrica* 66(1): 47–78.

Bai, Jushan, and Pierre Perron. 2003. Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics* 18(1): 1–22.

Balvers, Ronald, Yangru Wu, and Erik Gilliland. 2000. Mean Reversion across National Stock Markets and Parametric Contrarian Investment Strategies. *Journal of Finance* 55(2): 745–772.

Berlin, Statistics. 1897. *Statistisches Jahrbuch der Stadt Berlin, 24. Jahrgang, 1897*. Selbstverlag Stadt Berlin.

Bianchi, Daniele, Matthias Büchner, and Andrea Tamoni. 2021. Bond Risk Premiums with Machine Learning. *Review of Financial Studies* 34: 1046–1089.

Bochow, Nils, Anna Poltronieri, Martin Rypdal, and Niklas Boers. 2025. Reconstructing historical climate fields with deep learning. *Science Advances* 11: 1–11.

Bolt, Jutta, and Jan Luiten van Zanden. 2020. Maddison style estimates of the evolution of the world economy. A new 2020 update. *MPD University of Groningen release* .

Borup, Daniel, David Rapach, and Erik Schuette. 2023. Mixed-frequency machine learning: Nowcasting and backcasting weekly initial claims with daily internet search volume data. *International Journal of Forecasting* 39(3): 1122–1144.

Bracke, Philippe, Edward Pinchbeck, and James Wyatt. 2018. The Time Value of Housing: Historical Evidence on Discount Rates. *Economic Journal* 128: 1820–1843.

Braeuer, Karl. 1915. *Studien zur geschichte der Lebenshaltung in Frankfurt a.M. während des 17. und 18. Jahrhunderts* . Joseph Baer.

Brückner, Franz. 1989. *Häuserbuch der Stadt Dessau, 24 vols.* Rat der Stadt, Stadtarchiv.

Bundesamt, Statistisches. 2023. Wohnungsbestand im Zeitvergleich. *Zeitreihen, Stand 28. Juli 2023* .

Burgmaier, Andreas, and Gustav Schneider. 1966. *Häuserbuch der Stadt München, 5 Vols.* Oldenbourg.

Campbell, John Y., and Joao Cocco. 2015. A Model of Mortgage Default. *Journal of Finance* 70(4): 1495–1554.

Campbell, John Y., and Robert Shiller. 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Rates. *Review of Financial Studies* 1(3): 195–228.

Campbell, Sean, Morris Davis, Joshua Gallin, and Robert Martin. 2009. What moves housing markets: A variance decomposition of the rent–price ratio. *Journal of Urban Economics* 66: 90–102.

Case, Carl, and Robert Shiller. 1989. The Efficiency of the Market for Single-Family Homes. *American Economic Review* 79(1): 125–137.

Case, Karl, and Robert Shiller. 1990. Forecasting Prices and Excess Returns in the Housing Market. *Real Estate Economics* 18(3): 253–273.

Chambers, David, Christophe Spaenjers, and Eva Steiner. 2021. The Rate of Return on Real Estate: Long-Run Micro Evidence. *Review of Financial Studies* 34: 3572–3607.

Chen, Luyang, Markus Pelger, and Jason Zhu. 2024. Deep Learning in Asset Pricing. *Management Science* 70(2): 714–750.

Clark, Gregory. 2002. Shelter from the Storm: Housing and the Industrial Revolution. *Journal of Economic History* 62(2): 489–511.

Clark, Gregory. 2010. The macroeconomic aggregates for England, 1209–2008. *Research in Economic History* 27: 51–140.

Cochrane, John H. 2011. Presidential Address: Discount Rates. *Journal of Finance* 65(4): 1047–1108.

Cong, Lin, Ke Tang, Jingyuan Wang, and Yang Zhang. 2021. Deep Sequence Modeling: Development and Applications in Asset Pricing. *Journal of Financial Data Science* 3(1): 28–42.

Coppola, Antonio, Arvind Krishnamurthy, and Chenzi Xu. 2024. Liquidity, debt denomination, and currency dominance. *NBER Working Paper 30984* .

Corelogic, SPGlobal. 2024. SP CoreLogic Case-Shiller Home Price Indices Methodology, April 2024. *available at https://www.spglobal.com/spdji/en/documents/methodologies/methodology-sp-corelogic-cs-home-price-indices.pdf* .

Couperie, Pierre, and Emmanuel Roy Ladurie. 1970. Le mouvement des loyers parisiens de la fin du Moyen Âge au XVIIIe siècle. *Annales. Economies, Societes, Civilisations* 25(4): 1002–1023.

Das, Abhimanyu, Weihao Kong, Andrew Leach, Shaan Mathur, Rajat Sen, and Rose Yu. 2024. Long-term Forecasting with TiDE: Time-series Dense Encoder. *Transactions on Machine Learning Research* .

Dimsdale, Nicholas, and Ryland Thomas. 2016. A Millennium of Macroeconomic Data, Version 3.1. *Bank of England Data Repository* .

Dincecco, Mark. 2013. *Political Transformations and Public Finances: Europe, 1650-1913*. Cambridge University Press.

Dirlmeier, Ulf. 1978. *Untersuchungen zu Einkommensverhaeltnissen und Lebenshaltungskosten in oberdeutschen Staedten des Spaetmittelalters*. Carl Winter.

Duarte, Victor, Diogo Duarte, and Dejanir Silva. 2024. Machine Learning for Continuous-Time Finance. *Review of Financial Studies* 37(11): 3217–3271.

Earle, Peter. 1989. *The Making of the British Middle Class*. U California Press.

Eberstadt, Rudolf. 1903. *Rheinische Wohnverhältnisse und ihre Bedeutung für das Wohnungswesen in Deutschland*. G. Fischer.

Eichholtz, Piet. 1997. A Long Run House Price Index: The Herengracht Index, 1628–1973. *Real Estate Economics* 25(2): 175–192.

Eichholtz, Piet, Matthijs Korevaar, Thies Lindenthal, and Ronan Tallec. 2021. The Total Return and Risk to Residential Real Estate. *Review of Financial Studies* 34: 3608–3646.

Eichholtz, Piet, Thies Lindenthal, and Matthijs Korevaar. 2023. Limited Growth: 500 Years of Urban Rents. *Working Paper* .

Eisfeldt, Andrea, and Andrew Demers. 2022. Total Returns to Single Family Rentals. *Real Estate Economics* 50(1): 7–32.

Elliott, Graham, Thomas J. Rothenberg, and James H. Stock. 1996. Efficient Tests for an Autoregressive Unit Root . *Econometrica* 64(4): 813–836.

Flamm, Hermann. 1905. *Der wirtschaftliche Niedergang Freiburgs i. Br. und die Lage des städtischen Grundeigentums im 14. und 15. Jahrhunderts*. G. Braunsche Hofbuchdruckerei.

Ghysels, Eric, Alberto Plazzi, Rossen Valkanov, and Walter Torous. 2013. Chapter 9 - Forecasting Real Estate Prices. In *Handbook of Economic Forecasting*, edited by Elliott, Graham, and Allan Timmermann, 509–580. Elsevier.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2015. Very Long-Run Discount Rates. *Quarterly Journal of Economics* 130(1): 1–53.

Gillissen, Manfred. 2016. *Geschichte der Koblenzer Profanbauten*. Stadtarchiv Koblenz.

Goldsmith, Raymond. 1985. *Comparative National Balance Sheets: A Study of Twenty Countries, 1688-1979*. University of Chicago Press.

Golez, Benjamin, and Peter Koudjis. 2018. Four centuries of return predictability. *Journal of Financial Economics* 127(2): 248–263.

Greenwald, Daniel, and Adam Guren. 2025. Do Credit Conditions Move House Prices? *American Economic Review (forthcoming)* .

Gu, Shihao, Bryan Kelly, and Dacheng Xiu. 2020. Empirical Asset Pricing via Machine Learning. *Review of Financial Studies* 33: 2223–2273.

Hanauer, Auguste. 1878. *Etudes économiques sur l'Alsace ancienne et moderne*. A. Durand Pédone-Lauriel.

Hartung, Bernhard. 1861. *Die Häuser-Chronik der Stadt Erfurt*. Hennings und Hopf.

Himmelberg, Charles, Christopher Mayer, and Todd Sinai. 2005. Assessing High House Prices: Bubbles, Fundamentals and Misperceptions. *Journal of Economic Perspectives* 19(4): 67–92.

Hoffman, Philip T, Gilles Postel-Vinay, and Jean-Laurent Rosenthal. 2000. *Priceless Markets*. Princeton University Press.

Homer, Sidney, and Richard Sylla. 2005. *A History of Interest Rates*. University of Rutgers Press.

Hoover, Ethel D. 1960. Retail Prices after 1850. In *Trends in the American Economy in the Nineteenth Century*, 141–190. Princeton University Press.

Isenmann, Eberhard. 2012. *Die deutsche Stadt im Mittelalter 1150-1550, 4th edition*. Klett Cotta.

Jorda, Oscar, Katharina Knoll, Dimitri Kuvshinov, Moritz Schularick, and Alan Taylor. 2019. The Rate of Return on Everything. *Quarterly Journal of Economics* 134(3): 1225–1298.

Jorda, Oscar, Moritz Schularick, and Alan M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. In *NBER Macroeconomics Annual 2016*, edited by Eichenbaum, Martin, and Jonathan A. Parker. University of Chicago Press.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2019. Credit Supply and the Housing Boom. *Journal of Political Economy* 127(3): 1317–1350.

Kain, John F., and John Quigley. 1970. Measuring the Value of Housing Quality. *Journal of the American Statistical Association* 65(330): 532–548.

Kaspar, Fred. 1985. *Bauen und Wohnen in einer alten Hansestadt. Zur Nutzung von Wohnbauten zwischen dem 16. und 19. Jahrhundert, dargestellt am Beispiel der Stadt Lemgo*. Aschendorff.

Kelly, Bryan, and Dacheng Xiu. 2023. Financial Machine Learning. *Foundations and Trends in Finance* 13(3-4): 205–363.

Keneshloo, Yaser, Tian Shi, Naren Ramakrishnan, and Chandan Reddy. 2020. Deep Reinforcement Learning for Sequence-to-Sequence Models. *IEEE Transactions on Neural Networks and Learning Systems* 37(7): 2469–2489.

Kindleberger, Charles. 1991. The Economic Crisis of 1619 to 1623. *Journal of Economic History* 51(1): 149–175.

Knoll, Katharina, Moritz Schularick, and Thomas Steger. 2017. No price like home: Global house prices, 1870–2012. *American Economic Review* 107(2): 331–352.

Koch, Philip, Viktor Stojkoski, and César A. Hidalgo. 2024. Augmenting the availability of historical GDP per capita estimates through machine learning. *PNAS* 121(39).

Koijen, Ralph, and Stijn van Nieuwenburgh. 2011. Predictability of Returns and Cash Flows. *Annual Review of Financial Economics* 3: 467–491.

Korevaar, Matthijs. 2021. Reaching for yield and the housing market: Evidence from 18th-century Amsterdam. *Journal of Financial Economics* 148(3): 273–296.

Korevaar, Matthijs, and Marc Francke. 2024. Baby Booms and Asset Booms: Demographic Change and the Housing Market. *Journal of Finance* forthcoming.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2012. The aggregate demand for treasury debt. *Journal of Political Economy* 120(2): 233–267.

Landesamt, Berlin-Brandenburg Statistisches. 2022. Zensus 2022. Statistisches Landesamt Berlin-Brandenburg, Basistabellen Berlin, Stichtag: 15. Mai 2022 .

Landvoigt, Tim. 2017. Housing Demand During the Boom: The Role of Expectations and Credit Constraints . *Review of Financial Studies* 30(6): 1865–1902.

LeBris, David, William Goetzmann, and Sebastien Pouget. 2019. The present value relation over six centuries: The case of the Bazacle company. *Journal of Financial Economics* 132(1): 248–265.

Lesger, Cle. 1986. *Huur en conjunctuur: De woningmarkt in Amsterdam, 1550-1850*. Amsterdam Historic Series. Amsterdam: University of Amsterdam.

Lichtenberger, Elisabeth. 1977. *Die Wiener Altstadt. Von der mittelalterlichen Buergerstadt zur City*. Franz Deuticke.

Liebeherr, Irmtraut. 1971. *Der Besitz des Mainzer Domkapitels im Spätmittelalter*. Selbstverlag der Gesellschaft für Mittelrheinische Kirchengeschichte.

Lim, Bryan, Sercan Arık, Nicolas Loeff, and Tomas Pfister. 2021. Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting* 37(4): 1748–1764.

Lustig, Hanno, Zefeng Chen, Mindy Xiaolan, Zhengyang Jiang, and Stijn van Nieuwenburgh. 2023. Exorbitant Privilege Gained and Lost: Fiscal Implications. *Working Paper* .

Lyons, Ryan, Rowena Gray, Allison Shertzer, and David Agoratsos. 2024. The Price of Housing in the United States, 1890-2006. *NBER Working Paper No. 32593* .

Lüdicke, Reinhard. 1933. *Geschichte der Berliner Stadtgrundstücke seit der Einführung der Grundbücher Ende des 17. Jahrhunderts*. Historische Kommission Berlin-Brandenburg.

Mankiw, Gregory, and David Weil. 1989. The Baby Boom, the Baby Bust, and the Housing Market. *Regional Science and Urban Economics* 19: 235–258.

Margo, Robert A. 1996. The Rental Price of Housing in New York City, 1830-1860. *Journal of Economic History* 56(3): 605–625.

Masini, Ricardo P., Marcelo C. Medeiros, and Eduardo F. Mendes. 2023. Machine learning advances for time series forecasting. *Journal of Economic Surveys* 37: 76–111.

Mayr, Anton. 1931. *Die Grossen Augsburger Vermoegen in der Zeit von 1618 bis 1717*. Selbstverlag der Stadt Augsburg.

McFarlane, K.B. 1957. The Investment of Sir John Fastolf's Profits of War. *Transactions of the Royal Historical Society* 7: 91–116.

Metrick, Andrew, and Paul Schmelzing. 2025. Banking Crisis Interventions Across Time and Space. *Review of Financial Studies (forthcoming)* .

Meyer, Johannes August. 1903. *Die wirtschaftlichen verhaltnisse des grund und bodens der stadt Giessen in den letzten 25 jahren.* Bonifacius Druckerei.

Mian, Atif, and Amir Sufi. 2009. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics* 124(4): 1449–1496.

Middelhoven, P.J. 1981. Auctions at Amsterdam of Northern European Pinewood. *The Low Countries History Yearbook 1980* 65–89.

Mullainathan, Sendhil, and Jann Spiess. 2017. Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives* 31(2): 87–106.

Müller, Gustav. 1881. *Werth-Karte des Grundbesitzes in Berlin mit Rentabilitäts-Berechnung und Taxation der Gebäude und Grundstücke*. Alexius Kiessling.

O'Brien, Patrick. 1985. Agriculture and the Home Market for English Industry, 1660-1820. *English Historical Review* 100(397): 773–800.

Paker, Meredith, Judy Stephenson, and Patrick Wallis. 2025. Predictive Modeling the Past. *LSE Economic History Working Paper* .

Pfister, Ulrich. 2022. Economic growth in Germany, 1500–1850. *Journal of Economic History* 82(4): 1–42.

Piazzesi, Monika, and Martin Schneider. 2016. Chapter 19 – Housing and Macroeconomics. In *Handbook of Macroeconomics, Vol. 2*, edited by Taylor, John B., and Harald Uhlig, 1547–1640. Elsevier.

Razum, Matthias, Harald Sack, and Gerhard Weilandt. 2023. Transferring the Nuremberg topographical and temporal model into the public (TRANSRAZ). *Abschlussbericht der Leibnizgesellschaft / IRS* .

Renaud, Joseph Ritter. 1904. *Beiträge zur entwicklung der grundrente und wohnungsfrage in München*. Hirschfeld.

Rogers, James Thobald. 1902. *A History of Agriculture and Prices in England, Vols. IV-V*. Clarendon Press.

Rogoff, Kenneth S., Barbara Rossi, and Paul Schmelzing. 2024. Long-Run Trends in Long-Maturity Real Rates, 1311-2022. *American Economic Review* 114(8): 2271–2307.

SAM. 2022. Statistisches Amt München (SMA) – Indikatorenatlas München, accessed November 1, 2024 .

Schaab, Karl Anton. 1841. *Geschichte der Stadt Mainz, 3 vols.* F. Kupferberg.

Scheiber, Artur Maria. 1948. *Häuserchronik der Marktgemeinde Strengberg (Bezirk Amstetten, Niederösterreich)*.

Schmelzing, Paul. 2026. *Eight Centuries of Global Real Rates, 1311-2025*. Yale University Press (forthcoming).

Schuette, Grete. 1930. Was geschieht fuer Frau und Kind in der Neuen Siedlung. In *Das Wohnungswesen der Stadt Frankfurt*, edited by Schoenberger, Guido, 158–168. Englert.

Schwarz, Klaus. 1968. Der Bremer Wohnungsmarkt um die Mitte des 18. Jahrhunderts. *VSWG* 55(2): 193–213.

Shiller, Robert J. 2015. *Irrational Exuberance. 3rd edition*. Princeton University Press.

Solleder, Fridolin. 1938. *München im Mittelalter*. Oldenbourg.

Stahleder, Helmuth. 2006. *Aelteres Haeuserbuch der Stadt Muenchen. 2 vols.* Ph. C. Schmidt.

Taschner, Michael. 2013. Die Sieben Zeilen in Nürnberg – in vielerlei Hinsicht für die Bauforschung von Bedeutung. In *Nürnberger Altstadtberichte Nr. 38/2013*, 49–68. Altstadtfreunde Nürnberg e.V. – Osterchrist Druck und Medien GmbH.

Voigt, Paul. 1901. *Grundrente und Wohnungsfrage in Berlin und seinen Vororten : eine Untersuchung ihrer Geschichte und ihres gegenwartigen Standes*. G. Fischer.

von Hippel, Wolfgang. 1978. Bevoelkerung und Wirtschaft im Zeitalter des Dreissigjaehrigen Krieges: das Beispiel Wuerttemberg. *Zeitschrift fuer Historische Forschung* 4: 413–448.

Weidenbacher, J. 1926. *Die Fuggerei in Augsburg. Die erste deutsche Kleinhaus-Stiftung*. Selbstverlag.

Wenderoth, Thomas. 2025. *Mietwohnen in der Stadt Verbreitung, Entwicklung und Typologie am Beispiel Frankens 1400–1800*. Volk Verlag.

White, Eugene. 2009. Lessons from the American Real Estate Boom and Bust of the 1920s. In *Housing and Mortgage Markets in Historical Perspective*, edited by White, Eugene, Kenneth Snowden, and Price Fishback, 115–158. University of Chicago Press.

Wilde, Manfred. 1999. *Häuserbuch der Stadt Wolfen, Kreis Bitterfeld, 1550-1990*. Degener and Co.

Wolter, Helmut. 2010. *Das Häuserbuch der Stadt Coburg, 1400-1945, 5 Vols.* Peter Morsbach Verlag.

APPENDIX

1. Details on data sources, and *Häuserbücher*

In Table A.1, I provide full details on the underlying sources used for the repeat sales observations across the ten municipalities in our German sample.

I also provide contextual details for each city – whether accompanying information such as house transfers and inheritances, details on tax assessment values (as opposed to sales prices), and quality improvements on the property level are available in each case.

Table A.1: Overview of sources for Housing Index, repeat sales

	# Source	period	Source type	information
Erfurt	Hartung (1861)	1300-1700	Secondary, cadaster	F
Stengberg	Scheiber (1948)	1465-1860	Secondary	F, N
Mainz	Schaab (1841)	1604-1808	Secondary	
Frankfurt	ISG	1378-1845	Primary, archival	F, N
Koblenz	Gillissen (2016)	1408-1792	Primary, archival	
Berlin	Lüdicke (1933)	1681-1850	Secondary, cadaster	
Coburg	Wolter (2010)	1400-1945	Primary, archival	N, A
Munich	Stahleder (2006)	1368-1571	Secondary, cadaster	F, N, Q
	Burgmaier and Schneider (1966)	1572-1904	Secondary, cadaster	F, N, A, Q
Freiburg	Flamm (1905)	1572-1904	Secondary	
Dessau	Brückner (1989)	1693-1933	Secondary, cadaster	Q
Nuremberg	Topo N	1492-1910	Primary, archival	F, N, A, Q

Notes: The table reports underlying sources for repeat sales observations. For column 5, the following codes are used: F = contains explicit foreclosure events; N = contains additional information on non-sale transactions (such as inheritances or transfers within families); A = contains details on assessment values; Q = contains details on quality improvements.

Tables A.2.1-2 provide the corresponding summary statistics for the house price and rent observations, by city and period. Specifically, Table A.2.2 reports details on the rental observations for the German sample. As emphasized, the rental observations are often – but by no means always – matching the property level price observations. This is in line with the majority of existing literature that has typically abstained from such property level matches (an exception being the recent Chambers et al. (2021)).

Table 2.1: Summary Statistics for Housing Index (RHPI), repeat sales

	# properties	# transactions	Property-years	city size
Erfurt				
1300-1670	5	6	564	med
Stengberg				
1465-1860	17	91	4,240	sma
Mainz				
1604-1808	2	5	140	med
Frankfurt				
1358-1864	27	119	6,280	lar
Koblenz				
1408-1806	18	40	2,742	med
Berlin				
1658-1858	40	281	4,312	lar
Munich				
1391-1904	55	624	15,038	lar
Freiburg				
1572-1904	3	3	55	med
Dessau				
1693-1933	3	9	281	med
Nuremberg				
1387-1943	32	295	10,397	lar
Total				
1300-1943	202	1,473	44,052	—

Notes: The table reports the sample composition of the repeat-sales observations for residential real estate in Germany. Property-years reports the sum of all property level spans of years between first and final sales transaction. City sizes are grouped into three categories: ′"sma" (small cities averaging under 50,000 population across sample period), "med" (medium size cities averaging between 50,000 and 80,000 population across sample period), "lar" (large size cities averaging over 80,000 population across sample period).

Table A.2.2: Summary Statistics for rental yield observations

	# properties	# rental obs	Property-years	av yd
Basel				
1352-1566	5	6	215	7.1
Freiburg				
1407-1564	17	94	4,240	9.1
Konstanz				
1604-1808	2	5	140	6.5
Munich				
1401-1529	4	4	65	3.0
Warendorf				
1772-1863	7	100	186	7.9
Nuremberg				
1529-1805	30	31	39	4.6
Bremen				
1650-1765	2,435	2,438	38,960	2.9
Frankfurt				
1378-1902	28	31	37	10.9
Berlin				
1754-2019	>500	29	n/a	4.6
Total				
1352-1863	2,529	2,702	43,844	7.4

Notes: Rental yields are expressed as a percentage of sales price throughout, taking the price observation that is closest to the nominal rent observation in cases where no direct matching is available in the underlying source itself. Bremen data records city-wide averages across all rental properties assessed via Schwarz (1968). Munich also includes Strassburg and Augsburg data; Frankfurt includes Giessen data. Nuremberg data is partially from Wenderoth (2025). Column 5 records the average gross rental yield in the respective sample, in nominal terms, as a percentage of the house market value (the total number being the arithmetic average of city levels). The modern Berlin data is based on census statistics via Bundesamt (2023), which does not provide the underlying sample size for the aggregate figures – hence we designated a "n/a" value to column 4 for Berlin.

Table A.3: Overview of housing expenditure share estimates

	Source	period	EXP share	coverage
England	Allen (1988)	1600-1646	7.4%	M+T
England	Earle (1989)	early 1700s	up to 40%	M+T
Paris	Eichholtz et al. (2021)	1809-1943	25%	
Amsterdam	Eichholtz et al. (2021)	1900-1979	31.8%	M+C
U.K.	Chambers et al. (2021)	1901-1983	32%	M+C+T
U.S.	Eisfeldt and Demers (2022)	1986-2014	35%	
DMs	Jorda et al. (2019)	1870-2018	"About one-third"	
Elberfeld	Eberstadt (1903)	1891-1901	32.6%	M+C+I
Berlin	Müller (1881)	1870s	20.7%	M+C+I+T

Notes: The table reports estimates for the non-tax housing expenditure share for various geographies and historical periods in secondary literature, together with details. The EXP share in column 4 is expressed as a percentage of the property income, excluding taxes and vacancy adjustments, unless otherwise noted in column 5. Column 5 denotes the components included in the expenditure share by the respective authors: "M" = Maintenance; "I" = Insurance; "T" = Taxes.

2. Details on input estimates: expenditures, taxes, rents

Here I provide details about key inputs in the housing total return series, in particular the non-price components, on which even recent literature relies more on unobservables and generated informed estimates, as opposed to relying on direct property level data.

Table A.3 begins by presenting estimates of total expenditures associated with residential housing, as a share of gross rent. We observe that literature dealing with 19th and 20th century data has operated with relatively similar assumptions on this expenditure component across advanced economy cities, with shares ranging from 21-35% of rental income. Figure A.1 displays the resulting net total return estimate (nominal and real) for Germany, over the period 1466-1866. As with the gross total return basis shown in the main part of the paper, the "U-shape" of returns is maintained – with returns bottoming out around the early 1600s and rising rather continuously from then to the late 1800s.

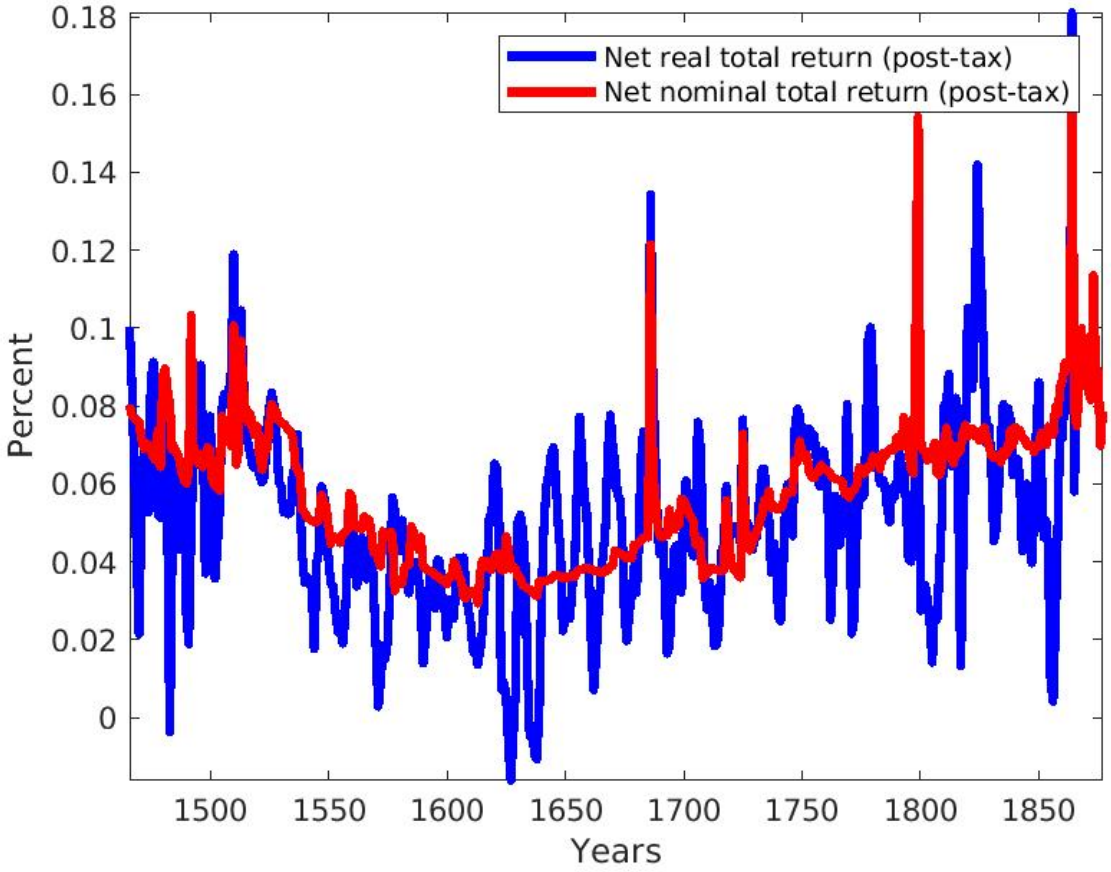
3. Mortgage rates, 1311-2022

Here I elaborate on the construction of the new German mortgage rate series, drawn primarily from newly digitized mortgage contracts (for the pre-1870 era). This archival data is then fused with modern effective mortgage rate data, also secured, as recorded by the Bundesbank.

.1 The new mortgage interest rate sample, Germany 1685-2022

In section 7.1 of the main paper, I introduced a new time series for multi-century mortgage rates. Here I elaborate on further details, with Schmelzing (2026) featuring an extensive discussion. The underlying data set spans the period of 1310-2022 (though I only plot 1685-2022 in the Figures, as informed by the structural break test).

Figure A.1: Total Net Returns, nominal and real (ex post), German housing, 1466-1866.



Notes: The picture displays the German housing total return series, net of taxes and maintenance costs (per estimates detailed in the text and the appendix), in nominal and real terms, all ex post, over 1465-1866. Inflation basis here (as elsewhere) via Schmelzing (2026).

The mortgage dataset features more than 1,100 individual mortgage observations drawn from both urban and rural areas across Germany (and its legal predecessor, the Holy Roman Empire). A representative archival mortgage contract from this sample was shown via Figure 7 of the main paper, where we see a mortgage involving the South German monastery of Pfullendorf (South-West Germany) on the creditor side, dated April 30, 1624.[55]

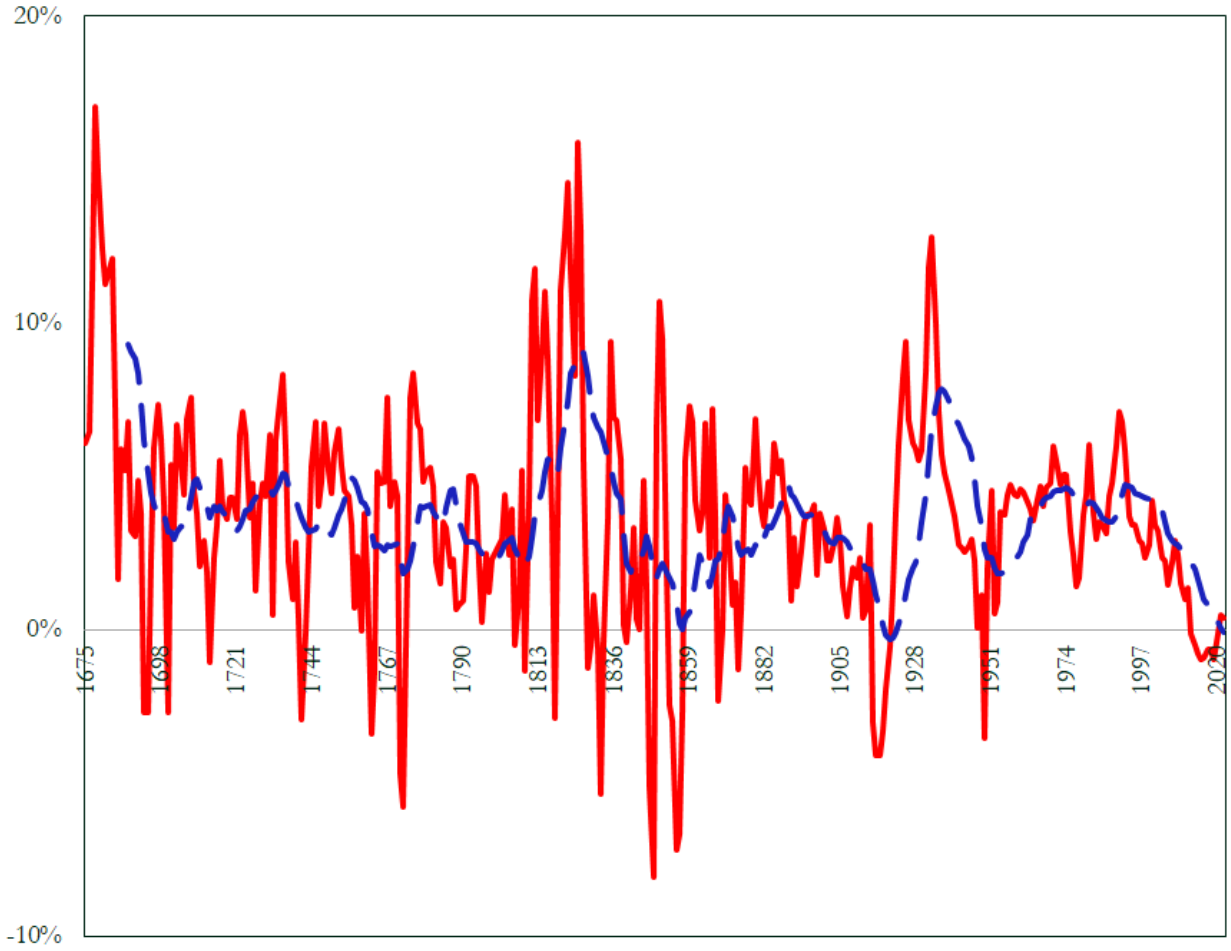
A key early modern source are the digitized mortgage contracts as compiled via Monasterium.net (https://www.monasterium.net/mom/home): this project to this date (July 2025) digitized more than 500,000 early modern primary contracts, across more than 100 European archives, of which rental and mortgage contracts are a key class of sources. This project mainly features ecclesiastical Italian, Dutch, Austrian, and German collections – which account for a substantial share of early modern market activity,

[55]The overall mortgage sample covers 18 different cities and areas (both urban and rural areas), with observations spread relatively evenly across the current geographic borders of the Holy Roman Empire from the inception of the data set: on the creditor side, the sample features a mix of private and institutional creditors (monasteries), on the debtor side exclusively private individuals and households: virtually all contracts are secured by the underlying residential real estate asset.

including real estate transactions. Beyond these digitized sources, I use a variety of historical secondary literature featuring mortgage contracts and associated data: this includes, for instance, the collection of early modern urban contracts in Arnold (1861), or the property sources compiled for the powerful Mainz Cathedral chapter, via Liebeherr (1971).

German mortgage interest rates represent a trend stationary series displaying a general downward trend, and per Table 5 of the main paper record a structural break on the nominal basis in the year 1685. This break in 1685 is echoed when testing the real mortgage basis, where breaks are revealed for 1666 and 1799.[56] I visualize German real mortgage rates in greater detail via Figure A.2, showing this secular downward trend over 1685-2022. Over the period of 1666-1799, real German mortgage rates fall by 7.7 basis points per annum, over the period 1799-2022 real German mortgage rates moderate their decline slightly, and display an average 7.3 basis point decline.

Figure A.2: Real mortgage rates, 1685-2020.



Notes: The Figure displays inflation-adjusted German secured mortgage rates. Nominal mortgage rates are adjusted by progressively lagged seven-year inflation rates for Germany, as sourced from Rogoff et al. (2024).

[56]The real basis uses the seven-year progressively lagged German inflation rates in Rogoff et al. (2024) which follows constructions in related literature. See Appendix tables A.6 and A.7 for ADF-GLS tests on German mortgage rates, as well as sovereign rate-mortgage rate spreads over time.

Table A.5: Correlations new German house price index (GHPI) and returns, other indices

	period	correlation	source
Total			
GHPI – Amsterdam, annual index change	1689-1748	.181	Korevaar (2021)
GHPI – Paris, annual index change	1810-1910	.093	Eichholtz et al. (2021)
GHPI – Germany, annual index change	1871-1910	.014	Knoll et al. (2017)
GHPI – U.S., annual index change	1891-1910	.099	Lyons et al. (2024)
GHPI – Paris, log total returns	1810-1910	.109	Korevaar (2021)

Notes: The table reports correlation statistics between the new German House Price Index (GHPI) introduced in this paper, with several existing historical indices. All index correlations in nominal terms, except U.S. which reports U.S. real correlation with German real basis.

4. Comparison to existing indices, variations of weightings and inflation, and targeting real capital gains

How does my new German price index compare to existing work? Via Table A.5, I report correlation statistics for the new German House Price Index (nominal year-on-year change) to other historical house price indices in recent literature. It is of course not to be expected that very high degrees of correlation should exist – but such measures at least provide a first indication to what extent our results could be representative for broader samples of advanced economies over longer periods of time. Per the table, we generally observe low but positive correlation values when comparing the GHPI nominal changes to historical indices for Paris, Amsterdam, and Germany.

Prominent repeat sales house price indices are constructed on a price-weighted basis, including the prominent Case-Shiller index for U.S. cities (Corelogic, 2024). For the Case-Shiller index, in the words of the creators follows the following approach (Corelogic, 2024): "The composite home price indices are analogous to a cap-weighted equity index, where the aggregate value of housing stock represents the total capitalization of all of the metro areas included in the composite. The numerator of the previous formula is an estimate of the aggregate value of housing stock for all metro areas in a composite index".

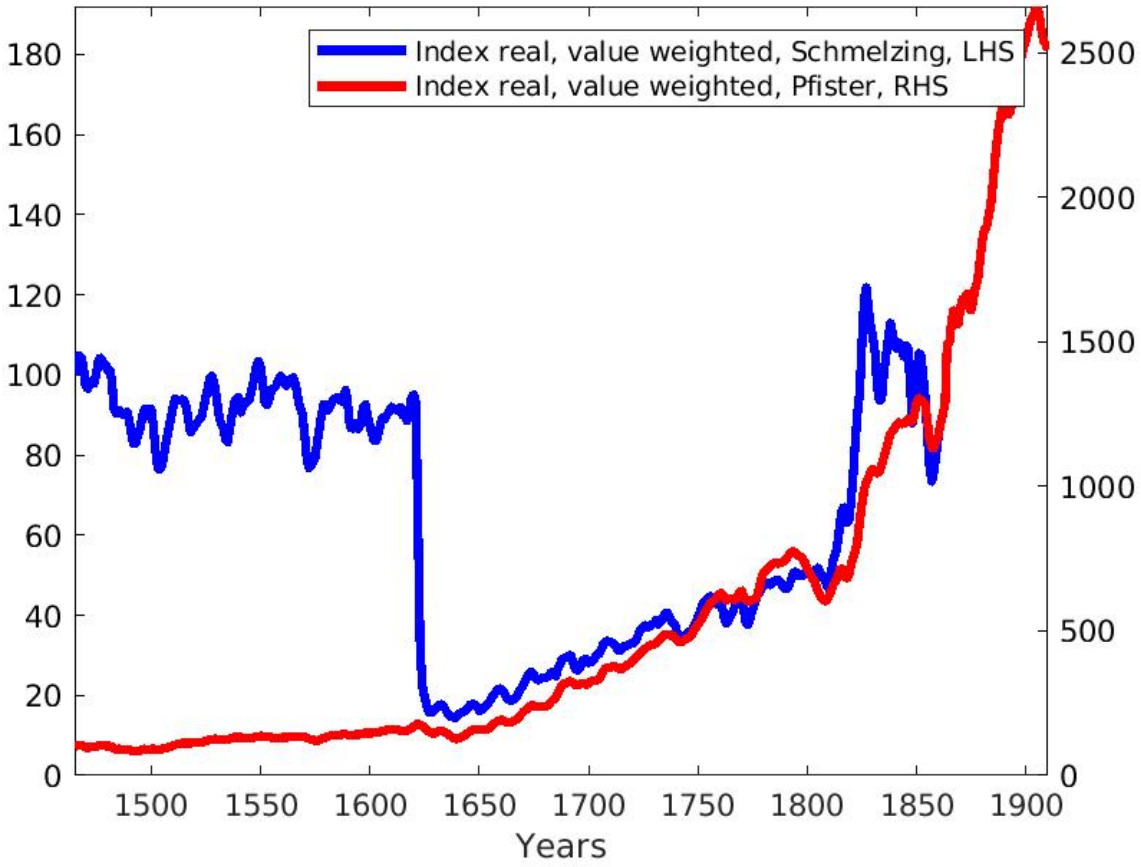
While such a methodology has a different set of potential drawbacks, it could be relevant to test whether high level differences exist to the arithmetically weighted index presented in the main part of the paper. In this subsection, therefore, I present the main German index on such a comparable price-weighted basis – with an index rebasing every year based on shifting value weights. Each individual home enters the aggregate index sample with its own weight – that is, there is no arithmetic weighting on the city level – thus also analogous to the Case-Shiller methodology. Interpolated home values are used for all annual observations in between actual transactions.

Figure A.3 displays this value-weighted version, for the two inflation adjustment approaches. When comparing this value-weighted version to the arithmetically weighted version in the main part of the paper, we observe that once more, the index adjusted with the lagged inflation figures in Rogoff et al. (2024) and Schmelzing (2026) for Germany (blue line),[57] a steep dropoff in real housing prices is observable at the time

[57]The lagged inflation basis in Rogoff et al. (2024) uses the Schmelzing (2026) basis, subsequently I will refer to this series as Schmelzing (2026).

of the outbreak of the Thirty Years War and the associated "Kipper- and Wipper" inflation shock (here from a 1620 index value peak of 95.1 down to a bottom of 15.6 for the year 1628). The general secular trends before and after this inflection point, meanwhile, are comparable across both weighting approaches, but they suggest that the arithmetically weighted version could lead to an over-representation of lower value housing, for which price growth appears to be higher over time. Indeed, for both inflation approaches. For instance, in the period of 1466-1617, the arithmetically weighted versions suggest a real house price index increase of 50-150% (Schmelzing and Pfister inflation bases, respectively) – while the value-weighted versions suggest cumulative real price growth over the period of -7 to +46%. We reach a terminal value for the Schmelzing (2026) inflation adjusted basis in the year 1910 of 197.8 compared to 375.5 for the arithmetically weighted version; the Pfister (2022) inflation adjusted version (red line) on the other hand records a terminal value of 5,226 for the arithmetically weighted version, compared to 2,746 for the value-weighted version. In all cases across all versions, we take the index inception as 1465=100.

Figure A.3: Inflation variations for value-weighted German RHPI, 1465-1910.

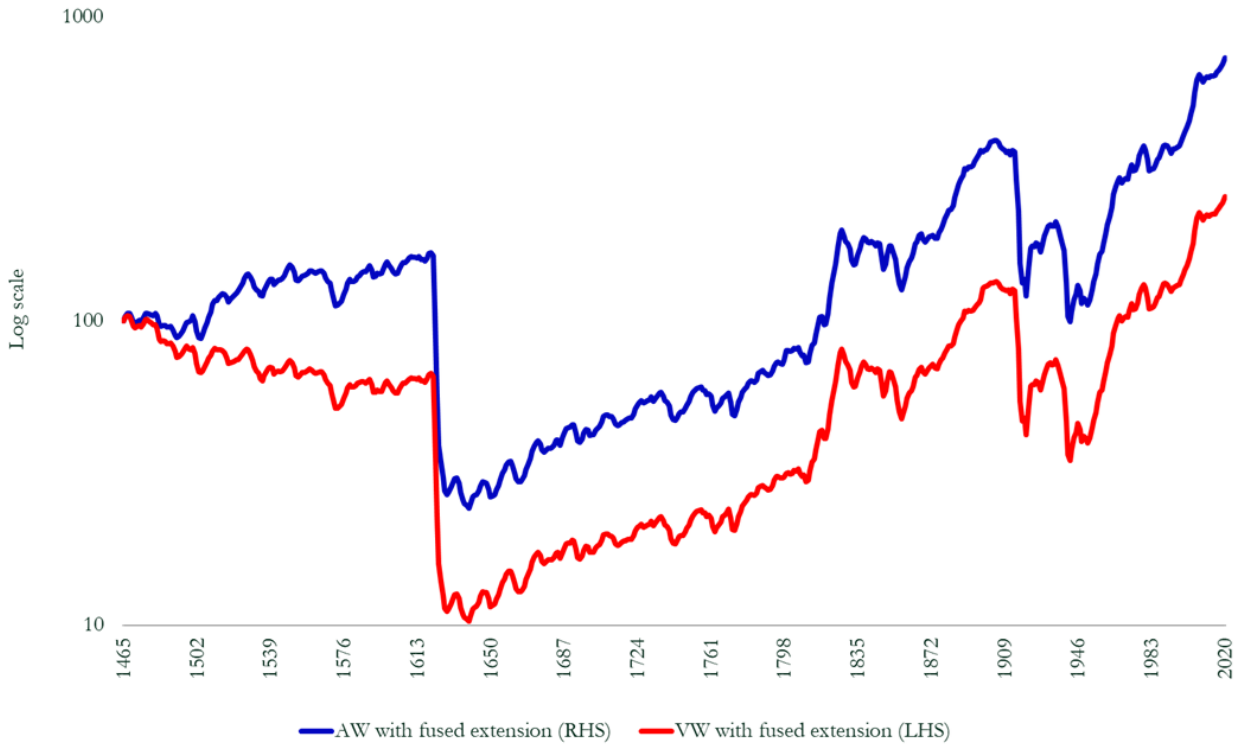


Notes: The Figure displays value-weighted real housing index, taking 1465=100. Two inflation bases are shown, the Pfister (2022) basis (in red, RHS), and the Schmelzing (2026) basis (in blue, LHS). Rebasings are every year, with interpolated housing values for each house in the sample outside of transaction years. Figure is non-logged.

Figure A.4 meanwhile compares the value-weighted to the arithmetically weighted German RHPI, log scale. Here, value-weighted ("VW") refers to weighting the price change of each house in the sample by the absolute value of each house in the total house sample, while "AW" takes the same weight for each house in each year. We observe that the main differences in the two bases are apparent during the pre-1618 era

– when the AW basis shows a moderate *upwards* trends, as opposed to the moderate downward trend of the baseline VW basis; and secondly, we note that the terminal index value for AW is moderately higher, standing at 2020=699, versus 2020=257 for VW. The post-1618 shape and average real price gains y-o-y, however, are closely aligned on both indices.

Figure A.4: AW versus VW German RHPI, 1465-2020.

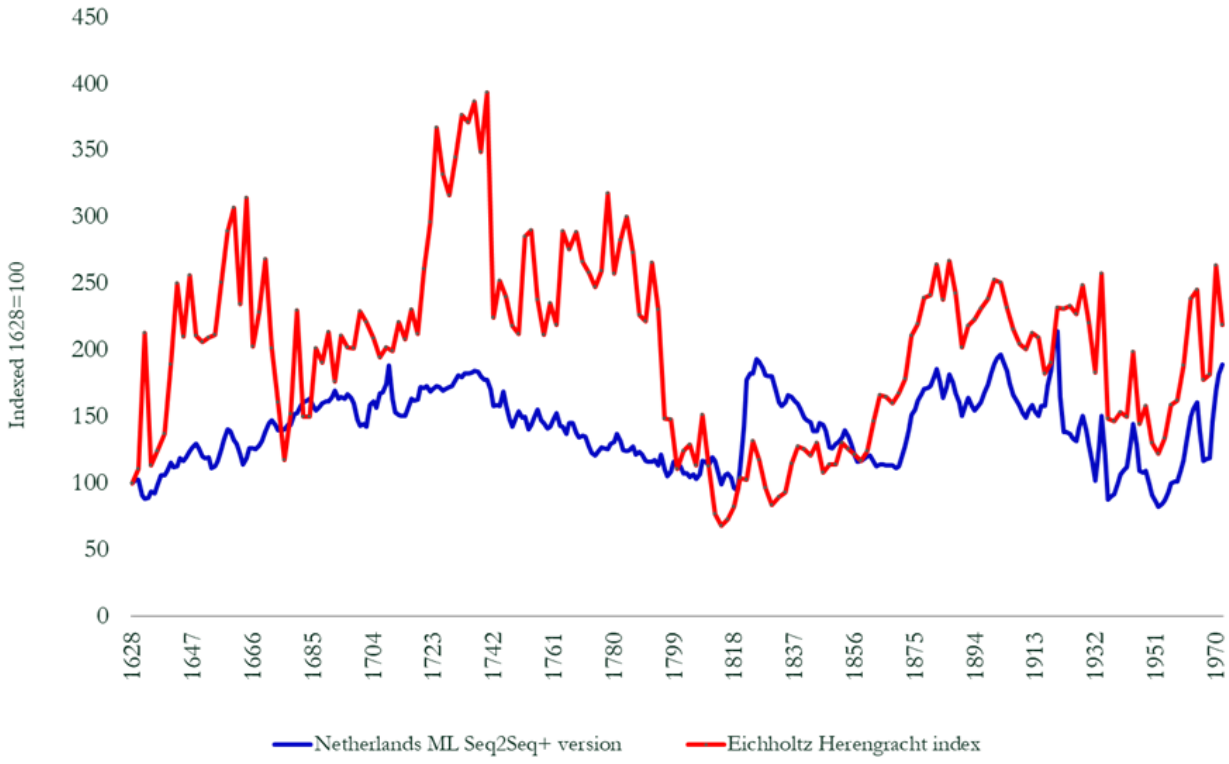


Notes: The Figure displays value-weighted (VW) versus the arithmetically weighted (AW) German RHPI, on log scale, and both indexed to 1465=100, and both adjusted with the same seven-year progressively lagged realized inflation figure for Germany.

Figure A.5 compares the ML-generated real house price index for the Netherlands to a benchmark real house price index for Amsterdam presented by Eichholtz (1997), the "Herengracht Index" spanning 1628-1972. As we are training our post-1870 models on national house price gains – in this case for JST's Netherlands index – we should not expect to see any perfect overlap between the two indices, even with optimal training outcomes. However, the exercise reveals that our ML-generated real house price index in fact mirrors the Herengracht Index in numerous regards: it is more volatile overall – an observation consistent with Piazzesi and Schneider (2016), who document a house price volatility on the city level roughly double that of aggregate national volatility. However, at multiple key inflections, the two indices are closely aligned, including the downturn after 1740, the bottom around 1816, and comparable terminal values in 1972 of 199 (ML) and 218 (Herengracht). The overall correlation between the two indices stands at 0.36.

Overall, these exercises illustrate that the ML process generates time series within plausible ranges, as benchmarked against existing primary source approaches in the literature. Of course, there are multiple conceptual reasons why we should not expect any significant overlap in these indices: for instance, the Herengracht index measures a street-specific repeat sales sample – while the ML index is based off the country-wide "JST" reconstruction for Holland. In any case, the moderately positive correlation to existing city- and street-specific data is a plausible observation, especially in light of agreement on key (visual)

Figure A.5: Machine-learning (ML) Netherlands RHPI index, versus Eichholtz Herengracht index, 1628-1973, with 1628=100.



Notes: The Figure compares the machine learning-generated Netherlands real house price index, compared to the Herengracht Amsterdam real house price index presented by Eichholtz (1997).

inflection points.

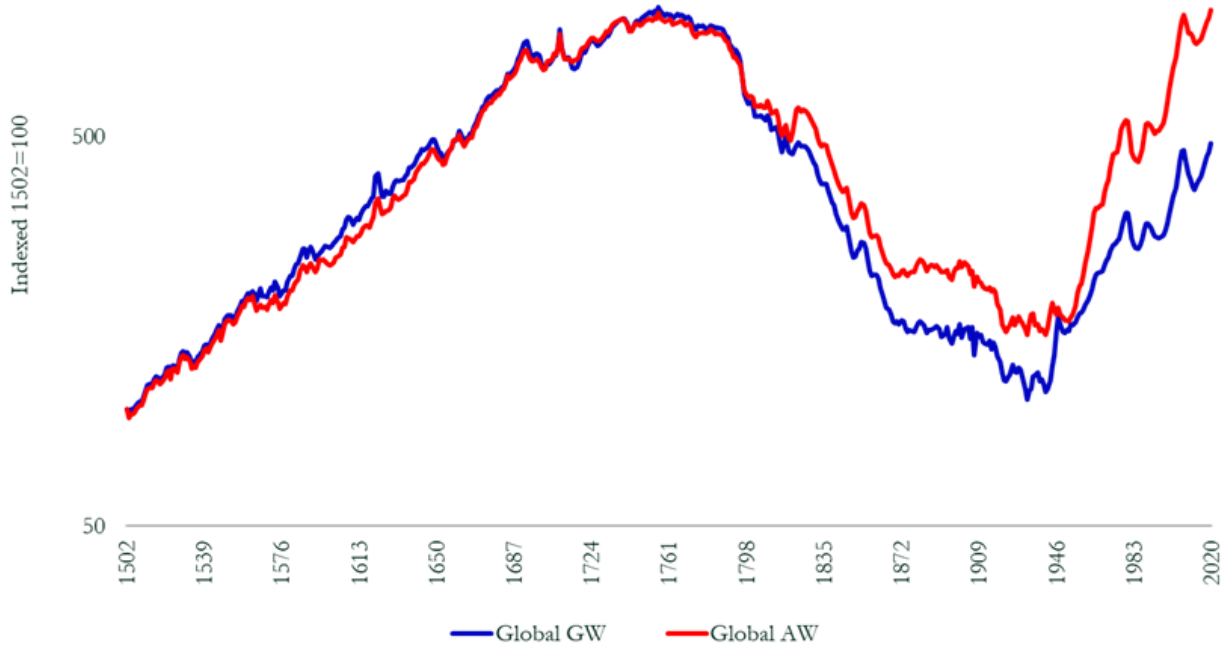
.1 Alternative ML index – real housing capital gains as target variable

Here I present the main alternative construction of ML-based "global" indices, paralleling Figure 5 from the main text. While I used nominal capital gains as the target variable in the main part of the paper – Table 2 in the main paper indicated that selecting real housing capital gains as the target variable itself might on some models yield at least comparable, if not superior, results – though it imposes some more stringent assumptions on the variable construction approach.

Figure A.6 displays the second best-performing ML model across all models tested in Table 3 of the main paper, the Seq2Seq+ model targeting *real* (as opposed to nominal) capital gains, for both the AW and GW global weighing bases. We observe that in terms of the general contours, there is actually a considerable degree of consistency between the nominal and real-derived RHPI indices: our real-derived RHPI also records the major turnaround point in the final years of the 18th century – and equally agrees with the generational reversal around the 1930s. A major difference is that on the GW basis here, real house price levels by the year 2020 remain *lower* than the previous high watermark: this is not the case for our benchmark basis (nominal cap gains as target variable) in the main part of the paper, where GW RHPI by 2020 are 3% higher than the previous high watermark in 1786, and at new all-time peaks.

This alternative real capital gains series is strictly agreeing on the key econometric properties outlined for the benchmark nominal capital gains model, across ADF-GLS (trend stationarity), structural break, and half-life results (all unreported here).

Figure A.6: Global RHPI, AW and GW, 1502-2020.



Notes: The indices display AW and GW weighted global real price indices based on the Seq2Seq+ prediction model for the five individual countries – analogous to Figure 4 in the main text – but this time taking "real capital gains" as the target variable, and fusing predicted values with modern data in "JST". In log terms (1502=100).

.2 Source details – building costs

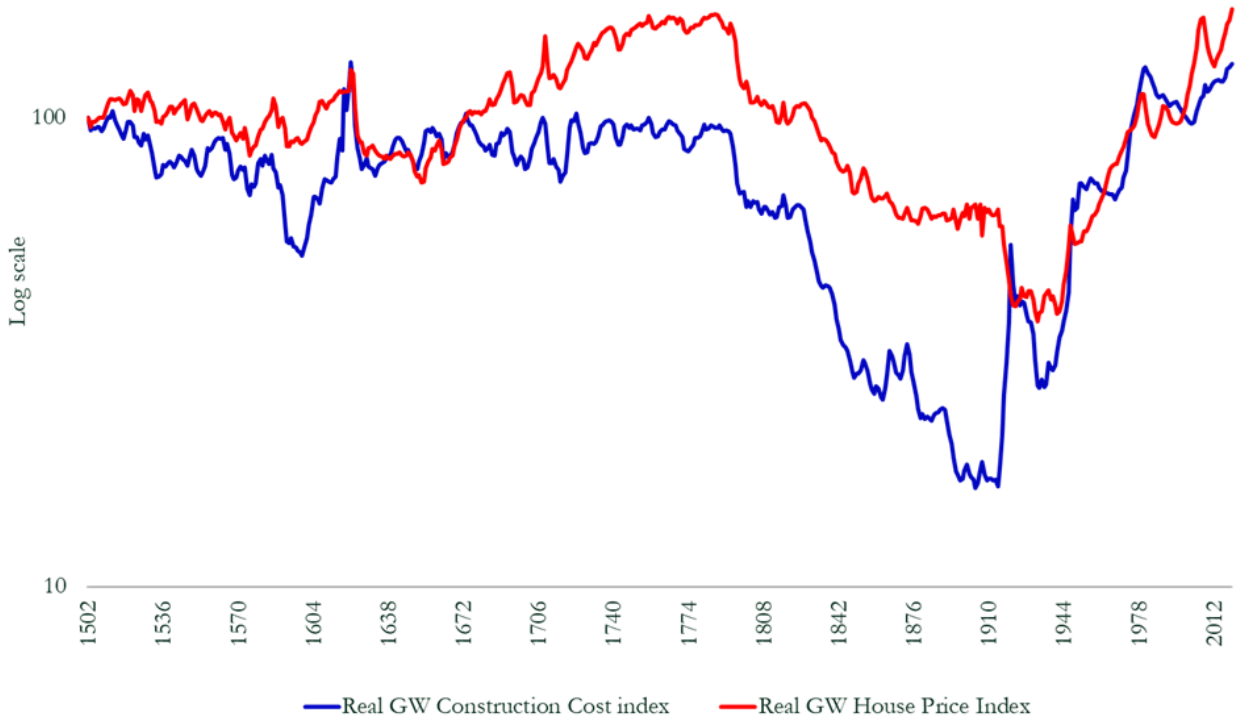
Next, I list the sources and methodology for our building cost indices, by country. First, I note that all post-1870s building cost figures used for the training data are sourced via respective national statistical agencies: note that I do not use the limited (post-1950) data available via Knoll et al. (2017). We recall generally, per Figure 3 of the main paper reporting variable importance, that building cost indices – several of which we newly construct in consistent ways from raw commodity data – exhibit key relevance for the machine-learning approaches presented. The early modern, pre-1870 data is obtained as follows:

- Germany: we use the building cost index constructed by Pfister (2022), who reports index values on a quinquennial basis over 1500-1860 – we linearly interpolate annual frequencies.
- U.S.: we use the U.S. construction cost index in Adams (1975), "Variant B", reported with annual frequency over 1785-1860. The author's "Variant B" is one of several indices created for the same period – we choose this particular variant as it allows for index weight changes from substitution effects over time, rather than operating with fixed base weights ("Variant A").

- U.K.: Spanning 1661-1820, O'Brien (1985) constructs a building cost index for Britain by weighting equally (50:50) material cost changes, and construction wages. We replicate this approach for the pre-1661 data using the price components in Rogers (1902): Rogers records annual price changes for timber, bricks, lime, board, laths, crests, and slates – these seven components are arithmetically weighted to constitute the materials component; for construction wages pre-1661, we use the building craftsmen and building laborer wage data in Allen (2001).
- Holland: we use wood prices in Middelhoven (1981) for 1717-1865. To these, we add building labor costs obtained via Allen (2001) for Amsterdam. Commodities and labor costs are weighted arithmetically (50:50), identical to the approach in O'Brien (1985) for the U.K.
- France: we use Hanauer (1878), who records building material prices in Alsace over 1407 to 1875, covering categories including bricks, tiles, lime, and wood.

Below we directly compare our new building cost indices to the aggregate global real house price indices – both on the identical sample and weightings on the GW (Figure A.7), and the AW (Figure A.8) bases.

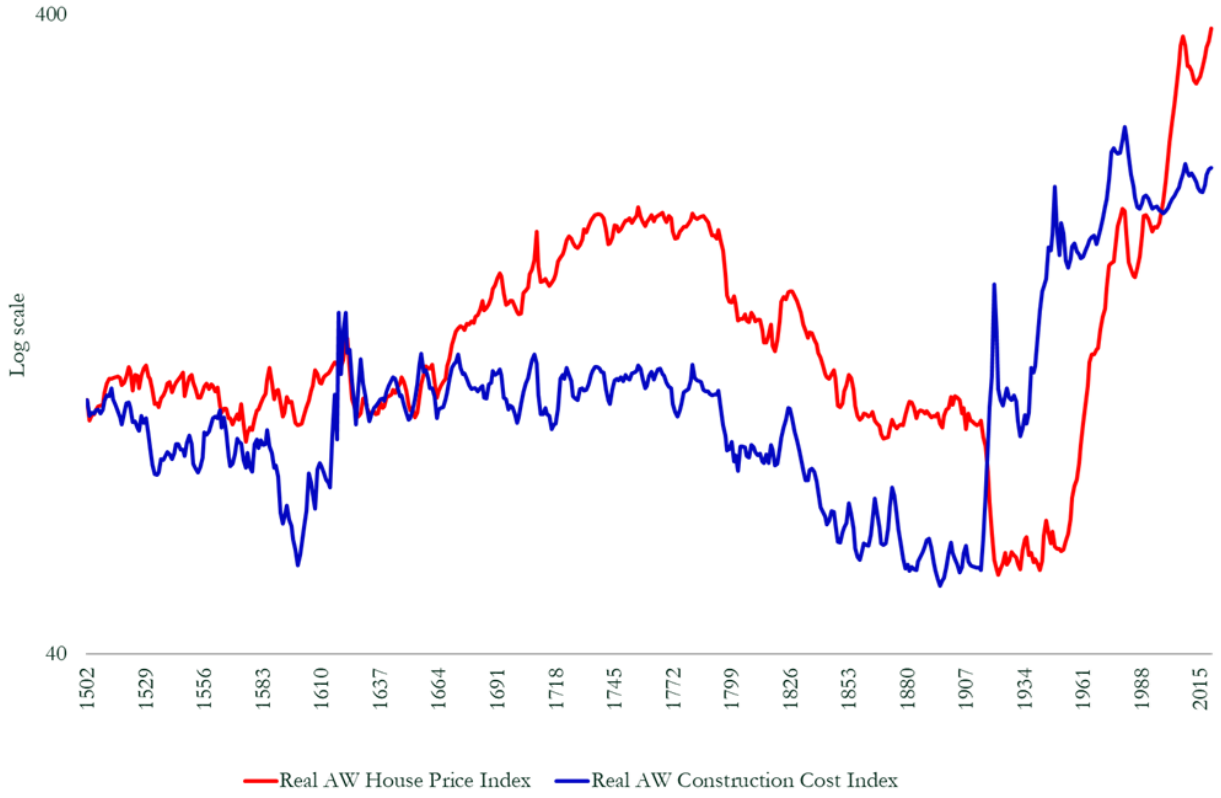
Figure A.7: Building Cost Index versus RHPI on GW basis, with 1502=100.



Notes: The Figure compares the building cost index less the general (seven-year progressively lagged) inflation rate (the latter as sourced by Rogoff et al. (2024)) relative to the Real House Price Index - both on the identical weighting basis (GW = GDP-weighted) covering the same five countries.

Finally, Figure A.9 compares directly the real construction cost index presented in Knoll et al. (2017) with the indices constructed (for the post-1870 period used for training purposes) in this paper, and as based on the national statistical agency releases. We observe that the Knoll index (green) – an index that arithmetically weights 11 countries but is only available in decomposed form over 1950-2012 – is extremely smooth and only implies a 2.5x increase of construction costs over 1880-2012.

Figure A.8: Building Cost Index versus RHPI on AW basis, with 1502=100.



Notes: The Figure compares the building cost index less the general (seven-year progressively lagged) inflation rate (the latter as sourced by Rogoff et al. (2024)) relative to the Real House Price Index - both on the identical weighting basis (AW = arithmetically weighted) covering the same five countries.

Also added is a real U.S. building materials index (here 1891=100, dashed black line). We observe that the Schmelzing indices are much closer to this U.S. index, as opposed to the Knoll reconstruction. The building cost index being an "input cost" index, as opposed to the Knoll preference for "output cost" indices, this suggests that the differences in measurement might be driven by the inclusion of output cost-specific items – such as profit margins for builders, and labor costs in the construction industry.

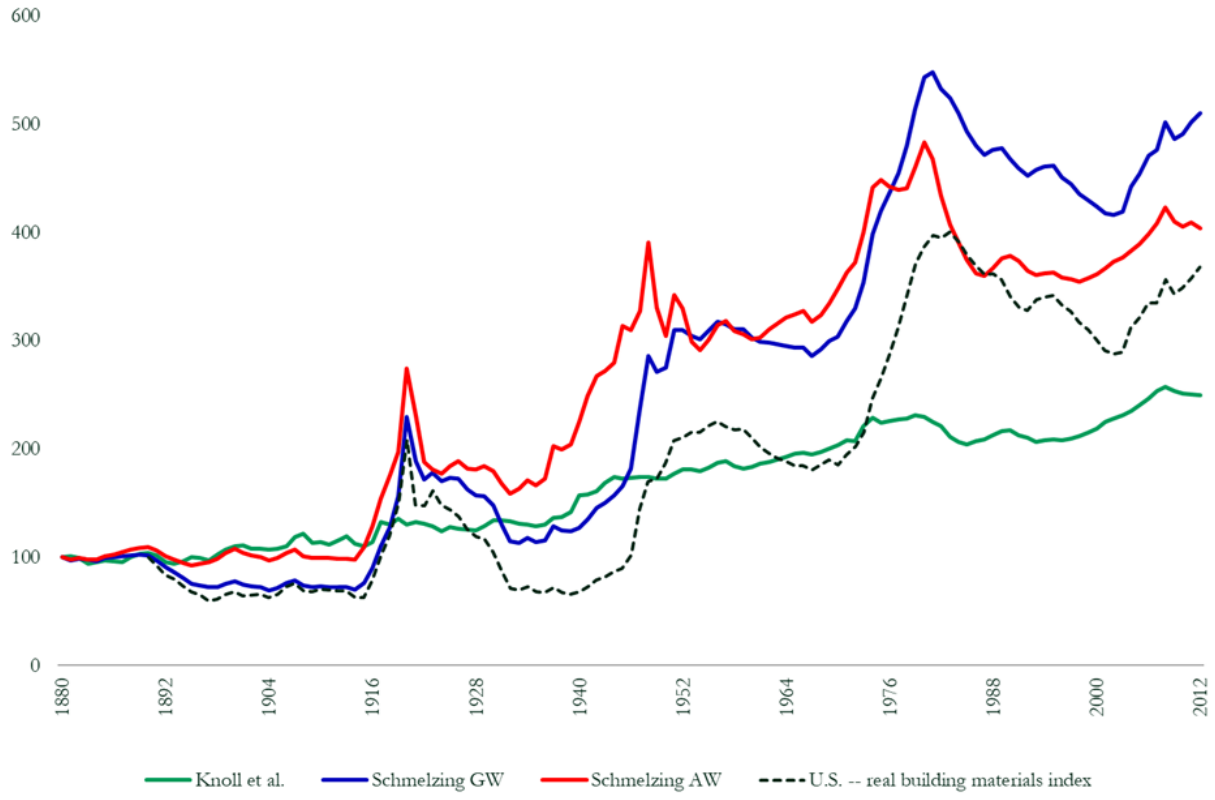
I leave the precise explanation out since for the ML exercises it is not at all necessary that the Knoll and Schmelzing bases agree: my aim in the ML constructions is to identify *any* variables that *both* display high variable importance in the post-1870 era, and can be reconstructed for the pre-1870 era.

.3 Source details – rental data

Here I detail the sources for the rent components used in the total return and rental yield aggregations.

- Germany: I use the primary source series as detailed in this paper. The underlying sources are per Table 2 of the main text.
- U.S.: For 1830-1860, I use Margo (1996) who covers rental yields for New York. Afterwards, I use Hoover (1960) to 1880, and then I use Lyons et al. (2024) rental yield series from 1890. The 1881-1890 annual data points are linearly interpolated, and from 2007 I rely on "JST".

Figure A.9: Comparison Real Building Cost Index, Knoll et al. (2017) vs Schmelzing vs U.S. real building materials index, 1880-2012, with 1880=100.



Notes: The Figure compares the real building cost index in Knoll et al. (2017) versus the series constructed for this paper ("Schmelzing"), GW and AW, over 1880-2012, indexed to 1880=100. The black dashed line represents the U.S. real building materials index, as sourced via Bureau of Labor Statistics (BLS) / FRED, and here indexed to 1891=100. All indices are in real terms, that is I subtract general inflation rates from the building cost rate, as Knoll et al. (2017) also do.

- U.K.: I use Clark (2002), who reports decadal-frequency rental indices for "London" and "outside London", starting in 1550. We focus on the "outside London" series given longer coverage, and linearly interpolate annual observations. From 1890, I switch to "JST" rental yield observations.
- Holland: I use Lesger (1986)'s composite index from 1578. Index values are reported in the source on an annual level, and run through 1850. Afterwards, I use the Amsterdam data in Eichholtz (1997), and from 1974, I switch to "JST".
- France: I use Couperie and Ladurie (1970), who construct Parisian rental prices from 1402 in tri-annual frequency, to 1788, for which I linearly interpolate annual observations. From 1809, I use Parisian data from Eichholtz et al. (2021), who report annual-level data.

5. Housing quality – and living space estimates

As mentioned briefly in a preceding section, it is crucial to note that all existing housing indices of relevance here continue to acknowledge the challenges with regards to quality changes in housing – and that remains true for both "hedonic" and "repeat sales" constructions of such indices. That is, it might well be that

much if not most of the residual positive index changes over time could still capture quality improvements. Representative for most of the existing literature, Knoll et al. (2017) state that:

- "researchers using our dataset in the future should take into account that accurate measurement of quality-adjustments remains a challenge."

These authors – in line with preceding work – undertake no specific quality adjustments for price indices in the early part of their sample, and instead splice early observations with quality adjusted series in the second half of the sample. My basis is therefore consistent in that no additional quality adjustments (beyond choosing the repeat sales methodology in the first place) are made: conceptually, that is of no concern for any of the preceding or subsequent discussions, though of course it would be *desirable* to know precisely the underlying decomposition in quality-related price changes and the residual components.

Yet, we can at least sketch the general contours and focus on one recognized major quality measures for which (tentative) long-run approximations are feasible: specifically, floor space per capita.

Key contributions decomposing the "quality" input into residential housing have posited that the floor area per capita and the number of rooms constitute major ingredients for housing quality, for both renters and owner-occupiers – and indeed a statistically more significant one compared to factors such as neighborhood, schooling, or crime indicators (Kain and Quigley, 1970). This is in particular true for existing long-run studies which try to devise constant-quality indices, and which emphasized floor space per capita as a dominant input (Clark, 2002; Eichholtz et al., 2023).[58] Given that the *Häuserbücher* frequently inform us about house level parcel and floor sizing, we can derive some long-run statistics on representative quality changes on this key indicator: we know, for instance, that the Berlin houses in our sample run at an average 417 square meter parcel size, and know that in the two districts from which our houses are drawn (Dorotheenstadt and Friedrichstadt), persons per house are calculated at 14.9 for the year 1711, yielding a square meter living space figure of 28 per capita.[59] Overall, however, the data suggests that it is not implausible to assume an increase of representative floor space per person in German urban areas by a factor of 2-2.5x between the early 1500s and the early 21st century – from levels between 15-20 square meters in the era prior to the Thirty Years War, to contemporary averages of 38-40 square meters based on most recent census data for key German cities. That is an increase that falls far short of the increases in real property level house prices over the same period when using arithmetically weighted indices – however, interestingly, value-weighted indices suggest that living space growth could account for a not insignificant share of overall real price growth. Specifically, our value-weighted RHPI that fuses JST data post-1910 records a full-sample increase in real house prices of 3.9x over 1465-2020. To the extent that the sample over time captures "representative" housing samples, this suggests that living-space adjusted real house price growth on such a basis would only account for the residual 1-1.5x in real price growth.

Repeat-sales housing indices are considered as the best construction method to incorporate quality changes in housing assets – but they still remain subject to potential quality-related distortions. In general, the

[58] As Eichholtz et al. (2023) state, "housing quality per capita [is] most relevant when discussing the evolution of the standard of living. Such concerns are relevant: Adam Smith explains the relative affordability of London housing as the result of mass-scale sub-letting of parts of dwelling homes (Smith, 1776). Smith's observation suggests that the number of people per house may have varied substantially across cities (and potentially over time). To address these concerns, we need to look at housing space per capita." And similarly, via Clark (2002), "to estimate housing rental values while controlling for quality, I have used the subset...of the observations where there are multiple observations on the same property at different times, and no indications of change in size or quality of the structure, or of change in lot".

[59] Earlier house and parcel sizes in Germany are typically measured in "Ruten" (1 Quadratrute = 14.1846 square meters) or "Fuss" (1 Fuss = 0.313 meters).

literature has settled on preferring repeat-sales indices to control for many quality-related changes in housing, but acknowledges that no perfect methodology exists to fully separate "true" housing index changes from housing quality changes. Indeed, some key literature dealing with long-run construction has resorted to simply acknowledging the potentially distorting role of quality changes, but gone ahead with splicing quality-controlled with non-quality controlled sub-indices (Knoll et al., 2017).[60]

So it appears relevant to get a sense of the potential size of the quality effect, to the best extent possible. For this purpose, the following Table A.5 begins by systematically compiling available data for living space in early modern and modern German cities, converted to a per capita basis. I only include data points where authors refer to "typical" or "representative" housing units in the respective city, and have included precise space data.

Here we observe that there appears to have been a clear upwards trend in living space per capita in German urban areas over the very long run – perhaps an intuitive general observation. Our earliest data points indicate typical living space of around 15 square meter – a figure which then rises over the course of the early modern period, and by the eve of the Thirty Years War rose above 20 square meters for the first time. From there, another doubling of average living space takes place by the early 21st century, where current official city surveys for Berlin and Munich indicate living spaces per capita of 38-40 square meters.

This means that, indeed, some of our housing price appreciation is likely driven by house quality improvements related to size: given homes can be enlarged, either by adding "attached houses" on the same address, or through adding additional stories, say. That both effects took place is confirmed by anecdotal evidence in our cities. At the same time, at least on the basis of the limited number of concrete data points, it is not evident that such space dynamics themselves had any decisive role per se: assuming a doubling of living space for a representative housing inhabitant in our sample over half a millennium would translate into an annual increase of living space of just 0.2%.

To provide some additional nuance to quality improvements, Figures A.10 and A.12 actually display the evolution of representative residential housing units that are featured in our sample. These are typical units behind our numbers over time and space. In the first, we exhibit a artisan building in Nuremberg, originally built in the year 1489 for newly arriving weavers from the Nuremberg region. The buildings have undergone an eventful history, but are still located at the same place in Nuremberg today, keeping the same address at the *Sieben Zeilen*. Thanks to in-depth studies by the Nuremberg Historical Association (Taschner, 2013) we have extensive detail about the precise quality adjustments and architectural details of the buildings.

The building contains two stories, with an additional roof area used for storage purposes, as well as a cellar containing work areas. Two households lived in such a building from the inception date. Per the architectural drawing displayed on the right handside of Figure A.10 – itself produced for a renovation project in the year 1912 – we know that the living space per story stands at 59 square meters (8.2 meters width, 7.2 meters depth, including staircases). Such dimensions are comparable with corresponding information from other German urban areas in our sample and beyond. Adjusting the living space for household size developments for urban middle class households yields several relevant data points for living space per capita.

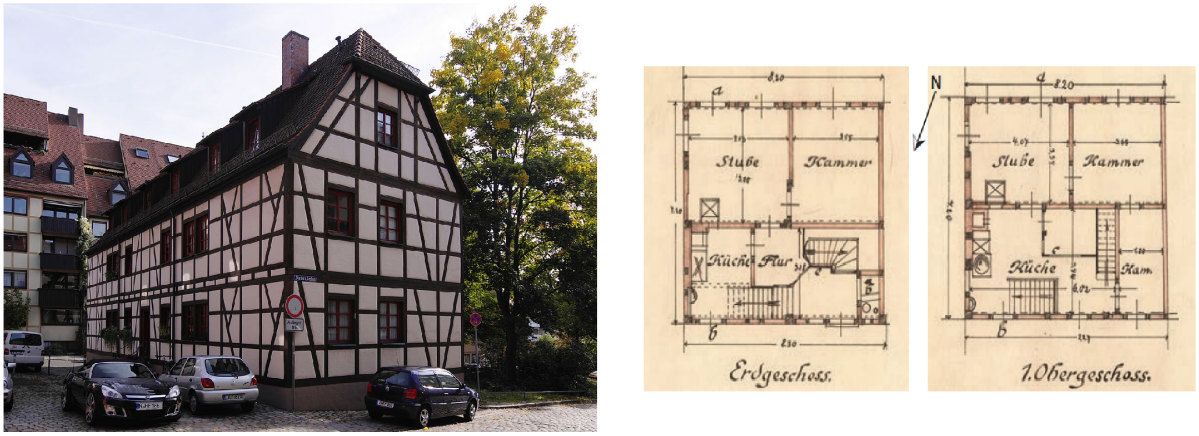
[60] Knoll et al. (2017) acknowledge the fact that their indices may insufficiently control for quality dynamics, especially in the pre-1945 era, and write that "researchers using our dataset in the future should take into account that accurate measurement of quality-adjustments remains a challenge." For Germany, the paper claims that the pre-1945 data is non-quality controlled, while from 1962, their nationwide index controls for quality changes.

Table A.6: Average or representative living space estimates, German cities

	Source	period	sqm per capita	type detail
		Early Modern Era		
Nuremberg	Dirlmeier (1978)	1488	15	representative M and LM
Augsburg	Weidenbacher (1926)	1516	16	representative LM
Lemgo	Kaspar (1985)	1400-1650	17.6	
Cologne	Eberstadt (1903)	ca. 1620	20.3	"typical three-window house"
Berlin	Voigt (1901)	1711	28.1	representative UM and M
Frankfurt		1855	25.1	representative M
Vienna	Lichtenberger (1977)	1856	16.8	weighted avg, inner city
		Modern Era		
Berlin	Berlin (1897)	1890-5	31.4	all-city average
Wolfen	Wilde (1999)	1872-1908	27.5	representative M and LM
Munich	Renaud (1904)	1903	9.3	all-city average
Giessen	Meyer (1903)	1900	20.4	representative M and LM
Frankfurt	Schuette (1930)	1930	37.0	Housing office Frankfurt
Berlin	Grundstuecks-MB	1946	11.8	all-city, census data
Berlin	Grundstuecks-MB	1996	34.8	all-city, census data
Berlin	Grundstuecks-MB	2000	37.9	all-city, census data
Berlin	Landesamt (2022)	2022	40.7	all-city, census data
Munich	SAM (2022)	2022	38.4	all-city, census data

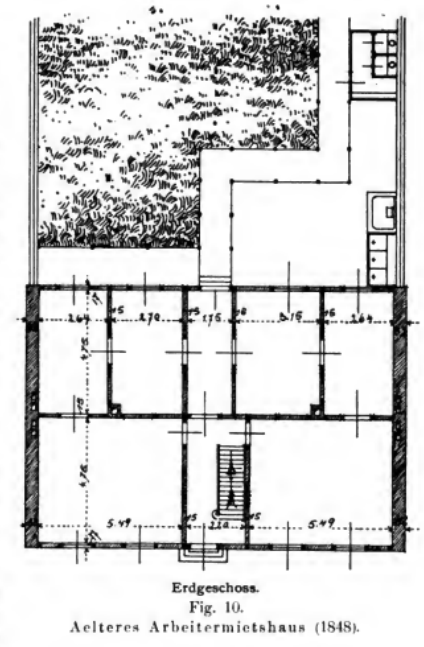
Notes: The table reports data for the residential living space per capita in contemporary and historical sources. Data prior to 1700 operates with average household size of 3.3 persons. um = upper middle class; M = middle class ; LM = lower middle class. The Vienna figure takes midpoints of the range in Lichtenberger (1977), and inner city population figure of 1834.

Figure A.10: Representative middle class residential house, built 1489, Nuremberg.



Notes: The Figure's left handside displays the exterior of a representative Nuremberg middle class home, located at the *Sieben Zeilen*; the left handside displays the corresponding architectural drawing, sourced via Taschner (2013).

Figure A.11: Representative lower middle class residential house, Rhineland, ca. 1848.

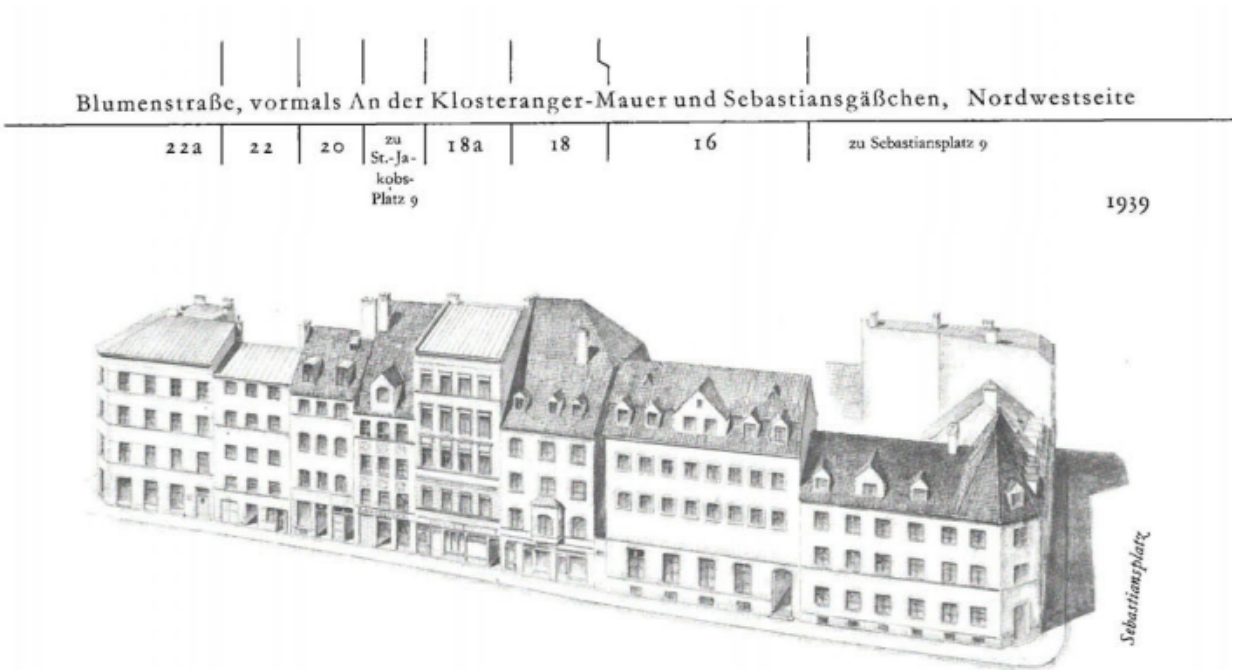


Notes: The architectural drawing from Eberstadt (1903) displays two representative middle class house types in the Rhineland, 1848. The house contains three storys, with two units per story. Displayed is the ground story, with the added backyard. The measurements per unit are 5.48 meters by 9.52 meters = 52.3 square meters of living space per unit.

Next, Figure A.12 displays a representative residential street in Munich, shown in its condition in the year 1939. We observe that by the 20th century, residential houses in larger German cities feature four to five stories, including attics. At this time, one can assume three to four households per residential unit, that is 12-20 persons in a representative residential building, per Burgmaier and Schneider (1966). These residential units again contain unit-level heating, hot water, and kitchens. In the ground floors, we often

find commercial real estate, and shared cellars are typically included below ground floors. The street length is about 80 metres, and the depth is 11 metres. Today, the buildings in the *Blumenstrasse* remain in place, and can be located at
geoportal.muenchen.de/portal/plan – coordinates 11° 34 16 | 48° 07 55.

Figure A.12: Representative middle class residential street, 1939, Munich.



Notes: The architectural drawing from displays a typical street of residential houses in Munich, shown in their condition as of 1939 (*Blumenstrasse 16-22a*). Architectural drawing by Gustav Schneider in Burgmaier and Schneider (1966).

6. Econometrics, further details

Here we undertake additional tests to analyze the main paper results with regards to robustness, especially the main regressions and time series properties of the housing and associated variables.

.1 Regression variations

In the main text, I alluded to a few simply OLS regressions for the German data, as a first cut to assess commonly invoked drivers of house prices and returns.

In Table A.7.1, we use German gross total returns (non-logged) as the dependent variable, and include a constant. We observe in the regression that "PUB-PRI" (the spread between public and private German long-maturity rates, the latter consisting of the mortgage rates) and population growth rate both correlate significantly with total gross returns. Inflation rates, real GDP growth (aggregate or per capita) and real interest rates are never significant.

Table A.7.2 presents the same simple OLS regression, now with log housing excess returns as the dependent variable. Independent variables again include real GDP growth, population growth, the inflation rate, and

Table A.7.1: OLS Regression, German gross housing total returns (TR)

	Gross TR	Gross TR	Gross TR	Gross TR
German inflation	0.0238	0.0238	-0.0236	-0.0236
	(0.34)	(0.34)	(-0.35)	(-0.35)
Real aggregate GDP growth	-0.00154		-0.0170	
	(-0.03)		(-0.38)	
Real per capita GDP growth		-0.00134		-0.0167
		(-0.03)		(-0.38)
Population growth	0.981***	0.980***	0.974***	0.957***
	(6.24)	(6.45)	(6.52)	(6.63)
German nominal rates	-0.676***	-0.676***	-0.175	-0.175
	(-7.78)	(-7.78)	(-1.52)	(-1.52)
German real rates	0.0122	0.0121	-0.0098	-0.0098
	(0.45)	(0.45)	(-0.38)	(-0.38)
PUB-PRI spread	0.217***	0.217***	0.197***	0.197***
	(5.42)	(5.42)	(5.18)	(5.18)
year			0.000093***	0.000093***
			(6.29)	(6.29)
constant	0.111***	0.111***	-0.0727*	0.0726*
	(23.10)	(23.10)	(-2.46)	(-2.46)
observations	355	355	355	355
Adj. R-squared	0.254	0.254	0.329	0.329

Notes: t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Columns 1 and 2 are not controlling for a year trend, while Columns 3 and 4 do. Real total gross return is not included as a variable given that it is mechanical. Columns 1 and 3 use German real aggregate GDP growth and Columns 2 and 4 use German real per capita GDP growth, as both measures are highly correlated and therefore cause multicollinearity issues. "PUB-PRI" spread here uses German mortgage rates for PRI basis.

Table A.7.2: OLS Regression, German log excess housing returns (er)

	log er	log er	log er	log er
German inflation	-0.0992	-0.0992	-0.0714	-0.0714
	(-1.36)	(-1.36)	(-1.25)	(-1.25)
Real aggregate GDP growth	0.0307		-0.00961	
	(0.62)		(-0.25)	
Real per capita GDP growth		0.0315		-0.00903
		(0.64)		(-0.23)
Population growth	1.216***	1.247***	0.981***	0.971***
	(7.38)	(7.87)	(7.55)	(7.74)
German real rates	-0.144***	-0.144***	-0.0895***	-0.0896***
	(-5.38)	(-5.38)	(-4.20)	(-4.20)
PUB-PRI spread	-0.0632	-0.0632	0.0198	0.0198
	(-1.58)	(-1.58)	(0.62)	(0.62)
year			0.000138***	0.000138***
			(14.86)	(14.85)
constant	0.000791	0.000788	-0.236***	-0.236***
	(-0.42)	(-0.42)	(-14.84)	(-14.84)
observations	355	355	355	355
Adj. R-squared	0.197	0.197	0.507	0.507

Notes: t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Columns 1 and 2 are not controlling for a year trend, while Columns 3 and 4 do. Real total gross return is not included as a variable given that it is mechanical. Columns 1 and 3 use German real aggregate GDP growth and Columns 2 and 4 use German real per capita GDP growth, as both measures are highly correlated and therefore cause multicollinearity issues. "PUB-PRI" spread denotes the spread between German sovereign long-maturity yields (data per Rogoff et al. (2024)) and German mortgage yields, with details described in text and further in Schmelzing (2026). All data on annual basis and for the period 1466-1910.

German real and nominal sovereign long-maturity interest rates.[61] Here, we now observe that real interest rates are strongly significant, as is (again) population growth.

.2 ADF-GLS variations

Next, we assess the time series properties of key variables that we have introduced and discussed in the main part of the paper (section 6). Table A.8 reports results for a standard ADF-GLS trend stationarity test, via Elliott et al. (1996), this time *not* assuming a time trend. Table A.9 repeats this exercise, this time assuming a time trend.

Across all these variations, we observe that German mortgage rates, as well as excess returns (PUB-PRI), are all recording clear trend stationarity, all at the 1% significance level.

[61] GDP growth and population growth data is via Pfister (2022). As in Table A.7.1, all financial variables are via Schmelzing (2026).

Table A.8: ADF-GLS, housing and credit variables, no time trend			
	no of lags	t statistic	optimal lag
German primary sources			
German mortgage rates, nominal, 1311-2022	3	-6.803	Seq, MAIC
	2	-7.662	
	1	-8.940	SIC
Real gross housing return, 1495-1910	3	-2.244	
	2	-2.375	Seq, MAIC, SIC
	1	-2.884	
Real net housing return, 1495-1910	3	-2.676	MAIC
	2	-2.832	Seq, SIC
	1	-3.429	
Real house price change, 1495-1910	3	-4.286	
	2	-4.480	Seq, MAIC, SIC
	1	-5.328	
German PUB-PRI, 1311-2022	3	-3.671	Seq, MAIC, SIC
	2	-4.393	
	1	-5.018	
ML variables, 1578-2020			
Holland RHPI	3	-10.163	
	2	-11.105	MAIC
	1	-13.865	Seq, SIC

Notes: The table applies the methodology of Elliott et al. (1996) and reports the ADF-GLS test statistic for several choices of the number of lags k (with a maximum of three lags). The regression includes a constant. The test assumes no time trend. The critical values at the 1, 5, and 10 percent significance levels are the following: -2.58 (1 percent); -1.95 (5 percent); -1.62 (10 percent). "Optimal lag" indicates the optimal number of lags according to the sequential procedure ("Seq"), the Bayesian Information Criterion (SIC), or the Modified Information Criterion (MAIC). The test rejects when the test statistic is negative and larger (in absolute value) than the critical value.

Via Table A.13, I also show the benchmark ADF-GLS series from the main part of the paper on a non-log basis (with time trend) – we observe that all trend stationarity results pass on this basis, too, with 1% significance.

.3 Half-lives, primary source basis

In the main part of the paper, we observed half-lives of approximately 1-3 years for our global total return series. Table A.12 now displays further results for our half-live exercise: here we report half-lives for the primary source series covering Germany – once more nominal and real gross total return bases. We observe a somewhat tighter range of these half-lives, standing between 1.3 and 1.9 years overall. In addition, the observation of a gradually rising adjustment speed over time is not visible on this basis.

Overall, however, half-lives and confidence intervals are all well within the general results on the global basis. Not least, the German-specific adjustment speeds for real estate can be directly compared to the

Table A.9: ADF-GLS, housing and credit variables, with time trend			
	no of lags	t statistic	optimal lag
German primary source			
German mortgage rates, nominal, 1311-2022	3	-7.835	MAIC
	2	-8.709	Seq
	1	-10.022	SIC
German PUB-PRI, 1311-2022	3	-7.019	Seq, MAIC
	2	-8.139	
	1	-8.985	SIC
ML variables			
France RHPI change	3	-6.031	
	2	-6.178	Seq, MAIC
	1	-7.158	SIC
Holland RHPI change	3	-10.061	
	2	-11.016	MAIC
	1	-13.599	Seq, SIC

Notes: The table reports the ADF-GLS test statistic for several choices of the number of lags k (with a maximum of three lags). The regression includes a constant. The test assumes a time trend. The critical values at the 1, 5, and 10 percent significance levels are the following: -2.58 (1 percent); -1.95 (5 percent); -1.62 (10 percent). "Optimal lag" indicates the optimal number of lags according to the sequential procedure ("Seq"), the Bayesian Information Criterion (SIC), or the Modified Information Criterion (MAIC). The test rejects when the test statistic is negative and larger (in absolute value) than the critical value.

Table A.12: Half-Lives of Real Rates, Germany primary source, Full Sample and Subsample				
	α	Confid. Interv. α	h	Confid. Interv. h
		Nominal total returns		
1560–2020	0.63	(0.55; 0.70)	1.36	(1.19; 1.55)
1750–2020	0.58	(0.49; 0.68)	1.30	(1.07; 1.52)
1914–2020	0.59	(0.46; 0.75)	1.32	(0.98; 1.75)
		Real total returns		
1578–2020	0.66	(0.59; 0.73)	1.90	(1.71; 2.24)
1750–2020	0.61	(0.52; 0.69)	1.62	(1.39; 1.89)
1914–2020	0.65	(0.52; 0.79)	1.57	(1.21; 2.24)

Notes: This table reports median unbiased estimates and 90 percent confidence intervals of α based on Hansen's (1999) grid-bootstrap as well as median unbiased estimates and 90 percent confidence intervals of the half-life (h) based on Steinsson (2008). The regression is $y_t = \mu_0 + \mu_1 t + \alpha y_{t-1} + \sum_{j=1}^{p} \gamma_j \Delta y_{t-j} + \varepsilon_t$, where α is the largest root. The row with the country name reports the full sample estimates, while the rows with subsamples report the subsample estimates.

German sovereign real rate adjustment speeds for the same subperiods as reported by Rogoff et al. (2024): here, the authors found adjustment speeds ranging from 2.4-5.2 years over the same horizon. Once more, it is suggested therefore that real estate total returns show faster adjustment speeds compared to sovereign real rates.

.4 Rent growth stationarity

In section 7 of the main paper, I proceeded with a present value model that assumes forecastability of income (rent) growth, leaning on related literature (e.g. LeBris et al. (2019)). In Table A.15, I demonstrate that such assumptions are plausible from an econometric point of view, given stationarity of rental growth rates in independently created historical datasets.

Independently created data sets exist for instance for Holland (Eichholtz, 1997); the U.K. (Clark, 2010); France (Couperie and Ladurie, 1970); and the U.S. (Lyons et al., 2024). Per the table, for these and all other series, ADF-GLS tests (here with time trend), on log bases, clearly confirm stationarity, typically at the 1% significance level.

Also reported is the rent growth series derived from new German primary series, here for 1466-1910 (but not sensitive to the precise period chosen), as well as the Global GW aggregation on the rent growth (log) level. We see that for these two series created anew in this paper, stationarity equally holds.

Taken together, the evidence suggests that the rent growth process exhibits similar econometric features to the dividend growth process, as for instance summarized in Koijen and van Nieuwenburgh (2011).

Table A.13: ADF-GLS, housing and credit variables, non-lagged, with time trend			
	no of lags	t statistic	optimal lag
Germany primary source, 1465-1910 (2020)			
Real gross housing return, 1465-1910	3	-3.396	
	2	-3.556	Seq, SIC, MAIC
	1	-4.262	
Real net housing return, 1465-1910	3	-4.037	
	2	-4.216	Seq, SIC, MAIC
	1	-5.027	
RHPI change, 1465-2020	3	-7.519	MAIC
	2	-8.616	Seq, SIC
	1	-11.111	
ML generated, real price y-o-y change, 1560-2020			
U.K. RHPI change	3	-3.396	
	2	-3.556	Seq, SIC, MAIC
	1	-4.262	
U.S. RHPI change	3	-4.037	
	2	-4.216	Seq, SIC, MAIC
	1	-5.027	
Global GW RHPI change	3	-7.019	Seq, MAIC
	2	-8.139	
	1	-8.985	SIC
Global AW RHPI change	3	-7.019	Seq, MAIC
	2	-8.139	
	1	-8.985	SIC
Total gross real returns, 1560-2020			
Global GW, log return	3	-3.396	
	2	-3.556	Seq, SIC, MAIC
	1	-4.262	
Global AW, log return	3	-4.037	
	2	-4.216	Seq, SIC, MAIC
	1	-5.027	
Rental yield, 1560-2020			
Global GW	3	-2.551	
	2	-2.664	
	1	-2.498	Seq, SIC, MAIC

Notes: The table reports the ADF-GLS test statistic for several choices of the number of lags k (with a maximum of three lags). The regression includes a constant. The test assumes a time trend. The critical values at the 1, 5, and 10 percent significance levels are the following: -2.58 (1 percent); -1.95 (5 percent); -1.62 (10 percent). "Optimal lag" indicates the optimal number of lags according to the sequential procedure ("Seq"), the Bayesian Information Criterion (SIC), or the Modified Information Criterion (MAIC). The test rejects when the test statistic is negative and larger (in absolute value) than the critical value.

Table A.14: Full sample of spread between actual and implied PV, by DR est.

	1560-1650	1650-1820	1820-1914	1914-1980	1980-2020	**Full P**
(1) – Bris et al. 1, IRC	-.076	-.126	-.075	-.275	.126	-.106
(2) – Bris et al. 2, Bazacle	.247	-4.292*	-.622	-.295	–	-1.975*
(3) – Mortgage rates	.155	-.153	-.223	-.029	.049	-.065
(4) – GW Sov. rates	.637	-.037	-.121	-.207	.241	.073
(5) – Campbell-Shiller v3	-.906	-.899	-.884	–	–	-.911
(6) – Campbell-Shiller v4	-.533	-.710	-.622	-.295	-.324	-.569
(7) – GMS-GK	-.102	-.220	-.126	-.485	-.151	-.211
(8) AVG actual house price, relative to PV, discount (-) premium (+)	-.028	-.830	-.229	-.241	-.012	-.427

Notes: The table reports discount rate estimates (rows 1-7, in %). The final row reports the implied discount of the actual Global GW house price value for each individual period relative to the present value model that uses the median DR estimate, as derived from the sample of DR estimates. For all sources of the DR estimates, see the text. "Full P" denotes average for entire period, 1560-2020.

Table A.15: Rent growth rate variations, log bases, ADF-GLS results			
	no of lags	t statistic	optimal lag
Germany primary source, 1465-1910			
German primary source rents, LOG, 1466-1910	3	-5.7321	Seq, MAIC
	2	-6.6466	SIC
	1	-9.2973	
Secondary source indices			
U.S. RPI, spliced, LOG, 1831-2020	3	-6.7200	
	2	-8.3971	
	1	-9.5367	SIC, MAIC
U.S. RPI, unspliced, LOG, 1891-2006	3	-4.7681	
	2	-4.9485	MAIC
	1	-6.2770	Seq, SIC
Holland, Lesger LOG	3	-4.7232	Seq, SIC, MAIC
	2	-6.3187	
	1	-7.5356	
Paris rent LOG	3	-3.1949	Seq, SIC, MAIC
	2	-5.7810	
	1	-5.7580	
U.K., Clark LOG	3	-5.4660	
	2	-5.2918	
	1	-5.1505	SIC, MAIC
Global GW			
Global GW LOG	3	-9.6381	MAIC
	2	-11.311	
	1	-12.767	SIC

Notes: The table reports the ADF-GLS test statistic for several choices of the number of lags k (with a maximum of three lags). All series are tested on the log bases. The regression includes a constant. The test assumes a time trend (see appendix for ADF-GLS results without time trend). The critical values at the 1, 5, and 10 percent significance levels are the following: -3.48 (1 percent); -2.88 (5 percent); -2.59 (10 percent). "Optimal lag" indicates the optimal number of lags according to the sequential procedure ("Seq"), the Bayesian Information Criterion (SIC), or the Modified Information Criterion (MAIC). The test rejects when the test statistic is negative and larger (in absolute value) than the critical value.