Housing *Is* the Financial Cycle: Evidence from 100 Years of Local Building Permits*

Gustavo S. Cortes† — *University of Florida*
Cameron LaPoint‡ — *Yale SOM*

June 2026

Abstract

Does the housing market lead the financial cycle, and if so, why? We address these questions by creating a new hand-collected database spanning a century of monthly building permit quantities and valuations for all U.S. states and the 60 largest MSAs. We show that the option to build embedded in permits renders volatility in residential building permit growth (BPG) a strong predictor of aggregate and cross-sectional stock and corporate bond returns and volatility. This predictability remains even after conditioning on corporate and household leverage, mortgage access, commodity price risk, and firms' exposure through their network of plants to other localized physical risks like natural disasters. Cities and states with more elastic housing supply consistently predict financial market downturns at 12-month horizons, admitting new trading strategies to hedge against overbuilding risk. A noisy rational expectations framework in which local building permits serve as a quasi-public signal for dividends explains these empirical patterns.

Keywords: Housing supply, building permits, real estate, option value, volatility, financial crises, equities, corporate bonds

JEL Classifications: E32, G01, G12, N22, R31

*For helpful suggestions, we thank Leah Boustan, Igor Cunha, Zhi Da, Slava Fos, Carola Frydman, Stefano Giglio, Will Goetzmann, Paul Goldsmith-Pinkham, Joshua Hausman, Matt Jaremski, Theis Jensen, Dmitry Kuvshinov, Tom Nicholas, Jonathan Payne, Geert Rouwenhorst, Nathan Seegert, Allison Shertzer, Alp Simsek, Alan Taylor, Christoph Trebesch, David Weinstein, and Wei Xiong; and our discussants: Darren Aiello, Alex Chinco, Sonia Gilbukh, Daniel Greenwald, Rawley Heimer, Seohee Kim, Elena Loutskina, Jonathan Rose, Chester Spatt, and Jiro Yoshida. We thank seminar participants at the Chicago Fed, Erasmus University, RIT Saunders, Trinity College Dublin, University of Luxembourg, University of Michigan, University of Southern California, Yale, and audiences at the ASU Sonoran Winter Finance Conference, Mountain West Economic History Conference, Midwest Finance Association, University of Kentucky Finance Conference, FSU–Truist Beach Conference, the Federal Reserve Monetary and Financial History Workshop, SFS Cavalcade, the EIEF Rome Junior Finance Conference, AREUEA Webinar Series, Barcelona Summer Forum, Pre-WFA Real Estate Symposium, Red Rock Finance, the Holden Conference at Indiana Kelley, Columbia–NYU–Yale Housing Day, and the NBER Summer Institute for comments and feedback. We are grateful to Bill Abriatis and Cheryl Cornish of the Economic Indicators Division of the Census Bureau for providing scanned tables from the historical Building Permits Survey and for lending us archived Census construction volumes to digitize. We thank Gwyneth Crowley at Yale University Library and Kris Abery at the Connecticut State Library for assistance in locating historical reports from the Building Permits Survey. We thank William Creech at the National Archives for answering our questions about historical Census Bureau records and BLS Building Construction reports. The Cotality Building Permits data used in this paper are licensed through Yale. This research was financially supported by the Yale International Center for Finance and the Tobin Center for Economic Policy. We thank a diligent team of RAs for their excellent work, including Will Eckert, Shawn Frazier, Luiz Ghiraldelli, Amit Kamma, Jakob Reinhardt, Fernando Reyes De La Luz, Aurel Rochell, Gabriel Thomaz Vieira, and Franziska Yost. Finally, we are grateful to Ed Leamer, whose contributions to economic forecasting inspired our work. First draft: May 2024.

†Warrington College of Business, University of Florida. 306 Stuzin Hall, PO Box 117168, Gainesville, FL 32611–7168. Email: **gustavo.cortes@warrington.ufl.edu**. Web: **sites.google.com/site/cortesgustavos**.

‡Yale School of Management. 165 Whitney Avenue, New Haven, CT 06511. Email: **cameron.lapoint@yale.edu**. Web: **cameronlapoint.com**.

"My goal is to provide unforgettable images that leave a lasting impression regarding the importance of housing to what we call the business cycle."
— Ed Leamer (2007). "Housing Is the Business Cycle,"
2007 Jackson Hole Economic Symposium.

1 Introduction

Economists have long sought to understand the relationship between housing markets and the economy. Almost a century ago, Long (1939b) wrote that *"the building industry is probably the most strategic single factor in making or breaking booms and depressions."* From the Great Depression of the 1930s to the Global Financial Crisis in 2008, fluctuations in the housing market often precede (Moore, 1961) and even predict significant economic downturns (Green, 1997; Leamer, 2007, 2015). This idea is reinforced by the observation of "twin bubbles" in asset markets, where real estate prices and sales volume reach their apex several months before similar peaks occur in the stock market during economic expansions.[1] The recurring pattern of housing market distress preceding broader economic turmoil presents a compelling challenge for economists and policymakers. Yet, our understanding of these macro-financial links remains limited, primarily due to a persistent shortage of granular, long-term data tracking *local* housing market dynamics.[2] This data gap has forced researchers to rely heavily on aggregate time series evidence, leaving relatively unexplored how local housing market conditions ripple through financial markets and affect the real economy.

In this paper, we provide a comprehensive examination of the links between *local* housing markets and financial market conditions. By constructing a new monthly dataset of U.S. historical local building permits from 1919 to 2019, we investigate whether housing cycles predict the financial cycle. Applying recent advances in optical character recognition (OCR) and deep learning, we digitize a wealth of archival records from Census surveys and industry reports to create continuous time series of residential building permit activity for all 50 states and the 60 largest metropolitan areas. Our geographically disaggregated data on permits allow us to examine the predictive power of housing market fluctuations for local economic activity,

[1] Pronounced examples of the "twin bubbles" time series phenomenon include the 1920s Florida Land Boom, which preceded the Great Depression (Knowlton, 2020; Calomiris and Jaremski, 2023); 1980s Japan — where Tokyo area land prices flat-lined in 1987 before the Nikkei crash in January 1990 (LaPoint, 2021); the 2000s U.S., where cities with elastic housing supply like Las Vegas experienced downturns in 2006 before Lehman's 2008 fall (Nathanson and Zwick, 2018); and more recently with Chinese residential property prices dropping in mid-2014 before the Shanghai Stock Exchange crash of June 2015 (Liu and Xiong, 2018).

[2] Earlier attempts to highlight boom-bust cycles in the U.S. construction sector focused on residential permit series aggregated from different sets of cities at annual frequency. See Riggleman (1932), Long (1939a), and Isard (1942) for early examples. An exception is Long (1936), who decomposes a permit series for Manhattan into residential and non-residential buildings. Colean and Newcomb (1952) and Abramovitz (1964) highlight that these early indices suffer from sample selection bias due to non-comprehensive geographic coverage.

FIGURE 1. Total Real Value of U.S. Monthly Building Permits Issued, 1919–2019

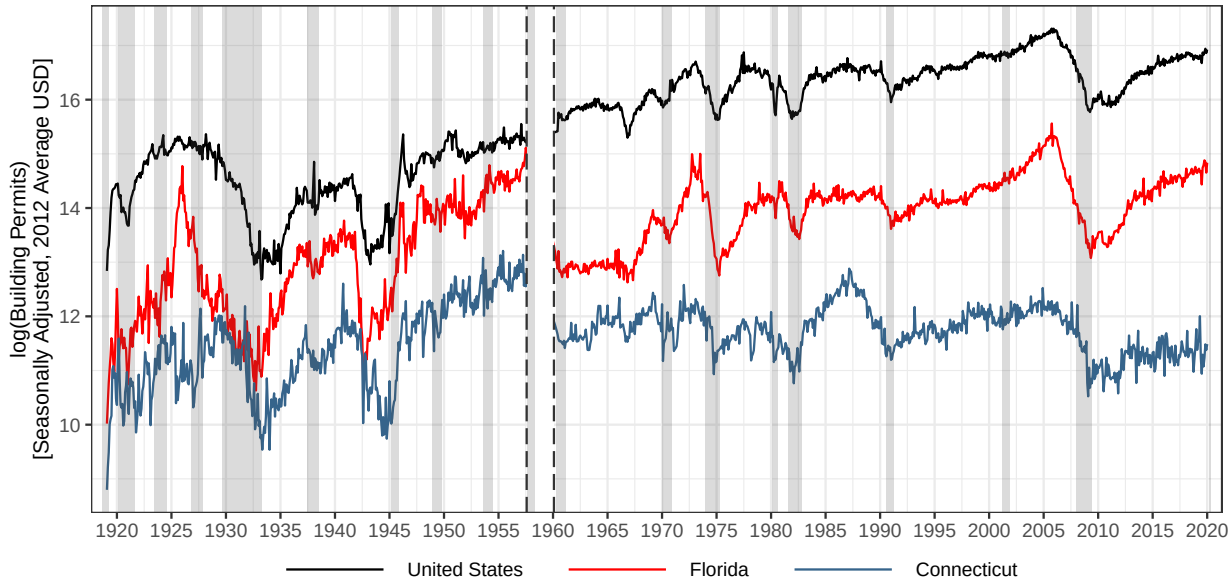

Notes: The figure plots the monthly total log real value of building permits issued in the U.S. and separately for Florida and Connecticut from 1919M1 to 2019M12. The vertical dashed lines indicate the gap between the two sources used in our full sample. Data for the first half of the sample is from *Dun & Bradstreet's*, and the second half is from the U.S. Census Bureau. The gap in data availability between 1958 and 1959 reflects the discontinuation of the relevant *Dun & Bradstreet's* publication. We deflate permit values to real 2012 average USD using the monthly historical CPI from Shiller (2015) before seasonally adjusting each series. We rank states from least to most housing supply elastic due to regulation in terms of court cases filed mentioning the phrase "land use," as compiled by Ganong and Shoag (2017). Connecticut ranks first in land-use regulatory stringency on this measure, with eight times as many average annual land-use court cases per capita as Florida between 1940 and 2000. Grey-shaded areas indicate NBER-dated recessions. See Appendix B for details on how we perform the seasonal adjustment and harmonize data sources over the full hundred-year period.

aggregate stock and corporate bond market volatility, and the volatility of individual firms' securities over a period spanning 20 recession episodes.

As illustrated by the century of U.S. real total building permit values plotted in Figure 1, our main finding is that the volatility of local building permit growth (BPG) is a strong predictor of future asset market volatility, even after controlling for a wide range of factors such as corporate and household leverage ratios, natural disaster risk, population growth, and other macroeconomic and financial fundamentals like industrial production and the VIX (or NVIX of Manela and Moreira, 2017). We consistently observe this pattern over time, showing that building permits are not just a harbinger of particular types of financial crises, such as the subprime mortgage episode in the years leading up to the Great Recession. Moreover, we document that this predictability holds not only for the stock market but also the corporate bond market, indicating that the housing market's influence extends across asset classes.

The richness of our permits data allows us to geographically disaggregate the predictability of housing market volatility. A handful of geographic submarkets drives this volatility; the predictive power of BPG volatility is particularly strong and positive for more housing supply elastic cities and states in the South (e.g., Florida and Georgia), which lead the stock and

bond markets by longer horizons relative to supply-inelastic regions like New England. For instance, for every 10 percentage point increase in cumulative BPG volatility over the prior 12 months in Florida, stock return volatility spikes by 0.2 percentage points, and bond return volatility spikes by 0.7 percentage points. In contrast, BPG volatility has no clear predictive power for stock return volatility over the same 12-month horizon in states with restrictive policies towards new residential development (e.g., Connecticut). Figure 1 shows that while housing supply elastic and inelastic areas largely comove, boom-busts in building permits are more pronounced in elastic areas like Florida around crisis periods.

To further illustrate the type of soft information building permit volatility conveys to investors in financial markets, we zoom in on the Global Financial Crisis. We observe that single-family home BPG volatility has stronger predictability in states and metros with a large share of subprime MBS loans on the eve of the crisis (Mayer and Pence, 2008). Since housing markets co-move, the signals offered by BPG volatility may be collinear across locations. We therefore conduct a principal components analysis (PCA). The first component alone delivers a 17% R^2 for stock return volatility and a 20% R^2 for bond return volatility. Because jumps in the first component anticipate well-known events in the narrative of the Global Financial Crisis, we label the first component as the "subprime" factor. These jumps include the failure of Bear Stearns, the Lehman Brothers bankruptcy, and the peak volume of foreclosure auction sales in the Summer of 2010. Interestingly, the subprime factor only emerges when we conduct PCA in a time window around the GFC. When we use the full post-1960s time series, the first principal component is instead dominated by input supply frictions around the oil crises of the 1970s. Hence, the subprime factor we identify is distinct from a more general leverage build-up factor.

Although we emphasize that building permits predict return volatility across many episodes, using recent decades for case studies helps us rule out confounding mechanisms by merging in available corporate data. To gain further insights at more granular levels of observation, we match listed firms to a registry of their plant locations to construct measures of their physical exposure to BPG volatility weighted by their plants' sales or employment. For every 10 percentage point increase in cumulative BPG volatility exposure, a firm's stock return volatility increases by 0.1 percentage points, even after controlling for standard balance sheet variables like EBITDA, Tobin's Q, size, age, and leverage. Predictability in the cross-section of equities gives rise to the possibility of using BPG volatility as a factor to design trading strategies hedging against the geographic exposure of firms' operations to real estate markets which face a glut of new residential development.

We next investigate the mechanisms through which volatility in building permits predicts the financial cycle. Building permits are especially informative because they sit at the intersection of housing supply decisions and beliefs about local fundamentals. Developers exercise the option to build when expected conditions are favorable and delay or abandon projects when prospects deteriorate. The "skyscraper wave" in New York City around the Great Depression, when prominent projects were delayed, scaled down, or never completed, offers one historical

illustration (Barr, 2010; Nicholas and Scherbina, 2013; Cortes and Weidenmier, 2019). We introduce a simple real-option model of housing development embedded in a noisy rational expectations framework à la Grossman and Stiglitz (1980), which we use to rationalize four features of the evidence: (*i*) permits proxy for local fundamentals; (*ii*) permit movements predict aggregate asset markets; (*iii*) the sign and magnitude of BPG-volatility loadings vary across places; and (*iv*) predictability is strongest where physical and regulatory constraints are weaker.

We extend Grossman and Stiglitz's (1980) framework along two dimensions. First, we introduce a spatial component by allowing traders in the unitary financial asset market to receive potentially different signals about dividends depending on their locality. Second, we introduce physical supply constraints into the model via a cap on the number of permits that can be filed in a locality, which is independent of construction costs; physical constraints on new construction consist of mountainous terrain and water or wetlands coverage (Saiz, 2010; Lutz and Sand, 2023). We distinguish physical constraints from regulatory constraints (Gyourko et al., 2008, 2021; Bartik et al., 2024), which we model as a profit wedge that raises construction costs, rendering development more expensive in already built-up housing markets (Favilukis et al., 2023). The presence of housing supply constraints introduces differences across locations in the informativeness of building permits as a signal to financial markets investors. Importantly, the existence of land use regulations alone cannot explain the cross-sectional patterns we observe in the longitudinal data given that many restrictions, other than minimum lot size and zoning rules (Shertzer et al., 2022), were passed by homeowners in the post-1970s period, in part, to preserve rising home equity values (Fischel, 2005).

Our paper makes several key contributions to the literature on the real estate sector and its links with the macroeconomy. First, we provide new evidence from a century of geographically disaggregated data in favor of the hypothesis that housing explains a large portion of the financial cycle. Second, we introduce a new monthly factor, BPG volatility, for forecasting stock and bond returns, which is orthogonal to other physical sources of risk at the firm level. Finally, our newly constructed dataset offers a valuable resource for future research on local housing supply and macroprudential housing policy.

We extend existing housing market research by constructing a comprehensive, long-run dataset of U.S. building permits. Unlike previous studies that focused on house price indices at the country level or individual cities (e.g., Knoll et al., 2017; Korevaar, 2023), our data provide coverage across both granular geographies and time. We expand on the work of Stock and Watson (2010), who digitized state-level historical Census Building Permits Survey (BPS) data from 1969–2007, by extending the BPS data back to its 1959 inception for both states and MSAs, as well as for single-family and multifamily units. More importantly, we splice the complete BPS series with *Dun & Bradstreet's* records dating back to 1919, accounting for changes in geographic boundaries over time.

Combined with recent work on historical housing price indices (Lyons et al., 2026) and construction sector productivity (D'Amico et al., 2024), our complete dataset offers new facts about the evolution of U.S. housing supply over the twentieth century. In most states, single-family permitting per capita peaked in the early 1970s and plummeted after the 2008 Financial Crisis. We show using panel microdata from Cotality that over 80% of permits issued for new residential units since 1990 are ultimately exercised. The microdata show large differences in completion probabilities across states depending on the rigidity of local planning laws. The high completion rates of residential permits — especially for single-family units — mean that swings in permits translate into shifts in the housing supply curve. However, completion rates still well below 100% indicate that building permits convey additional information about investors' beliefs beyond simply serving as a housing market indicator that leads to new building starts. We apply this insight to understand the recent house price cycle in key cities around the COVID-19 pandemic. We show that declines in permits precede housing market corrections, as measured by sharp declines in prices and rents in 2022–2023 for cities like Austin, Texas, which experienced a large influx of work-from-home migrants during the initial COVID-19 period of 2020–2021 (Hansen et al., 2023).[3]

We also provide novel evidence supporting the hypothesis that housing explains a significant portion of the financial cycle. An important innovation is introducing a new monthly factor, BPG volatility, for forecasting stock returns and return volatility. This factor is orthogonal to other physical sources of risk at the firm level, offering a new perspective on the relationship between real estate and financial markets. While previous research using modern Census BPS data has primarily focused on forecasting indicators of the real economy, our work bridges the gap to financial forecasting, expanding on studies like Ludvigson and Ng (2009) which explore macro factors' predictive power for government bond returns.

We contribute to understanding the role of housing markets in financial crises and recessions (Goetzmann and Newman, 2010; Brocker and Hanes, 2014; Fishback and Kollmann, 2014; Gjerstad and Smith, 2014; White, 2014; Cortes and Weidenmier, 2019). We show that the extent to which permits capture macroeconomic risk varies depending on pre-existing local housing supply frictions, complementing recent evidence on how real estate markets transmit capital flow shocks to local economic outcomes (Bednarek et al., 2021). We thus offer a theoretical foundation for the heterogeneity in permit predictability observed in the geographic cross-section. Our approach provides alternatives to traditional leverage and credit growth explanations (Schularick and Taylor, 2012; Jordà et al., 2013; Greenwood et al., 2022; Müller and Verner, 2024), demonstrating that fluctuations in local housing markets help link "Main Street" to "Wall Street" by providing informed, rational investors with a signal about local economic activity. This perspective contrasts with studies focused on non-rational beliefs (Shiller,

[3] Austin received substantial news coverage as one of the most prominent boom-and-bust house price cases (see, e.g., *"Once America's Hottest Housing Market, Austin Is Running in Reverse,"* The Wall Street Journal, March 18, 2024).

1981, 2015) or over-optimism (Baron and Xiong, 2017), highlighting building permits as a forward-looking indicator in the real estate market.

Beyond improvements to the measurement of links between real estate and financial markets, our work helps resolve the "co-movement" puzzle. Kuvshinov (2025) finds using rent-price ratios in the Jordà et al. (2019) database covering 17 countries over 150 years that cross-asset predictability is limited, even after accounting for risk factors like bank leverage and real credit. This lack of co-movement arises despite a robust prediction of macro-finance theories that housing, equity, and corporate bond markets should positively co-move (e.g., Lettau and Ludvigson, 2001; Piazzesi et al., 2007). Fama and French (2023) argue that the predictability of house prices for future rents can be improved by across-area demeaning, because local housing markets co-move. We propose a partial resolution to the co-movement puzzle by honing in on a segment of the real estate market — permits for new construction — which is more forward looking than house prices and transaction volume, and which is priced more like a financial asset than pre-existing structures for which hedonic demand characteristics matter more.

The rest of the paper proceeds as follows. Section 2 describes the construction of our longitudinal permit database and presents new stylized facts. Section 3 lays out the empirical designs. Section 4 reports the main longitudinal and historical results. Section 5 discusses the mechanisms and alternative interpretations. Section 6 develops a rational disagreement theory to shed light on our empirical findings. Section 7 concludes.

2 Data and Motivating Facts

In this section, we describe how we combine several sources of historical building permit counts, valuations, and residential price indices to construct forward-looking measures of local housing market activity that we later interpret through the framework in Section 6. Our resulting database spans over a century (1919–present) for the entire U.S., all 50 states, and 60 MSAs with continuous *monthly* time coverage.

2.1 Constructing Long-Run Building Permit Data Series

We combine hand collection and deep learning OCR techniques to digitize the data and create a continuous time series of building permit activity at the state and metropolitan area levels. We provide further details on data limitations, including the fact that land values are not capitalized into the permit series, and assumptions needed to splice the series in Appendix B.2; discuss the technicalities of our OCR data collection in Appendix B.3; and describe the seasonal adjustment methods in Appendix B.4.

1919–1957 Period. We combine several data sources to build our long-term monthly panel of building permit values. For the earlier decades, we follow Cortes and Weidenmier (2019)

and use historical building permit data reported in issues of *Dun & Bradstreet's Review* — renamed in 1937 to *Dun's Statistical Review*. The real-time data are assembled from building inspector reports. This allows us to collect the data for a growing number of cities between 1919M1 and 1957M10.[4] The first volume in 1919 had 164 cities, which remained constant until 1923. After further expansions in the mid-1920s, the sample stabilized at 215 cities included within each annual volume starting in 1927.

This expansion in the number of cities reporting building permit data reflects the growing comprehensiveness and scope of the *Dun's* publications in capturing the dynamics of urban development during this period. In the end, we obtain a balanced panel of 21 cities from 1919M1 and 65 cities with a complete time series starting from 1927M1. Appendix Figure B.1 shows the number of cities with permit valuations in the *Dun's* data over time.

To ensure the accuracy and reliability of the data in this earlier period, we cross-validate our *Dun's* series with annual data from the Bureau of Labor Statistics (BLS) Construction Reports, which cover construction occurring over most years between 1920 and 1953. The BLS reports contain data on the annual total number and valuation of permits, as well as construction costs, broken down by new residential and non-residential buildings vs. permits for alterations and repairs. Tables report information covering 1,790 cities from 1940 onward. Survey methodology used by the BLS informed the later monthly surveys formally taken over from BLS by the Census Bureau in July 1959 (Census Bureau, 1959). The BLS reports, therefore, provide a valuable check on the consistency and quality of our primary data sources prior to Census coverage.

1957–1959 Period: Data Limitations. Dun & Bradstreet ceased publication of *Dun's Statistical Review* in 1957 before the Census Survey of Construction and Building Permits Surveys began to be regularly published. For the short gap covering November 1957 up to and including April 1959, we only observe building permit data for cities in New York. The State of New York published the *Quarterly Summary of Business Statistics*, obtained from HathiTrust. We use the richness of our long-run time series to extrapolate from New York State's data, applying VARMA models to interpolate the missing data for locations outside New York (cf. Appendix B).

1959–1987 Period: Historical Census Building Permit Survey. The most comprehensive source of information on U.S. local building permit activity for newly constructed buildings is the Census Bureau's Building Permit Survey (BPS). In conjunction with the Department of Commerce, the Census has continuously administered and reported data from the BPS each month starting in May 1959. For each level of geography — State, MSA, and "Place" (i.e., a county, town, or village) — the BPS includes monthly tables consisting of privately-commissioned permit

[4]We begin our sample in 1919 rather than earlier years, as the *Dun's* publications reported permit counts sporadically and for a smaller set of series before 1919.

quantities and valuations broken down by units in single-family houses, two-family buildings, three-and-four-family buildings, and five-or-more family buildings.[5]

While the modern BPS data are digitized and readily downloadable, there is no repository for data or tables in the underlying reports prior to 1988. We obtained a subset of the permit tables for states and MSAs directly from the Economic Indicators Division of the Census Bureau. To source the place tables and remaining months for the MSAs and states, we either downloaded the reports from HathiTrust or contacted the network of regional Federal Depository Libraries (FDLs). In some cases, scanned tables from HathiTrust were too deprecated to apply OCR techniques, and we thus commissioned fresh, high-quality scans of the tables from the Connecticut Federal Depository Library to limit measurement error in transferring text to data.

We focus on locations within the 48 contiguous states in our analysis, as permitting activity is generally too lumpy in smaller jurisdictions for us to obtain stable estimates from the GARCH models we adopt in Section 3 to extract building permit growth (BPG) volatility. Part of this may be due to measurement error and non-response error from the BPS sampling frame, which requires the Census to provide imputed estimates that can greatly differ from reported numbers. Cities in small states like Alaska, Hawaii, Montana, and Nevada are also not included in the earlier portion of the *Dun's* sample.

When the survey started in 1959, there were no separate series with permits aggregated by state and MSA. The Census later added in May 1960 a new set of reports to the BPS containing tables by state and MSA. Before May 1960, we obtained counts at the state and MSA level by aggregating from the place level. We describe the steps involved in this aggregation procedure and provide a taxonomy for the complete set of series available in the BPS reports in Appendix B. In the initial 1959 survey year, the Census surveyed over 7,300 places spanning 174 MSAs and provided totals for 42 MSAs; this expanded to 61 MSAs starting in January 1964. To create a geographically harmonized series, we aggregate numbers from the place tables to adjust for changes over time in MSA definitions.

1988–2019 Period: Modern Census Building Permit Survey. We download the modern data period starting in 1988M1 for the sub-national level from the raw master text file available through the U.S. Census Bureau website. Updates to the BPS series can be automatically downloaded with a one-month lag via the Federal Reserve Bank of St. Louis's FRED API. We opt to end our estimation sample at 2019M12 to avoid the COVID-19 crisis due to the unique mismatch between housing demand and supply it created, alongside lockdown and public health restrictions which may have led to deferred construction (Ghent et al., 2024). The number of MSAs with permit series in the modern BPS stabilized in 2008, with data for 113 MSAs currently downloadable.

[5]The BPS stopped listing separate numbers for public permits with the modern version of the survey in 1988, and in the years before 1988 moved to reporting them only quarterly. We exclude permits associated with public contracts from our analysis as these do not represent arms-length market transactions.

The Census Bureau relies on imputations using historical survey response rates and population proportionality-based estimates in periods where data from certain places are unavailable.[6] In the modern survey, the Census provides both the raw survey variables and versions obtained via imputing missing information. We use the imputed series from the modern period to match conventions of earlier years of BPS data for which we only have the counts inclusive of any imputation procedures conducted by the Census.[7]

2.2 Stylized Facts about Historical Local Residential Permits

We plot in Figure 2 seasonally adjusted permit counts per 10,000 inhabitants for total residential units and single-family units for all 50 states covering the entire span of the Building Permit Survey. Before scaling by population, we seasonally adjust each building permits series using the Census X-13 ARIMA–SEATS filter, described in further detail in Appendix B.4. We create similar figures for our seasonally adjusted *Dun's Review* state-level permit value series in Appendix B.1 and plot the raw Census permit counts in Appendix B.2. Scaling permits by population offers a measure of planned new residential units relative to housing demand in an area.

Several stylized facts emerge from plotting our newly collected permit data, for which we provide formal statistical evidence in subsequent analyses. First, permits are procyclical and lead recessions; this is true both in aggregate (as argued by Leamer, 2007; Stock and Watson, 2010), and for particular geographic areas. For example, Florida permits peak 5 months before the 1973 OPEC recession and lead the Great Recession by almost two years. Our main results in Section 4 highlight why certain states like Florida predict stock and bond return volatility across multiple crash episodes.

Second, in most states single-family permit counts per capita peaked in the 1970s. In sunbelt states permitting continued on trend after the oil crises of the 1970s to meet demand from new residents. The data also displays a boom in permitting in New Jersey and New York in the 1980s under Edward Koch's mayorship of New York City, during which zoning restrictions were relaxed in dense parts of Manhattan in response to the "construction drought of the 1970s."[8] Somerville (2002) finds in Canadian permits data from 1972–1997 that 95% of permits are exercised within three months and 99% of started single-family construction is completed within 15 months. We update these numbers for the U.S. by isolating permits filed for new residential construction in the Cotality (formerly, CoreLogic) *Building Permits* data. In every year since 1990, over 80% of started residential buildings have been completed within 12 months of

[6] Imputations predominantly impact small places accounting for less than 1% of national permitting activity. For this reason, the place tables are annual for all places in the U.S. and monthly for a subsample of roughly half of all places. Under the current sampling scheme, the "certainty sample" of places with counts for each month consists of places with an average of at least 6 or more permits in the last 3 years (Census Bureau, 2022).

[7] The Census BPS series from the FRED API also include imputed permits.

[8] See the NYC government's history of the Midtown Development Project: https://www.nyc.gov/assets/planning/download/pdf/about/city-planning-history/midtown_zoning.pdf.

FIGURE 2. Building Permits Per 10,000 Inhabitants in U.S. States, 1961–2019

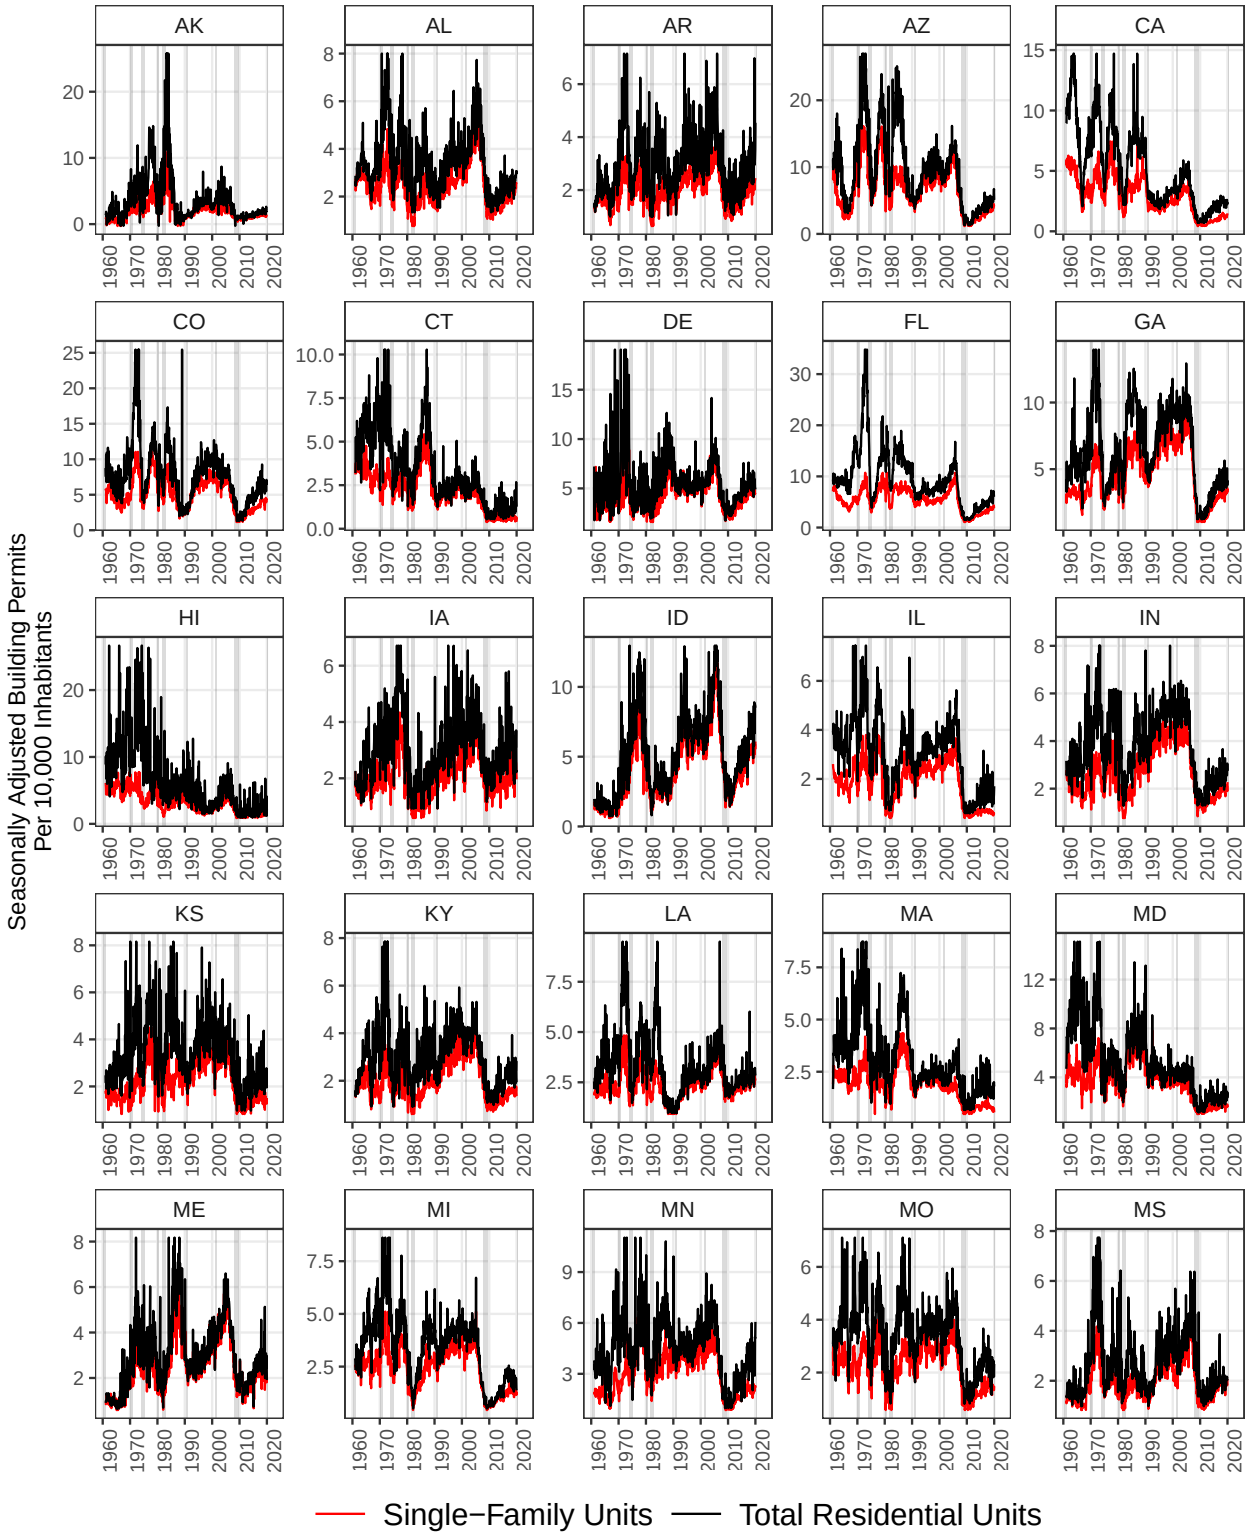

Notes: The figure plots for each state the seasonally adjusted number of total private residential building permits (black) and private single-family home permits (red) per 10,000 inhabitants filed in that state for a given year according to the Census Building Permits Survey (1961–2019). Permit counts exclude those filed for public construction contracts. Series are seasonally adjusted using the Census X-13 ARIMA-SEATS filter (cf. Appendix B.4). We start the sample at 1961M1, since 1961 is the first year the survey follows its standard format. Grey-shaded areas indicate NBER-dated recessions.

FIGURE 2. Building Permits Per 10,000 Inhabitants in U.S. States, 1961–2019 (Continued)

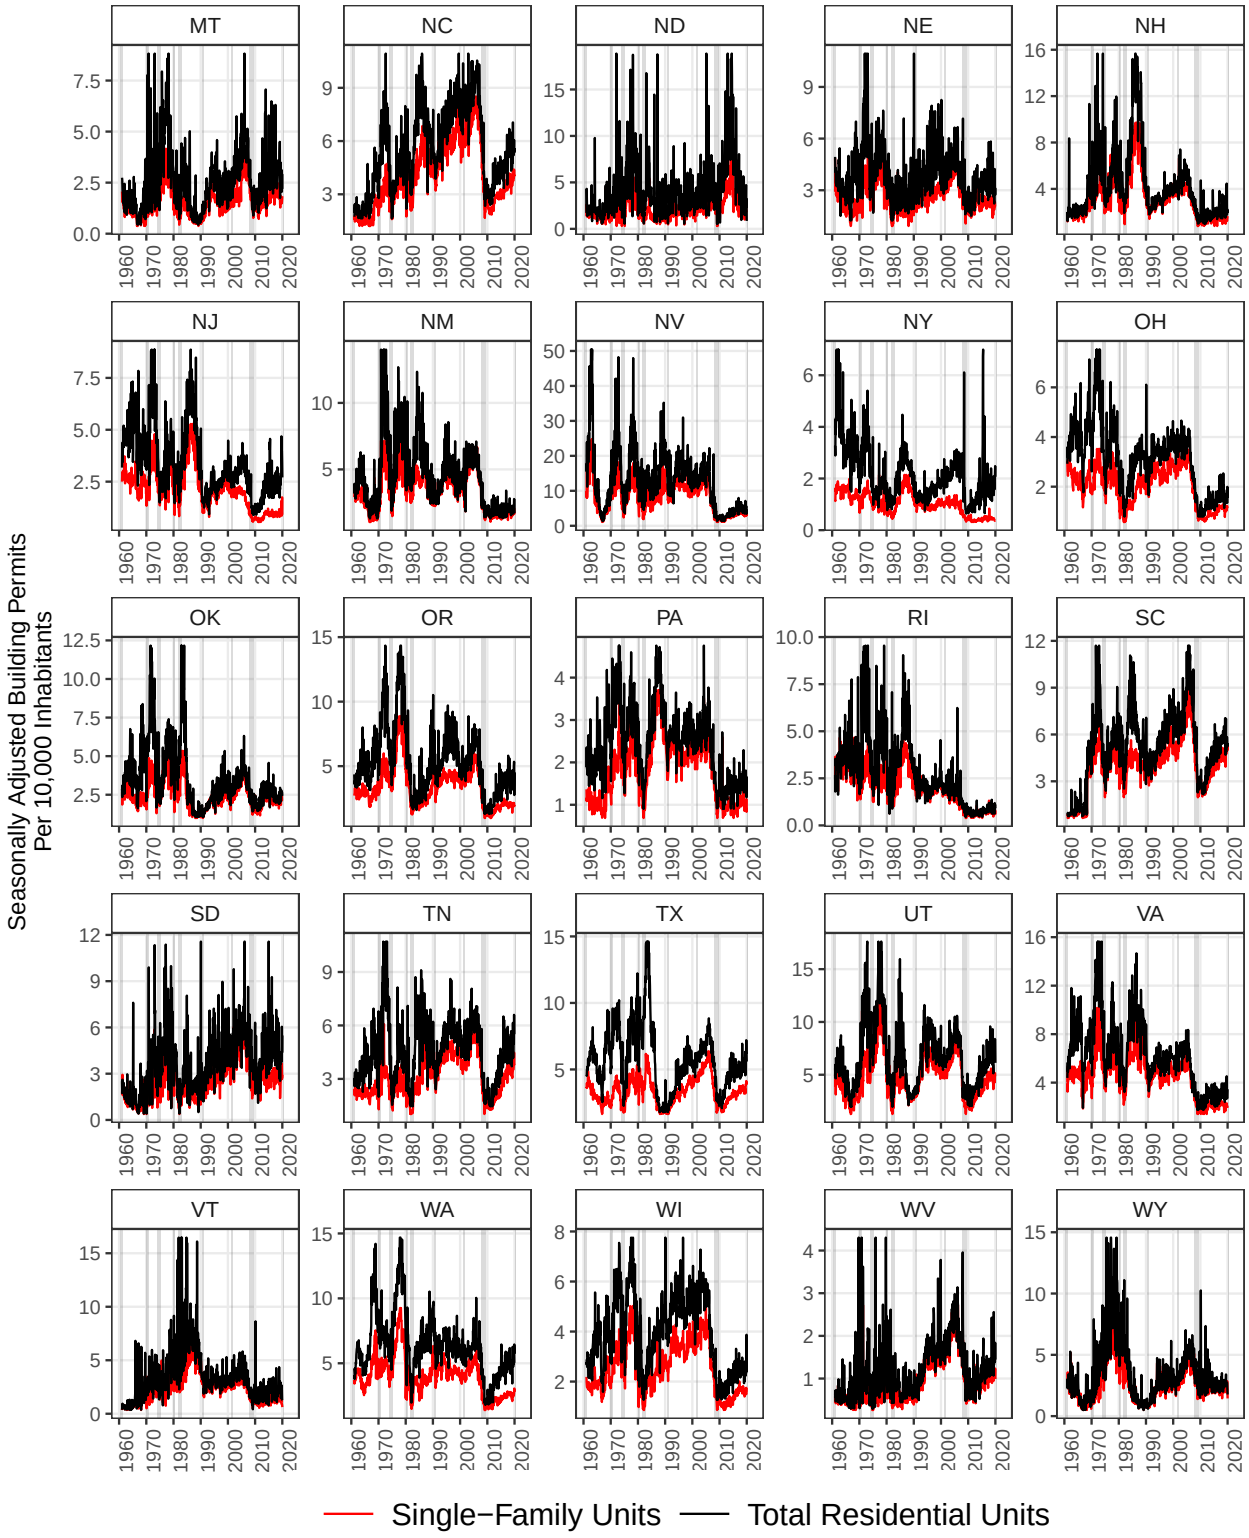

Notes: The figure plots for each state the seasonally adjusted number of total private residential building permits (black) and private single-family home permits (red) per 10,000 inhabitants filed in that state for a given year according to the Census Building Permits Survey (1961–2019). Permit counts exclude those filed for public construction contracts. Series are seasonally adjusted using the Census X-13 ARIMA-SEATS filter (cf. Appendix B.4). We start the sample at 1961M1, since 1961 is the first year the survey follows its standard format. Grey-shaded areas indicate NBER-dated recessions.

the permit being issued, with an average unconditional completion rate of 81% over that time. Permits for new residential units are therefore a valid proxy for new housing supply at one to two-year horizons. In Appendix E, we describe the Cotality data and report statistics about the relationship between permitting, time-to-build, and housing supply.

Third, there has been a clear collapse in single-family permitting activity since the GFC, which has only partially rebounded in some states. This collapse and non-recovery is more pronounced in states with stringent regulatory restrictions on new housing development, such as the high minimum lot size requirements in New England (Bronin, 2023; Song, 2025). Indeed, we find in Section 5 that the Wharton Residential Land Use Regulatory Index alone produces a 22% R^2 at the state level and 36% R^2 at the MSA level, and has a negative correlation with local permits issued in the post-GFC period. Taken together, these facts demonstrate geographically heterogeneous changes over time in new housing supply, and, through the lens of our model in Section 6, heterogeneity in the usefulness of building permits as a signal to financial investors about the macro-financial outlook.

2.3 Background on Local Permitting Process

Builders of new residential units in most U.S. jurisdictions are required to receive permit approval and pay a permit filing fee before construction can begin. Records of permits are retained by a local Building Department, City Planning Office, Zoning and Code Enforcement Office, or similar entity. Many of the earliest codes governing new buildings in the U.S. date back to the colonial era and are predicated on notions of preventing fire hazards or limited damages to neighboring properties (Oster and Quigley, 1977). Data on features of historical building codes are scarce, but President Lyndon B. Johnson's National Commission on Urban Problems review of land use regulation stated that, as of 1968, over half of all zoning decisions in Connecticut were handed down since 1958 (American Society of Planning Officials, 1968). This increasing sophistication of building codes coincides with the expansion of local governments' legal abilities to regulate building under the police power clause of the Tenth Amendment.[9]

More recently, local governments levy permit fees as a substitute for revenues from property taxes (Altshuler and Gomez-Ibanez, 1993). Levies on new construction — including permit, inspection, and certificate of occupancy fees — are referred to as "impact fees," because governments tie the fee schedule to the expected strain new development will place on the jurisdiction's ability to provide local public goods (e.g., traffic on roads might increase with new housing). Permit fees in the modern era thus take the form of a Pigouvian tax.

Horton et al. (2024) map the prevalence and burdens imposed by permit fees on new construction in recent years, documenting that virtually all counties with available permits data

[9] Glaeser et al. (2005) point to several watershed court cases, such as the U.S. Supreme Court case *Nectow v. City of Cambridge* (1928) and *Southern Burlington County N.A.A.C.P. v. Mount Laurel Township* (1975), which prescribed limited avenues for landowners to prevent local governments from restricting new development on their land.

levy fees on permits for new single-family units, yet in the median U.S. county, the amount is under 1% of forecasted construction costs. Hence, there are transaction costs to permitting, which are relatively small for the majority of areas, but there is some variation in the cross-section, with fees above 10% of project costs in the top decile of counties. Transaction costs prevent indiscriminate permitting that would attenuate the informativeness of building permits as a signal to informed traders in our model of financial markets.

2.4 Defining Building Permit Growth

Using our database of permits, we construct empirical analogs to aggregate permit valuations $V_{s,t}$, aggregate permit counts $Q_{s,t}$, and permit growth rates $q_{s,t}$ that later map into the model objects in Section 6. We focus on building permit growth (BPG) rather that levels due to the non-stationarity of local building permits, confirmed by previous studies using state-level permits (Stock and Watson, 2010; Strauss, 2013) but also by our own unit root tests applied to longer time series and finer geographies. We define log local building permit growth (BPG) based on permit values as:

$$v_{s,t+1} = \Delta \log(V_{s,t+1}) \tag{2.1}$$

$$V_{s,t} = P_{s,t} \times Q_{s,t} = \sum_{i=1}^{N} p_{i,s,t} \tag{2.2}$$

where building permit value $V_{s,t}$ depends on both the total quantity $Q_{s,t}$ of permits issued at time t in geography s and how each permit i is valued at a price $p_{i,s,t}$ by investors at the time it gets approved by the local government. $P_{s,t}$ is an index capturing average value per permit. Permit quantities depend on several factors, including demand for new properties and local supply-side factors such as the availability of developable land and the stringency of land use regulations. Our choice of geographic unit s is dictated by data availability over multiple boom-bust cycles for each test of the four empirical predictions of the model in Section 6.

The fact that land values are not capitalized into extant building permit series is a limitation of the data. While simple, the option value theory model in Section 6 illustrates that key information about macroeconomic fundamentals may be occluded by excluding land from the value of exercising a permit. There are several ways to value permits both at the level of the property and local market. Valuation methods vary depending on how the data are collected. Most permit time series attach a value based on versions of the replacement cost method. In *Dun & Bradstreet's* permit data, the valuations include additions, alterations, and repairs, but not land prices (see, e.g., *Dun's Review*, September 1935).

Modern Census data from the Survey of Construction or Building Permits Survey elicit similar notions of permit value to *Dun's*, with two main caveats. One is that there are 12 non-disclosure states that outlaw direct reporting of dollar values tied to real estate transactions,

including permits (Wall Street Journal, 2019).[10] Another issue is that not all permits reported by a local government authority pertain to entirely new builds on vacant land. For example, some permits might refer to major property remodels where the budget could change over the project timeframe due to input cost variability.[11]

We combat these issues with measuring BPG in valuation terms, $v_{s,t+1}$ by merging in other series capturing housing values. The Census disaggregates permits series into total quantities and valuations, meaning that we can observe $Q_{s,t}$ for all geographic levels beginning from 1959. We can therefore define BPG purely in terms of changes in quantities via:

$$q_{s,t+1} = \Delta \log(Q_{s,t+1}) \tag{2.3}$$

For all states via the Zillow Home Value Index from 2000 onward and for 20 Case-Shiller MSAs from 1990 onward, we can inflate up $Q_{s,t}$ using the relevant price index $P_{s,t}$ to obtain $V_{s,t}$ and $v_{s,t}$. Importantly, house price indices incorporate the value of structures and land, respectively denoted $B_{i,s,t}$ and $L_{i,s,t}$ in our model environment. We then compare our results computed using $q_{s,t}$ and $v_{s,t}$ for the modern Census segment of our data and find that inflating up quantities by a housing price index has little bearing on our cross-sectional results for predicting financial market movements. This is unsurprising given that $q_{s,t}$ and $v_{s,t}$ have a correlation of over 99% across all 50 states from 2000–2019 and across the 20 Case-Shiller MSAs from 1990–2019. Based on this information and the relative ease with which investors can research and interpret permit quantities, we consider $q_{s,t}$ as our main measure of BPG. We explore alternative deflators for converting between permit quantities and valuations in Appendix B.6.

2.5 Supplemental Data Sources

To complement our building permit data, we use several additional data sources. For stock market information, we use the CRSP Stock Database (1926–2019) accessed through WRDS, including the CRSP/COMPUSTAT merge for firm-level accounting fundamentals. We use the value-weighted CRSP index throughout our analysis. Our corporate bond market data come from two sources: the Dow Jones Corporate Bond Total Return Index from Finaeon Global Financial Data (1915–2019); and issue-level data from SDC Refinitiv (1990–2019). For housing price indices, we use the S&P Case–Shiller Index (1988–2019) for 20 MSAs and the state-level Zillow Home Value Index (2000–2019). We also construct a panel of single-family home permits using the data from *Cotality Building Permits* (formerly, CoreLogic) which tracks individual permits from

[10] Current non-disclosure states include Alaska, Idaho, Kansas, Louisiana, Mississippi, Missouri (some counties), Montana, New Mexico, North Dakota, Texas, Utah, and Wyoming.

[11] In their official documentation, the Census Bureau cautions researchers against using their permit value series, stating: *"Because of the nature of the building permit application process, valuations may frequently differ from the true cost of construction. Any attempt to use these figures for inter-area comparisons of construction volume must, at best, be made cautiously and with broad reservations."* (Census Bureau, 2022)

application to completion. We use the Cotality data to calculate completion times and determine the extent to which building permits translate to finished units in Appendix E.

Previous studies note that housing market indicators like permits might proxy for consumption or wealth indicators (Ghent and Owyang, 2010). Building permits capture hiring in the construction sector, which has a direct effect on local job growth and unemployment rates (Strauss, 2013; Howard et al., 2024). We attempt to isolate the additional signal that permits provide to informed investors on top of any other observable signs of local economic conditions by controlling for state and metro area GDP and personal consumption expenditures (PCE) compiled by the Bureau of Economic Analysis (BEA).[12] We download quarterly county and state-level employment statistics from the Quarterly Census of Employment and Wages (QCEW) produced by the BLS.

In our firm-level tests, we use Dun & Bradstreet's DUNS Marketing Identifier (1969–2019) to match firms to their physical locations. Beyond plant-level locations, D&B provides data on establishment-level employment and sales, which we use to apportion out each traded firm's exposure to building permit risk across location in our analyses at the individual securities level. We crosswalk the DUNS id to the *gvkey* firm identifier in COMPUSTAT by matching the parent company name strings across the two databases. This crosswalk allows us to obtain a firm-level panel with sales and employment-share weighted building permit market exposure measures matched to balance sheets and returns.

3 Empirical Strategies

Macro Return Prediction Regressions. We begin by validating our core conjecture that recent building permit growth positively predicts near-future risky asset returns. To do so, we regress CRSP total returns r_t on lags of seasonally adjusted month-to-month building permit growth (BPG), as defined in (2.3):

$$r_t = \beta_0 + \underbrace{\delta_t}_{\substack{\text{seasonal}\\ \text{dummies}}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_\tau \cdot r_{t-\tau}}_{\text{autocorrelation}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_\tau^{BPG} \cdot \Delta q_{t-\tau}}_{\text{U.S. MOM BPG}} + \underbrace{\gamma' \cdot \sum_{p=1}^{p*} \mathbf{X_{t-p}}}_{\text{aggregate controls}} + \varepsilon_t \tag{3.1}$$

[12] The BEA data can be downloaded at https://www.bea.gov/data/by-place-states-territories for states and at https://www.bea.gov/data/by-place-county-metro-local for MSAs and counties.

δ_t is a set of quarterly or monthly dummies to strip out return seasonality.[13] We include τ^* lags of the dependent variable volatility and σ^{BPG} to account for serial autocorrelation.[14] Aggregate controls in $\mathbf{X_t}$ include: the lagged corporate leverage ratio (Schwert, 1989; Cortes and Weidenmier, 2019); population growth (Mankiw and Weil, 1989; Francke and Korevaar, 2025); Shiller (2015)'s CAPE, household debt-service coverage ratio (Mian et al., 2017); growth in the industrial production index; and Manela and Moreira's (2017) news-implied volatility index (NVIX). To address the Smets (2007) critique of Leamer (2007) that building permits may simply predict macroeconomic outcomes due to the sensitivity of residential investment to monetary policy rates, we also consider the excess CAPE yield relative to the 10-year U.S. Treasury yield.[15]

Conditional Building Permit Volatility. Our rational disagreement framework shows how the precision, or inverse volatility, of building permits as a signal of risky cash flows maps to housing supply elasticities. Following Cortes and Weidenmier (2019), we extract volatility from the resulting monthly BPG series in equations (2.1) and (2.3) using GARCH models to compute one-period ahead conditional volatility. We focus on the workhorse GARCH(1,1) model introduced by Bollerslev (1986), to obtain local BPG volatility $\sigma^{BPG}_{s,t}$ for a state or MSA s and BPG measure defined by $x_{s,t} \in \{v_{s,t}, q_{s,t}\}$.

$$\begin{aligned} x_{s,t} &= \theta_0 + \theta_1 \cdot x_{s,t-1} + \varepsilon_{s,t}, \text{ with } \varepsilon_t \sim \mathcal{N}(0, (\sigma^{BPG}_{s,t})^2) \\ (\sigma^{BPG}_{s,t})^2 &= \alpha_0 + \alpha_1 \cdot \varepsilon^2_{s,t-1} + \alpha_2 \cdot (\sigma^{BPG}_{s,t-1})^2 \quad \text{s.t. } \alpha_i > 0;\ \alpha_1 + \alpha_2 < 1 \end{aligned} \tag{3.2}$$

where $x_{s,t}$ is BPG as constructed from data on permits issued and proxies for market values of newly constructed buildings, adjusted to capitalize land values (e.g., via housing price indices) whenever possible. We estimate (3.2) using quasi-maximum likelihood (QMLE) methods. In Appendix C we show robustness of our results to using other GARCH models to estimate BPG conditional volatility, such as GJR-GARCH (Glosten et al., 1993) and exponential GARCH (Nelson, 1991). In our simulation exercises, for most building permit series and time periods the, workhorse GARCH model exhibits less parameter instability while attaining similar convergence rates to analogous GJR-GARCH and E-GARCH specifications.[16]

[13] Since we construct σ^{BPG} from seasonally-adjusted permits, further adjustments for seasonality by incorporating either monthly or quarterly dummies into equation (3.3) has little quantitative impact on our results. In our baseline specifications we include monthly dummies to account for possible cyclicality in asset return volatility arising from quarterly earnings announcements.

[14] We select a lag order of $\tau^* = 12$ months to obtain comparable estimates to the literature (e.g., Schwert, 1989) but consider Akaike (AIC) and Bayesian-selected criteria (BIC) τ^*. We also consider the optimal lag order selected by allowing it to be asymmetric between σ and σ^{BPG} (Ozcicek and McMillin, 1999). In all cases we obtain $\tau^*_{AIC} = \tau^*_{BIC} = 1$.

[15] 30-year fixed rate mortgages, the most common loan type in the postwar period, are indexed to the 10-year U.S. Treasury rate due to average household tenure in the home.

[16] Further, GARCH(1,1) converges for more building permit series across geographic areas and subperiods than textbook GJR-GARCH and E-GARCH models.

FIGURE 3. Conditional U.S. Aggregate BPG *vs.* Financial Market Volatility (1919–2022)

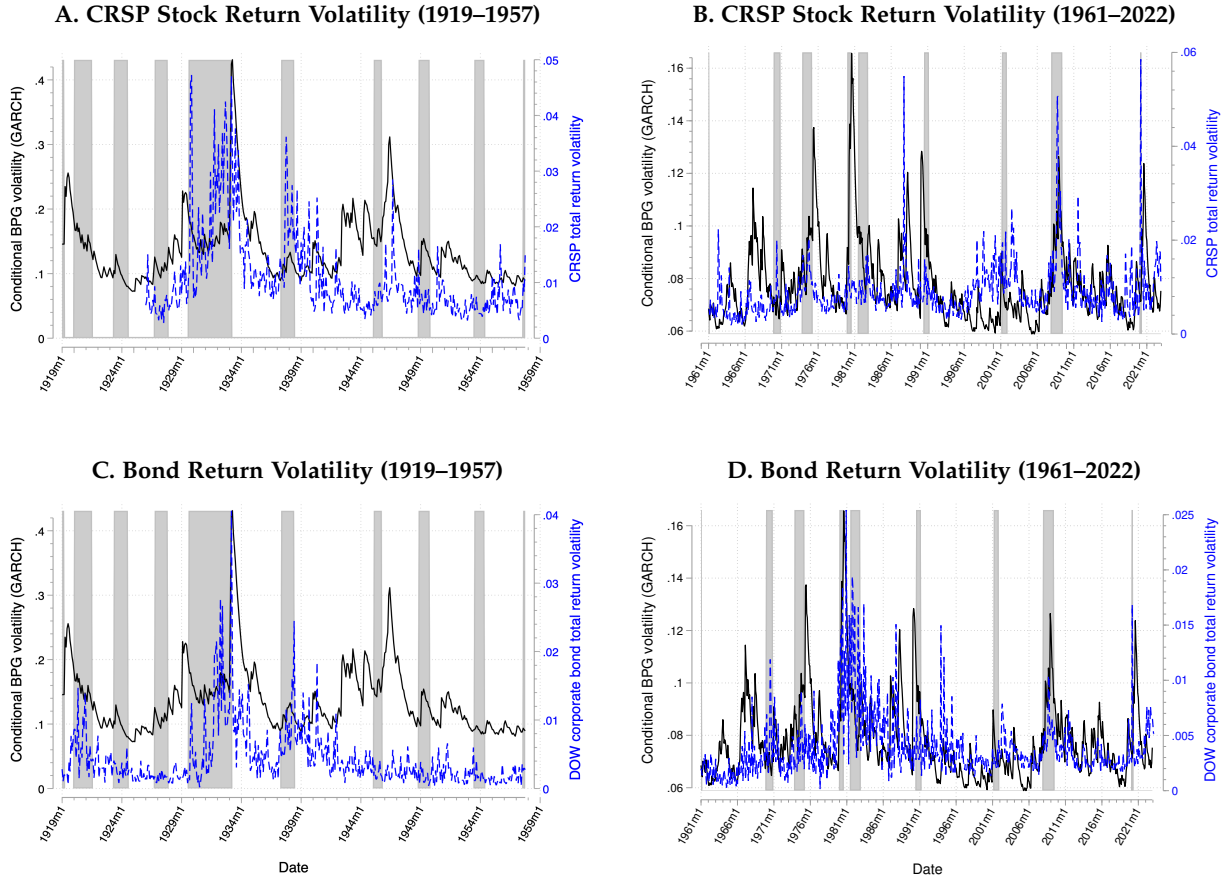

Notes: The left-hand axis in each panel plots conditional building permit growth volatility σ^{BPG} obtained by estimating the GARCH(1,1) model in equation (3.2) over the series of seasonally adjusted total private residential permits across the entire U.S. Panels A and C use the *Dun's Review* sample, spanning 1919M1–1957M10, while Panels B and D use the Census Building Permits Survey, covering 1961M1–2022M12. The right-hand axis in Panels A and B plots CRSP total return volatility, while the right-hand axis in Panels C and D refers to total return volatility on the Dow Jones corporate bond index. We compute CRSP and Dow return volatility as the monthly standard deviation from daily total returns on the value-weighted index. Grey-shaded areas indicate NBER-dated recessions.

As a proof of concept, we plot in Figure 3 the time series of σ_t^{BPG} for total private residential permits and total return volatility for stocks (Panels A and B) and corporate bonds (Panels B and C). Following Schwert (1989), we define total return volatility σ_t as the monthly standard deviation calculated from daily returns. Conditional BPG volatility spikes within a 6-month lead relative to the stock and bond markets in 12 out of 15 NBER recessions. BPG and bond total returns have lower average volatilities following the Savings & Loan Crisis of the late 1980s.[17] Stock and Watson (2010) find that the decline in nationwide BPG and real volatility (i.e., the Great Moderation) coincided with a convergence in mortgage rates across regions. As we show in Section 4, our finding that BPG volatility predicts financial return volatility holds longitudinally over the last one-hundred years even conditional on proxies for housing demand — including the prewar period when mortgage credit was scarce — indicating that BPG volatility is not merely a proxy for the leverage cycle.[18]

Aggregate Volatility Regressions. Analogous to the return on return regression in (3.1), our baseline specification for testing for links between asset market fluctuations and building permit volatility is the following time series regression:

$$\sigma_t = \beta_0 + \underbrace{\delta_t}_{\substack{\text{seasonal}\\ \text{dummies}}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_\tau \cdot \sigma_{t-\tau}}_{\text{autocorrelation}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_\tau^{BPG} \cdot \sigma_{t-\tau}^{BPG}}_{\text{U.S. BPG volatility}} + \underbrace{\gamma' \cdot \sum_{p=1}^{p*} \mathbf{X_{t-p}}}_{\text{aggregate controls}} + \varepsilon_t \tag{3.3}$$

where σ_t is total return volatility for an asset class (e.g., stock or bond total returns), and $\sigma_{s,t}^{BPG}$ denotes one-period conditional volatility for locality s, as obtained via the GARCH model in (3.2). We follow the same BIC-optimal lag order selection as in our return regressions.

Geographic Cross-Sectional Volatility Regressions. A natural extension of equation (3.3) involves replacing aggregate BPG volatility σ_t^{BPG} on the RHS with local BPG volatility $\sigma_{s,t}^{BPG}$ to test for heterogeneity in the informativeness of building permit movements across geography:

$$\sigma_t = \beta_0 + \underbrace{\delta_t}_{\substack{\text{seasonal}\\ \text{dummies}}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_\tau \cdot \sigma_{t-\tau}}_{\text{autocorrelation}} + \underbrace{\sum_{\tau=1}^{\tau^*} \beta_{s,\tau}^{BPG} \cdot \sigma_{s,t-\tau}^{BPG}}_{\substack{\text{BPG volatility}\\ \text{for locality } s}} + \underbrace{\gamma_s' \cdot \sum_{p=1}^{p*} \mathbf{X}_{s,t-p}}_{\text{local controls}} + \varepsilon_t \tag{3.4}$$

[17] Average BPG volatility declines from 0.083 to 0.073 after the nationwide trend break in 1992M6, with a two-sided t-stat on the difference of 9.40. We present results from applying Bai and Perron (1998) one-at-a-time break date tests to the aggregate and local building permit series in Appendix D.1.

[18] Appendix Figure D.3 depicts the incremental-R^2 exercise of Cortes and Weidenmier (2019), adding lags of BPG volatility *versus* corporate leverage to the same autoregressive volatility model. Leverage is mechanically related to stock volatility (Schwert, 1989), making it a natural candidate for a predictive-power horse race. Nonetheless, BPG volatility explains more incremental variation than leverage in 27 of 42 states. In the industry breakdown, it also exceeds leverage for Consumer Durables and comes close for Automobiles, two of the most cyclical industries.

where now the vector of controls $\mathbf{X_{s,t}}$ includes local economic factors within geographic unit s, such as recent population, employment, and GDP growth. To the extent that such local economic conditions are readily discernible to investors, failing to account for their influence on financial markets could lead us to potentially ascribe too much predictability to BPG volatility.

Principal Components Analysis: Accounting for Cross-Regional Collinearity. We also consider versions of (3.4) in which we include a subset of localities $s \in \mathcal{S}$ in a panel regression or use principal component analysis (PCA) to account for collinearity across local real estate markets. We test the hypothesis admitted by our model in Section 6 that $\beta_{s,\tau}$ will strongly predict future asset return volatility in areas of the country where building permits form a more precise signal to investors about the strength of the local economy. Such areas are ones with fewer regulatory restrictions but more land availability constraints which keep BPG volatility relatively low, on average. We report the cross-state correlation matrix for BPG in Appendix A.3. Generally, states like Florida for which building permits offer a more precise signal of financial markets have BPG exhibiting a higher average correlation with all other states.

Firm-Level Analysis: Cross-Section of Equity and Bond Volatility. To examine the relationship between firms' exposure to local construction cycles and their own-securities return volatility, we extend our time series analysis in (3.3) and (3.4) to the cross-sections of equities and bonds:

$$\sigma_{j,t} = \delta_t + \eta_j + \underbrace{\sum_{\tau=1}^{\tau_j^*} \beta_{j,\tau} \cdot \sigma_{j,t-\tau}}_{\text{own autocorrelation}} + \underbrace{\sum_{\tau=1}^{\tau_j^*} \varphi_{j,\tau} \times \left(\sum_{k \in \mathcal{J}} \omega_{k,t-\tau-1} \cdot \sigma^{BPG}_{k,t-\tau} \right)}_{\text{share-weighted exposure to BPG Vol}} + \underbrace{\gamma' \cdot \mathbf{X}_{j,t-1}}_{\text{firm-level controls}} + \varepsilon_{j,t} \quad (3.5)$$

where we include autoregressive terms $\sigma_{j,t-\tau}$ and a vector of controls $\mathbf{X_{j,t}}$ commonly used in the literature (e.g., size/age bins, leverage ratios, EBITDA, Tobin's Q). The $\varphi_{j,\tau}$ coefficients capture how each firm j's exposure to real estate market risk through its network of operating locations $\mathcal{J}$ predicts its own asset return volatility. We use Dun & Bradstreet's DUNS Marketing Identifier data (1969–2019) to construct (lagged) plant-level shares $\omega_{k,t-1}$ based on sales or employment, resulting in a share-weighted average measure of each firm's exposure to local BPG risk. This Bartik-style shock with possibly time-varying weights on BPG volatility exposure allows us to examine how firms' connections to local housing cycles through their plant networks affect their stock or bond return volatility.

4 Main Results: BPG Volatility Predicts Financial Cycles

This section documents the main empirical channels through which building permits appear to lead macro-financial movements. We first present longitudinal results using the full timespan of

our data. We then probe these channels through the lens of historical episodes to help establish whether our σ^{BPG} factor is a general predictor of financial boom-bust cycles throughout history regardless of the nature of risks the economy faces at particular points in time. We return in Section 6 to a model that organizes these patterns.

4.1 Longitudinal Analysis of Building Permits' Predictive Power

We begin by testing whether building permit growth positively predicts near-future stock returns, a relationship that we later rationalize in Section 6 as *Hypothesis 2*. Table 1 estimates regressions of total returns on nationwide BPG based on equation (3.1) for the post-1960s Census BPS data. We validate that this hypothesis holds and is robust to controlling for autocorrelation in returns and to the inclusion of a large set of macro factors such as past commodity future returns ($CFRI$), population growth ($PopGrowth$), price-earnings ratios ($CAPE$), yield curve movements ($CAPEYield$), industrial production growth ($IPGrowth$), corporate and household leverage ratios, and the news-implied VIX ($NVIX$). Average month-to-month BPG is 40 basis points; however, on the eve of recessions, there are large swings of between 10-20%. Outside of the Global Financial Crisis, our results imply that on the eve of a typical recession, 6-month lagged rolling average drawdowns in BPG predict a 50 basis point drop in aggregate stock market returns. While the results in Table 1 pertain to building permits in the entire U.S., we always estimate $\widehat{\beta}_s^{BPG} > 0$ if we run a separate set of regressions for each state s.[19]

Table 2 presents results from estimating the aggregate on aggregate volatility specification of equation (3.3) with a BIC-optimal lag of $\tau^* = 1$. In each specification, we regress total value-weighted CRSP return volatility or total Dow Jones Corporate Bond return volatlity on one-month lagged conditional BPG volatility, σ_{t-1}^{BPG}. We perform this exercise in Panel A using all private residential permits to construct BPG according to (2.3), and instead subset to permits for new single-family home (SFH) units in Panel B. We focus on growth rates in permit quantities rather than valuations to first establish the predictive power of permits independent of any assumptions about how to value the permits at the point they are filed. Our sample covers the entire U.S. over the run of the Census BPS, up to 2019M12.

Our broad conclusion from Table 2 is that over the full period 1960–2019, BPG volatility has strong predictability for both equity return and corporate bond return volatility, even conditional on a host of macroeconomic factors which might simultaneously drive asset market fluctuations and investors' appetite for developing real estate. Stock volatility is more predictable around the Great Recession period but the reverse is true for bonds. Incremental R^2 is very high for corporate bonds but weaker for equities over the full sample. The R^2 improves to 16% for CRSP around

[19]Several of our macro control variables are only available at a lower frequency than the monthly permits data. With month fixed effects included in every specification, the results are virtually identical regardless of whether we choose a lag order to reflect the frequency at which a low-frequency control variable updates (e.g. include $DSCR_{t-4}$ instead of $DSCR_{t-1}$ to reflect the quarterly frequency) or linearly interpolate between each updated value.

TABLE 1. Regressions of Stock Total Returns on Nationwide U.S. Building Permit Growth (BPG)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$r_{CRSP,t-1}$	0.059	0.048	0.057	0.056	0.050	0.052	0.051	0.063	0.086
	(1.331)	(1.083)	(1.344)	(1.321)	(1.167)	(1.186)	(1.160)	(1.187)	(1.603)
$US\ BPG_{SFH,t-1}$		0.055***	0.057***	0.057***	0.055***	0.054***	0.054**	0.044*	0.049*
		(2.576)	(2.719)	(2.720)	(2.630)	(2.576)	(2.591)	(1.650)	(1.692)
Time sample	1960-2019	1960-2019	1960-2019	1960-2019	1960-2019	1960-2019	1960-2019	1980-2019	1980-2016
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓
$CFRI_{t-p}$			✓	✓	✓	✓	✓	✓	✓
$PopGrowth_{t-p}$			✓	✓	✓	✓	✓	✓	✓
$CAPE_{t-p}$				✓					
$CAPE\ Yield_{t-p}$					✓	✓	✓	✓	✓
$IPGrowth_{t-p}$						✓	✓	✓	✓
$Leverage_{t-p}$							✓	✓	✓
$DSCR_{t-p}$								✓	✓
$NVIX_{t-p}$									✓
N	715	715	715	715	715	715	715	480	434
R^2	0.025	0.035	0.036	0.040	0.046	0.049	0.050	0.055	0.068

Notes: The table presents estimates from equation (3.3) relating total returns to lagged single-family home (SFH) building permit growth (BPG). We include in most specifications a set of controls for other macroeconomic observables directly related to construction sector expansions, which might in turn drive aggregate stock market performance. *CFRI* is the commodity futures return index of Janardanan et al. (2024). CAPE is the cyclically adjusted total return price-earnings ratio, and CAPE yield is the inverse CAPE ratio less the 10-year Treasury yield, both provided by Shiller (2015). *PopGrowth* refers to the annual population growth rate from the Census. *IPGrowth* is the month-on-month growth rate in the industrial production index (INDPRO). *Leverage* is the aggregate corporate leverage ratio based on firms' annual filings in COMPUSTAT, computed as the sum of long-term debts (*dltt*) and debts in current liabilities (*dlc*), divided by total stockholders' equity (*seq*). *DSCR* is the quarterly household debt service coverage ratio from the Federal Reserve, defined as household debt service payments as a fraction of disposable income. In the final column, we add the News Implied Volatility Index (NVIX) of Manela and Moreira (2017). We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. The time sample varies depending on the data availability of covariates, with *DSCR* available starting in 1980, and the NVIX available only up to 2016. *t*-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

TABLE 2. Regressions of Asset Return Volatility on Nationwide U.S. BPG Volatility

A. Total Private Residential Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.058**	0.029***	0.036**	0.030***	0.107***	0.069***	0.035***	0.030***	0.024*	0.019**
	(1.97)	(2.75)	(2.53)	(2.68)	(3.58)	(4.32)	(3.18)	(2.59)	(1.84)	(2.30)
Time sample	1960-19	1960-19	1980-19	1980-16	2000-16	1960-19	1960-19	1980-19	1980-16	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$Leverage_{t-p}$			✓	✓	✓			✓	✓	✓
$DSCR_{t-p}$			✓	✓	✓			✓	✓	✓
$IP\ Growth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
N	714	707	479	435	195	714	707	479	435	195
R^2	0.049	0.478	0.461	0.470	0.613	0.139	0.358	0.437	0.428	0.541

B. Single-Family Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.033*	0.022***	0.026**	0.023**	0.063**	0.076***	0.044***	0.041***	0.038***	0.023***
	(1.75)	(2.76)	(2.55)	(2.35)	(1.97)	(6.05)	(4.95)	(4.74)	(4.00)	(3.77)
Time sample	1960-19	1960-19	1980-19	1980-16	2000-16	1960-19	1960-19	1980-19	1980-16	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$Leverage_{t-p}$			✓	✓	✓			✓	✓	✓
$DSCR_{t-p}$			✓	✓	✓			✓	✓	✓
$IP\ Growth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
N	714	707	479	435	195	714	707	479	435	195
R^2	0.031	0.476	0.458	0.467	0.594	0.201	0.379	0.454	0.446	0.539

Notes: The table presents estimates from equation (3.3) relating total return volatility to lagged building permit growth (BPG) volatility. In Panel A we use total private residential permits as the quantity measure $Q_{s,t}$, but instead use permits attached to single-family units in Panel B. We include in most specifications a set of controls for other macroeconomic observables directly related to BPG volatility, which might also drive aggregate financial market volatility. *PopGrowth* refers to the annual population growth rate from the Census. *Leverage* is the aggregate corporate leverage ratio based on firms' annual filings in COMPUSTAT, computed as the sum of long-term debts (*dltt*) and debts in current liabilities (*dlc*), divided by total stockholders' equity (*seq*). *DSCR* is the quarterly household debt service coverage ratio from the Federal Reserve, defined as household debt service payments as a fraction of disposable income. *IPGrowth* is the month-on-month growth rate in the industrial production index (INDPRO). In some specifications, we add the natural disaster component of the News Implied Volatility Index (NVIX) of Manela and Moreira (2017). We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. The time sample varies depending on the data availability of covariates, with *DSCR* available starting in 1980, and the NVIX available only up to 2016. *t*-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

the GFC period (2000-19), even without including the vector of controls.[20] In Appendix D.4, we show robustness of the results in Table 2 to adding additional controls for commodity price risk and household mortgage access from monthly loan originations.

Building permits for new SFHs appear less closely linked to equities than to corporate bonds. The opposite is true for total private residential units, which includes permits for both SFHs and multifamily housing (MFH) units. The estimated elasticity for equities is roughly 40% greater in the specifications with the full set of controls when we use this total BPG volatility measure. From an investor's perspective, there are tradeoffs to acting on signals from the SFH vs. MFH market. Income-generating properties behave more like financial assets than single-family homes, because rental cash flows can be discounted to produce a valuation, whereas single-family homes are valued in an hedonic fashion to impute the unobserved dividends consumed by owner-occupiers. New multifamily properties may contain more soft information about local economies to the extent that they attract institutional investors who can develop units at scale (Gurun et al., 2023). Larger properties (i.e., those with at least 5 units) are more likely to be highly levered with floating rate debt (Glancy et al., 2023), leading to more volatility since developers' appetites will ebb and flow with interest rate conditions (Glancy and Kurtzman, 2022). Indeed, the average value of $\sigma^{BPG}_{US,t}$ for MFH permits is 12.0%, compared to 7.6% for SFH permits.[21]

However, some of the increased volatility arising from MFH permits may be noise, as it captures behavioral responses to idiosyncratic local regulatory changes rather than beliefs about economic fundamentals. There are clear spikes in MFH or total permits filed in months corresponding to the state's tax year end date, even after plotting the seasonally adjusted permit counts in Figure B.3. Developers filing permits to lock in preferential property tax incentives explains these anomalies.[22] Since it is more difficult to separate signal from noise in the MFH segment of the market, we use SFH BPG volatility as our preferred measure in empirical tests using the post-1960s data.[23]

Figure 4 decomposes the aggregate predictability by state and at short (1-month) vs. medium-run (12-month) horizons, according to equation (3.4). Each regression in the figure

[20]Note that since we already seasonally adjust the permits data and there is little monthly seasonality in stock or bond returns on the aggregate index, almost all of this explanatory power is coming from BPG volatility itself rather than time fixed effects. For example, regressing post-2000 CRSP return volatility on only a full set of monthly dummies results in an R^2 of just 1.6%, with σ^{BPG} adding an incremental R^2 of 14.4%.

[21]Moreover, due to the high unconditional volatility in multifamily permits, our GARCH specification in (3.2) applied to multifamily BPG only converges for 31 out of the 50 states plus the aggregate U.S.

[22]For instance, we observe a large spike in multifamily permits (which is included in the total permits series in Figure B.3) in California on the eve of Proposition 13, which passed in June 1978 and set property tax assessment limits indexed to a 1976 fiscal base year value. Similarly, there are large spikes in multifamily permits in New York City in December 2007 and December 2015 before pre-announced rollbacks of the 421a property tax exemption which created incentives for developers to file permits to lock in more generous tax abatements (Soltas, 2024).

[23]From the perspective of building coherent long-run time series, Census total residential permits correspond better to the composition of permits surveyed by *Dun's Review* during the pre-1960s time period.

FIGURE 4. Predictability of State-Level Permit Growth Volatility for Asset Return Volatility (1960–2019)

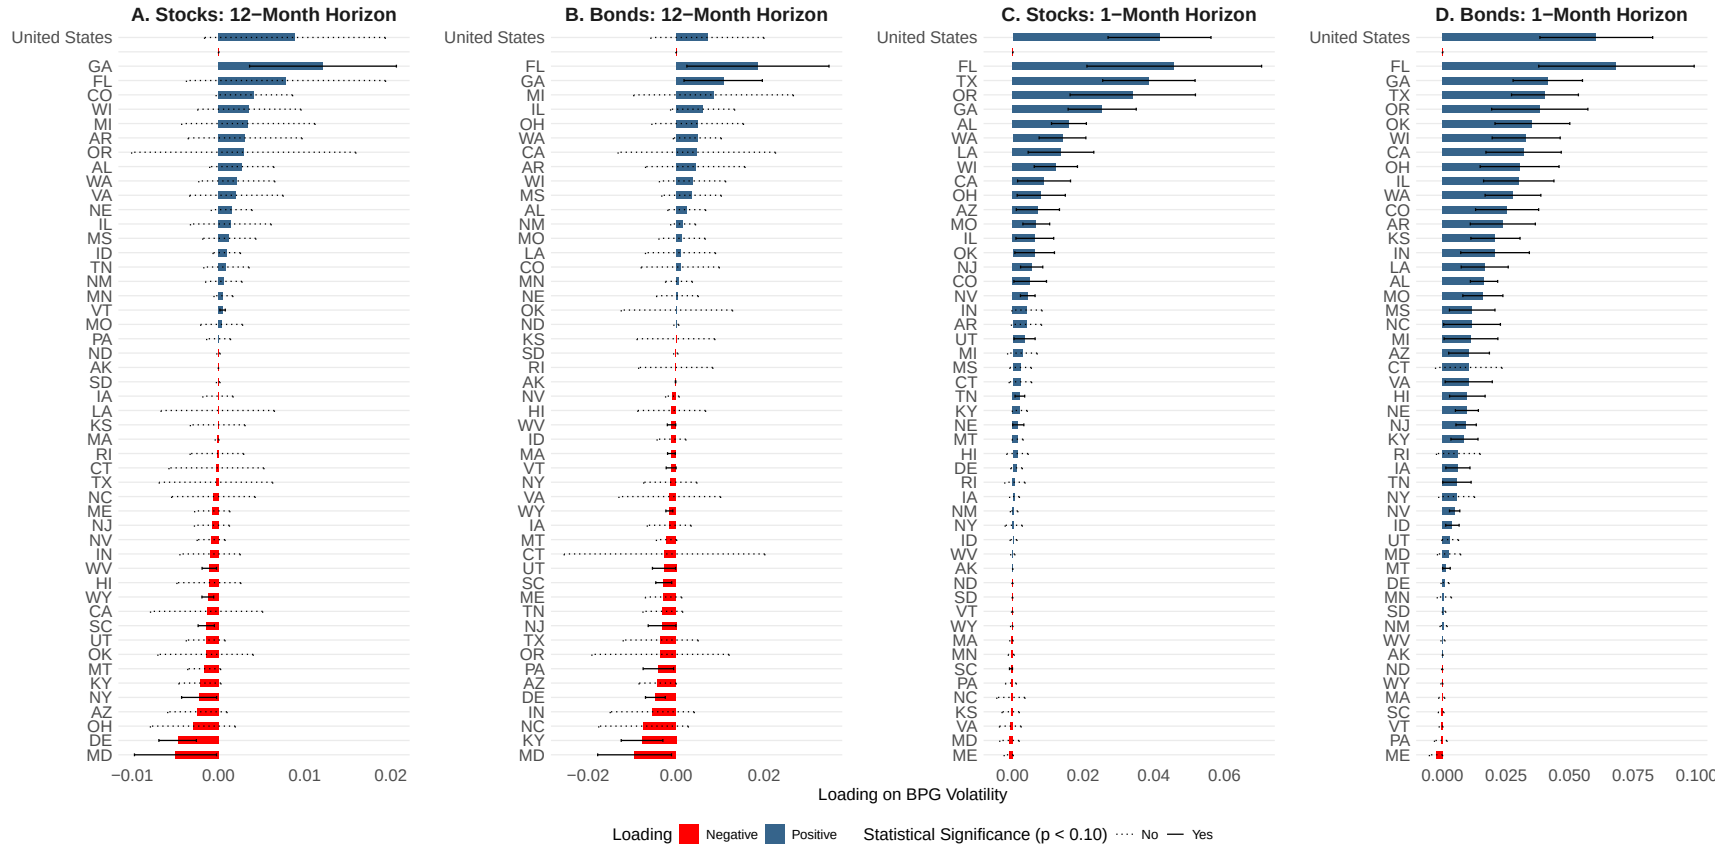

Notes: We estimate equation (3.3) for the entire U.S. and equation (3.4) separately for each state s. Each regression includes monthly dummies, the BIC-optimal autocorrelation term σ_{t-1}, but no other controls. We omit the regressions for Washington, D.C. and New Hampshire, as the GARCH model does not converge for those series. In panels A and B, we use $\tau^* = 12$ lags of $\sigma_{s,t}^{BPG}$ and report the sum of the estimated lagged coefficients $\sum_{\tau=1}^{\tau^*} \widehat{\beta}_{s,\tau}$. In panels C and D, we show predictability at the one-month horizon. Our sample period covers the Census Building Permit Survey from its start in 1960M5 to 2019M12. We use permits attached to single-family units as the quantity measure $Q_{s,t}$. Panels A and C perform this exercise for CRSP stock return volatility as the outcome, and total return volatility of the Dow Jones Corporate Bond Index in Panels B and D. Our sample period covers the Census Building Permit Survey from its start in 1960M5 to 2019M12. We use permits attached to single-family units as the quantity measure $Q_{s,t}$. We plot 90% confidence intervals obtained via autocorrelation and heteroskedasticity-robust standard errors, selecting the minimum lag order such that the estimator for the covariance matrix is consistent (Newey and West, 1987). Solid confidence bars indicate significance at the 90% level, while dashed bars indicate statistical insignificance.

includes monthly dummies, σ_{t-1}^{BPG}, and one-month lagged asset return volatility, but no controls.[24] A clear picture emerges of a small number of states, particularly Florida and Georgia, driving the predictability of BPG across both asset markets. This heterogeneity in the loadings is follows *Hypotheses 3 and 4* of our asymmetric information framework, whereby the predictive power of building permits depends on the extent to which local housing supply constraints bind. The figure reinforces the conclusions from Table 2 by showing that BPG volatility is a stronger signal for the bond market over the 60-year time period in the cross-section of states as well as for the aggregate U.S. In Appendix D.6, we reproduce the same figure showing the predictability of BPG volatility conditional on proxies for local housing demand, including state-level corporate leverage and population growth.

Predictability grows in magnitude as we move from the short-run to the medium-run, with some non-linearities present when comparing cumulative loadings over the 1-month vs. 6-month horizons (unpictured). However, due to large confidence interval bands, particularly for stocks, we cannot statistically distinguish the signals offered by most states from each other.[25] We show that our results in Figure 4 are stronger and more precisely estimated for both the entire U.S. and the cross-section of states — while largely preserving the ordinal ranking of predictability of BPG volatility across states — when we use GJR-GARCH instead of the more standard GARCH(1,1) specification in equation (3.2). Appendix C.3 documents that using GJR-GARCH addresses the skewness of the distribution of monthly building permit growth observations. The extent of this skewness varies across states and time periods; we fail to reject the null of a symmetric BPG distribution for the aggregate U.S. in the post-1960s period for either overall residential permits or SFH permits. We therefore face a tradeoff between the stability and convergence properties of GARCH(1,1) for smaller states and maintaining the precision of our estimates for states with large swings in permitting activity.

What explains the stronger predictability of BPG volatility for corporate bonds relative to stock volatility over the full modern sample time period? One possibility is that bond returns are systematically more predictable since they are non-callable and offer fixed coupon payments. We indirectly test the role of cash flow predictability in Appendix D.3 by decomposing CRSP total returns into the capital gain vs. dividend components (i.e., by taking the difference between `vwretd` and the ex-dividend total return index `vwretx`), computing stock dividend volatility, and then estimating analogous regressions to those in Table 2. We find that BPG volatility is a better predictor of dividend volatility, in an R^2 sense, than total return volatility since 1960. However, the opposite is true in the post-2000s period, during which dividends were

[24] In all regressions, we obtain nearly identical loadings on σ^{BPG} if we instead interact the autocorrelation return σ_{t-1} with a set of monthly dummies to account for the seasonality in the predictability of lagged stock and bond returns (Ogden, 2003; Heston and Sadka, 2008).

[25] We use Newey and West (1987) standard errors in all specifications pictured in the figure and hereafter, as we find the Newey-West estimator delivers t-statistics which are more conservative, on average, than those obtained via bootstrapping. We obtain more precisely estimated coefficients if in our GARCH models (3.2) we impose errors following a t-stat distribution instead of a Gaussian one.

less volatile. Consistent with our Grossman–Stiglitz modeling framework in Section 6, BPG volatility works well at forecasting dividend volatility, particularly during times when the cash flow risk component of returns dominates.

4.2 The Great Depression and Postwar Boom

The Great Depression is a pivotal episode in macroeconomic and financial history, with stock market volatility reaching an unprecedented 60% per annum — two to three times higher than any other period in American financial history (Schwert, 1989). Cortes and Weidenmier (2019) show that building permit growth volatility and financial leverage help rationalize such high degrees of volatility in an incremental R-squared sense.

In this subsection, we extend and enhance the findings of Cortes and Weidenmier (2019) in several ways. Most notably, while Cortes and Weidenmier (2019) focused on Schwert's (1989) Depression volatility puzzle period (i.e., 1928–1938), our study spans a much longer timeframe. Our data cover the period from 1919 to 1957, spanning nearly four decades, including the Roaring Twenties, the Great Depression, World War II, and the early post-war boom. This expanded dataset allows us to gain statistical power by examining the relationship between building permit growth volatility and asset return volatility across multiple economic cycles and regimes. Second, our analysis provides a more comprehensive understanding of this relationship over time for both stocks and corporate bonds rather than just stocks. Furthermore, our application of modern Census X-13 ARIMA-SEATS seasonal adjustment methods provides a robust treatment of seasonality.

Another explanation for the predictive power of building permits in aggregate and in the geographic cross-section is that local land use regulation ebbs and flows with macroeconomic sentiment, influencing the project's probability of success and its value upon completion, through the expected future cash flows generated by the property. Imposing supply constraints raises the value of holding permitted land, contributing to price growth during periods of high property demand (Glaeser et al., 2005).[26] In the model of Section 6, we assume local residential supply restrictions are static. Studying the historical period in which zoning and permitting systems were less complex (Fischel, 2004; Shertzer et al., 2022) is therefore useful for trying to rule out shifts in land use constraints as the main driver of the predictive power of our BPG volatility factor across crisis episodes.

Figure 5 illustrates the predictive power of building permit growth volatility for both stock (Panel A) and bond (Panel B) return volatility in the U.S. and across different states over the complete *Dun's Review* sample period, 1919–1957. We adopt specifications identical to those in Figure 4. The figure reveals significant heterogeneity in predictability across states, with some

[26]Comparing the earlier version of the Wharton Residential Land Use Regulatory Index from 2006 (Gyourko et al., 2008) to the 2018 version (Gyourko et al., 2021), there is little change in the ranking of states and 44 CBSAs.

FIGURE 5. Predictability of Building Permit Growth Volatility for Asset Return Volatility (1919–1957)

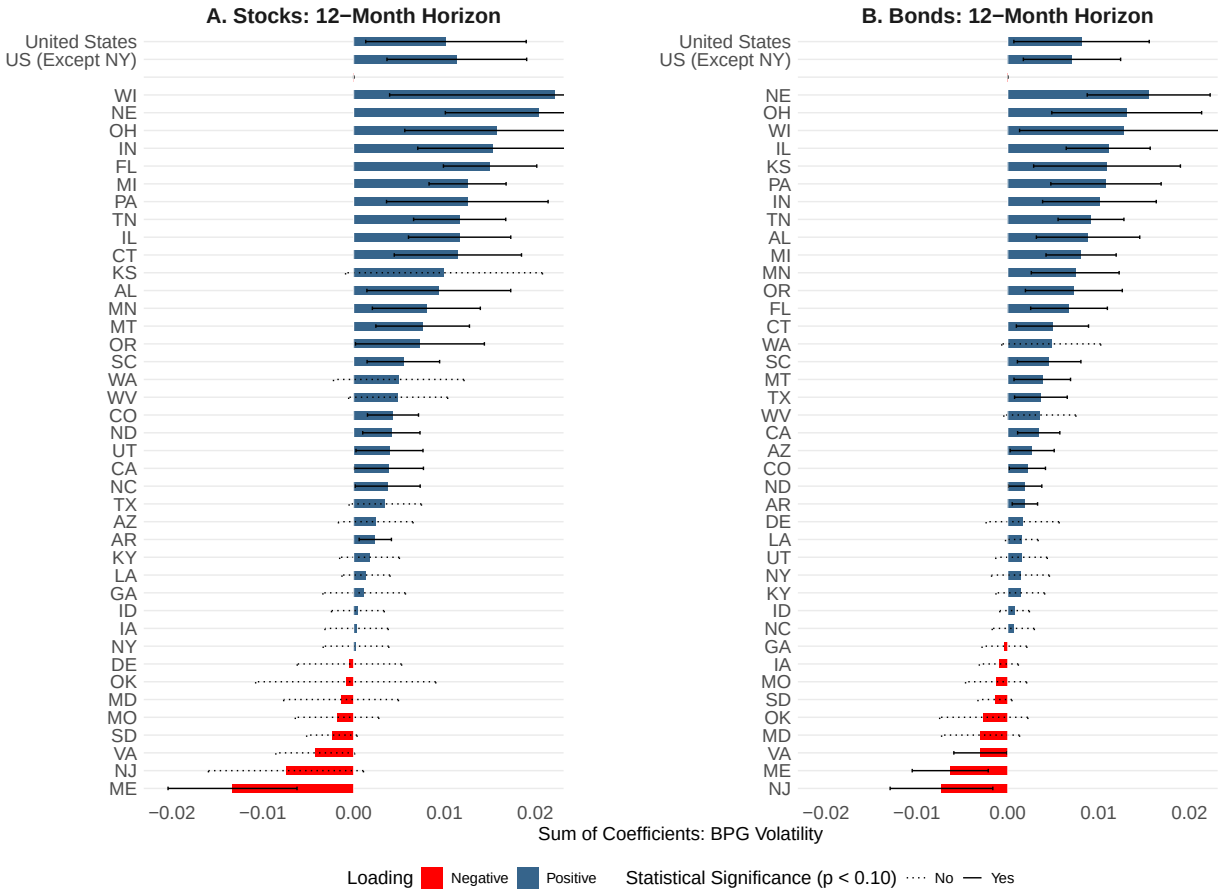

Notes: This figure shows the sum of coefficients for building permit growth (BPG) volatility predicting stock and bond return volatility over a 12-month horizon, using the sample of permit valuations from *Dun's Review*. We estimate versions of equation (3.3) separately for each state s with $\tau = 12$ lags of $\sigma^{BPG}{}_{s,t}$ and report the sum of the estimated lagged coefficients $\sum_{\tau=1}^{12} \widehat{\beta}_{s,\tau}$. The sample period is 1919M1 to 1957M10. Panel A shows results for stock return volatility, while Panel B shows results for bond return volatility. Results are shown for the United States at the top of each panel, then the United States excluding New York State, and then individual states. Blue (red) bars indicate positive (negative) coefficients. Solid error bars represent statistical significance at the 10% level, while dashed lines represent insignificant coefficients. We truncate the x-axis in Panel A to allow better visualization of the estimates and their confidence intervals. We plot 90% confidence intervals obtained via standard errors to correct for autocorrelation and heteroskedasticity, selecting the minimum lag order such that the estimator for the covariance matrix is consistent (Newey and West, 1987). See Appendix B.1 for details on how we aggregate the *Dun's Review* data into state-level permits series.

showing strong positive relationships (e.g., Wisconsin, Nebraska, Ohio, Indiana, and Florida) and others showing negative or insignificant relationships.

Notably, Florida BPG volatility exhibits strong cumulative predictability at the 12-month horizon for both asset classes, particularly if we re-estimate (3.4) restricting to dates within five years of the onset of the Great Depression. Narrative evidence points to the colossal magnitude of the Florida Land Boom of the 1920s, although no land price indices exist (White, 2009, 2014). Total building permits issued in Miami peaked at $65.7 million in 1925M3 and plummeted to $1.2 million (both in real 2012 dollars) on the eve of the 1928 Okeechobee Hurricane, the deadliest in Florida's history. More than 20 million lots were being developed for sale in Florida over the boom (Knowlton, 2020, p. xiv), a scale which could have accommodated half of the entire United States population planning to move to Florida (Calomiris and Jaremski, 2023, p. 2). The Florida land boom was predicated on the development of land made newly accessible by expansions of railroad networks along the Eastern Seaboard in the early twentieth century (Turner, 2015, pp. 90–113). Our model predicts that a context like 1920s Florida in which there are minimal regulatory restrictions on real estate development is one in which BPG volatility is a strong predictor of asset market volatility.

Figure 5 also highlights the importance of considering the United States both with and without New York, demonstrating how a single state with significant economic weight can influence national-level results. This nuanced approach provides a more comprehensive view of the relationship between building permit volatility and asset return volatility across the country. We discuss how we aggregate the *Dun's* permits data from cities to states in Appendix B.1. Our results for the aggregate predictability of BPG volatility in the earlier *Dun's* period hold even if we winsorize permit valuations at the tails within each state before aggregating the series from cities up to states and then up to the national level.

Analogously to Table 2 for the modern time period, we estimate aggregate BPG volatility on aggregate return volatility regressions for the *Dun's* sample period in Table 3. In line with Cortes and Weidenmier (2019), Panel A of Table 3 shows that the one-month predictability of BPG volatility for stock volatility is stronger around the Great Depression era. A novel finding shown by Panel B is that the strong predictability of BPG is present for both stock and bond volatility — even conditional on controls for market leverage, population growth, industrial production growth, and the NVIX (e.g., the war and natural disaster components). Both panels indicate that the predictability remains quantitatively and qualitatively significant over the full sample from 1919 to 1957. This suggests that our findings for the modern period are not specific to recent developments in the U.S. corporate or household lending environment, or contingent on the definition of BPG volatility.[27]

[27] In contrast to the Census Building Permits Survey, *Dun's Review* reports building permit valuations rather than counts, and includes both income-generating properties and single-family homes.

TABLE 3. Regressions of Asset Return Volatility on Nationwide U.S. BPG Volatility: Pre-1960s Sample

A. Stock Return Volatility and BPG Volatility (1926–1957)

Sample Period:	Full Time Period					Great Depression Era				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.036**	0.013**	0.013**	0.013**	0.017***	0.037**	0.020***	0.021***	0.021***	0.020***
	(2.52)	(2.37)	(2.39)	(2.41)	(2.94)	(2.46)	(2.85)	(3.18)	(3.13)	(3.00)
Time sample	1926-57	1926-57	1926-57	1926-57	1926-57	1928-38	1928-38	1928-38	1928-38	1928-38
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$MktLeverage_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$IPGrowth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
$WarNVIX_{t-p}$					✓				✓	✓
N	381	381	381	381	381	131	131	131	131	131
R^2	0.102	0.618	0.618	0.620	0.631	0.147	0.613	0.614	0.615	0.629

B. Bond Return Volatility and BPG Volatility (1919–1957)

Sample Period:	Full Time Period					Great Depression Era				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.021***	0.009**	0.009**	0.009**	0.011**	0.030***	0.017**	0.021***	0.021***	0.021***
	(2.72)	(2.11)	(2.10)	(2.12)	(2.54)	(3.31)	(2.39)	(3.02)	(2.95)	(2.97)
Time sample	1919-57	1925-57	1925-57	1925-57	1925-57	1928-38	1928-38	1928-38	1928-38	1928-38
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$MktLeverage_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$IPGrowth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
$WarNVIX_{t-p}$					✓				✓	✓
N	465	393	393	393	393	131	131	131	131	131
R^2	0.090	0.515	0.516	0.518	0.525	0.142	0.527	0.541	0.542	0.543

Notes: The table presents estimates from equation (3.3) relating total return volatility to lagged building permit growth (BPG) volatility over the sample of permit valuations from *Dun's Review*. In Panel A we present results using stock return volatility as the outcome, and instead examine bond return volatility in Panel B. We include in most specifications a set of controls for other macroeconomic observables directly related to BPG volatility, which might also drive aggregate financial market volatility and are available historically. *PopGrowth* refers to the annual population growth rate from the Census. *MktLeverage* is the aggregate market leverage ratio (i.e. debt-to-capital ratio) for CRSP firms from Graham et al. (2015). *IPGrowth* is the month-on-month growth rate in the industrial production index (INDPRO). In some specifications, we add the natural disaster and war components of the News Implied Volatility Index (NVIX) of Manela and Moreira (2017). We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. The time sample varies depending on the data availability of covariates; for instance, *MktLeverage* is available starting in 1925. t-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

4.3 The 2008 Global Financial Crisis

In this section, we find that the information embedded in building permits contains soft information about the timing of the onset and severity of the Financial Crisis, rather than simply proxying for a loosening of mortgage credit access for households during that period. Another reason to study the modern time period is that monthly frequency house price indices are available for states from the 2000s onward, and for the largest MSAs starting in the late 1980s.[28] This allows us to capitalize land values into permits at the same frequency as the permits data. Based on our option value theory framework in Section 6, the expected value of the exercised permitting option, $\mathbb{E}_t[V^*_{i,s,t+1}]$, is a function of the completed building, plus the land evaluated at its highest and best use. To the extent that construction costs are divorced from land values, relying on project cost measures reported with permit filings may reduce the informativeness of value-based BPG volatility as a signal.

Figure 6 displays the results from estimating 12-month lag versions of equation (3.3) separately by state and by MSA over the 2000s boom-bust cycle. Panels A and B do this for stock return volatility while Panels C and D use total bond return volatility as the outcome. For both bonds and equities the cumulative loadings of $\sum_{\tau=1}^{\tau^*} \widehat{\beta}_{s,\tau}$ are greater for states which had larger shares of mortgages classified as subprime during the 2000s boom, as based on the geographic subprime ranking of Mayer and Pence (2008).[29] Seven out of the top ten states, as ranked by their cumulative predictability of future stock market volatility, are also top ten subprime loan states. Similarly, all 20 MSAs in the Case-Shiller set of MSA-level indices are ranked in the top 60 by subprime loan activity, and half of these have cumulative loadings which are statistically significant at the 90% level. The ordering of states and MSAs by their cumulative loadings is similar across both asset classes.

We isolate via PCA the time series of a "subprime" factor suggested by the results of Figure 6. This factor appears as the first principal component (PC) in the time series plotted in Figure 7, and it alone accounts for 24% of the variation in $\sigma^{BPG}_{s,t}$.[30] The subprime factor jumps in lock-step with several key events of the Great Recession: the Bear Stearns failure (March 2008), the Lehman Brothers bankruptcy (September 2008), the NBER recession declaration (December 2008), and the peak volume of foreclosure auction sales in summer 2010 (Fout et al., 2017). Reassuringly, the subprime factor remains stable and negatively contributes to overall $\sigma^{BPG}_{s,t}$ during the 2001 recession, which was widely viewed as the result of the dot-com crash (Kliesen, 2003; Pástor and Veronesi, 2006), and therefore unrelated to the strength of the real estate sector.

[28] Price indices from the FHFA are available at all geographies from 1975, but only at quarterly frequency. The Freddie Mac House Price Index — whose underlying data form the basis for the FHFA indices — is available at monthly frequency for states and CBSAs starting in 1975, but relies on refinancing appraisal values. The need to match geographies and the sample of single-family homes across datasets guides our choice of indices for the post-2000s.

[29] Mayer and Pence (2008) define their ranking in terms of the fraction of residential single-family and 2-4 unit multifamily loans within a geography which are included in a subprime mortgage pool as of 2005.

[30] The seven components with eigenvalues greater than one together account for 63% of the variation in $\sigma^{BPG}_{s,t}$.

FIGURE 6. 12-Month Predictability of Permit Value Growth Volatility around the Global Financial Crisis

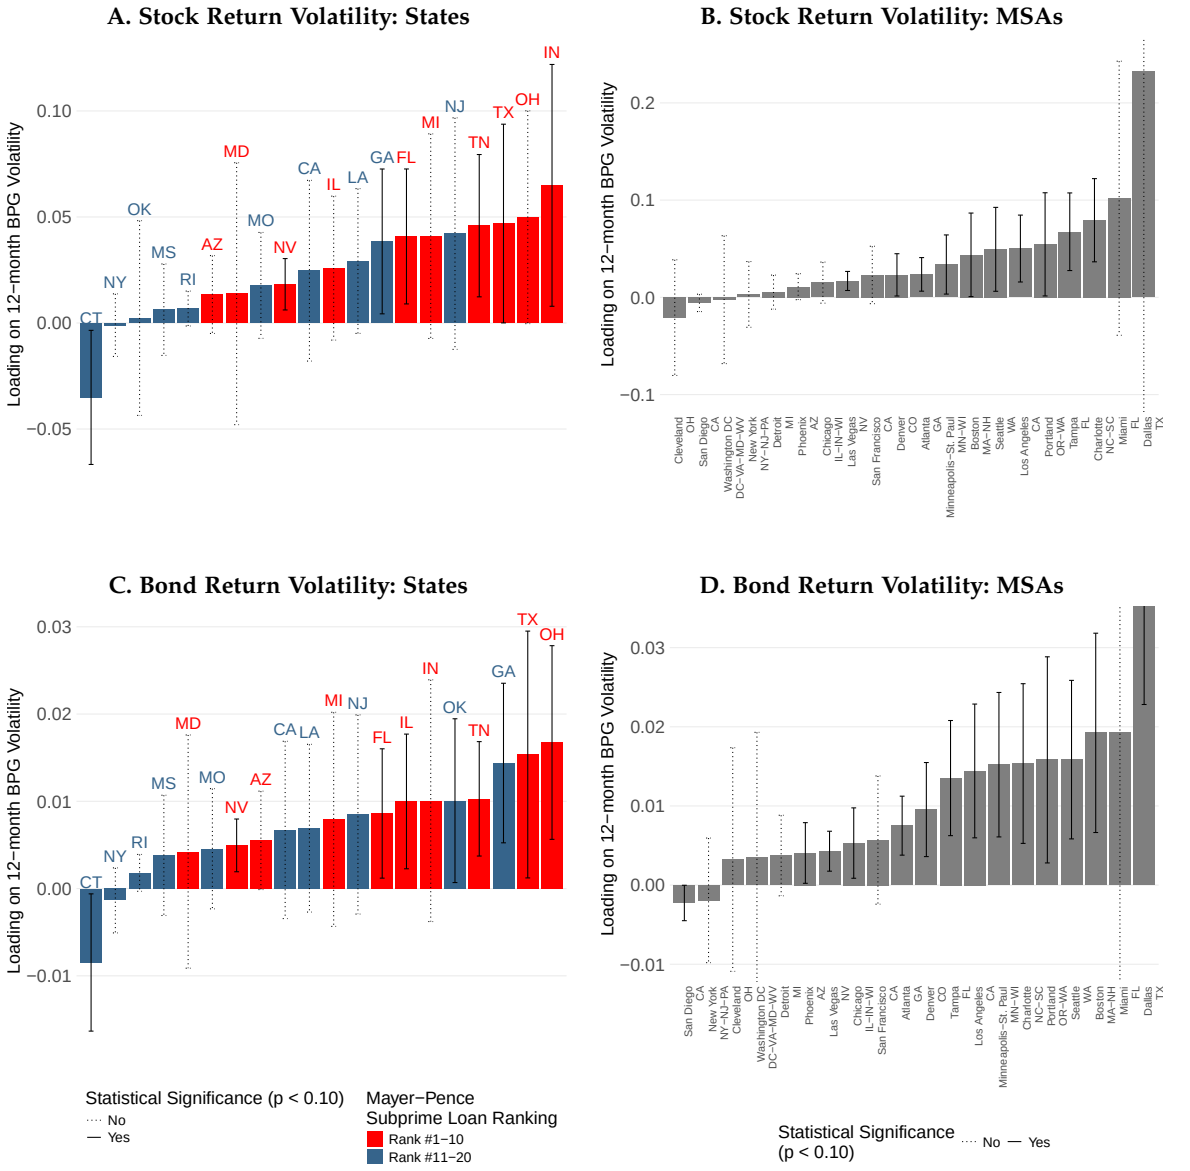

Notes: We estimate versions of equation (3.3) separately for each locality s with $\tau^* = 12$ lags of $\sigma^{BPG}_{s,t}$ and report the sum of the estimated lagged coefficients $\sum_{\tau=1}^{\tau^*} \widehat{\beta}_{s,\tau}$. We truncate the y-axis in Panels B and D to allow better visualization of the estimates and their confidence intervals. Our sample period in each specification is 2000M1 to 2019M12, for which Zillow price indices are available for each state and for each MSA included in the S&P Cotality Case-Shiller home price indices. We use the Zillow single-family residences indices (excluding condos and co-ops) to conform to the sampling conventions of the Case-Shiller indices. We use permits attached to single-family units as the quantity measure $Q_{s,t}$ in (2.2). Panels A and C perform this exercise for the top 20 states according to the Mayer and Pence (2008) subprime loan share ranking, while Panels B and D perform this exercise for the 20 Case-Shiller MSAs. We use CRSP stock return volatility as the outcome in the top two panels, and total return volatility of the Dow Jones Corporate Bond Index in the bottom two panels. We plot 90% confidence intervals obtained via standard errors to correct for autocorrelation and heteroskedasticity, selecting the minimum lag order such that the estimator for the covariance matrix is consistent (Newey and West, 1987).

FIGURE 7. Principal Components of BPG Volatility around the Global Financial Crisis

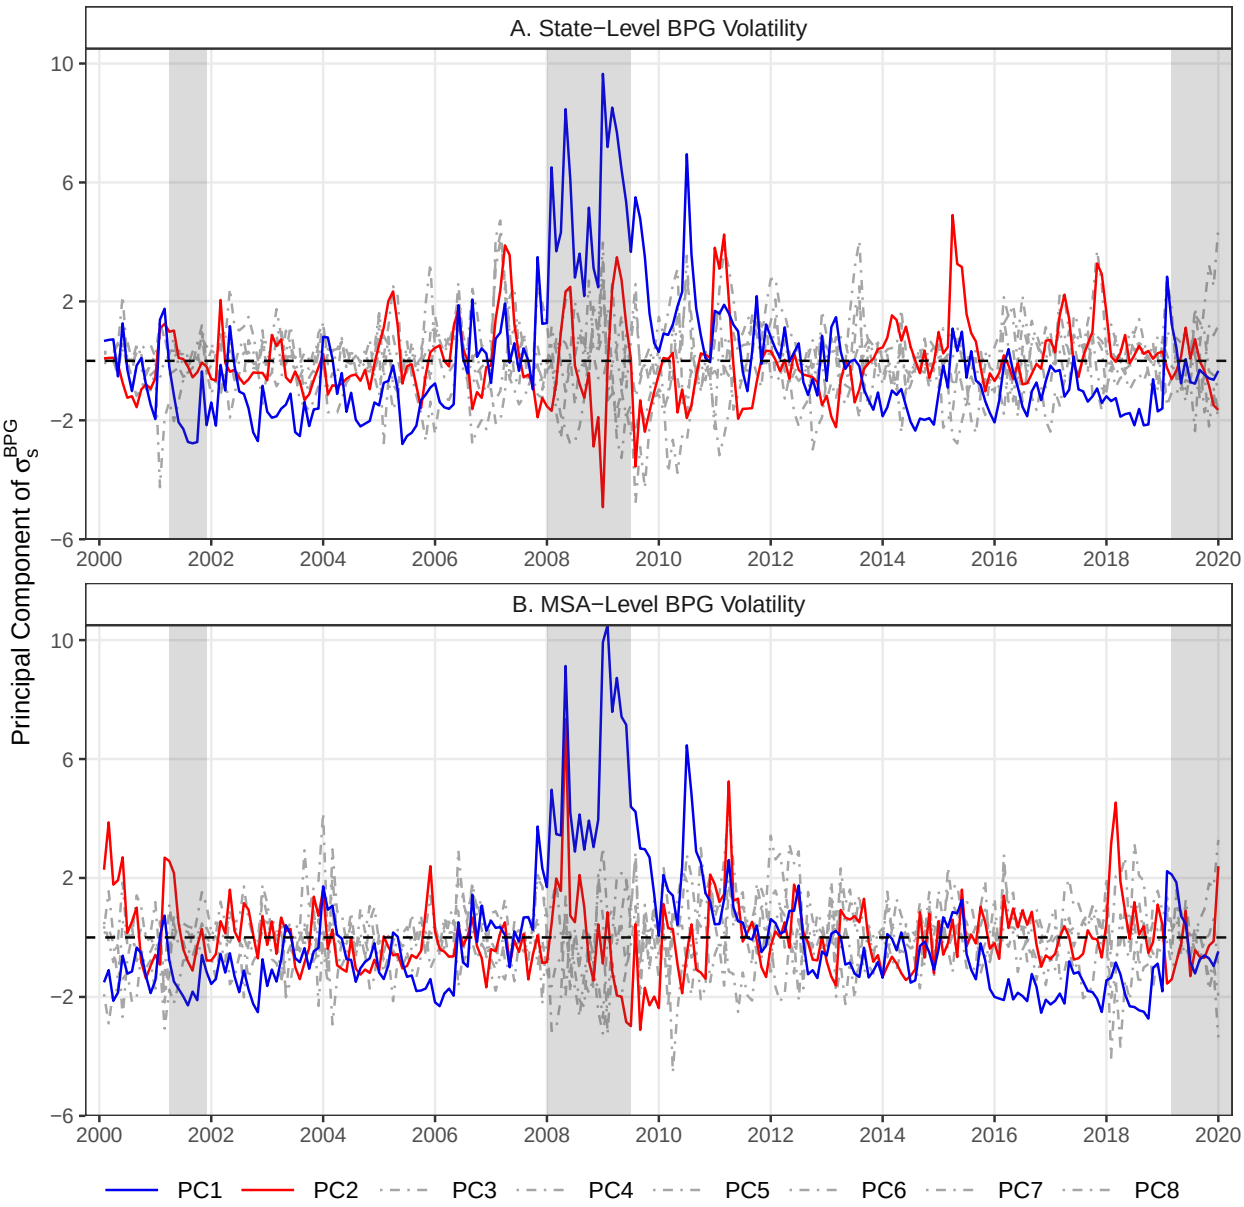

Notes: The figure plots the time series of principal components of state-level (Panel A) and MSA-level (Panel B) monthly building permit growth volatility with an eigenvalue greater than unity. Monthly building permit growth volatility is defined by $\sigma_{s,t}^{BPG}$ in equation (3.2). Our sample in Panel A includes the top 20 states in the Mayer and Pence (2008) ranking of states by their 2005 subprime loan share. We identify the subprime factor even if we pool all 47 states with sufficient permit volumes to estimate our GARCH models. In Panel B, we repeat the exercise with $\sigma_{s,t}^{BPG}$ for the 20 Case-Shiller MSAs. Grey-shaded areas indicate NBER-dated recessions.

TABLE 4. Regressions of Principal Components of Permit Value Growth Volatility on Return Volatility

Asset Market:	Equities				Corporate Bonds			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$PC_{t-1}^{(1)}$ ["subprime" factor]	0.0012***	0.0003**	0.0003**	0.0003**	0.0003***	0.0001***	0.0001***	0.0001***
	(2.78)	(2.09)	(2.06)	(2.27)	(4.45)	(2.51)	(2.44)	(2.64)
$PC_{t-1}^{(2)}$			−0.0003	−0.0003			−0.0001	−0.0001
			(1.41)	(1.35)			(1.54)	(1.63)
$PC_{t-1}^{(3)}$				0.0002				0.0001
				(0.82)				(1.36)
$PC_{t-1}^{(4)}$				0.0001				0.0000
				(0.28)				(0.55)
$PC_{t-1}^{(5)}$				−0.0002				−0.0001
				(0.77)				(1.47)
$PC_{t-1}^{(6)}$				0.0001				0.0001
				(0.53)				1.10
$PC_{t-1}^{(7)}$				0.0003				−0.0001
				(0.99)				(1.12)
Sample period	2000–2019	2000–2019	2000–2019	2000–2019	2000–2019	2000–2019	2000–2019	2000–2019
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓		✓	✓	✓
ΔHMDA $ originations			✓	✓			✓	✓
R^2	0.173	0.563	0.565	0.569	0.202	0.488	0.493	0.504
N	239	239	239	239	239	239	239	239

Notes: The table displays results from estimating regressions of the form in equation (3.3), but with just one lag for each of the principal components plotted in Figure 7 to avoid collinearity. Columns (1) to (3) use CRSP total stock return volatility as the outcome, while Columns (4) to (6) use Dow Jones Corporate Bond Index total return volatility as the outcome. Our sample period is 2000M1 to 2019M12, for which Zillow price indices are available for each state. Each regression includes a full set of month dummies to strip out seasonality. All columns with the exception of (1) and (5) include a one-month lag of asset return volatility to account for serial autocorrelation. To control for monthly growth in mortgage credit demand based on applications, or growth in equilibrium credit supply based on originated mortgages, we append nationwide totals provided by Neil Bhutta using an aggregated version of the confidential Home Mortgage Disclosure Act (CHMDA) data, downloadable at: `https://sites.google.com/site/neilbhutta/data`. We include both for-purchase and refinancing loan growth rates as separate variables in the regressions. t-statistics obtained from Newey-West standard errors where we select the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

In Table 4 we regress stock market total return and Dow Jones Corporate Bond Index total return volatility σ_t on one-month lagged PCs.[31] For both asset classes, the subprime factor is the only factor which exhibits a robust and statistically significant relationship with return volatility, regardless of the inclusion of the other principal components.[32] Importantly, controlling for growth in mortgage originations does not eliminate the ability of the subprime factor to predict the financial crisis. This suggests building permits contain soft information about impending waves of mortgage default rather than simply proxying for overall loose credit conditions.

Our results around the Great Recession period suggest that speculation explains heterogeneity in the geographic cross-section where corporate exposure to subprime loans or household leverage does not. Chinco and Mayer (2015) show that speculative house purchases by out-of-town buyers were the highest in Las Vegas, at 11% of sales, while that city's subprime loan rank is only #10. Importantly, in Las Vegas Case-Shiller single-family house prices peaked in April 2006, well before the October 2008 stock market crash. In contrast, cities like Miami (subprime loan rank #6), where speculative transactions comprised only 5% of total sales, prices continued to increase up until the Lehman failure.

After conducting a battery of robustness checks with different lag orders, we find that around the subprime episode the loading on σ^{BPG} for Las Vegas (and Nevada) is always statistically significant at the 99% level, pointing to other determinants $\mathbf{X_{s,t}}$ in the valuation model given by (6.1)–(6.3). Out-of-town speculators from the North buying in cash were an important source of price momentum in the 1920s Florida boom (Knowlton, 2020, pp. 176-177); only 40% of owner-occupied homes were mortgage-financed in 1920, with much higher average downpayments than today (Fishback et al., 2013, pp. 10-11). The predictability of BPG volatility across episodes and in pockets of the geographic cross-section where household and corporate leverage is relatively unimportant points to mechanisms highlighted in our theoretical framework, whereby building permits summarize investors' beliefs about local economic fundamentals in a way that is not perfectly collinear with build-ups in credit.

To offer further evidence of the mechanisms through which fluctuations in local permitting activity influence financial markets, we show that firms differentially exposed to BPG volatility around the Great Recession based on the location of their operations experience greater fluctuations in their returns. We estimate the shift-share version of our specification at the individual stock level from equation (3.5) and report the results in Table 5. The thought experiment underlying this specification is to compare monthly return volatility of firms with a network of operations in local markets which have larger vs. smaller recent changes in residential permits filed. To measure each stock's exposure, we take a weighted average of

[31] Given that the objective of PCA is to extract linearly independent factors, we restrict to one-month lags. Including additional lags of the contemporaneous PCs would obscure the economic interpretation, because one PC might lead a combination of the other PCs.

[32] The first principal component (i.e. our "subprime" factor) has similar time series characteristics whether we run PCA at the state level, as pictured in Figure 7, or at the MSA level using the Case-Shiller set of MSAs.

TABLE 5. Predictability of BPG Volatility in the Cross-Section of Equities

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\sigma^{BPG}_{j,t-1}$	0.0046**	0.0029**	0.0031**	0.0019*	0.0048**				
	(2.12)	(2.26)	(2.36)	(1.70)	(2.08)				
$\sum_{\tau=1}^{12} \sigma^{BPG}_{j,t-\tau}$						0.0079**	0.0057**	0.0062***	0.0100**
						(2.29)	(2.04)	(2.71)	(2.43)
Time sample	1989-2019	1989-2019	1989-2019	1989-2019	2000-2019	1989-2019	1989-2019	1989-2019	2000-2019
Share weights ω_k	Emp	Emp	Emp	Sales	Emp	Emp	Emp	Sales	Emp
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓
Firm controls			✓	✓	✓		✓	✓	✓
# of firms	2,067	2,066	1,865	1,865	1,280	1,865	1,713	1,713	1,174
N	157,040	156,907	135,808	135,808	73,832	132,342	117,345	117,345	65,348
Adj. R^2	0.31	0.40	0.43	0.43	0.35	0.33	0.42	0.42	0.35

Notes: The table displays results from estimating regressions of the form in equation (3.4), where the outcome in each regression is total return volatility for a stock associated with company j appearing in the matched CRSP/COMPUSTAT/DnB sample described in Section 2.5. We restrict attention to post-1989 DnB observations, since the number of firms included in the dataset stabilizes in that year. Columns (4) and (8) use shares of sales across states within each plant's network, while all other columns use employment shares. All columns with the exception of (1) and (6) include a one-month lag of asset return volatility to account for autocorrelation. The vector of firm controls includes deciles of firm age based on the listing year and balance sheet size, and one-year lags of EBITDA, Tobin's Q, and the leverage ratio. We follow standard procedures in the corporate finance literature to construct EBITDA, the Q ratio, and the leverage ratio from COMPUSTAT balance sheet items, winsorizing all items at $\pm 5 \times IQR$. t-statistics obtained from robust standard errors clustered by firm in parentheses. Clustering standard errors at the stock level results in more conservative (i.e. wider) confidence intervals. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

state-level BPG volatility, $\sigma^{BPG}_{j,t} \equiv \sum_{k \in \mathcal{J}} \omega_{k,t-\tau-1} \cdot \sigma^{BPG}_{k,t-\tau}$, for firm j in month t. We use *ex ante* notions of firm state-level locations k for the sales or employment-based weights $\omega_{k,t-\tau-1}$ to account for the possibility that firms may shift their operations to less-exposed parts of the country due to either supply chain risks or reduction in local demand picked up at high frequency by σ^{BPG} (Giroud and Mueller, 2019).

Table 5 documents that firms physically located in markets with sharper jumps in permitting activity in recent months experience greater own-stock return volatility; this is true regardless of whether we use employment or sales-based weights to apportion firms' BPG risk across locations. While we lack data on firm locations prior to 1989, over the last three decades, this relationship between BPG volatility and return volatility is driven by events around the Global Financial Crisis, as shown by columns (5) and (9), which restrict to a tighter window around 2008. The relationship is robust to autocorrelation in returns, seasonality, and *ex ante* controls for corporate fundamentals, including firm size, age, EBITDA, Tobin's Q, and the leverage ratio. The pass-through of building risk to return volatility is cumulative over longer time horizons, with the elasticity doubling from the one-month to the twelve-month horizon of BPG risk exposure. Overall, examining the stock cross-section points to corrections in building activity as a harbinger

of decline in local demand for a firm's products. When aggregated up, individual firms' exposure to future declines in local demand translate to volatility in the overall stock market.

5 Discussion: Why Do Local Building Permits Matter for Financial Market Fluctuations?

Our results demonstrate that fluctuations in local building permit quantities predict stock and bond return movements at short and medium-run horizons both in terms of elasticities and incremental R-squared. This finding persists across recession episodes, including during time periods when credit markets were less integrated and mortgages were less common or required households to make large downpayments. Further, the aggregate predictability of building permit growth (BPG) volatility holds conditional on measures of corporate and household leverage and debt service ratios. Hence, the predictability of BPG volatility for asset market fluctuations arises independently of the credit cycle documented in the literature (e.g., Schularick and Taylor, 2012; Jordà et al., 2013; Müller and Verner, 2024). In this section, we aim to understand why that is the case.

5.1 Permits as a Proxy for Investors' Beliefs about Fundamentals

One interpretation, which we formalize in Section 6, is that BPG volatility should predict stock and bond return movements because it encapsulates real estate investors' beliefs about local fundamentals which may be otherwise difficult for investors in risky asset markets to fully observe at high frequency. For instance, state-level employment statistics from the Bureau of Economic Analysis (BEA) are available at the quarterly frequency starting in 2018, yet only at the annual frequency prior to 2018. Moreover, the data are released with lags as long as a year and often revised in between release dates, indicating possible data quality issues as local government offices report information in a staggered fashion. The Quarterly Census of Employment and Wages (QCEW) produced by the BLS is subject to preliminary release lags of five months after the end of a quarter.[33] Similarly, corporate investment rates, while forward looking, can only be constructed at the quarterly frequency using Form 10-Q releases, which firms release with lags of up to 45 days.

One threat to this interpretation is that local BPG volatility proxies for risks to corporate assets and/or physical operations. This could be the case if, for example, natural hazards or the potential outbreak of conflict result in a shutdown of building or reallocation of resources from private towards public objectives. At an aggregate level, sources of physical risk to firms' operations are contained within components of the NVIX of Manela and Moreira (2017), and the aggregate predictability of BPG volatility remains intact conditional on these components.

[33] See the QCEW release calendar over the last decade: https://www.bls.gov/cew/release-calendar.htm.

To see where in the industry cross-section this predictability operates, we re-estimate equation 3.3 separately for each of the 17 Fama and French (1997) industry portfolios. We estimate the regression over the 1961M02–2019M12 estimation sample and within twelve non-overlapping five-year windows, labeled by the lustrum to which they belong, using SFH BPG volatility on the right-hand side to match Panel B of Table 2 (see Appendix D.5). In the full sample, lagged SFH BPG volatility positively predicts return volatility for all 17 industries, and the coefficient is statistically significant at the 10% level for 15 of them.

The largest loadings are uniformly on cyclically-sensitive industries: motor vehicles, mining, consumer durables, oil, transportation equipment, residential construction, fabricated metal inputs, mortgage finance, and chemicals. More defensive industries, such as retail, food products, nondurables, and utilities, display smaller loadings. The prominence of motor vehicles and transportation equipment is particularly notable. Cortes and Weidenmier (2019) show that the growth rate of U.S. truck production is the single aggregate leading indicator whose inclusion most attenuates BPG volatility's predictive power for stock volatility around the Great Depression. More broadly, the cross-sectional ordering aligns with the signal extraction mechanism we propose in Section 6. If local residential permits serve as a forward-looking public signal about real economic fundamentals, BPG volatility should load most strongly on industries whose cash flows are most sensitive to the underlying cyclical state.

5.2 Role of Housing Supply Constraints

Regulatory reforms or political upheavals are likely too slow-moving to explain the predictive power of permit activity in our one-hundred-year time series. The Wharton Residential Land Use Regulatory Index (WRLURI), which is a survey-based measure of local political constraints developers might face in seeking permits for new construction, has changed little between the first version produced in 2006 (Gyourko et al., 2008) and the 2018 update (Gyourko et al., 2021); there is virtually no change in the ranking of states and 44 CBSAs included in the Wharton survey panel. Similarly, Ganong and Shoag (2017) compute the total annual number of per capita court cases mentioning the phrase "land use" from 1940 to 2010. This measure of regulatory stringency exhibits a high degree of autocorrelation in the post-1960s period, with the cross-decade correlation in the ranking of states based on per capita land use court cases steadily increasing over time.

A key implication of the framework in Section 6 is that the signal-to-noise ratio of local BPG volatility is related to local constraints on real estate development. Intuitively, if constraints on the number of new permits are binding in jurisdiction s, then the signal $q_{s,t} = \Delta \log Q_{s,t}$ will be right-censored, leading to many periods of negative or near-zero growth in permitting. Without conditioning on any other observable information about the path of economic fundamentals in s, this reduces the informativeness of BPG in s as a signal to investors, since it becomes more

FIGURE 8. Correlation between Stringency of Land Use Regulations and Cumulative Permit Volume

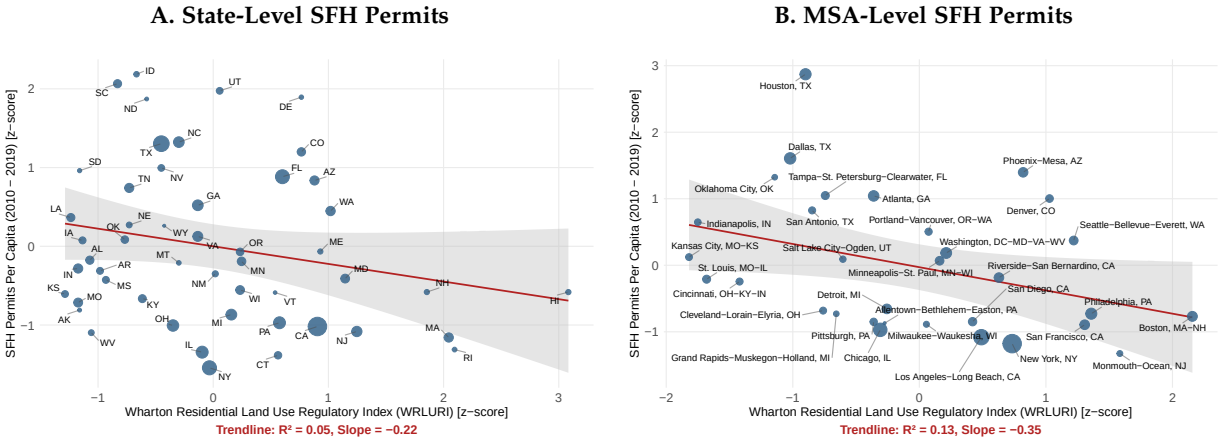

Notes: Each panel in the figure shows the relation between the Wharton Residential Land Use Regulatory Index (WRLURI) on the x-axis and cumulative single-family home (SFH) building permits per capita issued during the post-GFC recovery from 2010M1 to 2019M12. We scale each building permit series by 2010 decennial Census population, downloaded from IPUMS. We normalize both variables to z-scores such that the trendline slope in each plot represents the cross-sectional correlation. Panel A tabulates the data at the state-level, while Panel B tabulates at the MSA level. We downloaded the state and MSA-level WRLURI from Tables 10 and 11 of Gyourko et al. (2008). Of the 47 MSAs with a WRLURI value, we include the 35 MSAs with available data on permits from the Census BPS.

difficult to know whether developers forecast reduced economic activity in s or if they would otherwise invest in s but are unable to, and therefore decide to develop elsewhere.

To test this intuition, we use the first edition of the WRLURI, as the literature lacks a comprehensive measure of local constraints on housing development going back further in time. Figure 8 shows a strong negative correlation across states and MSAs between an area's regulatory stringency, as measured by the first Wharton Index, and per capita permit issuance during the post-Great Recession boom from 2010 to 2019. The resulting correlation of the precision (inverse standard error) on the one-month lagged loadings on σ^{BPG} plotted in Figure 4 for stocks with the WRLURI is -17%; for bonds, this same correlation is -22%. For the precision on the cumulative 12-month lagged loadings the correlations are -19% for stocks and -21% for bonds.

Bartik et al. (2024) show that the first and second principal components in their more recent generative AI-based measure of the stringency of local land use regulation — corresponding to regulatory complexity and exclusionary zoning (e.g., minimum lot sizes), respectively — are also strongly negatively correlated with average residential permits filed at the municipality level in the Census BPS during 2019 to 2023. This is true despite the fact that the WRLURI and generative AI-based index have 33% correlation for the first PC but only 11% for the first PC when computed against the overall Wharton Index. In Appendix E, we compare permit completion rates across states using the permit panel microdata from Cotality *Building Permits*; completion rates are lower in states facing more constraints on new building according to both types of regulatory indices as well as the buildable land share measure of Lutz and Sand (2023).

FIGURE 9. Recent Permitting Activity in Top Cities for Work-From-Home Migration

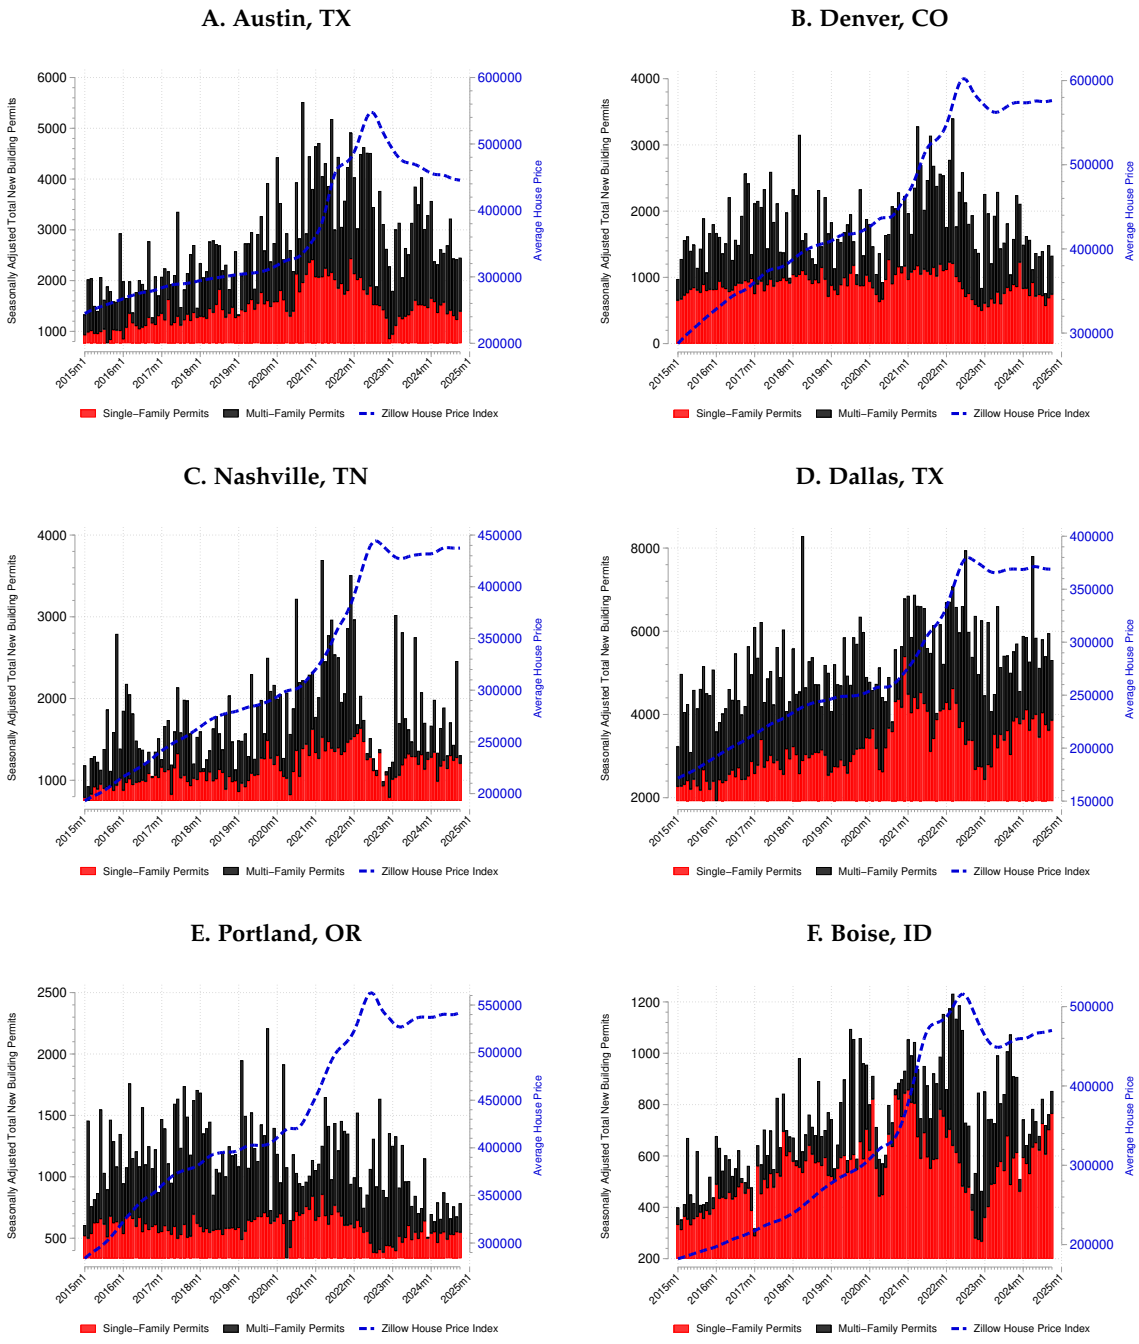

Notes: Each panel plots for a metropolitan statistical area (MSA) seasonally adjusted residential permits from the Census Building Permit Survey (left-hand y-axis) against the seasonally adjusted Zillow Home Value Index for the mid-tercile of the price distribution of single-family homes (right-hand y-axis). We restrict the time period to 2015M1–2024M10 to focus on market corrections from net migration flows of work-from-home residents. Each of the cities listed experienced large net migration of remote workers based on American Community Survey data, with cities ordered by their rank in terms absolute net migration during the period 2020–2021.

In our specifications the predictive power of BPG volatility holds conditional on past local population growth. This is important to the extent that permitting may reflect developers responding to *realized* rather than forecasted demand for new housing units. In U.S. historical booms, uncertainty about future demand stemming from the discovery of new territory results in over-development and subsequent housing busts (Glaeser, 2013), even without overly optimistic beliefs. Overbuilding risk is exemplified by recent reversals in house prices and rents in markets like Austin, Texas (Wall Street Journal, 2024), which were initially major beneficiaries of migration induced by the work-from-home boom. In Figure 9, we show that collapses in seasonally adjusted residential permit volume in the latter half of 2022 precede similar dips in house prices among cities which initially saw an influx of digital nomads, as mobile workers began to leave these cities amid waning COVID-19 incidence and the rollout of return-to-office policies (Flynn et al., 2024).[34]

Finally, by examining the predictability of permitting activity, we focus on the very initial stage of housing development. A possible source of uncertainty captured by our BPG factor is uncertainty about the time between exercising of the option to build and the completion of housing units. Time-to-plan lags in commercial real estate are long and variable across location, averaging 1.5 years (Glancy et al., 2024), a finding we replicate for multifamily units. Gabriel and Kung (2024) show in project approval microdata for Los Angeles that lengthy approval times for new housing stifle housing supply. Consistent with those studies, we uncover a cross-sectional pattern of low-BPG areas having greater predictability for financial volatility, as well as more relaxed regulatory constraints on new building and reduced lags between residential permitting and project completion.

5.3 Alternative Interpretations

An alternative explanation is that building permits chiefly reflect extrapolative beliefs in local housing markets. This interpretation is plausible in boom-bust environments where land and new construction are especially sensitive to optimistic beliefs about capital gains (Nathanson and Zwick, 2018), and it connects our evidence to broader work on extrapolative expectations in asset markets (Greenwood and Shleifer, 2014; Barberis et al., 2018). Yet the patterns in our data are not naturally explained by a predominantly price extrapolation story. In the modern period, the predictive content of permits for financial markets is driven much more by the quantity component of permitting than by local house price growth alone, suggesting that BPG captures shifts in planned real activity and beliefs about local fundamentals rather than simply extrapolation from recent capital gains.

A second interpretation comes from the real options literature, where greater uncertainty can raise the value of waiting and delay construction (Bulan et al., 2009). Under that view, periods

[34]Qualitatively similar patterns showing a predictable dip in prices emerge if we replace the Zillow Home Value Index in Figure 9 with the Zillow Observed Rent Index (ZORI). Each of the cities pictured also ranks highly in terms of net migration flows based on tax filers (Berube, 2024).

of high BPG volatility should be accompanied by lower permit issuance or longer gaps between permit issuance and project completion. We do see some evidence that completion delays matter more for multifamily projects, where planning horizons are longer. However, this mechanism does not fit the single-family segment that anchors much of our main analysis.

Using Cotality permit microdata, Figure E.1 shows that residential permit completion shares at fixed horizons are broadly flat around the 2008–2009 Global Financial Crisis; despite sharply rising uncertainty, around 80% of new residential permits issued since 1990 are completed within 12 months on average. While we lack individual permit-level information over many boom-bust episodes, this evidence suggests that cyclical delays in completing already-permitted single-family projects are unlikely to be the primary driver of the predictive relation emphasized in our main results. Taken together, the microdata suggest that permits proxy for both near-term housing supply and forward-looking beliefs about local fundamentals, rather than primarily measuring developers' option value of waiting.

6 Theoretical Framework

We close the paper by showing how a simple model of housing development and information aggregation rationalizes the empirical patterns established above. The model features two asset markets and two periods, t and $t+1$. In period t developers make decisions on whether or not to acquire permits to construct housing in each *local* market. Investors trade in nationwide financial markets after observing in t prices and signals about future period $t+1$ dividends. Building permit volume forms a local quasi-public signal observed by some investors playing an investment game with information aggregation à la Grossman and Stiglitz (1980). We use the framework to organize the evidence into four testable hypotheses which are validated by our data.

Housing development. Let $i \in [0,1]$ index a unit mass of potential housing market investors, and $s \in \{1,\ldots,\mathcal{S}\}$ index the locality (e.g., a state, MSA, or county) where the property is located. Developable land is in fixed supply $\bar{L}_s < 1$, and each investor can hold a permit on at most one parcel. Standard real option value theory (OVT) says that the value of holding entitled land is determined by the earnings potential of the underlying parcel less any construction costs required to deploy the land at its highest and best use (Geltner et al., 2014). Therefore, the expected value of a permit option if exercised depends on the probability of the project's success $f(\mathbf{X_{s,t}})$, the construction cost $C_{i,s,t+1}$, and the market value of the land $L_{i,s,t+1}$ plus the new building $B_{i,s,t+1}$ on top of it.

Suppose construction costs are paid in period $t+1$ but known in t. If the project is successful, then the property will be valued at its prevailing market price $B_{i,s,t+1} + L_{i,s,t+1}$. According to standard development option pricing models, this land *vs.* building distinction is important for

extracting forward-looking sources of macroeconomic risk because uncertainty about payoffs from exercising the option is positively linked to land values (Titman, 1985; Cunningham, 2006). This setup leads us to the expected value of an exercised permit $\mathbb{E}_t[V^*_{i,s,t+1}]$:

$$\mathbb{E}_t[V^*_{i,s,t+1}] = f(\mathbf{X_{s,t}}) \cdot \mathbb{E}_t[B_{i,s,t+1} + L_{i,s,t+1}] - C_{i,s,t+1}, \tag{6.1}$$

The success probability $f(\mathbf{X}_{s,t})$ depends on a vector of potentially time-varying factors $\mathbf{X_{s,t}}$, such as local macroeconomic fundamentals (e.g., unemployment) or the history of local weather and regulatory shocks which could lead to delayed or rescinded approvals. A key feature of our model — and supported by the data, as we show in Section 4 — is that developments in housing markets (Main Street) are predictive of financial market (Wall Street) movements due to potentially unobserved elements of $\mathbf{X_{s,t}}$.[35]

We can further simplify (6.1) by invoking a replacement cost approach to valuing buildings, as adopted by *Dun & Bradstreet's* for their permit series that we use in our pre-1960s analysis. If buildings are always valued at their replacement cost, assuming the construction cost is inclusive of any teardown costs or administrative fees, then $B_{i,s,t+1} = C_{i,s,t+1}, \forall i$.[36] Suppose further that housing production is Cobb–Douglas, so land values are proportional to the attached structure's value: $L_{i,s,t+1} = \varphi_{i,s} \cdot B_{i,s,t+1}$. Equivalently, the fraction of land in the housing production function is $\varphi_{i,s}/(1 + \varphi_{i,s})$, and it varies by parcel. The land value fraction is known to the landowner, since each local assessor office splits the value of a parcel into building and land components for levying property taxes.[37]

With these two simplifying assumptions, we can write the expected value of the exercised option and current permit value $V_{i,s,t}$ as:

$$\mathbb{E}_t[V^*_{i,s,t+1}] = \Big(\varphi_{i,s} \cdot f(\mathbf{X_{s,t}}) + (f(\mathbf{X_{s,t}}) - 1)\Big) \cdot C_{i,s,t+1} \tag{6.2}$$

$$V_{i,s,t} = \max\{0, \mathbb{E}_t[V^*_{i,s,t+1}]\} \tag{6.3}$$

Equation (6.3) follows the principle that the value of a permit depends only on the outcomes that result in the call option with a strike price of $C_{i,s,t+1}$ finishing in the money, otherwise it

[35]The assumption that construction costs are known in advance simplifies the exposition, but does not affect the model's main empirical predictions. If we relax the assumptions that construction costs are known in advance and project success depends only on realized fundamentals today, the expected value of the option in (6.1) becomes:

$$\mathbb{E}_t[V^*_{i,s,t+1}] = \mathbb{E}_t\Big[f(\mathbf{X_{s,t+1}}) \cdot (B_{i,s,t+1} + L_{i,s,t+1} - C_{i,s,t+1})\Big],$$

with the discount rate normalized to $r = 0$. Hence, it is without loss of generality to model the developer's problem as either an NPV decision rule or exercising a contingent claim on future states.

[36]In practice, investors also incur small fees for filing permit applications or obtaining certificates of occupancy, both of which we have implicitly rolled into $C_{i,s,t+1}$. We discuss these fees in Section 2.3.

[37]We have assumed that housing investors have rational expectations with respect to building values and that they use the replacement cost approach. If it is the case that $\varphi_{i,s} = \varphi_s, \forall i$, then in a given area all housing market investors either permit or do not permit, up to the land endowment constraint $\bar{L}_s$.

has zero value (Natenberg, 2014). Combining (6.2) with (6.3), we learn that investors demand a permit whenever a project has positive expected value — that is, if and only if $f(\mathbf{X_{s,t}}) > 1/(1+\varphi_{i,s})$. We also observe that standard building permit series which incorporate construction costs are proportional to $V_{i,s,t}$, but may miss the component of expected value originating from the risk factors embedded in $f(\mathbf{X_{s,t}})$.[38]

Aggregating up equation (6.3) and the decision embedded within it, we obtain aggregate local permit values and permitting activity, respectively:

$$V_{s,t} = \int_i V_{i,s,t} \cdot di \tag{6.4}$$

$$Q_{s,t} = \int_i \mathbb{1}\{V_{i,s,t} > 0\} \cdot di \leq \bar{L}_s \tag{6.5}$$

where $Q_{s,t}$ is equivalent to the total number of expected NPV positive projects in any period t. $Q_{s,t}$ is bounded above by the developable land endowment in s.[39] We will posit that $Q_{s,t}$ and permit volume growth rates $q_{s,t} = \Delta \log Q_{s,t}$ form quasi-public signals observed by informed investors in financial markets, while permit values $V_{s,t}$ are not readily discernible from public data. We justify our focus on $Q_{s,t}$ rather than $V_{s,t}$ as a signal based on the data compilation efforts we describe in Section 2. From equation (6.2), $Q_{s,t}$ aggregates local housing market investors' beliefs about local economic fundamentals $f(\mathbf{X_{s,t}})$, which are not directly observable to investors in financial markets.

Financial markets. We adopt the standard Grossman and Stiglitz (1980, hereafter GS) setup to characterize information aggregation in financial markets. There is a risky asset, such as corporate bonds or equities, that pays an unknown dividend d in period $t+1$, where $d_{t+1} \sim \mathcal{N}(\bar{d}, \sigma_d^2)$. There is a unit mass of investors $j(s) \in [0,1]$ in each locality s who trade the risky asset in period t at price p_t. To resolve the Grossman–Stiglitz paradox, the supply of the risky asset A is random so that prices are not perfectly informative:

$$A = m + u, \quad \text{with } m > 0 \text{ and } u \sim \mathcal{N}(0, \sigma_u^2) \tag{6.6}$$

where u denotes noise due to either unmodeled noise traders dumping shares into the market or firms issuing a random number of securities. There is also a riskless asset with a rate r between t and $t+1$. To limit the notational burden, in what follows we suppress the time subscripts.

[38] Davis and Heathcote (2007) show that the share of land in aggregate U.S. housing value has increased since the 1970s, primarily due to declining substitutability of new homes for old homes. The assumption of a static φ is more likely to hold for a context like 1920s Florida, where most housing was newly built. On average, between 1975 and 2006, 36% of aggregate U.S. housing value is land, implying $\varphi = 0.56$.

[39] We normalize the land endowment (and, consequently, permit volume) to a unit scale. One can interpret this convention as $\bar{L}_s$ is the share of developable land akin to the Saiz (2010) physical geography-based measures. $Q_{s,t}$ then represents the share of developable land newly permitted for development in period t.

Investors in the risky asset market have constant absolute risk aversion (CARA) utility over consumption with coefficient of absolute risk aversion γ:

$$-\mathbb{E}[\exp(-\gamma \cdot c_j)]$$

There is asymmetric information in the risky asset market. A fraction λ_s of investors are informed and observe a signal about the dividend, and this signal is common to all informed investors operating in s. The other $(1-\lambda_s)$ fraction do not observe the signal, and are therefore uninformed. All investors observe the price p. Note at this stage that while information is aggregated at a local level, the asset market is national, so there is only one price p which clears the market. This distinction is key for the model's ability to explain the heterogeneous effects of local housing markets on asset markets that we find in the data.

Let building permit growth (BPG) q_s serve as this quasi-public signal, aggregated from the local housing market according to equation (6.5). BPG relates to dividends according to:

$$q_s = d + \varepsilon_s, \quad \text{with } \varepsilon_s \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2_{q(s)}) \tag{6.7}$$

where ε_s is independent of d. The idea behind local BPG as a quasi-public signal for future dividends is that positive swings in permitting activity relative to trend reflect developers' positive beliefs about a project's future values given local fundamentals $\mathbf{X_{s,t}}$ which determine the project's probability of success via equation (6.2). For exposition, we assume the ε_s are i.i.d. across localities. In Appendix A.3, we present a version of the model which allows for permit signals $q_{s,t}$ to be correlated across local markets — in which case $q_{s,t}$ is drawn from a joint distribution.

A key reason why we define the quasi-public signal in growth rates rather than levels is that q_s has full support despite the fact that Q_s must be non-negative and is bounded above by the local land endowment. To see this, note that the maximum growth rate in local permits is $\log(\bar{L}_s) - \log(0) = +\infty$ and the minimum growth rate is $\log(0) - \log(\bar{L}_s) = -\infty$, implying that q_s follows a standard normal. In contrast, the *level* of building permits Q_s as a quasi-public signal would instead follow a truncated normal distribution in each locality with two truncation points at 0 and $\bar{L}_s$. Several authors have shown that extensions of the canonical [GS] paradigm featuring CARA utility with truncated normality produce the usual demand function with a mean-variance term plus a correction term for the truncation (Yuan, 2005; Pálvölgyi and Venter, 2015). The correction term is a non-linear function of the truncation points, mean, variance, and risk tolerance.[40] Hence, the scope for multiple equilibria is more limited by specifying BPG rather than permit levels as the quasi-public signal in an asymmetric information context.

[40]The penalty term arises because there is a mass of ε_s which deliver the same equilibrium asset price, meaning the uninformed investors' demand function is no longer fully linear. We relegate to Appendix A the derivations for versions of our model with permit levels as the signal.

Noisy Rational Expectations Equilibrium. Under these conditions, we can define a noisy rational expectations equilibrium (NREE) as a price function $p(\{q_s\}_{s=1}^{\mathcal{S}}, u)$ and set of demand functions $x_{j(s)}$ for the informed (I) and uninformed (U) investors $j(s)$ with information set $\omega_{j(s)}$ satisfying:

$$\textit{Portfolio optimization:} \quad x_{j(s)} = \frac{\mathbb{E}[d|\omega_{j(s)}] - (1+r)\cdot p}{\gamma \cdot Var[d|\omega_{j(s)}]}, \quad j \in \{I, U\} \tag{6.8}$$

$$\textit{Market clearing:} \quad \sum_{s=1}^{\mathcal{S}} \Big[\lambda_s \cdot x_I(q_s, p(q_s, u)) + (1-\lambda_s)\cdot x_U(p(q_s, u))\Big] = m + u \tag{6.9}$$

$$\textit{No cross-market arbitrage (law of one price):} \quad p_s = p, \forall s \tag{6.10}$$

Proposition 1. *The price function which satisfies the three conditions for a noisy rational expectations equilibrium is linear in the local signal q_s and noise u and follows:*

$$p = \phi_0(s) + \phi_q(s)\cdot(q_s + \phi_u(s)\cdot u), \forall s \tag{6.11}$$

Moreover, $\phi_q(s) > 0$ and $\phi_u(s) < 0$, regardless of the coefficient of absolute risk aversion γ, so the asset price loads positively on building permit growth in each locality and negatively on noise.

Proof. We derive the expressions for the coefficients $[\phi_0(s), \phi_q(s), \phi_u(s)]$ in Appendix A.1. □

The coefficients $[\phi_0(s), \phi_q(s), \phi_u(s)]$ vary by locality through the fraction of informed investors λ_s and BPG volatility $\sigma_{q(s)}$. This means that even if all areas experienced the same permit growth rate q_s, there would still be heterogeneity in the informativeness of the local BPG signal for asset prices (and returns). The coefficients in equation (6.11) are functions of the precisions of the dividend, the signal, and the transformed signal $\widetilde{p}$. Uninformed investors only observe the price, but in equilibrium this is a transformed version of the local quasi-public signal and noise:

$$\widetilde{p} = \frac{p - \phi_0(s)}{\phi_q(s)} = q_s + \phi_u(s)\cdot u \tag{6.12}$$

The transformed signal based on the price therefore has variance $\widetilde{\sigma}_p^2 = \sigma_{q(s)}^2 + \phi_u^2(s)\cdot\sigma_u^2$.

The derived relation between equilibrium asset prices and local housing markets leads to four main testable implications which we take to the data:

Hypothesis 1: Building permits proxy for local economic fundamentals. This follows from the option value theory principles underlying the market for building permits. Developers apply for permits when the project is NPV > 0, and based on (6.2), this is more likely to be true when local fundamentals $\mathbf{X_{s,t}}$ deliver a higher probability the project comes to fruition. If there is a large enough improvement in local fundamentals between periods, then $q_{s,t,t+1} > 0$.

Hypothesis 2: Movements in permit volumes positively predict aggregate asset price and return movements. It follows from the equilibrium price function and proof to **Proposition 1** that $\partial p/\partial q_s = \phi_q(s) > 0, \forall s$. Hypothesis 2 is validated by our results for total returns in Table 1.

Hypothesis 3: The sign of comovement between BPG volatility and asset price or total return volatility is heterogeneous across localities. A related corollary says:

Corollary 1.1 (Comparative Statics). *Given the equilibrium price function in equation (6.11) and the definition of the transformed price signal in (6.12):*

1. *Let σ_p denote the volatility of the equilibrium risky asset price. $\partial\sigma_p/\partial\sigma_{q(s)}$ has an ambiguous sign, but is positive for sufficiently small local BPG volatilities $\sigma_{q(s)}$.*
2. *Normalize the* ex ante *risky asset price to be $p_t = 0$, so that the total return can be written as $r_A = p_{t+1} + d_{t+1}$, with variance $\sigma_r^2 = \sigma_p^2 + (1 + 2\phi_{q(s)}) \cdot \sigma_d^2$. Then $\partial\sigma_r/\partial\sigma_{q(s)}$ has an ambiguous sign, but is positive for sufficiently small local BPG volatilities $\sigma_{q(s)}$.*

Proof. We show the full comparative statics with respect to σ_q in Appendix A.1. □

Corollary 1.1 indicates that we should expect some heterogeneity across locations in the signs of the loadings of risky asset price and asset return volatility on local BPG volatility.[41] Intuitively, adding volatility to an ultra-precise q_s raises price variance because it amplifies the impact of noise in the risky asset market. Indeed, we find in our main results of Section 4 that localities with relatively large average BPG volatility (e.g., Connecticut) have low values of $\partial\sigma_r/\sigma_{q(s)}$, while states and MSAs with small average BPG volatility (e.g., Georgia) have high predictability. Probing further why areas differ in equilibrium permits leads us to the final testable prediction offered by our model.

Hypothesis 4: All else held equal, the precision of BPG as a signal for aggregate asset price movements is negatively correlated with physical geography and regulatory constraints on local real estate development. In the baseline version of the model $\bar{L}_s$ is independent of the information structure. We relax this assumption in Appendix A.4 by endogenizing the construction cost $C_{i,s,t+1}$ term from the profit-maximization problem of a developer who faces a production function with decreasing returns to scale in local labor inputs. Heterogeneity in local housing supply elasticities manifests both via $\bar{L}_s$ and a "profit wedge" in the developer's production function arising from local building restrictions based on regulatory barriers to new construction. This addition to the model captures the observed negative relationship between building permit activity and the stringency of local regulations on new housing supply at both the state and sub-state levels (Gyourko et al., 2008; Gyourko et al., 2021; Bartik et al., 2024).

[41] We resolve some of this ambiguity in Appendix A.2 by making two adjustments to the model. One is to feature geography-linked supply noise $\sigma_u^2(s)$ in which supply noise is greater in more populous areas. Another is to endogenize the information mass by allowing the number of informed traders to fall as permits become noisier. This latter version can be rationalized by informed investors ignoring a sufficiently noisy signal altogether.

We give here an intuitive example of what happens when housing supply constraints enter into the information structure. Consider again a high BPG volatility state like Connecticut and a low volatility state like Georgia. Connecticut also faces more physical constraints on new housing development than Georgia; according to the buildable land share measure produced by Lutz and Sand (2023), Connecticut has $\bar{L}_{CT} = 0.51$ (primarily driven by prior housing development) and Georgia has $\bar{L}_{GA} = 0.68$.[42] Suppose the local economies of the two states experience prolonged booms, leading to an increase in the probability of project success, $f(\mathbf{X_{s,t}})$. This will result in an expansion of demand for permits in both states, but the constraint on permit *levels* in (6.5) is always more likely to bind for CT, thus increasing $\sigma_{q(CT)}$. Informed investors extract more information from swings in permits in Georgia given that there is a smaller steady-state inaction region. Hence, BPG volatility will be a more precise signal relative to noise for states like Georgia than Connecticut.

Accounting for both physical and regulatory sources of supply restrictions is essential because our results on the linkages between real estate and financial markets hold longitudinally. Since the 1970s represent an inflection point for the complexity of laws governing new building approvals (Shertzer et al., 2022), regulatory constraints cannot fully explain the patterns we observe in the data for earlier financial crises such as the Great Depression. Further, physical and modern regulatory restrictions are positively correlated. The MSA cross-sectional correlation between the Saiz (2010) undevelopable land share and the 2006 version of the Wharton Residential Land Use Regulatory Index is 33%.

The model provides a compact interpretation of the evidence. The option-value block rationalizes why permit issuance and completion speak to local fundamentals. The Grossman–Stiglitz block explains why that information can forecast national stock and bond volatility, while the comparative statics explain why loadings differ across places. Finally, adding supply constraints helps explain the weaker signal we document in more regulated or physically constrained markets.

7 Conclusion

We provide evidence that housing market activity, as measured by building permit growth (BPG) volatility, strongly predicts future asset market volatility. By constructing a novel dataset of U.S. historical local building permits from 1919 to 2019, we demonstrate that BPG volatility forecasts stock and corporate bond market volatility, even after controlling for various factors such as leverage ratios, natural disaster risk, and macroeconomic conditions. We find this predictive power is particularly pronounced in more housing supply elastic regions, indicating that the

[42] Relative to the earlier undevelopable land share of Saiz (2010), the Lutz and Sand (2023) buildable land measure uses satellite data to account for existing construction and public spaces. For instance, Florida's buildable land share would be higher if not for a large fraction of the state being covered by wetlands unsuitable for housing. We describe how we construct this measure at the state level in Appendix E.

housing market's influence on financial markets varies geographically. Our analysis also reveals that during specific events like the Great Recession, BPG volatility in areas with high subprime mortgage exposure exhibited stronger predictability for asset market volatility.

Our research contributes significantly to understanding the relationship between housing and financial markets, introducing BPG volatility as a new monthly factor for forecasting stock and corporate bond returns. By first documenting these facts in new historical data and then interpreting them through a simple model of housing development and information aggregation, we show why building permits are especially informative in markets where supply constraints are weaker and completion lags are shorter. We highlight the risk of overbuilding, showing how uncertainty about future economic conditions captured by building permit volatility often results in excessive development and subsequent housing busts. This insight offers valuable implications for macroprudential policy, suggesting that policymakers should consider measures tailored towards property developers to prevent runaway housing market booms. Our comprehensive longitudinal database of local permits opens new avenues for future research on local housing supply and its impact on financial markets.

References

Abramovitz, Moses, *Evidences of Long Swings in Aggregate Construction Since the Civil War*, Cambridge, MA: National Bureau of Economic Research, 1964.

Altshuler, Alan and Jose A. Gomez-Ibanez, *Regulation for Revenue: The Political Economy of Land Use Exactions*, Washington, DC: Brookings Institution, 1993.

American Society of Planning Officials, "Problems of Zoning and Land-use Regulation," Technical Report, Prepared for the consideration of the National Commission on Urban Problems 1968.

Baerlocher, Diogo, Gustavo S. Cortes, and Vinicios P. Sant'Anna, "Finding Home When Disaster Strikes: Dust Bowl Migration and Housing in Los Angeles," Working Paper 2024.

Bai, John (Jianqiu), Douglas Fairhurst, and Matthew Serfling, "Employment Protection, Investment, and Firm Growth," *Review of Financial Studies*, 2020, *33* (2), 644–688.

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

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, "Extrapolation and Bubbles," *Journal of Financial Economics*, 2018, *129* (2), 203–227.

Baron, Matthew and Wei Xiong, "Credit Expansion and Neglected Crash Risk," *Quarterly Journal of Economics*, 2017, *132* (2), 713–764.

Barr, Jason, "Skyscrapers and the Skyline: Manhattan, 1895–2004," *Real Estate Economics*, 2010, *38* (3), 567–597.

Bartik, Alexander, Arpit Gupta, and Daniel Milo, "The Costs of Housing Regulation: Evidence from Generative Regulatory Measurement," *Available at SSRN 4729511*, 2024.

Bednarek, Peter, Daniel Marcel te Kaat, Chang Ma, and Alessandro Rebucci, "Capital Flows, Real Estate, and Local Cycles: Evidence from German Cities, Banks, and Firms," *Review of Financial Studies*, 2021, *34* (10), 5077–5134.

Berube, Alan, "How the Pandemic Changed-and Didn't Change-Where Americans are Moving," *New York Times*, 2024. Online; accessed 11/27/2024.

Bollerslev, Tim, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 1986, *31* (3), 307–327.

Brocker, Michael and Christopher Hanes, "The 1920s American Real Estate Boom and the Downturn of the Great Depression: Evidence from City Cross-Sections," in "Housing and Mortgage Markets in Historical Perspective," University of Chicago Press, 2014, pp. 161–201.

Bronin, Sara C., "Zoning by a Thousand Cuts," *Pepperdine Law Review*, 2023, *50* (4), 719–784.

Bulan, Laarni, Christopher Mayer, and C. Tsuriel Somerville, "Irreversible Investment, Real Options, and Competition: Evidence from Real Estate Development," *Journal of Urban Economics*, 2009, *65* (3), 237–251.

Calomiris, Charles W. and Matthew S. Jaremski, "Florida (Un)chained," *Journal of Financial Intermediation*, 2023, *55* (101043).

Campbell, Sean D., Morris A. Davis, Joshua Gallin, and Robert F. Martin, "What Moves Housing Markets: A Variance Decomposition of the Rent-Price Ratio," *Journal of Urban Economics*, 2009, *66* (2).

Census Bureau, "Construction Reports Building Permits: New Dwelling Units Authorized by Local Building Permits, May and June 1959," Technical Report 1959. Online via HathiTrust; accessed December 14, 2023.

_ , "Residential Building Permits Survey Documentation, Master Compiled Dataset," Technical Report, https://www2.census.gov/econ/bps/Master%20Data%20Set 2022. Online; accessed August 5, 2023.

Chinco, Alex and Christopher Mayer, "Misinformed Speculators and Mispricing in the Housing Market," *Review of Financial Studies*, 2015, *29* (2), 486–522.

Colean, Miles L. and Robinson Newcomb, *Stabilizing Construction: The Record and the Potential*, New York, NY: McGraw-Hill, 1952.

Cortes, Gustavo S. and Marc D. Weidenmier, "Stock Volatility and the Great Depression," *Review of Financial Studies*, 2019, *32* (9), 3544–3570.

_ **, Angela Vossmeyer, and Marc D. Weidenmier**, "Stock Volatility and the War Puzzle: The Military Demand Channel," NBER Working Paper 29837, National Bureau of Economic Research 2024.

Cunningham, Christopher R., "House Price Uncertainty, Timing of Development, and Vacant Land Prices: Evidence for Real Options in Seattle," *Journal of Urban Economics*, 2006, *59* (1), 1–31.

D'Agostino, Ralph B., Albert Belanger, and Ralph B. D'Agostino Jr., "A Suggestion for Using Powerful and Informative tests of Normality," *American Statistician*, 1990, *44* (4), 316–321.

D'Amico, Leonardo, Edward L. Glaeser, Joseph Gyourko, William Kerr, and Giacomo A.M. Ponzetto, "Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation," 2024. SSRN Working Paper 4679195.

Davis, Morris A. and Jonathan Heathcote, "The Price and Quantity of Residential Land in the United States," *Journal of Monetary Economics*, 2007, *54* (8), 2595–2620.

Fama, Eugene F. and Kenneth R. French, "Industry costs of equity," *Journal of Financial Economics*, 1997, *43* (2), 153–193.

_ **and** _ , "House Prices and Rents," *Chicago Booth Research Paper*, 2023, *No. 20-48.*

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh, "Affordable Housing and City Welfare," *Review of Economic Studies*, 2023, *90* (1), 293–330.

Fischel, William A., "An Economic History of Zoning and a Cure for its Exclusionary Effects," *Urban Studies*, 2004, *41* (2), 317–340.

_ , *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*, Cambridge, MA: Harvard University Press, 2005.

Fishback, Price and Trevor Kollmann, "New Multicity Estimates of the Changes in Home Values, 1920-1940," in "Housing and Mortgage Markets in Historical Perspective," University of Chicago Press, 2014, pp. 203–244.

Fishback, Price V., Jonathan Rose, and Kenneth Snowden, *Well Worth Saving: How the New Deal Safeguarded Home Ownership*, University of Chicago Press, 2013.

Flynn, Sean, Andra C. Ghent, and Vasudha Nair, "Determinants and Consequences of Return to Office Policies," *Available at SSRN 4757876*, 2024.

Fout, Hamilton, Mark R.A. Palim, Yi Song, and Vincent W. Yao, "Foreclosure Timelines and House Prices," Technical Report, Fannie Mae Working Paper 2017.

Francke, Marc and Matthijs Korevaar, "Baby Booms and Asset Booms: Demographic Change and the Housing Market," *Journal of Finance*, 2025.

Gabriel, Stuart and Edward Kung, "Development Approval Timelines, Approval Uncertainty and New Housing Supply: Evidence from Los Angeles," *Available at SSRN 4872147*, 2024.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the U.S. Declined?," *Journal of Urban Economics*, 2017, *102*, 76–90.

Gao, Mingze, Henry Leung, and Buhui Qiu, "Organization Capital and Executive Performance Incentives," *Journal of Banking & Finance*, 2021, *123*, 106017.

Geltner, David M., Norman G. Miller, Jim Clayton, and Piet Eichholtz, "Chapter 27: Real Options and Land Value," in "Commercial Real Estate: Analysis and Investments, 3rd edition," Mason, OH: OnCourse Learning, 2014, pp. 706–731.

Ghent, Andra C. and Michael T. Owyang, "Is Housing the Business Cycle? Evidence from US Cities," *Journal of Urban Economics*, 2010, *67* (3), 336–351.

_ , Paige Rowberry, and Matthew I. Spiegel, "The Impact of COVID Restrictions on Business Dynamics," *Available at SSRN 4834348*, 2024.

Giroud, Xavier and Holger M. Mueller, "Firms' Internal Networks and Local Economic Shocks," *American Economic Review*, 2019, *109* (10), 3617–3649.

Gjerstad, Steven D. and Vernon L. Smith, *Rethinking Housing Bubbles: The Role of Household and Bank Balance Sheets in Modeling Economic Cycles*, Cambridge, UK: Cambridge University Press, 2014.

Glaeser, Edward L., "A Nation of Gamblers: Real Estate Speculation and American History," *American Economic Review*, 2013, *103* (3), 1–42.

_ , Joseph Gyourko, and Raven E. Saks, "Why Have Housing Prices Gone Up?," *American Economic Review*, 2005, *95* (2), 329–333.

Glancy, David and Robert Kurtzman, "How Do Capital Requirements Affect Loan Rates? Evidence from High Volatility Commercial Real Estate," *Review of Corporate Finance Studies*, 2022, *11* (1), 88–127.

_ , _ , and Lara Loewenstein, "On Commercial Construction Activitys Long and Variable Lags," *Federal Reserve Board Finance and Economics Discussion Series*, 2024, *2024* (016).

_ , _ , _ , and Joseph Nichols, "Recourse as Shadow Equity: Evidence from Commercial Real Estate Loans," *Real Estate Economics*, 2023, *51* (5), 1108–1136.

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 1993, *48* (5), 1779–1801.

Goetzmann, William N. and Frank Newman, "Securitization in the 1920s," NBER Working Paper 15650, National Bureau of Economic Research 2010.

Graham, John R., Mark T. Leary, and Michael R. Roberts, "A Century of Capital Structure: The Leveraging of Corporate America," *Journal of Financial Economics*, 2015, *118* (3), 658–683.

Green, Richard K., "Follow the Leader: How Changes in Residential and Non-residential Investment Predict Changes in GDP," *Real Estate Economics*, 1997, *25* (2), 253–270.

Greenwood, Robin and Andrei Shleifer, "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 2014, *27* (3), 714–746.

__ , **Samuel G. Hanson, Andrei Shleifer, and Jakob Ahm Sørensen**, "Predictable Financial Crises," *Journal of Finance*, 2022, *77* (2), 863–921.

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Gurun, Umit G., Jiabin Wu, Steven Chong Xiao, and Serena Wenjing Xiao, "Do Wall Street Landlords Undermine Renters' Welfare?," *Review of Financial Studies*, 2023, *36* (1), 70–121.

Gyourko, Joseph, Albert Saiz, and Anita Summers, "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 2008, *45* (3), 693–729.

__ , **Jonathan S. Hartley, and Jacob Krimmel**, "The Local Residential Land Use Regulatory Environment across U.S. Housing Markets: Evidence from a New Wharton Index," *Journal of Urban Economics*, 2021, *124*.

Hansen, Stephen, Nicholas Bloom, Steven J. Davis, Raffaella Sadun, and Bledi Taska, "Remote Work across Jobs, Companies, and Space," Working Paper 31007, National Bureau of Economic Research 2023.

Heston, Steven L. and Ronnie Sadka, "Seasonality in the Cross-Section of Stock Returns," *Journal of Financial Economics*, 2008, *87* (2), 418–445.

Horton, Emily, Cameron LaPoint, Byron Lutz, Nathan Seegert, and Jared Walczak, "Property Tax Policy and Housing Affordability," *National Tax Journal*, 2024, *77* (4), 861–901.

Howard, Troup, Mengqi Wang, and Dayin Zhang, "How Do Labor Shortages Affect Residential Construction and Housing Affordability?," *Available at SSRN 4729511*, 2024.

Isard, Walter, "A Neglected Cycle: The Transport-Building Cycle," *Review of Economics and Statistics*, 1942, *24* (4), 149–158.

Janardanan, Rajkumar, Xiao Qiao, and K. Geert Rouwenhorst, "The Flipside of Financial Innovation: Why Contracts Fail," Technical Report, Available at SSRN 4730746 2024. Working Paper.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan Taylor, "The Rate of Return on Everything, 1870-2015," *Quarterly Journal of Economics*, 2019, *134* (3), 1225–1298.

__ , **Moritz Schularick, and Alan M. Taylor**, "When Credit Bites Back," *Journal of Money, Credit and Banking*, 2013, *45* (s2), 3–28.

Kliesen, Kevin L., "The 2001 Recession: How Was It Different and What Developments May Have Caused It?," *Federal Reserve Bank of St. Louis Review*, 2003, *85* (5), 23–38.

Knoll, Katharina, Moritz Schularick, and Thomas Steger, "No Price Like Home: Global House Prices, 1870–2012," *American Economic Review*, 2017, *107* (2), 331–353.

Knowlton, Christopher, *Bubble in the Sun: The Florida Boom of the 1920s and How it Brought on the Great Depression*, New York: Simon & Schuster, 2020.

Korevaar, Matthijs, "Reaching for Yield and the Housing Market: Evidence from 18th-Century Amsterdam," *Journal of Financial Economics*, 2023, *148* (3), 273–296.

Kulka, Amrita, Aradhya Sood, and Nicholas Chiumenti, "Under the (neighbor)Hood: Understanding Interactions Among Zoning Regulations," *Review of Economics and Statistics, forthcoming*, 2026.

Kuvshinov, Dmitry, "Asset-Class-Specific Discount Rates," Technical Report, Available at SSRN 3289584 2025. Working Paper.

LaPoint, Cameron, "You Only Lend Twice: Corporate Borrowing and Land Values in Real Estate Cycles," Technical Report, Available at SSRN 3633606 2021. Working Paper.

Leamer, Edward E., "Housing Is the Business Cycle," in "Proceedings of the Jackson Hole Economic Policy Symposium" Federal Reserve Bank of Kansas City 2007, pp. 149–233.

_ , "Housing Really Is the Business Cycle: What Survives the Lessons of 2008–09?," *Journal of Money, Credit and Banking*, 2015, *47* (S1), 43–50.

Lettau, Martin and Sydney Ludvigson, "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 2001, *56* (3), 815–849.

Ling, David C., Chongyu Wang, and Tingyu Zhou, "A First Look at the Impact of COVID-19 on Commercial Real Estate Prices: Asset-Level Evidence," *Review of Asset Pricing Studies*, 2020, *10* (4), 669–704.

Liu, Chang and Wei Xiong, "China's Real Estate Market," NBER Working Paper 25297, National Bureau of Economic Research 2018.

Long, Clarence D., "Seventy Years of Building Cycles in Manhattan," *Review of Economics and Statistics*, 1936, *18* (4), 183–193.

_ , "Long Cycles in the Building Industry," *Quarterly Journal of Economics*, 1939, *53* (3), 371–403.

_ , "The Building Industry: Maker and Breaker of Booms and Depressions," *Dun's Review*, 1939, *December*, 17–50.

Ludvigson, Sydney C. and Serena Ng, "Macro Factors in Bond Risk Premia," *Review of Financial Studies*, 2009, *22* (12), 5027–5067.

Lutz, Chandler and Ben Sand, "Highly Disaggregated Land Unavailability," *Available at SSRN 3478900*, 2023.

Lyons, Ronan C., Allison Shertzer, Rowena Gray, and David N. Agorastos, "The Price of Housing in the United States, 1890–2006," *Quarterly Journal of Economics*, 2026, *141* (1), 559–603.

Manela, Asaf and Alan Moreira, "News Implied Volatility and Disaster Concerns," *Journal of Financial Economics*, 2017, *123* (1), 137–62.

Mankiw, N. Gregory and David N. Weil, "The Baby Boom, the Baby Bust, and the Housing Market," *Regional science and Urban Economics*, 1989, *19* (2), 235–258.

Mayer, Christopher J. and Karen Pence, "Subprime Mortgages: What, Where, and to Whom?," NBER Working Paper 14083, National Bureau of Economic Research 2008.

Mian, Atif, Amir Sufi, and Emil Verner, "Household debt and business cycles worldwide," *Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

Moore, Geoffrey H., "Business Cycle Indicators, Vol. 1," in "Statistical Indicators of Cyclical Revivals and Recessions," NBER, 1961, pp. 184–260.

Müller, Karsten and Emil Verner, "Credit Allocation and Macroeconomic Fluctuations," *Review of Economic Studies*, 2024, *91* (6), 3645–3676.

Natenberg, Sheldon, *Option Volatility and Pricing: Advanced Trading Strategies and Techniques*, 2nd ed., McGraw Hill, 2014.

Nathanson, Charles G. and Eric Zwick, "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market," *Journal of Finance*, 2018, *73* (6), 2587–2633.

Nelson, Daniel B., "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 1991, *59* (2), 347–370.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, *55* (3), 703–708.

Nicholas, Tom and Anna Scherbina, "Real Estate Prices During the Roaring Twenties and the Great Depression," *Real Estate Economics*, 2013, *41* (2), 278–309.

Ogden, Joseph P., "The Calendar Structure of Risk and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 2003, *70* (1), 29–67.

Oster, Sharon M. and John M. Quigley, "Regulatory Barriers to the Diffusion of Innovation: Some Evidence from Building Codes," *The Bell Journal of Economics*, 1977, *8* (2).

Ozcicek, Omer and W. Douglas McMillin, "Lag Length Selection in Vector Autoregressive Models: Symmetric and Asymmetric Lags," *Applied Economics*, 1999, *31* (4), 517–524.

Pálvölgyi, Dömötör and Gyuri Venter, "Multiple Equilibria in Noisy Rational Expectations Economies," 2015. SSRN Working Paper 2524105.

Pástor, Lubos and Pietro Veronesi, "Was There a Nasdaq Bubble in the Late 1990s?," *Journal of Financial Economics*, 2006, *81* (1), 61–100.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 2007, *83* (3), 531–569.

Riggleman, John R., "Building Cycles in the United States, 18751932," *Journal of the American Statistical Association*, 1932, *28* (182), 174–183.

Royston, Patrick, "Comment on sg3.4 and an Improved D'Agostino Test," *Stata Technical Bulletin*, 1992, *1* (3), 23–25.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Schularick, Moritz and Alan M. Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *American Economic Review*, 2012, *102* (2), 1029–1061.

Schwert, G. William, "Why Does Stock Market Volatility Change over Time?," *Journal of Finance*, 1989, *44* (5), 1115–1153.

Shen, Zejiang, Ruochen Zhang, Melissa Dell, Benjamin Charles Germain Lee, Jacob Carlson, and Weining Li, "LayoutParser: A Unified Toolkit for Deep Learning Based Document Image Analysis," Technical Report, Available at arXiv:2103.15348 2021.

Shertzer, Allison, Tate Twinam, and Randall P. Walsh, "Zoning and Segregation in Urban Economic History," *Regional Science and Urban Economics*, 2022, *94* (103652).

Shiller, Robert J., "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 1981, *71* (3), 421–36.

__ , *Irrational Exuberance, 3rd edition*, Princeton, NJ: Princeton University Press, 2015.

Smets, Frank, "Commentary: Housing is the Business Cycle," in "Proceedings of the Jackson Hole Economic Policy Symposium" Federal Reserve Bank of Kansas City 2007, pp. 235–243.

Soltas, Evan J., "The Price of Inclusion: Evidence from Housing Developer Behavior," *Review of Economics and Statistics*, 2024, *106* (6), 1588–1606.

Somerville, C. Tsuriel, "Permits, Starts, and Completions: Structural Relationships Versus Real Options," *Real Estate Economics*, 2002, *29* (1), 161–190.

Song, Jaehee, "The Effects of Residential Zoning in U.S. Housing Markets," *Journal of Urban Economics*, 2025, *149* (103784).

Stock, James H. and Mark W. Watson, "The Evolution of National and Regional Factors in U.S. Housing Construction ," in "Volatility and Time Series Econometrics: Essays in Honor of Robert Engle," Oxford: Oxford University Press, 2010, pp. 35–61.

Strauss, Jack, "Does Housing Drive State-Level Job Growth? Building Permits and Consumer Expectations Forecast a States Economic Activity," *Journal of Urban Economics*, 2013, *73*, 77–93.

Titman, Sheridan, "Urban Land Prices under Uncertainty," *American Economic Review*, 1985, *75* (3), 505–514.

Turner, Gregg M., *The Florida Land Boom of the 1920s*, Jefferson, NC: McFarland & Company, 2015.

U.S. Bureau of Labor Statistics, "Construction During Five Decades: Historical Statistics, 1907-1952," Technical Report 1954. Online; accessed July 16, 2024.

U.S. Census Bureau, "X-13ARIMA-SEATS Reference Manual, Version 1.1," Statistical Research Division, U.S. Census Bureau, U.S. Department of Commerce 2022.

Wall Street Journal, "The States Where Home Prices are Secret," https://www.wsj.com/articles/the-states-where-home-prices-are-secret-11560956939 2019. Online; accessed 06/19/2019.

_ , "Once Americas Hottest Housing Market, Austin Is Running in Reverse," https://www.wsj.com/economy/housing/once-americas-hottest-housing-market-austin-is-running-in-reverse-94226027 2024. Online; accessed 08/04/2024.

White, Eugene N., "Lessons from the Great American Real Estate Boom and Bust of the 1920s," NBER Working Paper 15573, National Bureau of Economic Research 2009.

_ , "Lessons from the Great American Real Estate Boom and Bust of the 1920s," in "Housing and Mortgage Markets in Historical Perspective," University of Chicago Press, 2014, pp. 115–158.

Yuan, Kathy, "Asymmetric Price Movements and Borrowing Constraints: A Rational Expectations Equilibrium Model of Crises, Contagion, and Confusion," *Journal of Finance*, 2005, *60*, 379–411.

Online Appendix to

Housing *Is* the Financial Cycle: Evidence from 100 Years of Local Building Permits

Gustavo S. Cortes (University of Florida) & Cameron LaPoint (Yale SOM)

Contents

A Theoretical Framework: Derivations and Extensions

In this appendix, we provide proofs for the key propositions in the theoretical framework developed in Section 6, which we use to interpret the empirical findings in the main text. We then offer two extensions: one where we allow signals based on local building permit growth, q_s, to be correlated across regions, and a second where we endogenize construction costs while incorporating *ad valorem* permit fees and regulatory compliance costs.

A.1 Model Proofs

Proof of Proposition 1

Step 1 (linearity). Because $(d, \{q_s\}, u)$ are jointly Gaussian and all investors have CARA utility, posterior means and variances are linear/constant functions of the sufficient statistics they observe. With any conjectured measurable price p, informed investors s observe (q_s, p) and the uninformed observed p. Under joint normality, optimal demands are linear in the posterior mean of d and in price. Hence, aggregate demand is affine in $(\{q_s\}, u, p)$. With market clearing holding in each locality, the unique solution for p must therefore be affine in the primitive shocks. There then exist coefficients $\alpha, \{\beta_s\}, \kappa$ such that:

$$p = \alpha + \sum_{k \in \mathcal{S}} \beta_k \cdot q_k + \kappa \cdot u \tag{A.1}$$

Re-writing this identity relative to a particular locality s yields:

$$p = \underbrace{\left(\alpha + \sum_{k \neq s} \beta_k \cdot q_k\right)}_{\phi_0(s)} + \underbrace{\beta_s}_{\phi_q(s)} \cdot \Big(q_s + \underbrace{\kappa/\beta_s}_{\phi_u(s)} \cdot u\Big) \tag{A.2}$$

which is the stated form in the proposition. Here $\phi_0(s)$ collects the effects of other localities' signals.

Step 2 (sign of the loading on q_s). Let $x_I(s)$ and $x_U(s)$ be the respective informed and uninformed demands in locality s. With CARA-normal preferences,

$$x_{j(s)} = \frac{\mathbb{E}[d|\omega_{j(s)}] - (1+r) \cdot p}{\gamma \cdot Var[d|\omega_{j(s)}]}, \quad j \in \{I, U\}$$

Holding p fixed, $\partial \mathbb{E}[d|q_s]/\partial q_s > 0$ and $\partial \mathbb{E}[d|p]/\partial q_s \geq 0$ (the price carries the signal), so aggregate demand $D(p, \{q_k\}) = \sum_s \Big(\lambda_s \cdot x_I(s) + (1 - \lambda_s) \cdot x_U(s)\Big)$ satisfies $\partial D/\partial q_s$. The aggregate market clearing condition then implies:

$$F(p, \{q_k\}, u) := D(p, \{q_k\}) - (m + u) = 0$$

By the implicit function theorem,

$$\frac{\partial p}{\partial q_s} = -\frac{\partial F/\partial q_s}{\partial F/\partial p} = -\frac{\partial D/\partial q_s}{\partial D/\partial p}$$

Each individual demand is strictly decreasing in price, hence $\partial D/\partial p < 0$. Therefore $\partial p/\partial q_s > 0$. Noting that $\phi_q(s) = \partial p/\partial q_s$ based on equation (A.2) shows that $\phi_q(s) > 0$.

Equivalently, we can write the slopes of the aggregate demand curve explicitly:

$$\frac{\partial D}{\partial q_s} = \frac{1}{\gamma} \cdot \Big(\lambda_s \cdot \tau_q(s) + (1-\lambda_s) \cdot \widetilde{\tau}_p(s)\Big) > 0 \tag{A.3}$$

$$\frac{\partial D}{\partial p} = -\frac{1+r}{\gamma} \sum_{k \in \mathcal{S}} \left(\lambda_k(\tau_d + \tau_q(k)) + (1-\lambda_k)(\tau_d + \widetilde{\tau}_p(k))\right) < 0 \tag{A.4}$$

where $\tau_d, \tau_q(s), \widetilde{\tau}_p(s)$ are the precisions for the dividend, local building permit growth, and the transformed price signal (6.12), respectively. This shows that the positive effect of an increase in building permit growth on risky asset prices does not depend on risk aversion γ.

Step 3 (sign of the loading on supply noise u). Again applying the implicit function theorem:

$$\frac{\partial p}{\partial u} = -\frac{\partial F/\partial u}{\partial F/\partial p} = \frac{1}{\partial D/\partial p} < 0 \tag{A.5}$$

where the sign follows from (A.4). In the local representation of $p = \phi_0(s) + \phi_q(s) \cdot (q_s + \phi_u(s) \cdot u)$ we have $\partial p/\partial u = \phi_q(s) \cdot \phi_u(s)$. Because $\phi_q > 0$, this implies $\phi_u(s) < 0$. This establishes that the stated affine form and sign restrictions hold independently of risk aversion γ.

Step 4 (closed-form pricing coefficients). For shorthand, define the aggregate informed "precision mass" T_Q and associated moments as:

$$T_Q \equiv \sum_s w_s \lambda_s \cdot \tau_q(s) \qquad \Lambda \equiv \sum_s = w_s \cdot \lambda_s \qquad H_Q \equiv \sum_s w_s^2 \lambda_s^2 \cdot \tau_q(s)$$

where w_s refers to the population or market weight of investors in located in s. H_Q can be interpreted as the dispersion-adjusted idiosyncratic noise from locality signals inside price, while Λ is the total mass of informed investors across areas.

Let τ_p denote the precision that the price contributes about d. In equilibrium, this object is equivalent to:

$$\tau_p = \frac{T_Q^2}{H_Q + \gamma^2 \cdot \sigma_u^2} \tag{A.6}$$

which follows from the linear-Gaussian filtering implied by the definition of the noisy rational expectation equilibrium in equations (6.8)–(6.12). Define market depth D_0 as:

$$D_0 \equiv \frac{R \cdot \Big(\tau_d + T_Q + (1-\Lambda) \cdot \tau_p\Big)}{1 + \frac{(1-\Lambda) \cdot \tau_p}{\tau_q}} \tag{A.7}$$

Then we can write the loadings on each locality's signal q_s as:

$$\beta_s = \frac{w_s \lambda_s \tau_q(s)}{D_0} = \frac{w_s \lambda_s \tau_q(s) \Big(T_Q + (1-\Lambda)\tau_p\Big)}{R \cdot T_Q \Big(\tau_d + T_Q + (1-\Lambda)\tau_p\Big)} \tag{A.8}$$

We can similarly write the loading on noise in the aggregate pricing function (A.1) as:

$$\kappa = -\frac{\gamma}{D_0} = -\frac{\gamma\Big(T_Q + (1-\Lambda)\tau_p\Big)}{R \cdot T_Q\Big(\tau_d + T_Q + (1-\Lambda)\tau_p\Big)} \tag{A.9}$$

Hence, we have solved for the coefficients for the conjectured local pricing function:

$$\phi_q(s) = \beta_s \qquad \phi_u(s) = \frac{\kappa}{\beta_s} = -\frac{\gamma}{w_s \lambda_s \tau_q(s)} \tag{A.10}$$

Note that $\phi_u(s)$ only depends on the local information mass. Better informed trading in s makes the price-revealed signal $\widetilde{p} = q_s + \phi_u(s) \cdot u$ less contaminated by supply noise.

To close the system, we can solve for the precision of $\widetilde{p}$ as:

$$\widetilde{\tau}_p(s) = \frac{w_s^2 \lambda_s^2 \tau_q(s)^2}{w_s^2 \lambda_s^2 \tau_q(s) + \gamma^2 \sigma_u^2} \tag{A.11}$$

Corollary 1.1 Comparative Statics

We prove that the sign of the slope estimates in regressions of price σ_p or return volatility σ_r on BPG volatility $\sigma_q(s)$ is generally ambiguous and isolate parameter spaces where the sign is either predicted to be positive or negative. First, note that the sign will be preserved if we consider variances instead of volatilities. This allows us to write the derivatives in closed form in terms of signal precisions, using the notation in the preceding proofs. As shorthand, denote $m_s = w_s \cdot \lambda_s$ the number of informed investors in locality s.

Note also that we can rewrite the coefficients defining the linear price equilibrium as:

$$\beta_s = c \cdot m_s \tau_q(s) \qquad \kappa = -\gamma \cdot c$$

$$\text{with } c = \frac{T_Q + (1-\Lambda)\tau_p}{(1+r) \cdot T_Q(\tau_d + T_Q + (1-\Lambda)\tau_p)}$$

and the price precision for dividends τ_p defined as in (A.6). Using only these objects, we can then simplify the price and return variance to:

$$\sigma_p^2 = c^2 \cdot (H_Q + \gamma^2 \sigma_u^2) + \sigma_d^2 (c \cdot T_Q)^2 \tag{A.12}$$

$$\sigma_r^2 = \sigma_p^2 + \sigma_d^2 \cdot (1 + 2c \cdot T_Q) \tag{A.13}$$

These two identities follow from inserting the previous definitions for β_s and κ into the aggregate price equation in (A.1).

The first comparative static then yields:

$$\frac{\partial \sigma_p^2}{\partial \tau_q(s)} = \sigma_p^2 \cdot \left\{ \frac{2(m_s + (1-\Lambda)\tau_p')}{(T_Q + (1-\Lambda)\tau_p)(\tau_d + T_Q + (1-\Lambda)\tau_p)} - \frac{\tau_p'}{\tau_p(1+\sigma_d^2\tau_p)} \right\} \tag{A.14}$$

$$\text{with } \tau_p' = \frac{\partial \tau_p}{\partial \tau_q(s)} = \frac{2m_s T_Q}{H_Q + \gamma^2\sigma_u^2} - \frac{m_s^2 T_Q^2}{(H_Q + \gamma^2\sigma_u^2)^2} \tag{A.15}$$

The sign of (A.14) is generally ambiguous. Both terms are positive functions of τ_p'. Depending on how large supply noise $\gamma^2\sigma_u^2$ is relative to the information mass H_Q, τ_p' changes, which can flip the sign.

To illustrate this, we examine several edge cases and discuss the relevance of each vis-á-vis our empirical results.

1. **Suppose we take the limit of an extremely precise local signal** $\sigma_q(s) \to 0$, or $\tau_q(s) \to \infty$. In this scenario we need to further distinguish between two cases. If the informational mass dominates price precision, $H_Q \gg \gamma^2\sigma_u^2$ then $\tau_p \sim T_Q^2/H_Q$ grows linearly with $\tau_q(s)$, the second (negative) term in (A.14) dominates at the margin, leading to:

$$\frac{\partial \sigma_p^2}{\partial \tau_q(s)} > 0 \implies \frac{\partial \sigma_p^2}{\partial \sigma_q^2(s)} < 0 \text{ near } \sigma_q^2(s) = 0$$

   If supply noise dominates price precision, or $\gamma^2\sigma_u^2 \gg H_Q$, then $\tau_p \approx T_Q^2/(\gamma^2\sigma_u^2)$ grows quadratically with $\tau_q(s)$. In this regime the first positive term in (A.14) dominates, and:

$$\frac{\partial \sigma_p^2}{\partial \tau_q(s)} < 0 \implies \frac{\partial \sigma_p^2}{\partial \sigma_q^2(s)} > 0 \text{ near } \sigma_q^2(s) = 0$$

   That is, adding noise to an ultra-precise q_s raises price variance because it amplifies the impact of supply noise in equilibrium. The second of these two cases accords with what we find in the data, suggesting that supply noise is a strong force.

2. Alternatively, **suppose we take the limit of a very noisy local signal** $\sigma_q(s) \to \infty$, or $\tau_q(s) \to 0$. In this case $\beta_s = cm_s\tau_q(s) \to 0$, while $\kappa = -\gamma \cdot c$ and the other β_k barely move. In this limit $\sigma_p^2/\sigma_q^2(s) > 0$. In other words, making an already poor signal even noisier raises the price's exposure to noise (the small decline in the direct $\beta_s^2\sigma_q^2(s)$ term is dominated by the equilibrium reweighting towards noise.

A.2 Sources of Cross-State Heterogeneity in the Comparative Statics

Given that states with lower $\sigma_q^2(s)$ tend to yield better return and volatility prediction, the parameter space described in the first case of **Corollary 1.1** is the empirically relevant one. We propose two extensions of the model to bring it closer to the data and derive the new comparative statics. We formalize the conclusions here as two additional propositions.

Proposition 2 (Geography-Linked Supply Noise). *Keeping all features of the baseline model intact, redefine the risky asset supply shock as* $u = \sum_j \theta_j \cdot u_j$ *with* $u_j \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{u,j}^2)$ *for locations* $j \neq s$. *Then the*

equilibrium is still linear and defined by:

$$p = \alpha + \sum_{j \in \mathcal{S}} \beta_j \cdot q_j + \sum_{j \in \mathcal{S}} \kappa_j \cdot u_j \qquad \textit{with} \qquad \kappa_j = -\gamma \cdot c\theta_j$$

Further, there exist thresholds $(\bar{\theta}_p(s), \bar{\theta}_r(s))$ *such that:*

$$\gamma^2 \sigma^2_{u,eff} \geq \bar{\theta}_p(s) \cdot H_Q \implies \frac{\partial \sigma^2_p}{\partial \sigma^2_q(s)} > 0$$

$$\gamma^2 \sigma^2_{u,eff} \geq \bar{\theta}_r(s) \cdot H_Q \implies \frac{\partial \sigma^2_r}{\partial \sigma^2_q(s)} > 0$$

where $\sigma^2_{u,eff} := \sum_j \theta_j^2 \sigma^2_{u,j}$ *is the total price-noise variance.*

Proof. First, note that all statements and expressions in the previous version of the model hold with $\sigma^2_{u,eff}$ instead of σ^2_u. Therefore, going forward, we replace σ^2_u with $\sigma^2_{u,eff}$ in all expressions in the preceding version of the model.

Fix $\{m_k, \tau_q(k)\}_{k \neq s}, \tau_d > 0, \gamma > 0, \sigma^2_u > 0$. Evaluate the derivatives at a point where $\sigma^2_q(s)$ is small enough such that $\gamma^2 \sigma^2_u \geq \theta \cdot H_Q$ for some $\theta > 1$. When this condition holds, supply noise locally dominates information mass in $\mathcal{H} = H_Q + \gamma^2 \sigma^2_u$. Then price thresholds satisfying the statement can be derived as:

$$\bar{\theta}_p = 1 + \frac{m_s \cdot T_Q}{H_Q} \qquad \bar{\theta}_r = \bar{\theta}_p \cdot \left(1 + \frac{\sigma^2_d}{\tau_d + T_Q}\right)$$

With $\theta \geq \bar{\theta}_p$, equation (A.15) reduces to $\tau'_p \approx 2 m_s T_Q / \mathcal{H}$, and (A.14) has a negative sign. To see this, note that the $c^2 m_s^2$ term is second order in $1/\theta$. Thus, $\partial \sigma^2_p / \partial \tau_q(s) < 0$ and $\partial \sigma^2_p / \partial \sigma^2_q(s) > 0$. For returns, the same logic holds, but the larger threshold $\bar{\theta}_r$ ensures the negative term from $\partial \sigma^2_p / \partial \tau_q(s)$ dominates.

The thresholds depend only on readily computed objects at the evaluation point (T_Q, H_Q). Hence, when we observe the "positive near small σ_q" pattern in a group of states, the model attributes this to regions where asset supply noise dominates. These are states with a smaller mass of informed investors m_s. In this version of the model, cross-state variation in the sign is therefore driven by information-side differences and not by state-specific supply noise. □

Proposition 3 (Endogenous Informative Mass). *Keeping all features of the baseline model intact, suppose now the informed investor mass at locality s is increasing in the precision of building permits* $\tau_q(s)$*, holding fixed* m_j *for all* $j \neq s$*. Hence,* $m'_s(\tau_q) \geq 0$*. Then* $\partial \sigma^2_p / \partial \tau_q(s) > 0$ *and* $\partial \sigma^2_r / \tau_q(s) > 0$ *for sufficiently small* $\sigma^2_q(s)$*.*

Proof. Define the elasticity of the number informed investors with respect to signal precision: $\psi_s := \tau_q(s) \cdot m'_s(\tau_q) / m_s \geq 0$. For shorthand, denote $x_s := m_s \cdot T_Q / \mathcal{H}$. Then $\tau'_p = \partial \tau_p / \partial \tau_q(s)$ reduces to:

$$\tau'_p = \frac{T_Q}{\mathcal{H}} \cdot m_s \Big(2(1 + \psi_s) - x_s(1 + 2\psi_s) \Big) \tag{A.16}$$

Since τ_p' defines the size of the positive-slope region in equation (A.14), based on (A.16), as ψ_s increases, the parameter space where $\partial\sigma_p^2/\partial\tau_q(s)$ is positive shrinks. To see this, note that:

$$\tau_p' > 0 \iff 2(1+\psi_s) > x_s(1+2\psi_s) \iff x_s < \frac{2(1+\psi_s)}{1+2\psi_s} \in (1,2]$$

Note in the baseline case where $\psi_s = 0$, the threshold is $x_s < 2$. With endogenize information, the threshold falls towards $x_s < 1$ as $\psi_s \to \infty$. Making m_s sensitive to precision makes it harder to keep τ_p' positive, because H_Q grows faster (more concentration of informed investors) as $\tau_q(s)$ rises. We can derive a similar conservative lower bound θ_s as in **Proposition 2**, which reprises the supply-noise vs. information-mass tradeoff:

$$\theta_s := \frac{\gamma^2\sigma_u^2}{H_Q} \implies x_s = \frac{m_s T_Q}{H_Q \cdot (1+\theta_s)}$$

$$x_s < \frac{2(1+\psi_s)}{1+2\psi_s} \iff \theta_s > \frac{m_s T_Q}{H_Q} \cdot \frac{1+2\psi_s}{2(1+\psi_s)} - 1$$

An endogenous information mass raises the required supply-noise dominance $H_Q \gg \gamma^2\sigma_u^2$ required to get the "positive for small $\sigma_q^2(s)$" comparative static. □

A.3 Allowing Signals to be Correlated across Geography

We retain the same setup as in the baseline version of the GS framework in Section 6, except now the BPG signal follows:

$$q = \mathbf{1}d + \varepsilon, \quad \text{with } \varepsilon \sim \mathcal{N}(0, \Sigma_\varepsilon) \tag{A.17}$$

with a full $\mathcal{S} \times \mathcal{S}$ covariance matrix Σ_ε which need not be diagonal. Let the informed mass in state s be $m_s := w_s \cdot \lambda_s$; $M := diag(m_1, \ldots, m_{\mathcal{S}})$; and $\Lambda := \sum_s m_s$. Denote prior precision $\tau_d = 1/\sigma_d^2$. As before, we conjecture a linear noisy rational expectation equilibrium price:

$$p = \alpha + b^\top q + \kappa \cdot u \tag{A.18}$$

Define matrix analogs to the scalar precision masses:

$$v := M\Sigma_\varepsilon^{-1}\mathbf{1} \qquad T := \mathbf{1}^\top v \mathbf{1}^\top M\Sigma_\varepsilon^{-1}\mathbf{1} \qquad H := v^\top \Sigma_\varepsilon v$$

where v is an $\mathcal{S} \times 1$ vector. Now T is the aggregate precision mass that can reach the price when permit signals are correlated across areas. H is the dispersion-adjusted idiosyncratic noise injected through signals. H collapses to $H_Q = \sum_s w_s^2\lambda_s^2\tau_q(s)$ when Σ_ε is diagonal.

We can now derive closed-form coefficients corresponding to the standard CARA-normal linear pricing equilibrium, but now with this richer information structure. First, the price precision (i.e., what prices tell us about dividends) takes the form:

$$\tau_p = \frac{T^2}{H + \gamma^2\sigma_u^2} \tag{A.19}$$

The slope vector on the state signals and the supply-noise loading is a constant:

$$c = \frac{T + (1-\Lambda)\tau_p}{(1+r)T\Big(\tau_d + T + (1-\Lambda)\tau_p\Big)} > 0$$

Then the vector of coefficients defining the aggregate equilibrium in (A.18) is $b = c \cdot v$, and $\kappa = -\gamma \cdot c$. The coefficients in the state-level pricing functions are then:

$$\phi_q(s) = b_s \qquad\qquad \phi_u(s) = \kappa / b_s = -\gamma / v_s = -\frac{\gamma}{w_s \lambda_s [\Sigma_\varepsilon \mathbf{1}]_s} < 0$$

We can also derive the incremental precision from the price (i.e., the precision of the transformed signal $\widetilde{p}$) beyond observing q_s:

$$\widetilde{\tau}_p(s) = \frac{(b^\top \mathbf{1})^2}{Var(b^\top \varepsilon + \kappa u | \varepsilon_s)} = \frac{T^2}{H + \gamma^2 \sigma_u^2 - \frac{\Big((\Sigma_\varepsilon v)_s\Big)}{\sigma_q^2(s)}} \tag{A.20}$$

which collapses to the baseline case in (A.11) when Σ_ε is diagonal.

In practice, Σ_ε is not diagonal, although building permit growth is only weakly correlated across most states. We present a version of the empirical analog to Σ_ε in Figure A.1. For each state pair, 12-month BPG has a positive correlation. This finding agrees with another prediction of the GS framework, whereby for each locality s, permits are positively related to stock prices and returns (as in **Proposition 1**). There is some evidence of geographic clustering by region, and two states are more likely to have BPG series which co-move if they share a border.

A.4 Incorporating Regulatory Constraints

We endogenize the construction costs in the housing development stage of the model. Construction costs impact the value of the exercised permit via equation (6.1); greater construction costs lower the number of equilibrium permits filed, hampering permit growth q_s relative to its historical baseline. Features of the regulatory environment, such as *ad valorem* permitting fees, administrative delays, and costs of complying with land use statutes.

We incorporate these elements of construction costs by supposing in each location s and period t, a representative developer facing perfect competition hires local labor $N_{s,t}$ to produce next-period net new housing $Q_{s,t}$, subject to decreasing returns to scale:

$$Q_{s,t} = A_s(H_{s,t}, \bar{L}_s) \cdot N_{s,t}^{\rho} \quad \rho \in (0,1) \tag{A.21}$$

The shifter $A_s(H_{s,t}, \bar{L}_s)$ embeds the land endowment $\bar{L}_s$ capturing physical constraints. The factor $H_{s,t}$ accounts for the existing housing stock in an area. Favilukis et al. (2023) model zoning constraints via $A_s = 1 - H_{s,t}/\overline{H}_s$, where $\overline{H}_s$ represents the square footage available for residential use. Zoning may interact with physical geography via $A_s = \bar{L}_s \cdot (1 - H_{s,t}/\overline{H}_s)$. A greater land endowment raises the scale of housing that is feasible, whereas restrictive zoning (i.e., a lower $\overline{H}_s$) renders development more expensive. This particular parameterization captures the idea that producing an additional unit of housing is more expensive if any area is already "built up." For instance, if an area is already nearly at residential capacity, construction may

FIGURE A.1. Cross-State Correlation of Building Permit Growth (Empirical Analog to Σ_ε)

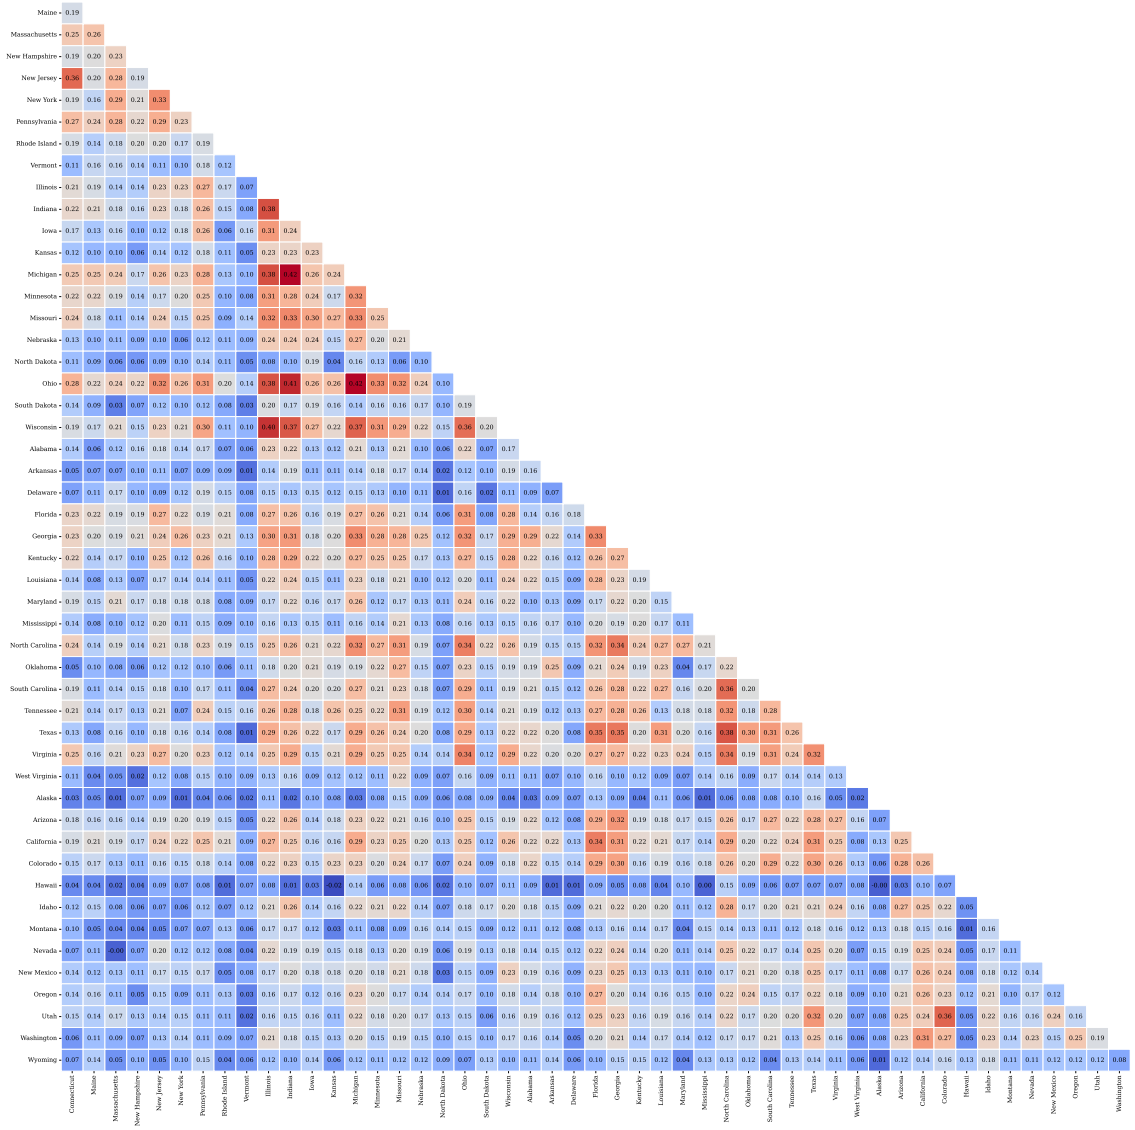

Notes: We report the correlation matrix showing how each , with states organized according to Census region (Northeast, Midwest, South, and West). To construct this matrix, we compute for each state a series of 12-month cumulative building permit growth (BPG) using total residential permits and then take pairwise correlations for each state. Cells in the matrix appearing in red and in darker shades indicate higher correlations, while those in blue and darker shades indicate lower correlations.

require more creative solutions such as building taller structures or on terrain with a greater incline. In general, $\partial A_s/\partial \overline{H}_s > 0$ and $\partial A_s/\partial \overline{L}_s > 0$.

On top of having to abide by zoning and geographical constraints, regulatory wedges include per-unit permit fees $\vartheta_s \geq 0$ and variable, *ad valorem* fees χ_s arising from delays (e.g., forgone rental revenue) or permitting fees that scale with the size of the project.[1] The final component of the housing supply elasticity is compliance convexity. Extra costs may grow faster than linearly with output due to risks stemming from the public process for gaining approval, design iterations, or environmental risks (e.g., discovering contamination).

Under these physical and regulatory constraints on new housing supply, the developer's cost function is:

$$C_s(Q_{s,t}) = \underbrace{W_{s,t}\Big(\frac{Q_{s,t}}{A_s}\Big)^{1/\rho}}_{\text{labor costs}} + \underbrace{\vartheta_s Q_{s,t}}_{\text{per-unit fees}} + \underbrace{\frac{\xi_s}{1+\delta_s} Q_{s,t}^{1+\delta_s}}_{\text{convex compliance}} \tag{A.22}$$

The per-permit cost function $C_{i,s,t}$ is defined analogously for any individual permit $Q_{i,s,t}$. Developers choose to produce housing units $Q_{s,t}$ (or equivalently $N_{s,t}$) to maximize expected profits:

$$\max_{Q_{s,t}\geq 0} \Pi_{s,t} = (1-\chi_s)\cdot \overline{P}_{s,t+1} Q_{s,t} - C_s(Q_{s,t}) \tag{A.23}$$

where $\mathbb{E}_t[P_{s,t+1}] := \overline{P}_{s,t+1}$ is the expected price of completed housing in $t+1$.[2]

We can write the marginal cost curve as:

$$\begin{aligned} &MC_s(Q_{s,t}) = a_s Q_{s,t}^{\alpha} + \vartheta_s + \xi_s Q_{s,t}^{\delta_s} \\ &a_s := W_{s,t} A_s^{-1/\rho}/\rho \quad \text{and} \quad \alpha := 1/\rho - 1 > 0 \end{aligned} \tag{A.24}$$

The developer's zero-profit condition sets $MC_s(Q_{s,t}) = (1-\chi_s)\cdot \overline{P}_{s,t}$.

The price elasticity of housing (permit) supply holding fixed wages, technology A_s and the regulatory parameters is then:

$$\epsilon_s(Q) = \frac{d\log Q_{s,t}}{d\log \overline{P}_{s,t+1}} = \frac{MC_s(Q)}{QMC_s'(Q)} = \frac{a_s Q^{\alpha} + \vartheta_s + \xi_s Q^{\delta_s}}{\alpha a_s Q^{\alpha} + \delta_s \xi_s Q^{\delta_s}} \tag{A.25}$$

If there were no land use regulation ($\vartheta_s = \delta_s = \xi_s, \forall s$, then the elasticity would reduce to the standard decreasing returns elasticity $\epsilon_s(Q) = \rho/(1-\rho)$.

For simplicity, we have assumed all permits applications are approved, so that $Q_{s,t}$ units supplied is equivalent to the number of permits filed. However, this assumption can be easily relaxed by instead supposing permits $\widetilde{Q}_{s,t} = \pi_s(\vartheta;\xi,\delta)\cdot Q_{s,t}$, where the approval probability

[1] See Horton et al. (2024) for a discussion of permit fees for new single-family housing across U.S. counties. Permit fee schedules can have fixed and variable components, depending on the jurisdiction. The variable component is embedded in χ_s, reflecting that most jurisdictions set a progressive schedule of permit fees that increases with the size of the home.

[2] For tractability, in our baseline setup in Section 6 we invoke a replacement cost approach such that the value of a completed unit is always equal to the building's construction cost plus the current value of the underlying land. Effectively, this means developers have rational expectations of sale prices, or $\overline{P}_{s,t+1} = P_{s,t}$.

$\pi_s(\cdot)$ depends on other local land use constraint parameters. If permits are ever rejected ($\pi_s < 1$), then the elasticity formula (A.25) goes through with a multiplicative π_s.

Embedding the developer's block into the construction cost $C_{i,s,t}$ underlying the permit valuation in (6.1)–(6.3) leads us to the following proposition, establishing what we refer to as *Hypothesis 4* in the main text of Section 6.

Proposition 4 (BPG Precision and Housing Supply Elasticity). *Consider a version of the model described in Section 6 in which construction costs determining the probability of permit exercise are pinned down by the developer's cost function in (A.22). If expected house price growth is affine in the risky asset dividend d_t, the precision of local building permit growth as a signal for aggregate asset price movements is negatively correlated with local physical and regulatory constraints on new housing development.*

Proof. Before proving **Proposition 4**, we note that expected house price growth being affine in the risky asset dividend d_t follows from a standard log-affine stochastic discount factor and a log-affine rent process. The procedure follows Campbell et al. (2009), who show via standard Campbell-Shiller log linearization that the rent-price ratio for housing behaves like the stock market's dividend-price ratio in that it forecasts returns and rent growth.[3] This is true at both the national and metro area levels. Lettau and Ludvigson (2001) make a similar argument in showing that the consumption-wealth ratio (with wealth including housing assets) predicts real stock returns. A large class of optimal models of consumer behavior implies that consumption, asset holdings, and labor income are cointegrated.

That is, if the stochastic discount factor is $\log m_{t+1} = \mu_m + \eta_m \cdot d_t + noise$ and the rent process is log-affine, then the Euler equation implies:

$$\overline{P}_{s,t} = \mathbb{E}_t\Big[m_{t+1}\Big(rent_{s,t+1} + (1-\delta_{struct})\cdot P_{s,t+1}\Big)\Big] \tag{A.26}$$

where δ_{struct} is the physical depreciation rate of housing. This follows from the gross housing return identity:

$$R^{(H)}_{s,t+1} = \frac{rent_{s,t+1} + (1-\delta_{struct})\cdot P_{s,t+1}}{P_{s,t}}$$

Under log-normality, expected home sale price growth follows:

$$\Delta \log \overline{P}_{s,t+1} = \lambda_d \cdot d_t + \eta_{s,t}, \quad \text{with } \mathbb{E}[d_t \eta_{s,t}] = 0$$

Up to a first-order log-linearization:

$$\Delta \log Q_{s,t} \approx \epsilon_s \cdot \Delta \log \overline{P}_{s,t+1}$$

So the BPG signal satisfies:

$$q_{s,t} := \Delta \log Q_{s,t} = (\epsilon_s \cdot \lambda_d) d_t + \underbrace{\epsilon_s \eta_{s,t} + u}_{\text{noise}}$$

[3]The affine relationship between house price growth and dividends also follows from a model in which households derive separate utility from consuming housing services and receiving labor income in the form of dividends (i.e., households own the firms).

Rescaling both sides by $\epsilon\lambda_d$ yields the familiar signal format of $z_{s,t} = d_t + \varepsilon^*_{s,t}$. Then the result follows from:

$$Var(\varepsilon^*_{s,t}) = \frac{Var(\text{noise})}{(\epsilon_s \cdot \lambda_t)^2} \implies \tau_q(s) \propto \frac{(\epsilon_s \cdot \lambda_d)^2}{Var(\text{noise})}$$

Hence, holding the noise term fixed across places, a lower supply elasticity ($\epsilon_s \downarrow$) lowers the precision of building permits as a signal of risky dividends. □

B Details on Building Permit Data Construction

We offer further details on how we compile our longitudinal database of U.S. building permit quantities and valuations (described in Section 2.1), including how we sourced the archival reports, digitized the reports and conducted quality control, and how we adjusted the raw series for seasonality.

B.1 Dun & Bradstreet's Building Permits

Figure B.1 illustrates the number of cities reporting building permits each year from 1919 to 1957. The initial count was 164 cities, remaining constant until 1923. Subsequently, significant increases were observed: 177 cities in early 1923, rising to 180 by year-end, and further to 185-187 in 1925. In 1926, the count reached 200, stabilizing at 215 from 1927 onward. Dashed lines indicate years with changes in the number of reporting cities. Data labels above each bar show the exact number of cities for each year. The final set of cities included in the *Dun's* survey covers 44 states, with no cities surveyed in Alaska, Hawaii, Mississippi, Nevada, New Hampshire, and Wyoming. Additionally, we drop New Mexico, Rhode Island, and Vermont from the panel, as those states have incomplete time series for their cities.

We aggregate the city-level *Dun's* permit value totals up to the state level to use in our analysis. Aggregating to the state level helps reduce the scope for measurement error arising from the staggered entry of smaller cities into the survey and the fact that permits can be reported in a lumpy fashion within the year, with different cities within each state operating on different filing calendars due to differences in how the local town clerks report records to the *Dun's* inspector. We aggregate to the state level in two ways. First, we simply total all permit values across cities within the same state-month; we plot this as the "unweighted" series in Figure B.2. Second, we compute a weighted state-level total by weighting each city's permit count in a given month in proportion to its Census population relative to the statewide population in that year; we plot this as the "population-weighted" series in Figure B.2. Since city-level population estimates are missing for some cities in the earlier part of the sample, we cannot compute population-weighted permit totals for 11 states, reducing the number of states in our panel to 30. However, population weighting helps account for the fact that city boundaries may have shifted over time, particularly during episodes of mass migration during the 1920s and the Dust Bowl period of the 1930s.

We emphasize a few broad patterns in the permits data for the pre-1960s period, as pictured in Figure B.2. We plot the X-13 seasonally adjusted per capita series, where we run the X-13 ARIMA-SEATS filter separately for the population-weighted and unweighted series. We describe the seasonal adjustment procedures in Appendix B.4. First, the data pick up the Roaring 1920s and postwar baby boom periods—with permits skyrocketing across most states during

FIGURE B.1. Number of Cities Reporting Building Permits: 1919–1957.

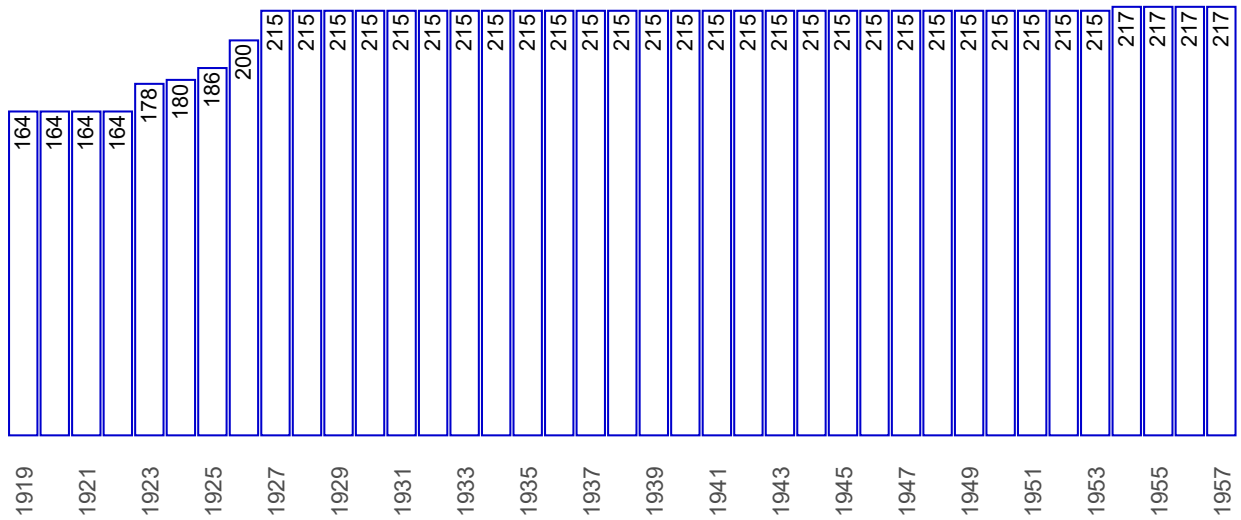

Notes: This figure illustrates the number of cities reporting building permits each year in Dun & Bradstreet's publications. The initial count was 164 cities, remaining constant until 1923. Subsequently, significant increases were observed: 177 cities in early 1923, rising to 180 by year-end, and further to 185-187 in 1925. In 1926, the count reached 200, stabilizing at 215 from 1927. Two cities, Albuquerque, NM and Billings, MT were added in 1954M1 until *Dun's Statistical Review* ceased publication in 1957M10.

those decades. Second, comparing the population-weighted vs. unweighted series reveals the influence of inter-state migratory patterns. For instance, the wedge between the weighted and unweighted series is largest in states like California and Florida, which experienced mass migration during the 1930s Dust Bowl period (Baerlocher et al., 2024) and 1920s Florida Land Boom (Knowlton, 2020; Calomiris and Jaremski, 2023), respectively. Third, there are idiosyncratic spikes in permitting activity for some states which do not seem to correspond to any imminent recession risk. A notable example of this is that New York State permits jump in 1957M1, corresponding to the state legislature's passage of an overhaul of the property tax system which effectively increased tax rates levied on new commercial properties. For this reason, in our analysis using the *Dun's* data in Section 4.2, we winsorize the resulting BPG volatility series at the 1% tails within each state to minimize the influence of such outliers in some regressions.

B.2 U.S. Census Bureau's Building Permits Reports

Origins of Census BPS Reports (1959–1969). Our data from 1959 onward originate from the Census Building Permit Survey (BPS). The BPS consists of four main series of interest with publication dating back to May 1959 when the survey was piloted (note that the series names have changed slightly over time):

1. **C20: Housing Starts and Building Permits:** Monthly counts of housing starts and building permits filed for single-family and multifamily housing units and mobile homes. The Census ceased this publication in 2001. Since the information in this series is not geographically disaggregated and aggregated permit counts are provided in the C40 and C42 reports (see description below), we do not use this series in our main analysis.

FIGURE B.2. Real Building Permit Values Per 10,000 Inhabitants in the Cross-Section of U.S. States, 1919–1957

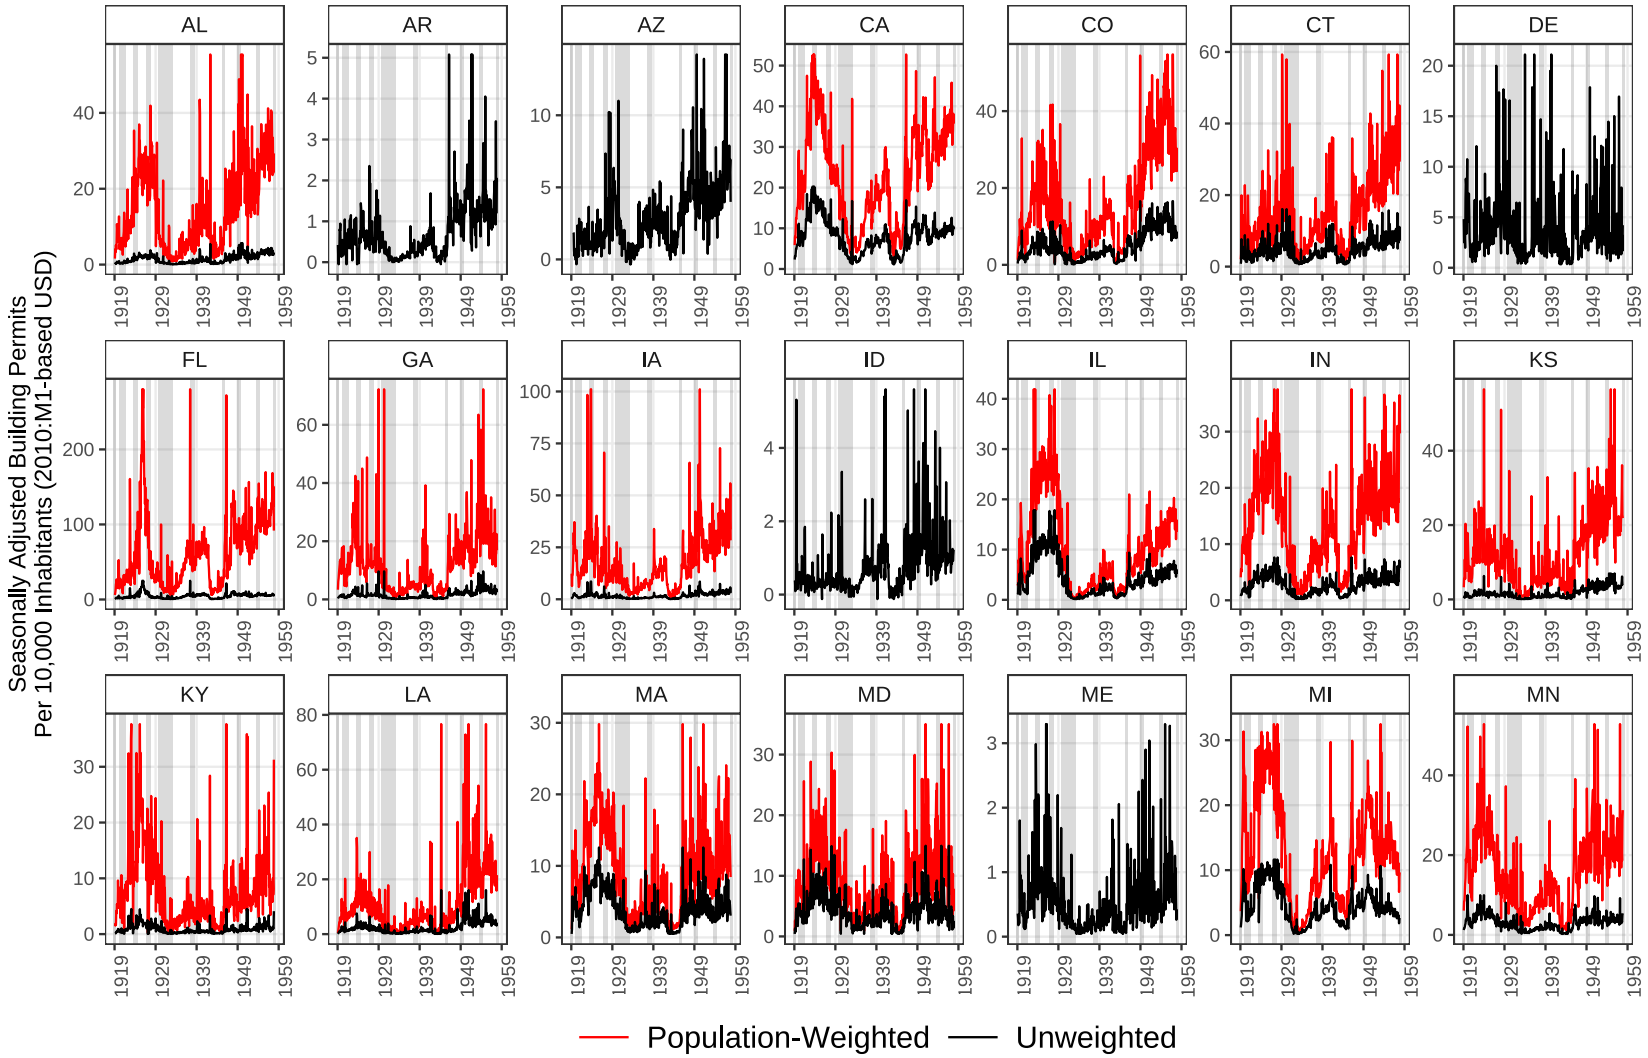

Notes: The figure plots the seasonally adjusted real value of building permits per 10,000 inhabitants for each state between 1919M1–1957M10. Real values refer to 2010M1-denominated U.S. dollars. Black lines depict unweighted series as reported in *Dun's Review* while the red lines depict population-weighted series. To construct an unweighted series, we simply total all permits across all cities within a state. The weighted series is constructed using U.S. Census city-level population counts to weight each city's data according to inverse-proportionality, weighting its share of the state's population as of that year. We omit states with only one city and/or large time series gaps in the *Dun's* survey, including NM, RI, and VT. We seasonally adjust each series using the Census X-13 ARIMA-SEATS filter (cf. Appendix B.4). Grey-shaded areas indicate NBER-dated recessions.

FIGURE B.2. Real Building Permit Values Per 10,000 Inhabitants in the Cross-Section of U.S. States, 1919–1957 (Continued)

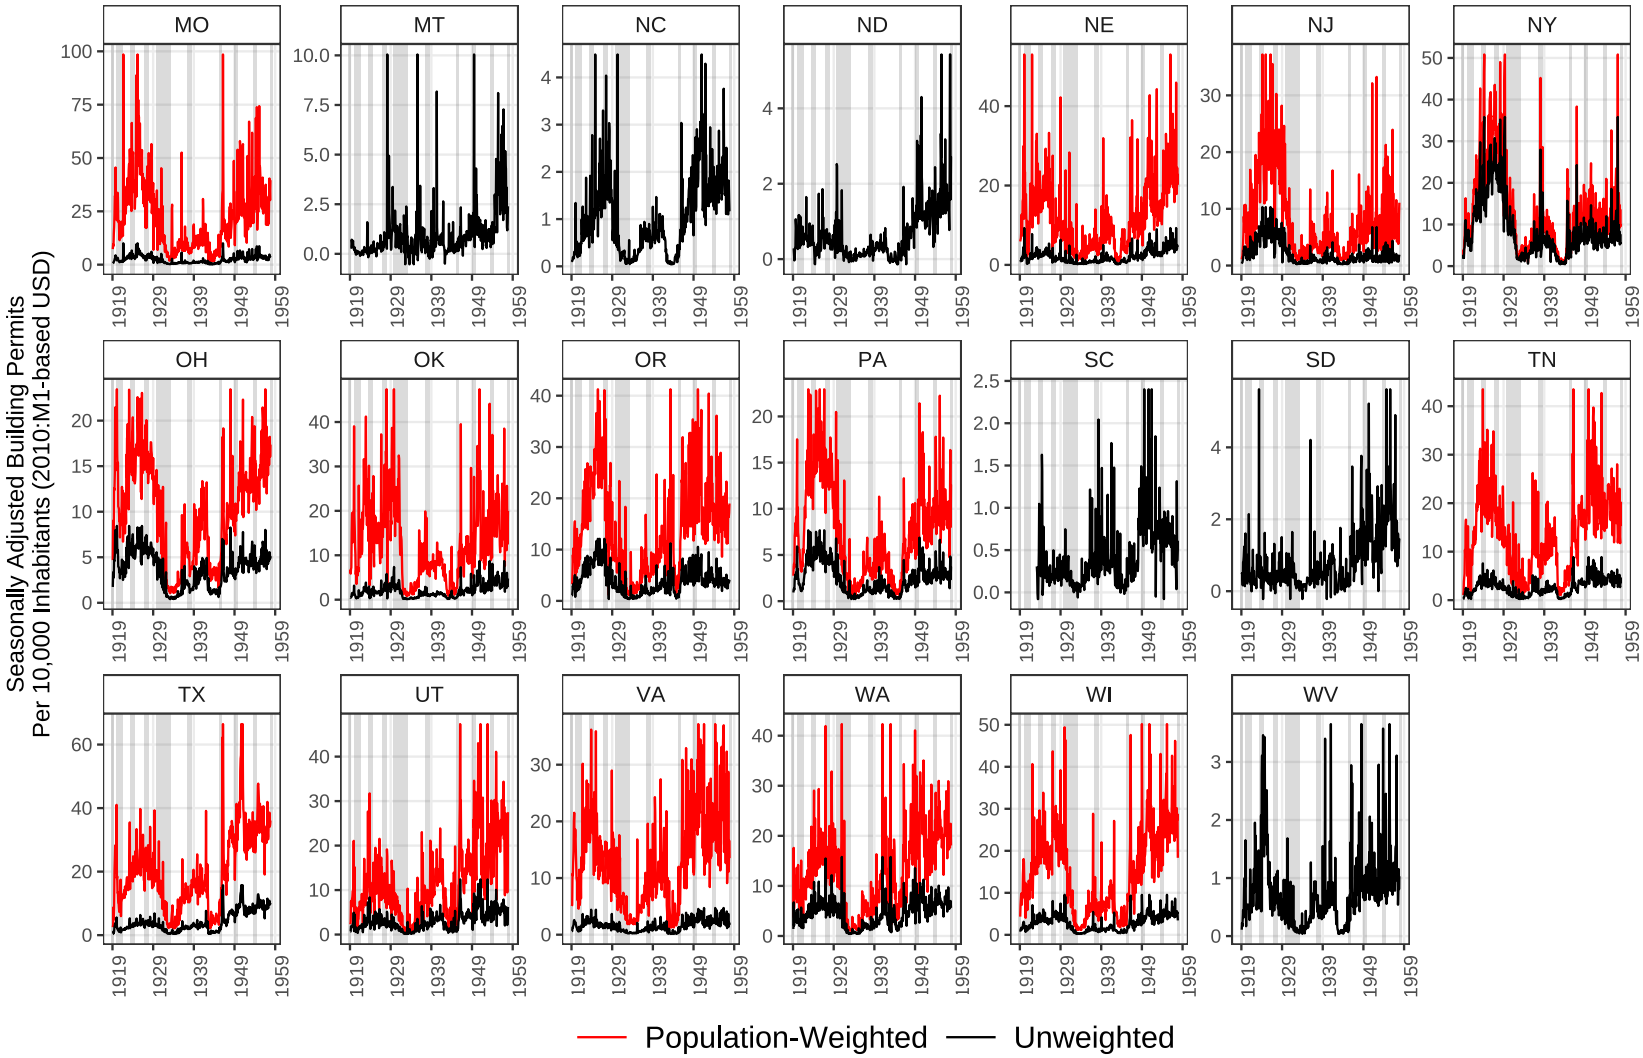

Notes: The figure plots the seasonally adjusted real value of building permits per 10,000 inhabitants for each state between 1919M1–1957M10. Real values refer to 2010M1-denominated U.S. dollars. Black lines depict unweighted series as reported in *Dun's Review* while the red lines depict population-weighted series. To construct an unweighted series, we simply total all permits across all cities within a state. The weighted series is constructed using U.S. Census city-level population counts to weight each city's data according to inverse-proportionality, weighting its share of the state's population as of that year. We omit states with only one city and/or large time series gaps in the *Dun's* survey, including NM, RI, and VT. We seasonally adjust each series using the Census X-13ARIMA-SEATS filter (cf. Appendix B.4). Grey-shaded areas indicate NBER-dated recessions.

2. **C40: Housing Authorized in Individual Permit-Issuing Places:** Monthly permit counts and valuations reported at the county and town level. See description below for more details.

3. **C42: New Residential Construction Authorized in Permit-Issuing Places:** Monthly permit counts and valuations reported at the MSA and state level. See description below for more details.

4. **C50: Residential Alterations and Repairs:** quarterly data on expenditures towards residential additions, alternations, maintenance and repairs, and replacements. The information in the C50 reports is sourced from household surveys and geographically disaggregated only to the Census region level. We do not use this series in our main analysis. The C50 series forms the basis for the Survey of Residential Alterations and Repairs, which the Census discontinued in 2007Q4.[4]

Most of the data on permit counts and valuations used in our main analysis are covered by the C40 and C42 series, depending on the month-year date and level of geography.

Post-1966, the Census also publishes annual summary statistics combining information from the C40 and C42 series, with the C40 and C42 labels used interchangeably for these annual reports. For example, the disclaimer from the 1968-1969 C40 summary statistics volume informs us that:

> *"Data on housing authored in permit-issuing places are published by the Bureau of the Census in two reports, C.40, Housing Authorized by Building Permits and Public Contracts: Individual Places, which is sold by the Government Printing Office, and C.42, Housing Authorized by Building Permits and Public Contracts: States and Selected Metropolitan Statistical Areas, which is sold by the Bureau of the Census. The 1966 annual summaries of these reports were consolidated and issued through the Government Printing Office under the joint designation C.40/C.42."*

The Census collects each of the four series at a monthly frequency, and the naming convention of the periodical follows the year-month. Example: C42-68-12 would contain state and MSA-level residential building permits issued in December 1968. Each set of monthly tables is preceded by a cover page that indicates the month and year in which the numbers pertain. This information on the cover page needs to be used in place of the publication date, since the publication date lags the actual information collection. The tables also sometimes, but not always, list the date the information was reported in the title caption.

For most of the 1960s, the tables in C40 are disaggregated to a "place" (i.e., a "town" or "county" level), whereas the C42 series contain state and MSA-level tables. The Census cautions that the county and place-level statistics may not aggregate to match the state or MSA-level totals due to rounding and the imputation procedures the Census employs to deal with survey non-response in creating state or MSA-level totals. Under the consolidation between the Government Printing Office and Census Bureau mentioned in the above quote, starting in 1969 the labeling reverted back to C40 being used to refer to the monthly reports for all three levels of geography. The title of C40 was renamed to "Housing Authorized by Building Permits and Public Contracts." After 1969, the Census retired the C42 series label.

[4]Data for the modern Survey of Residential Alterations and Repairs can be found on the Census website: https://www.census.gov/construction/c50/c50index.html.

1970s BPS Series. Starting in 1970, the information previously contained in C42 was incorporated into an expanded C40 series. As a result of this consolidation, the MSA-level tables in the 1970s contain more MSAs than in the 1960s, since the totals were apparently cross-walked between the "place" to the relevant MSA definition of the day. The state-level information is contained in "Table 2. Region, Division State, and Puerto Rico: New Housing Units Authorized in Permit-Issuing Places." The MSA-level information is contained in "Table 3. Selected Standard Metropolitan Statistical Areas: New Housing Units Authorized in Permit-Issuing Places" in both the 1970s and 1980s.

1980s BPS Series. In the 1980s, the format stayed largely the same as in the 1970s, except the December monthly tables are published together with the annual reports. The content of reforms to the survey, which resulted in the modern data beginning in 1988, is alluded to on the cover page of the January 1987 report (C40-87-1):

> *"Beginning with data for April 1987, two changes will be made to the procedures used to seasonally adjust building permits estimates. First, the seasonal adjustment factors will be recomputed each month using all available data rather than using projections based on data through March of the previous year. Second, except for the total and one-unit estimates, published figures will be adjusted directly rather than being derived by summing adjusted components. The one-unit (single-family) estimate will be derived by adjusting and summing its regional components. The total will be derived by summing the adjusted estimates of units by type of structure (1, 2-4, and 5 or more). Each regional total will be adjusted separately. Studies have shown that the revised methodology will usually result in small revisions."*

Monthly data at the state and sub-state levels only start in January 1988 in the master BPS file downloadable directly from the U.S. Census website, or for individual geographies via the FRED API.

Sourcing the reports. We received a collection of the pre-1988 scanned C42 reports directly from the Economic Indicators Division of the Census Bureau. We then downloaded from HathiTrust all non-duplicate scans of the *.PDF and *.TXT files from the publications for the series in the month-years other than the ones we received directly from the Census Bureau. We inspect the *.TXT files in cases where the scan is deprecated to the extent that we cannot read specific entries. However, we use our own OCR procedures to digitize the data, as described in Appendix B.3. For years before 1970, we focus on the C42 scanned PDFs, which contain many monthly tables appended within the same year or a subset of adjacent years. For the 1970–1988, the state-level and MSA-level tables are in the C40 series. We then screened all relevant tables within each scanned report for quality using the OCR tools. In cases where an entire scanned table was too deprecated in the HathiTrust volumes, we contacted the Regional Federal Depository Librarian at the Connecticut State Library to obtain fresh, high-quality scans.

In Figure B.3 we plot the raw seasonally adjusted permit counts we collected from the Census BPS, starting in 1961, the first year where the survey follows its now-standard format. The displayed trends are similar relative to Figure 2 which plots per capita permits. Most states experienced a large boom in new residential permitting during the 2000s boom, but single-family permit counts had only recovered in a small handful of states in the South (NC, TN, TX) on the eve of the first wave of COVID-19 lockdowns in March 2020. Due to local public health

restrictions (Ghent et al., 2024) and weakened demand for leased space due to COVID-19 exposure risk (Ling et al., 2020), permitting activity precipitously fell for both single-family and multifamily units during 2020 and 2021.

B.3 Digitization Procedures

To create our geographically disaggregated database of building permits, we rely on several optimal character recognition (OCR) techniques to scale up the digitization from scanned tables in PDFs. Collecting all the information used in our main analysis would be excessively time consuming given the monthly frequency of the reports and our use of state and MSA-level (or city) data. For the post-1960s segment of our sample from the Census Building Permits Survey, each annual volume consisting of the twelve monthly reports and annual summary statistics regularly comprises over 500 pages of tables and text. There is less information to process in the *Dun's Review* publications, since there is only one level of aggregation (city) and no separate tables for permit counts vs. valuation totals and single-family vs. total residential permits as there is in the modern Census survey.

We combine two sets of OCR tools to digitize all the permit tables. First, we use the standard ABBYY FineReader PDF software, combined with a customized Excel VBA code to purge the output of any extraneous characters (e.g. "@"). Second, we use the *Layout Parser* Python package designed by Shen et al. (2021) to apply modern deep learning algorithms to the digitization of large-scale corpuses of historical text. We use the `Tesseract` OCR engine to implement the *Layout Parser* routine. `Tesseract` provides a confidence level, or "score" for each rectangular block it identifies in the data; the score is set to -1 if `Tesseract` identifies a block but fails to detect any characters in that cell. We drop any output from blocks with a score less than 90 and filter out any extraneous characters from the remaining output.

After implementing both methods, we then compare the fraction of cells on each page populated data to the fraction obtained from ABBYY. For over 70% of the pages, the ABBYY plus VBA method outperforms *Layout Parser*, but we use both methods because use of ABBYY requires more "point and click" mechanical monitoring. For each table, we take the output file which covers a greater fraction of cells. We then perform quality control by comparing reported subtotals within each table to totals implied by cells within each row. For instance, in the Census BPS, we check whether the "total" column matches the total number of permits obtained from summing across the 1-unit, 2-units, 3 and 4-units, and 5-units or more totals. We set a tolerance threshold such that cases where the row totals do not match can only be due to errors in the digitization process rather than rounding. Tagging rows where the difference in row totals fall outside this tolerance, we obtain an error rate of 1% to 2% of table rows, depending on the sample period. We then assign hand-correction of the errors to our team of RAs. Finally, we standardize the place names across tables within each underlying source publication, checking the Census survey documentation to ensure the geographic span of the permit totals are consistent over time.

Figure B.4 provides an example of how *Layout Parser* identifies rectangular blocks of text (highlighted in red) to convert to machine-readable text. The example is a scan from the March 1986 MSA-level table with permit counts from the Census Building Permits Survey. This "token" structure helps isolate place names in the table rows, but often struggles with particular formatting conventions, such as the large type spaces between numerical characters for larger counts and the ellipses in the row headings. Figure B.5 displays the resulting output from *Layout Parser* on the left-hand side opposite the raw PDF table scan on the right-hand side. Comparing

FIGURE B.3. Permit Counts in the Cross-Section of U.S. States, 1961–2022

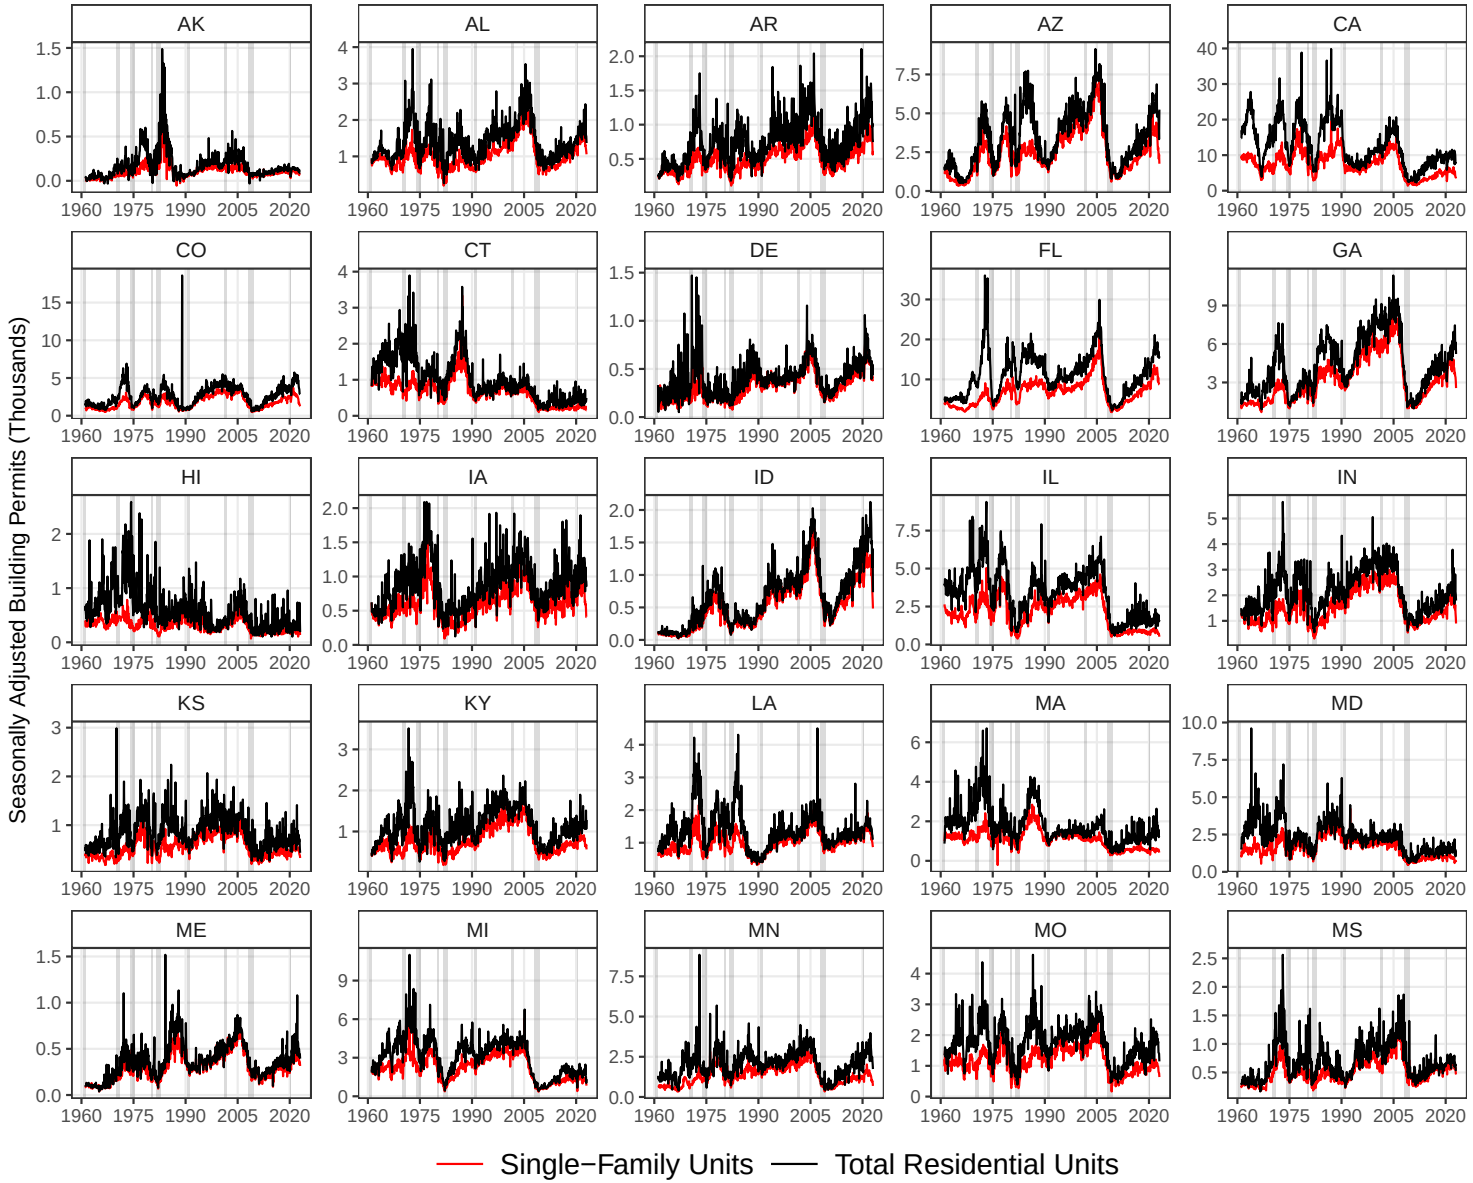

Notes: The figure plots for each state the seasonally adjusted number of total private residential permits (black) and private single-family home permits (red) filed in that state for a given month according to the Census Building Permits Survey (1961M1–2022M12). The time series of permit counts excludes permits filed for public construction contracts. We seasonally adjust each series using the Census X-13 ARIMA-SEATS filter (cf. Appendix B.4). Grey-shaded areas indicate NBER-dated recessions. We start the sample at 1961M1, since 1961 is the first year the survey follows its standard format.

FIGURE B.3. Permit Counts in the Cross-Section of U.S. States, 1961–2022 (Continued)

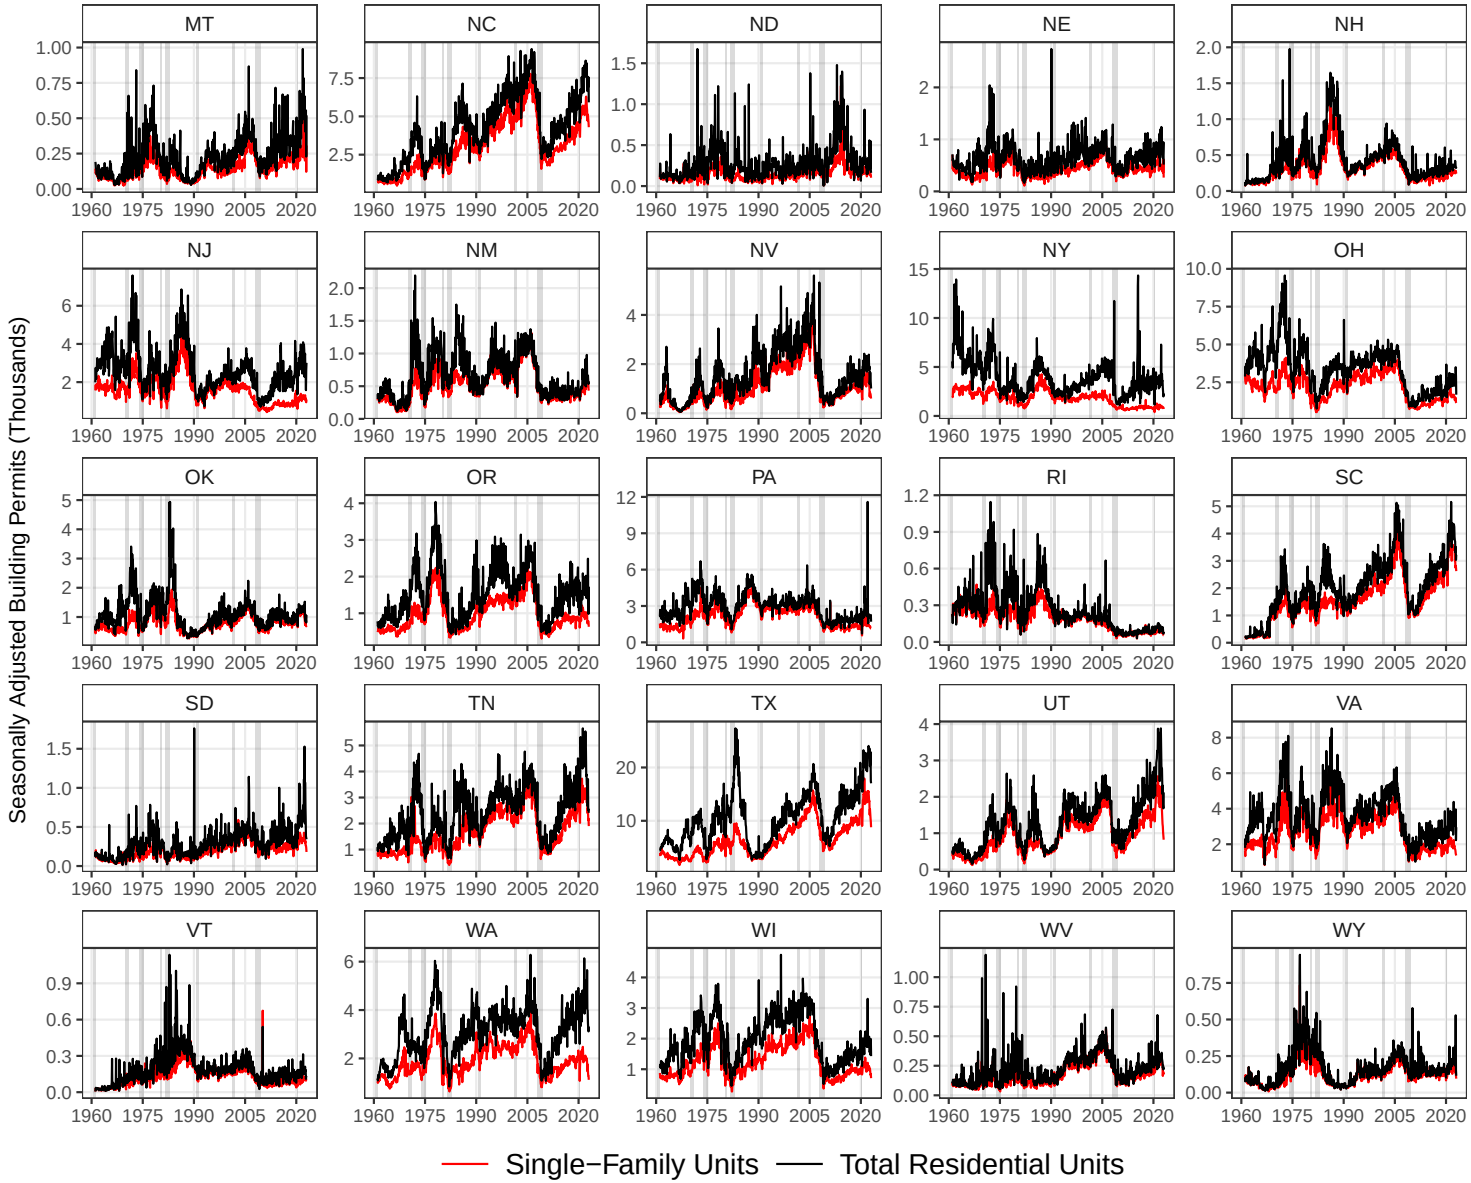

Notes: The figure plots for each state the seasonally adjusted number of total private residential permits (black) and private single-family home permits (red) filed in that state for a given month according to the Census Building Permits Survey (1961M1–2022M12). The time series of permit counts excludes permits filed for public construction contracts. We seasonally adjust each series using the Census X-13 ARIMA-SEATS filter (cf. Appendix B.4). Grey-shaded areas indicate NBER-dated recessions. We start the sample at 1961M1, since 1961 is the first year in which the survey follows its standard format.

the two tables, the package is able to correctly identify most of the numerical and missing entries (indicated by an en-dash). However, extraneous characters appear in the place names, and there is no easy way to automatically fix the formatting so that row headings are attached to the correct set of data entries. We also cannot retroactively add back in the row labels using a common alphabetical ordering given that the set of MSAs (or counties and towns) included in the changes month-to-month. For states, this drawback to *Layout Parser* is less of an issue.

Layout Parser's performance improves relative to ABBYY during the earlier period of our sample in which we source the permits data from scans of the tables in *Dun's Review*, as described in Appendix B.1. Figure B.6 shows sample output from a scan of part of the table of annual permit valuation totals reported in the 1939 volume of *Dun's Review*. Relative to the modern data source, *Layout Parser* produces fewer typos for the numerical entries in the pre-1960s tables. The training dataset used in the "Fast" version of the Shen et al. (2021) base model we use consists of early 20th century newspapers, resulting in more accurate renderings of the permits tables. While this could be due to the differences in contrast across training datasets (i.e. the pages are more or less yellowed due to age), the gains in performance are limited even if we experiment with the contrast settings for reading the modern Census tables into the package. We use GPUs to run the routine in an attempt to best mimic the computing environment used to train the digitization algorithm.

B.4 Seasonal Adjustment Methods

Given the high degree of intra-year cyclicality in the real estate sector relative to the seasonality of stock and bond returns, it is critical to seasonally adjust our building permits series before computing growth rates and volatilities. The Census provides its X-13ARIMA-SEATS program to strip macro time series of their seasonal components.[5] We set up a Linux machine to run the X-13 routine and adapt the underlying source code to accommodate longer time series so that we can apply the filter to our entire sample. To ensure that we are able to match the seasonal adjustments applied by the Census to the modern time sample, we run the X-13 filter on the unadjusted permit counts series obtained through the FRED API and compare our resulting seasonally adjusted series to the seasonally adjusted series publicly available through FRED.

Appendix B.4 shows that in the monthly state-level panel of building permit levels, our X-13 filtered series and the seasonally adjusted series downloadable from the Census via FRED line up almost perfectly, with an R^2 of 99.99%. One limitation to the X-13 filter routine provided by the Census is that it cannot accommodate missing values. Hence, for a small number of cases where the Census monthly report lists an "NA" value for permit counts of valuations, we interpolate using the midpoint between adjacent months with non-missing values for a given locality. For instance, permit counts are missing in May 1976 for Massachusetts, so we fill in the value for May 1976 using the average of the values reported for Massachusetts in April and June of that year. Our use of midpoint interpolation leads to the negligible differences between seasonally adjusted permit levels visible in Appendix B.4.

To maintain consistency over our full permits database spanning 100 years, we use the X-13 filter with this midpoint convention for filling in missing values in both the post-1960s Census period of our sample and the pre-1960s sample sourced from *Dun's Review*. Our results are qualitatively similar, albeit estimated with larger standard errors, if we instead

[5] The documentation and interface programs can be downloaded here: https://www.census.gov/data/software/x13as.html.

FIGURE B.4. Example: *Layout Parser* Fields from March 1986 Building Permits Survey

Notes: Example of the *Layout Parser* package of Shen et al. (2021) applied to Table 3 from the March 1986 Census Building Permits Survey. We received a PDF copy of this report directly from the Economic Indicators Division of the Census Bureau. The red boxes indicate "blocks" of text identified by *Layout Parser*.

FIGURE B.5. Example: *Layout Parser* Output from March 1986 Building Permits Survey

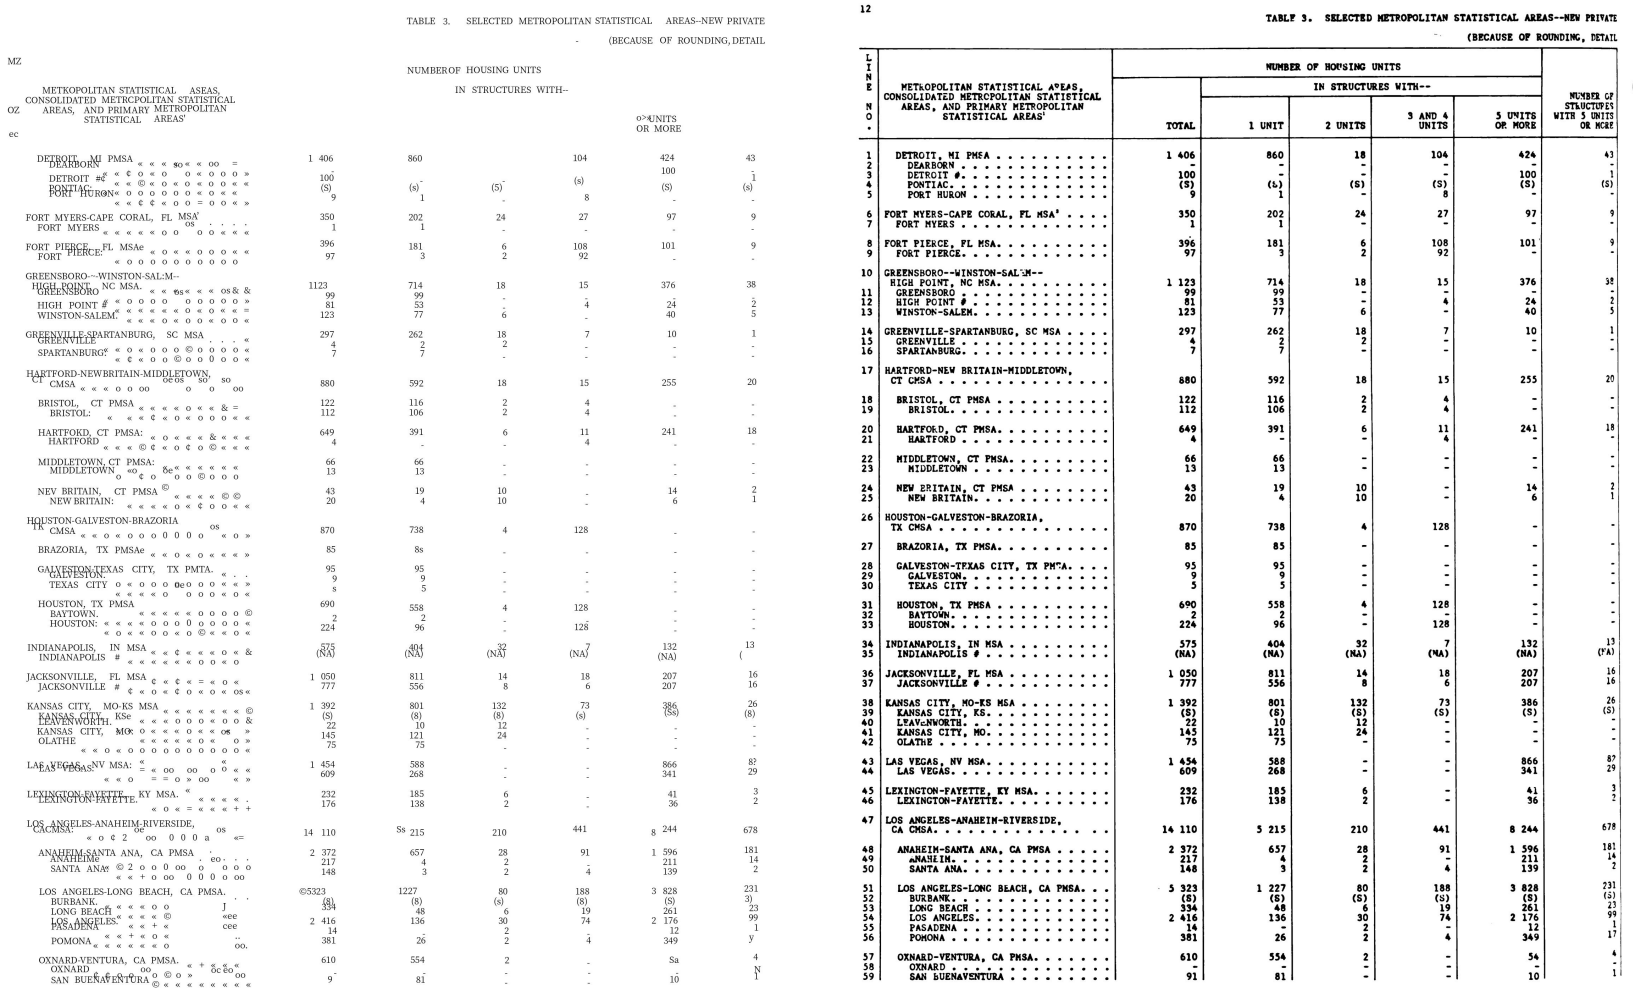

TABLE 3. SELECTED METROPOLITAN STATISTICAL AREAS--NEW PRIVATE

(BECAUSE OF ROUNDING, DETAIL

	METROPOLITAN STATISTICAL AREAS, CONSOLIDATED METROPOLITAN STATISTICAL AREAS, AND PRIMARY METROPOLITAN STATISTICAL AREAS'	NUMBER OF HOUSING UNITS					
			IN STRUCTURES WITH--				
						5 UNITS OR MORE	
	DETROIT, MI PMSA	1 406	860		104	424	43
	DEARBORN					100	1
	DETROIT	100	-	-	(s)	(S)	(s)
	PONTIAC	(S)	(s)	(5)			
	FORT HURON	9	1	-	8	-	-
	FORT MYERS-CAPE CORAL, FL MSA'	350	202	24	27	97	9
	FORT MYERS	1	1	-	-	-	-
	FORT PIERCE, FL MSAe	396	181	6	108	101	9
	FORT PIERCE.	97	3	2	92	-	-
	GREENSBORO--WINSTON-SAL:M--						
	HIGH POINT, NC MSA.	1123	714	18	15	376	38
	GREENSBORO	99	99	-	-	-	-
	HIGH POINT	81	53	-	4	24	2
	WINSTON-SALEM.	123	77	6	-	40	5
	GREENVILLE-SPARTANBURG, SC MSA	297	262	18	7	10	1
	GREENVILLE	4	2	2	-	-	-
	SPARTANBURG.	7	7	-	-	-	-
	HARTFORD-NEWBRITAIN-MIDDLETOWN, CT CMSA	880	592	18	15	255	20
	BRISTOL, CT PMSA	122	116	2	4	-	-
	BRISTOL:	112	106	2	4	-	-
	HARTFOKD, CT PMSA:	649	391	6	11	241	18
	HARTFORD	4	-	-	4	-	-
	MIDDLETOWN, CT PMSA:	66	66	-	-	-	-
	MIDDLETOWN	13	13	-	-	-	-
	NEV BRITAIN, CT PMSA	43	19	10	-	14	2
	NEW BRITAIN:	20	4	10	-	6	1
	HOUSTON-GALVESTON-BRAZORIA TX CMSA	870	738	4	128	-	-
	BRAZORIA, TX PMSAe	85	8s	-	-	-	-
	GALVESTON-TEXAS CITY, TX PMTA.	95	95	-	-	-	-
	GALVESTON	9	9	-	-	-	-
	TEXAS CITY	s	5	-	-	-	-
	HOUSTON, TX PMSA	690	558	4	128	-	-
	BATTOWN.	2	2	-	-	-	-
	HOUSTON:	224	96	-	128	-	-
	INDIANAPOLIS, IN MSA	575	404	32	7	132	13
	INDIANAPOLIS #	(NA)	(NA)	(NA)	(NA)	(NA)	(
	JACKSONVILLE, FL MSA	1 050	811	14	18	207	16
	JACKSONVILLE #	777	556	8	6	207	16
	KANSAS CITY, MO-KS MSA	1 392	801	132	73	386	26
	KANSAS CITY, KSe	(S)	(8)	(8)	(s)	(Ss)	(8)
	LEAVENWORTH.	22	10	12	-	-	-
	KANSAS CITY, MO:	145	121	24	-	-	-
	OLATHE	75	75	-	-	-	-
	LAS VEGAS, NV MSA:	1 454	588	-	-	866	8?
	LAS VEGAS.	609	268	-	-	341	29
	LEXINGTON-FAYETTE, KY MSA.	232	185	6	-	41	3
	LEXINGTON-FAYETTE.	176	138	2	-	36	2
	LOS ANGELES-ANAHEIM-RIVERSIDE, CACMSA:	14 110	Ss 215	210	441	8 244	678
	ANAHEIM-SANTA ANA, CA PMSA	2 372	657	28	91	1 596	181
	ANAHEIME	217	4	2	-	211	14
	SANTA ANA:	148	3	2	4	139	2
	LOS ANGELES-LONG BEACH, CA PMSA.	©5323	1227	80	188	3 828	231
	BURBANK.	(8)	(8)	(s)	(8)	(S)	3)
	LONG BEACH	334	48	6	19	261	23
	LOS ANGELES.	2 416	136	30	74	2 176	99
	PASADENA	14	-	2	-	12	1
	POMONA	381	26	2	4	349	17
	OXNARD-VENTURA, CA PMSA.	610	554	2	-	Sa	4
	OXNARD	-	-	-	-	-	N
	SAN BUENAVENTURA	9	81	-	-	10	1

12

TABLE 3. SELECTED METROPOLITAN STATISTICAL AREAS--NEW PRIVATE

(BECAUSE OF ROUNDING, DETAIL

LINE NO.	METROPOLITAN STATISTICAL AREAS, CONSOLIDATED METROPOLITAN STATISTICAL AREAS, AND PRIMARY METROPOLITAN STATISTICAL AREAS'	NUMBER OF HOUSING UNITS					
			IN STRUCTURES WITH--				
		TOTAL	1 UNIT	2 UNITS	3 AND 4 UNITS	5 UNITS OR MORE	NUMBER OF STRUCTURES WITH 5 UNITS OR MORE
1	DETROIT, MI PMSA	1 406	860	18	104	424	43
2	DEARBORN	-	-	-	-	-	-
3	DETROIT #.	100	-	-	-	100	1
4	PONTIAC.	(S)	(s)	(S)	(S)	(S)	(S)
5	PORT HURON	9	1	-	8	-	-
6	FORT MYERS-CAPE CORAL, FL MSA'	350	202	24	27	97	9
7	FORT MYERS	1	1	-	-	-	-
8	FORT PIERCE, FL MSA.	396	181	6	108	101	9
9	FORT PIERCE.	97	3	2	92	-	-
10	GREENSBORO--WINSTON-SAL:M-- HIGH POINT, NC MSA.	1 123	714	18	15	376	38
11	GREENSBORO	99	99	-	-	-	-
12	HIGH POINT #	81	53	-	4	24	2
13	WINSTON-SALEM.	123	77	6	-	40	5
14	GREENVILLE-SPARTANBURG, SC MSA	297	262	18	7	10	1
15	GREENVILLE	4	2	2	-	-	-
16	SPARTANBURG.	7	7	-	-	-	-
17	HARTFORD-NEW BRITAIN-MIDDLETOWN, CT CMSA	880	592	18	15	255	20
18	BRISTOL, CT PMSA	122	116	2	4	-	-
19	BRISTOL.	112	106	2	4	-	-
20	HARTFORD, CT PMSA.	649	391	6	11	241	18
21	HARTFORD	4	-	-	4	-	-
22	MIDDLETOWN, CT PMSA.	66	66	-	-	-	-
23	MIDDLETOWN	13	13	-	-	-	-
24	NEW BRITAIN, CT PMSA	43	19	10	-	14	2
25	NEW BRITAIN.	20	4	10	-	6	1
26	HOUSTON-GALVESTON-BRAZORIA, TX CMSA	870	738	4	128	-	-
27	BRAZORIA, TX PMSA.	85	85	-	-	-	-
28	GALVESTON-TEXAS CITY, TX PMTA.	95	95	-	-	-	-
29	GALVESTON.	9	9	-	-	-	-
30	TEXAS CITY	5	5	-	-	-	-
31	HOUSTON, TX PMSA	690	558	4	128	-	-
32	BATTOWN.	2	2	-	-	-	-
33	HOUSTON.	224	96	-	128	-	-
34	INDIANAPOLIS, IN MSA	575	404	32	7	132	13
35	INDIANAPOLIS #	(NA)	(NA)	(NA)	(NA)	(NA)	(NA)
36	JACKSONVILLE, FL MSA	1 050	811	14	18	207	16
37	JACKSONVILLE #	777	556	8	6	207	16
38	KANSAS CITY, MO-KS MSA	1 392	801	132	73	386	26
39	KANSAS CITY, KS.	(S)	(8)	(8)	(S)	(S)	(S)
40	LEAVENWORTH.	22	10	12	-	-	-
41	KANSAS CITY, MO.	145	121	24	-	-	-
42	OLATHE	75	75	-	-	-	-
43	LAS VEGAS, NV MSA.	1 454	588	-	-	866	87
44	LAS VEGAS.	609	268	-	-	341	29
45	LEXINGTON-FAYETTE, KY MSA.	232	185	6	-	41	3
46	LEXINGTON-FAYETTE.	176	138	2	-	36	2
47	LOS ANGELES-ANAHEIM-RIVERSIDE, CA CMSA.	14 110	5 215	210	441	8 244	678
48	ANAHEIM-SANTA ANA, CA PMSA	2 372	657	28	91	1 596	181
49	ANAHEIM.	217	4	2	-	211	14
50	SANTA ANA.	148	3	2	4	139	2
51	LOS ANGELES-LONG BEACH, CA PMSA.	5 323	1 227	80	188	3 828	231
52	BURBANK.	(8)	(8)	(S)	(S)	(S)	(S)
53	LONG BEACH	334	48	6	19	261	23
54	LOS ANGELES.	2 416	136	30	74	2 176	99
55	PASADENA	14	-	2	-	12	1
56	POMONA	381	26	2	4	349	17
57	OXNARD-VENTURA, CA PMSA.	610	554	2	-	54	4
58	OXNARD	-	-	-	-	-	-
59	SAN BUENAVENTURA	91	81	-	-	10	1

Notes: Example of the output (left) obtained from applying *Layout Parser* to Table 3 from the March 1986 Building Permits Survey, juxtaposed against the original scanned table (right). We filter out non-standard characters from this output, score the quality of the output, and then compare the result against output obtained from ABBYY FineReader applied to the same scan.

FIGURE B.6. Example: *Layout Parser* Output from 1939 *Dun's Review*

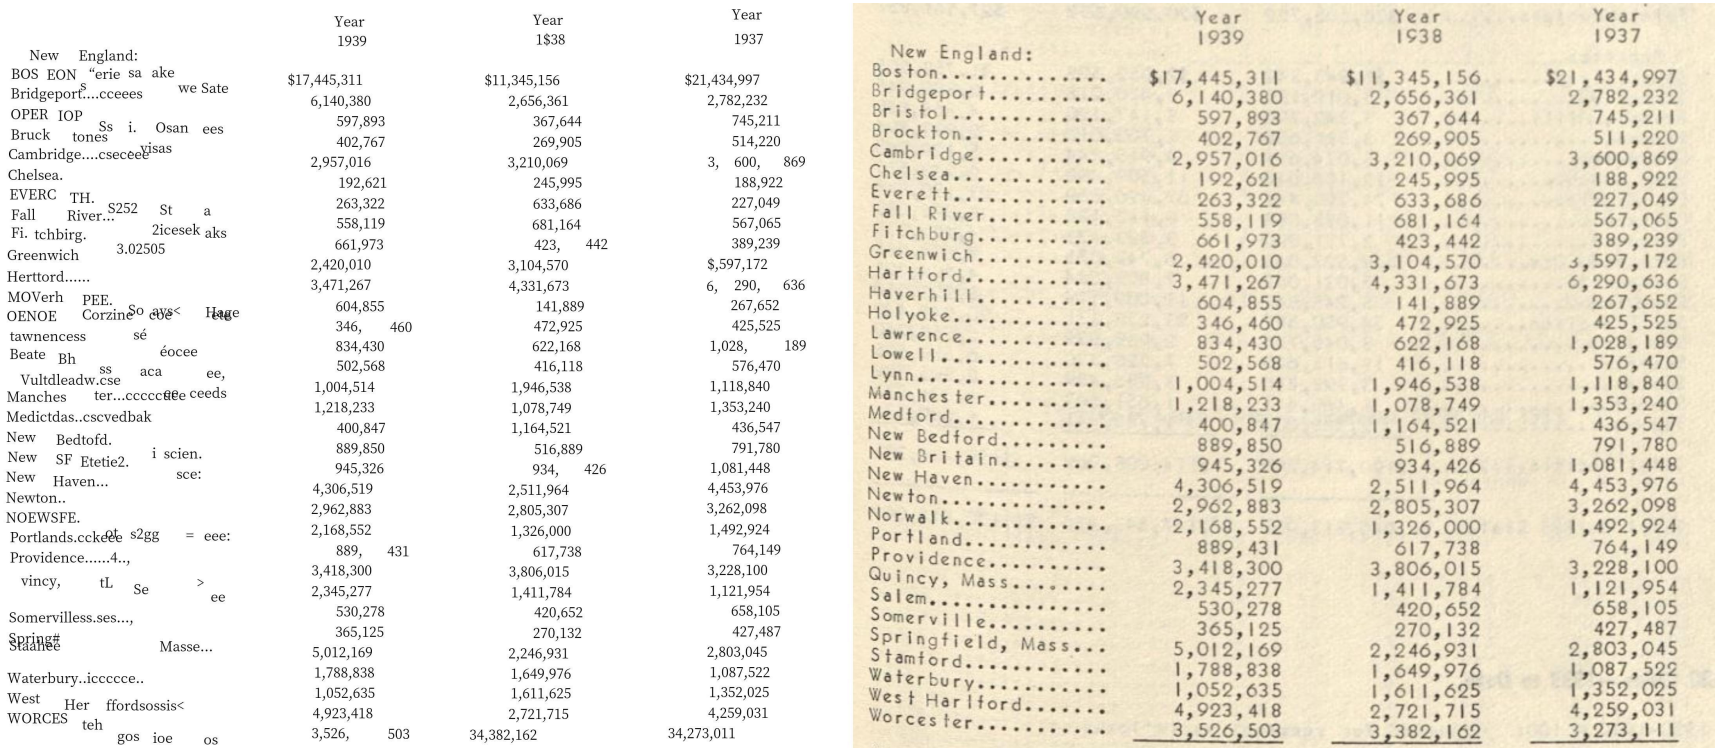

	Year 1939	Year 1$38	Year 1937
New England:			
BOS EON “erie sa ake	$17,445,311	$11,345,156	$21,434,997
Bridgeport....cceees we Sate	6,140,380	2,656,361	2,782,232
OPER IOP	597,893	367,644	745,211
Bruck tones Ss i. Osan ees	402,767	269,905	514,220
Cambridge....csecee visas	2,957,016	3,210,069	3, 600, 869
Chelsea.	192,621	245,995	188,922
EVERC TH.	263,322	633,686	227,049
Fall River... S252 St a	558,119	681,164	567,065
Fi. tchbirg. 2icesek aks	661,973	423, 442	389,239
Greenwich 3.02505	2,420,010	3,104,570	$,597,172
Herttord.....	3,471,267	4,331,673	6, 290, 636
MOVerh PEE.	604,855	141,889	267,652
OENOE Corzine eve So ays< Hage	346, 460	472,925	425,525
tawnencess sé	834,430	622,168	1,028, 189
Beate Bh éocee	502,568	416,118	576,470
Vultdleadw.cse ss aca ee,	1,004,514	1,946,538	1,118,840
Manches ter...ccccceee ceeds	1,218,233	1,078,749	1,353,240
Medictdas..cscvedbak	400,847	1,164,521	436,547
New Bedtofd.	889,850	516,889	791,780
New SF Etetie2. i scien.	945,326	934, 426	1,081,448
New Haven... sce:	4,306,519	2,511,964	4,453,976
Newton..	2,962,883	2,805,307	3,262,098
NOEWSFE.	2,168,552	1,326,000	1,492,924
Portlands.cckeee s2gg = eee:	889, 431	617,738	764,149
Providence.....4..,	3,418,300	3,806,015	3,228,100
vincy, tL Se > ee	2,345,277	1,411,784	1,121,954
Somervilless.ses...,	530,278	420,652	658,105
Spring#	365,125	270,132	427,487
Stamee Masse...	5,012,169	2,246,931	2,803,045
Waterbury..iccccce..	1,788,838	1,649,976	1,087,522
West Her ffordsossis<	1,052,635	1,611,625	1,352,025
WORCES teh	4,923,418	2,721,715	4,259,031
gos ioe os	3,526, 503	34,382,162	34,273,011

	Year 1939	Year 1938	Year 1937
New England:			
Boston.....	$17,445,311	$11,345,156	$21,434,997
Bridgeport.....	6,140,380	2,656,361	2,782,232
Bristol.....	597,893	367,644	745,211
Brockton.....	402,767	269,905	511,220
Cambridge.....	2,957,016	3,210,069	3,600,869
Chelsea.....	192,621	245,995	188,922
Everett.....	263,322	633,686	227,049
Fall River.....	558,119	681,164	567,065
Fitchburg.....	661,973	423,442	389,239
Greenwich.....	2,420,010	3,104,570	3,597,172
Hartford.....	3,471,267	4,331,673	6,290,636
Haverhill.....	604,855	141,889	267,652
Holyoke.....	346,460	472,925	425,525
Lawrence.....	834,430	622,168	1,028,189
Lowell.....	502,568	416,118	576,470
Lynn.....	1,004,514	1,946,538	1,118,840
Manchester.....	1,218,233	1,078,749	1,353,240
Medford.....	400,847	1,164,521	436,547
New Bedford.....	889,850	516,889	791,780
New Britain.....	945,326	934,426	1,081,448
New Haven.....	4,306,519	2,511,964	4,453,976
Newton.....	2,962,883	2,805,307	3,262,098
Norwalk.....	2,168,552	1,326,000	1,492,924
Portland.....	889,431	617,738	764,149
Providence.....	3,418,300	3,806,015	3,228,100
Quincy, Mass.....	2,345,277	1,411,784	1,121,954
Salem.....	530,278	420,652	658,105
Somerville.....	365,125	270,132	427,487
Springfield, Mass...	5,012,169	2,246,931	2,803,045
Stamford.....	1,788,838	1,649,976	1,087,522
Waterbury.....	1,052,635	1,611,625	1,352,025
West Hartford.....	4,923,418	2,721,715	4,259,031
Worcester.....	3,526,503	3,382,162	3,273,111

Notes: Example of the output (left) obtained from applying *Layout Parser* to annual totals reported in the 1939 volume of *Dun's Statistical Review*, juxtaposed against the original scanned table (right). We filter out non-standard characters from this output, score the quality of the output, and then compare the result against output obtained from ABBYY FineReader applied to the same scan.

FIGURE B.7. Comparison of X-13 Filter to FRED/Census Seasonally Adjusted Data

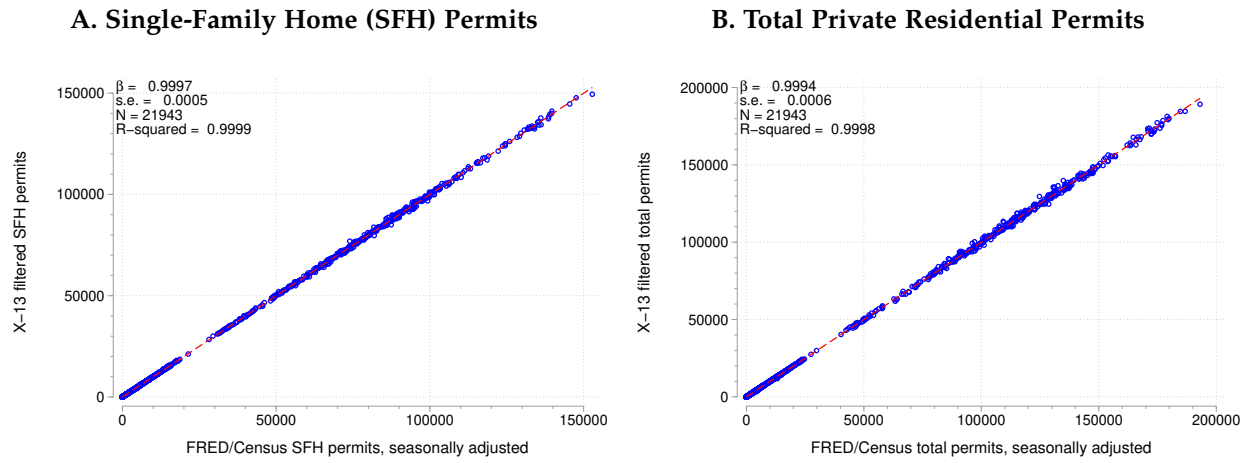

Notes: The figure compares seasonally adjusted Census building permit counts downloaded via the FRED API vs. the raw Census building permit counts that we seasonally adjusted using the X-13ARIMA-SEATS software provided by the Census. Each point in the scatterplot refers to a state-month observation. Panel A performs this exercise for single-family home permits, while Panel B does this for all total private residential permits (i.e. SFH + multifamily permits, excluding any permits tied to public contracts). We implement the X-13 filter on a Linux workstation.

interpolate missing values as zeroes. For low levels of raw unadjusted permits, the X-13 filter will produce negative values, which is not possible in reality. In such cases, we replace negative post-filtered values with zeroes.

B.5 Splicing Together Permit Valuation Series

We must address a data gap between our two primary sources to construct a continuous time series of building permit valuations. *Dun's Review* ceased publishing permit tables after October 1957, while the Census Bureau's Building Permits Survey (BPS) began in May 1959, with state-level information available at monthly frequency starting in May 1960. To bridge this 26-month gap, we digitize tables from the *New York State Construction and Real Estate Census*, which report comprehensive permit valuations covering both single-family and multifamily housing during this period. This source is particularly suitable for bringing the time series gap, as total valuations closely align with both Census and *Dun's Review* figures during overlapping months, suggesting consistent measurement methodology.

Following standard practices in the construction of long-run economic time series (e.g., Shiller, 2015; Jordà et al., 2019), we implement a three-step procedure to create a consistent series. First, we convert all nominal values to 2012 dollars using Shiller's (2015) long-run Consumer Price Index series, which accounts for changes in purchasing power over time. Second, we apply the X-13 ARIMA-SEATS seasonal adjustment filter to each component series, following the Census Bureau's methodology for construction statistics (U.S. Census Bureau, 2022) (see Appendix B.4 for details). Finally, we interpolate backwards using a Vector Autoregression (VAR) model of order one with the New York State data as an input. Figure B.8 plots the version of our one-hundred year time series in Figure 1 reflecting these interpolation procedures. The interpolation procedure effectively bridges the gap without pronounced discontinuities for the aggregate U.S. and housing supply inelastic states — which are more similar to

FIGURE B.8. Total Real Value of U.S. Monthly Building Permits Issued, 1919–2019 (Interpolated)

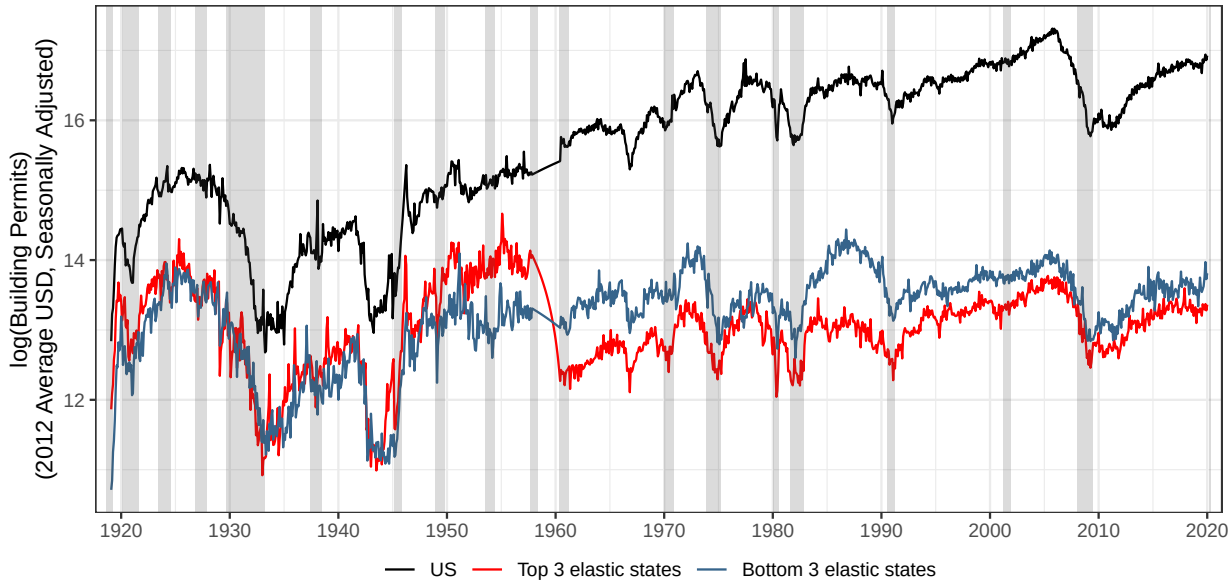

Notes: The figure replicates the century of real building permit valuation data from Figure 1. We interpolate the gap between our main data sources over the period 1957M11 to 1960M4 using the steps outlined above by feeding in data from the *New York State Construction and Real Estate Census* to estimate a VAR(1) model.

New York — but there is still a level shift across the gap for the more dissimilar supply elastic states. We unfortunately lack similar permits data from other state censuses to help mitigate this discontinuity, and therefore, in most of our analysis, we divide our dataset into post-1960s BPS and pre-1960s *Dun's* samples.

B.6 Going between Permit Quantities and Values

Despite our finding that permit quantities and valuations are nearly 100% correlated in the modern Census span of our data, it may be the case that volatility in permit quantities and valuations are more divorced in earlier time periods where housing price indices are not readily available at geographically disaggregated levels. We address these additional concerns by checking robustness of our results to using three sets of alternative series:

1. We use the historical housing price indices produced by Lyons et al. (2026) to extend $v_{s,t}$ further back to 1960 for 30 cities with home sale newspaper listings.

2. We deflate the *Dun's* series using annual state-level construction cost indices collected from volumes of *Building Construction Cost Data*, published by R.S. Means beginning in 1942 to present, to isolate a measure of permit quantities which can be spliced together with the Census $Q_{s,t}$ series. D'Amico et al. (2024) use a subset of the volumes from R.S. Means starting in 1998 to estimate the cost of economy quality homes. Conversely, we can convert Census $Q_{s,t}$ series to $V_{s,t}$ using the construction cost data from 1960 onward to improve the match with the replacement cost values collected in *Dun's*.

3. We digitize the rental price indices available from U.S. Bureau of Labor Statistics (1954) for 1919–1952 and assume a constant cap rate for multifamily structures to create another deflator for 20 cities' valuation series.

C Alternative GARCH Specifications

In this appendix, we compare the convergence properties of the three most common GARCH model specifications in the finance literature. We show that for most samples and time periods in the building permits data, GARCH outperforms GJR and E-GARCH in terms of convergence rates and unique solutions. As a robustness check, we present results using GJR-GARCH as the next-best alternative specification for BPG volatility.

C.1 Stability Tests of GARCH Models

We justify our use of the workhorse GARCH(1,1) specification in equation (C.2) for computing conditional volatility of building permit growth (BPG) by demonstrating that alternative GARCH models often converge to multiple solutions when applied to the building permit data or do not converge at all for some states-level permits series. There are three main classes of GARCH models used in the financial economics literature. For each model, we retain the same mean equation:

$$x_t = \theta_0 + \theta_1 \cdot x_{t-1} + \varepsilon_t, \text{ with } \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2) \tag{C.1}$$

1. **GARCH(1,1)**:

$$(\sigma_t^{GARCH})^2 = \alpha_0 + \alpha_1 \cdot \varepsilon_{t-1}^2 + \alpha_2 \cdot (\sigma_{t-1}^{GARCH})^2 \quad \text{s.t. } \alpha_i > 0; \alpha_1 + \alpha_2 < 1 \tag{C.2}$$

2. **GJR-GARCH**:

$$(\sigma_t^{GJR})^2 = \alpha_0 + \alpha_1 \cdot \varepsilon_{t-1}^2 + \alpha_2 \cdot (\sigma_{t-1}^{GJR})^2 + \gamma \cdot \varepsilon_{t-1}^2 \cdot \mathbb{1}\{\varepsilon_{t-1} < 0\} \text{ s.t. } \alpha_1 + \alpha_2 + \frac{\gamma}{2} < 1 \tag{C.3}$$

3. **E-GARCH**: An advantage to the E-GARCH augmentation is that it can capture a stylized fact about stock returns that the standard GARCH model cannot. Namely that negative shocks at period $t-1$ have a stronger effect on time t variance than positive shocks. Another advantage is that estimating a log-likelihood reduces computational time. There are no parameter restrictions imposed on the E-GARCH specification because the variance is always positive by construction.

$$\ln (\sigma_t^{EXP})^2 = \alpha_0 + \alpha_1 \cdot \left(\frac{\varepsilon_{t-1}}{\sigma_{t-1}^{EXP}} \right) + \alpha_2 \cdot \ln (\sigma_{t-1}^{EXP})^2 + \gamma \cdot \left(\left| \frac{\varepsilon_{t-1}}{\sigma_{t-1}^{EXP}} \right| - \sqrt{\frac{2}{\pi}} \right) \tag{C.4}$$

We estimate the three conditional volatility models for U.S. aggregate and state-level permit series and separately for single-family homes (SFH) and total private residential permits (TOT). We perform two versions of the same simulation exercise in which we vary the constraint on the starting values for the optimizer routine and the parameter domain on E-GARCH. In both versions, we adopt the textbook optimization constraints listed above for GARCH(1,1) and GJR-GARCH. Each simulation takes 10,000 draws from the feasible set of starting values with replacement and then runs variance targeting on the demeaned permit growth time series using the `basinhopping` routine in Python.[6]

[6]Documentation for `basinhopping` can be found here: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.basinhopping.html.

FIGURE C.1. Convergence of Coefficients for Conditional Volatility Models

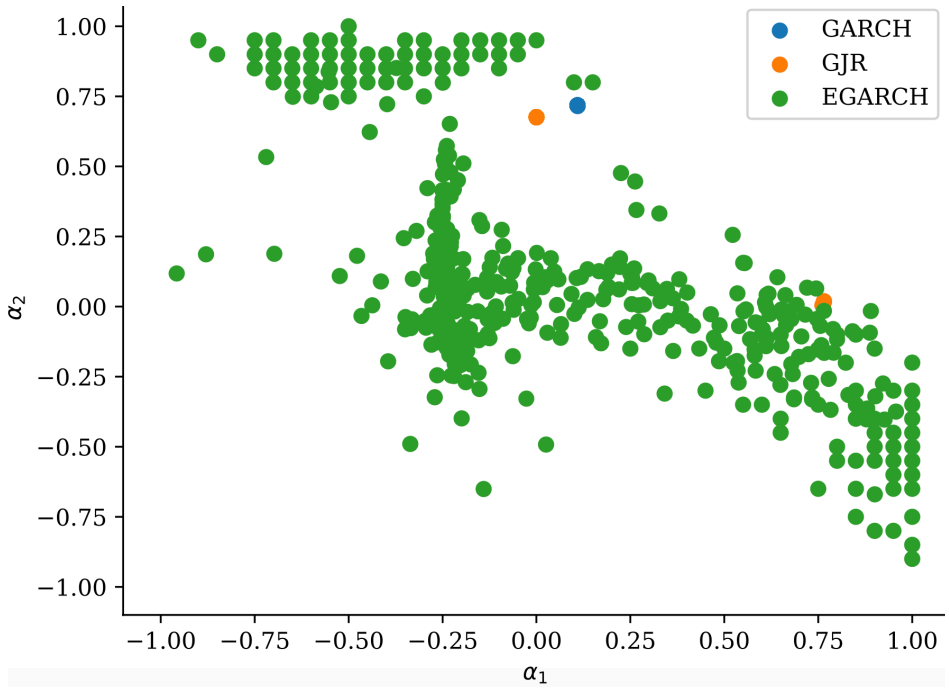

Notes: We illustrate the results of simulation version 1 by plotting the coefficients of convergence for the three main types of GARCH models considered in the literature. We take 500 random draws with replacement of starting parameter values of (α_1, α_2) on the interval $\alpha_i \in (-1, 1)$ for E-GARCH and $\alpha_i \in (0, 1)$ for the other two models. For each draw, we estimate via quasi-maximum likelihood estimation (QMLE) the GARCH(1,1) [blue] specification given by equation (C.2), the GJR-GARCH [orange] model from equation (C.3), and the E-GARCH [green] model from equation (C.4). We estimate each model for building permit growth (BPG) volatility, where we compute aggregate U.S. seasonally adjusted single-family home building permit growth covering the modern Census Building Permits Survey (1988M1 to 2019M12).

Simulation Version 1

- **GARCH specifications:**
  - Optimization constraint: $\alpha_1 + \alpha_2 < 1$
  - Starting values constraint: select two random non-negative values satisfying $\alpha_1 + \alpha_2 = 0.9$
  - Parameter domain: $\alpha_0 > 0; 0 < \alpha_1 < 1; 0 < \alpha_2 < 1$
- **GJR-GARCH specifications:**
  - Optimization constraint: $\alpha_1 + \alpha_2 + \gamma/2 < 1$
  - Starting values constraint: select three random non-negative values satisfying $\alpha_1 + \alpha_2 + \gamma = 0.9$
  - Parameter domain: $\alpha_0 > 0; 0 < \alpha_1 < 1; 0 < \alpha_2 < 1; 0 < \gamma < 1$
- **E-GARCH specifications:**
  - Optimization constraint: none
  - Starting values constraint: select three random values satisfying $\alpha_1 + \alpha_2 + \gamma = 0.9$
  - Parameter domain: $\alpha_0 \in \mathbb{R}; -1 < \alpha_1 < 0.9; -1 < \alpha_2 < 0.9; -1 < \gamma < 0.9$

SIMULATION VERSION 2

- **GARCH specifications:**
  - Optimization constraint: $\alpha_1 + \alpha_2 < 1$
  - Starting values constraint: select two random non-negative values satisfying $\alpha_1 + \alpha_2 = 0.999$
  - Parameter domain: $\alpha_0 > 0; 0 < \alpha_1 < 1; 0 < \alpha_2 < 1$
- **GJR-GARCH specifications:**
  - Optimization constraint: $\alpha_1 + \alpha_2 + \gamma/2 < 1$
  - Starting values constraint: select three random non-negative values satisfying $\alpha_1 + \alpha_2 + \gamma = 0.999$
  - Parameter domain: $\alpha_0 > 0; 0 < \alpha_1 < 1; 0 < \alpha_2 < 1; 0 < \gamma < 1$
- **E-GARCH specifications:**
  - Optimization constraint: none
  - Starting values constraint: select three random non-negative values satisfying $\alpha_1 + \alpha_2 + \gamma = 0.999$
  - Parameter domain: $\alpha_0 \in \mathbb{R}; 0 < \alpha_1 < 1; 0 < \alpha_2 < 1; 0 < \gamma < 1$

There are two key differences between the simulation versions. In version 1, we select starting values away from the boundaries of the parameter domain. In version 2, we allow for starting values at the boundary of the parameter domain but additionally require that the E-GARCH coefficients be strictly positive so that the parameter domain matches the domain for GARCH and GJR-GARCH.

Table C.1 summarizes how each of the GARCH models performs in terms of convergence rates (CR) and stability (i.e., number N of unique solutions up to four decimal places). We define two solutions to be identical if the estimated parameters of the GARCH model are the same up to five decimal places. In Panel A, we present results from simulation version 2 which imposes textbook optimization constraints and the same parameter domains across the models. We find that for the modern period GARCH outperforms GJR-GARCH in terms of the multiplicity problem. For the two non-unique solutions GARCH delivers for the TOT series, the parameter estimates only differ at the fourth decimal place. In Panel B, we compare simulation versions 1 and 2 and use the more recent years of data for which we can deflate permit quantities according to housing price indices. Convergence rates are broadly similar and close to 100% across all three GARCH types in version 2.

In cases of instability where $N > 1$, we assess how much two GARCH models deviate from each other on average. To do this, we compute all convergent conditional volatility series for each method. We then compute all pairwise correlations between the conditional volatility series estimated via two different specification (i.e., between GARCH and GJR-GARCH, between GARCH and E-GARCH, and between GJR and E-GARCH). Finally, we average across the pairwise correlations and report the results in Table C.2. Across different combinations of series (SFH for single-family homes vs. TOT for total residential permits) and GARCH models, instances of multiple solutions still result in conditional volatility estimates which are highly

TABLE C.1. Convergence and Parameter Stability across GARCH Models of BPG Volatility

A. Single-Family Homes vs. Total Private Residential Permits: Simulation Version 2

	Single-Family Homes Permits				*Total Private Residential Permits*			
	Convergence Rate	N. Unique Solutions	Convergence Rate	N. Unique Solutions	Convergence Rate	N. Unique Solutions	Convergence Rate	N. Unique Solutions
GARCH	0.9876	44	0.9984	4	0.9984	2	0.9999	2
GJR-GARCH	0.9457	7	0.9986	14	0.9976	5	0.9996	3
E-GARCH	0.9974	11	0.9998	7	0.9992	6	1	1
Sample	1960–2019	1960–2019	1980–2019	1980–2019	1960–2019	1960–2019	1980–2019	1980–2019

B. Comparing Simulation Version Results in the Post-2000s Period

	U.S. Building Permits: $P \times Q$			
	Simulation Version 1		Simulation Version 2	
	Convergence Rate	N. Unique Solutions	Convergence Rate	N. Unique Solutions
GARCH	0.9999	4	0.9999	4
GJR	0.9997	20	1	16
E-GARCH	0.3907	3859	0.9979	4
Sample	2000–2023	2000–2023	2000–2023	2000–2023

Notes: We estimate GARCH, GJR-GARCH, and E-GARCH models for each of the two aggregate U.S. building permits series: (*i*) Single-Family Homes (SFH); and (*ii*) Total Private Residential Permits. We consider the samples over the full Census period of 1960–2019 and the more recent sample from 1980–2019. For each model and series, we report two items: (*i*) the convergence rate, defined as the fraction of starting parameter draws for which the optimization routine converges to a solution; and (*ii*) the number of unique solutions obtained from each GARCH model, where uniqueness is defined up to five decimal places. See text for simulation details.

correlated, on average. Hence, despite any instability, our results on the predictability of BPG volatility are broadly robust to the choice of GARCH specification.

Importantly, we conduct our simulation exercises in using a segment of the data from the modern version of the Building Permits Survey, which consists of monthly building permit reports already digitized by the Census and made available to the public by download. This means that our finding of multiple solutions for GJR and E-GARCH models as applied to building permits cannot be due to any measurement error arising due to either our digitization methods—although, as discussed in Appendix B.3, we have multiple quality control measures in place—or due to low-quality versions of any extant scans of copies of the physical reports. Nonetheless, the instability of GJR and E-GARCH holds for different sample cuts, including when we fit each GARCH model to our U.S. aggregate BPG series covering the Census time period and when we estimate the GARCH models for each U.S. state.

C.2 Main Predictability Results using GJR-GARCH

We replicate the main analysis from Section 4.1 for the longitudinal sample of permits from the Census Building Permit Survey. Table C.3 shows that the loading on aggregate BPG volatility is quantitatively similar regardless of the choice of GARCH model, across various sets of controls for macroeconomic conditions. This is perhaps unsurprising given the findings from our simulation exercise that GARCH(1,1) and GJR-GARCH deliver highly correlated conditional volatility series when applied to nationwide permit series, even at the extremes

TABLE C.2. Average Pairwise Correlations between GARCH Conditional Volatility Estimates

Series	Sample Period	$\text{Corr}(\sigma_{\text{GAR}}, \sigma_{\text{GJR}})$	$\text{Corr}(\sigma_{\text{GAR}}, \sigma_{\text{EGR}})$	$\text{Corr}(\sigma_{\text{GJR}}, \sigma_{\text{EGR}})$
SFH Permits	1960–2019	0.8115	0.9538	0.8282
SFH Permits	1980–2019	0.8899	0.9754	0.8829
Total Permits	1960–2019	0.8590	0.6854	0.5439
Total Permits	1980–2019	0.9162	0.7866	0.6840

Notes: We estimate GARCH, GJR-GARCH, and E-GARCH models for each of the four aggregate U.S. building permits series over the full Census sample period of 1960–2019 and using the more recent sample from 1980–2019. For each unique solution obtained from each GARCH model, where uniqueness is defined up to five decimal places, we then report the average pairwise correlations across solutions between two models. See text for simulation details.

of the possible solution set. While the aggregate loading on BPG volatility is of a similar magnitude to the coefficients in Table 2, the BPG volatility estimated via GJR-GARCH has more explanatory power for both equities and corporate bonds. The incremental R^2 of single-family home volatility for equities in Table C.3 is 9.5% (Panel A, column 1) compared to 3.1% for GARCH-implied volatility; for bonds and single-family permits, the incremental R^2 is 5.7 p.p. greater for GJR-GARCH than for GARCH(1,1).

The cross-sectional results display a similar pattern when we compare GJR-GARCH to GARCH(1,1). Figure C.2 reproduces Figure 4 in which we regress equity and bond return volatility on BPG volatility for each state. While the ordinal ranking of states according to the magnitude of their loadings on BPG volatility is similar for GJR-GARCH vs. GARCH(1,1), the predictability is almost always greater and more precisely estimated for GJR-GARCH. For example, Florida single-family permits the elasticity on stock return volatility is significant at the 1% level and 30% greater when measured according to GJR-GARCH, compared to GARCH(1,1). For corporate bond return volatility, there are fewer differences in the state-by-state point estimates across the two GARCH models. Overall, we find that the use of the workhorse GARCH(1,1) specification in our main results results in conservative estimates of the predictability of BPG volatility for asset returns.

C.3 Tests of Normality of BPG Distributions

The GJR-GARCH specification was originally developed to account for the extreme skewness in the distribution of risky asset returns. Therefore, it is possible that for some permits series featuring fat tails in the permit growth (BPG) distribution, GJR-GARCH is the more appropriate model. We conduct separate tests for normality separately for each state and across different time periods and in the pooled panel of states in our permits database. We calculate the test statistic of D'Agostino et al. (1990) with the correction proposed by Royston (1992).[7]

Figure C.3 plots the distribution of monthly BPG for the entire U.S. (left-hand panels) and in the cross-section of states (right-hand panels). There are clear differences in the normality of BPG in aggregate vs. across states and for the pre vs. post-1960s period. In Panel A—regardless of whether we include New York State, which has outsize influence on national building permits during the prewar period—we can reject the null of a normal distribution on the skewness dimension for aggregate U.S. BPG in the *Dun's* data covering 1919 to 1957. Panel B shows that

[7] This corresponds to the default options for the `sktest` command in Stata.

TABLE C.3. Regressions of Asset Return Volatility on Nationwide U.S. BPG Volatility (GJR-GARCH)

A. Total Private Residential Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.088***	0.027**	0.026**	0.025**	0.064**	0.070***	0.036***	0.035***	0.033***	0.016***
	(2.82)	(2.45)	(2.47)	(2.39)	(2.57)	(4.68)	(3.76)	(3.40)	(3.18)	(3.77)
Time sample	1960-19	1960-19	1980-19	1980-16	2000-16	1960-19	1960-19	1980-19	1980-16	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$Leverage_{t-p}$			✓	✓	✓			✓	✓	✓
$DSCR_{t-p}$			✓	✓	✓			✓	✓	✓
$IPGrowth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
N	714	707	479	435	195	714	707	479	435	195
R^2	0.109	0.471	0.463	0.471	0.605	0.185	0.367	0.452	0.444	0.544

B. Single-Family Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.074***	0.024**	0.022**	0.022**	0.049**	0.076***	0.044***	0.040***	0.038***	0.015***
	(2.60)	(2.40)	(2.49)	(2.41)	(2.18)	(6.07)	(4.48)	(4.54)	(4.28)	(3.99)
Time sample	1960-19	1960-19	1980-19	1980-16	2000-16	1960-19	1960-19	1980-19	1980-16	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓		✓	✓	✓	✓
$PopGrowth_{t-p}$		✓	✓	✓	✓		✓	✓	✓	✓
$Leverage_{t-p}$			✓	✓	✓			✓	✓	✓
$DSCR_{t-p}$			✓	✓	✓			✓	✓	✓
$IPGrowth_{t-p}$			✓	✓	✓			✓	✓	✓
$DisasterNVIX_{t-p}$				✓	✓				✓	✓
N	714	707	479	435	195	714	707	479	435	195
R^2	0.095	0.470	0.462	0.471	0.599	0.258	0.391	0.471	0.463	0.543

Notes: The table presents estimates from equation (3.3) relating total return volatility to lagged building permit growth (BPG) volatility. The difference relative to Table 2 is that we produce conditional BPG volatility using the GJR-GARCH model in (C.3). In Panel A we use total private residential permits as the quantity measure $Q_{s,t}$, but instead use permits attached to single-family units in Panel B. We include in most specifications a set of controls for other macroeconomic observables directly related to BPG volatility, which might also drive aggregate financial market volatility. *PopGrowth* refers to the annual population growth rate from the Census. *Leverage* is the aggregate corporate leverage ratio based on firms' annual filings in COMPUSTAT, computed as the sum of long-term debts (*dltt*) and debts in current liabilities (*dlc*), divided by total stockholders' equity (*seq*). *DSCR* is the quarterly household debt service coverage ratio from the Federal Reserve, defined as household debt service payments as a fraction of disposable income. *IPGrowth* is the month-on-month growth rate in the industrial production index (INDPRO). In some specifications, we add the natural disaster component of the News Implied Volatility Index (NVIX) of Manela and Moreira (2017). We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. The time sample varies depending on the data availability of covariates, with *DSCR* available starting in 1980, and the NVIX available only up to 2016. *t*-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

FIGURE C.2. Predictability of State-Level Permit Growth Volatility (GJR-GARCH, 1960–2019)

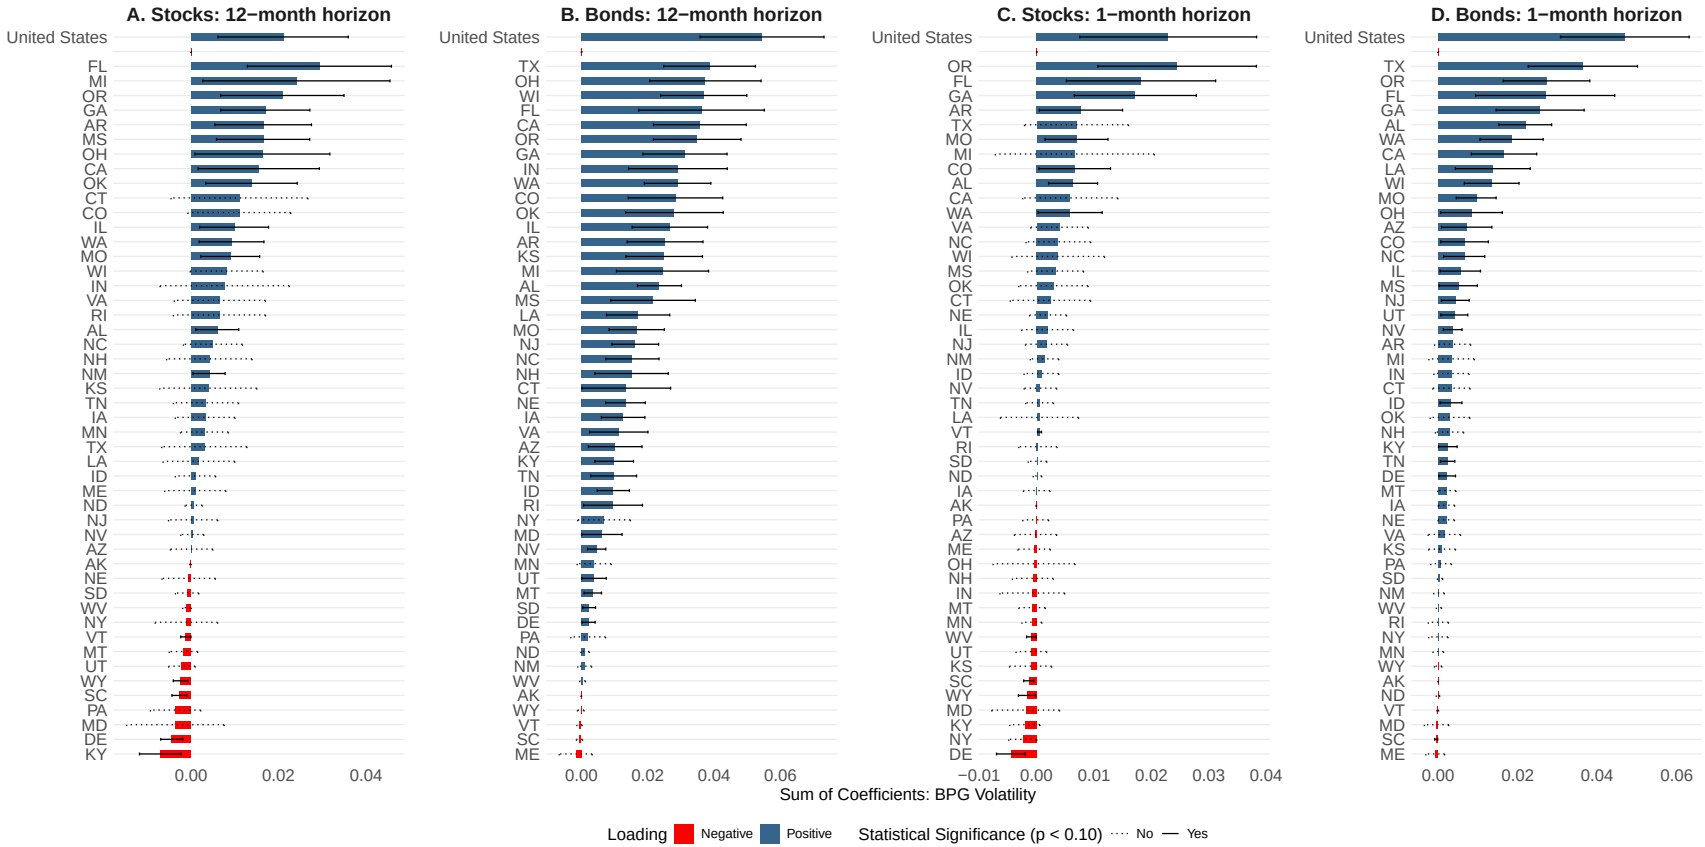

Notes: We estimate equation (3.3) for the entire U.S. and equation (3.4) separately for each state s using (C.3) as our conditional volatility model to obtain σ^{BPG}. Each regression includes monthly dummies, the BIC-optimal autocorrelation term σ_{t-1}, but no other controls. We omit the regressions for Washington, D.C. and Massachusetts, as the GARCH model does not converge for those series. In panels A and B, we use $\tau^* = 12$ lags of $\sigma^{BPG}_{s,t}$ and report the sum of the estimated lagged coefficients $\sum_{\tau=1}^{\tau^*} \widehat{\beta}_{s,\tau}$. In panels C and D, we show predictability at the one-month horizon. Our sample period covers the Census Building Permit Survey from its start in 1960M5 to 2019M12. We use permits attached to single-family units as the quantity measure $Q_{s,t}$. Panels A and C perform this exercise for CRSP stock return volatility as the outcome, and total return volatility of the Dow Jones Corporate Bond Index in Panels B and D. We plot 90% confidence intervals obtained via autocorrelation and heteroskedasticity-robust standard errors, selecting the minimum lag order such that the estimator for the covariance matrix is consistent (Newey and West, 1987). Solid confidence bars indicate significance at the 90% level, while dashed bars indicate statistical insignificance.

we can only reject the null of no skewness relative to a normal distribution at the 5% significance level, for both the population-weighted and unweighted permits series in *Dun's*, indicating that building permit growth in some states is more normally distributed than it is nationwide.

In contrast, in the modern Census sample of our data, aggregate U.S. BPG is almost perfectly symmetric, with a p-value on the skewness test of 0.91 for single-family home permits. Mean BPG is also much lower in the modern period, with zero average monthly growth in quantities post-1960s compared to an average monthly growth rate of 0.6% in the pre-1960s sample based on the valuations surveyed in *Dun's*. However, despite these differences, due to the fat tails in the BPG distributions, the joint Chi-squared test of normality combining the test statistics for skewness and kurtosis always rejects the null of a normal distribution across all time series and periods at the 1% significance level.

Table C.4 separately reports p-values from skewness tests for each sample time period and each state. There is a clear positive correlation between the extent to which BPG volatility implied by the GJR-GARCH model (as shown in Figure C.2 for single-family home permits) is strongly predictive of asset return volatility and the degree of skewness in the BPG distribution. This points to the appropriateness of the GJR specification for conditional BPG volatility in less population-dense parts of the country exhibiting very large swings in permitting activity. The results of this exercise also explain why for such states the magnitude and significance of the loadings on σ_{t-p}^{BPG} differ depending on the GARCH specification we use.

D Additional Results

We highlight some additional results in this appendix, including break dates implied by Bai-Perron tests in the building permit series, adding controls for local economic conditions, aggregate commodity risk, and household leverage, and principal component analysis over longer time periods.

D.1 Break Dates in the Geographic Cross-Section

We formally test for structural breaks in the seasonally adjusted state-level permit series and GARCH-implied building permit growth volatility using the "one break at a time" sup Wald statistics of Bai and Perron (1998). We adopt the standard symmetric trimming percentage of 15%, such that for each series we search for breaks over 1969M5 to 2014M1. This trimming percentage allows us to identify possible breaks around key economic episodes in the modern Census period, including the oil crises of the 1970s and stagflation and the Great Recession. We then search over this trimmed sample period for a maximum of seven breaks. The maximum number of breaks we identify in any state's permit levels or BPG volatility series is five. Only one state, Iowa, features no breaks in its building permit series, and two is the modal number of breaks.

Figure D.1 plots at each month over the trimmed sample period the number of states with breaks in their GARCH BPG volatility (left panels) or seasonally adjusted building permits (right panels). The breaks are concentrated around recession episodes and in the early 1970s around spikes in oil prices. For both volatility and permit levels, and regardless of whether we consider single-family permits or all residential permits, we uncover a large mass of states with breaks on the eve of the Great Recession. Level breaks are more prevalent than volatility breaks, and breaks are slightly less common in single-family permits (118 total breaks) than for

FIGURE C.3. Pooled Distributions and Skewness Tests for Building Permit Growth

A. U.S. BPG in *Dun's Review*

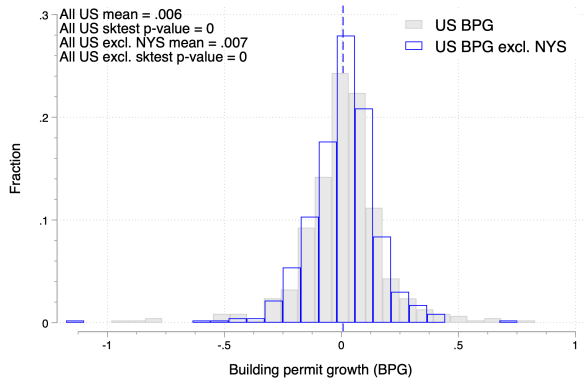

B. Cross-Sectional BPG in *Dun's Review*

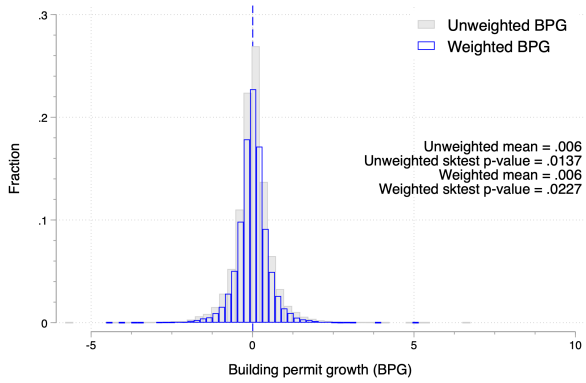

C. U.S. BPG in Census BPS

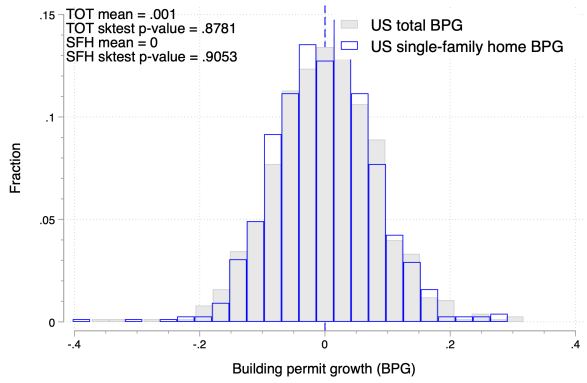

D. Cross-Sectional BPG in Census BPS

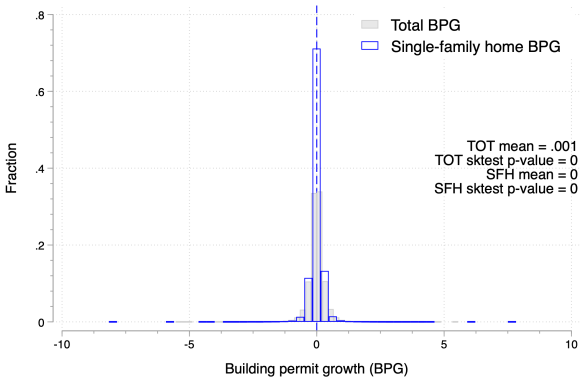

Notes: The figure plots the distribution of month-on-month building permit growth (BPG) for the *Dun's Review* data in Panels A and B and for the Census Building Permits Survey (BPS) in Panels C and D. The *Dun's Review* data cover the period 1919M1 to 1957M10, while the Census BPS data cover 1960M5 to 2019M12. In the left-hand panels we plot the distribution of monthly BPG computed using total nationwide permits. In the *Dun's* sample we report separate means and skewness test p-values for the entire U.S. and the entire U.S. excluding permits from New York State to examine the influence of New York City. In the Census BPS sample, we report separate statistics for the total residential permits (TOT) vs. single-family homes (SFH). The right-hand panels instead plot the cross-sectional distribution of monthly BPG, which includes observations from each state's permits time series. For the *Dun's* sample, we report separate statistics for population unweighted and weighted observations, while for the Census BPS sample we again distinguish between TOT and SFH permit counts in computing BPG.

TABLE C.4. Skewness Tests for Building Permit Growth by State and Subsample

Data Source:	*Dun's Review* (1919–1957)		*U.S. Census BPS* (1960–2019)	
Subsample:	Unweighted	Weighted	Total	Single-Family
United States	0.000***	–	0.878	0.905
Alabama	0.000***	0.000***	0.443	0.023**
Alaska	–	–	0.000***	0.560
Arizona	0.481	–	0.010***	0.571
Arkansas	0.929	–	0.833	0.261
California	0.000***	0.000***	0.032**	0.653
Colorado	0.000***	0.000***	0.000***	0.819
Connecticut	0.138	0.138	0.036**	0.429
Delaware	0.600	–	0.008***	0.002***
Florida	0.126	0.126	0.005***	0.132
Georgia	0.063*	0.063*	0.127	0.043**
Hawaii	–	–	0.340	0.743
Idaho	0.697	–	0.000***	0.000***
Illinois	0.058*	0.058*	0.133	0.050**
Indiana	0.216	0.216	0.366	0.389
Iowa	0.096*	0.096*	0.000***	0.163
Kansas	0.003***	0.003***	0.047**	0.000***
Kentucky	0.095*	0.095*	0.104	0.567
Louisiana	0.000***	0.000***	0.863	0.025**
Maine	0.219	–	0.001***	0.589
Maryland	0.216	0.216	0.177	0.953
Massachusetts	0.079*	0.079*	0.028**	0.960
Michigan	0.284	0.284	0.082*	0.153
Minnesota	0.008***	0.008***	0.009***	0.006***
Mississippi	–	–	0.182	0.000***
Missouri	0.410	0.410	0.023**	0.010***
Montana	0.000***	–	0.513	0.129
Nebraska	0.396	0.396	0.144	0.784
Nevada	–	–	0.081*	0.020**
New Hampshire	–	–	0.683	0.011**
New Jersey	0.323	0.323	0.152	0.000***
New Mexico	0.050**	–	0.949	0.350
New York	0.000***	0.000***	0.000***	0.003***
North Carolina	0.078*	–	0.623	0.075**
North Dakota	0.953	–	0.044**	0.001***
Ohio	0.000***	0.000***	0.209	0.444
Oklahoma	0.067*	0.067*	0.793	0.442
Oregon	0.016**	0.016**	0.696	0.191
Pennsylvania	0.962	0.962	0.219	0.000***
Rhode Island	0.401	0.401	0.000***	0.467
South Carolina	0.023**	–	0.000***	0.000***
South Dakota	0.017**	–	0.015**	0.001***
Tennessee	0.065*	0.065*	0.844	0.016**
Texas	0.167	0.167	0.378	0.372
Utah	0.028**	0.028**	0.001***	0.007***
Vermont	0.089*	–	0.127	0.467
Virginia	0.011**	0.011**	0.720	0.019**
Washington	0.031**	0.031**	0.842	0.054**
Washington, D.C.	–	–	0.974	0.928
West Virginia	0.379	–	0.419	0.004***
Wisconsin	0.025**	0.025**	0.162	0.354
Wyoming	–	–	0.006***	0.560

Notes: The table reports the p-value for D'Agostino et al. (1990) tests applied to the distribution of month-on-month building permit growth (BPG) observations within each state's time series for different subsamples. Within the *Dun's Review* sample, unweighted refers to monthly BPG based on raw X-13 seasonally adjusted permit valuations, while the weighted series proportionally weights cities surveyed within each state based on their annual population to create a state-level series. Within the Census Building Permits Survey (BPS) sample, we report p-values for the skewness test applied to total residential permit (TOT) and single-family home (SFH) permit counts. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

FIGURE D.1. Bai-Perron Break Date Tests for State-Level Building Permits

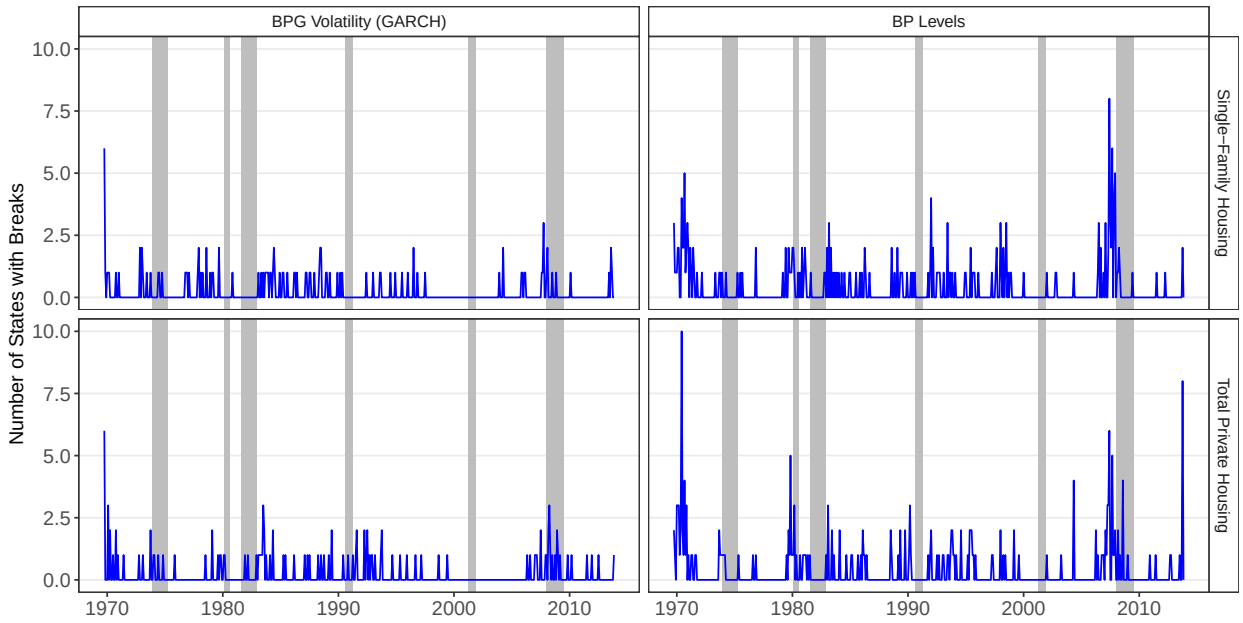

Notes: We conduct Bai and Perron (1998) break date tests for each building permit series using the conditional volatility of building permit growth computed via equation (2.3) in the left-hand panels, and levels in the right-hand panels. We start with the full sample from 1960M5 to 2022M12 and then trim the sample by 15% on either side to limit the dependence of the test on the start and end values of the time series. We then search over the trimmed sample period for a maximum of seven breaks, searching for one break at a time according to the default test procedures in the `strucchange` R package. Grey-shaded areas indicate NBER-dated recessions.

the series including all residential permits (125 total breaks), but breaks are more prevalent for single-family permits around the Great Recession, consistent with the results shown in Section 4.3 in which we extract the subprime factor as the first principal component. Overall, these tests complement our evidence in the main text that building permits in the geographic cross-section are leading indicators for real economic and financial volatility.

D.2 Principal Components of BPG over the Longrun

We conduct the same exercise as done in Figure 7 for the Great Recession episode, but instead using the extended sample covering the full Census BPS period. That is, we run principal components analysis (PCA) on the panel of $\sigma_{s,t}^{BPG}$ for the top 20 U.S. states and plot the principal components (PCs) with an eigenvalue greater than one.[8] The first PC explains 20% and 31% of the variation in BPG volatility for total private residential permits and single-family homes, respectively. However, unlike around the Great Recession period, when we conduct PCA over the full period starting in the 1960s, the first PC is dominated by input supply shocks rather than the subprime mortgage factor. For both total and SFH permits, largest spikes in the first PC correspond to the onset of the oil crises of the 1970s, including the first few months of 1973 (diplomatic fallout before the Yom Kippur War), the middle of 1975 when unemployment peaked during the stagflation era, and the summer of 1980 (the Iran-Iraq War).

[8]Here we select the top 20 states by their 1960 Census population, although which states we include in the sample ultimately does not matter for the general time series patterns in the principal components.

Cortes et al. (2024) show that the "war puzzle" identified by Schwert (1989) of lower stock volatility during wartime and conflict periods coincides with build ups in defense spending, which renders cash flows easier to forecast. On top of regressions in our main analysis in which we directly control for aggregate and state-level leverage, these longitudinal PCA results belie the argument that the predictability of BPG volatility originates from build ups in credit used to finance risky investments with deferred payoffs, such as real estate development.

D.3 Predictability of CRSP Dividends

We replicate our main analysis in Section 4.1 using CRSP dividend volatility as our main outcome variable. We construct CRSP dividend volatility as the monthly volatility of the sesaonally adjusted daily differences between the CRSP total return index (`vwretd`) and the CRSP ex dividend total return index (`vwretx`), keeping the same all other aspects of our research design. Table D.1 shows that aggregate U.S. BPG volatility is a strong predictor of dividends over various periods and conditional on a set of controls for macroeconomic conditions such as population growth, leverage, household debt service, industrial production growth, and the war component of the NVIX of Manela and Moreira (2017). We focus on a the war component of the NVIX as a possible confounding factor leading to spurious correlation between BPG volatility and dividend volatility given that large spikes in dividends over the post-1960s period correspond to defense build ups.

Notably, the predictive power of building permits declines in the post-2000s period around the Global Financial Crisis, with the loading on σ_{t-1}^{BPG} declining by almost one-half when we compare the two specifications with the full set of controls (columns 6 vs. 8). The post-2000s period corresponds to a 50% decline in dividend volatility over 2000-2019 relative to 1960–2000. Together with the evidence in Table 2, BPG volatility is a stronger predictor of stock volatility during cycles when the price component of returns accounts for the bulk of stock volatility. However, consistent with our Grossman–Stiglitz modeling framework in Section 6, the BPG volatility works well at forecasting dividend volatility, particularly during times when the cash flow risk component of returns dominates. Since dividend payouts are highly seasonal, much of the R^2 in Table D.1 comes from the time dummies. The incremental R^2 of σ_{t-1}^{BPG} generated by moving from a specification with only monthly dummies is 5.7% for single-family permits, three times greater than the incremental R^2 of 1.2% we obtain for BPG volatility on total stock return volatility over the full post-1960s period.

D.4 Controlling for Aggregate Commodity Risk and Household Leverage

We extend our baseline results from Table 2 in Section 4.1 by showing that aggregate building permit growth (BPG) volatility predicts asset return volatility even conditional on commodity price risk and changes in household leverage. To do so, we compute the conditional volatility of excess returns on an equally weighted commodity futures index constructed by Janardanan et al. (2024). Commodity risk is a potentially important omitted variable because it may simultaneously influence developers' decisions to engage in home building and firms' input costs, subsuming the predictive power of our BPG factor. Indeed, Figure D.2 demonstrates that the first principal component of BPG volatility for single-family homes is 41% correlated with commodity excess return volatility (20% if using GJR-GARCH), with both series spiking during the 1970s oil shocks.

Although our baseline regressions in Table 2 control for corporate leverage ratios and debt service coverage ratios, one concern — especially in the years leading up to the Great Recession

FIGURE D.2. Principal Components of BPG Volatility from Census Data (1961–2019)

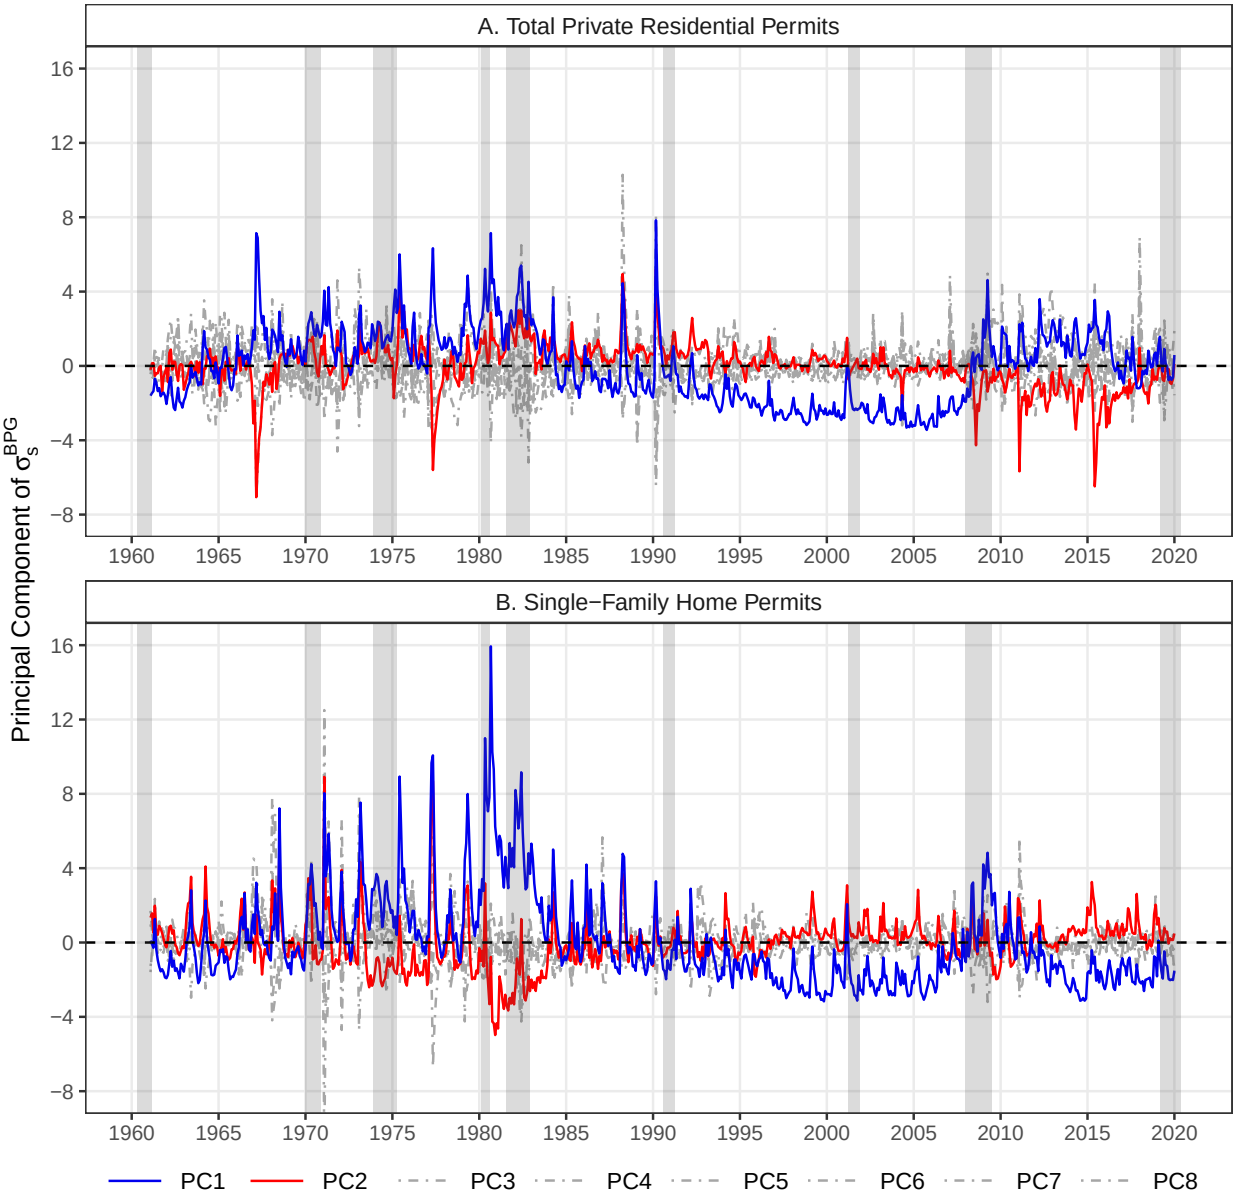

Notes: The figure plots the time series of principal components of state-level total private residential (Panel A) and single-family home (Panel B) monthly building permit growth volatility with an eigenvalue greater than unity. Monthly building permit growth volatility is defined by $\sigma_{s,t}^{BPG}$ in equation (3.2). Our sample in both panels includes the top 20 states ranked by 1960 decennial Census population. We conduct PCA over the time period 1961–2019, excluding the post-COVID-19 period to ensure convergence of the GARCH models. Grey-shaded areas indicate NBER-dated recessions.

TABLE D.1. Regressions of CRSP Dividend Volatility on Nationwide U.S. BPG Volatility

A. Total Private Residential Units BPG Volatility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
σ_{t-1}^{BPG}	0.0016***	0.0014***	0.0012***	0.0007***	0.0014***	0.0007***	0.0005**	0.0004*
	(6.51)	(6.08)	(5.18)	(3.95)	(5.60)	(3.74)	(2.10)	(1.91)
Time sample	1960-19	1960-19	1960-19	1980-19	1960-19	1980-16	2000-19	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓	✓	✓	✓
$PopGrowth_{t-p}$			✓	✓		✓		✓
$Leverage_{t-p}$			✓	✓		✓		✓
$DSCR_{t-p}$				✓		✓		✓
$IPGrowth_{t-p}$				✓		✓		✓
$WarNVIX_{t-p}$					✓	✓		✓
N	714	714	707	479	670	435	239	195
R^2	0.374	0.378	0.460	0.496	0.395	0.496	0.191	0.238

B. Single-Family Units BPG Volatility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
σ_{t-1}^{BPG}	0.0015***	0.0014***	0.0010***	0.0007***	0.0013***	0.0007***	0.0005**	0.0004
	(6.85)	(6.57)	(4.70)	(4.08)	(6.23)	(3.89)	(1.98)	(1.36)
Time sample	1960-19	1960-19	1960-19	1980-19	1960-19	1980-16	2000-19	2000-16
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.		✓	✓	✓	✓	✓	✓	✓
$PopGrowth_{t-p}$			✓	✓		✓		✓
$Leverage_{t-p}$			✓	✓		✓		✓
$DSCR_{t-p}$				✓		✓		✓
$IPGrowth_{t-p}$				✓		✓		✓
$WarNVIX_{t-p}$					✓	✓		✓
N	714	714	707	479	670	435	239	195
R^2	0.378	0.381	0.456	0.496	0.398	0.496	0.188	0.239

Notes: The table presents estimates from equation (3.3) relating CRSP dividend volatility (computed as the volatility of `vwretd` − `vwretx`) to lagged building permit growth (BPG) volatility. In Panel A we use total private residential permits as the quantity measure $Q_{s,t}$, but instead use permits attached to single-family units in Panel B. We include in most specifications a set of controls for other macroeconomic observables directly related to BPG volatility, which might also drive aggregate financial market volatility. *PopGrowth* refers to the annual population growth rate from the Census. *Leverage* is the aggregate corporate leverage ratio based on firms' annual filings in COMPUSTAT, computed as the sum of long-term debts (*dltt*) and debts in current liabilities (*dlc*), divided by total stockholders' equity (*seq*). *DSCR* is the quarterly household debt service coverage ratio from the Federal Reserve, defined as household debt service payments as a fraction of disposable income. *IPGrowth* is the month-on-month growth rate in the industrial production index (INDPRO). In some specifications, we add the war component of the News Implied Volatility Index (NVIX) of Manela and Moreira (2017). We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. The time sample varies depending on the data availability of covariates, with *DSCR* available starting in 1980, and the NVIX available only up to 2016. *t*-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

— is that residential building permits simply reflect changes in the demand and supply of mortgages. To address this concern, we download the monthly loan origination and application series created by Neil Bhutta using the confidential Home Mortgage Disclosure Act (CHMDA).[9] Relative to the public HMDA files, which only reveal the origination year, the collapsed version of the CHMDA data matches the frequency of our building permits data. We sum the mortgage totals across states to produce a nationwide total.

For this analysis, we focus on the post-1960s period, covering the modern Census BPS, because household and corporate leverage measures are only available in the modern period. We present results in Table D.2. There are two key takeaways: one is that the loading on σ_{t-1}^{BPG} is nearly quantitatively identical to those estimated in Table 2 after conditioning on commodity price risk, as we do in columns (2) and (3) for equities and in columns (7) and (8) for corporate bonds. Second, controlling for growth in household leverage demand, as proxied by the month-to-month change in HMDA mortgage applications, halves the predictability of σ^{BPG} around the Great Recession episode. However, the predictability remains substantial; a 1 p.p. in the subsequent month is associated with a 45% increase in stock volatility relative to its monthly average in the 2000s. This is true regardless of whether we control for mortgage application vs. origination growth, or instead control for the level or log level of mortgage dollars originated (unreported). Overall, easy credit access for homebuyers or commodity price risk does not eliminate the predictability of building permit volatility for financial markets.

D.5 Industry Heterogeneity in the Cross-Section of BPG Loadings

The main text establishes that BPG volatility predicts aggregate stock and bond return volatility conditional on macroeconomic controls. In this subsection, we ask where in the industry cross-section this predictability is concentrated. The question is motivated by the noisy rational expectations framework in Section 6, which treats local residential permits as a quasi-public signal about real economic fundamentals that are otherwise difficult to observe at high frequency. Under this interpretation, BPG volatility should load more strongly on the asset prices of firms whose cash flows are sensitive to the underlying state variable, and less strongly on the asset prices of firms whose cash flows are relatively inelastic to the cycle.

We re-estimate the baseline volatility regression separately for each Fama-French 17-industry portfolio, using single-family BPG volatility on the right-hand side to match Panel B of equation 2. For each industry j, we estimate

$$\sigma_t^{(j)} = \alpha_j + \beta_j \cdot \sigma_{t-1}^{BPG} + \rho_j \cdot \sigma_{t-1}^{(j)} + \Gamma_j' \mathbf{X_{t-1}} + \lambda_m^{(j)} + u_t^{(j)}, \tag{D.1}$$

where $\mathbf{X_{t-1}}$ contains one-period lags of U.S. population growth, aggregate corporate leverage, and industrial production growth. $\lambda_m^{(j)}$ is a month fixed effect, and $\sigma_{t-1}^{(j)}$ is an autoregressive term in sector return volatility. We choose this control set to maximize sample coverage. We estimate the regression over the 1961M02–2019M12 estimation sample and within twelve non-overlapping five-year calendar windows, labeling each window by the lustrum to which it belongs, using Newey–West standard errors with a lag length of $\lceil T^{1/4} \rceil$.

Figure D.4 reports the within-window rank of each industry's estimated β^j, where a rank of 1 corresponds to the largest coefficient in the window. Industries are sorted vertically by their full-sample coefficient, with the highest-loading industries at the top.

[9] The data are downloadable at https://sites.google.com/site/neilbhutta/data.

TABLE D.2. Regressions of Asset Return Volatility on U.S. BPG Volatility: Additional Controls

A. Total Private Residential Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.028***	0.028**	0.031**	0.046**	0.045**	0.036***	0.037***	0.040***	0.017***	0.017***
	(2.59)	(2.52)	(2.45)	(2.12)	(2.10)	(3.87)	(3.59)	(3.31)	(3.25)	(3.29)
Time sample	1960-19	1960-19	1980-16	2000-19	2000-19	1960-19	1960-19	1980-16	2000-19	2000-19
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged comm. return vol.		✓	✓				✓	✓		
Other controls$_{t-p}$			✓					✓		
ΔHMDA applications				✓					✓	
ΔHMDA $ originations					✓					✓
N	714	714	435	239	239	714	714	435	239	239
R^2	0.469	0.469	0.474	0.568	0.567	0.370	0.370	0.444	0.506	0.507

B. Single-Family Units BPG Volatility

Asset Market:	Equities					Corporate Bonds				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
σ_{t-1}^{BPG}	0.024**	0.024**	0.026**	0.047***	0.046***	0.044***	0.044***	0.047***	0.016***	0.016***
	(2.48)	(2.41)	(2.29)	(2.66)	(2.62)	(4.53)	(4.36)	(4.44)	(3.60)	(3.60)
Time sample	1960-19	1960-19	1980-16	2000-19	2000-19	1960-19	1960-19	1980-16	2000-19	2000-19
Monthly dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged asset return vol.	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lagged comm. return vol.		✓	✓				✓	✓		
Other controls$_{t-p}$			✓					✓		
ΔHMDA applications				✓					✓	
ΔHMDA $ originations					✓					✓
N	714	714	435	239	239	714	714	435	239	239
R^2	0.468	0.468	0.475	0.571	0.570	0.390	0.390	0.467	0.508	0.508

Notes: The table presents estimates from equation (3.3) relating total return volatility to lagged building permit growth (BPG) volatility implied by a GJR-GARCH model. In Panel A we use total private residential permits as the quantity measure $Q_{s,t}$, but instead use permits attached to single-family units in Panel B. All columns include a full set of monthly dummies and the asset return volatility lagged by one month. We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. We include in columns 3 and 8 a set of controls for other macroeconomic observables directly related to BPG volatility, which might also drive aggregate financial market volatility. This vector of other controls corresponds to those used in Table 2. Lagged commodity return volatility refers to the GARCH-implied volatility obtained from the commodity excess return index constructed by Janardanan et al. (2024). To control for monthly growth in mortgage credit demand based on applications, or growth in equilibrium credit supply based on originated mortgages, we append nationwide totals provided by Neil Bhutta using an aggregated version of the confidential Home Mortgage Disclosure Act (CHMDA) data, downloadable at: `https://sites.google.com/site/neilbhutta/data`. The CHMDA data are only available starting in January 1994. We include both for-purchase and refinancing loan growth rates as separate variables in the regressions. We include a BIC-optimal number of lags for each specification with control variables, for which we obtain a lag order of $p = 1$. t-statistics obtained from Newey–West standard errors where we select for each specification the minimum lag order such that the estimator for the covariance matrix is consistent. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

FIGURE D.3. Incremental R^2: BPG Volatility versus Corporate Leverage

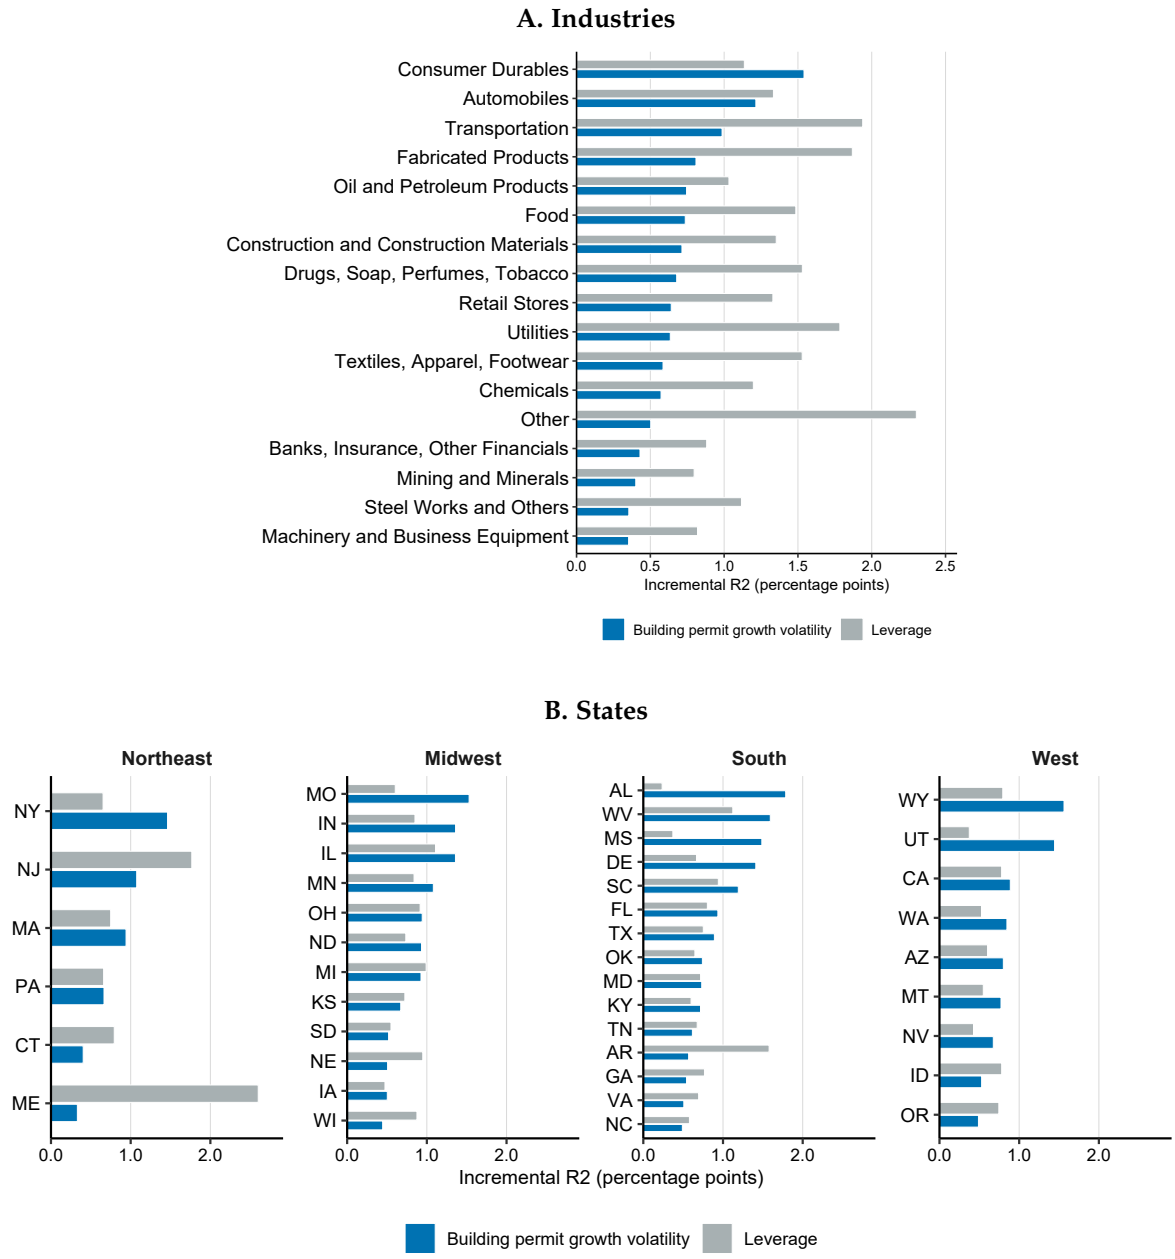

Notes: The figure reports the incremental value of the R^2, in percentage points, from adding lags 1–12 of building permit growth (BPG) volatility or corporate leverage to the same autoregressive stock-volatility model with month fixed effects. Panel A reports results for Fama–French 17 industries using aggregate U.S. BPG volatility. Panel B reports state-level results using own-state BPG volatility and aggregate corporate leverage. The state panel includes all available states with converged GARCH estimates after excluding Alaska, Hawaii, New Hampshire, Vermont, New Mexico, and Rhode Island.

FIGURE D.4. Fama–French 17 Industry Cross-Section: Rank of Loadings on Lagged SFH BPG Volatility

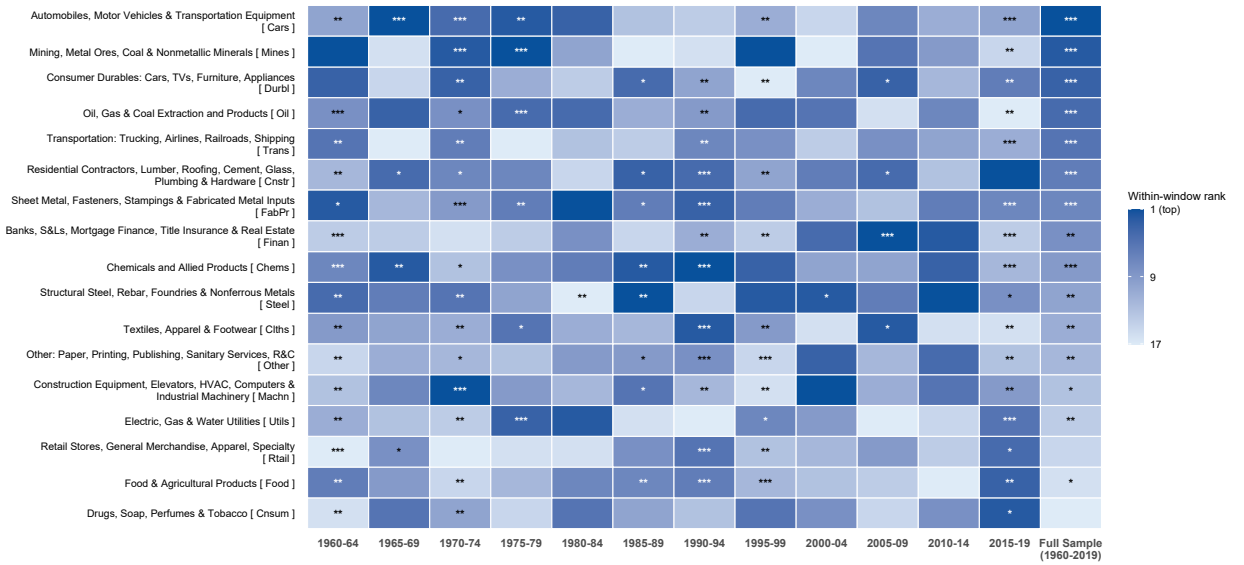

Notes: The figure reports the within-window rank of each Fama-French 17-industry portfolio's estimated coefficient on lagged U.S. single-family BPG volatility from equation D.1, where a rank of 1 corresponds to the largest coefficient in the window. Darker cells indicate higher within-window ranks; lighter cells indicate lower ranks. Significance labels use $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. Controls include month fixed effects, a lagged dependent variable, and one-period lags of U.S. population growth, aggregate corporate leverage, and industrial production growth. We choose this control set to maximize sample coverage. Rows are ordered from top to bottom by each industry's full-sample coefficient. Non-overlapping five-year calendar windows use monthly observations within each block and are labeled by the lustrum to which they belong. Because the estimation sample begins in 1961M02, the first cell labeled 1960–64 is estimated on the 1961M02–1964M12 observations. The final column uses the full 1961M02–2019M12 sample.

Two features of the cross-section are worth describing. First, in the full 1960–2019 sample, lagged U.S. SFH BPG volatility positively predicts return volatility for all 17 industries, and the coefficient is statistically significant at the 10% level for 15 of them. The uniform positive sign and the number of significant loadings indicate that the aggregate predictability documented in the main text is not concentrated in a narrow subset of firms.

Second, the industries with the largest full-sample coefficients are predominantly cyclical. In descending order of the estimated loading, the top half of the ranking contains motor vehicles (*Cars*), mining (*Mines*), consumer durables (*Durbl*), oil (*Oil*), transportation equipment and services (*Trans*), residential construction (*Cnstr*), fabricated metal inputs (*FabPr*), banks and mortgage finance (*Finan*), and chemicals (*Chems*). More defensive industries, including retail, consumer nondurables, food products, and utilities, appear in the bottom half of the ranking with smaller coefficients. This ordering is consistent with the signal-extraction interpretation of the baseline result. Because residential permits embed forward-looking information about real economic fundamentals, BPG volatility should be absorbed most strongly by the asset prices of firms whose cash flows are most sensitive to the underlying cyclical state, and our cross-section displays this gradient.

The concentration of the signal in motor-vehicle and transportation-equipment industries echoes an earlier finding from the aggregate time series in Cortes and Weidenmier (2019). They show that the growth rate of U.S. truck production is the single aggregate leading indicator whose inclusion most attenuates BPG volatility's predictive power for stock returns, and they interpret this correlation as evidence that industrial activity related to producing and moving

construction inputs is a close co-channel with residential permits for signaling real economic fundamentals. Our industry-level results are consistent with that interpretation. *Cars* contains the SIC codes for motor-vehicle assembly and parts, *Trans* contains the SIC codes for truck and bus bodies, truck trailers, and freight transportation services, and both industries appear among the top five in the full-sample ranking.

D.6 State-Level Results Controlling for Local Economic Conditions

In this appendix, we show that our results from Section 4.1 showing the predictability of BPG volatility in the geographic cross-section during the Census survey period hold after conditioning on proxies for housing demand such as state-level leverage and population growth. We download annual state-level population estimates from the Census and linearly interpolate within the year to match the monthly frequency of building permits.

We construct state-level corporate leverage ratios by assigning the leverage of each COMPUSTAT firm i in a given year t to its headquarters (HQ) state s and then compute the average leverage ratio for each state. That is, we compute:

$$Leverage_{s,t} = \frac{\sum_{i \in s}(dltt_{i,t} + dlc_{i,t})}{\sum_{i \in s} seq_{i,t}} \tag{D.2}$$

where `dltt` is total long-term debt outstanding, `dlc` is debt in current liabilities, and `seq` refers to total parent stockholders' equity. To assign each firm to its HQ state, we splice the list of COMPUSTAT HQs matched to a `gvkey` from Bai et al. (2020) covering 1969–2003 to an annex file provided by Gao et al. (2021) for 2004 to 2022. The latter set of authors created the annex file by scraping SEC 10-K filings to collect the business address over the later period. The HQ state is distinct from the state of incorporation, which is Delaware for the majority of firms. We use the HQ state to proxy for location because it reflects the largest fraction of firms' physical resources and business operations.

Gao et al. (2021) report that based on SEC filings, between 1995 and 2018, 2% to 3% of COMPUSTAT firms change their HQ state each year. Given the low probability of year-to-year cross-state changes to the HQ location, we fill in missing values for state HQs within each firm's history by carrying forward and backcasting for firms which have only ever listed a single HQ state in their available SEC filings; we do not fill in missing locations for firms which ever change their HQ state.[10] After filling in missing values, we then collapse to a state-level panel of leverage ratios according to (D.2). For nine observations in AK, NM, and WY, reported `seq` is negative due to a small number of firms with HQs located in those state-years; we set those leverage observations to missing. We then winsorize leverage ratios at the 1st and 99th percentile across the entire panel to minimize the role of measurement error arising from a combination of missing values for the firm balance sheet variables in (D.2) and a small number of firms with headquarters in less-populated states. For including leverage ratios as a control in equation (3.4), we consider the leverage ratio to be as of the start of the following year, in keeping with the majority of firms' fiscal year end dates. Hence, the variation due to leverage

[10] While we could also carryback the state HQ location within each `gvkey` to 1960 when COMPUSTAT coverage starts, this results in erroneously high leverage ratios > 2 for several small states in the early 1960s due to missing values for shareholder equity. Hence, we start our sample in 1969 when the Bai et al. (2020) state HQ file begins.

is concentrated in the turnover from December to January; monthly dummies account for any base month effects that this convention creates.[11]

State-level corporate leverage is pro-cyclical, with large spikes in leverage observed on the eve of the Global Financial Crisis. There is high degree of geographic dispersion in leverage ratios that widens following recessions. For instance, the average annual standard deviation in state leverage ratios is 0.27, compared to 0.44 in 2002, 0.38 in 1991, and 0.37 in 2009.

Figure D.5 plots the estimated predictability of BPG volatility from the GJR-GARCH specification in equation (C.3), but including interpolated monthly population growth and annual state-level corporate leverage ratios as controls for local housing demand. We use GJR-GARCH for this exercise because the distribution of BPG residualized on leverage and population growth is more skewed than unconditional BPG. Comparing Figure D.5 to Figure C.2, we find the predictability of BPG volatility for asset return volatility is attenuated for the aggregate U.S. and most states, but the ordinal ranking of states according to their cumulative 12-month predictability coefficients is similar after conditioning on housing demand factors.

E Completion Rates: Evidence from Building Permit Microdata

In this section, we analyze building permit completion rates using microdata provided by Cotality through their *Building Permits* product. The dataset contains detailed information on individual residential building permits across the United States, tracking the status of permits from issuance to completion. We observe permit-level data for 48 states and the District of Columbia. Alaska and Mississippi are excluded from our analysis as their non-disclosure laws prevent access to building permit records. While several other states have non-disclosure laws for real estate transactions (Idaho, Kansas, Louisiana, Missouri, Montana, New Mexico, North Dakota, Texas, Utah, and Wyoming), Cotality maintains comprehensive permit records for these jurisdictions through direct relationships with local town planning departments.

We isolate permits for new residential construction to align the Cotality sample with permits counted towards the Census Building Permits Survey. In particular, we map Cotality permits to the Census survey categories: permits related to the construction of 1-unit, 2-4 unit, and $>$ 5-unit properties. To do so, we identify new construction permits using project type indicators provided by Cotality and then parse the land use description string from the county tax assessment records to sort permits into bins by the number of residential units. For instance, we sort new home construction permits with a "duplex" or "quadplex" description into the 2-4 unit category.

Cotality *Building Permits* provides the state and county land use descriptions as of the date the permits database was last updated. For example, consider a single-family home construction permit issued in 2010 and completed in 2011. If that property is then subsequently converted to office space, then filtering based on the land use description will cause us to drop this observation from our sample even though it pertains to a single-family unit in 2010. To address this possible issue, we merge *Building Permits* to Cotality *Tax* using the assessor's parcel number (APN) and parcel's location. Cotality *Tax* contains annual tax year histories of property characteristics for all parcels, including the land use descriptions as of each year. We then use the land use descriptions

[11] Alternatively, we could apportion leverage monthly using the fiscal year end date. This is feasible for states with a large number of corporate headquarters (e.g., New York), but would still require us to interpolate many state-month-year observations. 67% of COMPUSTAT firms adopt the calendar year as their fiscal year.

FIGURE D.5. Including Controls: Predictability of State-Level Permit Growth Volatility (GJR-GARCH, 1960–2019)

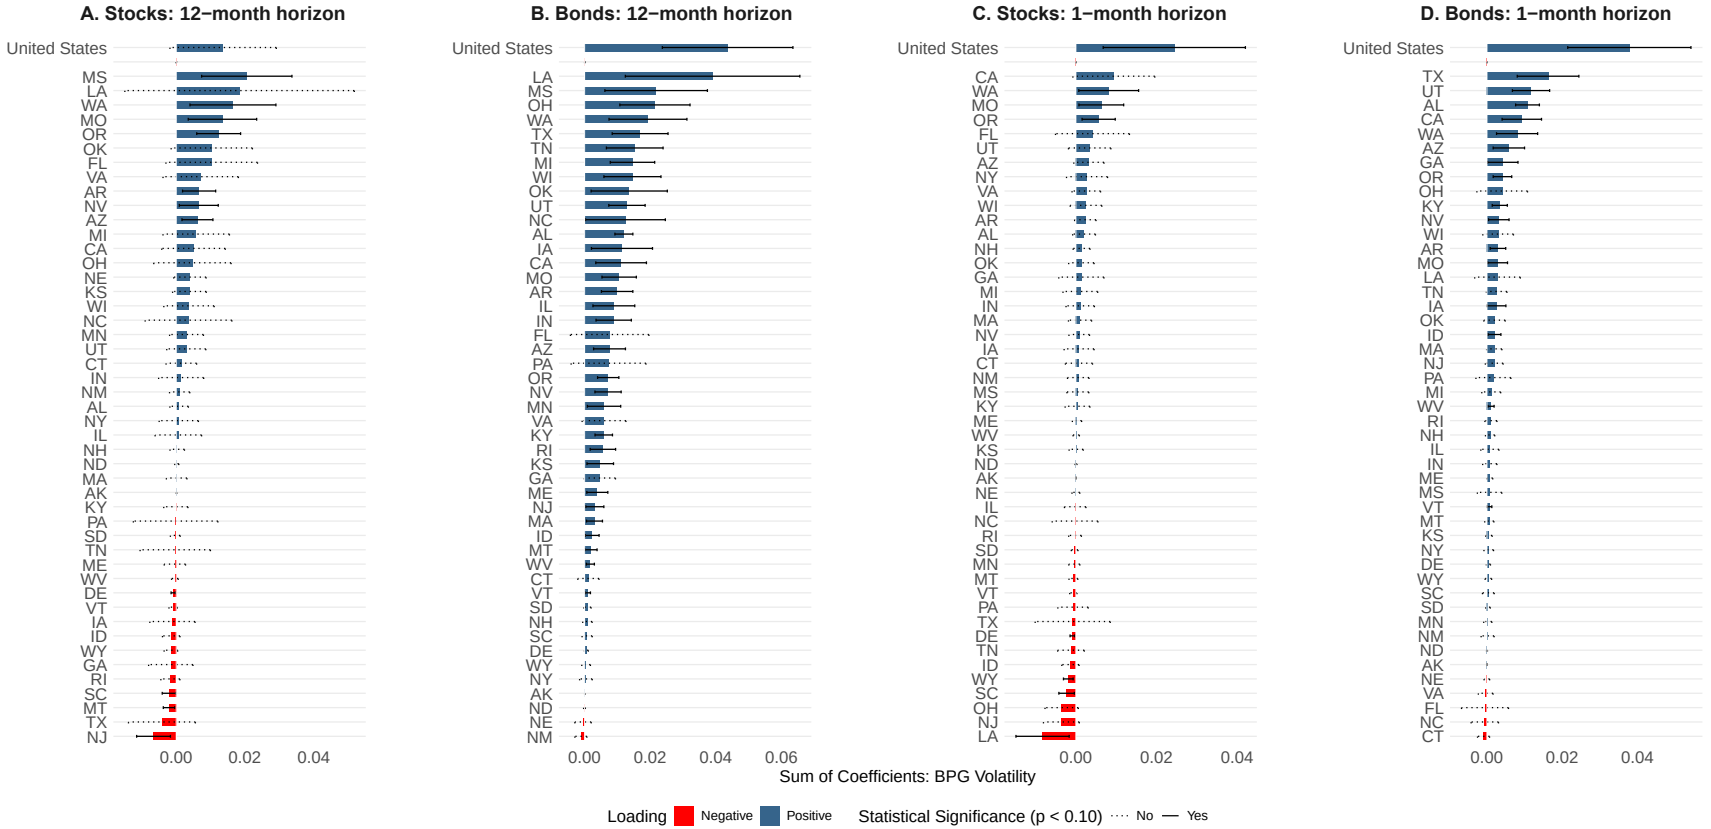

Notes: We estimate equation (3.3) for the entire U.S. and equation (3.4) separately for each state s using (C.3) as our conditional volatility model to obtain σ^{BPG}. Each regression includes monthly dummies, the BIC-optimal autocorrelation term σ_{t-1}, and controls for the state-level leverage ratio and interpolated monthly state-level population growth. We omit the regressions for Washington, D.C. and Massachusetts, as the GARCH model does not converge for those series. In panels A and B, we use $\tau^* = 12$ lags of $\sigma^{BPG}_{s,t}$ and report the sum of the estimated lagged coefficients $\sum_{\tau=1}^{\tau^*} \widehat{\beta}_{s,\tau}$. In panels C and D, we show predictability at the one-month horizon. Our sample period covers the Census Building Permit Survey from its start in 1960M5 to 2019M12. We use permits attached to single-family units as the quantity measure $Q_{s,t}$. Panels A and C perform this exercise for CRSP stock return volatility as the outcome, and total return volatility of the Dow Jones Corporate Bond Index in Panels B and D. We plot 90% confidence intervals obtained via autocorrelation and heteroskedasticity-robust standard errors, selecting the minimum lag order such that the estimator for the covariance matrix is consistent (Newey and West, 1987). Solid confidence bars indicate significance at the 90% level, while dashed bars indicate statistical insignificance.

FIGURE E.1. New Single-Family Home Permit Completion Shares by Application Year, 1990–2021

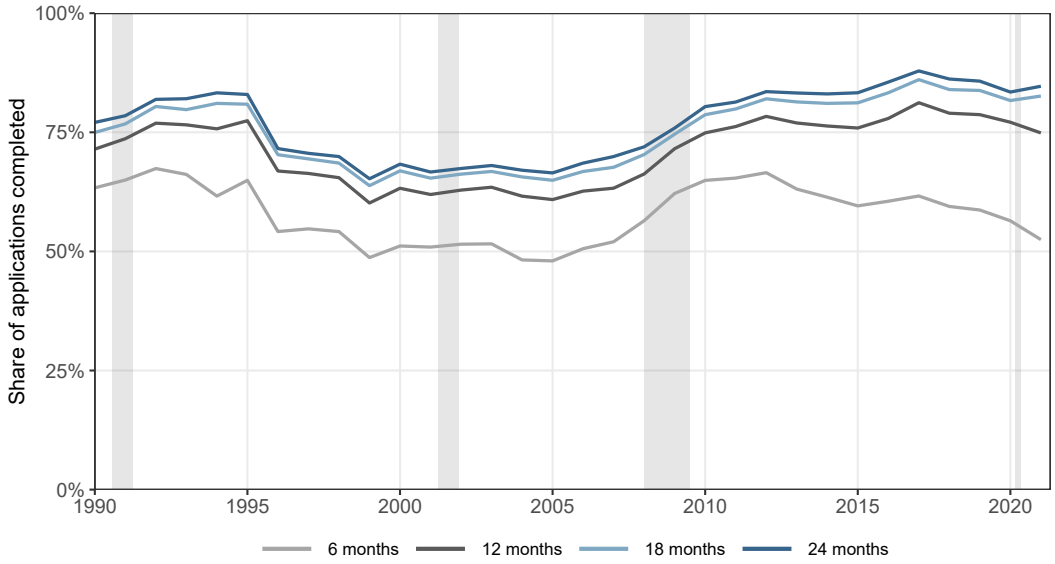

Notes: The figure plots the share of new single-family home permit applications in the Cotality *Building Permits* data. It shows permits recorded as completed within 6, 12, 18, and 24 months of the application year. The sample ends in 2021 to allow a complete 24-month follow-up window in data observed through 2023. Grey areas are NBER-dated recessions.

as of the tax year corresponding to the permit effective date to sort permits into residential vs. non-residential uses and by the number of housing units.

Figure E.1 plots residential permit completion shares by application year. The series move gradually over time and do not show a clear break around the 2008–2009 Global Financial Crisis.

Figure E.2 summarizes building permit unconditional completion rates for new single-family homes in a stylized map. Coastal states, which typically have more stringent land use regulations, tend to have lower completion rates (depicted in lighter-shaded quintiles), particularly in the Northeast. For instance, states in New England and the Tri-State Area fall into the lower completion rate quintiles. In contrast, several states in the Southeast and Mountain regions, which generally have fewer land use restrictions, exhibit higher completion rates (darker shades); notable examples of lax regulation states include North Carolina, South Carolina, and Colorado, which all fall in the highest completion rate quintiles. The pattern suggests a negative relationship between regulatory burden and permit completion rates, though this relationship is not uniform across all states. For example, California, despite its stringent land-use regulations, maintains moderate completion rates, possibly due to other factors such as strong market demand and established development processes.

Figure E.3 depicts how regulatory and physical constraints on new housing supply relate to permit completion rates in the cross-section of states. We compare 12-month and 24-month unconditional completion rates for new residential units, the x-axis in each bivariate map scale, to a constraint measure on the y-axis of the bivariate scale. Darkly-shaded states are ones for which a land use constraint measure is greater, while those shaded more red have greater completion rates. Both regulatory constraints, measured by either the Wharton Index (Gyourko et al., 2008) or the AI-based index component identify minimum lot size requirements (Bartik et al., 2024), and physical constraints based on topography or pre-existing development (Lutz

FIGURE E.2. Single-Family Building Permit Unconditional Completion Rates by State, 2012–2022

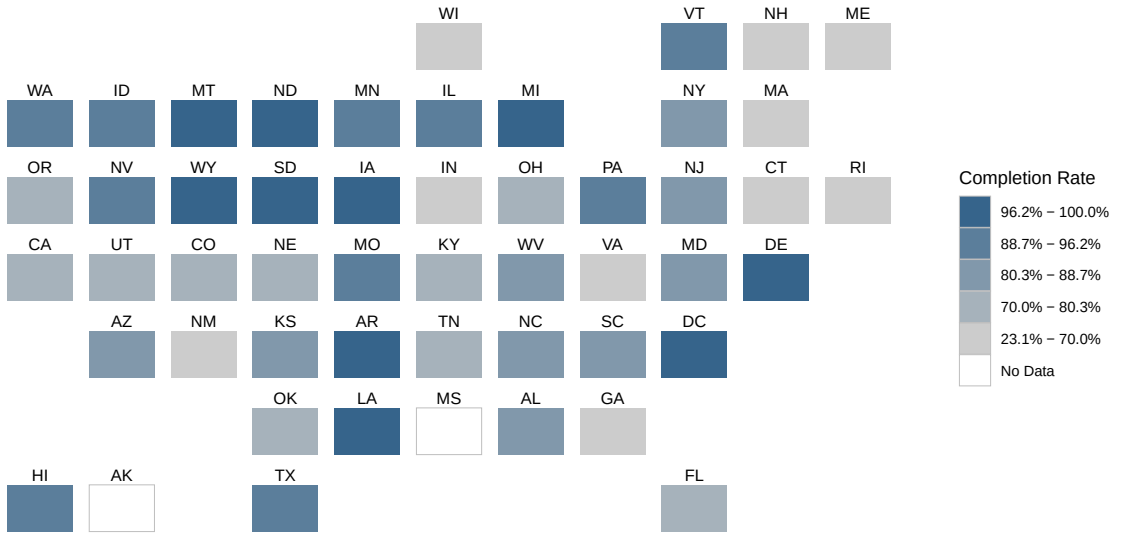

Notes: The map shows building completion rates for new single-family homes from Cotality (formerly, CoreLogic), depicted in a stylized map of the United States. Each state is represented in an equal-sized rectangle to allow for a clearer visualization. Building completion rates vary significantly across states, with states categorized by their unconditional completion rate quintiles. White cells with borders indicate non-disclosure states (AK, MS) where permit completion data is unavailable in most jurisdictions.

and Sand, 2023), are associated with lower short-term completion rates and therefore longer time-to-build and greater project failure rates.[12]

[12]Minimum lot size requirements are a binding constraint on the decision to build single-family homes and small multifamily properties (Kulka et al., 2026). Developers' decisions to build certain kinds of large multifamily and commercial properties are instead more likely to be dictated by maximum floor-to-area (FAR) ratio or height restrictions in non-residential zones of a city (LaPoint, 2021) For this reason, we compare this type of regulation to single-family unit permits rather than multifamily ones.

FIGURE E.3. Unconditional 12-Month and 24-Month Completion Rates for Residential Permits by State and Land Use Constraints, 1990–2022

A. 12-Month Rates by WRLURI

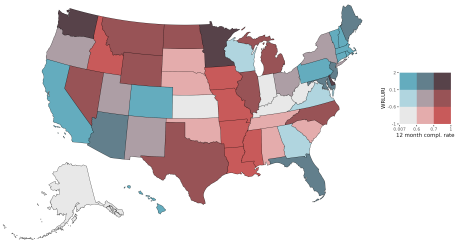

B. 24-Month Rates by WRLURI

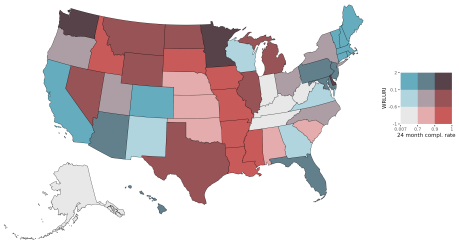

C. 12-Month Rates by Buildable Land

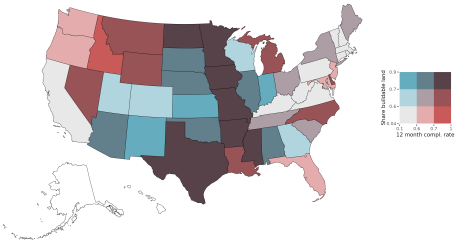

D. 24-Month Rates by Buildable Land

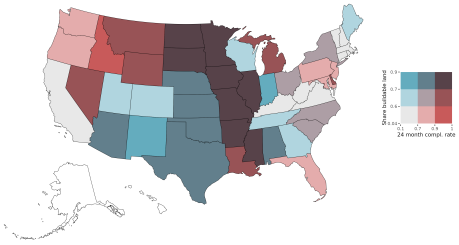

E. 12-Month SFH Rates by Min. Lot Size

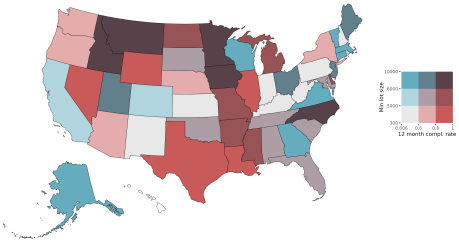

F. 24-Month SFH Rates by Min. Lot Size

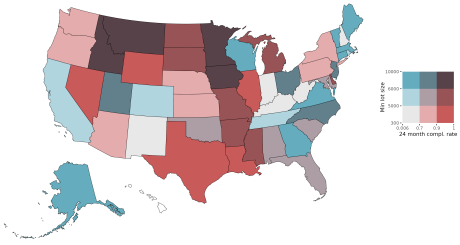

Notes: Each panel shows a bivariate heat map summarizing how 12-month (left-hand side) or 24-month (right-hand side) unconditional completion rates on residential building permits in each state vary with different measures of local constraints on new construction. We compute completion rates by counting the number of permits with time between permit effective date and the date when the permit's status was updated to "completed" of less than 12 or 24 months, and dividing by the number of permits issued in the Cotality (formerly, CoreLogic) *Building Permits* data. States depicted in white are ones where one of the measures is missing. Darkly-shaded states are ones where the land use constraint measure is greater, while states in deeper red shades have greater completion rates. Panels A and B compare completion rates to the earlier version of the Wharton Residential Land Use Regulatory Index (WRLURI) of Gyourko et al. (2008) based on surveys of local governments conducted in 2005. Panels C and D tabulate completion rates against the buildable land share of Lutz and Sand (2023), which refines the earlier land use availability measures of Saiz (2010) by using satellite imagery within a geographic polygon as of 2001 and accounting for existing construction and parks. We aggregate up the buildable land measure from the zip code to the state level by taking the ratio of the sum of buildable land divided by the sum of available land across all zip codes within the state. Panels E and F restrict to new single-family home (SFH) permits and tabulate completion rates against population-weighted average minimum lot size requirements, aggregated from the municipality to the state level, obtained from the generative AI-based index of Bartik et al. (2024).