Productivity Hysteresis from the Great Recession

Jose Carreño*
University of Oregon

April 2026

Abstract

The United States experienced a notable slowdown in labor productivity growth over the last decade. This paper exploits geographic variation across U.S. Metropolitan Statistical Areas (MSAs) to examine the relationship between the 2006–2012 decline in house prices—the housing bust—and this slowdown. Instrumental variable estimates support a causal link, indicating that a one standard deviation drop in house prices leads to a 6.3 percentage point increase in the local productivity gap—the shortfall in growth needed for an MSA to return to its pre-trend level—relative to an average gap of 11.53%.

Leveraging a newly constructed measure of capital expenditures at the MSA level, I show that the prolonged slump in investment following the Great Recession accounts for a substantial share of this effect. Further analysis documents that the housing bust triggered this investment decline, and ultimately the productivity slowdown, primarily through the collapse in consumption expenditures that followed the bust.

Finally, I develop a quantitative general equilibrium model that captures these empirical findings and estimate that the housing bust accounts for roughly half of the observed productivity slowdown.

---

* I am indebted to Marty Eichenbaum, Guido Lorenzoni, and Matthias Doepke for their invaluable guidance, advice, support and encouragement. A previous version of the paper circulated under the name "Housing Booms and the U.S. Productivity Puzzle". I am grateful to the editor and an anonymous referee for thoughtful comments and suggestions that substantially improved the paper. I also thank David Berger, Matt Notowidigdo, and Giorgio Primiceri for many useful discussions. I also wish to thank Rodrigo Adao, Cagri Akkoyun, Manuel Arellano, Bruno Barsanetti, Martin Beraja, Mark Bils, Gideon Bornstein, Ivan Canay, Gabe Chodorow-Reich, Nicolas Crouzet, Bob Gordon, Yuriy Gorodnichenko, Francois Gourio, Jessie Hanbury, Sasha Indarte, Omar Licandro, Martí Mestieri, Joel Mokyr, Alessandro Pavan, Matt Rognlie, Amir Sufi, Gabriel Ziegler, and seminar participants at Northwestern University, Bank of Spain, Bank of Portugal, Catolica-Lisbon School of Business and Economics, Bank of Canada, Federal Reserve Board, University of Washington, University of Oregon, Federal Reserve Bank of San Francisco, Florida State University, and the Philadelphia Workshop on Macroeconomics and Economic Policy at the University of Pennsylvania for helpful comments. The research in this paper was conducted while I was a Special Sworn Status researcher at the Chicago Census Research Data Center and the Berkeley Federal Statistical Research Data Center. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2133. (CBDRB-FY23-P2133-R10500SA). All errors are my own. Email: josecarreno.econ@gmail.com.

1 Introduction

The United States experienced a notable slowdown in labor productivity growth over the last decade. Figure 1 shows the post-1995 evolution of labor productivity, together with several trend lines based on the average growth rates of periods characterized by distinct productivity dynamics. Until the second half of the 2000s, productivity closely followed the prevailing trend. However, it then slowed sharply and, during the Great Recession (2008–2009), recorded a significant increase.[1] Strikingly, after this rebound, productivity did not resume its prior trend but instead stagnated, growing only very slowly for many subsequent years.[2] The resulting gap widened steadily over time and has been estimated to account for as much as a 16% loss in GDP, equivalent to $23,400 per household in the U.S. (Syverson, 2017). Although productivity growth has recovered significantly since 2016, a fundamental puzzle remains: what are the main drivers of this slowdown?

Previous literature has predominantly focused on linking this slowdown in productivity growth to long-run, acyclical forces, such as increasing challenges in measuring productivity, a decline in new business formation, or the possibility that recent innovations have been, for some reason, less transformative than those of earlier periods.[3] In contrast, this paper addresses the question of whether part of the observed productivity slowdown might instead reflect an endogenous response to the Great Recession.

In particular, this paper investigates the link between the 2006–2012 decline in house prices (the housing bust) and the subsequent labor productivity slowdown. To this end, I adopt an approach that exploits geographic variation across U.S. Metropolitan Statistical Areas (MSAs). This strategy offers several advantages. First, geographic cross-sectional analyses provide far greater variation than time-series approaches. Second, this variation helps isolate the effects of recession-independent ("secular") forces and nationwide shocks that are common across regions and that, over the multi-year horizon relevant for this study, could influence productivity dynamics and confound the results. Third, this type of analysis connects more directly to the structural primitives of the economic environment, as policy responses are largely "differenced out," yielding sharper theoretical predictions (Nakamura and Steinsson, 2014). This approach builds on the novel finding that the productivity slowdown

[1] This type of sharp productivity increase is a consistent pattern across recessions. One hypothesis is that it reflects changes in the composition of the labor force, driven by large job losses disproportionately affecting low-skill workers (see Charles et al., 2018b for evidence in the context of the Great Recession).

[2] This persistent weakness was largely unanticipated, as shown by the Congressional Budget Office (CBO)'s successive projections, some of which are displayed in Figure B.1.

[3] See, for example, Mokyr (2014a), Guvenen et al. (2017), Fernald (2015), Decker et al. (2017), and Gordon (2017).

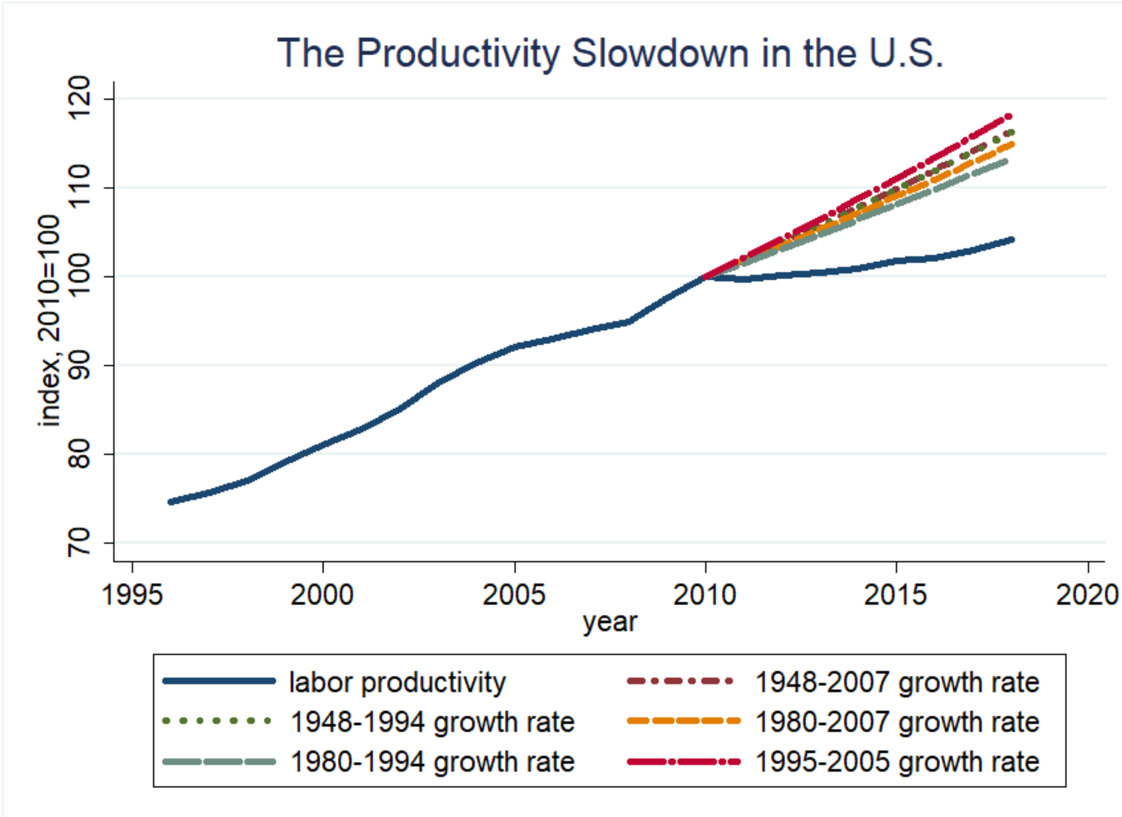

Figure 1: Labor productivity (solid line) is measured as real GDP divided by hours worked. The dashed lines display several trend paths extrapolated from 2010, based on the average growth rates of various post-1948 periods characterized by differing productivity dynamics. The period 1980–1995 was relatively slower than the long-run average, whereas 1995–2005 was relatively faster. Data source: Bureau of Labor Statistics.

exhibits substantial spatial heterogeneity—mirroring the considerable geographic variation in the magnitude of the 2000s housing cycle (Ferreira and Gyourko, 2011) and in the depth of the Great Recession (Mian and Sufi, 2014). Taken together, these features make U.S. MSAs a uniquely suitable laboratory for this investigation.

I measure productivity at the MSA level as real GDP per worker. I define the MSA-specific productivity slowdown as the difference between the log value of the pre-2006 trend extrapolated to 2015 and the actual log value of the series in that same year. I refer to this measure as the *productivity gap*, representing how much the series would need to grow—at a given point in time—to return to its underlying trend. Across MSAs, the average productivity gap is 11.53%.

I structure the paper in three main parts. The first part analyzes the direct relationship between the housing bust and the productivity slowdown. I first show that the productivity gap is, on average, substantially greater in MSAs that experienced larger house price declines. However, there may be confounding factors that could simultaneously influence both the housing bust and the productivity slowdown. To address endogeneity concerns, I implement an IV strategy drawing on the extensive literature that exploits geographic variation in housing supply elasticity to instrument for the housing boom-bust cycle.[4] To further mitigate concerns about self-selection related to pre-sample labor force and industrial composition, I include a comprehensive set of controls capturing MSA size and demographics, pre-sample labor force characteristics, pre-sample industrial composition, and pre-sample labor productivity levels.

The results indicate that a one-standard-deviation decline in house prices translates into an increase in the productivity gap of 6.3 percentage points. A simple back-of-the-envelope calculation suggests that the housing bust accounts for approximately 50% of the observed productivity gap. Moreover, a year-by-year analysis shows that this effect grew progressively over time, emerging around 2009–2010, peaking near 2014–2015, and gradually declining thereafter.

The second part of the paper explores the mechanisms linking the housing bust to the productivity slowdown. I proceed in two steps. First, I investigate whether the housing bust may have contributed to the productivity gap *via* the prolonged slump in investment that followed the Great Recession. To this end—and given the lack of comprehensive measures of capital expenditures below the national level—I construct a new measure of capital expenditures at the MSA level using confidential microdata from the Census Bureau's Annual Capital Expenditures Survey. I present evidence indicating that this investment slump represents an important channel through which the housing bust impacted the productivity slowdown.

Second, I examine the two primary channels through which the housing bust could have affected investment and, ultimately, the productivity gap. On one hand, the housing bust triggered a sharp decline in consumption expenditures (Mian et al., 2013), which may have persistently discouraged corporate investment. On the other hand, firms in MSAs that experienced larger house price declines might have faced tighter constraints on access to external funding—a mechanism I refer to as the "credit supply hypothesis." I provide empirical evidence pointing to a central role for the former and a more limited role for the latter.

In the third part of the paper, I develop a quantitative general equilibrium model that is directly informed and disciplined by the empirical findings. I use the model both to rationalize the empirical results and, importantly, to quantify the contribution of the documented

[4] See, among many others, Mian and Sufi (2009), Mian and Sufi (2011), Mian and Sufi (2014), Giroud and Mueller (2017), Stroebel and Vavra (2019), and Guren et al. (2018).

mechanism to the aggregate productivity slowdown.

The model builds on a New Keynesian framework with nominal rigidities and features two key ingredients. First, it incorporates a vintage capital structure in which each period introduces a new generation of machines that are more productive than those built previously. This mechanism creates a direct link between investment and productivity via technology diffusion and the standard capital deepening channel. Second, it includes a putty-clay production function with irreversible investment, allowing firms to freely choose the characteristics of capital goods at the time of installation, but fixing those characteristics thereafter. This feature, among other things, makes the persistence of productivity dynamics more consistent with the data. Together, these two ingredients shape the way productivity responds to investment fluctuations, yielding dynamics that align more closely with empirical evidence.

In the presence of nominal rigidities, a negative demand shock lowers the expected returns to investment, thereby slowing the pace at which new investment-specific (capital-embodied) technology is adopted and reducing the accumulation of capital per worker. This mechanism naturally generates a positive comovement between consumption and investment. Ultimately, both forces produce a slump in labor productivity, even in an environment where the technological frontier continues to advance as usual—an outcome consistent with evidence showing no significant decline in innovation indicators such as R&D expenditures and patents.[5]

The model matches several key moments in the data that were not explicitly targeted in the calibration and suggests that the housing bust accounts for approximately half of the measured labor productivity gap in the United States—a finding consistent with estimates from the empirical analysis. Taken together, these results stand in contrast to the prevailing narrative that attributes the productivity slowdown solely to secular forces.

Related Literature

This paper contributes to three main strands of literature.

First, it adds to the literature examining the causes of the U.S. productivity slowdown. One explanation posits that the slowdown is, to a large extent, illusory and reflects measurement problems.[6] The idea is that recent gains in productivity are not fully captured in official

[5] See Appendix H and, in particular, Figure H.1. Moreover, extensive literature has argued that it is not a lack of productivity-enhancing innovations that constrains productivity growth (see, for example, Brynjolfsson and McAfee, 2014 and Mokyr, 2014b).

[6] See, for example, Brynjolfsson and McAfee (2014); Mokyr (2014a); Guvenen et al. (2017); Byrne et al.

statistics for various reasons. However, Syverson (2017) provides extensive empirical evidence suggesting this hypothesis is largely at odds with the data, a view supported by findings in Nakamura and Soloveichik (2016), Byrne et al. (2016), and Cardarelli and Lusinyan (2015).

Another prominent hypothesis, rooted in Byrne et al. (2013), Fernald (2015), and Cette et al. (2015), attributes the slowdown to the reversal of productivity accelerations in manufacturing and the diffusion of information and communication technologies (ICTs) during the late 1990s and early 2000s. Gordon (2017) offers several potential explanations, linking the current slowdown to the one in the late 1970s, while considering the surge of the late 1990s and early 2000s as a one-off aberration. Although these hypotheses likely capture part of the story, reversal and aberration narratives appear insufficient to explain the full slowdown, as Figure 1 clearly shows that productivity growth slowed even relative to average rates observed during 1948–1994 and 1980–1994.

Finally, Cowen (2011) discusses various reasons why innovation may have decelerated, while Decker et al. (2017) and Alon et al. (2018) emphasize the slowdown in U.S. business dynamism. In contrast, this paper focuses on the link between the Great Recession and the productivity slowdown.

On the theoretical front, Anzoategui et al. (2019) develop an endogenous growth model based on R&D and technology adoption to examine how these processes contribute to business cycle persistence, concluding that technology adoption plays an important role, while exogenous TFP movements contribute little to explaining the recent productivity decline. My paper complements their approach by providing an empirical analysis that exploits geographic variation to document a causal relationship between the housing bust and the productivity slowdown and to shed light on the mechanisms linking the two.

Second, this paper contributes to the growing literature on the consequences of the 2006–2012 housing bust. Mian et al. (2013) and Guren et al. (2018) employ different empirical strategies to study the impact of the housing bust on consumption expenditures through housing wealth effects, while Berger et al. (2017) explores these findings from a theoretical perspective. In other contexts, Mian and Sufi (2014) examine employment effects; Charles et al. (2018a) document how the housing cycle influenced college attendance; and Stroebel and Vavra (2019) find a causal effect of house price changes on local retail prices during the boom and bust. To my knowledge, this paper is the first to investigate the effects of the housing bust on the productivity slowdown, and indeed, it is the first to analyze spatial variation in productivity itself.

(2018); Feldstein (2015); and Smith (2015).

Lastly, this study relates to the vast literature on the aftermath of the Great Recession. Lo and Rogoff (2015) reviews potential explanations for the sluggish post-recession growth and emphasizes demand-side factors—particularly the persistent problem of household debt, even years after the initial deleveraging—as major obstacles to recovery. The crucial role of weak demand has also been underscored by Summers (2016, 2017), who argues that if supply constraints were dominant, we would observe accelerating inflation, whereas subdued inflation and growth point instead to demand weakness. This paper complements these perspectives by demonstrating that the collapse and sluggish recovery of consumption expenditures following the Great Recession played a central role in linking the housing bust to the productivity slowdown. Moreover, it provides direct empirical evidence for a mechanism illustrating the so-called "inverse Say's Law," the idea that weak demand—when persistent—can erode the economy's potential output.

Hall (2015) conducts a time-series growth accounting decomposition, attributing the 13.3 percentage point shortfall in output relative to trend in 2013 to 3.5 points from total factor productivity, 3.9 points from capital input, and 5.9 points from labor input. My analysis complements Hall's work in that both underscore the crucial role of capital in the post-recession period. However, while Hall focuses on aggregate GDP growth, this paper centers specifically on the productivity slowdown and examines its connection to the housing bust, shedding light on the mechanisms involved.

This paper is also related to Bhattarai et al. (2021), who study the long-run local effects of the U.S. housing crisis and find persistent effects on employment and output, but no significant local response of labor productivity in their IV specification. The difference relative to my findings largely reflects differences in empirical object and design rather than a direct contradiction. Their analysis focuses on level changes over the housing cycle, whereas this paper measures the productivity slowdown as a deviation from each area's own pre-2006 trend. This distinction matters because the question here is not simply whether productivity levels changed after the crisis, but whether the housing bust left local productivity persistently below its previous trajectory. It also matters for timing: my analysis is designed to capture the post-2010 productivity shortfall, whereas measures that span the recession itself can mix that persistent effect with the temporary rise in labor productivity associated with the sharp collapse in employment. In that sense, the two papers speak to related but distinct margins of adjustment: Bhattarai et al. (2021) emphasize long-run scars in local employment and output, while this paper focuses on the persistent productivity gap that emerged after the Great Recession.

A fundamental contribution in this broader context is Yagan (2019), who exploits variation across local labor markets and linked employer-employee data to identify the long-term employment impacts of the Great Recession. Yagan's key finding is that exposure to severe local recessionary shocks significantly reduced local employment probabilities in 2015, years after the national recovery and despite falling local unemployment rates. While my paper differs in focus—examining the productivity slowdown rather than employment—both studies highlight the persistent, long-lasting economic consequences of the Great Recession.

Layout. The remainder of the paper is structured as follows. Section 2 describes the data. Section 3 introduces measurement details and documents the geographic heterogeneity of the productivity slowdown. Sections 4 and 5 present the empirical methodology and results. Section 6 describes the model and discusses its quantitative implications. Section 7 concludes.

2 Data

This section provides a brief overview of the main data sources, while a comprehensive discussion and variable definitions can be found in Appendix A. The unit of observation is the MSA, which is defined by the U.S. Office of Management and Budget as a contiguous area of at least 50,000 people, delineated around a central urban cluster. According to the 2010 census, 83.6% of Americans live in MSAs. The rationale for choosing MSAs is two-fold: they constitute local labor markets (Moretti, 2010)—unlike counties—and they are the smallest geographic delineation for which one can measure labor productivity at this time.

Labor productivity. Labor productivity is measured as real GDP per worker, sourced from the Bureau of Economic Analysis (BEA). Since GDP data at the MSA level is only available from 2001 onward, I also use personal income (from the BEA) to extend the trend back to 1969. This approach is validated in Appendix A, which documents that annual personal income growth closely tracks annual GDP growth at the state level over the entire sample. Furthermore, Section 4.3 provides extensive robustness checks confirming that the results are not sensitive to this choice of measure. While GDP data at the MSA level is limited in historical coverage, the combined use of personal income and robustness checks ensures that my measure captures the relevant dynamics of local labor productivity. Finally, a discussion of the distinction between GDP per worker and GDP per hour worked, and data limitations on hours worked at the MSA level, is provided in footnote 7 in Section 3.

House prices. The house price indexes are obtained from the Federal Housing Finance

Agency (FHFA). I compute the log difference between 2006 and 2012, and standardize the measure to facilitate the interpretation of the results. Further details and data sources are provided in Appendix A.

Consumption. There is no direct measure of consumption expenditures at the MSA level, as the most granular data available is at the state level. Some regional proxies used in prior literature include car purchases and credit card data (e.g., Mian et al., 2013), but these are proprietary and not accessible to me. Publicly available measures, such as retail sales (Fishback et al., 2005), are only collected at geographically disaggregated levels during Economic Census years (every five years), while retail sales tax data are often noisy and limited to a subset of regions (Garrett et al., 2005). My primary proxy for MSA-level consumption expenditures extrapolates national retail sales to MSAs using MSA retail GDP from the BEA. Appendix A details this methodology and shows that, when aggregated to the state level, the annual growth rate of this measure closely tracks changes in actual consumption expenditure growth rates. Furthermore, Section 4.3 provides additional evidence confirming that my results are robust to alternative specifications. The rationale for using retail sales, GDP, or employment as proxies for local consumption expenditures is that retail activity serves as an intermediate input for household consumption, since goods must be purchased before they can be consumed (Kaplan et al., 2016; Guren et al., 2018).

Capital expenditures. To the best of my knowledge, there is no publicly available—or even proprietary—comprehensive measure of capital expenditures at any geographic level finer than the nation as a whole. The data series that comes closest is the public statistics on capital expenditures from the Census of Manufacturers, which are only available during Economic Census years and limited to manufacturing. In this paper, I construct a novel measure of capital expenditures at the MSA level by leveraging confidential microdata from the Census Bureau's Annual Capital Expenditures Survey (ACES). Appendix A provides the methodological details and describes how this new variable is derived. Given the complexity of this construction, I summarize the main approach here but leave the full technical details to the appendix to maintain the flow of the main text. To my knowledge, this represents the first systematic attempt to measure capital expenditures at the MSA level, and may prove useful for future research on local investment dynamics.

Other variables. Other variables include population and the share of workers with a college degree from the decennial census; employment shares across different sectors, income per capita, and wages from the BEA; housing supply elasticity measures from Saiz (2010) and Baum-Snow and Han (2024); summary statistics on deposits and commercial and industrial loans from the Federal Deposit Insurance Corporation (FDIC); and additional controls constructed for robustness and mechanism exercises, including immigrant share, population density, Bartik growth, coastal and regional indicators, and a proxy for local exposure to intangible-intensive industries based on intellectual property products (IPP). Further details and data sources are provided in Appendix A. These variables are primarily used as controls or for robustness purposes throughout the empirical analysis.

3 Measuring the Productivity Slowdown across MSAs

This section explains the methodology used to measure the productivity slowdown at the MSA level and introduces the geographic variation that I exploit in the empirical analysis.

I measure labor productivity at the MSA level as real GDP (in constant 2009 dollars) per worker[7].

As Figure B.2 illustrates, there is substantial heterogeneity in the growth rates of labor productivity across MSAs. Therefore, measuring the productivity slowdown solely through simple growth rates could lead to misleading cross-sectional comparisons. A more accurate approach requires capturing deviations from local long-run trends, which provides a consistent benchmark for identifying genuine slowdowns in productivity dynamics.

A key challenge in constructing the trend is that GDP data for MSAs are only available starting in 2001. To address this limitation, I extend the series back to 1969 by using MSA personal income as a proxy during earlier years. Appendix A documents that annual growth in personal income closely tracks annual GDP growth at the state level throughout the entire sample period,[8] supporting the validity of this approach. Furthermore, Section 4.3 provides extensive robustness checks confirming that the main results are not sensitive to the choice of trend construction.

Figure 2 illustrates the construction of the measure of the productivity slowdown at the MSA level. First, I fit a flexible polynomial through the 1969–2006 values of the series to

[7] There are, to my knowledge, only two sources of hours worked data at the MSA level: the American Community Survey (ACS) and the Current Employment Statistics (CES). The problem with the ACS measure of hours is that it is only available for a subset of MSAs, as the smallest geographic area is the Public Use Microdata Area (PUMA), which contains at least 100,000 people and overlaps only with MSAs of that size or larger. Thus, PUMAs do not cover MSAs with 50,000 to 100,000 people. The problem with the CES measure of hours is that it is only available after 2006. These limitations motivate my use of output per worker as the baseline definition of labor productivity.

[8] They also track each other very well at the national level even at relatively high frequencies.

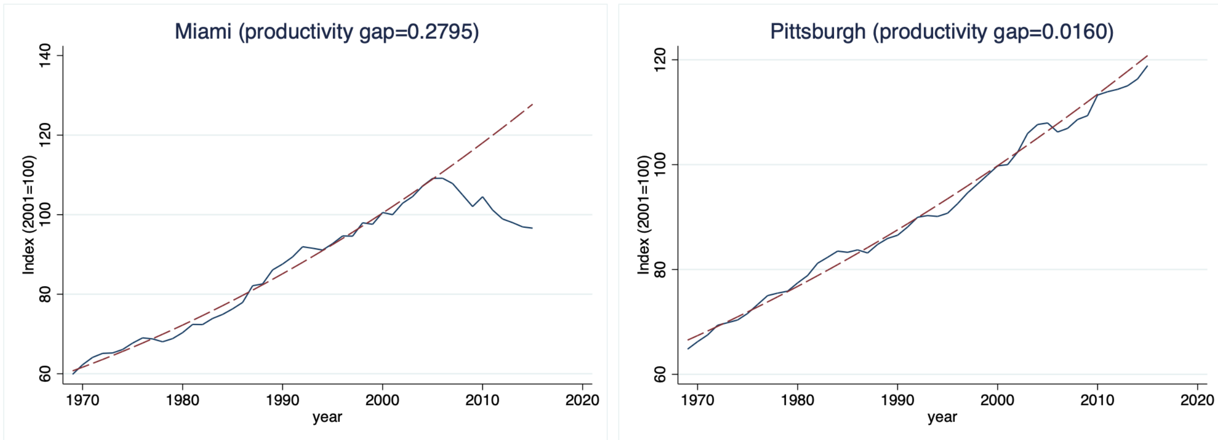

Figure 2: Illustration of the construction of the productivity gap for Miami (experiencing a large productivity slowdown) and Pittsburgh (experiencing a small productivity slowdown). The solid line depicts the real labor productivity series, while the dashed line represents the estimated long-run trend over 1969–2006. The gap between the two series serves as the measure of the local productivity slowdown used in the analysis.

approximate the long-run trend.[9] I then extrapolate this 1969–2006 trend forward to 2015. Finally, I define the labor productivity slowdown as the difference between the log of the extrapolated trend in 2015 and the log of the actual observed series in the same year.[10] I refer to this difference as the *productivity gap*, which captures how much labor productivity would need to grow, as of a given year, to catch up with its long-run trend. Figure 2 provides examples for two MSAs: Miami, on the left, which exhibits one of the largest productivity gaps at 27.95%, and Pittsburgh, on the right, which shows only a mild slowdown with no significant gap. Across all MSAs, the average productivity gap is 11.53%, with a standard deviation of 13.08%.

Figure 3 presents the geographic distribution of the productivity gap across MSAs. The spectrum ranges from areas shown in dark red and brown, indicating very large productivity gaps, through regions in yellow and light blue, where the gap is considerably smaller, to areas shaded in dark blue, where the productivity gap is negative—meaning that labor productivity is actually above its long-run trend. Figure B.3 displays a nonparametric probability density estimate of the productivity gap across MSAs, highlighting the substantial heterogeneity in local productivity dynamics.

[9]For the baseline results, I fit a third-degree polynomial. Results are virtually identical for any reasonable degree higher than one. I choose the period 1969–2006 because the series begins in 1969 and the Great Recession starts in 2007. Section 4.3 confirms that the main results remain robust to alternative specifications.

[10]I use 2015 as the benchmark year because it is the furthest point from the Great Recession before its effects begin to reverse, as discussed below. However, I also report results across all years in subsequent analyses.

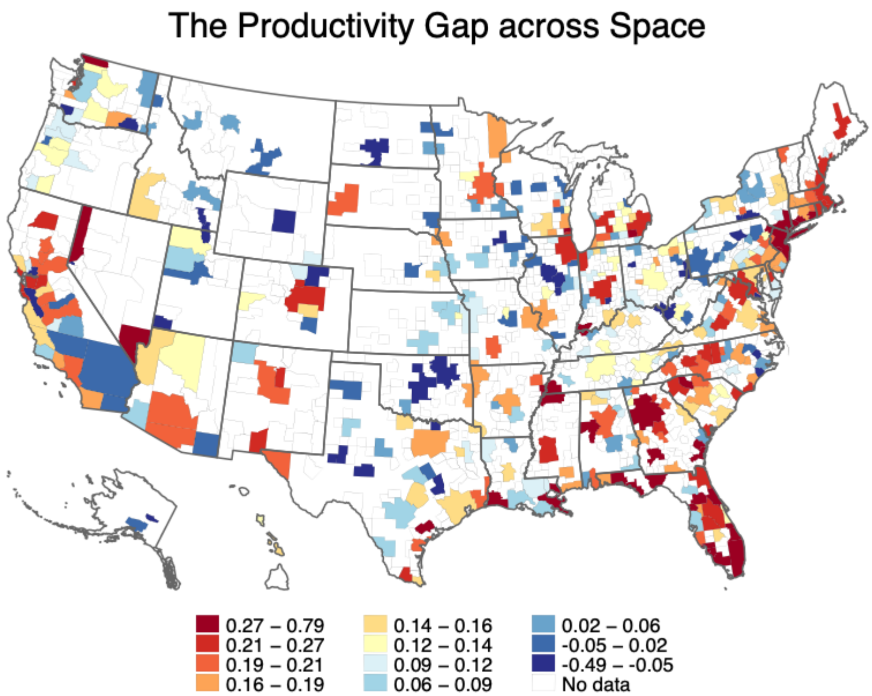

Figure 3: The productivity gap across MSAs in 2015. Colored areas represent MSAs, which encompass 83.7% of the U.S. population according to the 2010 Census. The remaining regions correspond to micropolitan areas (shown as uncolored, delimited areas, accounting for 10% of the population) and rural areas (uncolored, undelimited areas, covering 6.3% of the population). Darker colors indicate larger productivity gaps, while dark blue regions represent MSAs where labor productivity is above trend.

4 The Housing Bust and the Productivity Slowdown

This section investigates the direct relationship between the housing bust and the observed slowdown in productivity growth across MSAs, building upon the data sources and measurement framework introduced in the previous sections. The analysis is organized as follows: Subsection 4.1 describes the empirical strategy, Subsection 4.2 presents the main results, and Subsection 4.3 discusses robustness checks and alternative specifications.

4.1 Empirical Strategy

I begin by examining the cross-sectional relationship between the housing bust and the productivity slowdown across MSAs. My main empirical specification relates the 2015 productivity gap to the standardized log change in house prices between 2006 and 2012 (i.e., the housing bust):[11]

[11] I standardize the 2006–2012 log change in house prices to facilitate interpretation of the results, since the average decline across MSAs over this period is 18.45%.

$$\Delta\pi_i = \beta_0 + \beta_1 \Delta H_i + X_i \beta_2 + \epsilon_i, \tag{1}$$

where the unit of observation is the MSA, $\Delta\pi_i$ denotes the productivity gap in 2015, ΔH_i is the standardized housing bust measure described above, and X_i is a vector of controls.[12] Table C.1 reports the OLS results. Column (1) suggests that a one-standard-deviation decline in house prices between 2006 and 2012 is associated with an increase in the productivity gap of approximately 4 percentage points in 2015, relative to an average gap of 11.53%.

What are the productivity consequences of each standard deviation decline in house prices? Estimating this effect is challenging because there may be unobserved factors that simultaneously drive both the housing bust and persistent changes in productivity.[13] In such cases, the OLS estimates could suffer from omitted variable bias.

To address this challenge, I draw on the extensive literature that exploits geographic variation in housing supply elasticities to instrument for the housing boom-and-bust cycle (see, among many others, Mian and Sufi 2011; 2013; 2014; Giroud and Mueller, 2017; Stroebel and Vavra, 2019; Guren et al., 2018). The key insight is that heterogeneity in housing supply elasticities across MSAs causes differential local house price responses to a common housing demand shock, providing a source of plausibly exogenous variation in local exposure to the housing boom-and-bust cycle.

The specific measure of housing supply elasticity that I employ comes from Saiz (2010). This index quantifies the ease with which the supply of new housing can expand in MSAs, based on geographic characteristics. High elasticity scores are assigned primarily to areas with relatively flat topography, which facilitates urban development, while low elasticity scores correspond to regions where steep slopes or significant bodies of water constrain the availability of developable land.[14]

The first and second stages of the IV analysis correspond to the following regressions:

$$\Delta H_i = \gamma_0 + \gamma_1 HS_i + X_i \gamma_2 + \omega_i \tag{2}$$

[12]This vector includes population in 2000; the share of workers with a college degree and the average wage in 2001; the employment shares of services and manufacturing in 2001; and the level of labor productivity in 2001. The rationale for these choices is discussed below.

[13]It is relevant to emphasize the term *persistent*, because if the effects were purely transitory, they would not appear in the productivity gap measure—as this is computed in 2015 relative to the pre-2006 trend, while the housing bust occurred during the 2006–2012 period.

[14]Formally, the housing supply elasticity estimates are the predicted values from regression six in Table III, column 6 of Saiz (2010).

and

$$\Delta\pi_i = \beta_0 + \beta_1 \widehat{\Delta H}_i + X_i\beta_2 + \epsilon_i. \tag{3}$$

where the unit of observation is the MSA. All terms have been defined previously except for HS_i, which denotes the housing supply elasticity, and ω_i and ϵ_i, which are the error terms. The coefficient of interest is β_1, which captures the causal effect of local exposure to the housing bust on the productivity gap under the identifying assumptions discussed below.[15] The vector of controls, X_i, includes several pre-existing characteristics of the MSA, such as population in 2000; the composition of the labor force (measured by the share of workers with a college degree and the average wage in 2001); and industrial composition (captured by the employment shares of services and manufacturing in 2001). Importantly, X_i also includes the level of labor productivity in 2001.

Table C.2 reports the results of the first-stage regression, showing that the instrument strongly predicts the housing bust across MSAs.

The exclusion restriction requires that, conditional on the controls, housing supply elasticities affect the 2015 productivity gap only through their impact on the housing bust.[16]

To support the validity of Saiz (2010)'s housing supply elasticity as an instrument for the housing bust, Table C.3 shows that the instrument is uncorrelated with both pre-sample labor productivity levels and changes in labor productivity between 1998–2002 and 2002–2006. Furthermore, Mian and Sufi (2011) document that the instrument is not correlated with changes in wage growth over the same periods, nor with the employment share in the construction sector in 2006 or employment growth in construction from 2002 to 2006. Similarly, Stroebel and Vavra (2019) find no significant relationship between the instrument and income growth during the boom-and-bust cycle. These results lend support to the exclusion restriction underlying the IV strategy.

A residual concern is that the geographic characteristics contributing to housing supply inelasticity could also be associated, for different reasons, with systematic self-selection, leading to persistent differences in the pre-sample labor force composition (Davidoff et al.,

[15] In fact, the coefficient β_1 is expected to be negative, implying that its absolute value reflects the *increase* in the productivity gap resulting from a one-standard-deviation *decline* in house prices during 2006–2012.

[16] It is important to note that housing supply elasticity is a time-invariant characteristic, while the instrumented variable—the change in house prices—is time-variant. The elasticity index captures structural features intrinsic to each MSA (e.g., land availability, topography, proximity to water bodies) that remain relatively stable over time, rather than representing a shock per se. Thus, housing supply elasticity does not predict house price changes by itself but interacts with time-varying shocks—such as national housing demand shocks—to generate local house price responses.

2016) and industrial structure between elastic and inelastic MSAs. Such differences could bias the results if, for example, industry-specific shocks simultaneously affected both the housing bust and the productivity gap, and if the industries subject to those shocks were disproportionately concentrated in either more elastic or more inelastic MSAs.[17] To address this and other potential endogeneity concerns, the IV regressions include a rich set of controls summarizing pre-sample labor force composition (the share of workers with a college degree and the average wage in 2001) and industrial structure (captured by the employment shares of services and manufacturing in 2001).[18] Additionally, the set of controls includes labor productivity levels in 2001. This is an especially demanding control, since it helps absorb broad pre-existing differences across MSAs—including persistent local characteristics that may simultaneously shape both subsequent housing-market dynamics and long-run productivity outcomes. Section 4.3 further shows that the estimates remain stable when the specification is augmented with a broader Davidoff-style set of local demand controls.

4.2 Main Results

Table 1 reports the estimates from the second-stage regression. Column (1) presents the baseline results, yielding an estimate of -6.3. This implies that a one-standard-deviation decline in house prices is associated, in the IV specification, with an increase in the productivity gap of 6.3 percentage points, relative to an average productivity gap of 11.53%. These findings indicate a substantial impact of the housing bust on local productivity dynamics, a result that is further explored in subsequent robustness checks.

Column (2) reports estimates from a specification without controls. The similarity of the estimates across specifications suggests that the results are robust to the inclusion of controls, thereby alleviating concerns about omitted variable bias and lending support to the empirical strategy.

Figure 4 summarizes the IV estimates from the baseline specification (equation 3) for each year between 2009 and 2020. Specifically, each point estimate and its confidence interval derives from a separate regression in which the right-hand side of equation 3 remains unchanged, while the productivity gap is recalculated for each specific year.[19] The figure

[17] Appendix G specifically addresses this potential concern, as well as other related issues.

[18] Appendix G also presents a specification that controls for the employment share of each of the 23 two-digit industries.

[19] The procedure follows exactly the method described in Section 3, but applied to other years. For instance, the productivity gap for 2012 is calculated as the difference between the log value of the extrapolated trend component in 2012 and the log value of the actual series in the same year.

Table 1: Housing Bust versus Productivity Gap: IV analysis

	Productivity gap	
	(1) IV	(2) IV
House price 2006-2012	-0.063***	-0.071***
	(0.017)	(0.013)
Include baseline controls	Yes	No
First Stage F-statistic	72.22	34.51
First Stage (F-test) p-value	0.0000	0.0000
Mean dependent variable	0.1153	0.1153
N (Number of MSAs)	229	242

NOTE.—IV results, equation 3. House price change is standardized (standard deviation is 19.90, average is -18.45). Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

displays both the estimates and the 95% confidence intervals. The effect of the housing bust on the productivity gap has been increasing over time, becoming significant as early as 2009–2010 and reaching its peak around 2014–2015, with the previously noted point estimate of -6.3. Thereafter, the point estimates decline, suggesting that the effect may be gradually reversing.[20]

Finally, the effect of the housing bust on the productivity gap is substantial in economic terms. This can be seen from a simple back-of-the-envelope calculation. The log change in U.S. house prices between 2006 and 2012 is -21.33% (based on the FHFA annual house price index). Expressing this change in terms of cross-MSA standard deviations of the housing bust (with a standard deviation of 19.1%) and applying it to the benchmark estimate from the main specification implies that the housing bust increased the 2015 productivity gap by approximately 7.04 percentage points. For comparison, the national productivity gap in 2015 is 14.01%. Hence, this calculation suggests that the housing bust accounts for roughly half of the national productivity slowdown in 2015.[21]

[20] This interpretation is necessarily heuristic, as the confidence intervals widen substantially when moving farther away in time.

[21] See Beraja et al. (2016) and Chodorow-Reich (2019) for caveats regarding this interpretation and for further discussion.

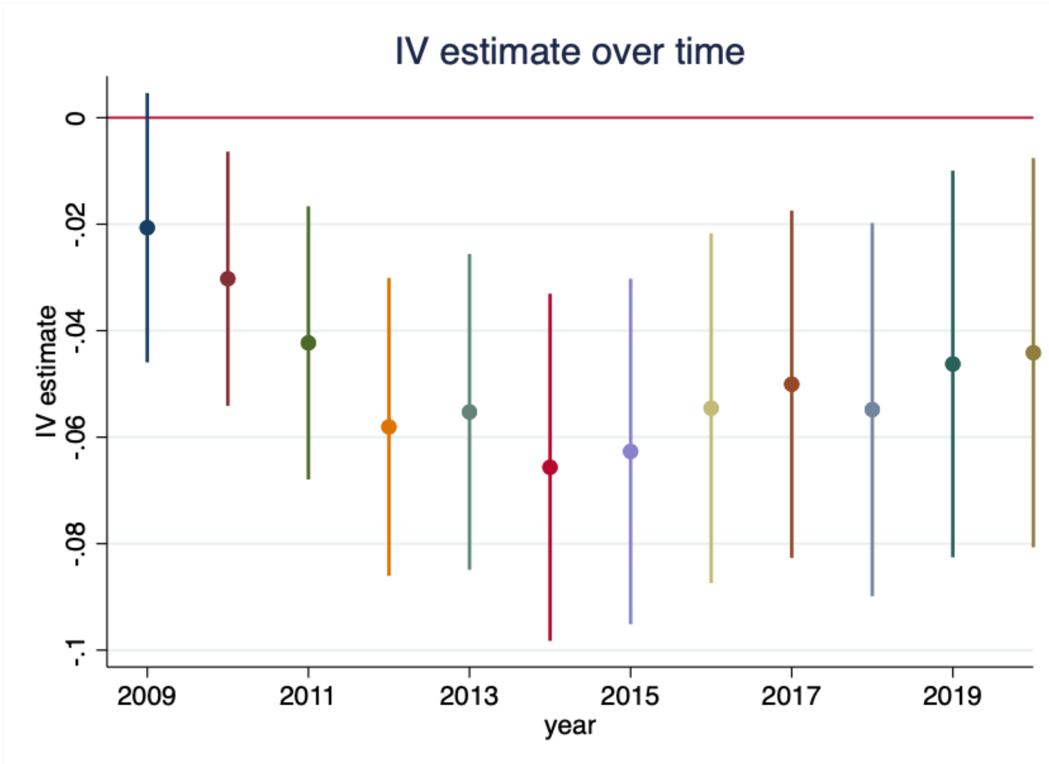

Figure 4: IV estimates from the main specification (equation 3), shown year by year (as described in the text). House price change is standardized. Baseline controls are enumerated in the text. Robust standard errors are used to construct 95% confidence intervals.

4.3 Robustness Checks

The estimated effect of the housing bust on the productivity slowdown is robust to a wide range of alternative specifications and empirical checks. Appendix D provides detailed descriptions and results of these analyses; I summarize the main findings here.

No Trend. Several authors (notably Gordon, 2017) have argued that the U.S. economy has experienced a structural shift, implying that the pre–Great Recession growth trend is no longer a valid benchmark, and that the economy has entered a "new normal" characterized by permanently lower growth rates. To address this concern, I estimate a specification in which the productivity gap is replaced by the log change in productivity between 2010 and 2015.[22] This alternative measure is inherently robust to changes in trend, as it focuses directly on post–Great Recession productivity growth, without requiring any trend extrapolation. However, this is a demanding test for reasons detailed in Section 3. The results remain qualitatively similar and even quantitatively larger in magnitude, suggesting that the

[22]The starting year of 2010 ensures that the measure is not contaminated by the contemporaneous effects of the Great Recession. Results are also similar when using annualized growth rates instead.

housing bust predicts an even greater share of subsequent productivity outcomes under this specification.

Trade Links. An additional concern, raised by Adao et al. (2019), is that an MSA may be indirectly affected by local shocks occurring in other regions with which it maintains trade relationships, in addition to being directly influenced by its own local shocks. A straightforward way to mitigate this issue is to focus the analysis on the non-tradable sector, where economic activity is less exposed to cross-regional spillovers.[23] The estimated effect in the non-tradable sector is significantly larger than in the baseline specification, suggesting that the benchmark results may represent a lower bound relative to an ideal estimate that fully accounts for trade linkages.

Placebo: The Tradable Sector. I conduct a placebo test focusing exclusively on purely tradable industries. Since tradable industries should be largely insulated from local shocks, the response of the productivity gap to local housing busts is expected to be minimal. Consistent with this hypothesis, the estimated effect in the tradable sector is statistically insignificant and, if anything, changes sign.

Housing-Supply Instruments. A natural concern is that the Saiz (2010) housing-supply elasticity may be correlated with broader local characteristics that also shape long-run productivity dynamics, such as migration patterns, density, regional composition, or pre-existing demand fundamentals. To address this concern, I progressively saturate the baseline specification with Davidoff-style controls, including immigrant share, population density, Bartik growth, region fixed effects, and a coastal-state dummy. As shown in Appendix Table E.5, the estimated coefficient remains remarkably stable across these increasingly demanding specifications. This provides reassurance that the main results are not driven by the omission of these long-run local demand correlates. At the same time, these exercises support a more cautious interpretation of the instrument: rather than isolating a purely orthogonal housing-wealth shock, it is more appropriate to view it as capturing a broader housing shock that is correlated with persistent local demand forces.

I also examine robustness to an alternative and independent housing-supply instrument based on the microgeographic estimates of Baum-Snow and Han (2024). Appendix Table E.6 shows that the second-stage coefficient is nearly identical to the baseline estimate, even though the first stage is weaker and the sample is smaller due to the conservative crosswalk required to map the measure into CBSAs. This close similarity in magnitudes provides additional reassurance that the core relationship between the housing bust and subsequent

[23] I follow Mian and Sufi (2014) in classifying industries as non-tradable or tradable and construct a corresponding measure of productivity for each sector.

local productivity dynamics is not specific to the Saiz measure.

Taken together, these results strengthen the interpretation of the housing bust coefficient as capturing a robust relationship between local housing-market distress and subsequent productivity dynamics, rather than an artifact of a particular elasticity measure or a narrow set of controls.

Other Robustness Checks. The results are also robust to several additional specifications and concerns. These include excluding sectors such as finance, insurance, real estate, construction, and public industries from the productivity measure, as well as clustering standard errors at the state level (based on the state in which each MSA has the majority of its population).[24] Together, these analyses further reinforce the robustness of the main findings.

Appendix G complements this section by carefully investigating four potential concerns that naturally arise when studying the relationship between housing market developments and local productivity. These include industry-specific shocks, the aggregate supply shock hypothesis, the business expectations channel, and the business uncertainty channel. Each of these represents a plausible factor that could potentially induce omitted-variable bias in any investigation seeking to relate local housing dynamics and productivity outcomes. While a full treatment of these analyses would make the main text excessively lengthy and technical, I provide detailed empirical examinations in the appendix, showing that none of these factors materially alters the main findings. Readers interested in further details are encouraged to consult this appendix.

5 Mechanisms

Section 4 establishes the existence of a robust, causal relationship between the housing bust and the productivity gap, demonstrating that the magnitude of this effect is economically significant. However, alternative theories could, in principle, also rationalize this relationship. Thus, before turning to the model, the crucial question arises: *what is the mechanism underlying this link?* This section approaches that question empirically in two steps, while remaining careful to distinguish suggestive evidence from causal proof, and to clarify that the mechanisms explored here relate specifically to explaining the empirical result documented in Section 4. Further analyses in Appendix H find little empirical support for alternative channels—such as R&D declines, industry-specific shocks, or local labor composition changes—as

[24] The reason I do not use state-level clustering in the main specification is that a significant fraction of MSAs cross state borders. This exercise assigns each MSA to the state where it has the largest share of its population.

primary drivers of the link between the housing bust and the productivity slowdown.[25]

First, I construct a novel measure of capital expenditures at the MSA level and present evidence indicating that the prolonged slump in investment following the Great Recession was a significant pathway through which the housing bust contributed to the productivity slowdown.

Second, I examine what are arguably the two primary channels through which the housing bust might have triggered this investment slump, thereby affecting the productivity gap. On the one hand, the housing bust led to a decline in consumption expenditures (Mian et al., 2013), which could have persistently depressed corporate investment. On the other hand, the bust may have tightened some firms' access to external finance—for instance, where housing or land served as collateral—potentially constraining investment decisions. I provide empirical evidence suggesting a central role for the former mechanism and a more limited influence of the latter. Importantly, the evidence presented here is intended to illuminate the principal mechanism behind the specific relationship documented in Section 4, and does not claim to rule out the broader relevance of other channels in the macroeconomic context of the Great Recession.

5.1 The Investment Slump

Figure 5 presents the time series of total capital expenditures from the Z.1 Financial Accounts of the United States. The figure shows two series, deflated using two alternative indices: the producer price index for private capital equipment and, for comparison, the GDP implicit price deflator. The data reveal that total capital expenditures peaked around 2006–2007 before falling back to levels last observed in the late 1990s. Moreover, the series did not return to pre–Great Recession levels until roughly 2014–2015 and remained below trend for several years afterward. This raises the question: *was this prolonged investment slump an important pathway through which the housing bust contributed to the productivity gap?*

A significant challenge in incorporating a measure of capital expenditures into a cross-MSA analysis is the lack of publicly available data at this geographic level. While numerous measures of investment exist at the national level—often highly disaggregated across various

[25] Naturally, other channels could also connect the housing bust to the productivity slowdown — such as scarring effects on workers, accelerated declines in business dynamism, or misallocation of resources. These mechanisms may help explain part of the observed relationship and could influence productivity through independent channels as well. This section focuses on providing empirical evidence that the investment slump constituted a major and robust channel linking the housing bust to the productivity slowdown, while acknowledging it is unlikely to be the only one.

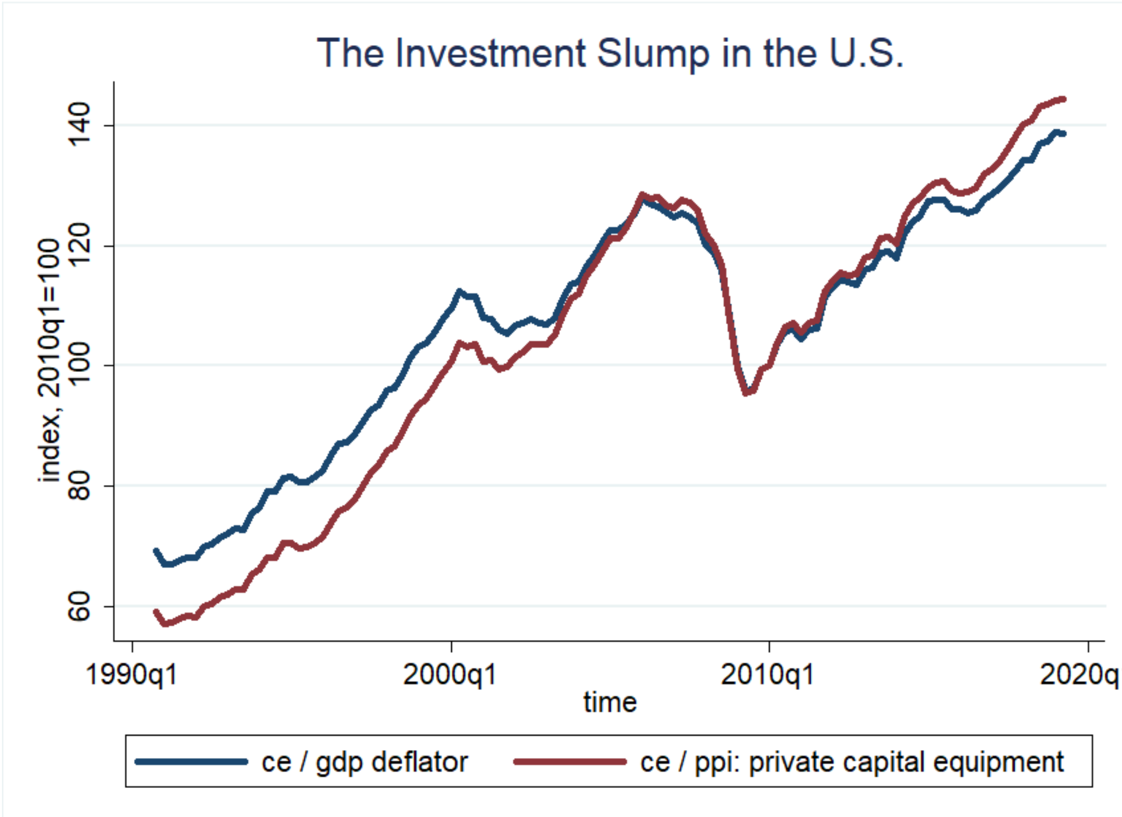

Figure 5: Total capital expenditures, flow. Z.1. Financial Accounts of the United States. The two series use two alternative deflators: the GDP implicit price deflator; and the Producer Price Index for Private Capital Equipment. Capital expenditures include business structures and equipment.

categories—there is, to my knowledge, no comprehensive, publicly available series of capital expenditures beyond the national scale, not even at the state level. The closest available public data at the MSA level comes from the Census of Manufactures, which reports capital expenditures at five-year intervals (2002, 2007, and 2012), but is limited exclusively to the manufacturing sector.

To address this limitation, I construct a novel measure of capital expenditures at the MSA level using confidential microdata from the Census Bureau's Annual Capital Expenditures Survey (ACES).[26] Full methodological details for this construction are provided in Appendix A. This new variable represents, to my knowledge, the first systematic attempt to capture capital expenditures across all sectors at the MSA level, offering an essential tool for exploring regional investment dynamics.

I begin by using the newly constructed series of capital expenditures at the MSA level to compute the 2007–2009 log change in investment. A two-year change is chosen to capture

[26] My benchmark measure of capital expenditures includes expenditures on both new and used equipment, although results are robust to including structures or restricting the measure to only new equipment.

the period spanning the onset and the end of the Great Recession—when most of the decline in investment occurred, as shown in Figure 5—while allowing control for heterogeneous pre-crisis trends through the inclusion of the 2005–2007 investment growth rate and the 2001 investment level as covariates.[27] I then examine whether the housing bust may have contributed to the productivity gap *via* this investment slump.[28] It is important to emphasize that the objective here is not to estimate a complete causal mediation channel, which would require an additional source of exogenous variation for the investment slump independent of housing supply elasticities—a task beyond the scope of this paper. Instead, I focus on providing suggestive evidence to help narrow the set of competing explanations and to inform the subsequent modeling choices.

To this end, I proceed in two steps. First, I extend the IV strategy developed in Section 4 to examine the causal relationship between the housing bust and the investment slump. Specifically, I regress the 2007–2009 log change in investment on the housing bust, instrumenting the housing bust with the measure of housing supply elasticity. Second, I estimate an OLS regression relating the productivity gap to the observed investment slump.

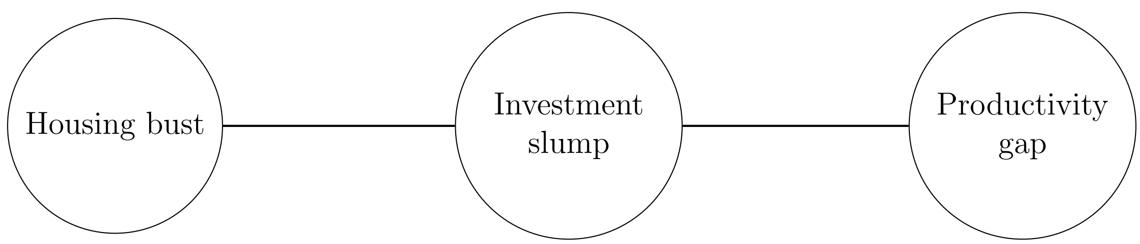

Figure 6: Illustration.

Table 2 reports the results. Column (1) shows the IV regression of the 2007–2009 log change in investment on the housing bust, indicating that a one standard deviation decline in house prices over 2006–2012 is associated with a 48.12 percentage point drop in investment. Column (2) presents the OLS regression of the 2007–2009 log change in investment on the productivity gap, suggesting that a one percentage point increase in the productivity gap is associated with a 2.3 percentage point decline in investment. Equivalently, this implies that every percentage point reduction in investment corresponds to an estimated 0.43 percentage point increase in the productivity gap.

[27] An alternative would be to construct a measure of the investment gap as a deviation from a long-run trend, similar to the approach used for productivity. However, the capital expenditures series is only available starting in 1997. Moreover, changes in MSA personal income are only weakly correlated with changes in capital expenditures, leaving no obvious way to extend the trend backward as done with MSA GDP. Nevertheless, Figure 5 suggests that the 2007–2009 log change in investment can also be interpreted as a proxy for a hypothetical investment gap.

[28] Appendix H presents evidence suggesting a limited role for two mechanisms alternative to the investment slump: the R&D hypothesis and the local labor hypothesis.

Table 2: Capital Expenditures Results

	2007-2009 log-change in investment	
	(1) IV	(2) OLS
House price 2006-2012	0.4812**	
	(0.2041)	
Productivity gap		-2.303**
		(0.9640)
Include baseline controls	Yes	Yes
First Stage F-statistic	36.8	
FS (F-test) p-value	0.0000	
Mean dependent variable	-0.2175	-0.2175
N (Number of MSAs)	150	150

NOTE.—Column (1) corresponds to the IV regression of the 2007–2009 log change in investment on the housing bust, where the housing bust is instrumented by the housing supply elasticity measure. Column (2) displays the results of the OLS regression between the 2007–2009 log change in investment and the productivity gap. House price change is standardized. Baseline controls are enumerated in the text and also include the 2005–2007 log change in investment and the 2001 investment level. Robust standard errors in parentheses. Following CES guidelines, all numbers are rounded to the nearest four significant digits except N, which is rounded to the nearest 50. Significance levels: * 10%, ** 5%, *** 1%.

Two important empirical relationships emerge from these findings. On the one hand, there is evidence linking the investment slump to the productivity gap. This empirical exercise remains agnostic about the specific forces underlying this connection, which the model later represents through two channels: a slowdown in the incorporation of investment-specific technological change (or, equivalently, a deceleration in technology diffusion) and a reduction in the accumulation of capital per worker. On the other hand, there is the link between the housing bust and the investment slump, which is examined in further detail in the following section.

5.2 Consumption versus Credit Supply

This section examines what are arguably the two primary channels through which the housing bust may have contributed to the subsequent slump in investment and, ultimately, to the productivity gap: the decline in consumption spending following the housing bust, and the credit supply hypothesis. The analysis provides empirical evidence suggesting a central role for the former and a more limited one for the latter.[29]

The Consumption Channel

The housing bust led to a significant decline in consumption spending (Mian et al., 2013), which may have persistently discouraged corporate investment. This subsection investigates whether the housing bust might have influenced investment—and, in turn, the productivity gap—*via* the slump in consumption expenditures that followed.

First, I construct a measure of the "consumption gap" by using the proxy for local consumption introduced in Section 2 and computing the log deviation in 2015 relative to its pre-crisis trend.[30] I then examine whether the housing bust could have contributed to the investment slump *via* this consumption gap. To this end, I adopt the same empirical approach as in Section 5.1. First, I extend the IV strategy developed in Section 4 to analyze the causal relationship between the housing bust and the consumption gap. Specifically, I regress the consumption gap on the housing bust, instrumenting the housing bust with the measure of housing supply elasticity. Second, I estimate an OLS regression examining the relationship between the investment slump and the consumption gap.[31]

Table 3 reports the results. Column (1) presents the IV regression of the consumption gap on the housing bust, indicating that a one standard deviation decline in 2006–2012 house prices is associated with an increase in the consumption gap of 7.6 percentage points. Column (2) shows the results from the OLS regression of the 2007–2009 log change in investment on the consumption gap, implying that a one percentage point increase in the consumption gap is associated with a 1.4 percentage point decline in investment.

To complement these findings, the next section presents evidence suggesting only a limited role for what is arguably the most prominent alternative channel: the credit supply

[29] It is important to emphasize that, even if the empirical evidence suggests a limited role for a particular channel in mediating this specific mechanism, such factors could still meaningfully influence the broader productivity slowdown through other pathways.

[30] As in the case of GDP, I use MSA personal income for years prior to 2001 to construct the trend, enabling the series to extend back to 1969. Consistent with the approach for GDP, the rationale is that growth rates of consumption and personal income closely co-move at the state level and are therefore expected to track each other similarly at the MSA level.

[31] As discussed in Section 5.1, the objective here is not to uncover a complete causal mediation channel, which would require an additional source of exogenous variation for the consumption gap independent of housing supply elasticities—a task beyond the scope of this paper. Instead, this analysis aims to provide suggestive evidence to narrow the set of competing explanations and guide the subsequent modeling framework. Accordingly, in principle, the relationship between the consumption gap and the investment slump could be bidirectional. Section 6 develops a general equilibrium model that formalizes and rationalizes these empirical findings.

Table 3: Consumption Expenditures Results

	Consumption gap	log-change in investment
	(1)	(2)
	IV	OLS
House price 2006-2012	-0.07591***	
	(0.02237)	
Consumption gap		-1.397**
		(0.5485)
Include baseline controls	Yes	Yes
First Stage F-statistic	71.77	
FS (F-test) p-value	0.0000	
Mean dependent variable	0.1992	-0.2175
N (Number of MSAs)	229	150

NOTE.—Column 1 reports the IV regression of the consumption gap on the housing bust, where the housing bust is instrumented using the housing supply elasticity measure. Column 2 presents the OLS regression results relating the 2007–2009 log change in investment to the consumption gap. House price changes are standardized. Baseline controls are described in the text, with column (2) including the 2005–2007 log change in investment and the 2001 investment level. Following CES guidelines, figures in column (2) are rounded to the nearest four significant digits except for N, which is rounded to the nearest 50. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

hypothesis.

The Credit Supply Hypothesis

The arguably most important alternative mechanism through which the housing bust could have affected the investment gap—and, ultimately, the productivity gap—is the credit supply hypothesis. This hypothesis posits that firms located in MSAs with larger house price declines might also have faced tighter access to external financing, potentially influencing their productivity dynamics[32]. For instance, this could occur if firms used housing as collateral and saw their borrowing capacity reduced due to falling house prices, or if banks tightened lending standards in areas with severe price declines as a consequence of rising defaults and increased financial risk. Any of these scenarios could affect investment and other firm decisions that ultimately shape productivity outcomes. This section presents several results suggesting that the credit supply hypothesis is unlikely to represent a major channel driving the specific mechanism uncovered in this paper.

On the one hand, business survey evidence from owners, presented in Figure B.4, shows that only 3% of respondents cited financing as their most pressing problem in 2007. Moreover, this share remained low and stable during and after the recession, while concerns about poor sales and regulatory issues rose substantially.

On the other hand, a pure credit supply shock would be expected to affect firms in both tradable and non-tradable sectors. This expectation conflicts with the evidence presented in Section 4.3, which shows significant heterogeneity between these sectors in their responses to the housing bust.

Furthermore, Table C.5 reports the results of two exercises designed to test the credit supply hypothesis more directly.

The first exercise, following Mian and Sufi (2014), uses FDIC data on deposit shares to proxy the geographic diversification of banks operating in each MSA. For each bank, I compute the share of its deposits held in a given MSA. I then average these shares across all banks active in each MSA, weighting by the size of each bank's deposits in that area. The resulting metric, termed the "average share of deposits," is MSA-specific. An MSA in which local banks hold only a small fraction of their deposits is characterized as a "national banking MSA" and should, in principle, be less sensitive to local credit supply conditions than MSAs

[32] There is a large literature analyzing financial factors during the Great Recession from many different perspectives. See, among others, Jermann and Quadrini (2012), Chodorow-Reich (2013), Greenstone et al. (2014), Christiano et al. (2015), Gertler and Gilchrist (2018).

where banks have more locally concentrated deposit bases. Column (2) of Table C.5 reports results for MSAs classified as "national banks" MSAs—those below the median average deposit share—while Column (3) focuses on "local banks" MSAs, above the median. The results show that the estimated effect for "local banks" MSAs is not statistically different from that in the baseline specification in Column (1) and, if anything, appears somewhat smaller. By contrast, the effect is slightly larger for "national banks" MSAs, which are presumed less exposed to local credit conditions. A caveat is that this test does not fully capture borrower-driven credit supply shocks—such as those caused by declining collateral values—which are addressed in the next exercise.

The second exercise, inspired by Stumpner (2019), leverages FDIC data to construct a measure of commercial and industrial loan growth at the MSA level. A key challenge is that loan data are available only in aggregated form at the bank level. To allocate these loans geographically, I use each bank's 2007 deposit share as a fixed proxy for its market presence across MSAs. Summing across banks yields an MSA-level estimate of commercial and industrial loans. From this, I compute a "loans gap" for each MSA in 2015. Column (4) reports results from a regression that includes this loans gap as a control for credit supply shocks, instrumented by the share of long-term illiquid assets held by banks in 1996[33]. The coefficient on the loans gap is statistically insignificant, and the estimated elasticity is not meaningfully different from the baseline result; if anything, it is slightly larger. Of course, this test also has limitations: deposit shares are only an imperfect proxy for banks' lending footprints, and physical distance between banks and firms may play a smaller role in credit relationships than in the past. Nonetheless, empirical studies suggest that borrower-lender proximity remains relevant[34].

Taken together, this evidence suggests that the credit supply channel does not appear to be a dominant mechanism behind the relationship documented in this paper between the housing bust and the local productivity slowdown. However, it is important to emphasize that these findings do not imply that financial frictions were unimportant during the Great Recession more broadly, nor that credit constraints did not play significant roles in other macroeconomic outcomes. Rather, the results here simply indicate that, in the specific context of explaining the causal mechanism identified in this study, the credit supply channel seems to play a limited role.

[33] The logic of this instrument, from Cornett et al. (2011), is that banks with higher shares of illiquid assets were more likely to cut lending during the Great Recession. Because this ratio is relatively persistent over time, I can use 1996 as a pre-boom measure while still achieving a strong first stage.

[34] See Petersen and Rajan (2002) and Brevoort et al. (2010), who find that borrower-lender distances remain low (typically around 5 miles) and have increased only modestly in recent decades.

5.3 Additional Evidence on Investment and Consumption Channels

This section provides additional evidence on the measurement and interpretation of the investment and consumption channels.

Intangible Investment

A natural concern is that the ACES-based measure of capital expenditures may understate investment in intangible assets, which could in turn matter for the estimated link between the investment slump and the productivity gap. It is important to clarify, however, that ACES is not a purely tangible-investment measure. In addition to structures and equipment, the survey covers capitalized computer software within equipment expenditures, including software developed or acquired for internal use under standard capitalization rules. Thus, the relevant concern is not that ACES entirely omits intangible investment, but rather that it does not provide a comprehensive measure of broader intangibles beyond those capitalized in the survey.

To the extent that omitted intangible investment is positively correlated with measured capital expenditures, this undercoverage would tend to attenuate the estimated relationship between the investment slump and the productivity gap. In that sense, the estimates reported in Section 5.1 are best interpreted, if anything, as a lower bound on the importance of the investment channel. The concern would be more serious only if unmeasured intangible investment were both quantitatively important and systematically decoupled from observed ACES investment across areas.

To assess this possibility, I exploit cross-MSA differences in pre-crisis industry composition to construct a proxy for local exposure to intangible-intensive sectors. Specifically, I compute an index based on each MSA's pre-crisis employment shares interacted with national industry shares of intellectual property products (IPP) intensity. This measure captures whether an area was, prior to the Great Recession, more exposed to sectors in which broader intangible investment is likely to be more important. I then re-estimate the main mechanism regressions controlling for this exposure measure. Table F.1 reports the results.

The estimates remain stable throughout. In the investment-slump regressions, the relationship between the housing bust and the investment decline becomes, if anything, slightly stronger once IPP exposure is included. The same is true for the consumption-gap specification and for the baseline housing-bust regression for the productivity gap. This pattern is consistent with the view that undermeasured intangible investment is not driving the main mechanism results and, if anything, that measurement error in observed investment may cause the benchmark estimates to understate the importance of the investment channel. Consistent with this interpretation, Appendix H also shows that adding a proxy for local innovative activity, based on patenting, leaves the main estimates essentially unchanged.

Triangulating the Investment and Consumption Channels

I next provide additional descriptive evidence on the investment and consumption channels to complement the main mechanism results in Sections 5.1 and 5.2. Appendix tables F.2 and F.3 report reduced-form and attenuation exercises that help triangulate the role of these mediators in the relationship between the housing bust and the productivity gap.

Table F.2 shows that the same housing-related shock that predicts a larger productivity gap also predicts both a sharper investment slump and a larger consumption gap. Panel A uses the predicted component of the housing bust obtained from the first-stage relationship with housing-supply elasticity, conditional on the controls, and delivers coefficients very close to the corresponding mechanism estimates in the main text. Panel B instead uses the Saiz elasticity directly. In both cases, the signs and magnitudes are consistent with the mechanism emphasized in the paper: areas more exposed to the housing shock experienced weaker investment, weaker consumption, and worse subsequent productivity outcomes.

Table F.3 provides additional suggestive evidence by introducing the mediators directly into the productivity-gap regressions. When the investment slump is added to the specification, the coefficient on the housing shock declines substantially. When both the investment slump and the consumption gap are included jointly, the coefficient declines further, by roughly forty percent relative to the baseline specification. The same attenuation pattern appears when the Saiz elasticity is used directly instead of the predicted housing shock. This is the pattern one would expect if investment and consumption account for a meaningful part of the relationship between the housing bust and the subsequent productivity gap.

These specifications should not be interpreted as formal causal mediation exercises, since the mediators are themselves endogenous outcomes of the housing shock. Nevertheless, taken together with the baseline mechanism results, they provide additional descriptive support for the view that the investment slump—and the consumption weakness associated with it—constituted important channels through which the housing bust translated into persistent productivity losses.

6 Model

This section builds a quantitative general equilibrium model that is directly informed by the empirical findings in Sections 4 and 5. The goal is to isolate and quantify the specific mechanism through which the housing bust affects productivity via investment dynamics.

The empirical analysis identifies a robust, causal relationship between the housing bust and the national productivity slowdown, mediated primarily through a decline in investment. To better understand these dynamics, the model abstracts from other mechanisms and focuses on this channel, reflecting both empirical evidence and the objective of tractability.

Although the empirical analysis exploits cross-MSA variation, the model is formulated at the national level for tractability. Incorporating regional heterogeneity would substantially increase complexity without altering the core mechanisms relevant for quantifying the effects documented empirically. Moreover, evidence from Section 4.3 suggests that the national estimates likely represent a conservative lower bound to those that might emerge from a multi-region model.[35]

The model starts from a New Keynesian framework with nominal rigidities and includes two key ingredients. The first is a vintage capital structure, where each period introduces a new generation of machines more productive than the previous ones; this feature connects investment and productivity via technology diffusion and capital deepening, consistent with the empirical findings. The second is a putty-clay production structure with irreversible investment, allowing firms to choose capital characteristics at installation but fixing them thereafter; this ingredient captures the persistence observed in productivity dynamics. Together, these elements generate productivity responses to investment fluctuations that align closely with empirical evidence.

The subsections below construct the model, describe the environment for each of its agents, and introduce the key equilibrium conditions.

[35] The empirical part of this paper exploits cross-MSA variation to inform aggregate dynamics. Modeling the economy as multi-region, while conceptually appealing, would significantly increase the complexity of incorporating vintage capital and putty-clay structures into a DSGE framework, without delivering first-order gains for quantifying the mechanisms of interest. Two key dimensions typically modeled in multi-region settings—trade and migration—are indirectly addressed in the empirical robustness analyses: Section 4.3 shows larger effects when restricting to purely non-tradable sectors, and Charles et al. (2018b) documents that job losses during the Great Recession were disproportionately concentrated among low-skill workers, who are less mobile. Taken together, this suggests that the national model likely provides conservative lower-bound estimates relative to a multi-region framework.

6.1 Households

Demographics and Preferences. There is a continuum of households of measure one.

Each household seeks to maximize its expected lifetime utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t S_t \, u(c_t, n_t), \tag{4}$$

where c_t and n_t denote consumption and labor supply, respectively, and where $\beta \in (0,1)$ is the subjective discount factor. The period utility function is separable and takes the form:

$$u(c, n) = \log(c) - \psi \, \frac{n^{1+\eta}}{1+\eta}. \tag{5}$$

The term S_t is a preference shifter, interpreted as a shock to the effective discount factor.[36] The process for S_t evolves as:

$$\log(S_t) = \rho_S \log(S_{t-1}) + \epsilon_S, \tag{6}$$

where ϵ_S is white noise.

Asset Structure. Households can trade a one-period nominal bond, denoted B_t. One unit of this bond costs \$1 in period t and pays \$$(1 + i_t)$ in period $t+1$. There is no government in this economy, so the net supply of bonds is zero. Households also own all shares in the economy.

Representative Household's Problem. The representative household chooses sequences $\{c_t, n_t, B_t\}$ to maximize (4) subject to:

$$P_t c_t + B_t \leq W_t n_t + (1 + i_{t-1}) B_{t-1} + \Pi_t,$$

[36] A previous version of this paper incorporated a richer household structure, building on Iacoviello (2005). That version featured two types of households (borrowers and lenders), housing services valued in utility with a fixed housing stock, and a Kiyotaki and Moore (1997)-style borrowing constraint collateralized by housing. In that framework, a shock to the marginal utility of housing services shifted preferences away from nondurable consumption and leisure toward housing. To simplify the framework—given the relative complexity of the supply side—I adopt here a simpler structure, modeling demand shocks as shocks to the effective discount factor. This captures the demand effects of the housing bust documented by Mian et al. (2013) and Berger et al. (2017). Both versions of the model yield very similar results for our purposes.

where Π_t denotes dividends.

6.2 Final Good Firm

There is a continuum of monopolistically competitive intermediate goods firms, indexed by $i \in [0, 1]$, each producing a differentiated output $Y_t(i)$ and selling it at price $P_t(i)$. These intermediate firms supply their differentiated goods to a single final good firm, which bundles them into a homogeneous output Y_t for consumption and investment purposes. The aggregation technology is given by:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} \, di \right)^{\frac{\epsilon}{\epsilon-1}}, \tag{7}$$

where $\epsilon > 1$ denotes the elasticity of substitution among differentiated goods.

The profit maximization problem of the final good firm is:

$$\max_{\{Y_t(i)\}} \left\{ P_t \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} \, di \right)^{\frac{\epsilon}{\epsilon-1}} - \int_0^1 P_t(i) Y_t(i) \, di \right\}.$$

The corresponding demand schedule for each intermediate firm i and the aggregate price index are:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t, \quad \text{and} \quad P_t = \left(\int_0^1 P_t(i)^{1-\epsilon} \, di \right)^{\frac{1}{1-\epsilon}}.$$

6.3 Intermediate Goods Firms

Environment. As described above, there is a continuum of monopolistically competitive intermediate goods firms, indexed by $i \in [0, 1]$. Each intermediate firm i produces a differentiated output $Y_t(i)$, sold at price $P_t(i)$.

Intermediate firms combine machines and labor to produce their output $Y_t(i)$. In each period t, there exists a distribution of heterogeneous machines, denoted $H_t(x)$, where $x \in [0, \bar{x}]$ represents the efficiency level of a given machine. The labor productivity of a machine with efficiency x is $A_t x$, where A_t captures disembodied aggregate technological change and grows at a gross rate γ_A.

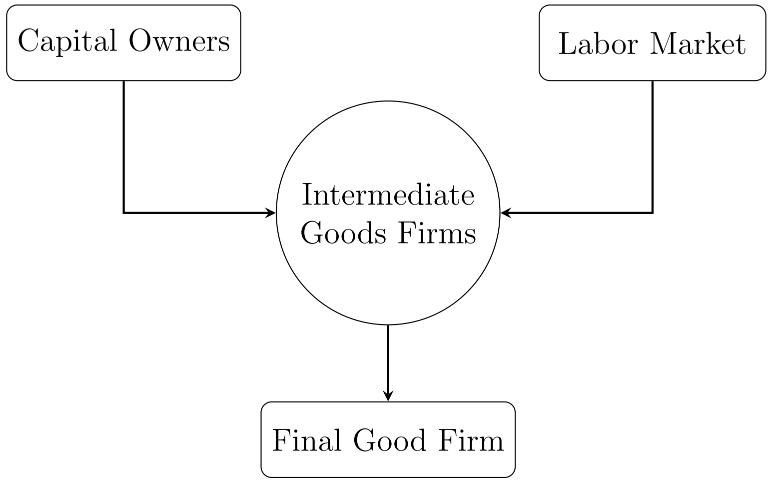

Figure 7: Structure of the supply side of the model.

There is a rental market for machines, where intermediate firms rent capital from capital owners. Let $h_t^i(x)$ denote the quantity of machines with efficiency x rented by firm i in period t, and $R_t(x)$ the corresponding rental price. Both the demand for machines and rental prices are functions of x, rather than scalars.

Without loss of generality, it is assumed that each machine requires one unit of labor to operate, incurring a wage cost W_t for the intermediate firm.[37]

Cost Minimization. Intermediate firms take input prices as given and choose their output prices subject to nominal rigidities. Each firm minimizes costs every period. The cost minimization problem of firm i is:

$$\begin{aligned} \min_{h_t^i(.)} \quad & \int_0^\infty [R_t(x) + W_t]\, h_t^i(x)\, dx \\ \text{s.t.} \quad & \int_0^\infty A_t x\, h_t^i(x)\, dx \geq Y_t(i). \end{aligned} \tag{8}$$

The normalization assumption that each machine requires one worker implies that the operating cost of a machine with efficiency x equals $R_t(x) + W_t$. Multiplying this by the number of machines of that efficiency rented by the firm, and integrating over x, yields total costs.[38] Meanwhile, $A_t x$ represents the labor productivity of a machine of efficiency x, so

[37] The reason this assumption entails no loss of generality will become clear in the next section with the introduction of the production function.

[38] The assumptions below imply that $H_t(x)$ and $h_t^i(x)$ are both continuous. Hence the integral notation. In more general contexts, this can be adapted to sums or other aggregation methods.

$A_t x\, h_t^i(x)$ is the total output produced by those machines. Aggregating across all efficiency levels gives total output for firm i. Formally:

$$Y_t(i) = \int_0^\infty A_t\, x\, h_t^i(x)\, dx \tag{9}$$

and

$$N_t(i) = \int_0^\infty h_t^i(x)\, dx. \tag{10}$$

Profit Maximization. The nominal flow profit for firm i is:

$$\Pi_t(i) = \underbrace{P_t(i)Y_t(i)}_{\text{revenues}} - \underbrace{\int_0^\infty \left[R_t(x) + W_t\right] h_t^i(x)\, dx}_{\text{operating costs}}.$$

Firm i sets its price $P_t(i)$ to maximize the present discounted value of future profits, subject to its demand curve. Prices are sticky due to nominal rigidities à la Calvo (1983): firms can adjust prices each period only with probability $1 - \phi$. Profits are discounted using the representative household's stochastic discount factor.

6.4 Capital Owners

There is a continuum of competitive capital owners, indexed by $\nu \in [0, 1]$. These agents make investment decisions, own capital, and rent it out to intermediate firms.

Technology. The technology side of the model builds on Gilchrist and Williams (2000) and has three essential features: vintage capital, embodied technological change, and a putty-clay production function.

Capital goods take the form of machines. A machine built in period t, denoted Q_t, is characterized by three elements:

1. The vintage level of embodied technology, θ_t. Each period, a new vintage is available, associated with an economywide embodied technology level θ_t, which grows at the rate $(1+g)$. This vintage factor is common across all machines produced in the same period.

2. The capital-labor ratio of the machine, k_t, which is machine-specific and chosen by

capital owners at the time of investment.

3. An idiosyncratic productivity shock, μ_t, which is machine-specific and independently drawn from a continuous distribution with pdf $f_\mu(\cdot)$, support $\mu_t > 0$, and mean $\mathbb{E}[\mu_t] = 1$.[39]

In each period t, a capital owner's investment decision involves two choices: the number of machines to build, Q_t, and the capital-labor ratio per machine, k_t. Total investment is defined as $I_t \coloneqq Q_t k_t$.

Investment is irreversible: once installed, capital cannot be transformed back into consumption goods or reconfigured into other types of capital. Each machine can either be operated or left idle, at no cost. Finally, machines have zero scrap value.

Capital goods take one period to become operational and face an exogenous failure rate δ each period.[40] As noted earlier, each machine requires one worker to operate.[41] The output produced in period t by a machine built in period $t-j$ is:

$$Y_t(A_t\, \mu_{t-j}\, \theta_{t-j}\, k_{t-j}^{\alpha}) = A_t\, \mu_{t-j}\, \theta_{t-j}\, k_{t-j}^{\alpha} \cdot \mathbb{1}[L_t(A_t \mu_{t-j} \theta_{t-j} k_{t-j}^{\alpha}) = 1], \tag{11}$$

where $\alpha \in (0,1)$, $L_t(\cdot)$ denotes the labor employed in the machine, and the indicator function captures the putty-clay constraint.[42] This means that, while machines can be operated or idled, if operated, they must maintain the ex-ante fixed capital-labor ratio. All terms in the

[39] The existence of a continuous idiosyncratic productivity shock greatly simplifies solving the equilibrium, ensuring a smooth and well-behaved investment decision. It also smooths aggregate outcomes and leads to a tractable aggregate production function, despite the putty-clay structure at the micro level. While the model could be solved without this assumption, it would require significantly more complex computation without meaningful conceptual gains.

[40] This implies that each machine has a constant probability δ of failing in any period, capturing physical wear and tear—a major component of observed depreciation rates. This assumption simplifies the model considerably. Setting $\delta = 0$ would not change qualitative results as long as economic growth is positive and some mechanism ensures machine turnover along the balanced growth path.

[41] This can be rationalized in a framework where, *ex ante*, before investment, the production function is Cobb-Douglas with constant returns to scale; but *ex post*, once the machine is built, the production function becomes putty-clay, with the capital-labor ratio fixed at its ex-ante choice. Constant returns imply scale indeterminacy at the machine level. For example, for a given capital-labor ratio, a firm would be indifferent between one machine employing N workers and N machines each employing one worker. To resolve this indeterminacy, machines are normalized to operate with one worker at full capacity, which involves no loss of generality. This is the implicit background structure, but such details are second-order and omitted here for clarity.

[42] As shown later, if a machine is used in equilibrium, it operates at full capacity. Hence, equation (11) can equivalently be written using the traditional Leontief form:

$$Y_t(A_t \mu_{t-j} \theta_{t-j} k_{t-j}^{\alpha}) = A_t \mu_{t-j} \theta_{t-j} k_{t-j}^{\alpha} \cdot \min\,[L_t(\cdot), 1].$$

production function are dated $t-j$, except A_t, which represents the current disembodied technological level. While A_t evolves annually, the embodied terms μ_{t-j}, θ_{t-j}, and k_{t-j}^{α} are fixed once the machine is built.

Define:

$$X_t \coloneqq \mu_t \, \theta_t \, k_t^{\alpha}, \tag{12}$$

which represents the efficiency level of a machine built at time t. Due to the exogenous failure rate δ, machines can be indexed solely by their efficiency level X, regardless of their construction date.[43] Let $H_t(X)$ denote the number of machines with efficiency X available at time t. The evolution of $H_t(X)$ follows:

$$H_{t+1}(X) = (1-\delta) H_t(X) + f_X(X; \bar{X}_t) \, Q_t, \tag{13}$$

where $f_X(X; \bar{X}_t) = \frac{1}{\bar{X}_t} f_\mu \left(\frac{X}{\bar{X}_t} \right)$ is the pdf of X_t, and $\bar{X}_t = \mathbb{E}[X_t] = \theta_t \, k_t^{\alpha}$ is the mean efficiency of newly built machines in period t. Thus, the stock of machines of type X in period $t+1$ equals surviving machines plus new machines constructed in t. The realized efficiency X_t of new machines is distributed around $\bar{X}_t$. The term Q_t denotes the aggregate quantity of new machines built in t.

Because $f_\mu(\cdot)$ is continuous, $H_t(X)$ is a continuous distribution.[44] Together with A_t and total labor, this distribution determines the economy's production possibilities in period t. The dynamics of $H_t(X)$ depend on $f_X(X; \bar{X}_t) \, Q_t$, which in turn is driven by θ_t, k_t, and Q_t. For any given paths of k and Q, $H_t(X)$ tends to shift rightward over time as the vintage technology θ_t improves. Larger growth in k_t raises average capital intensity, and faster growth in Q_t increases the share of newer, more efficient machines associated with higher θ_t.

Investment Decision. Each capital owner ν chooses the quantity of machines $Q_t(\nu)$ and the capital-labor ratio $k_t(\nu)$ to maximize:

$$\mathbb{E}_t \left[- P_t \, Q_t(\nu) \, k_t(\nu) + \left(\sum_{j=1}^{\infty} \beta^j \frac{c_t}{c_{t+j}} R_{t+j}(\mu_t \, \theta_t \, k_t^{\alpha}(\nu)) \right) Q_t(\nu) \right]. \tag{14}$$

[43] This is because under the constant-failure assumption, all machines—irrespective of vintage—face the same probability of surviving into the next period. Thus, the time subscript can be suppressed.

[44] Note, however, that $H_t(X)$ is not a probability density; it aggregates to the total number of machines available in the economy at time t.

Equation (14) describes the present discounted value of expected profits.[45] The first term on the right-hand side captures the cost of investment at time t. The discount factor $\beta^j \frac{c_t}{c_{t+j}}$ reflects the representative household's intertemporal marginal rate of substitution. Capital owners rent their machines to intermediate firms, with $R_{t+j}(X)$ denoting the rental price of a machine of type X in period $t+j$.[46] Note that the embodied characteristics μ_t, θ_t, and k_t are determined in period t and remain fixed for the machine's lifetime.

Because the idiosyncratic shock μ_t is realized only after investment decisions are made, the investment problem is symmetric across capital owners. Hence, in equilibrium, $Q_t(\nu) = Q_t$ and $k_t(\nu) = k_t$ for all $\nu \in [0,1]$.[47] Expectations in (14) are taken over the idiosyncratic shock μ_t and over future values of the discount factor and rental price function.

6.5 Resource Constraint and Monetary Policy

Resource Constraint. The economy's resource constraint is given by:

$$C_t + I_t = Y_t, \tag{15}$$

which follows from aggregating households' budget constraints and firm profits.

Monetary Policy. The monetary authority conducts policy via a Taylor rule:

$$i_t = (1-\rho_i)\, i + \rho_i\, i_{t-1} + (1-\rho_i)\left[\phi_\pi\,(\pi_t - \pi) + \phi_x\left(\log X_t^{gap} - \log X^{gap}\right)\right] + \epsilon_{i,t}, \tag{16}$$

where $\rho_i \in [0,1)$; i and π denote the steady-state nominal interest rate and inflation rate, respectively, and X^{gap} represents the output gap.

[45] Previous drafts included investment adjustment costs à la Christiano et al. (2005). Such costs, whether in Q, k, or both, influence short-run dynamics but have negligible impact on the medium run, which is the focus here. To maintain tractability, I omit them.

[46] In principle, capital owners can choose whether to rent their machines. However, it is always optimal to rent if $R_{t+j}(X) > 0$. A precise formulation could introduce a dummy variable for rental choice or directly write $\max\{R_{t+j}(X), 0\}$. For clarity, I omit this choice in the main text.

[47] This also implies that a capital owner never finds it optimal to choose different capital-labor ratios across machines, i.e., all machines $Q_t(\nu)$ built by owner ν have the same $k_t(\nu)$.

6.6 Equilibrium

I adopt a standard sequence-of-markets equilibrium concept. An equilibrium consists of a sequence of prices

$$\hat{P} = \left\{ \hat{i}_t, \hat{P}_t, \hat{W}_t, \left[\hat{P}_t(i) \right]_{i\in[0,1]}, \hat{R}_t(\cdot) \right\}_{t=0}^{\infty}$$

and allocations

$$\hat{A} = \left\{ \hat{c}_t, \hat{N}_t, \hat{B}_t, \hat{\Pi}_t, \hat{Y}_t, \left[\hat{Y}_t(i), \hat{N}_t(i) \right]_{i\in[0,1]}, \hat{Q}_t, \hat{k}_t, \left[\hat{Q}_t(\nu), \hat{k}_t(\nu) \right]_{\nu\in[0,1]} \right\}_{t=0}^{\infty},$$

such that:

1. Given $\hat{P}$, the allocations $\hat{A}$ solve the optimization problems of households, the final goods firm, intermediate goods firms, and capital owners.
2. All markets clear.
3. Aggregation conditions hold.
4. The price index equations are satisfied.
5. Dynamic consistency conditions hold.

Appendix I provides a variety of equilibrium results and characterizations.

6.7 Calibration

Following Gilchrist and Williams (2000), the idiosyncratic productivity shock μ_t is assumed lognormally distributed:

$$\log \mu_t \sim \mathcal{N}\left(-\frac{1}{2}\sigma_\mu^2,\ \sigma_\mu^2 \right),$$

which is a natural choice since μ_t is strictly positive, and—as shown in Appendix I—this greatly simplifies the analysis.

This section documents the chosen parameter values. Several follow standard conventions in the macroeconomic literature, while others are pinned down to match steady-state features of the relatively stable period of the 1990s.

Time is measured quarterly. Following common practice, I set $\beta = 0.98$, choose an inverse Frisch elasticity $\eta = 1$, and set $\psi = 1$. I select $\rho_S = 0.975$ and $\sigma_S = 1.5$ to capture the persistent dynamics of consumption following 2007.[48]

I adopt the conventional Calvo parameter value of 0.75, consistent with the evidence in Nakamura and Steinsson (2008).

On the production side, I follow standard practice, setting the capital share $\alpha = 0.3$, the depreciation rate $\delta = 0.026$ (equivalent to 10% annually), and the CES aggregator parameter $\epsilon = 6$. Next, I set the growth rate of embodied technology, $g = 0.0023$, to match the average quarterly growth of the inverse real price of investment during the 1990s.[49] To calibrate the growth rate of disembodied technological change, γ_A, I use the fact that the gross growth rate of real GDP along the balanced growth path is:

$$(1 + \gamma_A)^{\frac{1}{1-\alpha}} \times (1 + g)^{\frac{1}{1-\alpha}},$$

and that real GDP grew at an average quarterly rate of 0.82% during the 1990s. This yields $\gamma_A = 0.0037$. Finally, I set the standard deviation of the idiosyncratic productivity shock, σ_μ, to 0.2834 so that the steady-state capital utilization rate matches the observed average capacity utilization during the 1990s (82.39%).

Lastly, to parameterize the Taylor rule, I choose $\rho_i = 0.8$, $\phi_\pi = 1.5$, and $\phi_x = 0$, consistent with empirical estimates for the post-1984 period and prior macroeconomic literature.

6.8 Results

I use the calibrated model to assess the quantitative importance of the mechanism documented in the empirical analysis. Specifically, I perform an experiment in which I shock the preference shifter S_t so that the resulting consumption dynamics closely replicate those observed in the data following the housing bust. For clarity, Table 4 summarizes which quantitative objects are directly targeted in this experiment and which are left untargeted. The key targeted object is the post-2007 consumption shortfall, which disciplines the persistence and magnitude of the preference-shock process. By contrast, the responses of GDP, investment, and labor productivity are not targeted directly and therefore provide non-targeted quantitative implications of the mechanism emphasized in the model. This setup makes it

[48] These values are based on Lo and Rogoff (2015), who document the slow recovery of consumption after the Great Recession due to demand-side factors such as high household debt and significant deleveraging. The persistence of this consumption response is a crucial empirical feature motivating the modeling of the housing bust as a persistent demand shock.

[49] The real price of investment is measured as the ratio of the Producer Price Index: Private Capital Equipment (investment deflator) to the Personal Consumption Expenditures Index (nondurable consumption deflator).

possible to quantify how much of the observed productivity gap can be accounted for by the consumption shortfall alone, once investment and productivity are allowed to respond endogenously within the model.

Moment / object	Role	Discipline
Consumption gap at 40 quarters	Targeted	Matched through (ρ_S, σ_S)
GDP gap at 40 quarters	Non-targeted	Quantitative implication
Investment gap at 40 quarters	Non-targeted	Quantitative implication
Labor productivity gap at 40 quarters	Non-targeted	Quantitative implication

Table 4: Targeted and non-targeted moments in the quantitative experiment. The shock process for S_t is calibrated to reproduce the persistence and magnitude of the post-2007 consumption shortfall. The responses of GDP, investment, and labor productivity are not targeted directly and therefore serve as non-targeted implications of the model.

While I do not interpret preference shocks as the fundamental cause of the housing bust, this approach provides a parsimonious way to capture its demand-side consequences, consistent with empirical evidence in Mian et al. (2013), Berger et al. (2017), and Iacoviello (2005).[50]

Figure 8 presents the impulse response functions from this experiment, expressed as log deviations from trend. Time zero corresponds to 2008, and the final period—40 quarters later—represents the end of 2018. The shock to S_t triggers a decline in consumption. Simultaneously, this drop in consumption is accompanied by reductions in real GDP and labor. Although the cost of capital goods falls, the expected returns on investment decline sharply, leading to a significant contraction in investment and a relatively slow recovery thereafter. This dynamic produces a positive comovement between consumption and investment.

The path of productivity is shaped by three key forces:

1. Even though labor is homogeneous in the model, capital goods are heterogeneous. The downturn in economic activity lowers capacity utilization, which induces a compositional shift: the least productive machines—up to a certain threshold—are taken offline, thereby raising average machine productivity. This generates a temporary increase in aggregate productivity around the time of the shock, a pattern that closely

[50] As discussed previously, a prior version of this paper included a richer household structure with borrowers and lenders, housing services valued in utility, and a borrowing constraint à la Kiyotaki and Moore (1997) using housing as collateral. In that framework, the equivalent experiment consisted of a shock to the marginal utility of housing services, generating a persistent drop in consumption demand. Both versions of the model produce very similar results and are effectively equivalent for the purposes of this analysis, while modeling the shock via S_t offers substantial simplification on the demand side—particularly given the relative complexity of the supply side introduced by vintage capital and putty-clay structures.

mirrors the empirical data.[51] However, this positive effect is soon overtaken by other forces that push productivity downward.

2. The decline in the capital-labor ratio makes newly installed machines less productive, which gradually reduces average machine productivity—particularly as the share of machines installed after the shock grows over time. While this effect eventually fades as the capital-labor ratio converges back to its steady-state level, it exhibits considerable persistence due to the putty-clay structure of production.

3. Each period introduces a new vintage of machines that is more productive than the previous one, continually shifting the productivity distribution rightward. A collapse in the flow of new machines significantly slows this rightward shift, effectively reducing the pace of technological diffusion. The impact of this mechanism grows over time, because the counterfactual stock of new machines that would have been installed in the absence of the shock represents an increasingly larger share of the overall capital stock. Although this effect eventually diminishes as the flow of new machines returns to steady state, it continues to affect the productivity distribution for as long as those machines remain in use.[52]

The model performs reasonably well in replicating several key moments in the data that were not directly targeted in the quantitative experiment. Table 5 compares the model's predictions with the data for the gap relative to trend in 2018. By construction, the model is disciplined to match the consumption shortfall. The fit for GDP, investment, and labor productivity therefore provides a non-targeted assessment of how much of the broader post-recession dynamics can be accounted for by the investment/embodiment mechanism.[53]

[51] In reality, labor-force heterogeneity likely plays an important role, as suggested in Charles et al. (2018b). For simplicity, the model abstracts from this source of heterogeneity, effectively capturing it through capital heterogeneity instead.

[52] It is important to note that this force is somewhat muted under the current calibration. This remains true even in the presence of adjustment costs, which, as discussed earlier, influence short-run dynamics but have minimal effects beyond several years. As shown below, the model generates a decline in investment that, in the medium and long run, still falls short of the magnitude observed in the data (see Table 5). Introducing mechanisms that generate more sizable and persistent investment dynamics—such as adaptive expectations—would amplify this channel and further lower productivity. However, that lies beyond the scope of this model, which likely produces conservative estimates of the productivity effects compared to what might emerge if investment dynamics matched the data more closely.

[53] Note that the method for computing these gaps differs somewhat from the approach used in the empirical analysis, which explains minor discrepancies in levels. In Figure 1, the national trend is extrapolated from 2010 to illustrate deviations in productivity following the Great Recession. In the main empirical analysis, MSA trends are extrapolated from 2007 to remain conservative and avoid distortions from the crisis itself. In contrast, the model considers 2008 as the shock year, so the moment comparison extrapolates trends from that point forward, taking pre-2008 steady-state levels as the baseline. Small differences in levels arise

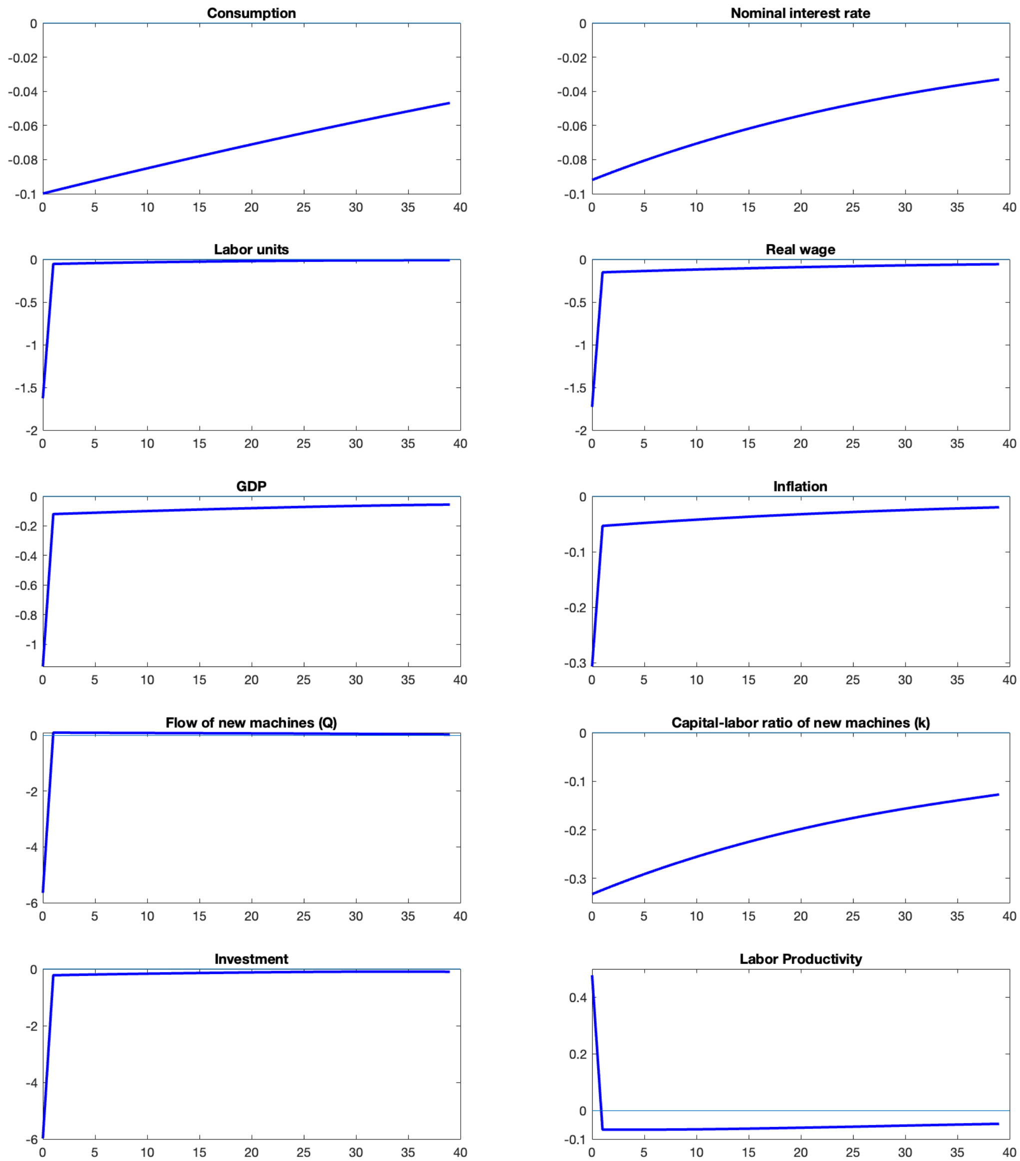

Figure 8: Impulse response functions from the S_t shock experiment.

The figures indicate the percentage growth each series would need in 2018 to return to its extrapolated trend (from 2008 to 2018). Appendix J reports the corresponding dynamic fit at 3-, 6-, and 10-year horizons, showing that the model tracks the medium-run behavior of consumption and investment reasonably well, but underpredicts the persistence of the investment slump at longer horizons, which helps explain why the implied long-run productivity

primarily from the pronounced productivity increase around 2008–2009 and the fact that the MSA analysis uses regional data, while both Figure 1 and the model exercise rely on national data.

	Model	Data
Consumption	0.0468	0.0450
GDP	0.0562	0.0727
Investment	0.0993	0.1275
Labor productivity	0.0464	0.0933

Table 5: Consumption is real personal consumption expenditures per capita (goods: nondurable goods, chained 2012 dollars). GDP is real gross domestic product per capita (chained 2012 dollars). Investment is the deflated flow of total capital expenditures for all sectors (Z.1. Financial Accounts of the United States). Labor productivity is nonfarm business sector real output per hour of all persons (chained 2012 dollars). The data numbers come from computing the pre-2008 average growth rates, extrapolating that growth from 2008 onward, and taking the difference between the logs of this extrapolated trend component and the actual series in 2018. The model numbers are the IRFs after 40 quarters (10 years). Each number is expressed in percentage points and indicates how much the series would need to grow in 2018 to catch up with trend (or with the steady state). Consumption is the targeted moment in the quantitative experiment; the GDP, investment, and labor productivity gaps are non-targeted model implications.

gap remains conservative relative to the data.

The model comes reasonably close for GDP, despite not being calibrated to target this moment directly. However, it somewhat underpredicts the magnitude of the investment decline, suggesting that the quantitative estimate of the productivity gap is likely conservative. The figures for labor productivity imply that the model can account for roughly 50% of the observed productivity gap. Appendix J also reports sensitivity exercises showing that this quantitative conclusion is robust across alternative calibrations. These exercises indicate that the persistence of the productivity response is shaped mainly by depreciation and nominal price rigidity, while adjustment costs affect the magnitude of the response but are not essential for the mechanism to operate. The appendix also compares the benchmark experiment with an alternative monetary disturbance and shows that, although the two shocks are not fully observationally equivalent, they generate very similar productivity implications once they produce a similar investment slump. Importantly, these results remain broadly similar regardless of the benchmark year chosen for computing the gaps.

7 Conclusion

The United States experienced a decade-long slowdown in labor productivity, which may account for an accumulated 16% gap in GDP relative to pre-crisis trends (Syverson, 2017). While this fact is widely recognized, its underlying causes remain an open puzzle. What is less well-known is that the productivity slowdown was not uniform across the U.S.: some MSAs, such as Los Angeles, CA, and Miami, FL, faced severe declines, whereas others, like Pittsburgh, PA, and San Antonio, TX, experienced only mild slowdowns—or even grew above trend. This rich spatial heterogeneity provides a valuable laboratory to investigate the link between the 2006-2012 housing bust and the subsequent productivity slowdown.

Exploiting this geographic variation through an IV strategy, I find evidence consistent with a causal relationship between house price declines and regional productivity slowdowns. My estimates indicate that a one standard deviation drop in 2006-2012 house prices increases the productivity gap by 6.3 percentage points—over half the average gap of 11.53%.

An empirical exploration of potential channels suggests that the housing bust affected productivity mainly through the protracted collapse in consumption spending, which, in turn, triggered a persistent slump in investment. To capture this mechanism precisely, I construct a new measure of investment at the MSA level using confidential Census microdata—a contribution that offers novel insights into local capital dynamics. By contrast, I document only a limited role for credit supply constraints or other alternative mechanisms.

To further quantify this mechanism, I develop a general equilibrium model that is explicitly disciplined by the empirical findings. The model generates a slump in productivity despite continued growth in the technological frontier—a feature that aligns with evidence showing stable R&D expenditure and patent activity during the period. Quantitatively, the model suggests that roughly half of the observed productivity slowdown may be attributable to the housing bust. This result closely matches the partial equilibrium back-of-the-envelope calculations from the empirical analysis and stands in contrast to views attributing the productivity slowdown solely to secular forces.

The mechanism documented here illustrates the so-called "inverse Say's Law": the idea that weak demand, when persistent, can impair potential output. This has important policy implications, as periods of chronically low demand risk being misinterpreted as an economy operating near its potential—leading to policies that could exacerbate, rather than resolve, the underlying problem.

Finally, these findings contribute to the broader policy debate on preventing housing bubbles. Housing busts can trigger recessions marked by heavy household debt burdens and severe asset price depreciation, leading to prolonged weakness in consumption, inflation, investment, productivity, and ultimately wages. Such dynamics can set in motion adverse feedback loops that further weaken demand and economic performance.

References

Adao, R., C. Arkolakis, and F. Esposito (2019). *Spatial linkages, global shocks, and local labor markets: Theory and evidence*. Tech. rep. National Bureau of Economic Research.

Alon, T., D. Berger, R. Dent, and B. Pugsley (2018). "Older and slower: The startup deficit's lasting effects on aggregate productivity growth". *Journal of Monetary Economics* 93, pp. 68–85.

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019). "Endogenous technology adoption and R&D as sources of business cycle persistence". *American Economic Journal: Macroeconomics* 11.3, pp. 67–110.

Baum-Snow, N. and L. Han (2024). "The Microgeography of Housing Supply". *Journal of Political Economy* 132.6, pp. 1897–1946. DOI: 10.1086/728110.

Beraja, M., E. Hurst, and J. Ospina (2016). *The aggregate implications of regional business cycles*. Tech. rep. National Bureau of Economic Research.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017). "House prices and consumer spending". *The Review of Economic Studies* 85.3, pp. 1502–1542.

Bhattarai, S., F. Schwartzman, and C. Yang (2021). "Local scars of the US housing crisis". *Journal of Monetary Economics* 119, pp. 40–57.

Brevoort, K. P., J. D. Wolken, and J. A. Holmes (2010). "Distance still matters: the information revolution in small business lending and the persistent role of location, 1993-2003".

Brynjolfsson, E. and A. McAfee (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Byrne, D. M., J. G. Fernald, and M. B. Reinsdorf (2016). "Does the United States have a productivity slowdown or a measurement problem?" *Brookings Papers on Economic Activity* 2016.1, pp. 109–182.

Byrne, D. M., S. D. Oliner, and D. E. Sichel (2013). "Is the information technology revolution over?"

Byrne, D. M., S. D. Oliner, and D. E. Sichel (2018). “How fast are semiconductor prices falling?” *Review of Income and Wealth* 64.3, pp. 679–702.

Calvo, G. A. (1983). *Staggered Prices in a Utility-Maximizing Framework*. Tech. rep. 3, pp. 383–398.

Cardarelli, M. R. and L. Lusinyan (2015). *US Total factor productivity slowdown: Evidence from the US States*. 15-116. International Monetary Fund.

Cette, G., J. Fernald, and B. Mojon (2015). “The Pre-Global-Financial-Crisis Slowdown in Productivity”. *Unpublished Paper*.

Charles, K. K., E. Hurst, and M. J. Notowidigdo (2018a). “Housing Booms and Busts, Labor Market Opportunities, and College Attendance”. *American Economic Review* 108.10, pp. 2947–94.

– (2018b). “Housing booms, manufacturing decline and labour market outcomes”. *The Economic Journal* 129.617, pp. 209–248.

Chodorow-Reich, G. (2013). “The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis”. *The Quarterly Journal of Economics* 129.1, pp. 1–59.

– (2019). “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?” *American Economic Journal: Economic Policy* 11.2, pp. 1–34.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015). “Understanding the great recession”. *American Economic Journal: Macroeconomics* 7.1, pp. 110–67.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). “Nominal rigidities and the dynamic effects of a shock to monetary policy”. *Journal of political Economy* 113.1, pp. 1–45.

Cornett, M. M., J. J. McNutt, P. E. Strahan, and H. Tehranian (2011). “Liquidity risk management and credit supply in the financial crisis”. *Journal of financial economics* 101.2, pp. 297–312.

Cowen, T. (2011). *The great stagnation: How America ate all the low-hanging fruit of modern history, got sick, and will (eventually) feel better: A Penguin eSpecial from Dutton.* Penguin.

Davidoff, T. et al. (2016). "Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors". *Critical Finance Review* 5.2, pp. 177–206.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2017). "Declining dynamism, allocative efficiency, and the productivity slowdown". *American Economic Review* 107.5, pp. 322–26.

Feldstein, M. (2015). "The US underestimates growth". *Wall Street Journal* 18, p. 2015.

Fernald, J. G. (2015). "Productivity and Potential Output before, during, and after the Great Recession". *NBER macroeconomics annual* 29.1, pp. 1–51.

Ferreira, F. and J. Gyourko (2011). *Anatomy of the beginning of the housing boom: US neighborhoods and metropolitan areas, 1993-2009.* Tech. rep. National Bureau of Economic Research.

Fishback, P. V., W. C. Horrace, and S. Kantor (2005). "Did New Deal grant programs stimulate local economies? A study of Federal grants and retail sales during the Great Depression". *The Journal of Economic History* 65.1, pp. 36–71.

Garrett, T. A., R. Hernández-Murillo, M. T. Owyang, et al. (2005). "Does consumer sentiment predict regional consumption". *Federal reserve Bank of st. Louis review* 87.2, pp. 123–135.

Gertler, M. and S. Gilchrist (2018). "What happened: Financial factors in the great recession". *Journal of Economic Perspectives* 32.3, pp. 3–30.

Gilchrist, S. and J. C. Williams (2000). "Putty-clay and investment: a business cycle analysis". *Journal of Political Economy* 108.5, pp. 928–960.

Giroud, X. and H. M. Mueller (2017). "Firm leverage, consumer demand, and employment losses during the Great Recession". *The Quarterly Journal of Economics* 132.1, pp. 271–316.

Gordon, R. J. (2017). *The rise and fall of American growth: The US standard of living since the civil war*. Vol. 70. Princeton University Press.

Greenstone, M., A. Mas, and H.-L. Nguyen (2014). *Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and 'normal'economic times*. Tech. rep. National Bureau of Economic Research.

Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson (2018). *Housing wealth effects: The long view*. Tech. rep. National Bureau of Economic Research.

Guvenen, F., R. J. Mataloni Jr, D. G. Rassier, and K. J. Ruhl (2017). *Offshore profit shifting and domestic productivity measurement*. Tech. rep. National Bureau of Economic Research.

Hall, R. E. (2015). "Quantifying the lasting harm to the US economy from the financial crisis". *NBER Macroeconomics Annual* 29.1, pp. 71–128.

Iacoviello, M. (2005). "House prices, borrowing constraints, and monetary policy in the business cycle". *American economic review* 95.3, pp. 739–764.

Jermann, U. and V. Quadrini (2012). "Macroeconomic effects of financial shocks". *American Economic Review* 102.1, pp. 238–71.

Kaplan, G., K. Mitman, and G. L. Violante (2016). *Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data*. Tech. rep. National Bureau of Economic Research.

Kiyotaki, N. and J. Moore (1997). "Credit cycles". *Journal of political economy* 105.2, pp. 211–248.

Lo, S. and K. Rogoff (2015). "Secular stagnation, debt overhang and other rationales for sluggish growth, six years on".

Mian, A., K. Rao, and A. Sufi (2013). "Household balance sheets, consumption, and the economic slump". *The Quarterly Journal of Economics* 128.4, pp. 1687–1726.

Mian, A. and A. Sufi (2009). "The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis". *The Quarterly Journal of Economics* 124.4, pp. 1449–1496.

Mian, A. and A. Sufi (2011). "House prices, home equity-based borrowing, and the US household leverage crisis". *American Economic Review* 101.5, pp. 2132–56.

– (2014). "What explains the 2007–2009 drop in employment?" *Econometrica* 82.6, pp. 2197–2223.

Mokyr, J. (2014a). "Secular stagnation? Not in your life". *Secular Stagnation: Facts, Causes and Cures* 83.

– (2014b). "The Next Age of Invention: Technology's future is brighter than pessimists allow." *City journal* 24.1, pp. 12–21.

Moretti, E. (2010). *Local labor markets.* Tech. rep. National Bureau of Economic Research.

Nakamura, E. and J. Steinsson (2014). "Fiscal stimulus in a monetary union: Evidence from US regions". *American Economic Review* 104.3, pp. 753–92.

Nakamura, E. and J. Steinsson (2008). "Five facts about prices: A reevaluation of menu cost models". *The Quarterly Journal of Economics* 123.4, pp. 1415–1464.

Nakamura, L. and R. Soloveichik (2016). "Capturing the Productivity Impact of the 'Free'Apps and Other Online Media". *Economics of Digitization conference, National Bureau of Economic Research.* Vol. 4.

Petersen, M. A. and R. G. Rajan (2002). "Does distance still matter? The information revolution in small business lending". *The journal of Finance* 57.6, pp. 2533–2570.

Saiz, A. (2010). "The geographic determinants of housing supply". *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.

Smith, N. (2015). "The Internet's hidden wealth". *Bloomberg View, June* 6.

Stroebel, J. and J. Vavra (2019). "House prices, local demand, and retail prices". *Journal of Political Economy* 127.3, pp. 1391–1436.

Stumpner, S. (2019). "Trade and the geographic spread of the great recession". *Journal of International Economics* 119, pp. 169–180.

Summers, L. (2016). "The age of secular stagnation: What it is and what to do about it". *Foreign Aff.* 95, p. 2.

– (2017). "Secular stagnation even truer today". *Wall Street Journal* 25.

Syverson, C. (2017). "Challenges to mismeasurement explanations for the US productivity slowdown". *Journal of Economic Perspectives* 31.2, pp. 165–86.

Yagan, D. (2019). "Employment hysteresis from the great recession". *Journal of Political Economy* 127.5, pp. 000–000.

Appendices

A Data

This section elaborates on the data sources used in this paper. The unit of observation is the MSA, defined by the U.S. Office of Management and Budget as a contiguous area with at least 50,000 people, delineated based on a central urban cluster. According to the 2010 Census, 83.6% of Americans live in MSAs. The rationale for choosing MSAs is twofold: they constitute local labor markets (Moretti, 2010)—unlike counties—and they are the smallest geographic delineation for which it is currently possible to measure labor productivity.

Labor Productivity

Labor productivity is measured as real GDP per worker and comes from the Bureau of Economic Analysis (BEA).[54]

Given that MSA-level GDP data is only available from 2001 onward, I also use personal income (from the BEA) to construct trends for earlier years, allowing me to go back to 1969.[55] Personal income is deflated using the national CPI.

	GDP	Consumption
Mean	0.6540	0.5354
Std. Dev.	0.1169	0.1523
Median	0.6790	0.5478
25th percentile	0.6057	0.4313
10th percentile	0.4940	0.3695

Table A.1: Column 'GDP' reports descriptive statistics across states for the R^2 from within-state, across-year (1963–2020) regressions of annual GDP growth on annual personal income growth (one regression per state). Column 'Consumption' does the same exercise for the annual growth of the constructed consumption proxy versus actual consumption expenditures (1997–2020).

Table A.1, column 'GDP', shows that annual personal income growth comoves closely with annual GDP growth at the state level over the entire sample period. The state-level regres-

[54] To my knowledge, there are only two sources of hours-worked data at the MSA level: the American Community Survey (ACS) and the Current Employment Statistics (CES). The problem with the ACS measure is that it is only available for a subset of MSAs, as its smallest geographic unit is the Public Use Microdata Area (PUMA), which covers at least 100,000 people and only overlaps with MSAs above that size. Thus, PUMAs do not fully encompass MSAs with 50,000 to 100,000 people. The CES measure of hours, in turn, is only available after 2006. These limitations motivate taking output per worker as the baseline definition. Also, there is no measure of GNP below the national level.

[55] Section D shows that the results are robust to various alternative approaches.

sions yield coefficients that are all statistically significant and close to one, while the table reports descriptive statistics for the R^2 values of those regressions.

House Prices

The house price indexes come from the Federal Housing Finance Agency (FHFA). I compute the log difference between 2006 and 2012 and standardize it to facilitate interpretation. The cross-MSA standard deviation is 19.90, while the mean is -18.45.

Consumption

There is no direct measure of consumption expenditures at the MSA level (the finest geographic aggregation available is the state). Some regional consumption proxies used in prior literature include car purchases and credit card data (see, e.g., Mian et al., 2013), but these data are proprietary and unavailable to me. Publicly available alternatives include retail sales (Fishback et al., 2005), which are only reported for disaggregated geographies during Economic Census years (every five years), and retail sales tax data, which are noisy and only cover a subset of regions (Garrett et al., 2005). My main measure of MSA-level consumption expenditures extrapolates national retail sales to MSAs using MSA-level retail GDP from the BEA as weights.

Table A.1, column 'Consumption', shows that the annual growth rate of this measure, when aggregated to the state level, comoves closely with the actual growth rate of consumption expenditures at the state level.[56] As before, state-level regressions across years yield statistically significant coefficients close to one, with the table reporting descriptive statistics for the R^2 values.

Capital Expenditures

There is, to my knowledge, no publicly available—and likely not even proprietary—comprehensive measure of capital expenditures at any subnational geographic level finer than the national aggregate. The closest existing dataset is the Census of Manufacturers' capital expenditures statistics, which is only available every five years and restricted to manufacturing. Therefore, I directly construct a new measure of capital expenditures (at both MSA and state levels)

[56] The rationale for proxying local consumption expenditures with retail sales, GDP, or employment—and for why these variables track consumption so closely across many independent studies—is that retail is an intermediate input to household consumption: one must purchase goods to consume them. See also Table A.1 and, e.g., Kaplan et al., 2016 and Guren et al., 2018.

using confidential Census Bureau microdata from the Annual Capital Expenditures Survey (ACES).

The ACES is an annual Census Bureau survey of about 50,000 employer firms, designed to provide comprehensive information on capital expenditures by U.S. domestic non-farm businesses. It reports expenditures on new and used structures and equipment, as well as detailed industry breakdowns. The equipment category includes, among other items, machinery, computers, transportation equipment, and capitalized computer software. Accordingly, while ACES captures some intangible-type investment through capitalized software, it does not provide a comprehensive measure of broader intangible investment. My benchmark measure focuses on equipment expenditures, although the results are also robust to including structures.

Investment is localized using establishment location rather than firm headquarters. Starting from firm-level capital expenditure data, I allocate investment geographically to MSAs (or states) by linking firms to their establishments in the LBD and then aggregating across establishments. I merge the ACES (whose unit of observation is the firm) with the Longitudinal Business Database (LBD), which contains the universe of establishments, including data on firm ownership, establishment locations, industry classifications, and economic details such as employment and payroll. Once firms are matched to establishments (and, in turn, establishments to MSAs or states), I distinguish between:

1. Single-unit firms and firms whose establishments are all located in the same MSA (or state).[57]

2. Firms with establishments spread across different MSAs (or states).

The first group can be directly assigned to its corresponding MSA, while the second requires further processing, as explained below.

For multi-establishment firms operating across multiple MSAs (or states), I proceed as follows. First, I restrict the sample to:

- Firms with at least 95% of their business concentrated in the same (2-digit) NAICS industry.[58] This drops firms with establishments that attend different lines of business and, hence, are hard to compare.

[57]Single-unit firms are firms with only one establishment. They are 96.07% of all firms in 2015, gathering 48.15% of employment. Conversely, multi-unit firms are 3.93% of firms and 51.85% of employment.

[58]The reason for using a 95% cutoff instead of 100% is that it's common for multi-establishment firms to have nearly all establishments in one industry, except for a few small outliers. Results are robust to thresholds of 100% or even 90%.

- Firms with fewer than 50 establishments.[59] This drops large firms with many establishments.

These restrictions apply to the analytical sample used to localize firm-level ACES investment through the LBD establishment structure. The resulting sample covers 99.49% of all non-farm employer businesses in the LBD, representing 65.54% of employment in 2015. Within multi-unit firms, it captures 86.94% of firms and 33.54% of employment.

Capital expenditures at the firm level are then distributed across establishments according to each establishment's share of total payroll within the firm.[60] Finally, I aggregate up to the MSA (or state) level using the weighted median. I restrict the analysis to cells (MSA $\times$ year) with at least 200 observations, covering all states and most MSAs with populations above 150,000. All figures are deflated using the national CPI, and investment growth is computed in log differences.

Two main concerns arise with this approach. First, some capital expenditures may be firm-wide in nature, serving all establishments collectively rather than being physically tied to a single location. Even so, such investments contribute economically to each establishment's activity. Second, payroll might not perfectly reflect the capital intensity of each establishment, as firms could have establishments with differing capital-labor ratios. While this is possible, it's less likely when comparing establishments within the same firm and NAICS classification.

The first two columns of Table A.2 compare our sample with the universe of firms in the LBD. Our sample includes 99.49% of firms and 65.54% of employment. Thus, while relatively large firms are excluded, the sample retains most employment, implying that the average firm in the sample is smaller but representative. The employment percentiles remain similar, suggesting a good match to the broader distribution once outliers are excluded. Payroll-per-employee figures also closely align, indicating that our sample is representative in terms of labor productivity.

Given that all single-unit firms are included in the sample, but not all multi-unit firms, the last two columns of Table A.2 further compare multi-unit firms in our sample to those in the full LBD universe. Again, the average firm size in our sample is smaller because the largest firms are excluded, but employment percentiles remain similar. Total payroll comparisons are limited by size differences, but payroll per employee remains very close, reinforcing the

[59] Stricter thresholds, such as fewer than 20 establishments, yield virtually identical results.

[60] Using employment count yields very similar results. Payroll has two advantages: it better captures worker heterogeneity within firms and relies on IRS-sourced data that is more reliable and less prone to missing values.

sample's representativeness.

Table A.2: Sample Statistics

	Overall Statistics		Within multi-unit firms	
	Sample	Universe	Sample	Universe
Mean employment	18.87	28.65	145.8	378
s.d. employment	206.4	963.5	843.9	4801
p-25 employment	1	1	14	15
Median employment	4	4	36	44
p-75 employment	10	10	101	135
Mean payroll	867.4	1416	8316	21310
s.d. payroll	13690	46350	64650	230600
p-25 payroll	37	37	458	535
Median payroll	112	113	1452	1780
p-75 payroll	349	356	4453	6000
Mean pay/emp	40080	40150	52010	52240
s.d. pay/emp	76980	76930	72180	71340
p-25 pay/emp	16000	16000	24170	24630
Median pay/emp	29000	29100	41400	42000
p-75 pay/emp	48000	48000	62550	63220
% firms	99.49	100	86.94	100
% employment	65.54	100	33.54	100

NOTE.—Mean, standard deviation, 25th percentile, median, and 75th percentile for employment (March 12 employment), annual payroll (in $1,000), and annual payroll per employee (dollars per worker). The first two columns compare the sample to the entire LBD universe. The last two columns compare multi-unit firms in the sample with all multi-unit firms in the LBD. The LBD universe covers all non-farm businesses with employees. % firms refers to the share of the universe represented by the sample; % employment reflects the combined employment of firms in the sample relative to the total universe. All data are from 2015. Per CES guidelines, figures are rounded to four significant digits, and percentiles are pseudo-percentiles.

Table A.3 analyzes the properties of our main dependent variable, investment growth.[61] The top of the table reports results from three separate regressions: investment growth on real GDP growth, employment growth, and personal income growth. These regressions are not intended to establish causality but to explore joint dynamics. Two key findings emerge: (i) investment growth is strongly positively correlated with real GDP growth, employment growth, and personal income growth, all with high statistical significance; and (ii) as ex-

[61] Given the high volatility of investment, I follow an approach inspired by Tornqvist et al. (1985), which has become standard after Doms and Dunne (1998) and in studies of firm dynamics following Davis et al. (1998). I compute investment growth as the growth rate relative to the average investment of the previous two periods. This reduces noise and increases statistical power.

pected, investment growth is substantially more volatile than these variables. For instance, a one-percentage-point increase in real GDP growth is associated with a 2.8pp increase in investment growth, rising to 6.2pp and 3.6pp for employment and personal income growth, respectively. The bottom panel reports summary statistics, revealing that investment growth has a mean and median close to zero but considerable volatility, with the 25th and 75th percentiles around -50% and +50%, and a standard deviation near 110%. Thus, investment growth displays a relatively symmetric distribution centered around zero but characterized by high volatility.

Table A.3: Exploring the Dependent Variable

	Investment growth		
	(1)	(2)	(3)
Real GDP growth	2.821***		
	(0.9776)		
Employment growth		6.243***	
		(1.781)	
Personal income growth			3.568***
			(1.273)
p-25 dependent variable	-0.4972	-0.4972	-0.4972
Median dependent variable	-0.0372	-0.0372	-0.0372
p-75 dependent variable	0.4645	0.4645	0.4645
N	1000	1000	1000

NOTE.—Standard errors are clustered at the MSA level. Per CES guidelines, figures are rounded to four significant digits except for N, which is rounded to the nearest 100 if above 1,000 or to the nearest 50 if below 1,000. Percentiles are pseudo-percentiles. Significance levels: * 10%, ** 5%, *** 1%.

Other Variables

Other controls include population and the share of adults with a college degree from the decennial Census; sectoral employment shares, labor productivity, and wages from the BEA; housing-supply elasticity measures from Saiz (2010); and FDIC data on deposits and commercial and industrial loans.

To address the concern that the Saiz elasticity may be correlated with broader local demand fundamentals, Appendix Table E.5 adds a set of Davidoff-style controls. Immigrant share and population density are constructed from the 2000 Census and aggregated to the CBSA level. Bartik growth is measured over 2001–2006 using local pre-bust industry shares interacted with national industry growth. I also include region fixed effects based on the four Census regions and a coastal-state indicator for metropolitan areas located in Atlantic, Pacific, or Gulf Coast states.

Appendix Table E.6 reports robustness to an alternative housing-supply instrument based on the microgeographic housing-supply estimates of Baum-Snow and Han (2024). Because their measure is defined at a finer geographic level than CBSAs, I map their regions into my CBSA sample using metro names and state identifiers, supplemented by conservative manual reconciliation in cases with transparent naming differences. I retain only clean and unambiguous matches, which reduces the sample size relative to the baseline specification.

Appendix Table F.1 reports an additional exercise aimed at bounding the role of undermeasured intangible investment in the ACES-based capital expenditures series. I construct a proxy for local exposure to intangible-intensive industries by interacting each MSA's pre-crisis industry employment shares with national industry-level measures of intellectual property products (IPP) intensity in 2001. This variable captures whether an area was more exposed, prior to the Great Recession, to sectors in which broader intangible investment is likely to be more important. Re-estimating the main investment, consumption, and baseline productivity specifications with this additional control leaves the results largely unchanged. An alternative version based on average IPP exposure over 1998–2005 yields virtually identical results.

B Additional Figures

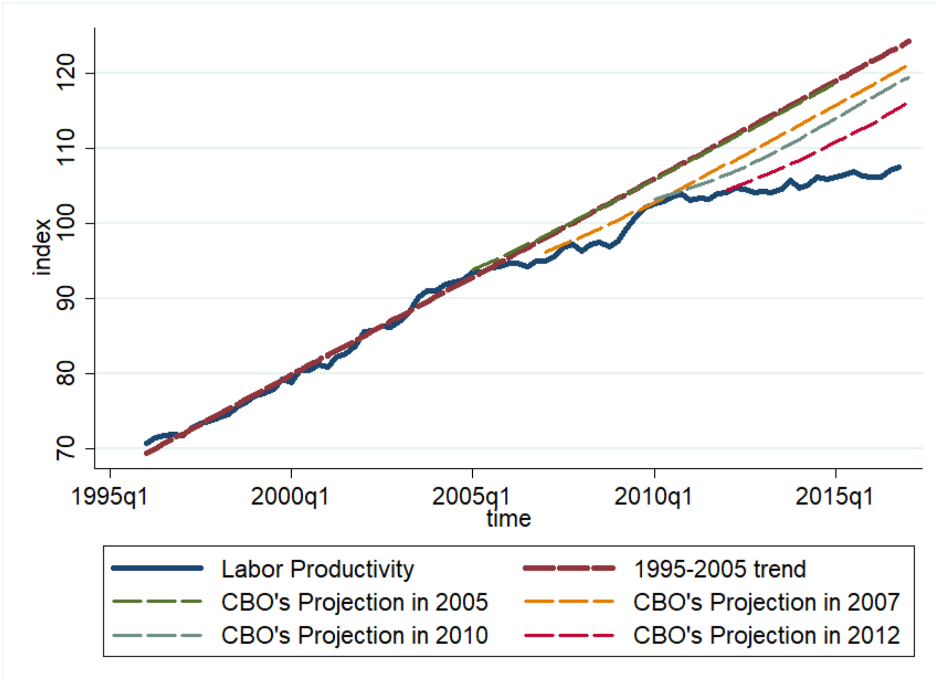

Figure B.1: Labor productivity (real GDP over hours worked) from the Bureau of Labor Statistics, and projections of labor productivity from the Congressional Budget Office.

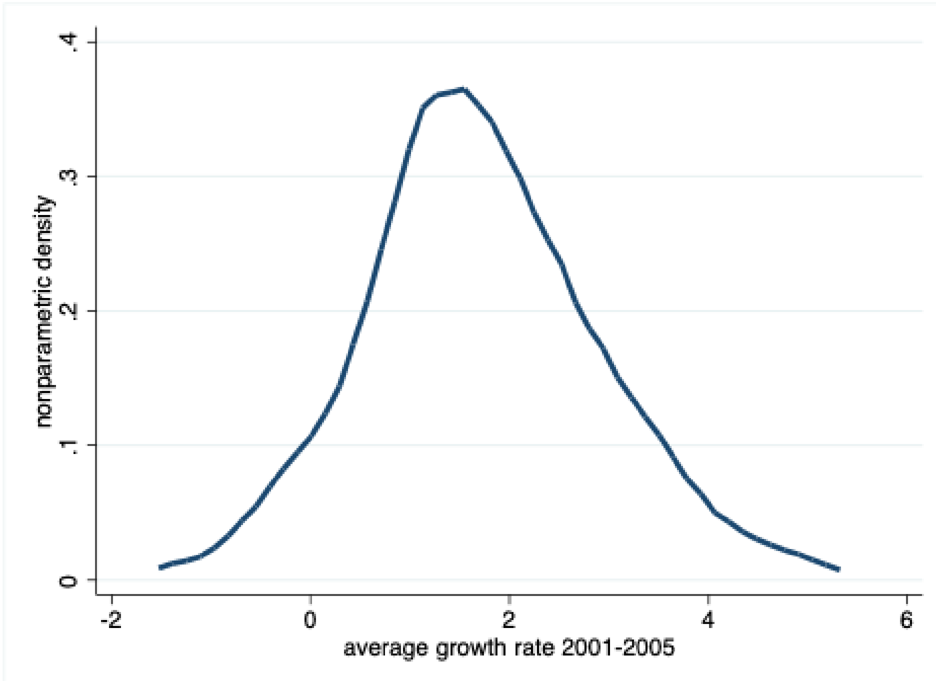

Figure B.2: Nonparametric kernel density estimate of the average annual growth rate (in percentage units) of labor productivity across MSAs, 2001–2005. Kernel = Epanechnikov, bandwidth = 0.2997. The average is 1.74, and the standard deviation is 1.62. The figure is truncated at the 2.5th and 97.5th percentiles for clarity.

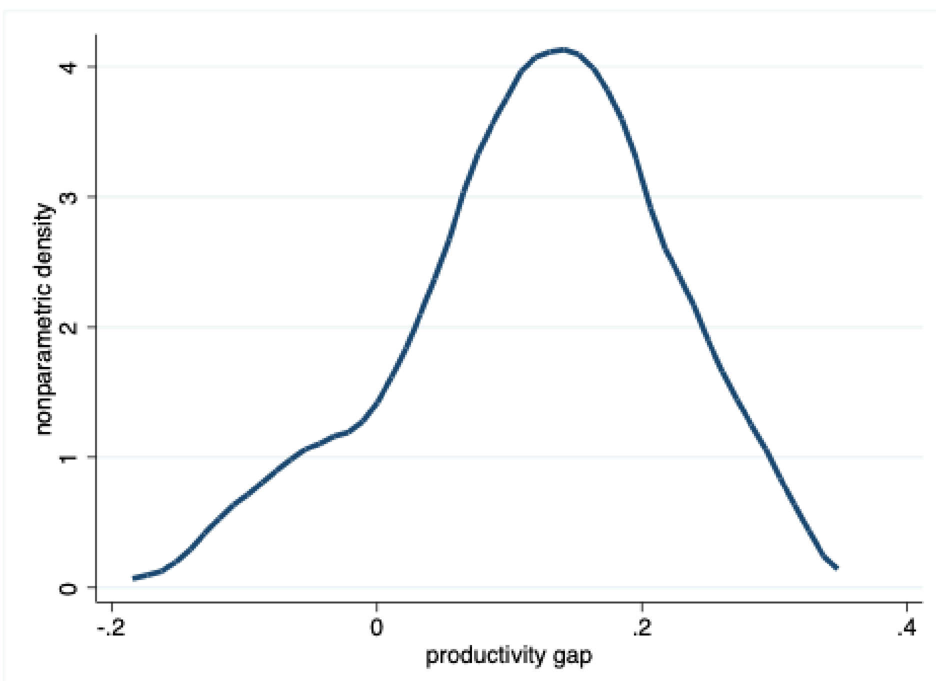

Figure B.3: Nonparametric kernel density estimate of the productivity gap across MSAs. Kernel = Epanechnikov, bandwidth = 0.0263. The figure is truncated at the 2.5th and 97.5th percentiles for clarity.

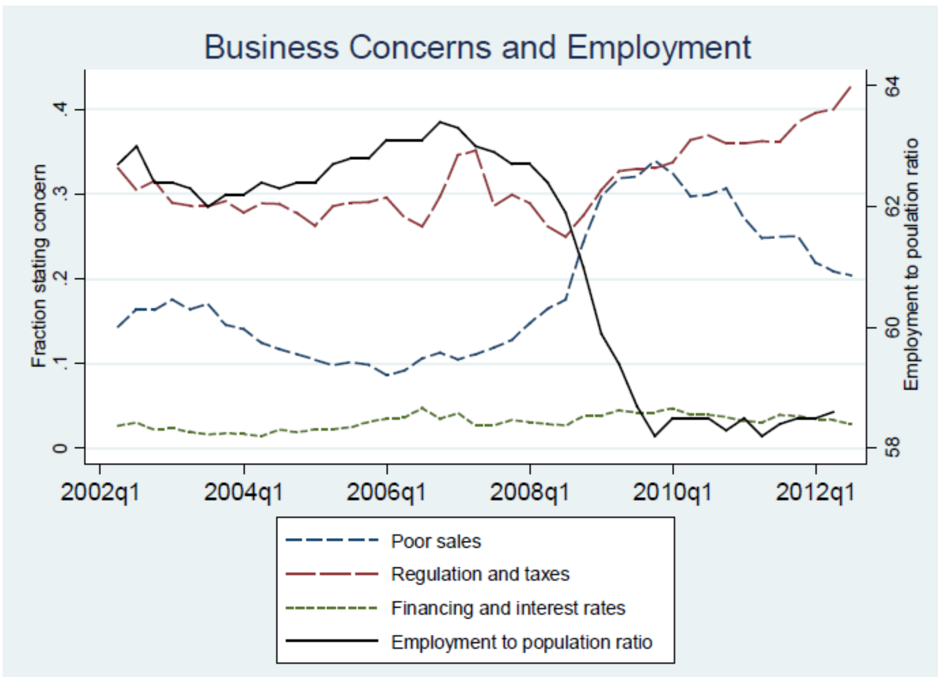

Figure B.4: Taken from Mian and Sufi (2014). National Federation of Independent Business survey. The survey asks small business owners: "What is the single most important problem facing your business today?" The figure shows the fraction of responses stating five of the ten possible concerns, with "regulation" and "taxes" combined into one category.

C Additional Tables

	Labor Productivity gap	
	(1)	(2)
	OLS	OLS
House price 2006-2012	-0.040***	-0.039***
	(0.009)	(0.008)
Include baseline controls	Yes	No
R-squared	0.1692	0.0907
Mean dependent variable	0.1153	0.1153
N (Number of MSAs)	328	380

Table C.1: Benchmark OLS estimates. House price change is standardized. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	House price 2006-2012	
	(1)	(2)
	OLS	OLS
Housing supply elasticity	0.411***	0.319***
	(0.048)	(0.054)
Include baseline controls	Yes	No
R-squared	0.3475	0.2094
N (Number of MSAs)	229	243

Table C.2: First stage regression. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	Labor productivity 2001	LP (2002-2006) wrt. (1998-2002)
	(1)	(2)
	OLS	OLS
Housing supply elasticity	-0.001	0.003
	(0.006)	(0.005)
Include baseline controls	Yes	Yes
R-squared	0.7047	0.0370
N (Number of MSAs)	229	229

Table C.3: Two additional regressions supporting the exclusion restriction. The dependent variables are the log of labor productivity levels in 2001, and the change in labor productivity growth (2002–2006 relative to 1998–2002). Baseline controls are enumerated in the text; I exclude the 2001 labor productivity level from the controls. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	Labor Productivity gap	
	(1)	(2)
	IV	IV
House price 2001-2006	0.065***	0.067***
	(0.018)	(0.012)
Include baseline controls	Yes	No
First Stage F-statistic	71.44	32.82
First Stage (F-test) p-value	0.0000	0.0000
Mean dependent variable	0.1153	0.1153
N (Number of MSAs)	229	242

Table C.4: Second stage regression using the housing boom as the main independent variable. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	labor productivity gap			
	(1) Baseline	(2) Local Banks	(3) National Banks	(4) Loans
House price 2006-2012	-0.063***			-0.075***
	(0.017)			(0.019)
House price x Local		-0.052***		
		(0.017)		
House price x National			-0.088**	
			(0.037)	
Loans gap				-0.051
				(0.031)
Include baseline controls	Yes	Yes	Yes	Yes
FS F-statistic: house prices	72.22	44.64	19.06	35.73
F-test p-value: house prices	0.0000	0.0000	0.0000	0.0000
FS F-statistic: loans				10.39
F-test p-value: loans				0.0000
N (Number of MSAs)	229	117	111	229

Table C.5: Credit supply hypothesis exercises (see text). Column 2 restricts the sample to "Local Banks" MSAs, whereas Column 3 restricts to "National Banks" MSAs. In Column 4, the loans gap is instrumented using the share of long-term illiquid assets in 1996. House price change is standardized. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

D Robustness examinations

This section describes and discusses a series of robustness and validity exercises for the main results.

No trend. Some authors (see, mainly, Gordon, 2017) have argued that the U.S. has experienced a structural change, implying that the levels of economic growth preceding the Great Recession are no longer attainable. Under this hypothesis, the pre-Great Recession trend would no longer serve as a meaningful benchmark, because the economy has entered a "new normal" oscillating around a different trend. To address this concern, I estimate a specification in which the productivity gap measure is replaced by the 2010–2015 log-difference in productivity[62]. This approach is robust to any structural change in the trend, as it relies directly on post-recession productivity growth. However, it is a demanding exercise for the reasons detailed in Section 3. Table E.1, Column (2), reports the results, indicating that a one-standard-deviation decline in house prices reduces the 2010–2015 growth in productivity by 1.6 percentage points, compared to an average of –1.59%[63]. The magnitude differs from the baseline because the dependent variable is different—in particular, a growth rate rather than a log-deviation relative to trend—but the economic effect is even larger than in the baseline (in the sense that the housing bust explains a greater fraction of the variation in the dependent variable).

Trade links. A valid concern, analyzed by Adao et al. (2019), is that an MSA is not only directly affected by local shocks but also indirectly affected by shocks in other areas with which it has trade links. A simple way to address this is to focus on the non-tradable sector[64]. The measured effect is notably larger, implying that the baseline estimates are likely a lower bound on the true effect that would be observed if trade linkages could be fully accounted for. Table E.3 shows the results.

Placebo: the tradable sector. I conduct a placebo test focusing exclusively on purely tradable industries. A purely tradable sector should be largely insulated from local demand shocks, so the response of the productivity gap to the local housing bust should be minimal. The measured effect is statistically insignificant and, if anything, reverses in sign. Table E.3 reports the results.

[62] The reason to begin in 2010 is to ensure that the measure is not affected by the contemporaneous effects of the Great Recession. Results are similar when using the annualized growth rate instead.

[63] The average growth rate from 2010 to 2015 is slightly negative when productivity is measured as GDP per worker, and slightly positive when measured using hours worked.

[64] I follow Mian and Sufi (2014) in categorizing industries as tradable or non-tradable to construct separate productivity measures.

Housing-supply instruments. A natural concern is that the Saiz (2010) housing-supply elasticity may be correlated with broader local characteristics that also shape long-run productivity dynamics, such as migration patterns, density, regional composition, or pre-existing demand fundamentals. To address this concern, I progressively saturate the baseline specification with a Davidoff-style set of controls, including immigrant share, population density, Bartik growth, region fixed effects, and a coastal-state dummy, on top of the already rich baseline controls in the main specification. Table E.5 reports the results. The estimated coefficient remains remarkably stable across these increasingly demanding specifications, which provides reassurance that the main results are not driven by the omission of these long-run local demand correlates.

These exercises also support a more cautious interpretation of the instrument. Rather than viewing the IV strategy as isolating a purely orthogonal housing-wealth shock, it is more appropriate to interpret it as capturing a broader housing shock that may be correlated with persistent local demand forces, while still delivering robust identifying variation for the empirical relationship of interest.

I also examine robustness to an alternative and independent housing-supply instrument based on the microgeographic estimates of Baum-Snow and Han (2024). Because their measure is defined at a finer geographic level than CBSAs, I map their regions into my sample using metro names and state identifiers, supplemented by conservative manual reconciliation in cases with transparent naming differences. I retain only clean and unambiguous matches, which reduces the sample size relative to the baseline specification. Table E.6 reports the corresponding results. The second-stage estimate is nearly identical to the baseline result, even though the first stage is weaker and the sample is smaller. This close similarity in magnitudes provides additional reassurance that the core relationship between the housing bust and subsequent local productivity dynamics is not specific to the Saiz measure.

Other robustness checks. The results remain robust to a variety of additional checks, including excluding finance, insurance, real estate, and construction sectors, as well as public sectors, from the productivity measure; and clustering standard errors at the state level (assigning each MSA to the state in which the majority of its population resides)[65]. Table E.4 reports these results.

Appendix G complements this section by addressing four alternative channels that merit separate examination: industry-specific shocks, the aggregate supply shock hypothesis, the business expectations hypothesis, and the business uncertainty hypothesis.

[65]The reason I do not use state clustering in the baseline specification is that many MSAs cross state boundaries. This exercise assigns each MSA to the state where most of its population is located.

E Robustness tables

	labor productivity gap	
	(1) Baseline	(2) No trend
House price 2006-2012	-0.063***	0.016**
	(0.017)	(0.007)
Include baseline controls	Yes	Yes
N (Number of MSAs)	229	229

Table E.1: Each column reports results using an alternative measure of the productivity gap. Column (1) shows the baseline results. Column (2) measures the productivity gap as the 2010–2015 log-difference in productivity (the average across MSAs is –0.0159). Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	labor productivity gap		
	(1) Baseline (degree 3)	(2) degree 2	(3) degree 4
House price 2006-2012	-0.063***	-0.063***	-0.063***
	(0.017)	(0.017)	(0.017)
Include baseline controls	Yes	Yes	Yes
N (Number of MSAs)	229	229	229

Table E.2: Each column reports results using alternative trend specifications. Column (1) uses a third-degree polynomial to construct the trend; Columns (2) and (3) use second- and fourth-degree polynomials, respectively. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	labor productivity gap		
	(1) Baseline	(2) non-tradables	(3) tradables
House price 2006-2012	-0.063***	-0.079***	0.033
	(0.017)	(0.019)	(0.065)
Include baseline controls	Yes	Yes	Yes
N (Number of MSAs)	229	161	122

Table E.3: Column (1) reports the baseline results. Column (2) restricts the productivity measure to non-tradable industries. Column (3) restricts it to tradable industries. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	labor productivity gap			
	(1) Baseline	(2) No FIRE	(3) No FIRE+public	(4) cluster se's
House price 2006-2012	-0.063***	-0.066***	-0.069***	-0.063***
	(0.017)	(0.018)	(0.018)	(0.023)
Include baseline controls	Yes	Yes	Yes	Yes
N (Number of MSAs)	229	229	229	229

Table E.4: Column (1) reports the baseline results. Column (2) excludes finance, insurance, real estate, and construction from the productivity measure. Column (3) further excludes public industries. Column (4) clusters standard errors at the state level, assigning each MSA to the state containing the majority of its population. Baseline controls are enumerated in the text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	Baseline	+ Immigration	+ Density	+ Bartik	Saturated
Housing bust	-0.063***	-0.071***	-0.060***	-0.068***	-0.072***
	(0.017)	(0.022)	(0.022)	(0.024)	(0.022)
Observations	229	225	225	225	225
KP first-stage F	72.22	49.18	43.16	39.91	45.18
Immigrant share	No	Yes	Yes	Yes	Yes
Density	No	No	Yes	Yes	Yes
Bartik growth	No	No	No	Yes	Yes
Region FE	No	No	No	No	Yes
Coastal dummy	No	No	No	No	Yes

Table E.5: Column (1) reports the baseline IV specification. Columns (2)–(5) progressively add Davidoff-style controls intended to absorb long-run local demand correlates of housing supply inelasticity: immigrant share, population density, Bartik growth, region fixed effects, and a coastal-state dummy. Immigrant share and density are constructed from the 2000 Census and aggregated to the CBSA level. Bartik growth is measured over 2001–2006 using pre-bust industry shares interacted with national industry growth. All regressions include the baseline controls from the main specification. The dependent variable is the local productivity gap. The endogenous regressor is the local house-price decline during the bust, instrumented with the Saiz housing supply elasticity measure. Robust standard errors are reported in parentheses.

Notes: Region fixed effects are defined using the four Census regions. The coastal dummy identifies metropolitan areas located in Atlantic, Pacific, or Gulf Coast states. Significance levels: * 10%, ** 5%, *** 1%. See Appendix A for additional details on variable construction.

	Baseline	Baum-Snow–Han
Housing bust	-0.0627***	-0.0630**
	(0.0166)	(0.0248)
Observations	229	137
KP first-stage F	72.22	12.27
Instrument	Saiz (2010)	Baum-Snow and Han (2024)

Table E.6: Column (1) reports the baseline IV specification using the Saiz (2010) housing-supply elasticity as the excluded instrument for the local house-price decline during the bust. Column (2) re-estimates the same specification using the neighborhood-level housing-supply measure from Baum-Snow and Han (2024), aggregated to the CBSA level through a conservative crosswalk. All regressions include the baseline controls from the main specification. The dependent variable is the local productivity gap. Robust standard errors are reported in parentheses.

Notes: The Baum-Snow–Han instrument is based on the authors' microgeographic housing-supply estimates and is matched to CBSAs only when the crosswalk is clean and unambiguous, which reduces the sample size relative to the baseline specification. Kleibergen–Paap first-stage F-statistics are reported in the table. Significance levels: * 10%, ** 5%, *** 1%. See Appendix A for additional details on variable construction.

F Additional Mechanism Evidence

Intangible Investment

To assess whether undermeasured intangible intensity may be driving the mechanism results, I augment the main investment- and consumption-channel regressions with a proxy for local exposure to intangible-intensive industries. Table F.1 reports the results.

	Investment channel		Consumption channel		Main result
	Investment slump IV	Investment slump OLS	Consumption gap IV	Investment slump OLS	Productivity gap IV
House price 2006–2012	0.6424**		-0.0837***		-0.0704***
	(0.2496)		(0.0243)		(0.0186)
Productivity gap		-2.459**			
		(0.9679)			
Consumption gap				-1.415**	
				(0.5605)	
IPP exposure (2001)	6.892*	2.721	-0.3174	1.977	-0.3645**
	(3.746)	(3.075)	(0.2767)	(3.553)	(0.1838)
Include baseline controls	Yes	Yes	Yes	Yes	Yes
N (Number of MSAs)	150	150	200	150	200

Table F.1: This table examines whether the investment and consumption mechanisms remain robust after controlling for local exposure to intangible-intensive industries. IPP exposure (2001) is constructed by interacting each MSA's pre-crisis industry employment shares with national industry-level investment in intellectual property products (IPP), measured in 2001. Columns (1) and (2) revisit the investment-channel results from Section 5.1, Table 2. Columns (3) and (4) revisit the consumption-channel results from Section 5.2, Table 3. Column (5) re-estimates the baseline IV specification for the productivity gap with the same additional control, corresponding to column (1) of Table 1. The endogenous regressor in the IV specifications is the local house-price decline during the bust, instrumented with the Saiz housing supply elasticity measure.

Notes: Baseline controls are the same as in the corresponding main-text specifications. Results using an alternative measure based on average IPP exposure over 1998–2005 are virtually identical. See Appendix A for additional details on the construction of the ACES-based investment series and the IPP exposure measure. Following CES guidelines, figures are rounded to the nearest four significant digits except for N, which is rounded to the nearest 50. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Triangulating the Investment and Consumption Channels

I next provide additional descriptive evidence on the investment and consumption channels. Tables F.2 and F.3 report reduced-form and attenuation exercises that help triangulate the role of these mediators in the relationship between the housing bust and the productivity gap.

	Productivity gap	Investment slump	Consumption gap
Panel A. Predicted housing shock			
Predicted housing shock	-0.06343***	0.4812**	-0.06918**
	(0.01609)	(0.2038)	(0.03063)
Include baseline controls	Yes	Yes	Yes
R-squared	0.2074	0.2993	0.2315
N	200	150	150
Panel B. Saiz elasticity			
Saiz elasticity	-0.02839***	0.2154**	-0.03096**
	(0.007201)	(0.09121)	(0.01371)
Include baseline controls	Yes	Yes	Yes
R-squared	0.2074	0.2993	0.2315
N	200	150	150

Table F.2: Reduced-form evidence on the investment and consumption channels. Panel A uses the predicted component of the housing bust, obtained by projecting the housing bust on the Saiz housing-supply elasticity and the controls. Panel B uses the Saiz elasticity directly. Columns report reduced-form regressions for the productivity gap, the 2007–2009 log change in investment, and the consumption gap. As in tables 2 and 3, all regressions include the baseline controls from the main specification, the 2001 investment level and 2005–2007 investment growth.

Notes: Following CES guidelines, figures are rounded to the nearest four significant digits except for N, which is rounded to the nearest 50. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Productivity gap	(2) Productivity gap	(3) Productivity gap
Panel A. Predicted housing shock			
Predicted housing shock	-0.06343***	-0.04805**	-0.03852**
	(0.01609)	(0.01841)	(0.01885)
Investment slump		-0.009441	-0.006423
		(0.007926)	(0.007943)
Consumption gap			0.1086
			(0.06706)
Include baseline controls	Yes	Yes	Yes
R-squared	0.2074	0.2194	0.2263
N	200	150	150
Panel B. Saiz elasticity			
Saiz elasticity	-0.02839***	-0.02150**	-0.01724**
	(0.007201)	(0.008239)	(0.008435)
Investment slump		-0.009441	-0.006423
		(0.007926)	(0.007943)
Consumption gap			0.1086
			(0.06706)
Include baseline controls	Yes	Yes	Yes
R-squared	0.2074	0.2194	0.2263
N	200	150	150

Table F.3: Bad-control-aware evidence on attenuation through investment and consumption. Panel A reports regressions of the productivity gap on the predicted housing shock, first without mediators, then adding the 2007–2009 log change in investment, and finally adding both the investment measure and the consumption gap. Panel B repeats the same exercise using the Saiz housing-supply elasticity directly. As in tables 2 and 3, all regressions include the baseline controls from the main specification, the 2001 investment level and 2005–2007 investment growth.

Notes: Following CES guidelines, figures are rounded to the nearest four significant digits except for N, which is rounded to the nearest 50. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

G Potential Concerns

This appendix complements Section 4.3 by discussing potential alternative factors that might underlie the measured effect. While some of these factors could indeed contribute to the productivity slowdown in their own right, the focus here is on the extent to which they might bias the estimated results.

Industry-Specific Shocks

A potential concern, already introduced in Subsection 4.1, is that industry-specific supply shocks could have simultaneously influenced both the housing bust and the productivity gap. This could occur if supply shocks disproportionately affected certain industries during the recession, leading MSAs more exposed to those industries to experience both a decline in house prices and persistent adverse effects on productivity.

Table G.1 controls for separate industry-specific effects by including the employment shares for each of the 23 two-digit industries. Column (1) presents the benchmark estimates, and Column (2) includes the 2001 employment shares as additional controls. The results remain similar.

	Labor Productivity gap	
	(1) IV	(2) IV
House price 2006-2012	-0.063*** (0.017)	-0.062** (0.031)
Include baseline controls	Yes	Yes
Include industry controls	No	Yes
First Stage F-statistic	72.22	19.36
First Stage (F-test) p-value	0.0000	0.0000
N (Number of MSAs)	229	141

Table G.1: Second-stage regression with additional industry controls. Baseline controls are enumerated in the text. Additional industry controls include the employment shares of the two-digit industries, excluding those with missing data for more than 100 MSAs. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

The Aggregate Supply Shock Hypothesis

Another concern is the hypothesis that a persistent aggregate supply shock could underlie both the housing bust and the productivity gap. However, this possibility is unlikely to bias the results for two main reasons.

First, an aggregate supply shock would typically affect all industries more uniformly. Yet Section 4.3 finds no correlation between the housing bust and the productivity gap within purely tradable industries. If an aggregate supply shock were the true driver of the observed relationship, it should also manifest in the tradable sector. There is no clear reason why such a shock would affect only the non-tradable sector, especially considering that the productivity slowdown spans both tradable and non-tradable industries (see, e.g., Manyika et al., 2017).

Second, an aggregate supply shock would not naturally produce the clear geographic patterns documented above—which are essential to qualify as an alternative explanation—unless it disproportionately affected the industries concentrated in MSAs with larger house price declines. Even in that case, much of the effect would likely be absorbed by the industry-specific controls included in the preceding exercise, which do not materially alter the estimates.

The Business Expectations/Uncertainty Hypothesis

This section also considers the hypotheses that either a shift in business expectations or an increase in business uncertainty could lie behind both the housing bust and the productivity gap.

As in the previous discussion, these shocks would be expected to exert broad effects across sectors, which contrasts with the finding in Section 4.3 showing no significant relationship between the housing bust and the productivity gap for tradable industries. If either hypothesis were the primary explanation, there is no clear reason why it would spare the tradable sector, which itself experienced a pronounced slowdown.

Even if an expectations or uncertainty shock affected the economy but left the tradable sector relatively untouched, for it to qualify as an alternative explanation of the results, it would have to hit more strongly those MSAs experiencing larger house price declines. This could occur for three reasons. First, such shocks might disproportionately impact certain industries, which in turn are more prevalent in the MSAs with severe housing declines. However, this effect would likely be captured by the industry-specific controls introduced earlier. Second, these shocks could be a consequence of the housing bust itself—for example, through negative business sentiment following reduced local demand. In this case, they would reinforce the main results, representing part of the transmission mechanism behind the estimated elasticity. Third, such shocks could directly influence house price declines. This final possibility merits closer examination.

Although one might expect such shocks to affect the tradable sector as well, the third scenario is potentially problematic because it could bias the results. However, for this to occur, the shock would need to have been significantly more severe in those MSAs that eventually experienced larger house price declines. While expectation or uncertainty shocks are often national, there are two plausible ways in which they could have exerted differential local effects. On one hand, a national shock might have a greater impact in MSAs with more inelastic housing supply. Examples include national housing demand shocks driven by sentiment (see Soo, 2018; Shiller, 2015) or fundamentals such as interest rate changes. In such cases, local effects stem from the interaction between the shock and local housing supply elasticity, implying that the shock would affect the productivity gap only indirectly through house price movements—and thus would not bias the results[66]. On the other hand, expectations or uncertainty shocks might be local in nature if they stem directly from the housing boom. In that case, they simply represent an alternative mechanism linking the boom to the productivity gap. The housing boom is identified similarly to the bust and has comparable implications, though it is challenging to empirically disentangle the consequences of one versus the other because the bust effectively represents the unwinding of the preceding boom[67]. Table C.4 reports the results using the housing boom as the primary explanatory variable.

[66] In other words, while such shocks could influence the productivity gap, they would not introduce bias because they would be effectively differenced out in the regressions.

[67] The main specification focuses on the housing bust rather than the boom because, as Section 5 explains, the bust is the more plausible trigger for the mechanisms examined in this paper.

H Alternative Mechanisms

The R&D Hypothesis

An alternative mechanism is the possibility that R&D declined more sharply in MSAs with larger house price drops, ultimately dampening productivity. This could arise if local demand effects stemming from the housing bust reduced the economic returns to R&D activities.

However, aggregate trends suggest this is unlikely to be a major driver. Figure H.1 shows that aggregate R&D expenditures as a share of GDP did not fall during the Great Recession; in fact, they increased significantly from 2005 onward and remained historically high in subsequent years. Meanwhile, Anzoategui et al. (2019) documents that while U.S. corporate R&D spending (in levels) declined relative to trend during the Great Recession, the drop was substantially smaller than during the 2001 recession, even though the 2001 downturn was milder.

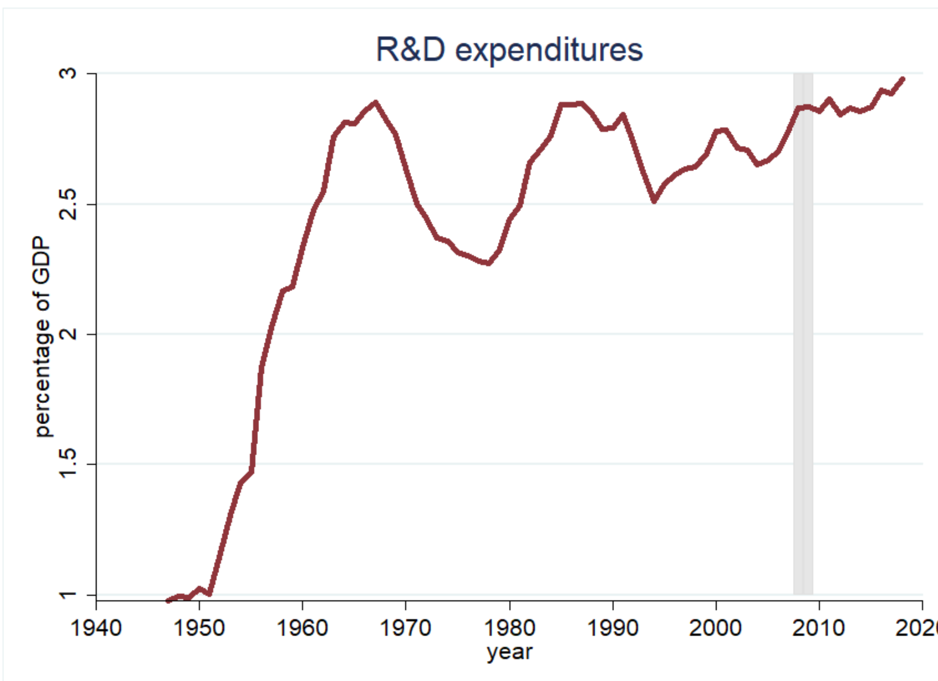

Figure H.1: Total R&D expenditures as a percentage of GDP. Source: U.S. Bureau of Economic Analysis. The shaded area corresponds to 2008–2009.

Table H.1 reports the results of an exercise that tests this hypothesis directly. It uses data on patent grants (USPTO) as a proxy for R&D and innovation. I compute patent grants per 10,000 people and calculate the 2015 log gap relative to trend. Column (2) presents estimates from a specification that includes this patents gap as an additional regressor. The coefficient on the patents gap is insignificant, and the main estimate remains statistically indistinguishable from the baseline.

	labor productivity gap	
	(1)	(2)
	2SLS	2SLS
House price 2006-2012	-0.063***	-0.061***
	(0.017)	(0.016)
Patents gap		-0.015
		(0.010)
Include baseline controls	Yes	Yes
First Stage F-statistic	72.22	78.22
First Stage (F-test) p-value	0.0000	0.0000
N (Number of MSAs)	229	229

Table H.1: See text. Robust standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

The Local Labor Hypothesis

Another hypothesis is that the housing bust altered the composition of local labor markets in ways that affected aggregate productivity. For example, the bust may have reduced employment disproportionately in MSAs experiencing the steepest house price declines (Mian and Sufi, 2014), thereby shifting the local employment mix.

However, this seems unlikely to be a meaningful mechanism, because employment losses during the recession were heavily concentrated among low-skilled workers (see, e.g., Charles et al., 2018). Since these workers tend to be less productive, their disproportionate exit from employment would arguably lead to an increase, rather than a decline, in measured productivity in MSAs hit hardest by the housing bust. Figure H.2 further supports this interpretation, showing a sharp rise in real wages during the recession, which is consistent with a compositional shift away from lower-productivity workers[68].

A related possibility is that the preceding housing boom attracted lower- and middle-skilled construction workers into MSAs that later suffered severe house price declines, potentially skewing local labor markets in ways that reduced productivity. While this selection effect might have influenced productivity dynamics during the boom, it is likely to have been

[68] Part of this sharp increase may reflect changes in price deflators due to falling oil prices during the Great Recession.

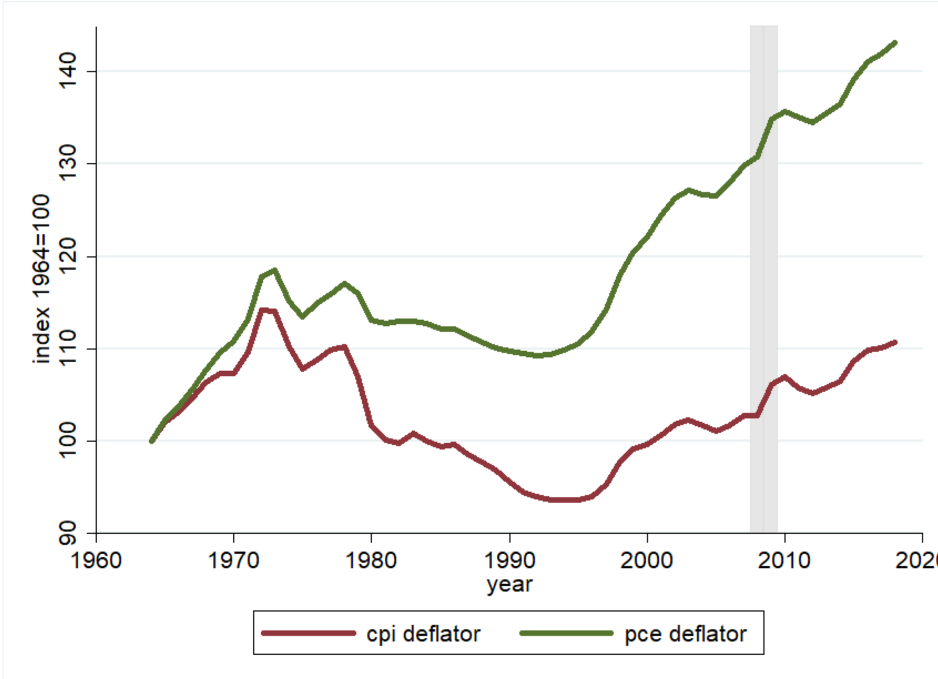

Figure H.2: Average hourly earnings for production and nonsupervisory employees: total private sector. "CPI deflator" denotes values deflated by the consumer price index; "PCE deflator" uses the personal consumption expenditure index. The shaded area corresponds to 2008–2009.

largely reversed during the Great Recession, as employment losses were heavily concentrated among lower-skilled workers[69]. Moreover, if such compositional changes were concentrated in specific sectors like construction, a substantial portion of any resulting effect should have been absorbed by the industry controls used in Section G and the construction-specific controls in Section 4.3.

[69]See Charles et al., 2018 and Yagan, 2019.

I Model Appendix

This section provides the proofs for the model's equilibrium results.

Let

$$\Omega := \underset{x}{\operatorname{argmin}} \frac{R_t(x) + W_t}{A_t x}.$$

Proposition 1. *For all $x \notin \Omega$, $R_t(x) = 0$.*

Proof. Suppose there exists an efficiency level $\tilde{x} \notin \Omega$ such that $R_t(\tilde{x}) > 0$. From the definition of Ω, we have $h_t^{(i)}(\tilde{x}) = 0$ for all $i \in [0, 1]$. However, if $R_t(\tilde{x}) > 0$, capital owners would wish to rent out all machines with efficiency $\tilde{x}$. Thus, the rental market would not clear, contradicting equilibrium.

Proposition 2. *For all $x_1, x_2 \in \Omega$,*

$$\frac{R_t(x_1) + W_t}{A_t x_1} = \frac{R_t(x_2) + W_t}{A_t x_2} = c$$

for some constant $c \geq 0$.

Proof. Follows directly from the definition of Ω. Suppose there exist $\tilde{x}_1, \tilde{x}_2 \in \Omega$ such that:

$$\frac{R_t(\tilde{x}_1) + W_t}{A_t \tilde{x}_1} > \frac{R_t(\tilde{x}_2) + W_t}{A_t \tilde{x}_2}.$$

Then, $\tilde{x}_1$ could not belong to Ω, contradicting its definition.

Corollary 1. From Propositions 1 and 2, $R_t(x)$ must take the following form:

- For all $x \notin \Omega$, $R_t(x) = 0$, implying

$$\frac{R_t(x) + W_t}{A_t x} = \frac{W_t}{A_t x}.$$

- For all $x \in \Omega$,

$$R_t(x) = cA_t x - W_t.$$

This follows from
$$\frac{R_t(x) + W_t}{A_t x} = c.$$

Proposition 3. $\Omega = [T_t, \infty)$, *where* T_t *is defined implicitly by*
$$\int_{T_t}^{\infty} A_t x \, H_t(x) \, dx = Y_t.$$

Proof. Suppose for contradiction that $\Omega \neq [T_t, \infty)$.[70]

From Corollary 1 and the definition of Ω, we have:
$$\frac{R_t(x_1) + W_t}{A_t x_1} = c < \frac{W_t}{A_t x_2}, \quad \forall x_1 \in \Omega, \quad \forall x_2 \notin \Omega.$$

Now, for any $\tilde{x}$ such that $h_t^{(i)}(\tilde{x}) > 0$, it must hold that $R_t(\tilde{x}) \geq 0$. Otherwise, capital owners would refuse to rent machines of efficiency $\tilde{x}$.

Hence, there must exist some $\hat{x} \in \Omega$ such that $R_t(\hat{x}) \geq 0$. Moreover, from Corollary 1, for any $x_1 < x_2$ in Ω, we have $R_t(x_1) < R_t(x_2)$.

Therefore, if $[T_t, \infty) \not\subset \Omega$, there must exist some $\hat{x} \in \Omega$ such that:

(i) $\hat{x} < T_t$,

(ii) $R_t(\hat{x}) \geq 0$.

Otherwise, $R_t(x) < 0$ for all $x < T_t$, implying $h_t^{(i)}(x) = 0$ for all $x < T_t$ and for all i. Given that $[T_t, \infty) \not\subset \Omega$, the goods market could not clear.

Now, choose any $\bar{x} \notin \Omega$ with $\bar{x} > T_t$ (such $\bar{x}$ exists in this case). From the above:
$$c < \frac{W_t}{A_t \bar{x}}.$$

Also, because $R_t(\hat{x}) \geq 0$, we have
$$c \geq \frac{W_t}{A_t \hat{x}}.$$

Together, this implies:
$$\frac{W_t}{A_t \hat{x}} \leq c < \frac{W_t}{A_t \bar{x}},$$

[70] I focus on the case where $[T_t, \infty) \not\subset \Omega$. Cases where $[T_t, \infty) \subset \Omega$ might be consistent with equilibrium but would imply $R_t(x) \leq 0$ for all $x < T_t$, meaning no machines with $x < T_t$ would be rented. The equilibrium outcome would then be identical to the case $\Omega = [T_t, \infty)$.

which gives $\hat{x} > \bar{x}$. But this contradicts $\hat{x} < T_t < \bar{x}$.

Thus, $\Omega = [T_t, \infty)$.

Proposition 4. $c = \frac{W_t}{A_t T_t}$.

Proof. Suppose not. Assume instead that $c = D < \frac{W_t}{A_t T_t}$.[71]

Note the following:

1. $R_t(x) \geq 0 \iff R_t(x) = DA_t x - W_t \geq 0 \iff x \geq \frac{W_t}{DA_t}$. Since $D < \frac{W_t}{A_t T_t}$, it follows that
$$x \geq \frac{W_t}{DA_t} > T_t.$$
Thus, $R_t(x) \geq 0$ implies $x > T_t$.

2. $R_t(T_t) = DA_t T_t - W_t < 0$.

3. $R_t(x) = 0$ implies $x = \frac{W_t}{DA_t}$.

4. The derivative of $R_t(x)$ with respect to x equals DA_t.

Define
$$\Phi := \left[T_t, \frac{W_t}{DA_t}\right).$$

For all $x \in \Phi$, we have $R_t(x) < 0$, so no machine with $x \in \Phi$ would be rented. Thus, only machines with $x \geq \frac{W_t}{DA_t} > T_t$ could be rented. But since $\Phi \subset \Omega$ and $\Omega = [T_t, \infty)$, the goods market cannot clear. This contradiction implies

$$c = \frac{W_t}{A_t T_t}.$$

[71] An equilibrium where $c > \frac{W_t}{A_t T_t}$ cannot exist because $\frac{W_t}{A_t T_t}$ is the infimum of $\frac{R_t(x)+W_t}{A_t x}$ over $x \notin \Omega$, contradicting the definition of Ω.

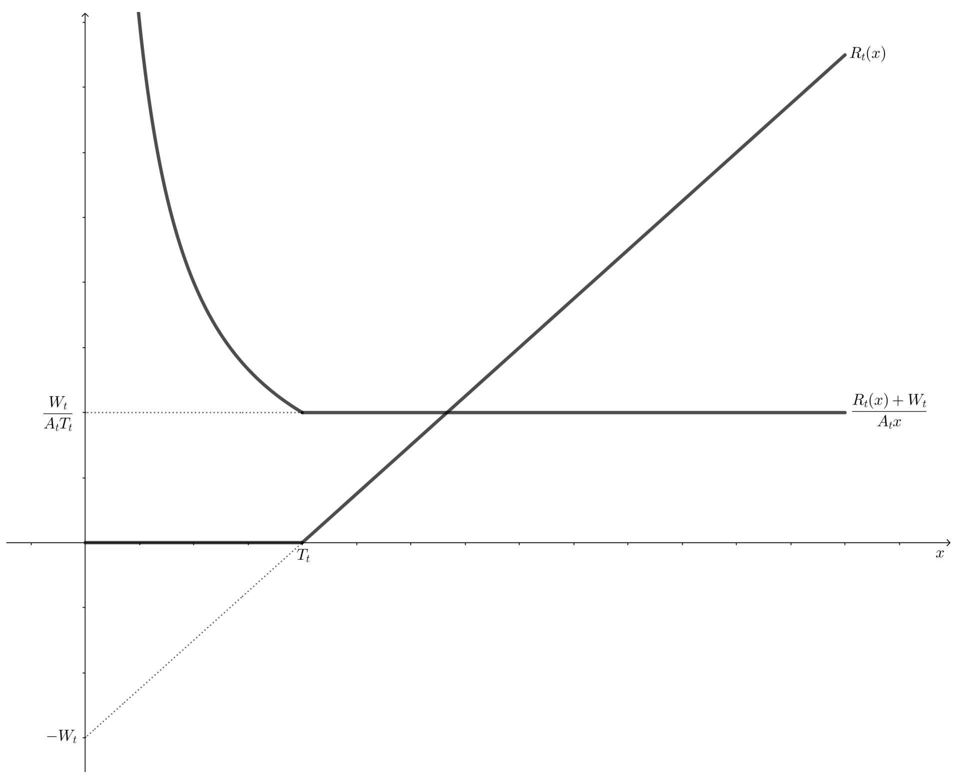

Figure I.1: Illustration.

J Quantitative Appendix

This appendix reports additional quantitative results for the model. It provides three complements to the main text. First, it documents the dynamic fit of the benchmark experiment at intermediate horizons. Second, it examines the sensitivity of the results to key structural parameters. Third, it compares the benchmark demand shock to alternative demand disturbances in order to clarify the extent to which the model's implications are shock-specific or instead reflect the broader investment/embodiment mechanism.

J.1 Dynamic Fit at Intermediate Horizons

Table J.1 reports the dynamic fit of the benchmark quantitative experiment at 3-, 6-, and 10-year horizons. The model is calibrated to match the long-run consumption shortfall, while the implied paths for GDP, investment, and labor productivity are left untargeted.

The table shows that the benchmark experiment tracks consumption well at medium and long horizons and also delivers a reasonable fit for GDP. For investment, the model performs reasonably well at short and medium horizons, but it underpredicts the persistence of the investment slump at the 10-year horizon. This is important for interpreting the productivity results, since the weaker long-run investment response mechanically limits the amount of persistent productivity decline that the model can generate through the investment/embodiment channel.

For labor productivity, the model is not designed to match the short-run increase observed in the data around the recession itself, which is likely influenced by compositional effects associated with the sharp contraction in employment. More importantly, the model generates a persistent productivity shortfall at medium and long horizons, but one that remains below the magnitude observed in the data. This pattern is consistent with the main quantitative message of the paper: the investment/embodiment channel can account for a substantial share of the post-recession productivity gap—roughly one half in the benchmark experiment—but not all of it. The fact that the model also understates the long-run investment slump suggests that this quantitative estimate is likely conservative.

J.2 Sensitivity Analysis

Table J.2 examines how the benchmark quantitative results change under alternative values of the key parameters governing persistence in the model: adjustment costs, depreciation,

	3 years		6 years		10 years	
	Model	Data	Model	Data	Model	Data
Consumption	0.0837	0.0581	0.0670	0.0727	0.0468	0.0450
GDP	0.0975	0.0629	0.0750	0.0718	0.0562	0.0727
Investment	0.1612	0.1744	0.1117	0.1258	0.0993	0.1275
Labor productivity	0.0653	0.0001	0.0579	0.0454	0.0464	0.0933

Table J.1: Dynamic fit of the benchmark experiment at intermediate and long horizons. Model entries report impulse responses at 12, 24, and 40 quarters after the shock. Data entries report the corresponding empirical gaps at 3, 6, and 10 years after 2008. Consumption is the only targeted moment, and only at the long horizon. GDP, investment, and labor productivity are non-targeted implications of the model. All entries are expressed in percentage points and indicate how much each series would need to grow to return to trend (or to steady state in the model).

and nominal price rigidity. In each case, I report the implied GDP, investment, and labor productivity gaps at the 40-quarter horizon, together with the fraction of the corresponding empirical gap explained by the model.

Three results stand out. First, the benchmark conclusion is robust across calibrations: the model continues to generate a sizable and persistent productivity shortfall in response to the demand shock. Second, depreciation plays a central role in shaping persistence. Lower depreciation increases the share of the observed productivity gap explained by the model, whereas higher depreciation substantially weakens it, consistent with the idea that faster capital turnover reduces the persistence of the vintage-capital mechanism. Third, greater nominal price rigidity also strengthens the mechanism by making the investment slump more persistent.

By contrast, adjustment costs mainly affect the quantitative magnitude of the response rather than the existence of the mechanism itself. For reasonable variations around the baseline, their effects are limited. Only when adjustment costs become sufficiently large do they materially dampen the investment response, thereby weakening the model's ability to generate a persistent productivity shortfall.

Taken together, these exercises suggest that the model's ability to account for the post-recession productivity gap does not rely on a knife-edge calibration. At the same time, they clarify which ingredients matter most for persistence: slow capital turnover and sufficiently persistent demand-side adjustment.

Calibration	Y_{40}	Inv_{40}	$Prod_{40}$	Share Y explained	Share Inv explained	Share $Prod$ explained
Baseline	0.0562	0.0993	0.0464	0.773	0.779	0.497
$\phi_Q = 25$	0.0585	0.0831	0.0493	0.805	0.651	0.528
$\phi_k = 50$	0.0533	0.1069	0.0367	0.734	0.839	0.394
$\phi_Q = 50,\ \phi_k = 50$	0.0577	0.1068	0.0413	0.793	0.838	0.443
$\delta = 0.020$	0.0589	0.1047	0.0483	0.810	0.821	0.518
$\delta = 0.060$	0.0460	0.0893	0.0366	0.633	0.700	0.392
$\delta = 0.100$	0.0383	0.0806	0.0282	0.526	0.632	0.302
$\phi = 0.66$	0.0507	0.0891	0.0438	0.698	0.699	0.470
$\phi = 0.85$	0.0739	0.1263	0.0544	1.016	0.990	0.583

Table J.2: Sensitivity of the benchmark quantitative results to key structural parameters. The table reports the model-implied GDP, investment, and labor productivity gaps at the 40-quarter horizon under alternative calibrations, together with the share of the corresponding empirical gap explained by the model. The baseline calibration is reported in the first row. The parameter ϕ denotes the Calvo price stickiness parameter, δ the depreciation rate, and ϕ_Q and ϕ_k the adjustment cost parameters on the number of machines and on the capital-labor ratio of new machines, respectively.

J.3 Alternative Demand Shocks

The benchmark quantitative experiment models the post-2007 contraction as a persistent demand shock through the preference shifter S_t. This is intended as a parsimonious device to capture the demand-side consequences of the housing bust rather than as a literal claim about the primitive source of the disturbance. To assess how shock-specific the quantitative results are, I compare the benchmark experiment with an alternative contractionary monetary shock, which arises naturally in the model through the Taylor rule.

A monetary shock is not structurally equivalent to the benchmark preference shock, since it affects demand not only through expenditure decisions but also through the interest rate, which directly influences investment. For that reason, the most informative comparison is not one that attempts to align the consumption path, but one that reproduces the benchmark investment slump. Table J.3 reports the results for a monetary-shock calibration chosen to bring Inv_{40} close to its benchmark value. In particular, the monetary-shock experiment sets $\rho_i = 0.935$, $\phi_\pi = 1.75$, and $\sigma_i = 0.35$.

The table shows that, once the investment response is aligned, the implied productivity gap is also very similar to the benchmark experiment. This suggests that, for the purposes of the quantitative exercise, the key object for productivity in the model is the persistence of the investment slump rather than the precise form of the underlying demand disturbance. At the same time, the two shocks are not fully observationally equivalent: they imply different paths for hours, wages, marginal costs, and especially the interest rate. Thus, the model is broadly agnostic about the ultimate source of the demand contraction for the purposes of the investment/embodiment channel, but it does not imply full equivalence across all observable margins.

	Benchmark S_t shock	Monetary shock matched on Inv_{40}
Panel A. Core moments at the 40-quarter horizon		
Consumption gap (C_{40})	0.0468	0.0319
GDP gap (Y_{40})	0.0562	0.0446
Investment gap (Inv_{40})	0.0993	0.1028
Labor productivity gap ($Prod_{40}$)	0.0464	0.0448
Panel B. Diagnostics at the 40-quarter horizon		
Hours (N_{40})	0.0098	-0.0002
Real wages (w_{40})	0.0566	0.0317
Marginal cost (mc_{40})	0.0096	-0.0011
Nominal interest rate (i_{40})	0.0330	-0.0025

Table J.3: Comparison between the benchmark preference-shock experiment and an alternative monetary shock calibrated to match the benchmark investment gap at the 40-quarter horizon. The benchmark experiment uses the preference shifter S_t. The monetary-shock calibration sets $\rho_i = 0.935$, $\phi_\pi = 1.75$, and $\sigma_i = 0.35$, which yields an investment gap at the 40-quarter horizon very close to the benchmark experiment while keeping the calibration reasonably parsimonious. Panel A reports the main quantitative moments. Panel B reports selected diagnostics that help distinguish the two disturbances. All entries are expressed in percentage points, with positive values indicating a shortfall relative to trend (or relative to steady state in the model).

Appendix References

Adao, R., C. Arkolakis, and F. Esposito (2019). *Spatial linkages, global shocks, and local labor markets: Theory and evidence*. Tech. rep. National Bureau of Economic Research.

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019). "Endogenous technology adoption and R&D as sources of business cycle persistence". *American Economic Journal: Macroeconomics* 11.3, pp. 67–110.

Baum-Snow, N. and L. Han (2024). "The Microgeography of Housing Supply". *Journal of Political Economy* 132.6, pp. 1897–1946. DOI: 10.1086/728110.

Charles, K. K., E. Hurst, and M. J. Notowidigdo (2018). "Housing booms, manufacturing decline and labour market outcomes". *The Economic Journal* 129.617, pp. 209–248.

Davis, S. J., J. C. Haltiwanger, and S. Schuh (1998). "Job Creation and Destruction". *The MIT Press.*

Doms, M. and T. Dunne (1998). “Capital Adjustment Patterns in Manufacturing Plants”. *Review of Economic Dynamics* 1(2), pp. 409–429.

Fishback, P. V., W. C. Horrace, and S. Kantor (2005). “Did New Deal grant programs stimulate local economies? A study of Federal grants and retail sales during the Great Depression”. *The Journal of Economic History* 65.1, pp. 36–71.

Garrett, T. A., R. Hernández-Murillo, M. T. Owyang, et al. (2005). “Does consumer sentiment predict regional consumption”. *Federal reserve Bank of st. Louis review* 87.2, pp. 123–135.

Gordon, R. J. (2017). *The rise and fall of American growth: The US standard of living since the civil war*. Vol. 70. Princeton University Press.

Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson (2018). *Housing wealth effects: The long view*. Tech. rep. National Bureau of Economic Research.

Kaplan, G., K. Mitman, and G. L. Violante (2016). *Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data*. Tech. rep. National Bureau of Economic Research.

Manyika, J., J. Remes, J. Mischke, and M. Krishnan (2017). “The Productivity Puzzle, A Closer Look at the United States”. *McKinsey Global Institute, Discussion Paper.*

Mian, A., K. Rao, and A. Sufi (2013). “Household balance sheets, consumption, and the economic slump”. *The Quarterly Journal of Economics* 128.4, pp. 1687–1726.

Mian, A. and A. Sufi (2014). “What explains the 2007–2009 drop in employment?” *Econometrica* 82.6, pp. 2197–2223.

Moretti, E. (2010). *Local labor markets*. Tech. rep. National Bureau of Economic Research.

Saiz, A. (2010). “The geographic determinants of housing supply”. *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.

Shiller, R. J. (2015). *Irrational exuberance: Revised and expanded third edition*. Princeton university press.

Soo, C. K. (2018). "Quantifying sentiment with news media across local housing markets". *The Review of Financial Studies* 31.10, pp. 3689–3719.

Tornqvist, L., P. Vartia, and Y. O. Vartia (1985). "How Should Relative Changes Be Measured?" *The American Statistician* 36(1), pp. 43–46.

Yagan, D. (2019). "Employment hysteresis from the great recession". *Journal of Political Economy* 127.5, pp. 000–000.