THE ECONOMIC JOURNAL

The Economic Journal, **133** (*July*), 2055–2067 https://doi.org/10.1093/ej/uead008
Advance Access Publication Date: 25 January 2023

LOCATION, LOCATION, LOCATION: MANUFACTURING AND HOUSE PRICE GROWTH*

Xiangyu Feng, Nir Jaimovich, Krishna Rao, Stephen J. Terry and Nicolas Vincent

Exploiting data on tens of millions of housing transactions, we show that (1) house prices grew by less in manufacturing-heavy US regions, (2) this pattern is especially present for the lowest-value homes and that (3) price declines coincided with worse labour market outcomes, consistent with an income channel. Counterfactual accounting exercises reveal that regional differences in the growth of these lowest-value homes are an important driver of the changes in overall house price inequality. Hence, the economic decline in manufacturing-heavy areas extends far beyond income and employment flows to house prices.

Over the past few decades, US manufacturing plunged from an aggregate employment share of over 21% in 1980 to just under 9% in 2010.[1] Concurrent with this aggregate decline, geographic locations where manufacturing used to account for a high share of employment, such as those in the Upper Midwest or Rust Belt, have seen lower wage and employment growth than their low-manufacturing peers.[2]

In this paper, we demonstrate that the relatively poor economic experience of manufacturing hot spots extends beyond labour market outcomes to the price of a key local asset: housing. Our analysis is based on a rich micro-dataset of prices and characteristics of tens of millions of recent home transactions with broad geographical coverage. We proceed in three steps.

First, manufacturing-heavy areas saw lower house price growth on average. The right panel of Figure 1 maps the manufacturing employment share in 2000 across regions, revealing substantial heterogeneity, while the left panel plots regional house price growth from 2001–06. The negative spatial correlation between the two variables is evident. As we show below, our estimates suggest that areas with a 10 percentage point higher manufacturing share experienced an average of 4.7% lower house price growth each year. To interpret the mechanism behind this relation, we present evidence consistent with the presence of an income channel: areas heavily exposed to

* Corresponding author: Nicolas Vincent, Department of Applied Economics, HEC Montréal, 3000 ch. de la Cote-Sainte-Catherine, Montreal, QC, Canada. Email: nicolas.vincent@hec.ca

This paper was received on 16 November 2020 and accepted on 17 January 2023. The Editor was Francesco Lippi.

The authors were granted an exemption to publish their data because access to the data is restricted. However, the authors provided a simulated or synthetic dataset that allowed the Journal to run their codes. The synthetic/simulated data and codes are available on the Journal repository. They were checked for their ability to generate all tables and figures in the paper; however, the synthetic/simulated data are not designed to reproduce the same results. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.7525610.

A previous version of this paper circulated under the title ‘Location, Location, Location: Industrial Structure and the Distribution of House Price Growth’. We thank numerous seminar participants as well as our discussants Evi Pappa and Linda Tesar for helpful comments. Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions do not reflect the position of Zillow Group. The results and opinions do not reflect the position of Landed, Inc. All remaining errors are our own. Vincent gratefully acknowledges financial support from the Foundation HEC Montréal and the Social Sciences and Humanities Research Council of Canada (grant #435-2022-0800).

[1] These figures are from the Bureau of Labor Statistics’ Establishment Survey.

[2] See the evidence in Ramey (2018) and Charles *et al.* (2019) as well as in our own data in Online Appendix A.

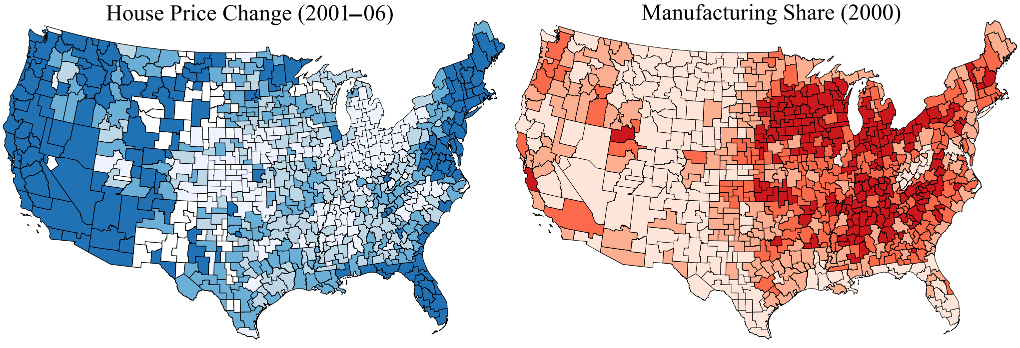

Fig. 1. *House Prices and Manufacturing.*
Notes: Both maps plot the contiguous United States in commuting zones. The left map in blues indicates average percentage house price appreciation from 2001–06, based on local Federal Housing Finance Agency indexes. The right map in reds indicates the manufacturing share of employment in 2000 based on US Census IPUMS micro-data. Darker shades indicate larger values.

manufacturing see worse labour market outcomes reflected in lower house price growth through depressed demand for housing.

While these differences in average growth rates are quantitatively significant, our second key finding is that they mask important distributional heterogeneity. Specifically, we leverage the strength of our micro-level data to show that the lowest-valued homes in manufacturing-heavy areas experienced substantially lower price growth than their peers in manufacturing-light locales. This predicted relationship between manufacturing and house price growth is significantly more muted for higher-valued homes, linking industrial structure to house price inequality both between *and* within regions.

Finally, we relate our findings to the overall evolution in cross-sectional house price inequality. To do so, we exploit our full micro-distribution of US house prices to document a strong positive contribution of the heterogeneity in house price growth across regions to the inequality in house values. Moreover, we show that most of this contribution is driven by the lowest-price segment.

Our paper contributes to multiple strands of work. First, it adds to the literature studying the relationship between manufacturing exposure and various labour market and social outcomes, such as the polarisation of the job market in the United States (Bárány and Siegel, 2018); the marriage market value of young men (Autor *et al.*, 2019); the recent increases in mortality and morbidity among white non-Hispanic Americans (Case and Deaton, 2017) and other phenomena (Kahn, 1999; Feyrer *et al.*, 2007; Alder *et al.*, 2017; Notowidigdo, 2020). Our contribution is to show that the exposure to manufacturing is also a factor linked to rising cross-sectional inequality in house prices. As such, we add to the broader understanding of the evolution of wealth and income inequality (Saez and Zucman, 2016; Ahn *et al.*, 2018; Kaplan *et al.*, 2018b; Song *et al.*, 2018). Finally, our work is related to the literature that links house price movements and the broader economy (Charles *et al.*, 2016; Piazzesi and Schneider, 2016; Berger *et al.*, 2017; Liebersohn, 2017; Howard and Liebersohn, 2018; 2021; Kaplan *et al.*, 2018a; Guren *et al.*, 2021).

1. Data

In this section we present the data used throughout the analysis.

1.1. *Housing Data*

Studying average and distributional shifts in house prices requires us to follow house price distributions within narrow geographic locations over time. We rely on a unique micro-dataset from Zillow, the ZTRAX dataset (Zillow Group, 2020), containing tens of millions of observations from 2001 to 2015 with a wide geographical coverage. ZTRAX combines two sources of information: local municipalities' transaction records, including sales prices, and tax assessment data featuring detailed home characteristics. Thus, an observation in our dataset combines both the sales price *and* home characteristics for a single transaction. For this study, our focus is on single-family homes.

1.2. *Geography*

Our geographical analysis is at the commuting zone (CZ) level, an area whose size is typically between that of a county and a state, and which corresponds to a locally unified economic agglomeration. This measure ensures comparability with other recent work on industrial structure and labour market outcomes (Autor *et al.*, 2013).

1.3. *Labour Market*

To measure local manufacturing employment shares, as well as various other labour market outcomes and controls, we use 1% decennial census and annual American Community Survey IPUMS micro-data extracts (U.S. Census Bureau, 2019). At the CZ level, this dataset provides universal geographic coverage within the United States, and sample weights attached to the micro data allow for the formation of representative measures.

1.4. *Additional Datasets*

To compute various ancillary statistics and provide cross checks of our main results, we also use several other additional sources: aggregated local house price indexes from the Federal Housing Finance Agency (FHFA, 2020), local housing supply elasticities from Saiz (2010) and local industry employment from the Census County Business Patterns (U.S. Census Bureau, 2020). See Online Appendix A for more details on our sample construction, exact variable definitions and summary statistics for our datasets.

2. Average House Prices and Manufacturing

Before presenting our results, we emphasise that our primary aim in this paper is to contrast home price dynamics in manufacturing-heavy locations versus their peers in manufacturing-light areas. Our overall goal is not causal identification. However, in addition to our baseline OLS specifications, we also exploit a widely used shift-share IV strategy and present evidence consistent with an income channel causally linking manufacturing exposure to subsequent income growth, and therefore to home prices via demand.

We begin by quantifying differences in the growth rate of average house prices across locations with different pre-existing manufacturing exposure. Since we are not yet interested in distributional shifts, we use ZTRAX average CZ-level house price indexes. Specifically, let

Table 1. *House Prices and Manufacturing.*

Percent change in house prices	(1) OLS	(2) OLS	(3) IV	(4) OLS	(5) OLS
Manufacturing share	−0.488***	−0.467***		−0.118**	0.005
	(0.142)	(0.137)		(0.048)	(0.042)
Percent change in wage income			0.771**		
			[0.0325]		
Controls	Yes	Yes	No	Yes	Yes
Fixed effects	No	Census division	Census division	No	Census division
Underlying house transactions	19,670,168	19,670,168	19,670,168	43,686,431	43,686,431
Commuting zone observations	179	179	179	179	179
Years	2001–06	2001–06	2001–06	2001–15	2001–15
Adjusted R^2	0.226	0.249		0.128	0.285

Notes: Regressions run at the commuting zone level with the average percentage house price growth over 2001–06 or 2001–15 on the manufacturing employment share in 2000 or wage income growth 2001–06. Controls include the Saiz (2010) housing supply elasticity, the percentage of routine cognitive jobs, the college educated working share, the female working share and the foreign working share. Standard errors in parentheses are clustered at the state level. ** and *** denote significance at the 5% and 1% levels. In column (3), we follow Adao *et al.* (2019) and instrument for CZ-level wage growth with a shift-share IV based on three-digit 1998 NAICS employment shares. The brackets under the coefficient report the *p*-value associated with the second-stage estimate computed using the Adao *et al.* (2019) procedure, and the main text reports first-stage results.

$p_{c,t}$ denote the log of the average house price in CZ c in year $t \in \{2001, 2006\}$. Then, we consider a regression

$$\Delta p_{c,t} = \alpha + \beta M_{c,2000} + \gamma X_{c,2000} + \delta_{div(c)} + \varepsilon_{c,t}, \tag{1}$$

where $\Delta p_{c,t}$ denotes the change in log prices between the two time periods, and $M_{c,2000}$ denotes the share of manufacturing in total employment for CZ c in the year 2000. We add standard CZ-level controls (denoted by $X_{c,2000}$) such as the (*i*) educational composition, (*ii*) share of female workers, (*iii*) share of foreign-born workers and (*iv*) share of workers in routine cognitive occupations. We also augment the regression with controls for housing supply elasticities (Saiz, 2010). Since Figure 1 suggests that the geographical distribution of manufacturing is not random, but rather concentrated in specific regions of the United States, we also include census division indicators in our regressions. As such, our identification of the variation predicted by exposure to manufacturing comes from changes *within* census division. Finally, because turbulence in the US housing market formed the epicentre of the financial crisis and Great Recession, we initially focus on the 2001–06 period, holding constant the various shares discussed above at their 2000 values. We later extend our analysis to the 2001–15 period.

Table 1 contains our first main results. In column (1), we see that home prices in areas with a higher manufacturing share in 2000 grew more slowly over the 2001–06 period than their peers, after controlling for the set of local labour market characteristics discussed above. In column (2), we verify that the negative association survives the inclusion of census division fixed effects. The coefficient estimates in both columns exhibit high statistical precision. In other words, even once we control for broad geographical trends and a host of potential confounding factors in labour and housing markets, our results indicate that manufacturing-heavy areas failed to see house price growth as high as their low-manufacturing peers.

The magnitudes of the differences are large. Consider a CZ at the 75th percentile of manufacturing exposure, which has a 2000 manufacturing employment share 6.7 percentage points higher than a CZ at the 25th percentile. Our estimates in column (2) predict that house price growth from 2001–06 in that same CZ will be $0.467 \times 6.7 \approx 3.1\%$ lower per year than in the manufacturing-light region. The drop of 3.1% in house price growth each year represents a variation of $3.1/7.0 \approx 44\%$ relative to the interquartile range of house price growth in the full sample.

While broad, the coverage of the ZTRAX dataset (Zillow Group, 2020) is not uniform across the United States. Hence, for robustness purposes, we run the same regressions using CZ-level house price indexes that we construct from locally aggregated data from the FHFA (see Online Appendix A for details). As evidenced in Online Appendix Table A2, we find very similar results using this alternative to the ZTRAX database, despite its wider geographical coverage.[3]

2.1. *Exposure to Manufacturing and Labour Market Outcomes in the Data*

Our results throughout the paper are consistent with the hypothesis that the relationship between initial manufacturing exposure and house price dynamics operates through an income channel: in regions with a higher proportion of employment in manufacturing, the secular industrial decline led to disproportionate job losses and income stagnation relative to other regions. This, in turn, morphed into relatively lower house price growth. In Online Appendix B, we formalise this intuition in a model. We note that a link between exposure to manufacturing and labour market outcomes has been studied before (Ramey, 2018; Charles *et al.*, 2019), and is confirmed in panel A of Online Appendix Table A3 for our level of spatial and time coverage.[4] As this table suggests, we find that CZs heavy in manufacturing in 2000 experienced a steeper drop in manufacturing employment over the 2001–06 period, a more pronounced positive *change* in the likelihood of not working, and lower growth of average wages.

To provide additional evidence in favour of the income channel for house prices, we rely on a standard shift-share industry employment IV approach. The underlying identifying assumption is that the past composition of industries in a CZ interacted with national employment dynamics does not impact local house price dynamics except through wage growth.[5] Specifically, in our approach we implement the algorithm advocated for by Adao *et al.* (2019), with details of our IV construction in Online Appendix A.[6] The results are presented in column (3) of Table 1. First we note that the first-stage results confirm that the initial industry composition has reasonable explanatory power for wage growth: the estimated coefficient is 0.75, significant at the 5% level with an F-statistic of 22.3. The second-stage coefficient in column (3) suggests that the instrumented variation in income is an economically and statistically significant determinant of house price dynamics. The IV coefficient indicates that a 1% change in wage income is associated with a 0.77% increase in house prices.

[3] To further validate the ZTRAX data, we note the following points. First, we note that our 179 CZs include a disproportionate fraction of the overall population in the FHFA data. Specifically, our 179 CZs account for 77% of the total population aged 15–65 from the 657 CZs used for regression (1) in Online Appendix Table A2. Hence, this reinforces the view that our Zillow CZs account for a vast majority of the US data. Second, in the FHFA data restricted to the CZs for which we have ZTRAX coverage, we estimate a manufacturing coefficient of −0.444 in column (2) of Online Appendix Table A2, which is almost identical to the −0.467 coefficient reported in column (2) of Table 1.

[4] The regressions in Online Appendix Table A3 rely on the 741 CZs in the United States, while for the regressions in Table 1, we limit ourselves to the CZs we cover in the Zillow data.

[5] In our view, the most relevant variable to include is labour income. One could imagine a world in which population and employment do not change, yet despite no change in employment, the fall in income would show up in house prices.

[6] See Table 6, panel A in Adao *et al.* (2019). We note that a robust applied econometrics debate over the properties and potential drawbacks of shift-share approaches has arisen (Borusyak *et al.*, 2018; Jaeger *et al.*, 2018; Adao *et al.*, 2019; Goldsmith-Pinkham *et al.*, 2020).

2.2. *Extending the Data: 2001–15*

Our analysis up to this point has focused on the 2001–06 period, in order to avoid conflating the role of manufacturing exposure with that of factors specific to the 2007–9 financial crisis that was associated with a major disruption in housing markets. Yet, one may wonder whether exposure to manufacturing in 2000 may still predict house price dynamics beyond the Great Recession. Columns (4) and (5) in Table 1 show results extending our sample to 2015. As expected, the results suggest that the predictive power of past manufacturing intensity for average house price growth dissipates as structural adjustment occurs at the local level in the aftermath of the Great Recession. Column (5) reveals that manufacturing exposure in 2000 does not predict house price growth once we include region fixed effects. Yet, as we discuss next in Section 3.2, when we analyse heterogeneity across the housing price distribution, we find that low-price segments of this distribution in manufacturing-heavy areas did continue to experience lower cumulative house price growth through 2015.

3. Manufacturing Exposure and the House Price Distribution

We now turn our attention to our main question of interest: is lower house price growth in manufacturing-heavy areas distributionally neutral? Or does exposure to manufacturing predict more pronounced reductions in house price growth for low-price homes than for high-value homes? Under the income channel, we would expect heterogeneous patterns if manufacturing workers, who have taken the brunt of the detrimental labour market shifts during deindustrialisation, live disproportionately in a specific portion of the house price distribution. See Online Appendix B for a simple model that formalises this idea.

Indeed, as is evident in Figure 2, manufacturing workers are disproportionately represented in the lowest tercile of the housing distribution and much less prevalent in the highest tier.[7] Moreover, under the income channel, we would expect the income of households living in lower-priced houses to be more sensitive to the manufacturing share than that of households in high-priced houses. Panel B of Online Appendix Table A3 confirms this: manufacturing exposure predicts both economically and statistically larger declines in wage growth for individuals in lower-priced homes.[8] Hence, under an income channel, we would expect lower-priced homes to grow even more slowly in value in manufacturing-heavy areas than their high-priced peers, a hypothesis that we test next.

3.1. *Location in the House Price Distribution*

To study distributional price dynamics at the CZ level, we need to first construct a distribution of local house prices and allocate each transaction to its relevant position in the distribution. We pursue two approaches.

The first solution, our baseline, uses a hedonic approach based on projecting house prices on a list of observable house characteristics. In a nutshell, we first estimate the loading of house

[7] To construct Figure 2 we exploit self-reported census house price valuations and compute the fraction of manufacturing workers in each of three equally sized home-price terciles defined at the CZ level. We use this source as we do not have information in the ZTRAX dataset regarding the occupation of the sellers, requiring this use of IPUMS home valuation data.

[8] Moreover, we note that column (1) in panel A, which uses all 741 CZs in the US data, and column (1) in panel B, which uses the ZTRAX geographic coverage, reveal very similar wage growth patterns after exposure to manufacturing. This similarity again reinforces the relevance and representativeness of the ZTRAX data (Zillow Group, 2020).

Fig. 2. *Manufacturing Workers by House Price Tercile.*
Notes: For a given house price tercile, the figure plots the distribution across commuting zones of the share of manufacturing workers living in that house price category. The underlying data is the US Census IPUMS micro-data in the year 2000.

prices on various amenities in 2001.[9] Next, using these loadings and the amenities of houses sold in 2006, we can rank each house within the CZ-specific distribution of 2001 home prices. We therefore ensure that statements made about 'high-value' or 'low-value' homes reflect consistent comparisons and valuations of home characteristics across time periods. We then construct for each CZ and segment of the 2001 housing distribution the growth rate of mean house prices over our 2001–06 sample period, which becomes our dependent variable. See Online Appendix C for a more detailed discussion of the approach.

The second solution follows the methodology of the Standard & Poor's Case-Shiller home price indexes and relies only on repeated sales of individual properties. In this approach, we can directly assign a house to a part of the CZ-level price distribution in the base year (in our case 2001). We then compute the average change in house prices for these repeat sales for each segment of the initial CZ-specific house price distribution. The repeat sales approach allows us to directly study the distributional dynamics of house prices while holding permanent unobserved characteristics fixed.

The trade-offs between the two approaches are clear. The repeat sales approach does not require us to control for a pre-specified list of house characteristics. The hedonic approach, on the other hand, allows for wider data coverage across time and space. Despite the large size of our dataset, the repeat sales approach results in a significantly reduced sample size. Hence, to present a baseline relying on as broad an underlying sample as possible, we first compute

[9] The amenities we include are square footage, year of construction, number of rooms, number of bathrooms, number of bedrooms, number of stories, the presence of a garage and a set of ZIP-code-level dummies. See Online Appendix C for further details.

Table 2. *House Prices and Manufacturing across the Distribution.*

Percent change in house prices Years	(1) 2001–06	(2) 2001–15
Manufacturing share	-0.690^{***}	-0.088^{*}
∗ Low-value houses	(0.151)	(0.047)
Manufacturing share	-0.583^{***}	-0.029
∗ Mid-value houses	(0.142)	(0.041)
Manufacturing share	-0.448^{***}	0.071
∗ High-value houses	(0.123)	(0.045)
Controls	Yes	Yes
Fixed effects	Census division	Census division
Underlying house transactions	19,670,168	43,686,431
Commuting zone × tercile obs.	535	535
Adjusted R2	0.286	0.366

Notes: Regressions run at the commuting zone × house price tercile level with the percent change in average house prices for the relevant cell on the manufacturing employment share in 2000. The terciles reflect 2001 home values. Controls include the Saiz (2010) housing supply elasticity, the percentage of routine cognitive jobs, the college educated working share, the female working share and the foreign working share. Standard errors in parentheses are clustered at the state level. * and *** denote significance at the 10% and 1% levels.

local house price changes across the distribution using the hedonic method. However, in Online Appendix Table A4 we show that the main conclusions of this section hold if we instead focus on repeat sales, despite a much smaller number of observations.

3.2. *Results*

We divide all transactions within each CZ in our 2001 benchmark year into three equally sized price terciles or segments: low value, mid value and high value. That is, for each CZ, we split the distribution of house prices in 2001 into three equally sized bins.[10] Using the hedonic pricing approach discussed above, we then map homes in later years into the same three segments based on a consistent valuation of their characteristics.

Our regression specification is given by

$$\Delta p_{c,s,t} = \alpha + \theta_s + \sum_s \beta_s \mathbb{1}_s \times M_{c,2000} + \gamma X_{c,2000} + \delta_{div(c)} + \varepsilon_{c,t},$$

where the growth of average prices from 2001 to 2006 within each CZ × segment cell provides our main outcome measure. In the equation above, $s \in \{1, 2, 3\}$ denotes the segment or tercile of the housing distribution, θ_s denotes the segment fixed effect and $\mathbb{1}_s$ is an indicator function that equals 1, when an observation belongs to the relevant housing segment and that is interacted with the manufacturing employment share. We continue to control for all the variables $X_{c,2000}$ discussed in the context of (1). This specification allows us to investigate whether the association between manufacturing shares and subsequent house price growth differs when moving from low-value (with $s = 1$) to mid-value (with $s = 2$) to high-value homes (with $s = 3$).

Column (1) of Table 2 reports the results. Houses in all terciles appreciated more slowly on average in high-manufacturing areas. These dynamics are however heterogeneous: while low-value homes display sharply lower growth in the face of manufacturing exposure, this difference

[10] We note that in unreported results we obtain similar findings when using a finer segmentation of the distribution, e.g., with quintiles or deciles.

is more muted for their high-value neighbours. In other words, manufacturing exposure predicts relatively *more* within-region house price inequality, not less.

The magnitudes at work here prove large once again. For a high-manufacturing CZ at the 75th percentile of manufacturing shares in 2000, the coefficient estimates in column (1) reveal that low-value homes experienced $0.690 \times 6.7 \approx 4.6$ percentage points lower yearly subsequent house price growth than the same segment in a light-manufacturing CZ at the 25th percentile of exposure. By contrast, high-value homes saw a relatively lower price growth of around $0.448 \times 6.7 \approx 3.0$ percentage points per year, i.e., a differential that is a third smaller in magnitude.

As discussed above, our main analysis covers 2001–06. Recall that column (5) of Table 1 showed that the predictive power of past manufacturing intensity dissipates over the longer 2001–15 horizon during a period spanning the Great Recession. What are the dynamics across the housing distribution during this longer horizon? Column (2) in Table 2 shows tercile-level results for 2001–15. Naturally, exposure to manufacturing predicts more muted differences over the longer horizon. But remarkably, and in contrast to the dynamics of average prices, our estimates reveal persistent and long-lasting differences in the price growth of the lowest-value homes in manufacturing-heavy areas. By contrast, over this period covering one and a half decade, their high-value neighbours in manufacturing-exposed areas had not experienced any significant difference in house price growth relative to their manufacturing-light peers.

All in all, these empirical patterns are consistent with a framework in which manufacturing exposure reduces income growth, feeding into declines in the price of both homes overall and especially the price of the lowest-value homes.

4. Shifts in House Price Inequality

In light of the outsized importance of housing in overall household wealth, we next exploit the full distribution of home prices in our ZTRAX micro-data to document some overall shifts in house price inequality over our sample period. To isolate the role of the lowest-value homes in the observed variation in house price inequality, we engage in a series of simple counterfactual accounting exercises.

4.1. *The Evolution of House Price Inequality*

We start by documenting the evolution of house price inequality over our sample periods and across the distribution. Comparing the first two columns in the top panel of Table 3, we find that the standard deviation of log house prices in fact declined between 2001 and 2006, from 87.9% to 84.9%. These dynamics, however, were strikingly different across the house price distribution: while the dispersion fell for high-value (86.2% to 79.4%) and mid-value homes (85.3% to 80.6%), it instead *increased* for the lower-priced houses (92% to 93.7%). Over the longer 2001–15 sample period, the overall dispersion rose from 87.9% to 92.8%. Again, the evolution is very different across the house price distribution: while the standard deviation fell significantly for the high-value homes (86.2% to 72.6%), and remained practically the same for the mid-value homes, it rose from 92% to 113.2% for the bottom tercile. Next, using counterfactuals we investigate the roles of region- and segment-level heterogeneity in driving aggregate variation in house price inequality.

Table 3. *House Price Inequality.*

Panel A: observed data			
SD of log prices in year:	2001	2006	2015
All houses	0.879	0.849	0.928
Low-value houses	0.920	0.937	1.132
Mid-value houses	0.853	0.806	0.876
High-value houses	0.862	0.794	0.726
Panel B: counterfactuals			
SD of log prices for all houses in year:	2001	2006	2015
Observed	0.879	0.849	0.928
Removing all regional and segment house price growth differences	0.879	0.784	0.854
Removing low-value house price growth differences	0.879	0.819	0.873
Removing manufacturing-predicted low-value house price growth differences	0.879	0.859	0.920
Underlying house transactions	2,664,242	2,255,561	3,196,915

Notes: The top panel reports observed inequality in house prices in various categories in the indicated year. The bottom panel reports the inequality for all homes in each year under various counterfactuals described in the text.

4.2. *A Simple Accounting Framework*

To account in more detail for these shifts in inequality, we introduce some additional notation and an accounting framework that can be mapped directly to our data. We write the log price $p_{h,c,s,t}$ of house h in CZ c in home value segment s in year t as

$$p_{h,c,s,t} = \mu_{c,s,t} + \sigma_{c,s,t}\varepsilon_{h,c,s,t}.$$

Here $\mu_{c,s,t}$ is the average price in the CZ $\times$ segment $\times$ year cell and $\sigma_{c,s,t}$ is the standard deviation (SD) of the same cell. The value $\varepsilon_{h,c,s,t}$ represents the normalised home price, featuring zero mean and unit standard deviation within a cell. We can directly and easily compute estimates of each of the values in the decomposition above from our micro-data. This simple accounting framework reveals that an increase in the dispersion or inequality of overall house prices can in principle stem from one or a combination of three channels: (*i*) an increase over time in within-cell dispersion $\sigma_{c,s,t}$ that is common across all $c \times s$ cells, (*ii*) heterogeneous shifts over time in the within-cell dispersions $\sigma_{c,s,t}$ or (*iii*) heterogeneity in the growth over time of average prices $\mu_{c,s,t}$ across cells.

Our empirical results so far, which document distinct growth rates for average prices $\mu_{c,s,t}$ in CZ $\times$ segment cells, map directly to the third channel. Next, we seek to quantify the importance of this growth rate heterogeneity through a series of simple counterfactual exercises.

4.3. *The Role of Heterogeneity across Regions and House Price Segments*

In our first scenario, we shut down all heterogeneity in the growth of mean house prices $\mu_{c,s,t}$ across a CZ (c) or house price tercile (s), and recompute the dispersion of house prices in both 2006 and 2015. The second row of the bottom panel of Table 3 reveals that this counterfactual generates a much stronger decrease in inequality between 2001 and 2006, from 87.9% to 78.4% instead of 84.9%. Over the 2001–15 period, which saw an increase from 87.9% to 92.8% in the SD, the counterfactual instead generates a *drop* to 85.4%. In other words, the changes in house price inequality are overwhelmingly driven by heterogeneity in house price growth *across* regions and price segments, instead of variation *within* region-segment cells.

4.4. *The Role of Low-Value Houses*

In our second counterfactual, we focus on the role of low-value homes by shutting down heterogeneity across regions in the growth rate of this segment alone, allowing for inter-regional differences in the mid- and high-price segments. The third row in the bottom panel of Table 3 reveals that the 2006 SD of log prices under this counterfactual scenario equals 81.9% (starting from 87.9% in 2001); hence, without the contribution from the lower-price segment, inequality in house values would have been significantly lower than in reality (84.9%). The crucial role played by the cross-regional divergence of the lowest-priced homes is also evident from the 2015 counterfactual: between 2001 and 2015, ignoring this margin generates a *decline* in overall inequality, from 87.9% to 87.3%, while the observed dispersion in 2015 is 92.8%. We can therefore conclude that differences in the average growth rates of *only* the lowest-valued homes across regions have been a significant positive contributor to the inequality in house prices over our sample period.

For our third exercise, we narrow the analysis even further: we aim to identify the portion of the change in the overall inequality in house prices that is predicted *solely by the manufacturing exposure of the lowest-value homes*. This effectively shuts down only the heterogeneity in the growth rates of the lowest-value segment of homes that is predicted by our regressions in Table 2. The last row of Table 3 shows that, for 2001–06, without the contribution of this very narrow source of heterogeneity, inequality in house values would have been only slightly lower than in reality, accounting for 6% of the observed change. If we instead consider the 2001–15 period, the resulting increase in SD would have been 16% smaller than the observed increase.[11] In other words, heterogeneity in average growth rates for the lowest-value homes predicted only by their exposure to manufacturing accounts for around one-sixth of the increase in total house price inequality over the 2001–15 period.

All in all, we conclude that our counterfactual exercises lead to the striking conclusion that regional differences, and in particular those for the very lowest-value homes, prove critical for understanding increased cross-sectional house price inequality.

5. Conclusions

Our analysis leverages a rich dataset of tens of millions of house price transactions tracked by Zillow. We show that areas with higher exposure to the US manufacturing sector experienced lower growth in home prices on average in recent years. Furthermore, the lowest-value homes in these regions experienced an even heavier decline in price growth relative to their higher-value neighbours. In other words, manufacturing exposure predicts shifts in both cross-region and within-region inequality in house prices.

In an exercise leveraging our full distribution of house prices at the microlevel, we show that a recent increase in house price inequality is fully accounted for by heterogeneity across regions in the growth of prices of the lowest-value homes, exactly those dwellings disproportionately predicted to grow more slowly in the face of manufacturing exposure.

Thus, we conclude that the relative decline of manufacturing-heavy areas extends far beyond income and employment flows to include shifts in important local asset prices.

[11] For 2001–06, the contribution of this channel is about 6% of the observed change, i.e., $(0.847 - 0.879)/(0.849 - 0.879)$. For 2001–15, it would have been $0.920 - 0.879 = 0.041$, which is 16% smaller than the observed increase to 0.928.

Paula and Gregory Chow Institute for Studies in Economics, Xiamen University, China
University of California San Diego & CEPR & NBER, USA
Landed, Inc., USA
Boston University & NBER, USA
HEC Montréal, Canada

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Adao, R., Kolesár, M. and Morales, E. (2019). 'Shift-share designs: Theory and inference', *The Quarterly Journal of Economics*, vol. 134(4), pp. 1949–2010.
Ahn, S., Kaplan, G., Moll, B., Winberry, T. and Wolf, C. (2018). 'When inequality matters for macro and macro matters for inequality', *NBER Macroeconomics Annual*, vol. 32(1), pp. 1–75.
Alder, S., Lagakos, D. and Ohanian, L. (2017). 'Labor market conflict and the decline of the rust belt', Working paper, Boston University.
Autor, D.H., Dorn, D. and Hanson, G.H. (2013). 'The China syndrome: Local labor market effects of import competition in the United States', *American Economic Review*, vol. 103(6), pp. 2121–68.
Autor, D., Dorn, D. and Hanson, G.H. (2019). 'When work disappears: Manufacturing decline and the falling marriage market value of young men', *American Economic Review: Insights*, vol. 1(2), pp. 161–78.
Bárány, Z.L. and Siegel, C. (2018). 'Job polarization and structural change', *American Economic Journal: Macroeconomics*, vol. 10(1), pp. 57–89.
Berger, D., Guerrieri, V., Lorenzoni, G. and Vavra, J. (2017). 'House prices and consumer spending', *The Review of Economic Studies*, vol. 85(3), pp. 1502–42.
Borusyak, K., Hull, P. and Jaravel, X. (2018). 'Quasi-experimental shift-share research designs', Working paper, London: University College.
Case, A. and Deaton, A. (2017). 'Mortality and morbidity in the 21st century', *Brookings Papers on Economic Activity*, vol. 2017, pp. 397−443.
Charles, K.K., Hurst, E. and Notowidigdo, M.J. (2016). 'The masking of the decline in manufacturing employment by the housing bubble', *Journal of Economic Perspectives*, vol. 30(2), pp. 179–200.
Charles, K.K., Hurst, E. and Schwartz, M. (2019). 'The transformation of manufacturing and the decline in US employment', *NBER Macroeconomics Annual*, vol. 33(1), pp. 307–72.
Feyrer, J., Sacerdote, B., Stern, A.D., Saiz, A. and Strange, W.C. (2007). 'Did the rust belt become shiny? A study of cities and counties that lost steel and auto jobs in the 1980s [with comments]', *Brookings-Wharton Papers on Urban Affairs*, vol. 2007, pp. 41–102.
FHFA. (2020). 'Federal Housing Finance Agency (FHFA) house price index', https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx (last accessed on 25 October 2017).
Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020). 'Bartik instruments: What, when, why, and how', *American Economic Review*, vol. 110(8), pp. 2586–624.
Guren, A.M., McKay, A., Nakamura, E. and Steinsson, J. (2021). 'Housing wealth effects: The long view', *Review of Economic Studies*, vol. 88(2), pp. 669–707.
Howard, G. and Liebersohn, J. (2018). 'The geography channel of house price appreciation', Working paper, University of Illinois.
Howard, G. and Liebersohn, J. (2021). 'Regional divergence and house prices', Working paper, University of Illinois.
Jaeger, D.A., Ruist, J. and Stuhler, J. (2018). 'Shift-share instruments and the impact of immigration', Working paper, NBER.
Kahn, M.E. (1999). 'The silver lining of Rust Belt manufacturing decline', *Journal of Urban Economics*, vol. 46(3), pp. 360–76.
Kaplan, G., Mitman, K. and Violante, G. (2018a). 'Consumption and house prices in the Great Recession: Model meets evidence', Working paper, Princeton University.
Kaplan, G., Moll, B. and Violante, G.L. (2018b). 'Monetary policy according to HANK', *American Economic Review*, vol. 108(3), pp. 697–743.

Liebersohn, J. (2017). 'Housing demand, regional house prices, and consumption', Working paper, Massachusetts Institute of Technology.

Notowidigdo, M.J. (2020). 'The incidence of local labor demand shocks', *Journal of Labor Economics*, vol. 38(3), pp. 687–725.

Piazzesi, M. and Schneider, M. (2016). 'Housing and macroeconomics', in (J.B. Taylor and H. Uhlig, eds.), *Handbook of Macroeconomics*, vol. 2, pp. 1547–640, Amsterdam: Elsevier.

Ramey, V.A. (2018). 'Comment on "The transformation of manufacturing and the decline in US employment"', in (M. Eichenbaum and J.A. Parker, eds.), *NBER Macroeconomics Annual 2018*, vol. 33, pp. 380–8, Chicago: University of Chicago Press.

Saez, E. and Zucman, G. (2016). 'Wealth inequality in the United States since 1913: Evidence from capitalized income tax data', *The Quarterly Journal of Economics*, vol. 131(2), pp. 519–78.

Saiz, A. (2010). 'The geographic determinants of housing supply', *The Quarterly Journal of Economics*, vol. 125(3), pp. 1253–96.

Song, J., Price, D.J., Guvenen, F., Bloom, N. and Von Wachter, T. (2018). 'Firming up inequality', *The Quarterly Journal of Economics*, vol. 134(1), pp. 1–50.

U.S. Census Bureau. (2019). 'U.S. census data for social, economic, and health research', https://usa.ipums.org/usa/ (last accessed on 28 February 2018).

U.S. Census Bureau. (2020). 'County Business Patterns (CBP)', https://www.census.gov/programs-surveys/cbp.html (last accessed on 28 April 2017).

Zillow Group. (2020). 'Zillow's transaction and assessment database (ZTRAX)', http://www.zillow.com/ztrax/ (last accessed on 19 April 2017).