

MEASURING GENTRIFICATION



Ron Bekkerman
Strategic Advisor
Cherre

Donal Warde
Entrepreneur in Residence
Tenney 110

Maxime C. Cohen
Professor
McGill University

Despite the fact that a lot of research has been done on gentrification at the macroeconomic level, it has not been extensively studied at the microeconomic level using large-scale granular data. This article attempts to correct that.

Gentrification is regarded as one of the main driving forces of the real estate industry: investors look for neighborhoods that start gentrifying, with the intention to invest in them and, ultimately, generate high returns.

Despite the fact that a lot of research has been done on gentrification at the macroeconomic level, it has not been extensively studied at the microeconomic level using large-scale granular data. This article attempts to correct that. Starting with an examination of how gentrification can be measured at a neighborhood level and utilized to compare between neighborhoods. We then investigate whether gentrification positively affects neighborhood development.

We propose a Normalized Gentrification Index (NGI) of a neighborhood and use it to compare different neighborhoods. In this initial study, we use U.S. counties for neighborhoods, while we plan fine-grained studies in the future. Depending on the underlying data availability, the NGI can be computed at any level of granularity: for US states, counties, metropolitan areas, submarkets, zip codes, census tracts, and so forth.

In this initial study, we chose to focus on the level of counties. We constructed a dataset of 630 US counties for which we had enough data to compute NGI values for each year between 2008 and 2021 and demonstrate their usefulness for real estate investors. We identified a statistical connection between the dynamics of NGI over time and the key performance metric of counties, measured by their residential property prices. We illustrated our findings on two use cases. While we chose residential property prices due to the data abundance, these findings may also be relevant for multifamily, which remains a dominant institutional asset class.

NICHE TRANSFORMATION

In recent years, multifamily investors have experienced a stronger bull market as compared to other asset classes.¹ However, as the rising tide has lifted (almost) all boats, a potentially receding tide necessitates a more selective investment strategy.

Just as asset allocation is a primary driver of portfolio returns in a mixed asset class portfolio (e.g., stocks, bonds, etc.),² selection decisions at the metropolitan statistical area (MSA) level or county level drive the bulk of returns within the real estate investment universe, with individual asset selection relegated to the second place in the return composition hierarchy. Our analysis indicates that—in the majority of cases—a residential property’s price can be predicted within a 20% margin solely based on its county, which validates the importance of county-level selection.

Real estate portfolio management can be enhanced by a proactive, “top-down” market selection approach highlighting locations with potential for above-market return profiles. “Bottom-up” individual investment-level strategies can be supplemented rather than replaced with these analyses. Knowledge of high-potential target counties within the investable universe can ultimately be an alpha enhancer in a competitive investment market.

But there are more than 3,000 counties in the US, so where shall we start?

Our goal is to go beyond trailing signals and momentum-driven strategies and to attempt to create forward-looking signals that may present above-average returns in future years. Specifically, can individual- or family-level wealth inflows predict residential real estate price patterns?

RETHINKING POPULATION DATA

Our primary model in this analysis can be characterized as the gentrification of a neighborhood, which is defined as the net wealth inflow relative to the wealth of the neighborhood. We propose an NGI of a neighborhood and use it to compare different neighborhoods.

The NGI is calculated as the difference between the estimated wealth excess of the in-migration and out-migration of the focal county. Counties with a high NGI have positive wealth inflows, meaning that when we estimate the wealth of in-migrants and out-migrants, there is a net increase. And vice versa for negative NGIs. We consider our NGI as a proxy for gentrification measure.


Obviously, NGI does not fully explain property price changes. In our dataset of 630 US counties, the NGI variation from 2010 to 2020 only weakly correlates with residential price change ($r = 0.293$), which has a marginally higher correlation compared to a simple demographic indicator such as a population increase estimate ($r = 0.177$). However, population changes are measured at a very low cadence (typically, once a decade) and cannot reliably be predictive. NGI, in contrast, can be measured quarterly. Moreover, we report a statistical connection between NGI and property price increase, which indicates a predictive power that NGI has on prices. While NGI is not yet fully studied, our first results look promising.

County-level housing supply can often be inelastic, which we assumed would increase prices as new residents bid up the existing real estate inventory. As expected, population growth is more correlated with residential real estate prices in high-barrier-to-entry markets that are characterized by low growth in new housing units.

However, as we analyzed the NGI of these counties based on the bifurcation between high- and low-barrier-to-entry counties, counterintuitively, we saw that the NGI had a stronger relationship with low-barrier-to-entry markets than high-barrier-to-entry markets (if we consider only low-barrier-to-entry counties, the correlation goes up from 0.293 to 0.456).

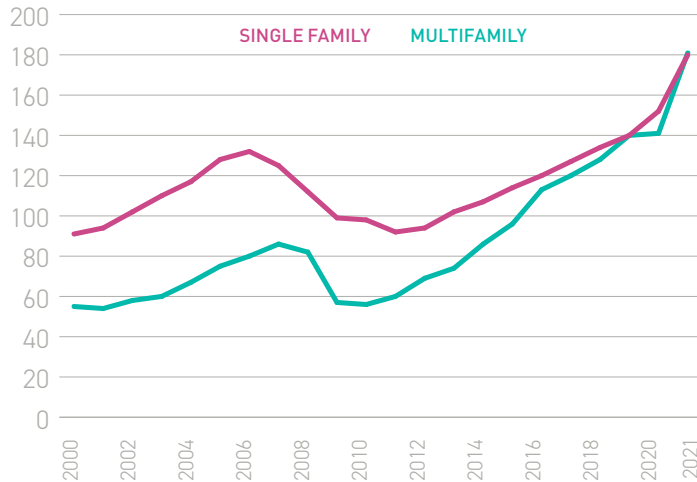
There are several potential explanations for this dynamic. Firstly, in supply-constrained markets, the utility of adding the wealth dimension to the net population growth may be reduced, because the primary effect will be on adding incremental demand, irrespective of wealth levels. Secondly, elastic supply markets may be lower-cost markets, which are more sensitive to incremental wealth inflows, because the economic effect of this net wealth addition might be more significant throughout the county. To illustrate these dynamics, we have analyzed two counties representing strictly divergent trends: Pinal County, AZ and Cape May County, NJ.

Our empirical study focuses on the statistical connection between NGI and single-family price dynamics, primarily because we have a large number of single-family transactions that ensure statistical significance. We claim that our results are likely to be valid for the multifamily asset class as well. While there are two orders of magnitude fewer multifamily transactions than single-family transactions (at the per-county per-year level), multifamily prices and single family prices highly correlate ($r = 0.9$). *Exhibit 1* illustrates of the similarity of median prices per square foot in the multifamily and single-family markets nationwide.



Knowledge of high-potential target counties within the investable universe can ultimately be an alpha enhancer in a competitive investment market.

EXHIBIT 1: MEDIAN PRICE PER SQUARE FOOT, MULTIFAMILY AND SINGLE-FAMILY



The NGI is calculated as the difference between the estimated wealth excess of the in-migration and out-migration of the focal county.

WHAT IS A NORMALIZED GENTRIFICATION INDEX?

Intuitively, the gentrification of a neighborhood depends on both inbound and outbound migration; namely, a neighborhood is gentrifying if higher-income people are moving into it and/or lower-income people are moving out. The notions of “richer” and “poorer,” for the sake of this article, are relative to the average wealth of the neighborhood’s residents.

Gentrification is cyclic: while richer people are moving in and becoming residents, the average level of wealth of the neighborhood residents rises until it becomes high enough so that people moving in are no longer richer than the average, and then the opposite process of “de-gentrification” starts until the average wealth goes down, for the gentrification to pick up again.

We measure gentrification by utilizing the notion of “wealth excess” of a household i with respect to the average wealth \widehat{W}_n of a household in neighborhood n . We define the wealth excess of household i in neighborhood n as $WE_{i,n} = S(W_i - \widehat{W}_n)$, where W_i is an estimated wealth of i , and S is a smoothing function that levels down extreme differences between W_i and \widehat{W}_n . Smoothing is necessary to take care of edge cases and data issues as explained below. We highlight that $WE_{i,n}$ might be negative if i is poorer than the average in n .

We use a propriety algorithm to estimate household wealth, which bases our estimation on home values, as homeownership is typically the primary source of wealth for a family. For the purposes of this analysis, if we could not obtain the value of a family’s home, we approximate it with an average home value in the same neighborhood.

For each neighborhood n and a certain time period (typically, one year), we define an incoming wealth excess as the sum of wealth excesses of families moving into the neighborhood: $WE_n^{in} = \sum_i k_n^{in} WE_{i,n}$, where k_n^{in} is the number of families that moved into the neighborhood over the year. Analogously, wealth excess of families moving out of the neighborhood is defined as $WE_n^{out} = \sum_i k_n^{out} WE_{i,n}$, where k_n^{out} is the number of families that moved out of the neighborhood over the year. The gentrification index of the neighborhood is then defined as the difference between the incoming wealth excess and the outgoing wealth excess: $GI_n = WE_n^{in} - WE_n^{out}$.

Because neighborhoods are different in size and overall wealth, GI values of neighborhoods are not necessarily on the same scale and therefore cannot be used for neighborhood comparisons. To enable such comparisons, we normalize the GI of a neighborhood by the collective wealth W_n of the neighborhood’s residents. We then define the NGI of neighborhood n as: $NGI_n = GI_n / S(W_n)$. Note that we apply the same smoothing function S to W_n , to ensure that both the numerator and the denominator are consistent with each other.

As mentioned, the NGI of a neighborhood will be positive if many richer families are moving in, and many poorer families are moving out. It will be around zero if few families are moving in and out, or those moving families are of average wealth. The index will be negative if many poorer families are moving in, and many richer families are moving out. Note that if the smoothing function S was not applied, the NGI value might have been primarily influenced by very few wealth excess values $WE_{i,n}$ if they are truly outstanding. For example, if the wealth of a family happens to be estimated at three orders of magnitude compared to the average wealth of a neighborhood’s residents, then this family’s wealth excess would be the dominant factor in the neighborhood’s NGI, which is obviously an undesired outcome.

DATASET

NGI can be defined for “neighborhoods” at any level of granularity, such as MSAs, zip codes, and census tracts. In this work, we analyze NGIs at the coarse level of counties.


Out of 3,243 US counties and county equivalents, we select only the counties for which we have enough assessed property value data, transaction price data, and relocation data. We also make sure that assessed property values are not too low when averaged over all properties in a county (low average values per square foot might be due to a data artifact) and that property values and transaction prices would be in the same ballpark (very large difference between them would be caused by another data artifact). The selection process resulted in a dataset with 630 counties, covering the majority of the largest US counties that are also attractive destinations for investment capital. Out of the 398 largest US counties reported by the World Population Review,³ 334 counties are present in our dataset. Those that are not present are either located in non-disclosure states (so they lack transaction price data) or their data is only partially populated in public data sources. For example, assessed property value data of Cook County,

IL (the second largest US county) is populated only from 2017. *Exhibit 2* shows our dataset of 630 counties displayed on the map, where the color represents the size of the county in terms of the number of residential properties.

Our relocation dataset consists of 14.4M relocation records spread over the 2000-2021 time period. Each record contains the name of a person, their previous address, their new address, and the year of the relocation. Names of people are not used in this research. The relocation dataset was constructed from public data sources (property sale transactions and property tax records) using a proprietary record-matching algorithm. Each address is associated with the name of the county it belongs to.

Our relocation dataset is only a subset of all US relocations. However, it does not show a selection bias: on the 630 counties, the Pearson correlation between the number of relocations into a county and the size of the county (measured in the number of residential properties) is high ($r = 0.855$). The correlation between the number of relocations out of a county and the size of the county is even higher ($r = 0.94$).

EXHIBIT 2: OUR DATASET OF 630 COUNTIES. THE COLOR REPRESENTS THE COUNTY SIZE AS MEASURED BY THE NUMBER OF RESIDENTIAL PROPERTIES



RESULTS

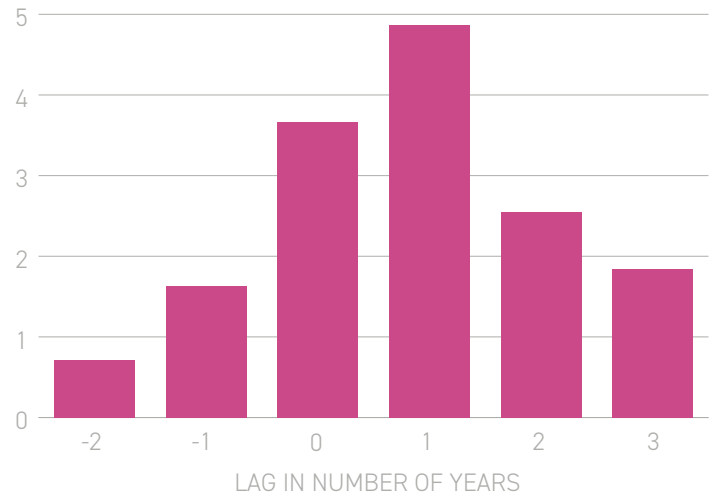
In this paper, we do not attempt to identify the causal effect that the NGI of counties might have on county development. While we leave that for future work, we focus here on measuring correlations, which would indicate statistical connections between NGI and county characteristics.

More specifically, our study aims at establishing a connection between NGI and single-family residential property price variation. For each county in our dataset, we built two time series: the yearly values of a county's NGI, and the yearly relative change of median property prices in the county. While the former time series is quite self-explanatory, we next elaborate on the latter. For county c and each pair of consecutive years $t-1$ and t , we calculate the median price of transactions P_{t-1}^c and P_t^c , and establish a yearly price change as $\Delta P_t^c = (P_t^c - P_{t-1}^c) / P_{t-1}^c$.

Analogously, we establish a nationwide median price change as $\Delta P_t^N = (P_t^N - P_{t-1}^N) / P_{t-1}^N$. The yearly relative change of median property prices for county c is then defined as $\Delta P_t^c - \Delta P_t^N$. For example, if prices in county c rose by 10% from one year to the next, while the nationwide prices rose by 18%, then the relative price growth in county c was 10% - 18% = -8%, that is, the growth was negative relative to the nationwide growth. If, however, the prices in county c rose by 2% from one year to the next, while the nationwide prices fell by 3%, then the relative price growth in county c was 2% - (-3%) = 5%, so the relative growth may be larger than the absolute growth.

For each county in our dataset, we computed the Pearson correlation between the two time series, and counted the number of counties for which the correlation was strictly positive ($r > 0.5$) as well as the number of counties for which it was strictly negative ($r < -0.5$). We then report on the ratio between those two numbers. We repeat this process for cases when the two time series lag relative to each other, both positively and negatively. A positive lag of 1 means that the NGI time series starts one year earlier than the price change time series, and also ends one year earlier. A negative lag of 1 means that the NGI time series starts one year later than the price change time series and also ends one year later. Lags of -2, 2, and 3 are defined analogously (*Exhibit 3*).

EXHIBIT 3: RATIO OF NUMBER OF COUNTIES WITH POSITIVELY TO NEGATIVELY CORRELATED NGI SCORES AND RELATIVE MEDIAN PROPERTY PRICE CHANGES



As we can see, for a lag of 0 (i.e., both time series start and end in the same year), there are 3.7 times more strictly positively correlated time series than strictly negatively correlated ones. This implies that property prices fluctuate similarly to the NGI values for 3.7 times more counties than those for which the prices fluctuate differently from the NGI values.

Remarkably, for a lag of 1 (i.e., NGI values precede price changes by one year), this ratio grows to 4.9: property prices follow NGI values almost five times more often than they don't follow NGI values. The growth of this ratio from 3.7 to 4.9 indicates some predictive power of the NGI.


For a lag of two years, the ratio drops significantly to 2.5, that is, although we can still see more counties in which prices follow NGIs two years later, the number of such counties is not much larger than the number of counties in which prices do not follow NGIs. For a lag of three years, the ratio falls even more, down to 1.8. Remarkably, it is still above 1 (i.e., there are more counties in which prices follow NGIs even three years after).

Negative lags (i.e., price changes precede NGIs by one or two years) show a different picture: while for a lag of -1 there are insignificantly more counties where NGIs follow prices than those where NGIs do not (the ratio is 1.6), for a lag of -2 this ratio drops to 0.7, which means that there are more counties for which NGIs do not follow prices than those for which NGIs do. This implies that property price changes have little to no predictive power over NGI values.

CASE STUDIES

Given the selected 630 counties, the histogram of their NGI values (averaged over the years) is reported below:

EXHIBIT 4: HISTOGRAM OF NGI AVERAGES FOR 630 US COUNTIES



As we can see, the majority of average NGI values fall around zero (with a mean of 0.09 and a standard deviation of 0.34), while some of the counties are truly outstanding. Specifically, Pinal County, Arizona has exceptionally high NGI values (with an average of 2.4), while Cape May County in New Jersey has very low NGI values (with an average of -1.3).

Yearly NGI values of Pinal County, AZ are shown on the top panel in *Exhibit 5*. We can see that they are consistently positive, with two peaks in 2009 and 2017. The bottom panel illustrates the relocation numbers, in and out of Pinal County (positive numbers represent inbound traffic, whereas negative numbers correspond to outbound traffic). We show the relocation numbers relative to the inbound relocation in 2008.

Arizona has exceptionally high NGI values (with an average of 2.4), while Cape May County in New Jersey has very low NGI values (with an average of -1.3).


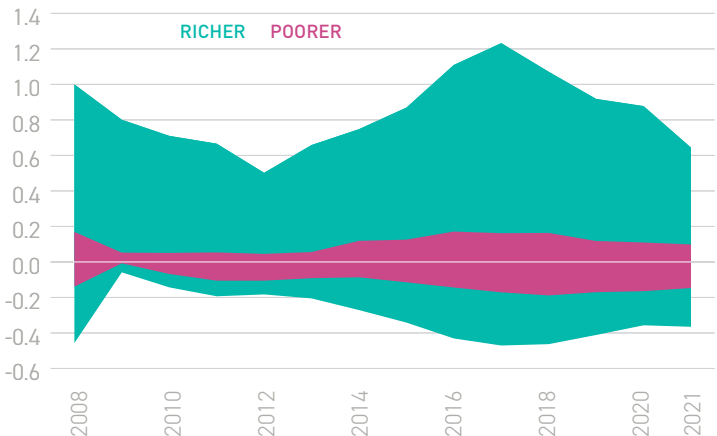



EXHIBIT 5: PINAL COUNTY, AZ: NGI AND THE RELATIVE NUMBER OF RICHER AND POORER PEOPLE MOVING IN AND OUT



From the bottom panel in *Exhibit 5*, it is clear that the inbound traffic is substantially stronger relative to the outbound traffic, which is primarily defined by the richer cohort: the number of richer people relocating into Pinal County is significantly greater than the number of richer people relocating out of it. As for the poorer population, the trend is reversed: the number of poorer people relocating into Pinal County is slightly smaller than the number of poorer people relocating out of it. This is consistent with our intuition on the factors that would make NGI values high. The majority of high-end inbound traffic comes from the nearby city of Phoenix, followed by California and (with a substantial gap) by the states of Washington and Colorado. The high-end outbound traffic is mostly within Arizona.

The peak in 2009 is explained by an almost non-existent outbound traffic. The peak in 2017 is associated with a very strong inbound traffic, specifically among the richer population. A certain decline of NGI from 2017 to 2021 is likely caused by an overall decline in relocations. We note, however, that there is no direct connection between the number of relocations and the value of NGI, because NGI takes into account the wealth excess those relocations bring in and out of the county. Specifically, if fewer people move in but they bring more capital, the NGI values may be higher than when more people move in, but they bring less capital.

In terms of residential property prices, we can see a very healthy growth starting from 2011, as shown in *Exhibit 6*.

While on the nationwide level property prices doubled since 2011, in Pinal County they trebled, and reached the national average pricing in 2021. Note that the NGI time series of Pinal County does not correlate with its property price time series ($r = -0.4$), however, exceptionally high NGI values suggest a strong gentrification, which is likely to lead to a high appreciation.

As we could expect, Cape May County, NJ behaves very differently from Pinal County, AZ. Cape May's NGI values are overly negative (*Exhibit 7*, upper panel, and this is due to a very strong inbound traffic of people who are less wealthy than Cape May residents (*Exhibit 7*, bottom panel). Cape May is a luxurious resort community on the Southern coastline of NJ, with home prices being significantly above the national medians (*Exhibit 8*). With the majority of inbound traffic coming from the nearby counties of Pennsylvania and New Jersey, Cape May's wealth is just greater than that of most of its neighbors and, thus, the county cannot be further gentrified.

EXHIBIT 6: PINAL COUNTY, AZ: MEDIAN RESIDENTIAL PROPERTY PRICES PER SQUARE FOOT




EXHIBIT 7: CAPE MAY COUNTY, NJ: NGI AND THE RELATIVE NUMBER OF RICHER AND POORER PEOPLE MOVING IN AND OUT





EXHIBIT 8: CAPE MAY COUNTY, NJ: MEDIAN RESIDENTIAL PROPERTY PRICES PER SQUARE FOOT



The lack of gentrification ultimately affects the home price dynamics. Overall, the prices rise in Cape May slower than the nationwide average. Nevertheless, they kept relatively unchanged during the 2009-2011 recession (when nationwide prices slipped), and then rose faster in 2020 during the pandemic. Their counter-cyclical behavior is in compliance with the resort, get-away nature of the county.

THE FUTURE OF GENTRIFICATION RESEARCH

Our proposed NGI provides a useful tool for real estate investment managers to optimize their investment decisions. National population and income levels can obfuscate significant variation between counties within the US, which offer abundant opportunities to identify and take advantage of county-level trends.

Just as prudent multifamily owners underwrite the creditworthiness of a potential renter (as opposed to immediately accepting), migration patterns can be subject to comparable financial underwriting of households who change counties. Such an analysis can impact our understanding of the materiality of each move and its potential influence on the real estate market of a given county.

Future research on this topic may include a deeper analysis of explanations for why the NGI seems to have predictive power, and the variation of such predictions across different market characteristics, such as differences between in-state and out-of-state moves, and comparing NGI values with real estate pricing forecasts.

ABOUT THE AUTHORS

Ron Bekkerman is a Strategic Advisor and a former Chief Technology Officer of Cherre, a leading real estate data platform. Donal Warde is an Entrepreneur in Residence at Tenney 110, a PropTech venture studio, and worked in private equity real estate for several years. Maxime C. Cohen is the Scale AI Chair Professor of Retail and Operations Management and Director of Research at McGill University.

NOTES

- ¹ CBRE. “Steady Investment Activity Shows Commercial Real Estate Resilience,” CBRE Insights, accessed April 24, 2023, <https://www.cbre.com/insights/viewpoints/steady-investment-activity-shows-commercial-real-estate-resilience>.
- ² Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. “Determinants of Portfolio Performance.” *Financial Analysts Journal* 51, no. 1 (January/February 1995): 133-138. <https://www.cfainstitute.org/en/research/financial-analysts-journal/1995/determinants-of-portfolio-performance>.
- ³ World Population Review. “List of Counties in the United States.” Accessed April 24, 2023. <https://worldpopulationreview.com/us-counties>.

REVIEWER RESPONSE


Based on evidence that allocation decisions represent the primary driver of return in mixed-asset portfolios, the authors put forth a top-down approach to achieving outperformance in multifamily investment, given the strong (0.90) correlation between home prices and those of multifamily assets. Getting to this involved the creation of a database with two primary elements: annual changes in the price of single-family homes in US counties from 2008 to 2021 compared with an index that measures changes in the wealth of a county. Their wealth index, which they call a Normalized Gentrification Index or NGI, is calculated by comparing the wealth of people moving into and out of a county with that of current residents. The authors back into a county’s wealth based on the assumption that home values are the primary source of wealth for the typical household. While this analysis uses county data, they posit that their methodology can be applied to any geography.

The authors find only a weak correlation between wealth inflows and residential price changes, with an r of 0.293,

but this well exceeds the correlation between population growth and residential price change of 0.177. That said, they did see utility in the use of the NGI in predicting home price movements, especially when lagged a year. Then “property prices follow NGI values almost five times more often than they don’t follow NGI values.”

This is a provocative paper on many levels, including the use of the term gentrification in their title and their index. That highly charged term aside, investors are indeed seeking ways to use data science to achieve outperformance and the potential for applying readily available data on home prices and some proxy for growth in a location (be it population increases, median household income gains, or the author’s net wealth index) is high.

– Mary Ludgin, PhD
Senior Managing Director,
Global Head Investment
Research, Heitman
Member, Summit Journal
Editorial Board



National population and income levels can obfuscate significant variation between counties within the US, which offer abundant opportunities to identify and take advantage of county-level trends.