



FEDERAL RESERVE BANK OF MINNEAPOLIS

Pursuing an economy that works for all of us

Minneapolis 2040 Housing Indicators: Technical Appendix

This technical appendix describes the synthetic control approach that we use to generate counterfactual Minneapolis outcomes and estimated statistical significance (p-value) in the [accompanying dashboard](#). The appendix also describes how we construct each measure featured in the dashboard, including both core indicators (for which we typically produce counterfactuals) and contextual figures.

The Synthetic Control Framework

The Minneapolis 2040 Comprehensive Plan (the Plan) comprises an array of policies that affect different economic and social outcomes in the City of Minneapolis (the City). Consequently, we will track a variety of City outcomes for many years after 2020, when the Plan took effect. Rather than attempt to disentangle the effects of disparate elements of the Plan—from each other and from simultaneous unrelated developments—we assess the Plan in its entirety, focusing on the housing-related impacts.

We do this using the synthetic control method (SCM), a methodology that allows the construction of counterfactual Minneapolis outcomes. The methodology was devised by Alberto Abadie and coauthors, deployed most famously in Abadie, Diamond, and Hainmueller (2010). It has quickly become a staple of program evaluation and an invaluable tool for assessing policy effects when experimental designs are not feasible.

We begin the description of our synthetic control implementation with the standard “potential outcomes” notation used in program evaluation to characterize treatment effects. The effect of the Minneapolis 2040 plan in period t can be described as:

$$\tau_t = Y_t^I - Y_t^N$$

where Y_t^I is the observed outcome for Minneapolis affected by the policy implementation, Y_t^N is the potential outcome in the absence of the policy for Minneapolis, and $t > T_0$ where T_0 is the period before the policy implementation.¹

¹ t may be of annual or monthly frequency depending on the data availability. For some indicators (e.g. Income Segregation Index), although the data are presented and inputted as annual data, they are calculated using three years or five years of data.

Y_t^N is estimated using a synthetic control, defined as a weighted average of donor unit outcomes. The pool of donor units consists of other cities in the U.S., subject to the following criteria²:

1. Total population between 150,000 and 2,000,000
2. No similar policy in effect around the same time (e.g., Portland, OR)
3. No cities that may experience spillover effects of the Minneapolis 2040 plan (e.g., St. Paul, MN)
4. Principal cities only (e.g., excluding Irving, VA)
5. Cities that are a census place during the whole study period (2010-current) (e.g., excluding Urban Honolulu census designated place, HI)

After removing cities that do not meet these criteria, we retain $J = 126$ donor cities across the U.S.

The synthetic control estimator of Y_t^N is then:

$$\hat{Y}_t^N = \sum_{j=1}^J w_j Y_{tj} \quad (1)$$

where j denotes each of the donor cities, $w_j \in [0,1]$ is the weight for city j and $\sum_{j=1}^J w_j = 1$, and Y_{tj} is the observed indicator for city j at time $t \in (T_0 + 1, \dots, T)$, i.e., during post-treatment periods.

Let $\mathbf{W}^* = (w_1^*, w_2^*, \dots, w_J^*)'$ be a vector of the synthetic control weights that minimizes:

$$\|\mathbf{X} - \mathbf{X}_0 \mathbf{W}\| = \left(\sum_{h=1}^k v_h (X_h - w_1 X_{h1} - w_2 X_{h2} - \dots - w_J X_{hJ})^2 \right)^{\frac{1}{2}} \quad (2)$$

subject to $\sum w_j = 1$ and $0 \leq w_j \leq 1$

where v_h is a given set of non-negative predictor weights, X_h is predictor h for Minneapolis, X_{hj} is predictor h for donor city j , and k is the number of predictors.

Before $\mathbf{W}^*$ can be estimated, we need to find $\mathbf{V} = (v_1, \dots, v_k)$ such that $\mathbf{W}(\mathbf{V}) = (w_1(\mathbf{V}), \dots, w_J(\mathbf{V}))'$ minimizes equation (2) and the weights in $\mathbf{W}(\mathbf{V})$ are positive and sum to one.

$\mathbf{V} \in (0,1)^m$, where m is the total number of predictors, is estimated based on a given set of predictors. In this particular project, the predictors include:

² All estimates are based on ACS 2014-2018 data and cities are Census places.

- Average of natural logarithm of total population for each city (Minneapolis and all donor cities) over $\{1, 2, \dots, T_0\}$
- Average of percent of owner-occupied housing units for each city over $\{1, 2, \dots, T_0\}$
- Average of natural logarithm of median housing cost for each city over $\{1, 2, \dots, T_0\}$
- Average of natural logarithm of median household income for each city over $\{1, 2, \dots, T_0\}$
- Odd-year lagged values of the observed indicator (e.g. if $\{1, 2, \dots, T_0\} = \{2010, 2011, \dots, 2019\}$, the values of the indicator at years 2011, 2013, 2015, 2017 and 2019 are included in the algorithm as predictors; if $\{1, 2, \dots, T_0\}$ is of monthly frequency, then the values of the indicator in December of those odd years would be included as predictors)

These variables are selected such that they are general enough to help predict many of the indicators we consider. The lagged indicator values are particularly important, given that they capture unobserved determinants of a given city's indicator.

We estimate the initial values of $\mathbf{V}$ using two methods:

1. Equal weights: $\mathbf{V} = \left(\frac{1}{k}, \dots, \frac{1}{k}\right)$
2. Regression-based:
 - a. Solve for $\boldsymbol{\beta}$: $(X'X)^{-1}X'Z'\boldsymbol{\beta} = \mathbf{0}$ where X is a $(J + 1) \times (k + 1)$ matrix with an intercept column and k columns of predictors as listed above, and Z is a $T_0 \times (J + 1)$ matrix where each column is the observed indicator value for Minneapolis and each donor city
 - b. Remove the intercept term from $\boldsymbol{\beta}$
 - c. Then, $\mathbf{V} = \boldsymbol{\beta}\boldsymbol{\beta}'$

For each of the two potential initial values for $\mathbf{V}$, we find the corresponding $\mathbf{W}$ by solving the quadratic programming problem:


$$\mathbf{W} = \operatorname{argmin} \left\{ c'w + \frac{1}{2} w'Hw \right\} \quad \text{subject to} \quad \sum w = 1 \text{ and } 0 \leq w \leq 1 \quad (3)$$

where $c = -X_1'V_{\text{diag}}X_0$ where X_0 is a $k \times J$ matrix containing predictor values for all J donor cities, X_1 is a $1 \times k$ matrix containing predictor values for Minneapolis, and V_{diag} is a diagonal matrix with diagonal values $\mathbf{V}$, and $H = X_0'^{V_{\text{diag}}}X_0$.

The loss is then calculated as $\frac{1}{|\{1, 2, \dots, T_0\}|} \|\mathbf{X} - \mathbf{X}_0\mathbf{W}\|$. The $\mathbf{V}$ that results in the lowest loss—i.e., whichever of methods (1) or (2) performs best—is $\mathbf{V}^*$, which then implies a corresponding $\mathbf{W}^*$ that solves equation (3).

Finally, the pre- and post-treatment counterfactual outcome of interest can be estimated using equation (1). The figure below shows an example—housing cost burden among extremely low-income households—with observed Minneapolis values in blue and the counterfactual in yellow.

Minneapolis housing cost burden among extremely low-income households



Assessing the Significance of Differences between Actual and Counterfactual Outcomes in the Post-treatment Period

It is important to distinguish genuine impacts of a policy from those differences in actual and counterfactual values that are likely due to chance. However, distinguishing between significant and insignificant effects in an SCM context is somewhat more complicated than it is for simpler econometric methods. The approach we take in this dashboard is to compare the model's fit in the post-treatment period to the fit in the pre-treatment period. Specifically, the root mean squared error (RMSE) of the model can be calculated in the two periods and then the ratio of the post-RMSE to pre-RMSE values can be constructed. Intuitively, the counterfactual should tightly fit the actual series in the pre-treatment period (producing a low value for the RMSE), while the fit for the treated unit in the post-treatment period should be worse if the policy has an effect (yielding a higher value for the RMSE). The resulting ratio of the post-to-pre RMSE values should also increase as the effect of the treatment becomes more pronounced. To construct a p-value for Minneapolis, we implement this process for every donor city in our sample.

For $0 \leq t_1 \leq t_2 \leq T$ and $j = \{1, \dots, J + 1\}$, the RMSE at periods $t_1, \dots, t_2$ is calculated as:

$$R_j(t_1, t_2) = \left(\frac{1}{|\{t_1, \dots, t_2\}|} \sum_{t=t_1}^{t_2} (Y_{jt} - \hat{Y}_{jt}^N)^2 \right)^{\frac{1}{2}}, \quad (4)$$

where $\hat{Y}_{jt}^N$ is the outcome on period t produced by a synthetic control when donor city j is coded as treated and using all other J donor units (excluding Minneapolis) to construct the donor pool. Then, the post-to-pre RMSE ratio is calculated as:

$$r_j = \frac{R_j(T_0 + 1, T)}{R_j(1, T_0)} \quad (5)$$

Because the donor cities have not undergone treatment, we expect the variation between the actual data and the synthetic controls to be random and the resulting RMSE ratios to be smaller than the value obtained for Minneapolis (if the policy has had an effect). The calculated r_j values can then be sorted and the position of Minneapolis, relative to the other donor cities, can be determined and reported as an empirical p-value. Specifically, the p-value based on the permutation distribution of r_j is calculated as:

$$p = \frac{1}{J+1} \sum_{j=1}^{J+1} I_+(r_1 - r_j), \quad (6)$$

where I_+ returns one if $r_1 \geq r_j$ and zero otherwise, and r_1 is the post-to-pre RMSE ratio for Minneapolis.

For example, if the ratio for Minneapolis was the 10th highest out of the 126 donor cities and itself then the resulting p-value would be $10 / 127 = 0.079$. See Abadie, Diamond, and Hainmueller (2010) or Abadie (2020) for more details.

When analyzing shares or rate measures, such as share of multifamily units and the rental vacancy rate, we calculate an alternative p-value to focus on the post-treatment effect in the most current year of data. Similar to the above methodology, we compute the squared prediction error (SPE) for each unit in the sample and the pre-treatment RMSE for each unit across the pre-treatment period.

For the most recent year, T , we take the square root of the SPE for each unit as:

$$Error_{j,T} = (SPE_{jT})^{\frac{1}{2}}$$

We then calculate the ratio of the most recent year's error to the pre-treatment RMSE for each unit:

$$ratio_j = \frac{Error_{j,T}}{RMSE_{j,pre}}$$

The point-in-time p-value is constructed as the proportion of donor units with a ratio greater than or equal to Minneapolis' ratio:

$$P_T = \frac{1}{J+1} \sum_{j=1}^{J+1} 1(ratio_j \geq ratio_{Minneapolis})$$

where J is the number of donor units and $1(\cdot)$ is the indicator function which equals 1 if the condition is met and 0 otherwise.

Non-convergence and Overfitting

Synthetic control, even when limiting consideration to its classic Abadie, Diamond, and Hainmueller (2010) version, allows for different implementation choices. We now detail some of those choices and the considerations that informed them.

Lagged dependent variables

The SCM literature is not settled on the question of how much emphasis to place on lagged dependent outcomes. On the one hand, lagged values of the dependent variable capture a range of otherwise unobserved factors that matter for constructing any reasonable counterfactual. This is immediately evident in the within-sample fit improvement that is achieved when including lagged values. However, “excessive” use of lagged values relegates the time-invariant observable factors to a position of little or no influence over the counterfactual. This may lead to undesirable overfitting that reduces the model’s out-of-sample accuracy, particularly if the researcher believes that control units with certain key observable characteristics deserve emphasis in the construction of the counterfactual. We therefore aimed to balance these considerations, in line with advice from Abadie and others in their recent publications, by using only odd-year lagged values of the dependent variable.

Numerical optimization

There is nothing about the Abadie, Diamond, and Hainmueller (2010) implementation of SCM that necessitates use of a specific numerical optimization routine for solving the optimization problems (see above for details of those minimization problems) that yield synthetic control weights. Unfortunately, different routines yield different solutions; these differences can be non-trivial in magnitude. (To be clear, this is not a situation that is unique to SCM applications.)

Broadly speaking, numerical routines use some combination of “guess-and-check” simplex methods and gradient-based approaches that necessitate intermediate calculations of numerical derivatives. Again speaking broadly, the former is generally considered more robust to non-smooth problems with multiple local minima, while the latter is more efficient with smoother problems.

We settled on the following approach: solve each problem using both the commonly used gradient-based BFGS method and the venerable Nelder-Mead simplex method, then select the minimum-distance solution. This approach is what the R implementation of Abadie, Diamond, and Hainmueller (2010), written by the authors of that paper, is set up to do by default.

However, in 2021 we explored many potential options for numerical optimization besides the approach described above. Those options included running the numerical methods in sequence, e.g., Nelder-Mead first and then BFGS if it does not converge, or vice versa. Our goals were a) to avoid non-convergence to the extent possible, b) to reliably find global minima, c) to generate “stable” results that did not fluctuate based on economically irrelevant matters like the ordering of the data or negligible changes in variables’ values, and d) to minimize computation time.

One note that applies throughout is that while the weights generated by a particular approach may differ visibly from the weights generated with a different approach, this does not necessarily mean that the generated counterfactual will differ meaningfully. There are in general multiple combinations of donor unit weights (particularly in applications like ours with many potential donors) that achieve similar model fit and could be “reached” with different numerical methods. Ex post, different synthetic controls will sometimes have different post-treatment trajectories, but the same identification assumptions must be applied to any estimated counterfactuals in order to interpret post-treatment differences as causal impacts of the treatment. Perhaps most important is that the numerical optimization approach be settled definitively before post-treatment data is observed, eliminating the possibility of “gaming” the approach to reverse-engineer a particular desired counterfactual.

Non-convergence due to extreme values relative to donor cities

As described in more detail above, to calculate p-values associated with post-treatment differences between Minneapolis and its counterfactual, we estimate the counterfactuals for each of the donor cities. While counterfactuals are generated without incident for Minneapolis in each of our indicators, this is not always the case for the donor cities.

In addition to potential non-convergence as described above, non-convergence due to extreme values can be a concern. Ultimately, the output of the SCM procedure is the weighted average of donor cities. Because of how we selected donor cities for Minneapolis—with pre-treatment characteristics that are very roughly similar to that city—Minneapolis outcomes are not at the bottom or top of the distributions of donor cities. The same cannot be said for some of the donor cities, of course. And when constructing counterfactuals for donor city outliers, the SCM procedure can fail to generate corresponding donor city weights. For example, if the housing cost burden for Cincinnati is 10 percent but it is between 15 and 50 percent for other donor cities, then there are no donor city weights that could yield a weighted average of 10 percent. This is because one of the constraints in our optimization is that weights have to add up to 1 and must be non-negative; in other words, the procedure does not permit extrapolation.

While this issue does not arise frequently, it does sometimes occur when we calculate p-values, but not so frequently as to undermine our confidence in the usefulness of the p-values and associated assessment of statistical significance.

Construction of Specific Indicators

1. New Dup/Tri/Quad

1.1 Core

The New Dup/Tri/Quad indicator is calculated as the trailing three-year average of units in new duplex, triplex, or fourplex construction for which the City issues building permits over the preceding 12 months.

Data source: [Metropolitan Council, Residential Permits](#); [U.S. Census Bureau, Building Permits Survey](#).

1.2 Contextual

There are three contextual charts for this indicator. The first displays the raw values of permitted duplex, triplex, or fourplex units between Minneapolis and the projected counterfactual outcome from the core indicator. The second shows the number of units in permitted buildings for duplex, triplex, or fourplex, broken down by whether the building is in the pre-2020 single-family zoning as provided by the City. The third contextual chart is a map showing the pre-2020 single-family zoning and the locations of permitted duplex/triplex/fourplex buildings from 2020 to the most recent year.

Data source: [U.S. Census Bureau, Building Permits Survey](#); [Metropolitan Council, Residential Permits](#).

2. New Multifamily Housing

2.1 Core

The New Multifamily Housing indicator is calculated as the number of units in new residential, multifamily (5 or more units) construction for which the City issues building permits over the preceding 12 months.

Data source: [U.S. Census Bureau, Building Permits Survey](#).

2.2 Contextual

There are two contextual charts for this indicator. One maps new residential, multifamily building permits at the address level for the most recent three years based on data availability where the base layer shows the most recent census-tract-level median household income. The other shows the ranking of Minneapolis compared to the donor cities used in the SCM by the most recent available trailing 12-month sum of new multifamily units.

Data sources: [Metropolitan Council, Building Permits Survey](#); [U.S. Census Bureau, Building Permits Survey](#); [American Community Survey 5-year files](#).

3. Housing Composition

3.1 Core

The housing composition indicator is calculated as the share of single-family detached units out of all housing units, excluding mobile and housing units labeled as “Other.”

Data source: [U.S. Census Bureau, American Community Survey 1-year files](#).

3.2 Contextual

There are two contextual charts for this indicator. The first contextual chart shows the share of existing housing stock by structure type over time. The second contextual chart shows the ranking of Minneapolis compared to the donor cities used in the SCM by the share of single-family detached units.

Data sources: [U.S. Census Bureau, American Community Survey 5-year and 1-year files; Metropolitan Council, Housing Stock; Open Data Minneapolis, Communities.](#)

4. Rental Vacancy Rate

4.1 Core

The Rental Vacancy Rate indicator uses the Vacancy Index from Apartment List. Properties are included in the sample if they meet two criteria: the property has a monthly vacancy rate of 15 percent or lower at least once and the property has been on the Apartment List platform for a minimum of six months. The sample of properties varies from month to month and does not include single-family rentals of multifamily housing with less than ten units and properties that are not on Apartment List's platform. Read more about Apartment List's methodology [here](#).

Data source: [Apartment List, Vacancy Index.](#)

4.2 Contextual

There is one contextual chart for this indicator. The chart shows the ranking of Minneapolis compared to the donor cities by the lowest to highest rental vacancy rate.

Data source: [U.S. Department of Housing and Urban Development, Comprehensive Housing Affordability Strategy data.](#)

5. New Affordable Housing

The New Affordable Housing indicator is calculated as the number of new income-restricted housing units at or below 30 percent of area median income (AMI), above 30 percent and at or below 50 percent of AMI, and above 50 percent and at or below 60 percent of AMI. Due to the lack of comparable data from other cities, we construct two comparisons to Minneapolis based on the regional data: St. Paul and the combination of other cities in the Twin Cities region that the Metropolitan Council designates as Urban Center and Urban areas. (Learn more about Community Designations from the [Metropolitan Council](#).)

Data source: [Metropolitan Council, Affordable Housing Construction.](#)

6. Housing Cost Burden

6.1 Core

The Housing Cost Burden indicator is calculated as percent of extremely low-income renting households who experience housing cost burden. Extremely low-income household is defined as a household whose household income is less than 30 percent of AMI. A household is cost-burdened if their housing cost-to-household income ratio is at or above 30 percent.

To smooth out the fluctuation due to measurement error from using ACS 1-year files, the three-year lagging average is reported as the final indicator.

Data sources: [U.S. Census Bureau, American Community Survey, Public Use Microdata Sample, 1-year files; U.S. Department of Housing and Urban Development.](#)

6.2 Contextual

There are five contextual charts for this indicator. The first one shows racial/ethnic composition of residents among Minneapolis extremely low-income renters, Minneapolis overall, Minnesota overall and U.S. overall. Race/ethnicity of residents is grouped into five categories:

1. Asian, non-Latino
2. Black or African American, non-Latino
3. Latinx, inclusive of Hispanic or Latino/a origin regardless of their race
4. White, non-Latino
5. All other races, inclusive of all residents who do not fall into one of the four categories above. This is done due to the lack of availability of the data for other groups such as Native Americans and individuals reporting two or more races.

The second contextual chart shows the time series of housing cost burden by racial/ethnic group and tenure. The third contextual chart shows the ranking of Minneapolis compared to the donor cities by the most recent available city-level housing cost burden among extremely low-income renters. The fourth shows the share of total rental units where the housing cost is affordable (less than or equal to 30 percent of income) to households earning 30 percent of HUD Area Median Family Income (HAMFI) and below and the share affordable to households earning above 30 percent of HAMFI to 50 percent of HAMFI. The last chart shows the ranking of Minneapolis compared to the donor cities by the most recent available city-level share of rental units that are affordable to extremely low-income households.

Data sources: [U.S. Census Bureau, American Community Survey, Public Use Microdata Sample, 5-year and 1-year files](#); [U.S. Department of Housing and Urban Development](#).

7. Price of Housing

7.1 Core

The Price of Housing indicator is the monthly median rent paid for new leases. To ensure the Price of Housing and Rental Vacancy Rate core indicators reflect the same underlying market methodology, we use Apartment List's rent estimates to compare Minneapolis and the projected counterfactual outcome.

Data source: [Apartment List, Rent Estimates](#).

7.2 Contextual

There are six contextual charts for this indicator, as detailed below:

1. A time series of two rent measures, the quarterly average rent from CoStar and the monthly median rent paid for new leases from Apartment List, to capture the typical rent charged in Minneapolis.
2. The time series of median asking rent for the apartment market and the shadow market (non-apartment rentals) by number of bedrooms to capture rent across different markets.
3. The map of census-tract-level share of rental units with monthly gross rent below \$1,000 for the most recent year as well as five years prior to allow for non-overlapping data.

4. The time series of [Zillow Home Value Index](#) (ZHVI) to capture the level and appreciation of home values across Minneapolis.
5. The time series of median monthly owner cost broken down by whether the household has a mortgage to capture the housing cost of owner-occupied households.
6. The ranking of Minneapolis compared to the donor cities by the median rent paid for new leases to capture Minneapolis' standing relative to the donor cities.

To adjust for inflation, we use the Consumer Price Index for All Urban Consumers (CPI-U) obtained from the FRED (Federal Reserve Economic Data) API (series ID: CPIAUCSL). The CPI-U is published monthly by the U.S. Bureau of Labor Statistics and represents changes in prices paid by urban consumers for a representative basket of goods and services. The base date is the most recent observation date of the data.

Data sources: [CoStar](#); [Apartment List](#), [Rent Estimates](#); [HousingLink](#), [Twin Cities Rental Revue](#); [Zillow, Inc.](#); [U.S. Census Bureau, American Community Survey, 1-year and 5-year files](#).

8. Housing Choice

8.1 Core

The Housing Choice indicator is calculated as percent of homes in Minneapolis that are affordable to purchase for households earning 50 percent of AMI or below. Given household income, we calculate the house-buying power as the price of the house the income could afford by assuming:

- 30 percent debt-to-income ratio
- 10 percent down payment
- 30-year term
- 1.36 percent property tax rate, annually
- 1 percent mortgage insurance, annually
- 150 dollars hazard insurance, monthly
- Annual average mortgage interest rate

To smooth out the fluctuation due to measurement error from using ACS 1-year files, the three-year lagging average is reported as the final indicator.

Data sources: [U.S. Census Bureau, American Community Survey, Public Use Microdata Sample, 1-year files](#); [U.S. Department of Housing and Urban Development](#); [Freddie Mac, 30-Year Fixed Rate Mortgage Average in the United States](#).

8.2 Contextual

There are four contextual charts for this indicator. The first one shows the most recent available share of homes that are affordable to buy by Minneapolis neighborhood at median household income of homeowners and renters by race and ethnicity. Only single-family homes, townhomes and condominiums are included in this sample to capture the one-unit homes that a household could buy. The second contextual chart shows census tracts where median rent is affordable at median household income by race and ethnicity. The data are binary for this chart due to data

availability. The third chart shows the time series of homeownership rate by race/ethnicity in Minneapolis. The final contextual chart shows the ranking of Minneapolis compared to the donor cities by the most recent available city-level share of homes affordable to households earning 50 percent of AMI or below. The racial/ethnic group is based on the same definition as in 6.2, and Minneapolis neighborhood is constructed in the same way as in 3.2.

Data sources: [MetroGIS Regional Parcel Dataset](#); [U.S. Census Bureau, American Community Survey, Public Use Microdata Sample, 5-year and 1-year files](#); [U.S. Census Bureau, American Community Survey 5-year files](#); [U.S. Department of Housing and Urban Development](#).

9. Income Segregation Index

9.1 Core

Income segregation is measured using the rank-order information theory index (H^R) as described in [Reardon \(2011\)](#) and implemented via the *OasisR* package in *R*. 5-year ACS data at the Census tract level are used to estimate H^R . Importantly, the household counts are reported in the ACS are for categories of income (less than \$10,000, \$10,000 - \$15,000, \$15,000 - \$20,000, etc.) as opposed to fixed percentiles of the income distribution. Construction of the estimate is done in the following manner:

1. Obtain the counts of households by income group (there are 16 of these in the ACS surveys starting in 2010) for all of the tracts that comprise a city. There are 116 tracts in the city of Minneapolis, so the resulting set of household counts comprises a matrix whose size is 116 x 16. Each column of the matrix corresponds to the number of households in tract (j) that have income in category (k).³
2. Convert the matrix of raw counts into one of cumulative counts up to each income category. For example, the third column of the matrix would contain the total number of households with incomes up to \$20,000 (so it would be the sum of the raw counts of columns 1 – 3 in the original matrix).
3. Calculate the traditional *information theory index* of segregation for each the first 15 columns of the matrix in 2 using the following formula:

$$H(k) = 1 - \sum_j \frac{t_j * E_j(k)}{T * E(k)}$$

where T is the population of the entire city (with incomes less than or equal to category k), t_j is the population of Census tract j (with incomes less than or equal to category k),

³ Census tract boundaries and census place boundaries do not always align well. To get the tract-to-place mapping to allow us to calculate this indicator for Minneapolis and all the donor cities, we use a 40 percent area overlap between tract and place cutoff. We did a sensitivity check by looking at how different cutoffs affect the indicators and found 40 percent to be a reasonable choice compared to all other options between 25 percent and 50 percent. This mapping is also used in indicators 9 and 10.


$E(k)$ is the entropy measure for the city, and $E_j(k)$ is the entropy measure for tract j . The entropy measures are given by the following formula:

$$E(k) = k * \log_2\left(\frac{1}{k}\right) + (1 - k) * \log_2\left(\frac{1}{1 - k}\right)$$

The value of k in the above equation corresponds to the percentile of the cumulative income distribution at each income category (measured at the individual Census tract level or for the entire city). For example, if the first Census tract in Minneapolis contained 1,055 households and there were 93 households with income in the first category (less than \$10,000), 38 households with income in the second category (\$10,000 - \$15,000) and 17 households with income in the third category (\$15,000 - \$20,000) then the value of k for this tract would be 0.139 $([93 + 38 + 17] / 1,055)$.

The 16th column corresponds to income greater than \$200,000. However, this category is redundant since the segregation measures are based on separating the households into two groups at each point (those with incomes below the category cutoff and those with incomes above).

The result from this step will be a vector of 15 values, each of which represents the *information theory index* of segregation at the various income category cutoffs for the city as whole. The chart below shows the result for Minneapolis using the 5-year ACS data from 2014-2018:




4. Finally, the individual H values across the entire income distribution (i.e., all of the income percentiles) can be aggregated into a single value as the *rank-order information theory index* and denoted by H^R . If the full, continuous H function were available then value of H^R would be given by the following equation:

$$H^R = 2 * \ln(2) \int_0^1 E(p) H(p) dp$$

However, we only have 15 discrete observations of H as opposed to the full function. Reardon et al. (2006) show that the H^R can be approximated by fitting a 4th order polynomial through the discrete H observations and then using the associated polynomial coefficients in the following equation:

$$\hat{H}^R = \hat{\beta}_0 + 0.5 * \hat{\beta}_1 + 0.305 * \hat{\beta}_2 + 0.208 * \hat{\beta}_3 + 0.152 * \hat{\beta}_4$$

The coefficients shown below that are associated with the 4th order polynomial for the Minneapolis data produce a value of 0.1189 for $\hat{H}^R$.



Prior to the 2026 update, income segregation index census tract data were harmonized to 2010 boundaries by aggregating 2020 tract definitions. This harmonization process has been discontinued due to errors in the aggregation methodology and because the primary analysis unit is at the city level, where boundaries remain consistent across census years. Beginning with the 2026 update, data are maintained in their native tract boundaries without crosswalks, as city-level aggregates are unaffected by underlying changes in tract geography.

Data Source: [U.S. Census Bureau, American Community Survey 5-year files.](#)

9.2 Contextual

There is one contextual chart for this indicator showing the ranking of Minneapolis compared to the donor cities by the most recent available city-level income segregation index.

Data source: [U.S. Census Bureau, American Community Survey 5-year files.](#)

10. Rent Dissimilarity Index

10.1 Core

This indicator is calculated as:

$$D_i = \frac{1}{2} \sum_{t \in T} \left| \frac{U_{1t}}{U_1} - \frac{U_{2t}}{U_2} \right|$$

Where T is a set of all census tracts in Minneapolis, U_{1t} is the number of rental units in tract t where rent is at or below the city-level 25th percentile rent, U_1 is 25 percent of Minneapolis

rental units, U_{2t} is the number of rental units in tract tt where rent is above the city-level 25th percentile rent, and U_2 is 75 percent of Minneapolis rental units.

Data source: [U.S. Census Bureau, American Community Survey 5-year files](#).

10.2 Contextual

There is one contextual chart for this indicator, and it shows the ranking of Minneapolis compared to the donor cities by the most recent available city-level rent dissimilarity index.

Data source: [U.S. Census Bureau, American Community Survey 5-year files](#).

11. Isolation of White Residents

11.1 Core

The isolation index is calculated as:

$$I = \sum_{t \in T} \left(\frac{W_t}{W_R} \times \frac{W_t}{P_t} \right)$$

where T is a set of all census tracts in Minneapolis, W_t is the White, non-Latinx population in tract t , W_R is the total White, non-Latinx population in the city, and P_t is the total population in tract t .

Prior to the 2026 update, rent dissimilarity index census tract data were harmonized to 2010 boundaries by aggregating 2020 tract definitions. This harmonization process has been discontinued due to errors in the aggregation methodology and because the primary analysis unit is at the city level, where boundaries remain consistent across census years. Beginning with the 2026 update, data are maintained in their native tract boundaries without crosswalks, as city-level aggregates are unaffected by underlying changes in tract geography.

Data source: [U.S. Census Bureau, American Community Survey 5-year files](#).

11.2 Contextual

There are three contextual charts for this indicator. The first shows the isolation index by Minneapolis neighborhood. Minneapolis neighborhood is constructed based on block groups as described in 2.2.

The second contextual chart shows the actual to predicted number of households of color based on income only. The ratio is calculated as:

$$\text{Actual-to-predicted ratio}_n = \sum_{i \in I} \frac{HOC_{iR}}{T_{iR}} \times T_{in}$$

where Actual-to-predicted ratio _{n} is the actual-to-predicted number of households of color in neighborhood n , I is the set of all income brackets available in the ACS data, HOC_R is the number of households of color in income bracket i in Minneapolis, T_R is the total households in

income bracket i in Minneapolis, and T_{in} is the total households in income bracket i in neighborhood n .

The last contextual chart shows the ranking of Minneapolis compared to the donor cities by the most recent available city-level isolation index.

Data source: [U.S. Census Bureau, American Community Survey 5-year files](#).

References

[Abadie 2020](#), “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects”

[Abadie, Diamond and Hainmueller 2010](#), “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program”

[McClelland and Gault 2017](#), “The Synthetic Control Method as a Tool to Understand State Policy”

[Reardon et al. 2006](#), “A New Approach to Measuring Socio-Spatial Economic Segregation”

[Reardon 2011](#), “Measures of Income Segregation”