WHAT IS THE SYNTHETIC CONTROL METHOD?

Robert McClelland and Livia Mucciolo

May 2022

In 2017, the Tax Policy Center published "The Synthetic Control Method as a Tool to Understand State Policy", a guide for using the synthetic control method (SCM) along with case studies. In a new report we update the guide to account for more recent research. The method, developed in Abadie and Gardeazabal (2003) and popularized in Abadie, Diamond and Hainmueller (2010) (hereafter ADH 2010), was an important breakthrough because it can evaluate policy changes that only affect one area, such as a city, state or country.

For example, Abadie, Diamond and Hainmueller (2015) study the economic cost of the reunification of East and West Germany. Hankins (2020) examines Nebraska's unique switch from a bicameral house legislature to a unicameral house. In both cases the treatment is completely unique, and in the case of the reunification, it is difficult to think of the treatment applying to any other European nation. Other studies, such as McClelland and Iselin (2019), examine relatively rare changes in policy, including large increases in alcohol taxes.

Without the SCM, analysts may try to use a nearby area as a control (that is, an untreated area used as a comparison), but there are at least two problems with relying on that method. First, proximity cannot be ensured, or the closest areas are clearly different. Consider Seattle, with a population of about 725,000. The nearest cities include Renton, with a population of about 100,000, and Redmond, with a population of about 65,000. The nearest similar-sized cities are Portland, which is in a different state, and Vancouver, which is in a different country. Because of the differences in size, state and national laws, and culture, for many studies these cities will not form a useful control for a policy change in Seattle. Second, potential controls that are close geographically may experience spillover from the treatment. For example, a treatment in Seattle may spill over into Redmond.

Instead of an actual control area, the SCM creates a synthetic control, so that the treatment can be evaluated by comparing the outcome of the treated area with the outcome of the synthetic control area. The synthetic control is created by selecting and then weighting together a small number of control areas drawn from a pool of potential candidates. The weights are selected based on a set of variables, called predictors, that are closely related to the outcome. The idea is that if the predictors of the synthetic control are near the predictors of the treated area, the outcome of the synthetic control should be similar to the outcome of the treated area before the policy change. Then if the outcomes diverge after the policy change it looks like the policy was effective. If they don't diverge, the policy does not appear to be effective.

Unlike many other statistical methods, the SCM is remarkably transparent and accessible. The small number of donors means that analysts can determine if the makeup of the synthetic control is sensible. Weights for the predictors can be similarly reviewed by analysts and published. By selectively removing suspect donors and re-estimating the synthetic control (i.e., running placebo tests), the results can be tested to ensure that they do not rely on any one area. Code for the method is freely available on several popular statistical packages.

Since TPC's publication, many articles have proposed improvements or extensions to the SCM. In this report we update McClelland and Gault (2017) by reviewing selected articles on some of these improvements. In this brief we review the assumptions necessary for the SCM and the article that made the method so popular.

A number of conditions are necessary for the method to accurately estimate the effect of a treatment. First, no area in the donor pool can have a similar policy change, because analysts cannot determine the effectiveness of a policy by comparing one area subject to the policy with another area also subject to the policy. Second, the treatment cannot bleed over into the pool of donor areas. For example, in the ADH (2010) study of the effect of an increase of cigarette taxes on cigarette purchases per capita, cigarette purchases in border states, such as Nevada or Arizona, are assumed not to change. Third, it must be possible to combine predictors of the control areas so that they are nearly or exactly the same as the predictors of the area that saw the policy change.

There are several additional conditions as well. Outcomes must be available for periods before the treatment (ideally as many as possible without including any periods in which potential donor states received similar treatments) and at least one period after treatment for both the treated area and the pool of potential donor areas. The policy change must have no effect before it is enacted. This may not always be the case. For example, California smokers might hoard packs of cigarettes prior to a cigarette tax increase, or alcohol drinkers in Illinois might stock up on bottles of alcohol before an increase in liquor tax increase. Finally, the predictors must have values for the donor pool regions that are similar to those for the affected region. While in theory it is possible to average together predictors from two very different areas and produce predictors that are close to those of the treated area, doing so can produce a subtle bias.

A good way to describe the synthetic control method is to use the example that popularized the method: ADH (2010). As described in McClelland and Gault (2017):

In that article, the authors examine how California's tobacco control program under Proposition 99, implemented in 1988, affected smoking by creating a synthetic control version of California. They estimate that by 2000, per capita sales of cigarette packs had fallen 26 packs because of the program. Here, we describe the steps in analyzing a policy change with the SCM using the Proposition 99 example when possible. We also use the original data in ADH (2010) to demonstrate the results' sensitivity to various modeling choices.

California voters approved Proposition 99 in 1988, an initiative that raised the cigarette excise tax by 25 cents a pack and implemented a large-scale antitobacco media campaign. The tax raised $100 million annually, and the revenue was initially directed toward antismoking efforts, including antismoking education budgets. Those efforts were substantially larger than efforts in other states, but the California assembly passed Assembly Bill 99 in 1991, which diverted a large share of the tax revenue for other purposes (Glantz and Balbach 2000). Beyond the tax increase, ADH (2010) report that Proposition 99 led to numerous local ordinances prohibiting smoking in indoor spaces such as restaurants and workplaces, and by 1993 almost two-thirds of employees in California worked in smoke-free environments. They also note that tobacco lobbyists in California responded by spending 10 times more in 1991 and 1992 than in 1985 and 1986.

ADH (2010) determine the impact of Proposition 99 on per capita cigarette packs using data from 1970 through 1988 as the period before the law took effect and followed cigarette sales from 1989 through 2000. They use four predictors averaged over the years before the policy: the average retail price of cigarettes, the share of the population between the ages of 15 and 24, per capita beer consumption, and the log of per capita gross domestic product. They also use per capita annual sales of cigarette packs for the years 1975, 1980 and 1988. ADH (2010) start with 49 states other than California and drop the District of Columbia. After eliminating states that also raised cigarette taxes by at least 50 cents before or after treatment, they were left with a pool of 38 potential donors. The method chose five states, Colorado, Connecticut, Montana, Nevada and Utah, to create a synthetic California.

The results are shown in Figure 1. Cigarette sales per capita in the synthetic California carefully match cigarette sales in the actual California until 1988, the year that cigarette taxes were increased. From that point forward, the two paths diverge. Although both the synthetic and actual California show a continued decline in cigarette sales per capita, the decline is faster in the actual California, with the difference providing an estimate of the effect of Proposition 99.

FIGURE 1

Synthetic California Per Capita Cigarette Sales (ADH 2010)

Before and after Proposition 99 passage in 1988

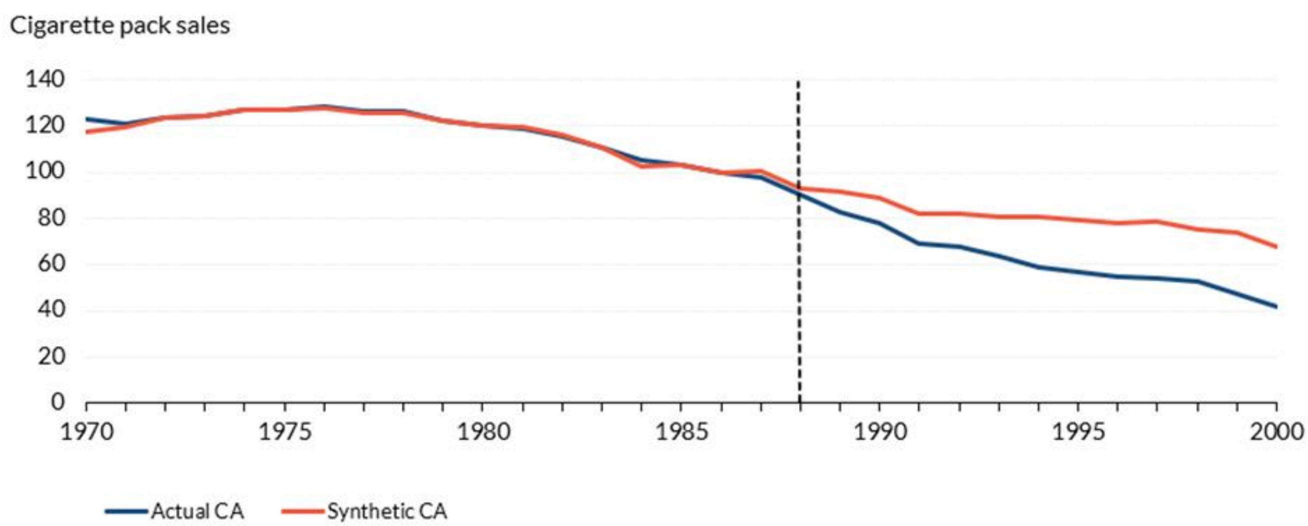

Source: Authors' calculations using the Synth package and data from ADH to replicate ADH's analysis.
Note: ADH = Abadie, Diamond, Hainmueller (2010).

To further demonstrate the fit, ADH (2010) provides two tables. In the first they compare the predictors of the actual California with its synthetic control. As shown in Table 1, the average values for the synthetic's predictors closely match those of California. If the predictors affect the outcome, this provides evidence that the synthetic control mimics the behavior of the treated area in the absence of treatment. Table 2 shows the weights of each state used in the synthetic control. Utah and Nevada contribute the most, and those two states make up more than 50 percent of the total. Montana make up an additional 20 percent, and Colorado and Connecticut make up about 22 percent combined.

TABLE 1

Actual and Synthetic California Predictor Means (ADH 2010)

Variable	Years	Actual CA	Synthetic CA
Beer consumption per capita	1984–88	24.28	24.21
Log state per capita GDP	1980–88	10.08	9.86
Retail price of cigarettes	1980–88	89.42	89.41
Share of state population ages 15-24	1980–88	0.17	0.70
Cigarette sales per capita, 1988	1988	90.1	91.64
Cigarette sales per capita, 1980	1980	120.2	120.45
Cigarette sales per capita, 1975	1975	127.1	127.06

Source: McClelland and Gault (2017) using the Synth package and data from ADH to replicate ADH's analysis.
Notes: ADH = Abadie, Diamond, and Hainmueller (2010). Units are gallons for beer consumption per capita, cents for retail price of cigarettes, and packs for cigarette sales per capita.

Figure 2 plots the estimated effect of Proposition 99 as the difference between the per capita cigarette sales in the actual California and the synthetic California, along with the analogous difference for 37 placebos.[1] These placebos are formed by running the SCM on each state as though they implemented Proposition 99 in 1988. Because they did not implement Proposition 99, the difference between cigarette sales per capita in the real and synthetic state is due to random chance. The line at the very bottom of the figure represents New Hampshire, which ADH (2010) note has the highest per capita cigarette sales in each year before 1988. After 1988, the estimated effect for California is greater than the estimate effect for most placebo states. This strongly suggests that the estimated effect for California is, in fact, the result of Proposition 99.

TABLE 2

Synthetic California Donor State Weights (ADH 2010)

State	Weight
Colorado	0.161
Connecticut	0.068
Montana	0.201
Nevada	0.235
Utah	0.335
Sum	**1.000**

Source: McClelland and Gault (2017) using the Synth package and data from ADH to replicate ADH's analysis.
Note: ADH = Abadie, Diamond, and Hainmueller (2010).

[1] Utah cannot be estimated as a synthetic control because it has the lowest per capita sales of cigarettes.

FIGURE 2

Difference in Synthetic and Actual Per Capita Cigarette Sales for California and 37 Placebo States

Before and after Proposition 99 passage in 1988

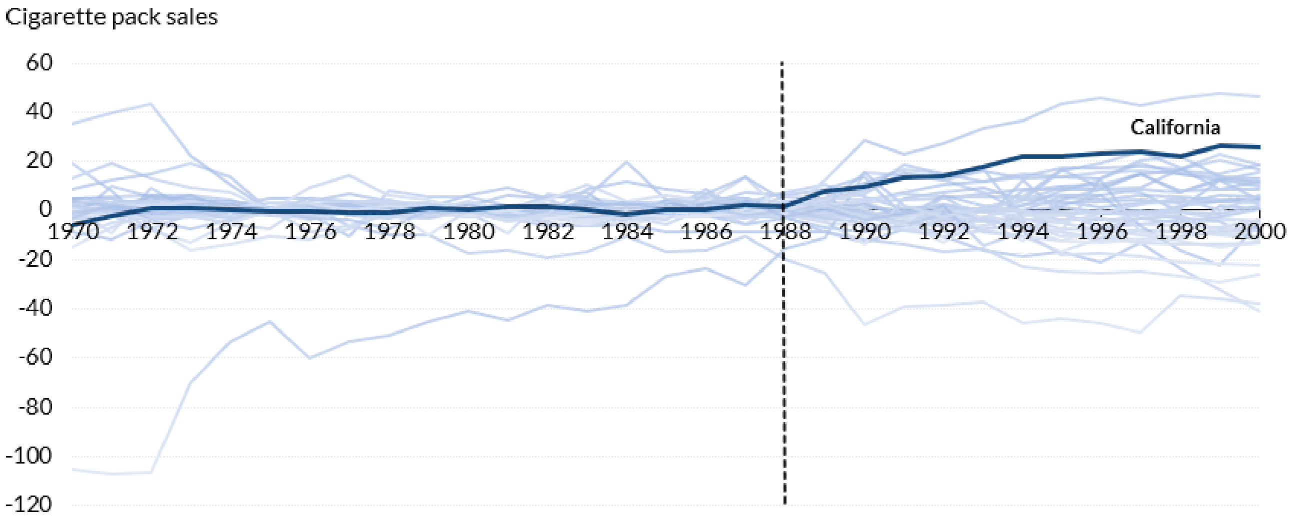

Source: Authors' estimations using the Synth package and data from Abadie, Diamond, and Hainmueller (2010).
Note: The Synth package cannot generate a synthetic placbeo for Utah, one of the 38 donor pool states.

In addition to visually inspecting plots such as Figure 2, ADH (2010) use a statistical measure of the effectiveness of Proposition 99 and find that it is unlikely that the difference between the real and synthetic California is due to chance.

In our new report, we review selected articles on some of the recent improvements to the SCM, update the step-by-step guidance for implementing the SCM, and provide links to available code. Some improvements and innovations to the SCM include: solutions to bias from overfitting (Abadie and L'Hour 2020; Ben-Michael et al 2020; Ferman and Pinto 2021), solutions to interpolation (and thus extrapolation) bias (Abadie and L'Hour 2020; Kellog et al 2020), guidance for reducing "cherry-picked" results (Ferman, Pinto, and Possebom 2020), and generalization of the method by allowing it to consider a wide range of testable hypotheses about the effect of the treatment (Firpo and Possebom 2018). Rather than overwhelm readers with mathematical notation, our goal is to make these topics more accessible by discussing them and how they apply to ADH (2010) with new visualizations and simplified intuitive frameworks.

REFERENCES

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." Journal of the American Statistical Association 105 (490): 493–505.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." American Journal of Political Science 59 (2): 495–510.

Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." American Economic Review 93(1): 113–32.

Abadie, Alberto and Jeremy L'Hour. 2021. "A Penalized Synthetic Control Estimator for Disaggregated Data" MIT manuscript

Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. 2020. "The Augmented Synthetic Control Method." Manuscript.

Ferman, Bruno, Cristine Pinto and Vitor Possebom. 2020. "Cherry Picking with Synthetic Controls" Journal of Policy Analysis and Management, 39(2): 510-532.

Firpo, Sergio and Vitor Possebaum. 2018. "Synthetic Control Method: Inference, Sensitivity Analysis and Confidence Sets". Journal of Causal Inference. 6(2): 1-26.

Glantz, Stanton A., and Edith D. Balbach. 2000. Tobacco War: Inside the California Battles. Berkeley: University of California Press.

Kellogg, Maxwell , Magne Mogstad, Guillaume Pouliot, and Alexander Torgovitsky, 2020, "Combining Matching and Synthetic Controls to Trade off Biases from Extrapolation and Interpolation" NBER Working Paper 26624.

ACKNOWLEDGMENTS

This brief was funded by Arnold Ventures. We are grateful to the foundation and to all our funders, who make it possible for the Urban-Brookings Tax Policy Center to advance its mission.

The views expressed are those of the authors and should not be attributed to the Urban-Brookings Tax Policy Center, the Urban Institute, the Brookings Institution, their trustees, or their funders.