IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 10605

The Labor Market Effects of a Refugee Wave: Synthetic Control Method Meets the Mariel Boatlift

Giovanni Peri
Vasil Yasenov

MARCH 2017

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 10605

The Labor Market Effects of a Refugee Wave: Synthetic Control Method Meets the Mariel Boatlift

Giovanni Peri
University of California, Davis, NBER and IZA

Vasil Yasenov
University of California, Davis and IZA

MARCH 2017

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Labor Market Effects of a Refugee Wave: Synthetic Control Method Meets the Mariel Boatlift*

We apply the Synthetic Control Method to re-examine the effects of the Mariel Boatlift, a large inflow of Cubans into Miami in 1980, first studied by David Card (1990). This method improves on previous studies by choosing a control group so as to best match Miami's labor market features before the Boatlift. We also provide reliable standard errors for the inference. Using data from the larger and more precise May-ORG Current Population Survey (CPS) one finds no significant departure of wages and employment of low-skilled workers between Miami and its control after 1979. The result is robust to several checks.

JEL Classification: J3, J61

Keywords: immigration, wages, mariel boatlift, synthetic control method, measurement error

Corresponding author:
Vasil Yasenov
Department of Economics
University of California, Davis
1 Shields Ave
Davis, CA 95616
USA
E-mail: viyasenov@ucdavis.edu

* We thank Colin Cameron, David Card, Michael Clemens, David Green, Pat Kline, Thomas Lemieux, Doug Miller, Joan Monras, Enrico Moretti, David Roodman, Shu Shen and participants in seminars at University of British Columbia, Georgetown University, NBER Labor Studies 2016 Meeting, IZA Summer School of Labor Economics 2016, and 2017 ASSA Annual Meeting for useful comments. We are grateful for the criticism contained in Borjas (2017) to the first draft of this paper, as it has helped our current revision. We have not received any financial support for this paper.

1 Introduction

How do receiving countries absorb sudden waves of immigrants? What are their immediate effects on wages and employment? How long do these effects last? Sudden and unexpected refugee waves in which immigrants have little ability to choose their destination provide cleaner exogenous supply shocks, relative to economic migrants whose inflow are more gradual, predictable and driven by local economic conditions.[1] One such episode has held a special place in the minds of the American citizens and economists as an example of an unexpected and large refugee inflow on American soil. On April 20, 1980, Fidel Castro announced he would open the ports of Mariel, in Cuba, enabling anyone who wanted to leave the country to do so. Consequently, between April and September of the same year, almost 125,000 Cubans fled to the United States' shores in what is known as the Mariel Boatlift. The majority of them settled in Miami, increasing its labor force by about 8 percent. Because most of these immigrants had little schooling, the relative increase of the low-skilled group was much larger (around 18 percent). This event provides a quasi-experimental environment to test theories of change in immigrants supply. If all other factors, such as technology, efficiency and physical capital did not respond, the nature of this sudden inflow generates high potential for short-run consequences on wages and for displacement of existing workers in the Miami local labor markets.

An early study by David Card (1990) analyzed the Mariel Boatlift. His results showed that the impact on employment and wages of low-skilled non-Cubans in Miami was insignificant.[2] This made the Boatlift a prominent example of how the predictions of the simplistic canonical model of labor demand and labor supply do not work well in analyzing the con-

[1] Several studies have examined the impact of sudden inflows, often of refugees, in Europe. Examples include Hunt (1992), Carrington and De Lima (1994), Friedberg (2001) and Borjas and Monras (2016) among others.

[2] Other studies have suggested how different channels for absorbing the Mariel Cubans might have worked, rationalizing the results within richer models. Lewis (2004) showed that less skilled Cubans were absorbed by industries that chose more "unskilled-intensive" technology and less automation. In addition, Bodvardson et al. (2008) argued that the immigrants increased significantly local demand for services, and hence also labor demand and not only labor supply.

sequences of immigration even in the short run. From a methodological point of view, the experimental design of Card's paper profoundly influenced the direction of research in labor economics (see Angrist and Pischke 2010). Card (1990) has represented the "final word" on this episode for 25 years. However, it has important econometric limitations and modern-day methods can significantly improve on them. First, when constructing standard errors Card (1990) did not take into account city-level idiosyncratic shocks that affect labor market outcomes and treated data from all workers within a city as independent observations, accounting only for classical independent measurement error on wages. Second, the choice of a control group for Miami consisting of Los Angeles, Houston, Atlanta, and Tampa Bay-St. Petersburg, was somewhat arbitrary and ad-hoc. Its validity was never formally tested. Finally only a very small set of labor market outcomes was analyzed while a more systematic and broader analysis can be performed.

Given the historical importance of the Mariel Boatlift and of the Card (1990) study, one reason to revisit it is that, since then, we have improved our methodological toolbox. The Synthetic Control Method (SCM), an econometric technique developed and used in a series of papers by Alberto Abadie and coauthors (Abadie and Gardeazabal 20003; Abadie, Diamond and Heinmueller 2010, 2015) is better suited than the original approach to addressing this type of case-studies. Its underlying idea is that a linear combination of labor markets is a better control group for Miami than any single one. The method, rather than arbitrarily choosing a single (or a group of) labor market as a control, identifies an "optimal" control group by minimizing the pre-1979 difference with Miami for a set of relevant labor market characteristics. This creates a single "Synthetic" city which serves as the control group.

This method has several advantages which address the shortcomings of the original study. First, this formalized procedure reduces the "ad-hoc" nature of choosing the control group. Second, it allows to validate the quality of the control group by checking the pre-treatment differences between the outcome variable in the treated and in the Synthetic units. Finally, by applying this method to each city and simulating a distribution of effects, we can construct a p-value for how significant is the post-treatment difference for Miami relative to the whole distribution, accounting for those idiosyncratic city-specific shocks that Card (1990) did not account for in his standard errors.

Two relevant data sources are available to analyze the impact of the Mariel Boatlift, the May-ORG (Outgoing Rotating Group) and the March extracts from the Current Population Survey (CPS). We analyze the statistical power of both datasets and we strongly conclude that ORG-CPS is far superior when measuring average weekly wage of groups in metropolitan areas. This is due to the much larger sample size and to the smaller measurement error in reporting "point-in-time" weekly wages (May-ORG), as opposed to retrospective yearly wages (March CPS), especially for low-skilled workers who are often paid by the hour. To further alleviate the measurement error problem, which is the largest concern in this analysis, we include the largest group potentially affected by the competition of Mariel immigrants. These are non-Cuban workers, with no high school degree between 19 and 65 years of age, in the labor force and not self-employed. We perform extensive analysis on other sub-groups too. As labor market outcomes we analyze log wages (annual, weekly and hourly) and unemployment rates for different groups of low-skilled workers.

Our results show no significant difference in the post-1979 labor market outcomes of high school dropouts between Miami and its Synthetic Control. Neither wages (annual, weekly or hourly) nor unemployment rates of high school dropouts differ significantly between years 1980, right after the boatlift, and 1983. The point estimate of the Miami-control departure in log wages is actually positive after 1979 in most samples. We reach the same conclusion when we consider wages in the bottom 15th percentile of the distribution for non-Cuban workers. We run "difference-in-differences" type of regressions for Miami and the Synthetic Control and we provide the simulated significance level of the Synthetic control method estimates to formally show the lack of statistically significant differences. With more reliable control group and a better inference method our results confirms the early results of Card (1990).

We then move on to address two studies that raise criticisms to the original Mariel Boatlift case study. The first one is Borjas (2017) who finds a large and long lasting negative deviation of wages of high school dropouts in Miami, between 1982 and 1985. Our analysis initially reproduces his finding, using his data and methods. Then we show that these results hold only when using the March CPS dataset and are strongest in the sub-sample obtained by eliminating women, non-Cuban Hispanics and selecting a shorter age range (25-59 years old) among high school dropouts. These very drastic choices leave the sample in Miami an in its control cities as small as 16-24 observations per year. We argue that the measurement error in the March CPS sample is so large that differences of 15 or even 20% in average wages can arise purely as consequence of that error. Second, we address the criticism of Angrist and Krueger (1999) who show the risk of drawing inference from analyzing an event with a small number of treatment and control units. They imagine a "non-existent" Mariel Boatlift in 1994 and analyze it the same way Card (1990) did, finding, erroneously, statistically significant results and arguing that in small samples deviations may arise by pure chance. Our Synthetic Control analysis obviates the problem of a small "control set" and in fact we find that a well-constructed control group for Miami in 1994 shows no significant deviations of relevant outcomes in the following years.

This paper also contributes to the literature on the effects of immigrants on wages and labor market outcomes of less educated native workers. Some recent papers analyzing US labor markets have found negative wage effects, as predicted by a simple canonical model (e.g., Borjas 2017, Monras 2015) while other have found no significant wage effects (e.g., Card 2001, 2009, Ottaviano and Peri 2012) and argued that several adjustment margins reduce the wage impact of immigrants and some degree of complementarity between natives and immigrants exists (Peri 2016). The results of this paper are in line with the second set of findings.

The rest of the paper is organized as follows. In Section 2 we begin by describing the main data source and documenting the supply shock in Miami. In Section 3 we briefly describe the Synthetic Control Method and how we implement it to study this episode. Section 4 contains our main results. Then, in Section 5 we consider the small and very imprecise March CPS sample and we show how the large negative wage effects estimated in Borjas (2017) are very noisy and sensitive to the choice of sample and outcome variable. Next, in Section 6 we analyze the Boatlift of 1994 that did not happen and we address the associated criticism of Angrist and Krueger (1999). Section 7 concludes the paper.

2 Data and First Stage

Our main data source is the combination of the May extracts from the Current Population Survey (CPS) (for years 1973 until 1978) with the Merged Outgoing Rotational Group (ORG) data from the same survey (starting from the first year it is available, 1979). Our preferred sample of workers whose earnings we analyze are non-Cuban high school dropouts, between 19 and 65 years of age, in the labor force and not self-employed. This is the group of existing workers in Miami which are likely to suffer the most severe impact from the competition of the newly-arrived Cubans. The arguments supporting our choice of sample and showing alternatives are laid out in detail in Section 5.

2.1 Number and Demographics of the Mariel Cubans

Table 1 shows the summary statistics as well as the aggregate count of immigrants and Miami workers around the time of the Boatlift. In the first column we present data relative to the labor force in Miami as of 1980. In the second one we display the same characteristics on all Mariel Cubans, identified as Cubans who arrived in the United States in 1980 and were 19 or older at arrival. In the last column we do the same for Mariel Cubans from Census 1990 who were still living in Miami as of 1990, and hence likely settled there permanently.

We obtain a total of 54,196 working-age Mariel Cubans, 56% of whom lacked a high school degree and 62% of whom were in still in Miami as of 1990. Hence either at arrival or in the successive years about two out of five Mariel Cubans relocated to other places. The share of Cubans in Miami, in fact, peaks in 1981 and then declines between 1981 and 1985. As we show below there is some evidence in the (noisy) CPS data for Miami that some of the Cubans who arrived in 1980 might have left the city in the following 2-3 years.

Overall, these statistics imply that the Mariel Boatlift produced an 18% increase in the number of high school dropouts in the Miami labor market, while for the other education groups the increase was only 5% and for the total labor force it was 8.4%. The Mariel population and the 1980 Miami labor force are similar in other demographic characteristics like gender and share of young individuals. This is true even conditional on high school dropouts.[3]

2.2 Measured Labor Supply Shift in Miami

Figure 1 shows the share of Cubans in Miami's population, age 19-65 (dark lines) and in the population with no high school degree age 19-65 (lighter line), between 1973 and 1985. We adopt the conventions in all figures, to call "1979-Pre" the data relative to the last observation before, and as close as possible to, the Mariel Boatlift and "1981-Post" the first data point after all Mariel Cubans had arrived.[4] Panel A shows the March CPS data, while Panel B uses the May-ORG CPS. Paying attention to the pre-post details around 1980 allows us to align our data precisely around the Miami Boatlift shock. Identifying a clear jump upwards of the Cuban share from "1979-Pre" to "1981-Post" would be the mark of the

[3] The numbers presented in this Section are in accordance with previous studies including Card (1990) and Borjas (2017). As this supply shock took place over the course of a few months it was certainly exceptional. The most significant change analyzed in other "quasi-experiments" literature is the inflow of Russians to Israel (Friedberg 2001) which was equal to 12% of initial population but took place over 5 years (between 1989 and 1994).

[4] The first one is usually year 1979. For wages and unemployment in the March CPS we use data collected in 1980, which is relative to the previous year. For the ORG-CPS we use data collected in year 1979. We also take the convention, in each figure, of showing a vertical bar exactly at the last pre-shock period (hence on "1979-Pre"). This notation helps to visually identify the last period of the status quo, right before the shock. To the immediate right of the bar we can see the impact of the sudden shock. To its left we can see the trend and variation during the pre-treatment period.

Boatlift on the CPS data.[5]

Three facts emerge from Figure 1. First, the March CPS and May-ORG CPS data show similar shares of Cubans in the total population but they are more noisy and less consistent with each other for the share among high school dropouts. This suggests significant noise and discrepancies between the May-ORG and the March CPS statistics when the sample size is small. Second, both time series and samples show an increase between the "1979-Pre" and "1981-Post" data points. Compared to trends and year-to-year movements before and after, the 1979-1981 increase, however, does not seem particularly large. Considering May-ORG CPS figures, the 1979-1981 increase as percentage of the population equals about six points and as percentage of high school dropouts the increase was around 12 points.[6] Third, after initially increasing between 1979 and 1981, the share of Cubans decreased in the following four years (in both samples) and this effect was stronger for the share among high school dropouts. In year 1985 the share of Cubans among high school dropouts is back at percentage levels comparable to those of the pre-Boatlift period. This emphasizes the temporary nature of the shock that was completely indiscernible in the CPS data by 1985.[7]

3 The Synthetic Control Method

The Synthetic Control Method (SCM), first introduced by Abadie and Gardeazabal (2003) and then further developed in Abadie et al. (2010), provides a systematic way of analyzing the impact of case-study events such as the Mariel Boatlift. Typically in these

[5] Notice that for the March CPS, as the data on demographics are relative to the month of March CPS, the last pre-treatment observation is the one collected in March CPS 1980, and it is called ("1979-Pre"), and it is differentiated from the 1979 (March CPS 1979). The 1981-Post is the observation for March CPS 1981, while "1980-shock" is simply the linear interpolation of 1979 and 1981. This is done only for this graph, due to the timing of the March CPS enumeration, that in 1980 was just before the Boatlift.

[6] These figures are consistent with the ones obtained from the Census and described in Section 2.1.

[7] One can also show that the inflow of Cuban high school dropouts did not move the non-Cuban high school dropouts out of Miami. Figure A1 in the Online Appendix shows that in the period 1979-83, non-Cuban high school dropouts (dashed-dot line) remained flat, as share of the Miami population, while Cuban dropouts as share of the population (dashed line) temporarily increased between 1979 and 1981, right after the boatlift. Hence, the local supply of all high school dropouts (solid line) increased between 1979 and 1981 as a result of Mariel, with virtually no change in the non-Cuban dropouts, revealing no offsetting outmigration of natives.

settings a single unit (often represented by a city or a region) experiences, at a point in time, an event (or treatment) while the rest do not. In order to evaluate whether the treatment had an impact on some outcomes in the treated unit, relative to what would have happened in its absence, the method formally identifies a control group called the Synthetic Control unit.

In our case, we consider $J+1$ metropolitan areas indexed by $j = 0, 1, 2 \ldots J$ and denote Miami as 0 while we call the group of all the rest "the donor pool". This is the group of other 43 cities in the United States for which data is consistently available in the Current Population Survey. Then we define a vector G_0 of dimension $k \times 1$ whose elements are equal to the values of variables that help predict the wages of high school dropouts in Miami, in the years 1972 to 1979. Then we similarly define a $k \times J$ Matrix, G_J in which row j is the sequence of values for the same variables and years relative to city j in the "donor pool".

The SCM then identifies the vector of weights $W^* = (w_1, \ldots, w_J)$ that produce a convex combination of variables in cities in the donor pool, G_J so as to approximate as close as possible, in terms of a quadratic error, the pre-treatment vector of variables chosen for metropolitan area 0, G_0. In other words, it minimizes the difference between G_0 and $G_J W$:

$$W^* = \arg\min (G_0 - G_J W)' V (G_0 - G_J W) \quad \text{s. to} \quad \sum_{j=1}^{J} w_j = 1, w_j \geq 0 \tag{1}$$

Once we have identified W^* we can use it to calculate the post-treatment outcome variables for the Synthetic Control unit by weighting each city appropriately.[8] The basis to evaluate if the treatment has had any effect is comparing the pre-post 1979 change in the outcome variable for Miami relative to the pre-post change for the Synthetic Control.

As there is discretion in choosing what variables to match in the pre-treatment period, it is important to validate the choice of the control group. To do so we can check the pre-

[8] V is a $k \times k$ diagonal, positive-definite matrix that determines the weight for each element of the vector in the objective function. We use STATA's default option for the matrix V which is chosen among all diagonal and positive definite matrices to minimize the average squared prediction error of the outcome variable during the pre-shock period.

intervention (1972-1979) levels and trends of the outcome variable to see how closely Miami and the Synthetic Control group track each other. Large differences in the pre-treatment path between the two time series would cast doubt on the validity of the chosen group as a control. As a more formal test we also check whether there are any statistically significant differences between outcomes, pre-event, in a linear regression environment.

We use the following variables in the pre-treatment distance minimization: the outcome variable itself for selected pre-1979 years, the share of low-skilled workers, the share of Hispanics and the share of manufacturing workers in the labor force. These are all important characteristics in predicting the labor market outcomes of low-skilled workers.[9] To allow for a reasonably long pre-treatment period we go far back to year 1973. Most results we show are robust to (small) variations in the selection of the Synthetic Control.[10]

4 Empirical Estimates

The Miami Boatlift was a one-time, unexpected shift in supply that took place between March and September 1980, and not a persistent policy change. In addition, as we already showed in Figure 1, the share of Cubans was back to pre-Boatlift levels by year 1985. The adjustment dynamics, then, would determine its persistence after year 1981, however the bulk of the effect on employment and wages should be detected in years 1980 and 1981. Moreover, several economic shocks (including a significant recession and a worsening of the war on drugs in 1982 which actively involved Miami) took place during the following decade and Miami could have responded differently from any control group. Therefore, deviations between Miami and the control city after year 1983 are likely caused by factors unrelated to the Boatlift.

[9] The share of Hispanic and pre-1979 employment are also among the variables used by Card (1990) to establish similarity between Miami and the control group.

[10] In Online Appendix Figure A2 we also show the analysis keeping the metropolitan areas and weights in the Synthetic Control fixed to the group that is identified to best match the log weekly wages of high school dropouts (1972-1979) between Miami and Synthetic Control.

4.1 Main Results

We show our main results in the Panels of Figures 2 and 3. As the effect of unskilled immigrants on native labor markets is predicted to be particularly strong at the low end of the wage distribution, we start analyzing those effects. Workers with no high school degree certainly belong to this segment of the labor market. Some economists (e.g., Dustmann, Frattini and Preston 2013) have argued that the best way to capture this competition effect is to look directly at natives in the bottom percentiles of the wage distribution.

Panels A and B of Figure 2 show the Synthetic Control results considering the 15th and the 20th percentiles of the natives' weekly wage distribution as outcome. An advantage of this method, relative to looking the average wage of a small group such as the high school dropouts (shown in Figure 3), is that the sample used in estimating the statistic is larger. Hence, the results are less sensitive to extreme values of wages and therefore somewhat less volatile. The footnotes to each figure indicate which cities enter in the Synthetic Control and their associated weights.[11] The solid line tracks the outcome for Miami and the dashed one represents the outcome for the Synthetic Control labor market. A reasonably good match of the pre-1980 trend is needed to consider the Synthetic Control as a good placebo group. Overall the control cities we construct in this Figure do a reasonable job in tracking the dropout wages and matching the 1972-1979 Miami trend in each of the Panels.[12] During the 1979-1981 period both Panels A and B show no notable deviation of the Miami wage trend relative to Synthetic Miami, with small positive differences in favor of Miami in the 1981-1985 period.

As a falsification check we also show, in Panel C, the time evolution of the 90th native

[11] Notice that in every Panel the Synthetic Control is constructed (at least partly) from cities not included in Card (1990)'s control group (i.e., Los Angeles, Houston, Atlanta, Tampa Bay-St. Petersburg). Because we match similar variables as the ones he based his decision on and we make no arbitrary decision regarding which cities enter the control, by construction we improve on his identification strategy.

[12] Still, due to noise in the data, deviations between Miami and Synthetic Control in the order of 0.01 to 0.05 logarithmic points are common. This implies that such level of noise could make it hard to identify deviations of average wages between Miami and the Control in the order of 1 to 4%. Nevertheless, we should be able to discern if differences (in the order of 7-8%) suddenly arise between Miami and its control in the aftermath of the Cuban inflow between 1979 and 1981.

wage percentile in Miami and the Synthetic Control. Classic labor economics theory would argue that this group of workers should not experience any effect (or possibly experiences a small positive effect from complementarity forces) on their wages. Again, we see no significant deviation between Miami and its control group. Neither at the bottom nor at the top of the wage distribution do we observe any notable departure between Miami and Control after 1979.

A possible explanation for the small wage effects is that wages were rigid downward in the years 1979-1981 and hence a negative demand shock for native workers did not translate into lower wages in Miami. However that would result in displacement of natives in the workplace rather than wage adjustment. If this explanation is correct the inflow of Mariel Cubans must be associated with an increase in the unemployment rate of non-Cuban high school dropouts in Miami. Panel D of Figure 2 shows the unemployment rate of this group for Miami and and the Synthetic Control. The year-to-year volatility of unemployment in Miami before year 1979 is quite large and in particular we see a spike for Miami in 1975 that is not matched by the Synthetic Control. With this caveat in mind we observe that the 1980-1985 behavior of Miami unemployment rate relative to the Synthetic Control does not show any significant departure. Even immediately after the shock in 1980 and 1981 no significant difference between the unemployment rate in Miami and Synthetic Control arises.

Next, in Figure 3, we present the results for average weekly wages (Panels A and C) and average hourly wages of non-Cuban high school dropouts (in Panels B and D). Hourly wages are calculated by dividing the weekly earnings by reported hours worked per week. One may argue that this measure is closer to capturing the marginal productivity (and hence price) of labor. Panels A[13] and B show average wages while Panels C and D use regression-adjusted wage measures. It is well-known that the average wages of different demographic groups such as women, men, young, old, Hispanic and non-Hispanic workers had different national

[13] Figure 5 in Doudchenko and Imbens (2016) replicates and extends the Figure in Panel A. The authors generalize the Synthetic Control estimator by considering various ways of relaxing some of the restrictions imposed in it. Their results concerning the Mariel Boatlift episode further confirm our findings.

trends in the 1970s and 1980s. Depending on the demographic composition, these trends may affect the labor markets differentially, introducing confounding factors in the analysis. The commonly used method for reducing the potential confounding effects of differential demographic characteristics (age, gender and ethnicity) is to adjust individual log wages by running the following regression, which controls for demographics:

$$\ln w_{ijt} = \beta + \beta_{AGE} * (Age)_{it} * (Year)_t + \beta_{female} * (Female)_{it} * (Year)_t + \beta_{Hisp} * (Hisp)_{it} * (Year)_t + \varepsilon_{ijt} \tag{2}$$

The wage of (high school dropout) individual i in metropolitan area j in year t is regressed on a set of five-year age dummies $(Age)_{it}$ interacted with year dummies, $(Year)_t$, on a female dummy, $(Female)_{it}$ interacted with year dummies and on a Hispanic dummy, $(Hisp)_{it}$ interacted with year dummies. This produces the residual ε_{ijt} that captures individual (and city) log wage variation once those aggregate trends are accounted for. We then implement the Synthetic Control Method on those residuals, averaged by city, showing the results in Panels C and D of Figure 3. These time series show significant noise, in the form of year to year fluctuations both before and after 1979 and the match of Miami and Synthetic Control before 1979 is not very accurate. This is reasonable as we are now trying to match "residualized" wages which are noisier measures than average wages. Overall, despite the noisier outcome measures, the results from this Figure confirm the ones from Figure 2.

The key takeaway from Figures 2 and 3 is that the average labor market outcomes of low-wage workers, and high school dropouts in Miami do not show any negative break or jump in correspondence of the Mariel Boatlift relative to the control group. Because the constructed control group mimics well the wages in Miami, before 1979, we are reasonably confident in the credibility of these results.

4.2 Subsamples

By restricting the focus to sub-samples of high school dropouts one faces the serious risk of introducing substantial measurement error as the sample size include only few dozens individuals in each metropolitan area. However, where possible, it can be interesting to separate the impact across various demographic groups to check whether there is evidence of heterogeneous effects. While the Mariel Cubans were divided between genders in similar proportions as the preexisting labor force, one may still think that their impact was differential between men and women. This may be the case if, for instance, they specialized in occupations that were in more direct competition with the male labor force. Panels A and B of Figure 4 show the Synthetic Control analysis, when separating the labor force by gender. While the time profiles of the wages of both groups are different we do not notice unusual deviation for Miami from control after 1979 in either one. Second, one may think that non-Cuban Hispanics in Miami were likely to be prior immigrants and they should be separated when evaluating the impact of Mariel Cubans. Moreover, Hispanic and African-American workers could be particularly exposed to the new immigrant competition as they may be more similar in terms of skills and occupational choices. Panels C and D of Figure 4 show the results for those two groups, separately. Even in these cases, Miami and control exhibit very similar trends of average log wages in the 1970s and negligible deviation post-1979.

4.3 Regression Analysis

A main drawback of the original analysis of the Mariel Boatlift by Card (1990) was neglecting aggregate city-specific shocks in conducting statistical inference. This is problematic because it treats contemporaneous observations from different workers within the same labor market (and hence subject to the same labor demand shocks) as independent data while in fact they could be strongly correlated due to city-specific idiosyncratic shocks. In the following two subsections we address this issue within the context of the SCM.[14]

A limitation of the Synthetic Control Method is the difficulty in constructing confidence intervals and hence conducting credible statistical inference. One approach we undertake in this direction is to use classic regression analysis for the two time series, Miami and its respective control city. The main goal of this exercise is to quantify the noise in the data and give a sense of statistically significant difference between the two. The pre-Boatlift differences Miami-Control in this regression will also provide a validity test for the comparison group selected by the Synthetic Control Method. We estimate the following regression:

$$\begin{aligned} y_{it} &= Miami_i + \sum_{P \in PRE-79} \alpha_P D_P + \sum_{P \in POST-79} \alpha_P D_P + \\ &+ \sum_{P \in PRE-79} \beta_P (D_P * Miami_i) + \sum_{P \in POST-79} \beta_P (D_P * Miami_i) + \varepsilon_{it} \end{aligned} \tag{3}$$

The variable y_{it} is the outcome of interest (e.g., average log of weekly wages of high school dropouts) in unit i which can take only two values, either Miami or its respective Synthetic Control, and year t, between 1972 and 1991. The variable $Miami_i$ is a dummy equal to one for Miami and zero for the Synthetic Control. Next, D_P is a set of 3-year dummies that span the whole period but omit 1979, which is absorbed in the constant and hence serves as reference year. In the pre-1979 period the dummies are D_{73-75} and D_{76-78} in the post-1979 period they are $D_{80-82}, D_{83-85}, D_{86-88}$ and D_{89-91} and they equal one in the years indicated in the subscript and zero otherwise. Next, α_P is the set of coefficients corresponding to the period dummies and β_P is the set of coefficients associated to the interaction between the dummy $Miami_i$ and the time period dummies. The term ε_{it} captures the classical city-level error term, uncorrelated with the observables that we interpret as residual measurement error for each metropolitan area. The method of estimation used is Feasible Generalized Least Squares allowing the measurement errors to be autocorrelated in an AR(1) process.

[14] For instance, in modern-day regression analysis it is common to apply a "cluster-robust" adjustment to the standard errors of the estimated coefficients to account for aggregate level shocks which are common to all observations within the "cluster" (e.g., Cameron and Miller 2015).

The coefficients of interest are β_P's. In particular, if the Mariel shock had any labor market effect, this should be captured by the coefficient β_{80-82}. It presents the average difference between Miami and Synthetic Control arising in 1980, 1981 and 1982 once the 1979 difference is standardized to zero. Just as importantly, our framework allows us to estimate the pre-1979 differences between the two cities. The estimates of β_{73-75} and β_{76-78} provide validation for how well the two cities track each other before the shock. Statistically significant pre-1979 differences would cast doubt on our control group as they will imply systematic deviations between the Miami and control, even before the treatment (Boatlift). The subsequent coefficients $\beta_{83-85}, \beta_{86-88}$ and β_{89-91} complete the picture.

Table 2 shows all the estimated β_p coefficients for a given outcome variable. The header of each column indicates the Panel and Figure corresponding to the estimated regression as well as the dependent variable. The estimates are simply a quantification of the deviations between the time series represented in Figures 2 and 3 with the provision that the regression standardizes this to 0 in 1979, (which the graphs do not do).

Some consistent features of Table 2 estimates are worth noting. First, the estimated coefficients $\hat{\beta}_{73-75}$ and $\hat{\beta}_{76-78}$ are, for the most part, not statistically distinguishable from zero. This is a more formal test that Miami and its Synthetic Control move together to a reasonable extent, in the pre-Boatlift period, validating our identification strategy. Second, the standard errors for the wage regressions are not small (between 0.04 and 0.06 log points) and deviations in the order of few percentage points would be difficult to measure. Year-to-year fluctuations of 5-6 per cent seem common both before and after 1979 and it is hard to say whether that are real or measurement error. Finally, none of the coefficients for the period after the Boatlift, is significantly different from zero. Moreover the point estimates of all β_{80-82} relative to a wage outcome (Columns 1-5) are positive and they reveal that Miami had a small departure upwards relative to its Synthetic Control after the Boatlift. Given the estimated coefficients and standard errors we can rule out a negative effect larger in absolute value than three or four per cent. In any case deviation of all labor market outcomes for less skilled workers after 1979 up to 1982 between Miami and Synthetic Control are all well within the margin of error.

4.4 Inference Using Permutations

While the regression approach has its appeal and simplicity, the small number of time series observations and the imprecision of the estimates affects its credibility. An alternative and perhaps more accurate way of doing inference with the Synthetic Control Method proposed by Abadie et al. (2010) is based on permutations. The core idea is to simulate a distribution of deviations between each city in the donor pool and its Synthetic Control and examine whether Miami shows a post-1979 deviation from its Synthetic Control that is large relative to the whole distribution of city-control deviations.

Panels A-D of Figure 5 do exactly this, analyzing non-Cuban high school dropouts log weekly and hourly wages (Panels A and B), log wages at the 15th percentile of natives' wage distribution (Panel C) and unemployment rate of high school dropouts (Panel D). The dark line in each Panel corresponds to Miami-Control deviations, while each of the lighter ones corresponds to one of the 43 cities' deviation from their Synthetic Control. All Panels reveal that Miami is a rather average city in the pre-1979 deviations from its Synthetic Control. regardless of the dependent variable. Then Panels A and B show that Miami's average wage had a positive deviation from its control, in the high end of the simulated range, in 1980-1982 while its deviations look within the range of idiosyncratic variation after that year. Panels C and D show that for the 15th wage percentile and for the unemployment rate Miami is well into the range of simulated deviations any year post-1979. Notice that the range of simulated idiosyncratic noise in the sample can be quite large. For instance log weekly and hourly wages show a range of noise spanning the interval between -20% and +20%. Let us emphasize once again that with this degree of noise it may be hard to identify effects in the order of five or six percentage points.[15]

[15] In Appendix Table A1 we show test statistics based on the simulations reported in Figure 5. We

5 March CPS Sample

In this Section we use the March CPS to verify our main results presented earlier. We also compare the measurement error in the March CPS sample to possible wage effects of the Mariel Boatlift, obtained from a canonical model. Because of the size of the error in this sample we argue that not much can be learned about possible wage effects in March CPS. Then we compare our estimates with those shown in Borjas (2017) which are based on the March CPS. In his analysis Borjas uses a smaller sample of non-Hispanic males, age 25-59 years old along with the March CPS and obtains a large negative deviation of their wages compared to the ones in the respective control city. We first replicate his results and then show their sensitivity to variation in outcome variable and sample definition.

5.1 Sample Size and Measurement Error: March CPS and May-ORG CPS

Table 3 shows the number of observations for our and Borjas (2017) preferred samples in Miami. Using the March CPS (first column) our sample includes only 60-80 observations per year which could certainly raise concerns of significant measurement error. The May-ORG CPS, instead, is relatively small between 1973 and 1978 (comparable to the March CPS), but beginning in 1979 it consists of usually around 150 observations. While still not too large, these numbers of observations are much closer to comfort for any empirical researcher. We will comment on the remaining two columns of Table 3 in Section 5.3 below.

Besides the larger sample size of the ORG-CPS, there are several additional reasons to

first calculate the Pre-Post ratio in the average absolute deviation of Miami from its control, considering 1980-1982 as the post-period and, alternatively, either the 1972-1979 (upper Panel) or the more recent 1977-1979 interval (lower Panel) as the Pre-period. This procedure adjusts the post-period differences for the idiosyncratic deviations experienced in the pre-1979 period. We then do the same for all other 43 cities in the sample. In the table we show the rank of Miami in the distribution of 44 cities and the probability that a random city in the distribution has a statistics larger than Miami (i.e., p-value). Along the way we implement the correction technique of Ferman and Pinto (2015) who derive conditions under which inference in Synthetic Control corrects for heteroskedasticity. A low value of the rank and a value of the p-statistics higher than 0.10 indicates that Miami deviations are not unusual relative to the other cities. The results are in accordance with the respective placebo simulation graphs.

believe that its hourly and weekly wages figures are superior to the March CPS.[16] First, the variation of measurement error in average wages across cities in the March CPS sample is much larger than in the May-ORG CPS sample. One way to show this is to assume that, for year 1979, the 5% sample of the Census 1980 allows to calculate the ”true” average city log wages for the considered group of high school dropouts. We then compare these with the ones estimated from March CPS data (from year 1980) and the ORG-CPS data (from year 1979). We first calculate average log weekly wages for our preferred sample in each of the 41 metropolitan areas available in all three datasets. Then, we calculate the difference of these average log wages in each metropolitan area, between the March CPS and the Census and we do the same between the ORG-CPS and the Census. The measurement error in the March CPS has a standard deviation of 0.12 logarithmic points (about 12%) while the one in ORG-CPS has a standard deviation of about half of that (only 0.07 log points, 7%). Consequently, in the March CPS, a difference in average wages in the order of 12-15% between two cities would arise, frequently, by simple measurement error.[17] As long as the reasonable wage differences to be identified are smaller than 15% the noise of the March CPS sample is simply too large to have any power in identifying such an effect.

A second reason to be skeptical about the use of the March CPS in measuring weekly wages, especially of low-skilled individuals, was proved by Lemieux (2006) and Bollinger (1998). They showed that the March CPS wage data, based on the recollection of previous year annual salary and previous year weeks worked, compound recall and division errors which are particularly severe for people who are paid by the hour. This includes a large fraction of the high school dropouts which are the main subject of this analysis. People

[16] Perhaps due to the reasons outlined in this Sections, virtually all previous studies analyzing labor market aspects of the Mariel Boatlift also choose the ORG-CPS. These include Card (1990), Angrist and Krueger (1999), Monras (2015), Bodvardson et al. (2008), Laing (2011), Borjas (2012), Cahuc et al. (2014), and Doudchenko and Imbens (2016) among others. To our knowledge the only study to prefer the March CPS is Borjas (2017). Moreover most of the studies we are aware of that analyze wage dispersion and the effects of shocks on low-skilled worker wages at the sub-national level (e.g. recently Autor, Manning and Smith 2016 analyzing minimum wage in US states) use the ORG-CPS.

[17] Assuming normal and independent measurement errors across cities, the mean absolute difference between measurement errors for two randomly chosen cities will be equal to $\approx 2/\sqrt{\pi} * Std.dev.$

often have a difficult time recalling the exact actual number of hours worked last year, a figure used in constructing weekly wages. The impact of this measurement error will certainly be magnified in the very small sample at the metropolitan area level. To the contrary, the May-ORG CPS sample based on weekly wage recall from the last week of work produces a less noisy and more reliable estimate of earnings for individuals paid by the hour.

Yet another argument in favor of May-ORG CPS is the fact that the March CPS makes available fewer cities to be included in the control group, 31, relative to the May-ORG CPS, 44. All else equal, this reduces the probability of finding a good control group for Miami pre-1979 even regardless of the employed method.[18] In summary, we are convinced that May-ORG CPS is far superior to March CPS and that, when analyzing labor market outcomes in metropolitan areas during the 1980's, the March CPS data should be avoided.

5.2 March CPS: Estimates

Despite the apparent flaws of the March CPS dataset we proceed with presenting estimates of the impact of the Mariel Boatlift using the Synthetic Control Method on this sample. Our main results are presented in Figure 6 using our preferred sample of non-Cuban Hispanic workers aged 19-65. Panels A and B present results for high school dropouts log wages, Panel C does the same for the 15th percentile of the weekly wage distribution (among the whole non-Cuban labor force) and Panel D shows the evolution of the unemployment rate among the unskilled non-Cuban workers. The solid line shows Miami and the dashed one is the constructed control labor market. The footnote lists all the cities with positive weights in the SCM control group.[19]

The obtained results are rather similar to the ones shown in Figures 2 and 3 using the

[18] Perhaps a fourth argument in this directions is that the small March CPS dataset seems to fail the falsification test presented in Panel C of Figure 2. In Figure A3 in the Online Appendix we show that using March CPS one finds a decrease of the 90th wage percentile in Miami compared to the control city, post-1979, which would be very hard to relate to the Boatlift.

[19] In Figures A3 and A4 in the Online Appendix we show more results using the March CPS dataset mirroring most of the Panels from Figures 2-4. In Table A2 in the same Appendix we show additional regression tables similar to Table 2 again using the March CPS.

ORG-CPS, in that, for the pre-Boatlift period, the constructed control tracks the dynamics of the labor market in Miami reasonably well. Panels A and B show a small dip in Miami's log wages in the period 1980-1982. However, the larger noise, especially in Panel B implies that the post-1979 changes look similar to the year-to-year fluctuations before and after the shock. A somewhat more substantial departure of Miami from control takes place around 1984-1985, four years after the Boatlift. It would be very difficult to justify any connection of this departure to the Boatlift itself. Panels C and D show absolutely no systematic deviations of the 15th wage percentile and of low-skill unemployment between Miami and its control after 1979. In spite of the smaller power of March CPS, and large fluctuations we do not find any systematic deviation post 1979 relative to the years preceding it. In this sense March CPS confirm the findings of Section 4.

5.3 The Sensitivity of Borjas (2017)'s Results

Before we get into any further empirical results, knowing the magnitude of the labor supply shift due to the Mariel Boatlift allows us to calculate the short-run prediction of a simple model with fixed capital and technology, perfect substitutability between immigrants and natives and reasonable elasticity of labor demand. This is an useful exercise because it will tell us the *the most negative* impact from the Mariel inflow that a partial equilibrium model with reasonable parameters can predict. It should serve as a benchmark for the empirical estimates.

We use a very standard model, as used in Autor Katz and Kearny (2008). We consider output as a function of total factor productivity, physical capital and a labor aggregate made of workers with less than high school diploma and workers with high school diploma or more. These two groups are combined in a Constant Elasticity of Substitution (CES) function.[20] We then calculate the effect produced by an increase in aggregate labor, by 8% and in the supply of high school dropouts by 18% which correspond to the changes in Miami's labor

[20] In an earlier version of this paper (Peri and Yasenov 2015) we write out and describe the model in greater technical detail.

supply due to the Mariel Boatlift as measured in Section 2.1.[21] One can then easily obtain the short-run (partial) wage effect from the model above, considering that the wage of low-skilled workers equals their marginal productivity and differentiating it with respect to percentage (logarithmic) change in the total and low-skilled labor forces. Capital and productivity are kept constant. Considering commonly used value for the share of income to physical capital (0.33) and the elasticity of substitution between high- and low-skilled workers (set equal to 2, in the range estimated by Katz and Murphy 1992, Angrist 1995, Johnson 1997, and Krusell et al. 2000) the predicted effect for "short-run" is equal to -7.6%. Even the largest predicted negative effect of the Boatlift from a canonical model is well within the average difference between two cities' measurement errors in March CPS, calculated to be about 15% in Section 5.1 above. Hence it is clear that any evidence from the March CPS dataset has to be taken with a very high degree of caution.

Next, let's notice here that the sample of workers in Borjas (2017) is male, non-Hispanic, in the 25-59 age range. Unfortunately, this choice of a sample excludes two thirds of low-skilled Miami workers and the very small March CPS sample adds to the problem. In the end, to calculate wages of native dropouts in Miami, a group that included about 120,000 individuals as of 1980 (as shown in Table 1) Borjas (2017) uses a sample between 16 to 24 individuals per year (as shown in Column 3 of Table 3). Just as importantly, when constructing a control group of cities, many of the available metropolitan areas have even much smaller sample sizes than that of Miami. Including these as a part of the Synthetic Control analysis will also contribute to an enormous amount of measurement error.[22]

Regardless of these concerns, we continue by replicating Borjas' main results and we analyze their sensitivity to the choice of a sample as well as a dataset. Figure 7 illustrates our findings. We begin in Panel A where we essentially reproduce exactly Borjas' main

[21] Using the data in Figure 1 the supply changes in the period 1979-1981 would rather be 6% for the total labor force and 12% for the low-skilled so the values chosen in our exercise are at the upper end of the likely range.

[22] A look at this dataset with this choice of sample in the period 1979-1981 reveals that, out of the 32 available control cities, half of them had a minimum sample size smaller than or equal to 23 observations! In fact, eight of these cities have a smaller sample size than that of Miami.

estimates, except for introducing two small modifications. First, Borjas (2017) smooths the time series using a 3-year moving average and we do not do this. The procedure seems to run contrary to the identification idea, based on exploiting the sudden occurrence of the temporary shock in April-June 1980, exploiting the pre and post outcomes. By using a moving average we confound data of the pre-shock observation with 1980 data.[23] Even with these two small changes we see very clearly in Panel A the downward departure of Miami from the control that reaches a peak in 1985. Second, Borjas (2017) only includes non-Hispanic prime-age (25-59) males. We add to these workers the Hispanic non-Cuban individuals and we extend the age range to include the more often used working-age period for high school dropouts (age 19-65). Hispanic workers were mainly US born and were more similar in their jobs and occupations to the newly arrived Cubans. Hence this broader choice should improve the precision and detect a stronger effect. Individuals with weaker job protection and shorter labor market attachment (young) could also be more vulnerable to new immigrant competition. In Panel B then we present the same results using yearly (instead of weekly) wages. Let us reiterate that weekly wages from the CPS contain, by definition, more noise. In constructing them we include the recollection bias in the number of weeks worked in the past year and add it to the recollection bias on yearly wages. We move on to Panel C where, even more directly testing the sensitivity, we keep the exact sample definition of male non-Hispanic dropouts age 25-59 as in Borjas (2017), but we use the larger and more accurate ORG-CPS sample. Finally, in Panel D we mirror the choices in Panel A but again we use the ORG-CPS dataset.

The results presented in the Panels of Figure 7 show clearly two facts. First, even Panel A reproducing Borjas (2017) almost exactly but with no smoothing, shows a departure between Miami and control opening in 1982 (rather than in 1980), and the major deviation occurring

[23] With the smoothing one also includes later dynamics in the "post" observation of 1981, which should capture just the response to the shock. Borjas (2017) argues that this procedure is adopted to increase the sample size. Certainly, the 16 to 24 observations used in his analysis need some improvement, but using a moving average mechanically builds autocorrelation in the time series and one cannot consider the post-1979 observations as independent. Hence, we will not use the moving average.

in 1985. Why would an increase in supply in 1980 take two years to have its early effects and five years to have its largest effect? This is never explained or justified in Borjas (2017). Second, even small variations in variable and sample definition (as done in Panels B-D) lead to big changes in the post-1979 behavior of Miami and control, completely undoing the dip of Miami relative to the control. This suggests a very large amount of imprecision and noise. We did several other checks on sub-samples and on other variables (not shown here). The main features of those are the presence of large noise but no systematic change in the pre- to post-1979 Miami-control difference. Overall the large negative deviation of the wages of high school dropouts in Miami arises only when using the March CPS data, it is significant only in one sub-sample obtained by eliminating women, non-Cuban Hispanics and selecting a short age range (25-59 years old) among high school dropouts and its timing does not quite coincide with the Boatlift.[24] We think this is far too little to argue in favor of a causal impact of the Boatlift on low-skilled wages in Miami.

6 The Boatlift That Did Not Happen

In 1994 Fidel Castro once again announced that Cubans who wanted to flee the country could leave. This time, however, the United States Coast Guard diverted the majority of the flow to the naval base in Guantanamo Bay. A big inflow of refugees to Miami was about to take place but did not. Nevertheless, Angrist and Krueger (1999) show that between 1993 (pre-non-shock) and 1995 (post-non-shock) the unemployment rate for Black workers in Miami increased by 3.6 percentage points, while in the control group of cities in Card (1990) it decreased by 2.7 percentage points. Hence, if a researcher were to analyze the impacts of this non-event, she would estimate a fake treatment effect of +6.3 percentage points. They argue that this "false positive" is a cautionary tale when utilizing a very small

[24] Table A3 in the Online Appendix presents a regression table for Figure 7 akin to Table 2. It further highlights the statistical significance of these statements.

number of units in event studies like this one.[25]

We are sympathetic to this methodological caveat and illustrate that the Synthetic Control Method can significantly reduce this problem. By construction, the SCM eliminates the arbitrary choice of a control group (as done by Angrist and Krueger 1999, who simply use the same controls that Card had used for 1980) and allows validation of the newly-constructed one by checking the fit in the pre-1994 period. Similarly to the results presented in Figures 2, 3, and 4, we apply the Synthetic Control Method using high school dropouts' weekly and hourly wages as well as weekly and hourly wages at the 15th percentile of the respective distribution. We use the ORG-CPS dataset and our preferred sample of workers.[26]

Figure 8 shows the results. Panels A and B show the behavior of hourly and weekly wages of high school dropouts in Miami and Synthetic Control between 1989 and 2001. The rest of the Panels do the same for the lower tail of the natives' hourly (Panel C) or weekly (Panel D) wage distribution. The vertical line is on year 1993, the last year before this non-shock. The solid line shows Miami and the dashed one is the constructed control labor market. The footnote lists all the cities with positive weights in the SCM control group. In all graphs of Figure 8 we see a good fit between Miami and its respective control labor market in the years leading up to 1993 and no significant departures immediately afterwards. Overall, these figures do not produce any false evidence of a downward wage movement. Looking at them we recognize significant noise in the data but we do not to identify erroneous signs of an effect on wage and employment from the non-existent 1994 Boatlift.[27]

[25] Angrist and Krueger's argument is somewhat crude and it is simply a cautionary tale. Any serious researcher will at least dig into validating the parallel trends of Miami and the control group assumptions. Reassurance that the control group and Miami had similar labor markets in the mid 1990s (.e.g., in terms of occupational structure, demographics, etc.) is also necessary for the credibility of this strategy. Importantly, let us also point out that in 1994 the CPS underwent a major redesign and several measures of employment, especially for males and subgroups were significantly affected (see Polivka and Miller, 1995). Hence focusing on changes exactly around 1994 can be very risky.

[26] In this case, to keep computational time within a reasonable amount, we limit the "donor pool" for the control group to cities with at least 20 observations in the relevant group of high school dropouts per year. This produces a pool of about 40 cities. Moreover, in order to have a balanced panel of control cities we keep the pre-1994 period to six years only. Regression analysis of the presented time series and placebo simulations (not shown here) further confirm the obtained results.

[27] We show the behavior of the unemployment rate of minorities (Black and Hispanics) vis-a-vis the Synthetic Control in Figure A5 in the Online Appendix. While the unemployment rate of Black workers still

7 Conclusions

In this paper we apply the Synthetic Control Method to the well-known Miami Boatlift episode with the goal of improving on Card's (1990) methods. We analyze a wide variety of labor market outcomes for high school dropouts and for low-wage non-Cubans in Miami, as well as for several sub-groups. We look for a significant and sudden deviation of Miami labor markets' outcomes from those of its Synthetic Control right in 1980, as a potential evidence of a causal effect of the Boatlift on local workers. We do not find any consistent evidence of a short-run depressing effect on low-skilled labor demand nor any lasting effect later on. While one may think that the contribution of this paper is limited, as we confirm Card's (1990) results, the importance of putting those estimates on sounder ground, showing their robustness and plausibility, is quite important, in the light of recent criticism by Borjas (2017) who finds, instead, a puzzling large and delayed wage effect.

Two findings stand out. The first is that we do not find any significant deviation of Miami from Synthetic Control post-1979 for any of the considered labor market outcomes. The lack of a significant wage effect, while in part attributable to large measurement error which contributes to the noise of average wage data, is also consistent with the recent literature emphasizing mechanisms that allow absorption of immigrants. These may take place through complementarity, technology adjustment, increase in demand and efficiency. Furthermore, they go beyond the naive canonical model that is too partial in the understanding of the labor market effects of immigrants (see Peri 2016). Our second contribution is to show that a careful application of the SCM and the performance of an extensive set of robustness checks, varying sample, outcome variables and sub-groups, exposes the fragility of the criticism by Borjas (2017) and Angrist and Krueger (1999). Their claim that exist evidence of an effect (of the Boatlift or of the non-Boatlift, respectively) are predicated on very narrow assumptions and choices. Once we relax some of those we are left with imprecisely estimated values and

shows an increase relative to the Synthetic Control in 1994 and 1995, that difference is less dramatic and it is reversed by 1996. The unemployment of Hispanic individuals experienced actually a decline relative to the Synthetic Control in 1994-1995.

usually non-significant effects. Moreover, in performing the analysis, we became aware of the fact that the measurement error in constructing the average wages of the analyzed groups is so large in the March CPS sample as to prevent any reasonable inference. This problem underscores the importance of using, at the very least, the largest possible set of low-skilled workers and the larger May-ORG CPS data.

In conclusion, we think that a reasonable re-assessment of the initial Card (1990) findings of no labor market effects of Mariel Cubans on native workers, is that when we account for idiosyncratic and measurement error, we cannot rule out small wage effects in the order of 3-4% of wages. However the point estimates in most of the samples using the May-ORG data are slightly positive and do not suggest any negative impact. Finally, our analysis convinced us that the March CPS should not be used to do any exercise on wages in Metropolitan Statistical Areas because of the very small nature of its sample and the resulting extremely noisy average wages.

References

[1] 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 (2010), 493–505.

[2] Abadie, Alberto, and Javier Gardeazabal, (2003) “The Economic Costs of Conflict: A Case Study of the Basque Country,” *American Economic Review* 93, 113–132.

[3] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. (2015) ”Comparative politics and the Synthetic Control Method.” *American Journal of Political Science* 59.2 : 495-510.

[4] Angrist, Joshua (1995). “The Economic Returns to Schooling in the West Bank and Gaza Strip.” *American Economic Review*, 85, 1065–1087.

[5] Angrist Joshua and Alan Krueger (1999) ”Empirical Strategies in Labor Economics” in O. Ashenfelter & D. Card (ed.), Handbook of Labor Economics, Elsevier, edition 1, volume 3, number 3.

[6] Angrist, J. D., and Pischke, J. S. (2010). ”The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics.” *Journal of Economic Perspectives*, 24(2), 3-30.

[7] Autor David, Larry Katz and Melissa Kearney (2008) “Trends in U.S. wage inequality: Revising the revisionists” Review of Economics and Statistics 90 (May) pp.23-30

[8] Autor, David, Alan Manning and Christopher Smith (2016) “The Contribution of Minimum Wage to US wage inequality over three Decades:A Reassesment” *American Economic Journal: Applied Economics* 2016, 8(1): 58–99http://dx.doi.org/10.1257/app.20140073.

[9] Bodvarson Orn, H.F. Van der Berg and J.J. Lewer (2008) ”Measuring Immigration’s effect on Labor Demand: A Reexamination of the Mariel Boatlift” *Labor Economics*, 15:4.

[10] Bollinger, Christopher R. (1998) ”Measurement error in the Current Population Survey: A nonparametric look.” *Journal of Labor Economics* vol. 16 issue3 : 576-594.

[11] Borjas, George (2012) ”Labor Economics” 6th Edition, McGraw-Hill Companies.

[12] Borjas, George (2017) ”The Wage Impact of the Marielitos: A Reapprisal” Industrial Relation and Labor Review, forthcoming.

[13] Borjas, George, and Monras, Joan. (2016). The Labor Market Consequences of Refugee Supply Shocks (No. w22656). National Bureau of Economic Research.

[14] Cameron, A. Colin, and Douglas L. Miller (2015) ”A practitioner’s guide to cluster-robust inference.” *Journal of Human Resources*, 50(2), 317-372.

[15] Card, David (1990) "The Impact of the Mariel Boatlift on The Miami Labor Market" *Industrial and Labor Relation*, Vol 43 n. 2.

[16] Cahuc, Pierre, Stephan Carcillo, A. Zylberberg and Andre' Zylberberg (2014) "Labor Economics" Second Edition, MIT Press, Cambridge Mass.

[17] Carrington, W. J., and De Lima, P. J. (1996). "The impact of 1970s repatriates from Africa on the Portuguese labor market." *Industrial and Labor Relations Review*, 49(2), 330-347.

[18] Doudchenko, N., and Imbens, G. W. (2016). Balancing, regression, difference-in-differences and Synthetic Control Methods: A synthesis (No. w22791). National Bureau of Economic Research.

[19] Dustmann, Christian, Tommaso Frattini and Ian Preston (2013) "The Effect of Immigration along the Distribution of Wages," *Review of Economic Studies*, Oxford University Press, vol. 80(1), pages 145-173.

[20] Ferman, Bruno, and Christine Pinto (2015) "Inference in differences-in-differences with few treated groups and Heteroskedasticity."

[21] Flood, Sarah, Miriam King, Steven Ruggles, and J. Robert Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 4.0. [Machine-readable database]. Minneapolis: University of Minnesota, 2015.

[22] Friedberg, R. M. (2001). "The impact of mass migration on the Israeli labor market." *Quarterly Journal of Economics*, 1373-1408.

[23] Hunt, J. (1992). "The impact of the 1962 repatriates from Algeria on the French labor market." *Industrial and Labor Relations Review*, 45(3), 556-572.

[24] Johnson, George E., (1997). "Changes in Earnings Inequality: The Role of Demand Shifts." *The Journal of Economic Perspectives*, 11, 41–54.

[25] Katz, Larry and Kevin Murphy (1992). "Changes in Relative Wages 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics*, 107, 35–78.

[26] Krusell, Per, Lee Ohanian, Victor Rios-Rull, and Giovanni Violante (2000). "Capital–Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68, 1029–1053.

[27] Laing, Derek (2011) "Labor Economics" W.W. Norton and Company, New York-London.

[28] Lemieux, Thomas. (2006) "Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?." *The American Economic Review* : 461-498.

[29] Lewis, Ethan (2004) "How did the Miami labor market absorb the Mariel immigrants?," Working Papers 04-3, Federal Reserve Bank of Philadelphia.

[30] Monras, Joan (2015) "Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis." Manuscript, Paris Science Po.

[31] Peri, Giovanni (2016) "Immigrants, Productivity, and Labor Markets" *Journal of Economic Perspectives*, 30(4), 3-29.

[32] Peri, Giovanni, and Yasenov, Vasil. (2015). The labor market effects of a refugee wave: Applying the synthetic control method to the Mariel boatlift (No. w21801). National Bureau of Economic Research.

[33] Polivka, Anne E. and Stephen M. Miller (1995) "The CPS After the Redesign: Refocusing the Economic Lens." United States Bureau of Census, 1995; available at "http://www.bls.gov/ore/abstract/ec/ec950090.htm" (accessed on November 20th, 2015)

Tables and Figures

Table 1: Demographics of Mariel Immigrants and of Miami Labor Force in 1980

	Miami Labor Force in 1980	Mariel immigrants, measured from the 1990 Census	Mariel Immigrants still in Miami as of 1990
Total in Labor Force (16 to 65 years of age)	644, 860	87, 347	54,196
Share with no HS degree	26.28	55.77	56.34
Share with HS degree	32.11	25.18	24.22
Share with some college	22.37	12.53	12.47
Share with college	19.24	6.52	6.97
Share of female	45.79	37.80	41.97
Share of young (<25 years old)	16.35	16.34	14.01
Only individuals with no High School degree			
Total in labor force	169,440	48,714	30,532
Percentage female	43.33	39.79	44.41
Percentage young (<25 years old)	11.36	12.90	10.24

Notes: The values for the Miami Labor force are obtained from the 1980 census. Those on the Mariel Immigrants are obtained from the 1990 census as people born in Cuba who arrived in the US in 1980 and 1981 and were at least 19 years of age at the time of arrival. Labor force is defined as individual 19-65, not in school, and working or looking for a job.

Figure 1: Cubans in Miami as share of total and of the high-school-dropout population

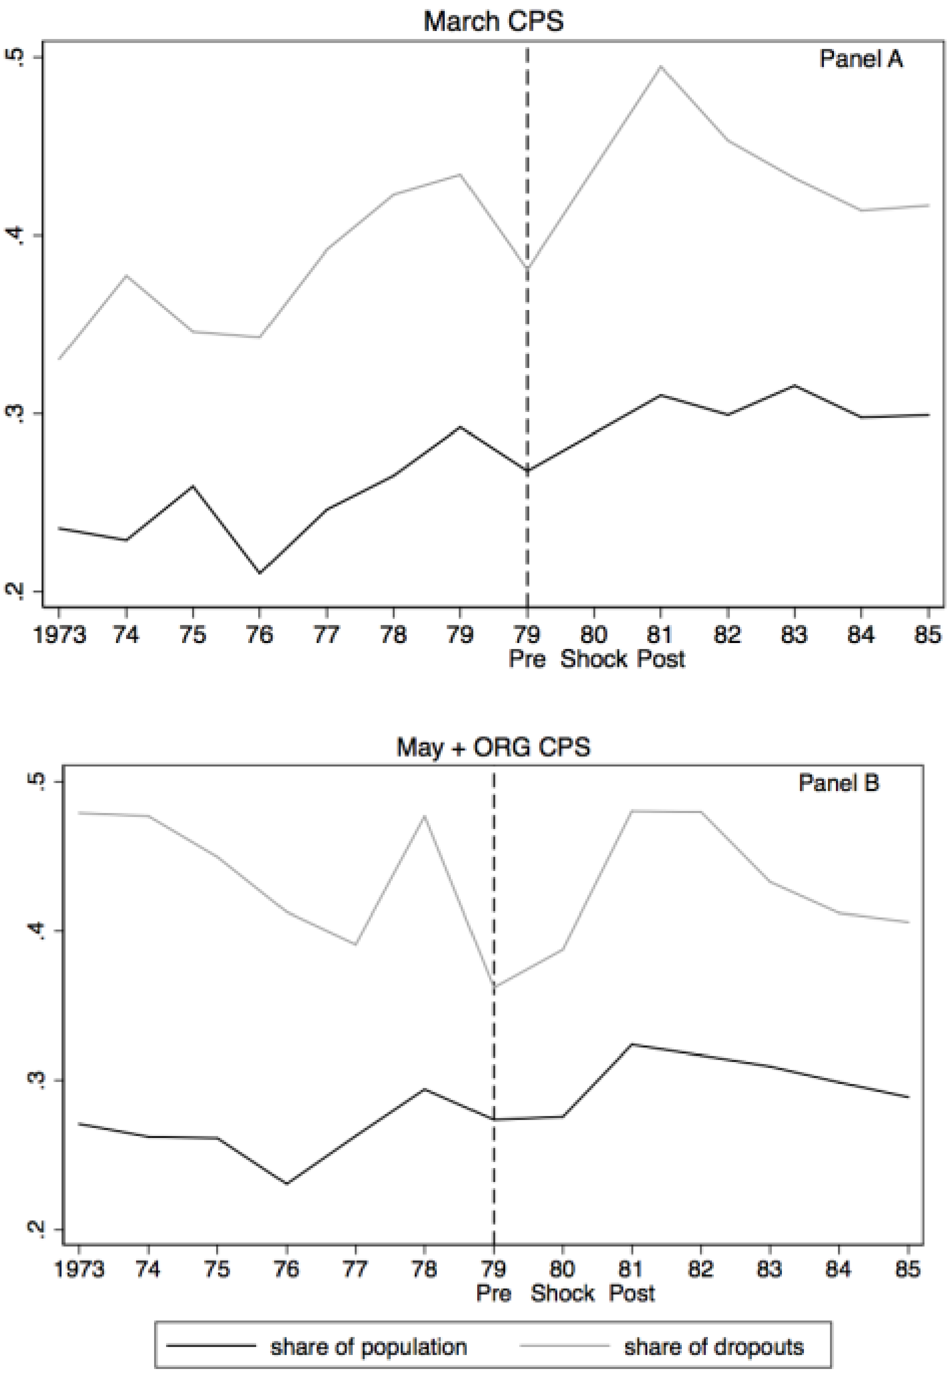

Notes: We calculate the share of all those who define themselves as "Cuban" in the ethnicity question of the CPS. The population considered is the total number of individuals between 19 and 65. The high school dropout population is constituted by those who do not have a high school degree in the age range 19 to 65. For the March CPS, we include the figure for March 1980 as "1979-Pre" and we interpolate the figure for "1980-Shock", between 1979-Pre and 1981-Post. The vertical dashed bar corresponds to the last observation before the Mariel Boatlift happened.

Figure 2: Log Weekly Wages at Different Percentiles and Dropouts' Unemployment Rate

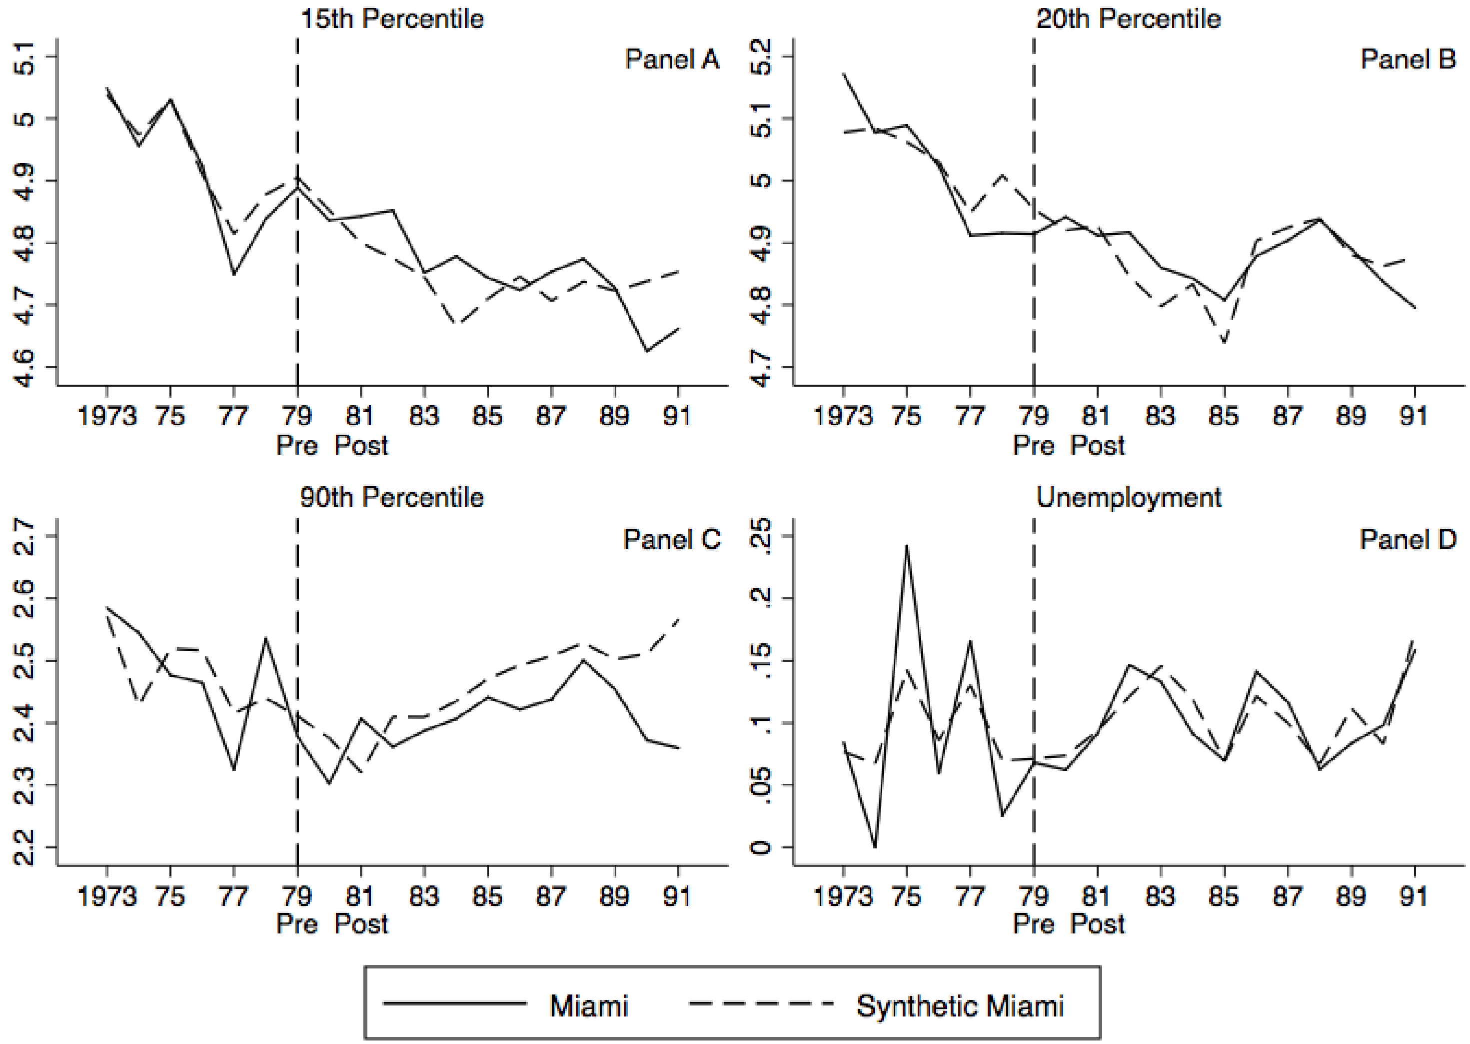

Note: The data source is May + ORG CPS. Each Panel shows the outcome variable for Miami (solid line) and Synthetic control (dashed line) in the period 1972-1991. The outcome showed for each Panel and the sample used are noted in the title of each panel. Preferred sample means: non-Cubans, not self-employed individuals, in the labor force, age 19-65 for Panels A, B and C and non-Cubans, high school dropouts not self-employed, in the labor force, age 19-65 for Panel D.-The cities with positive weight in the synthetic control are as follows. Panel A: Birmingham, AL 60.6%, Rochester, NY 28.6%, Nassau-Suffolk, NY 10.4%; Panel B: San Diego, CA 57.7%, Birmingham, AL 28.4%, Nassau-Suffolk, NY 13.8%; Panel C: Tampa-St Petersburg, FL 67.3%, Nassau-Suffolk, NY 32.7%; Panel D: New Orleans, LA 48.4%, New York City, NY 30.9%, Albany-Schenectady-Troy, NY 19.5% Cincinnati, OH 1.1%.

Figure 3: Log Wage Measures of High School Dropouts

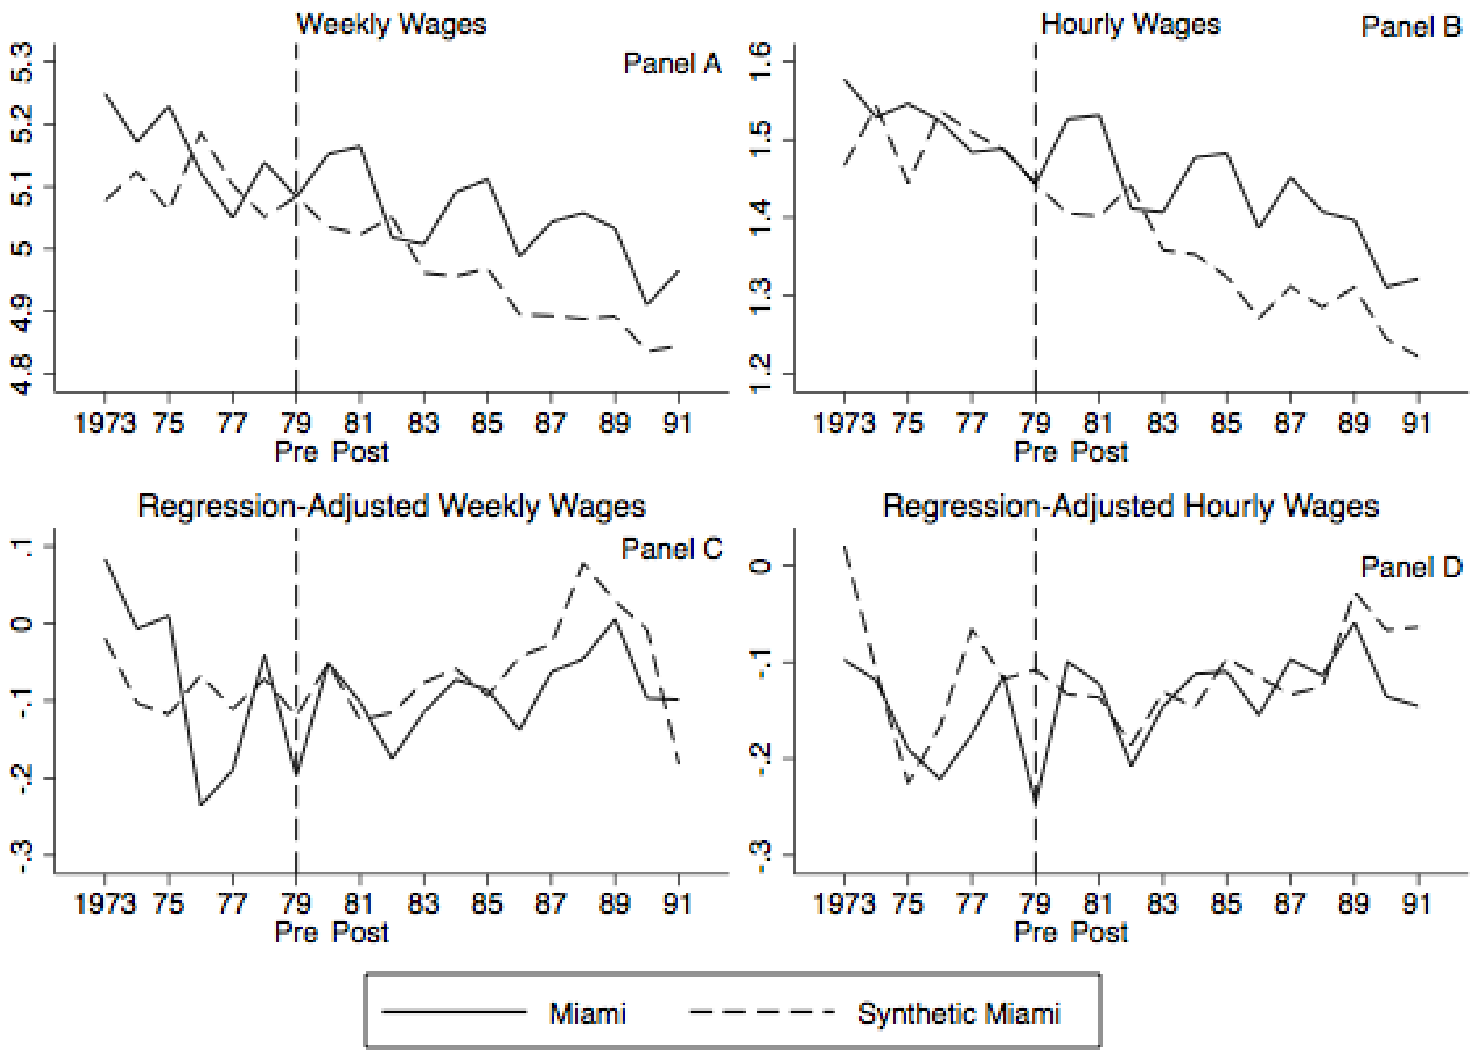

Notes: The data source is May + ORG CPS. Each Panel shows the outcome variable for Miami (solid line) and Synthetic control (dashed line) in the period 1972-1991. The outcome shown and the sample used are noted in the title of each panel (A through D). Preferred sample means: non-Cubans, high school dropouts, not self-employed, in the labor force. The age group varies upon panel. The vertical line is drawn between the data points of 1979 and 1980 and it identifies the interval in which the Mariel Boatlift took place. The cities with positive weight in the synthetic control are as follows. Panel A: New Orleans, LA 43.2%, New York City, NY, 30.1%, Baltimore, MD 24.9%; Panel B: New Orleans, LA 43.9%, New York City, NY, 29.9%, Baltimore, MD 24.8%; Panel C: Tampa-St. Petersburg, FL 45.4%, New York City, NY 24.6%, Sacramento, CA 20.8%, Albany-Schenectady-Troy, NY 9.2%; Panel D: Greensboro, NC 48.3%, Norfold-Portsmouth, VA 30.5%, Cincinnati, OH 21.2%.

Figure 4: Subsamples of High School Dropouts, Log Weekly Wages

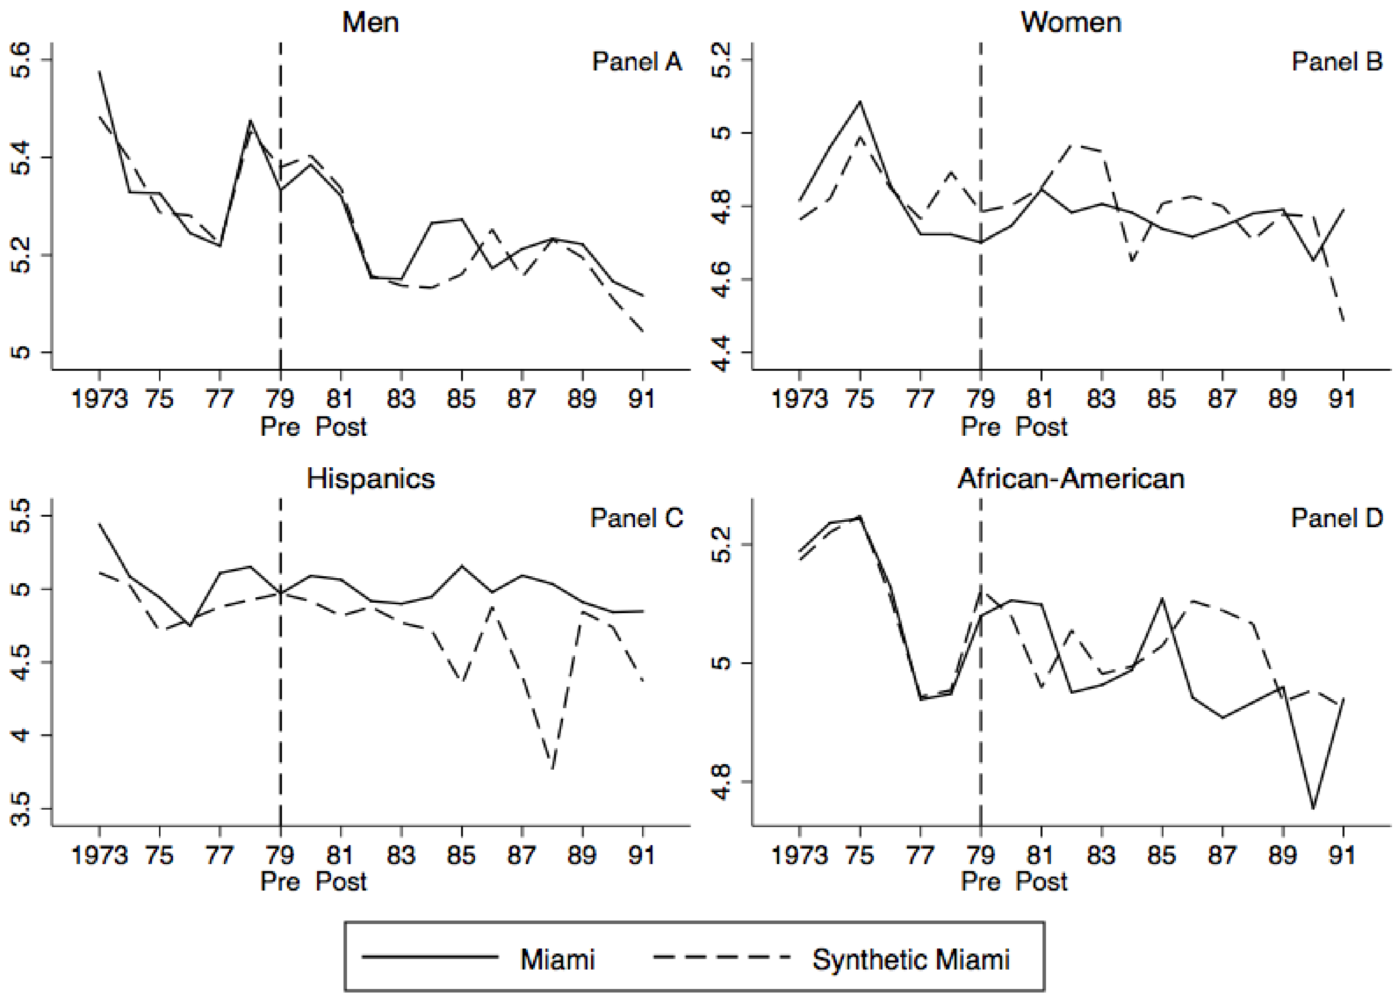

Notes: The data source is May + ORG CPS. Each Panel shows the outcome variable for Miami (solid line) and Synthetic control (dashed line) in the period 1972-1991. The outcome variable and the sample are noted in the title of each panel. Preferred sample means: non-Cubans high school dropouts in the labor force, not self-employed, age 19-65. Panel A and B restrict the preferred sample to males and females only, respectively. Panel C and Panel D further restricts the preferred sample to Hispanic and Blacks only. The corresponding minimum and maximum values of the sample sizes are: Panel A (May CPS 23-33; ORG 54-153), Panel B (May CPS 16-20; ORG 26-141), Panel C (May CPS 5-14; ORG 16-97), Panel D (May CPS 14-25; ORG 47-99). The cities with positive weight in the synthetic control are as follows. Panel A: Tampa-St Petersburg, FL 92.6%, Greensboro, SC 6.3%, New York City, NY 1.1%; Panel B: Cincinnati, OH 66.8 %, Pittsburgh, PA 29.4%, Indianapolis, IN 3.8%; Panel C: Sacramento, CA 49.8%, Houston, TX 30.1%, Philadelphia, PA 20.1%; Panel D: Greensboro, NC 39.6%, Cincinnati, OH 19.8%, New York City, NY 15.8%, Seattle, WA, 9.7%, Birmingham, AL, 8.5%, New Orleans, LA 6.7%.

Figure 5: Inference, Simulated permutations

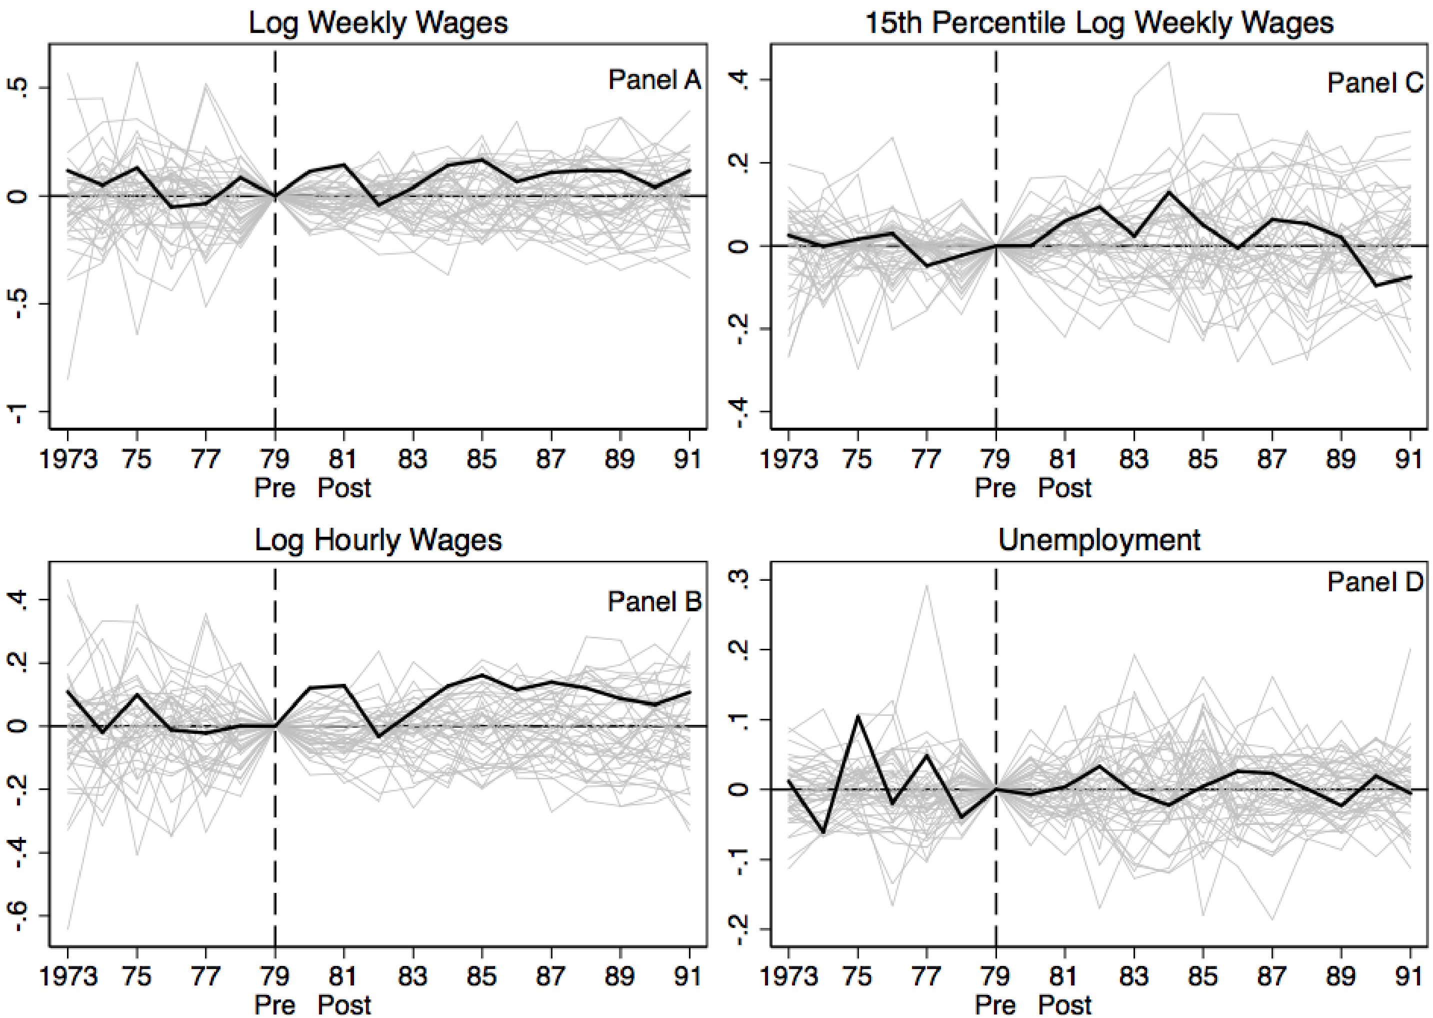

Notes: The data source is May + ORG CPS. Each graph reports deviations between synthetic control and treated group, assuming a treatment in 1980, for 44 metropolitan areas. The bold line represents Miami. Panel A shows the graph for the logarithm of weekly wages, Panel B shows it for the logarithm of hourly wages. Panel C for the 15th percentile of log weekly wages and Panel D the unemployment rate. The sample in Panel A, B and D includes non-Cuban, high school dropouts, 19-65 years old from the May and ORG CPS. In panel C the 15th percentile is calculated on all non-Cuban workers between 19 and 65 years old from the May and ORG CPS.

Table 2: Miami - Synthetic Control Regressions 1973-1991

	(1) Figure 3 Panel A	(2) Figure 3 Panel B	(3) Figure 2 Panel A	(4) Figure 2 Panel B	(5) Figure 2 Panel C	(6) Figure 2 Panel D
Dependent Variable:	**Ln (Weekly Wages) of HSD**	**Ln (Hourly Wages) of HSD**	**Ln Wages, 15th percentile**	**Ln Wages, 20th percentile**	**Ln Wages, 90^{th} percentile**	**Unemployment Rate of HSD**
Miami X ('73-'75)	0.089	0.044	0.037	0.089	0.047	0.028
	(0.055)	(0.045)	(0.043)	(0.055)*	(0.055)	(0.057)
Miami X ('76-'78)	-0.042	-0.027	0.006	0.008	-0.018	0.023
	(0.058)	(0.047)	(0.047)	(0.059)	(0.059)	(0.061)
Miami X ('80-'82)	0.043	0.057	0.071	0.077	0.007	0.027
	(0.058)	(0.047)	(0.047)	(0.059)	(0.059)	(0.061)
Miami X ('83-'85)	0.069	0.088	0.093	0.099	-0.020	0.010
	(0.055)	(0.045) *	(0.043)*	(0.055)*	(0.055)	(0.057)
Miami X ('86-'88)	0.111	0.114	0.064	0.040	-0.039	0.033
	(0.055) *	(0.045)*	(0.044)	(0.055)	(0.055)	(0.057)
Miami X ('89-'91)	0.080	0.066	-0.031	0.022	-0.115*	0.014
	(0.055)	(0.045)	(0.044)	(0.056)	(0.056)	(0.057)
Observations	38	38	38	38	38	38

Notes: "HSD" stands for high school dropouts. Each column represents a regression of annual observations for Miami and the corresponding synthetic counterfactual between 1973 and 1991. Each specification includes vectors of city and year bins dummies. Each period dummies extends for three years. The bin for 1979 is excluded so as to standardize the value of that interaction to 0. The interaction coefficients between a dummy variable for Miami and a corresponding year bin are reported. The method of estimation is FGLS with AR1 process for the error term assumed * p<0.05; ** p<0.01; *** p<0.001.

Figure 6: Log Wage Measures of High School Dropouts, March-CPS

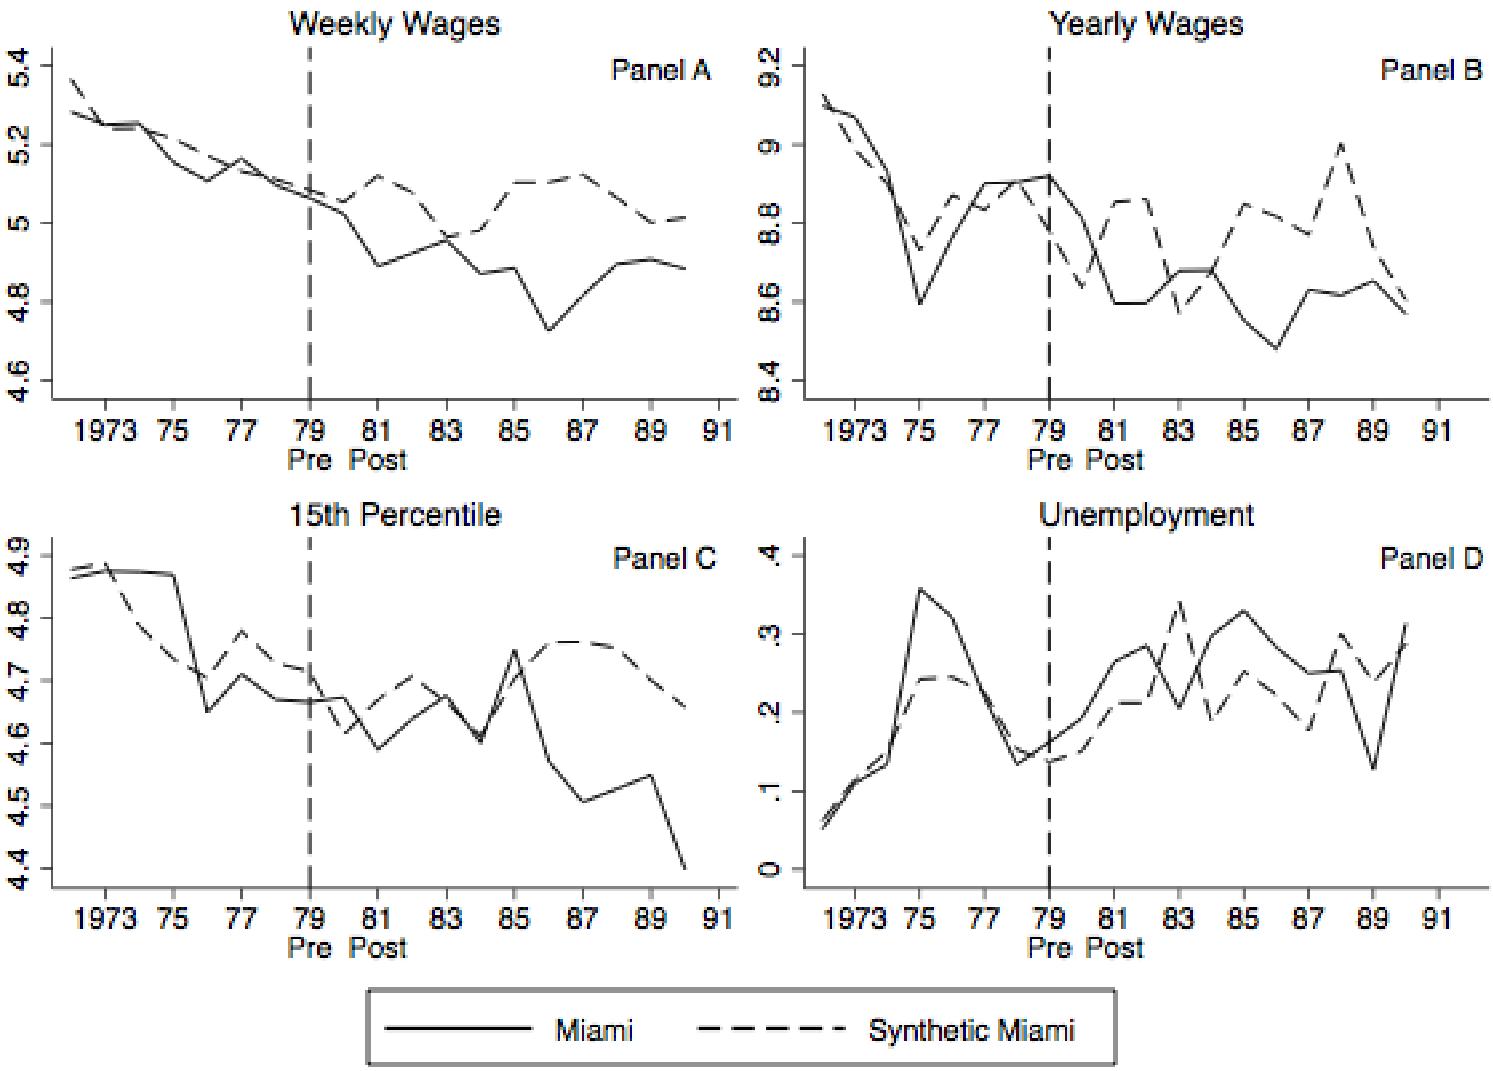

Notes: The data source is March CPS. Each Panel shows the outcome variable for Miami (solid line) and Synthetic control (dashed line) in the period 1972-1991. The outcome shown and the sample used are noted in the title of each panel (A through D). Preferred sample means: non-Cubans, high school dropouts, not self-employed, in the labor force, age 19-65. The cities with positive weight in the synthetic control are as follows. Panel A: Cincinnati, OH 37.1%, Atlanta, GA 32.2%, New Orleans, LA 18.5%, Tampa-St Petersburg, FL 12%; Panel B: Los Angeles, CA 80.1%, Dallas, TX 14.4%, New York City, NY 5.5%; Panel C: New York City, NY 51.8%, San Diego, CA 38.2%, Riverside-San Bernadino, CA 10%; Panel D: New Orleans, LA 39.3%, Denver-Boulder-Longmont, CO 30.4%, Tampa-St. Petersburg, FL 30.3%.

Table 3: Number of Observations for High School Dropouts in Miami

Year	**Our sample selection, March CPS**	**Preferred: Our sample selection, May-ORG**	**Borjas (2015)'s Preferred Sample, March CPS**	**Borjas (2015)'s Preferred Sample, May-ORG**
1973	70	42	30	17
1974	69	32	25	11
1975	66	41	16	13
1976	62	43	22	18
1977	64	39	25	19
1978	61	37	21	11
1979-Pre	62	**145**	17	56
1980-Shock	68	**161**	16	55
1981-Post	72	**145**	18	51
1982	62	**135**	24	39
1983	59	**149**	17	50
1984	55	**145**	15	48
1985	55	**72**	17	26
1986	61	**183**	17	61
1987	66	**221**	17	78
1988	86	**222**	17	72
1989	96	**222**	17	64
1990	74	**250**	4	63
1991	72	**175**	9	41

Notes: Our sample includes individuals with no high school degree, non-Cuban, with positive earnings, not self-employed, in the labor force in the age range 19-65. Borjas' sample includes individuals with no high school degree, with positive earnings, not self-employed, in the labor force, only male, only non-Hispanic and in the age range 25-59. A one year adjustment is made to the March CPS numbers as previous year earnings are reported.

Figure 7: Sensitivity of Borjas (2017)'s Results

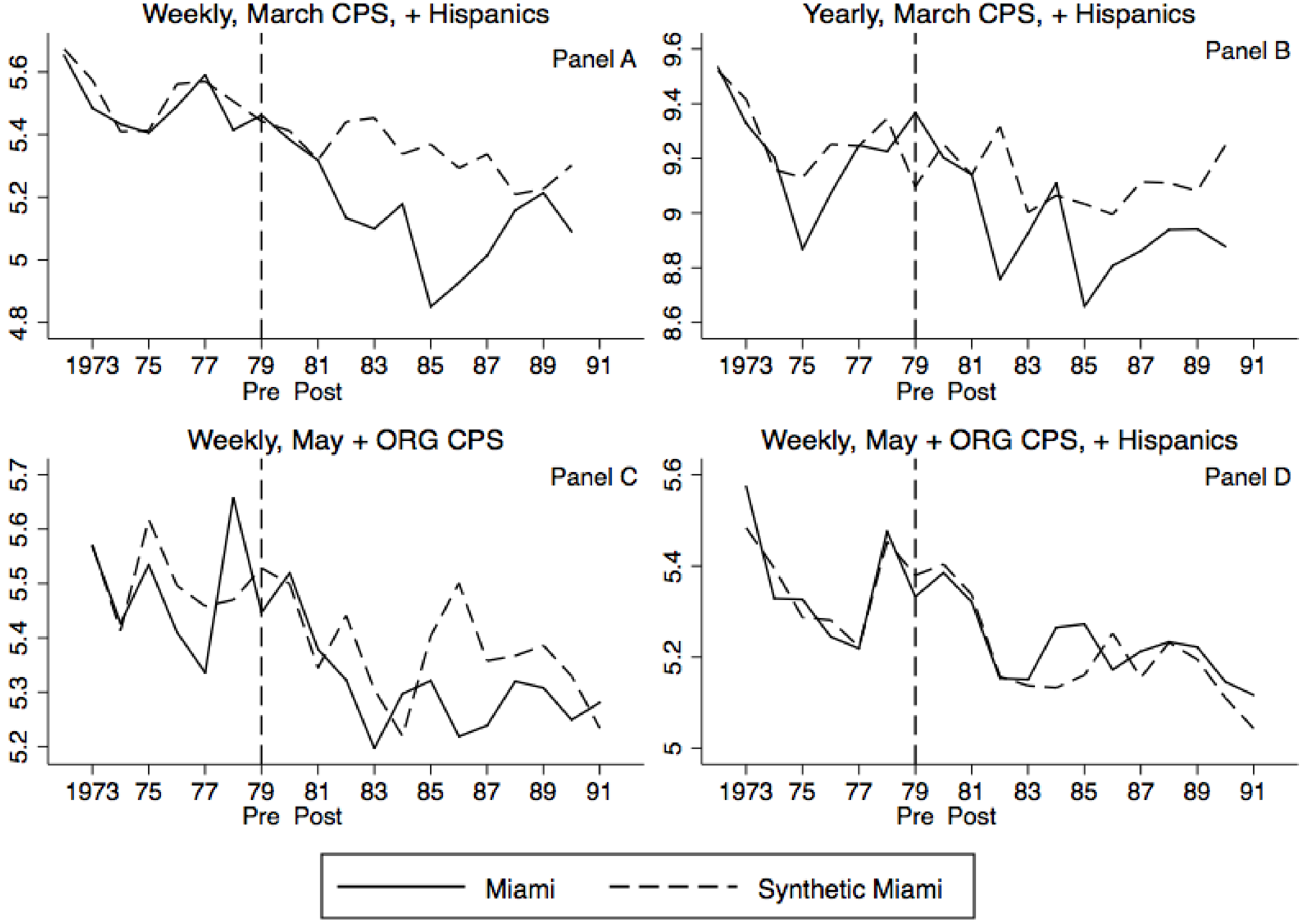

Notes: Panel A reproduces Figure 3A in Borjas (2015) with yearly wages. The logarithm of weekly wages for males, high school dropouts, non-Hispanic age 25-59, smoothed with a 3-years moving average in the March CPS sample for Miami and the "Employment control" are shown. Panel B extends the sample to non-Cuban, males 19-65, drops the smoothing and introduces the synthetic control. Panel C goes back to Borjas original sample, removes the smoothing and uses ORG-CPS data. Panel D uses the same group definition of Panel B, on ORG data. The cities with positive weight in the synthetic control are as follows. Panel A: Dallas, TX 67.8%, Baltimore, MD 18.8%, New York City, NY 13.5%; Panel B: Dallas, TX 60.4%, Atlanta, GA 39.6%; Panel C: Greensboro, NC 64.8%, Tampa-St. Petersburg, FL 35.2%; Panel D: Tampa-St. Petersburg, FL 92.6%, Greensboro, NC 7.4%

Figure 8: The Non-Event of 1994, Log Wage Measures

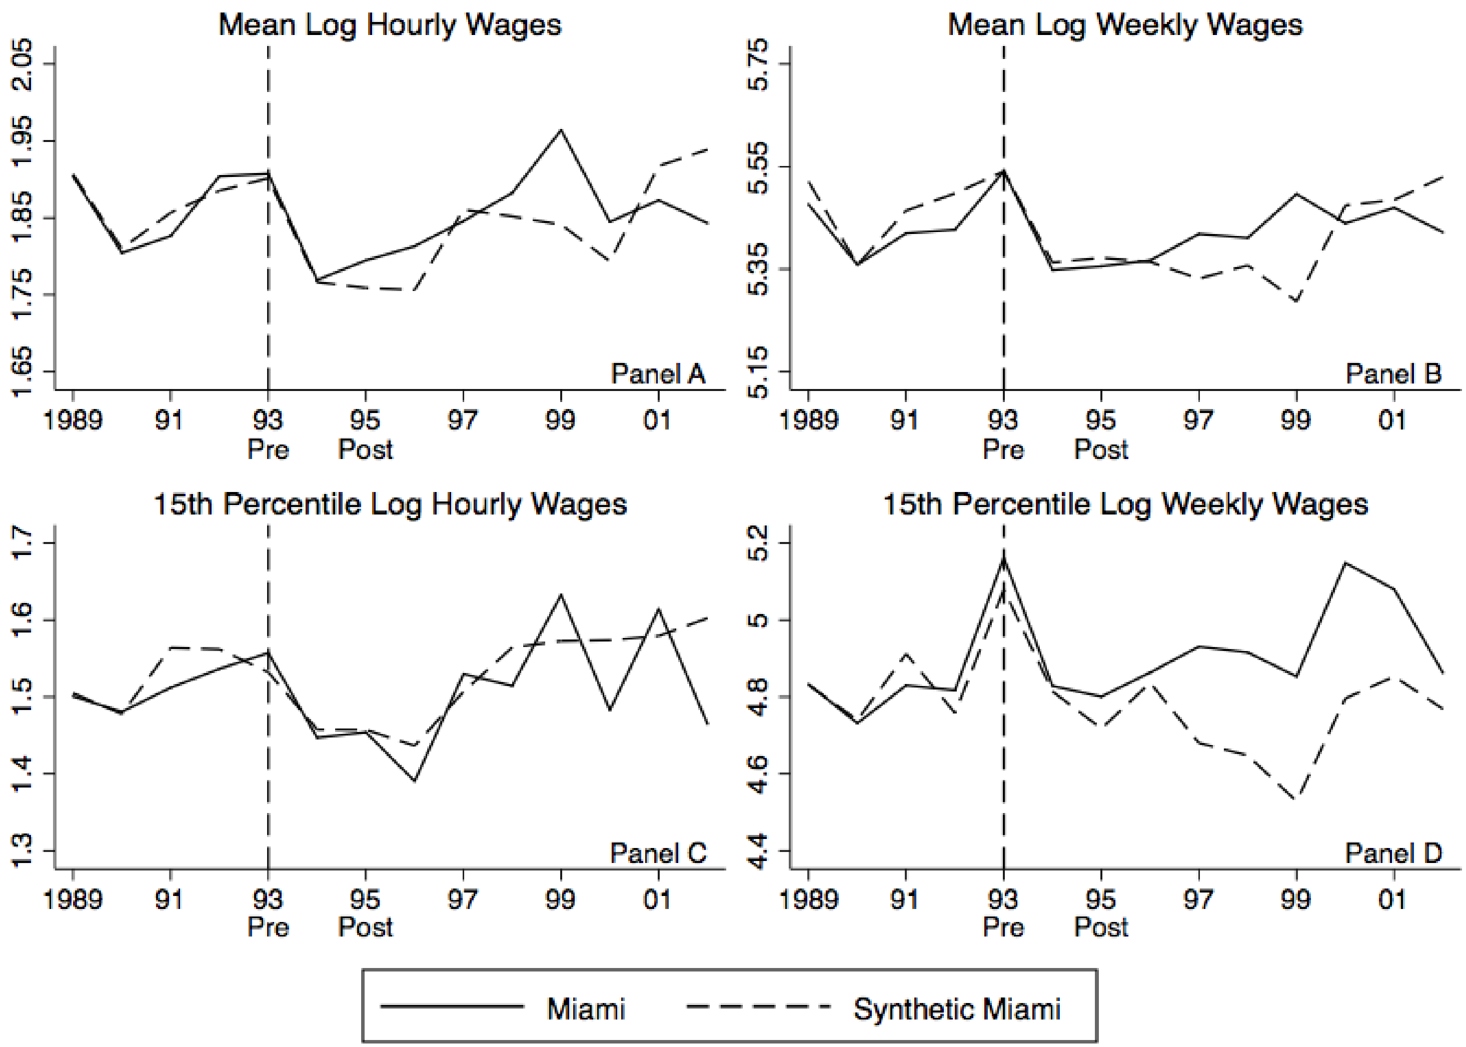

Notes: The data source is ORG CPS. Each Panel shows the outcome variable for Miami (solid line) and for the synthetic control (dashed line) in the period 1989-2001. The variable and sample are noted in the title of each panel. Preferred sample means: non-Cubans, high school dropouts in the labor force, age 19-61. The cities with positive weight in the synthetic control are as follows. Panel A: Bakersfield, CA 35.4%, New Orleans, LA 21.8%, El Paso, TX 16.8%, Jackson, MS 13.1%, Visalia-Tulare-Porterville, CA 12.9%; Panel B: Jersey City, NJ 31.8%, Sioux Falls, SD 30.4%, Bakersfield 27.6%, San Antonio, TX 10.2%; Panel C: Bakersfield, CA 38.4%, Lakeland-Winter Haven, FL 28.6%; New Orleans, LA 17.5%, El Paso, TX 10.9%, Jackson, MS 4.4%; Panel D: Jersey City, NJ 50.5%, Sioux Falls, SD 32.1%, San Antonio, TX 17.4%.