The Labour Market Effects of a Refugee Wave

A Replication Study of Peri and Yasenov (*Journal of Human Resources*, 2019)

Michael Christl*

International Journal for Re-Views in Empirical Economics, Volume 4, 2020-4, DOI: 10.18718/81781.19

JEL: J3, J61
Keywords: Immigration, Wages, Mariel boatlift, Synthetic control method, Measurement error, Replication study

Data Availability: The original data sets and codes of the replicated study are available from the author's website: giovanniperi.ucdavis.edu/data-and-codes. The Stata programs as well as the analysis data sets for this replication study can be downloaded at IREE's data archive (DOI: 10.15456/iree.2020136.181539).

Please Cite As: Christl, Michael (2020). The Labour Market Effects of a Refugee Wave: A Replication Study of Peri and Yasenov (2019). *International Journal for Re-Views in Empirical Economics*, Vol 4(2020-4). DOI: 10.18718/81781.19

Abstract

Peri and Yasenov (2019) exploit a natural experiment, the Mariel boatlift of 1980, to analyse the impact of immigration on wages and other labour market outcomes of natives. The authors find no impact of this (immigrant) labour supply shock on the wages of local workers. These results are heavily discussed in the literature, making it a good example for replication to check the robustness of the findings. This paper analyses the impact of the selection of control variables when choosing the synthetic control group and the influence of the sample choice on the findings of the original study. The replication exercise shows that the original results are very stable to several robustness checks. Even though restricting the analysis to a smaller sample influences the results slightly, there is still no evidence of a significant (negative) impact of the labour supply shock on wages of locals.

*European Commission JRC, Seville, Spain. Email: michael.christl@ec.europa.eu

Received May 15, 2020; Revised October 26; Accepted November 9, 2020; Published December 18, 2020.

1 Introduction

There has been an ongoing academic discussion in recent years about the effects of immigration on the wages of (local) workers, especially with respect to the Mariel boatlift. It is not easy to identify causal effects on the labour market of a migration inflow. Several attempts have been made to analyse those effects. One of the most prominent examples is the Mariel boatlift. It is of special interest because it provides a sudden, unexpected immigration wave that caused what was (to a large extend) a clean labour supply shock. This shock has been used as a natural experiment in several studies to determine the impact of labour supply shocks.

In April 1980, Fidel Castro announced the opening of the port of Mariel for anyone who wanted to leave the country. So, between April and September, over 125.000 Cubans fled to the US, and most of them came to the nearest big city, Miami. According to Peri and Yasenov (2019), this led to an 8% increase in the labour force, indicating a huge supply shock on the labour market of Miami. In addition, most of the immigrants where low skilled, implying that the increase of the low-skilled labour force was even stronger, namely 18 percent.

Several studies analysed this sudden inflow of foreigners on the labour market of Miami. Card (1990) found no significant impact from the labour supply shock. This was in strong contrast to the theoretical prediction of the canonical model of labour demand and labour supply. This contradiction with theoretical models led to a lot of related research about the reasons for those findings. For example, Bodvarsson, Lewer and van den Berg (2008) argue that the immigrant workers increased consumption in the region, which also caused an increase in labour demand, which led to a new equilibrium that could lead to similar wage levels.

The result of Card (1990) was undisputed for over 25 years, but with strong improvements in the econometric toolbox and some issues regarding the methods used in the original paper, the Mariel boatlift has recently enjoyed a revival. Borjas (2017) argue that one key lesson from the migration literature is that the 'effect of immigration on the wage structure depends crucially on the differences between the skill distributions of immigrants and [those of the] natives'. Therefore, the effect of immigration should be visible only for those that have a similar skill distribution as the immigrants. In his paper, he shows that wages of low-skilled groups (in his case, high school drop outs) fell dramatically, by about 10 to 30 percent, depending on the control placebo he chooses.

Later on, Peri and Yasenov (2019) used the synthetic control method to choose a control group of other areas that showed similar labour market characteristics in the time before the migration inflow. Using this method, the authors find no wage or unemployment impact on natives after the arrival of the immigrants to Miami.

The discussion about the impact of the Mariel boatlift on the labour market of Miami is still ongoing[1]. This makes it an interesting topic for a replication exercise. As already argued by several authors, using a placebo or synthetic control group can influence results substantially. In this short note, I replicate the main results of Peri and Yasenov (2019) using the synthetic control method. I also conduct some additional robustness checks to get a feeling for the validity and reliability of the method. I focus on three aspects: first, the chosen control group and the impact on the robustness

[1] See, e.g. Clemens and Hunt (2017), Borjas (2017), Peri and Yasenov (2019), ...

of the results; second, the robustness of the placebo tests, and third, the robustness regarding the chosen sample.

2 Data

The authors use several data sets but for the main data source, they use a combination of the current population survey (CPS) from May for the years 1973 through 1978 and merge that data set with the outgoing rotational group data (ORG) from the same survey from 1979 onwards[2]. The authors focus on the sample of non-Cuban high-school drop-outs between 19 and 65 years of age who are working in Miami because this should be the group of workers that is most likely influenced by the supply shock.

The impact of the supply shock is about 6 percentage points when considering the entire population. When looking only at high-school drop-outs, the share of Cubans increases by even more (12 percentage points). The jump of the share from 1979 to 1981 makes this event a good one for using the synthetic control method. In addition, the supply shock shows a temporary character. By 1985, the share of Cubans is back to the level of the pre-shock value for the share of Cubans in the total population and among high school drop-outs.

3 Methodology

After a descriptive analysis of the shock, the authors use the synthetic control method (SCM). This method goes back to Abadie and Gardeazabal (2003), and it was later further developed by Abadie et al. (2010). This method provides a systematic way to analyse the impact of certain events. Typically, this method is used when a certain event or treatment is restricted to a single unit, such as a region, while this event was not experienced in other regions. This method identifies a synthetic control group, a linear combination of units that had not experienced the event but that show a similar trend in the pre-treatment period. This allows us to identify treatment effects related to the event analysed.

Formally speaking, this approach considers several groups $J+1$, which are indexed by $j = 0, ..., J$. Group 0 refers to the group in which the event happens while the rest of the groups are defined as the so-called donor pool. Let G_0 be a vector with k elements that are equal to the number of variables used to predict our variable of interest. We define G_J as a kxJ-matrix in which each row represents the sequence of the same variables and years relative to city j in the donor pool.

The synthetic control method searches for a vector of weights W^* that constructs a convex combination of variables in cities in the donor pool G_J that has the lowest quadratic error compared with the pre-treatment vector G_0. Formally this means:

$$W^* = argmin(G_0 - G_J W)'V(G_0 - G_J W) \quad s.t. \quad \sum w_j = 1, w_j \geq 0 \tag{1}$$

[2]My special thanks are extended to Giovanni Peri and Vasil Yasenov for making these data freely available on their personal website.

Once we receive the weights W^*, we can identify post-treatment outcome variables for this synthetic control group and the treated group. This is the intuitive idea of the SCM. The composition of the synthetic control group is essential for the results. Therefore, validating the control group is very important when applying this method.

The SCM allows for a nice descriptive way of validation. Pre-treatment levels of the outcome variable of the control group and the treated group are usually compared. If they match closely in the years before the treatment, this is taken as a sign that the choice of the synthetic control group is reasonable.

Peri and Yasenov (2019) use the outcome variable itself, the share of low-skilled workers, the share of Hispanics and the share of manufacturing workers in the labour force to minimise the pre-treatment distance between the treated group and the synthetic control group. They analyse the period from 1973 until 1979 to see whether there are substantial differences.

Figure 1: Main results: Miami vs. Synthetic control group

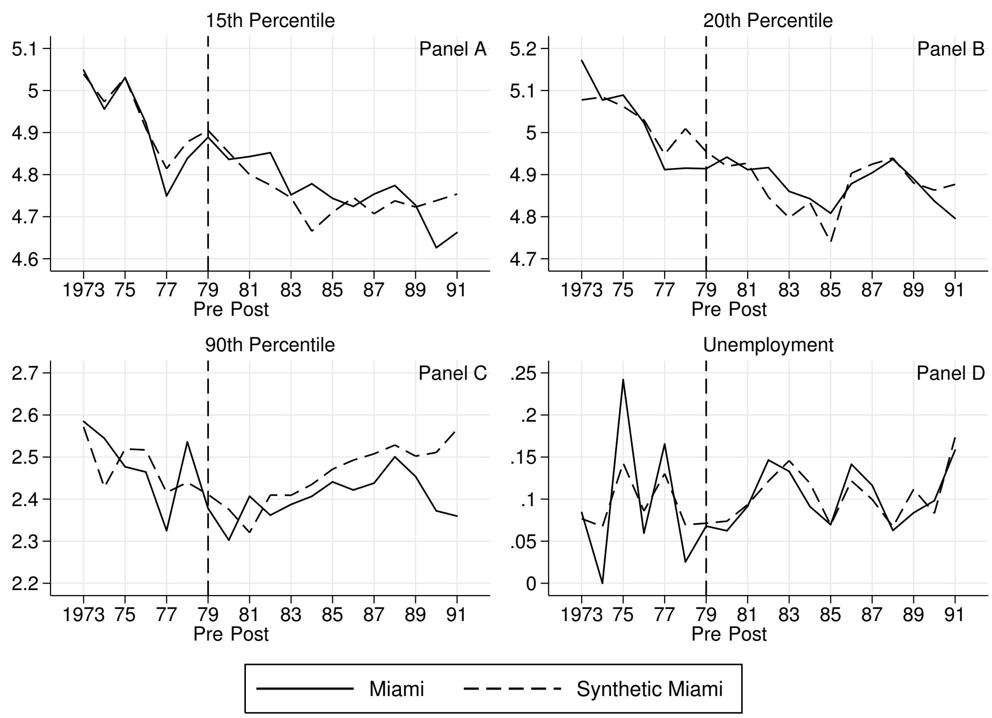

Source: Peri and Yasenov (2019)

From a theoretical point of view, we would expect an effect of the labour supply shock because of the strong inflow of Cubans on the labour market immediately after their arrival in 1980. As we have seen, the shock did not persist, meaning that the share of Cubans on the labour market in Miami reached pre-shock levels by 1985. We would expect the adjustment mechanisms on the labour market to take place soon after the arrival of the Cubans.

Figure 1 highlights the main result of the synthetic control method regarding wages and unemployment. It shows a very good fit in the pre-treatment period, especially for the low end of the wage distribution (15th percentile), on which this study focuses. However, the synthetic control group also mimics quite well the pre-treatment behaviour for the other outcome variables, such as the 20th percentile, the 90th percentile and the unemployment rate in Miami.

In both measures for the lower income distribution of the natives' hourly wages (Panels A and B), the control group seems to match Miami quite well in the pre-treatment period. But we do not really observe differences in the years after the treatment (between 1980 and 1982). The same holds true for the upper part of the income distribution (Panel C). Even though, from a theoretical point of view, we might expect a wage increase for this group of workers because of complementary forces, we do not see any effect on their wages. Likewise, the unemployment rate does not show any significant difference in comparison to the synthetic control group after the treatment occurs.

Taking a closer look on the synthetic control group (see Table 1 in the Appendix), we can see that for the wage variables (Panels A, B and C), the control group typically contains Nassau-Suffolk and Birmingham, mixed with Tampa-St. Petersburg or San Diego or Rochester. For the unemployment rate, a combination of New York City, New Orleans, Albany-Schenectady-Troy and Cincinnati was the one mimicking the unemployment rate best in the pre-treatment period. We can already see that typically only 3 to 4 of the regions are used to build the control group.

Strictly speaking, the wage measures used in the analysis so far might not measure the price of labour in a correct way, since they do not account for changes in working time. Therefore, Peri and Yasenov (2019) use hourly and weekly wages and regression-adjusted wage measures. Those also account for differences in certain characteristics (e.g. age, gender and ethnic group). The predicted measures typically show higher volatility, as it is also visible in Figure 2. Again, no significant difference can be found between Miami and the synthetic control group after the labour supply shock caused by the Mariel boatlift in 1980.

In the next step, the authors analyse the impact of the boatlift on several sub-samples, even though they warn that this involves quite small sample sizes. Again the authors do not find any negative break or jump corresponding to the labour supply shock caused by the Mariel boatlift, neither for males nor for females. In addition, they do not see any significant deviations in wages between any ethnic group and the control group.

A criticism of the early study by Card (1990) was that they were neglecting city-specific shocks. Peri and Yasenov (2019) try to overcome this problem within the SCM framework. A common problem with the SCM is the construction of confidence intervals. Therefore, the authors decided to use a regression framework to calculate statistical significance levels. As already expected by the graphical results, the differences in the pre-treatment period between Miami and the synthetic control group are statistically insignificant for all labour market outcomes used (weekly wages, hourly wages, 15th percentile, 20th percentile, 90th percentile and the unemployment rate).

Figure 2: Different wage measures for high school drop-outs

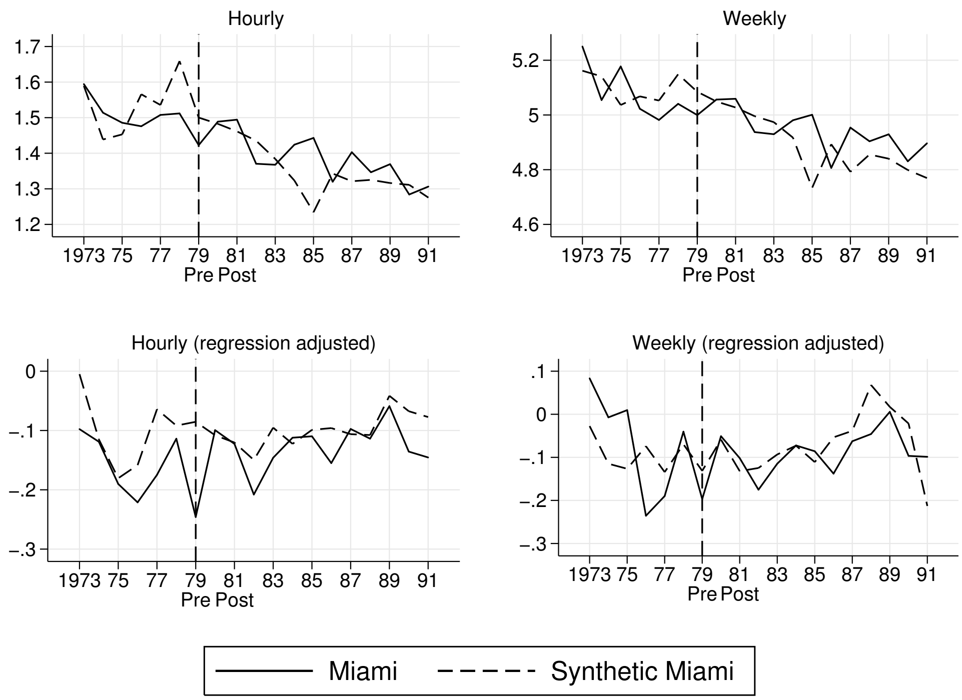

Source: Peri and Yasenov (2019)

None of the coefficients is significantly different from zero after the Mariel boatlift. Moreover, the coefficients on wage outcomes have a positive coefficient, indicating that wages rose slightly more in Miami than for the synthetic control group. This suggests that the supply shock had no significant effect on the analysed labour market variables (wages and unemployment) of the low skilled local workers.

4 Robustness of the results

In this section, I discuss some additional robustness checks in the synthetic control method. It is worthwhile to mention that Peri and Yasenov (2019) performed a lot of robustness checks in their original study. I still thought three additional checks could be of interest: First, the control variables used to construct the synthetic control group, second, changes in the so-called 'placebo tests' and, third, changes in the sample that was used for the analysis.

4.1 Choice of control variables

A common criticism of cross-country or cross-region regression is the fact that changes in the variables used can change the estimates substantially. This also holds true for choosing a synthetic control group. As argued by e.g. McClelland and Gault (2017), the synthetic control group can be very sensitive to the choice of control variables.[3]. We follow the argument of Dhungana (2011), who argues 'given that there is not one guideline on the choice of variables, we attempt to test whether the synthetic controls procedure we are using is more robust to similar criticism'. Obviously, and as argued by Kreif et al. (2016),[4] the choice of the variables must be justified.

Peri and Yasenov (2019) use different control variables when using the SCM for different labour market outcomes. Since there was no clear explanation for doing that, I replicate the exercise using the same control variables for all the estimates. In addition, I add new control variables to the estimations, to control for additional factors, such as the share of male workers and the average age of the workforce. All these factors could influence labour market outcomes and they might be different for the control group and for Miami.

The results are shown in Figure 3 and 4. Both figures show that the selection of control variables indeed influences the fit of the model in the pre-treatment period. However, when we change these control variables, there are still no strong deviations in the labour market outcomes for Miami and the synthetic control group after the treatment Moreover the regression results do not hint at any significant deviation between Miami and the control group after the labour supply shock caused by the Mariel boatlift.

4.2 Inference using permutations

As mentioned by the authors, the regression approach to check for significant deviations between Miami and the synthetic control group has its caveats. The small sample size of only 38 observations raises doubts about the credibility of such an approach. Therefore, the authors also try to construct a measure for statistical inference based on permutations suggested by Abadie et al. (2010).

The synthetic control method allows us to conduct a falsification exercise, which Abadie et al. (2010) call 'placebo studies'. This provides an alternative mode of inference that is based on the premise that 'our confidence that a particular synthetic control estimate reflects the impact of the

[3]They argue that 'Selecting regions from the donor pool can be sensitive to the choice of variables used to match the donor regions to the treated region'.

[4]They argue that the choice of control variables must be justified using substantive knowledge of the outcome process.

Figure 3: Robustness I: Same control variables in all estimations

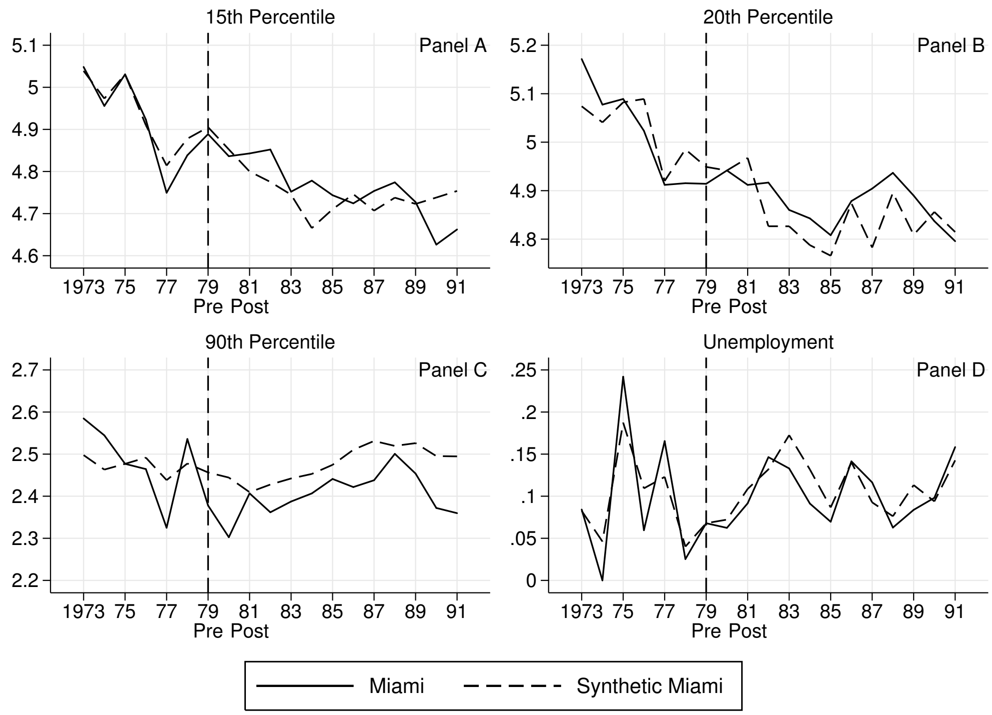

intervention under scrutiny would be severely undermined if we obtained estimated effects of similar or even greater magnitudes in cases where the intervention did not take place. Suppose, for example, that the synthetic control method estimates a sizeable effect for a certain intervention of interest. Our confidence about the validity of this result would dissipate if the synthetic control method also estimated large effects when applied to dates when the intervention did not occur'.

The idea of that approach is to simulate a distribution of deviation of all cities in the donor pool from their synthetic control group. In a second step, the deviation of Miami and its synthetic control after the treatment are constructed and compared to the deviations in other cities.

While the original study validates only the weekly and hourly wages, the 15th percentile and the unemployment rate, I also performed permutations for the 20th and 90th percentiles. In general, Figure 5 shows that Miami is an average city in the pre-treatment period regarding the deviation from the control group. We can also see a positive deviation after the treatment for log weekly and log monthly wages. Moreover, for the 15th percentile, there seems to be a negative deviation after the treatment. In general, the graphs show that the idiosyncratic variation can be substantially large, spanning a noise of up to 20%. This makes the identification of effects quite hard.

Figure 4: Robustness II: Adding control variables

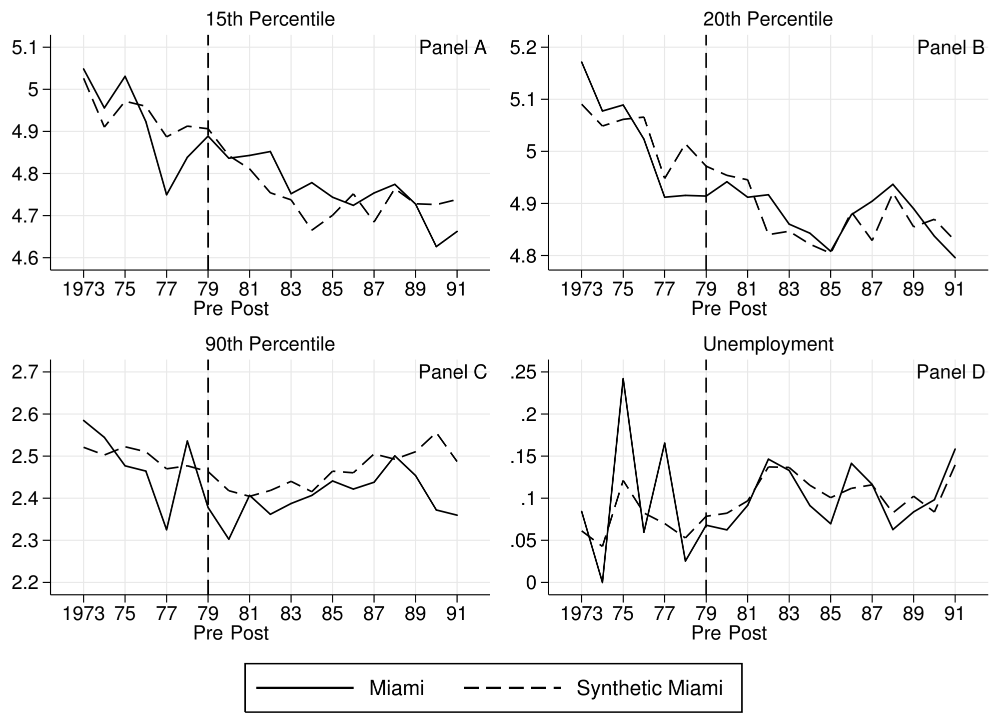

In a second test, I added some control variables (the share of males in the labour force and the age) to find the corresponding synthetic control group. There are of course changes in the outcome, as Figure 6 shows. But again, the additional tests show that also when using additional control variables, Miami remains an average city in the pre- and post-treatment period regarding the deviation from the control group.

4.3 Changing the sample

There is an ongoing discussion about the validity of the results of Peri and Yasenov (2019). Borjas (2017) defended his results by arguing that Peri and Yasenov (2019) do not consider any specific information on the labour market in their analysis and they use a sample of people that might hide the real influence on wages of low-skilled local workers. He argues e.g. that: 'It is tempting to increase sample size by including working women in the study, but female labour force participation was increasing very rapidly in the 1980s, so that wage trends are likely to be affected by the selection that marks women's entry into the labor market.

To conclude the examination of robustness, I try to replicate the results of Peri and Yasenov (2019) concentrating only on males between 25 and 59 in the synthetic control method analysis. In addition, as suggested by Borjas (2017), not only do I exclude Cubans from the sample, but also

Figure 5: Robustness III: Inference using permutations

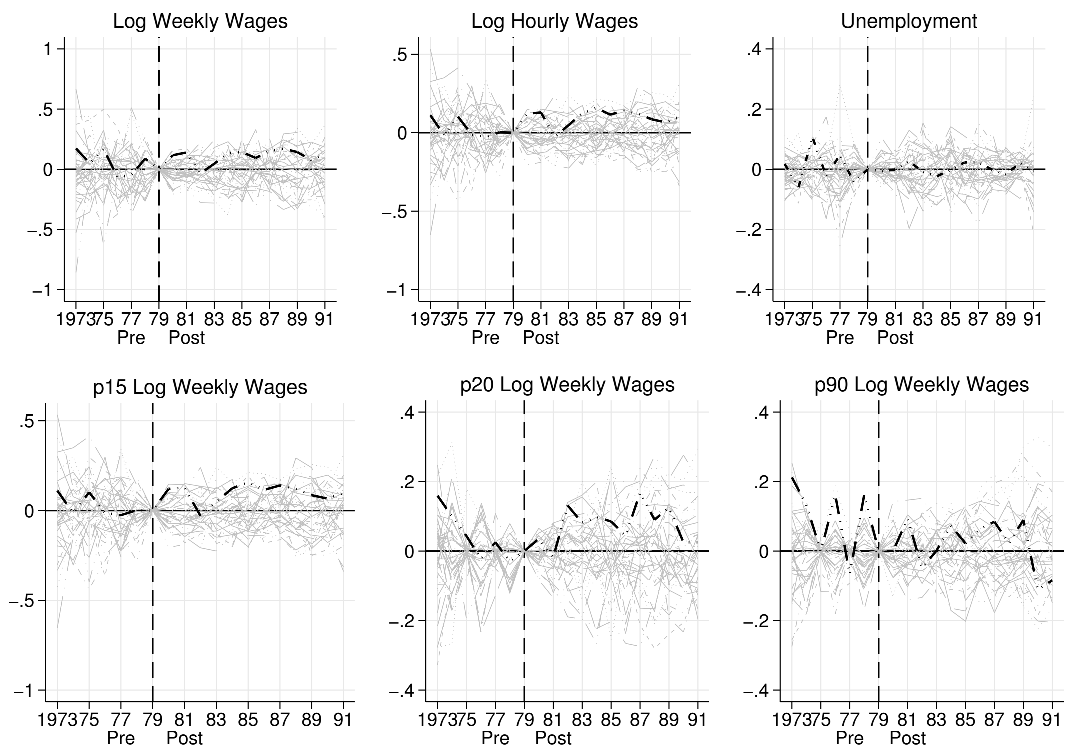

Hispanics. The author argues that the strong increase of Hispanic immigration after 1980 could influence the wage trend by systematically changing the composition of the sample each year, but it does not say anything about the impact of the Mariel boatlift.

By changing the sample, we can see some differences in the outcomes for our labour market variables of interest, as highlighted by Figure 7. Especially for the 15th and 20th percentiles, there are now differences between Miami and the synthetic control group. However, we also see that the fit in the pre-treatment period is poor. This indicates that using the synthetic control group does not offer a good fit, and the validity of the results might be questionable. In addition, these results show a slightly different behaviour in the labour market variables than in the sample used before. We can see that wages (15th and 20th percentile) seem to stay on a higher level after the Mariel boatlift in the synthetic control group, while they fall in Miami. Still, it is important to note that this effect vanishes after some years and that in general it is not very pronounced, as highlighted in Figure 7.[5] Applying the standard regression analysis as in Peri and Yasenov (2019) also shows that those negative effects on wages are not statistically significant, highlighting again that the findings of Peri and Yasenov (2019) are strongly robust, also across changes in the sample.

[5] Results for females are shown in Figure 8 in the Appendix.

Figure 6: Robustness IV: Inference using permutations

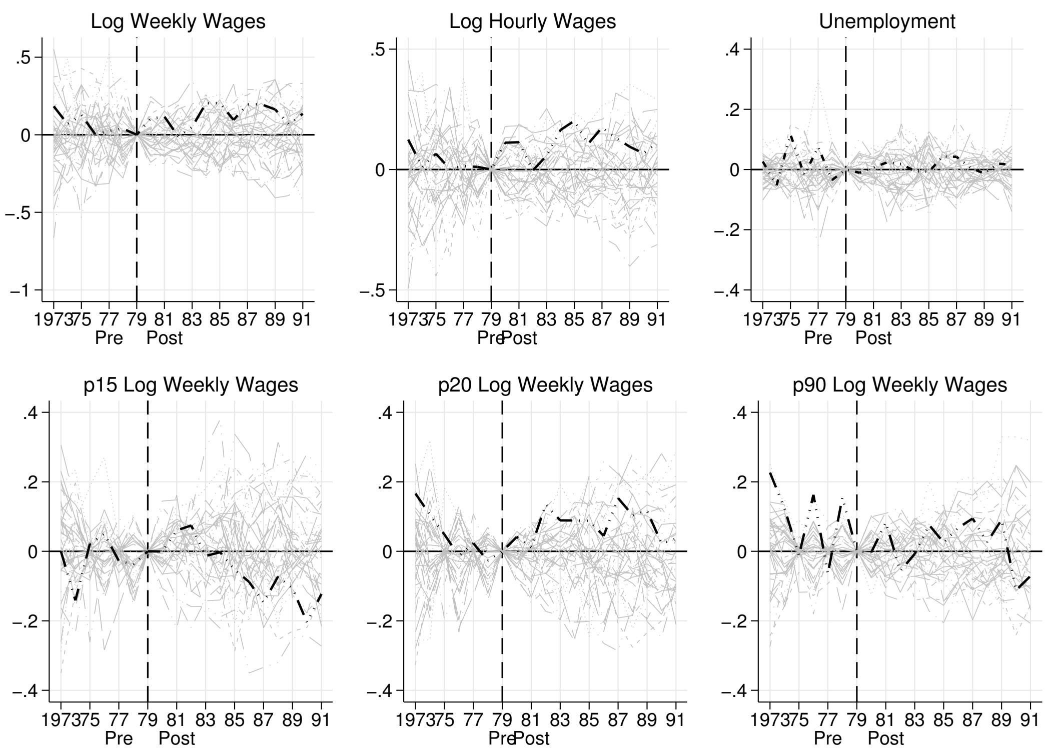

5 Conclusion

In this study I first replicate the main results of Peri and Yasenov (2019), who use a synthetic control method to study the impact of the supply shock. In a second step, I analyse the robustness of the results.

First, this study suggests that the results are very robust regarding additional or different control variables when choosing the control group. From a theoretical point of view, even though it could be reasonable to add age and the share of males as controls, the results indicate that using these additional controls does not change the main results. In addition, I added the same variables and checked the statistical inference of the results by using a placebo test. I find that additional controls do not change the statistical inference of the results.

Second, since there is an ongoing discussion about the sample of local workers used for the analysis, I tried to replicate the results for a sample that was suggested by Borjas (2017). There, we analyse only the effects on non-Hispanic males between 25 and 59 years old. The results indicate that there could be negative wage effects for the selected group compared to the synthetic control

Figure 7: Robustness: Sample adjustment (no Hispanics, only males)

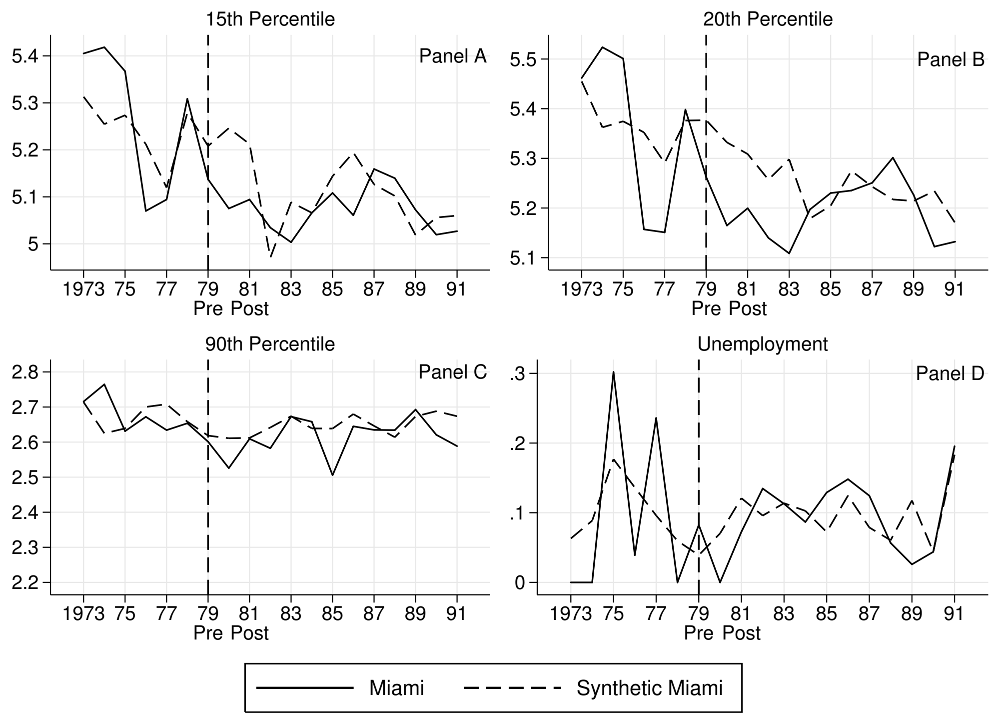

group, but these effects vanish after some years and they are not very pronounced. This raises questions about their statistical significance in general.

Overall, the results seem to be quite robust for the sample chosen by Peri and Yasenov (2019). However, looking more closely at a small group of the labour market (males between 25 and 59, who could be the ones that were the most affected by the supply shock), the results suggest that there could be a negative effect on wages for a (quite specific) group in the labour market. In addition, we can see substantial differences between wage developments of females and males, an additional aspect that has not been discussed. However, caution must be taken because the data set used in the analysis is quite small.

References

Abadie, Alberto and Javier Gardeazabal (2003). "The economic costs of conflict: A case study of the basque country." *American Economic Review* 93(1): 113–132.
DOI: 10.1257/000282803321455188.

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. DOI: 10.1198/jasa.2009.ap08746.

Bodvarsson, Örn B., Hendrik F. Van den Berg and Joshua J. Lewer (2008). "Measuring immigration's effects on labor demand: A reexamination of the mariel boatlift." *Labour Economics* 15(4): 560–574. URL: 10.1016/j.labeco.2008.04.003.

Borjas, George J. (2017). "The wage impact of the marielitos: A reappraisal." *ILR Review* 70(5): 1077–1110. URL: 10.1177/0019793917692945.

Card, David (1990). "The impact of the mariel boatlift on the miami labor market." *ILR Review* 43(3). URL: 10.1177/001979399004300205.

Clemens, Michael A. and Jennifer Hunt (2017). "The labor market effects of refugee waves: reconciling conflicting results." *ILR Review* 72(4) 818–857. DOI: 10.1177/0019793918824597.

Dhungana, Sandesh (2011). "Identifying and evaluating large scale policy interventions: what questions can we answer?" *The World Bank* Policy Research Working Paper 5918. URL: 10.1596/1813-9450-5918.

Kreif, Noémi, Richard Grieve, Dominik Hangartner, Alex J. Turner, Silviya Nikolova and Matt Sutton (2016). "Examination of the synthetic control method for evaluating health policies with multiple treated units." *Health Economics* 25(12): 1514–1528. URL: 10.1002/hec.3258.

McClelland, Robert and Sarah Gault (2017). "The synthetic control method as a tool to understand state policy." Washington DC: The Urban Institute.

Peri, Giovanni and Vasil Yasenov (2019). "The labor market effects of a refugee wave synthetic control method meets the mariel boatlift." *Journal of Human Resources* 54(2): 267–309. DOI: 10.3368/jhr.54.2.0217.8561R1.

Appendix

Figure 8: Robustness: Sample adjustment (no Hispanics, only females)

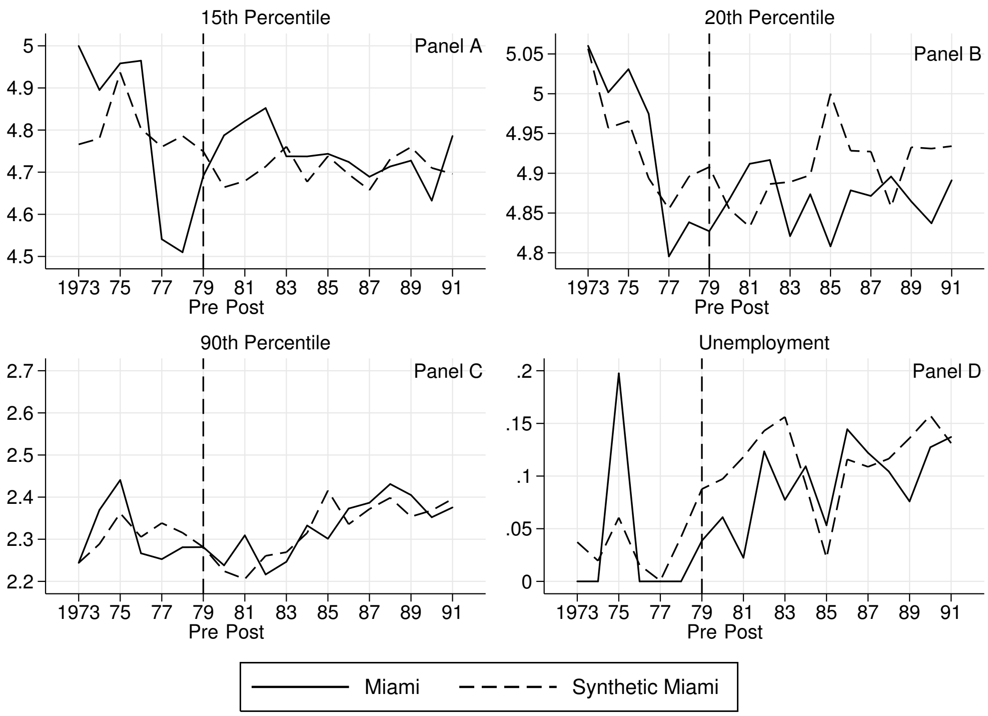

Table 1: Weights used to generate the synthetic control group in percent

	15th	20th	90th	uenmployment
New York City	0.0	0.0	0.0	30.9
Los Angeles	0.0	0.0	0.0	0.0
Chicago	0.0	0.0	0.0	0.0
Philadelphia	0.0	0.0	0.0	0.0
Detroit	0.0	0.0	0.0	0.0
San Francisco-Oakland	0.0	0.0	0.0	0.0
Washington	0.0	0.0	0.0	0.0
Boston	0.0	0.0	0.0	0.0
Nassau-Suffolk.	10.4	13.8	32.7	0.0
Pittsburgh	0.0	0.0	0.0	0.0
St Louis	0.0	0.0	0.0	0.0
Baltimore	0.0	0.0	0.0	0.0
Cleveland	0.0	0.0	0.0	0.0
Houston	0.0	0.0	0.0	0.0
Newark	0.0	0.0	0.0	0.0
Minneapolis-St Paul	0.0	0.0	0.0	0.0
Dallas	0.0	0.0	0.0	0.0
Seattle-Everett	0.0	0.0	0.0	0.0
Anaheim-Santa Anna-Garden Grove	0.0	0.0	0.0	0.0
Milwaukee	0.0	0.0	0.0	0.0
Atlanta	0.0	0.0	0.0	0.0
Cincinnati	0.0	0.0	0.0	1.1
Patterson-Clifton-Passaic	0.0	0.0	0.0	0.0
San Diego	0.0	57.7	0.0	0.0
Buffalo	0.4	0.0	0.0	0.0
Kansas City	0.0	0.0	0.0	0.0
Denver	0.0	0.0	0.0	0.0
San Bernardino-Riverside-Ontario	0.0	0.0	0.0	0.0
Indianapolis	0.0	0.0	0.0	0.0
San Jose	0.0	0.0	0.0	0.0
New Orleans	0.0	0.0	0.0	48.4
Tampa-St Petersburg	0.0	0.0	67.3	0.0
Portland	0.0	0.0	0.0	0.0
Columbus	0.0	0.0	0.0	0.0
Rochester	28.6	0.0	0.0	0.0
Sacramento	0.0	0.0	0.0	0.0
Fort Worth	0.0	0.0	0.0	0.0
Birmingham. AL	60.6	28.4	0.0	0.0
Albany-Schenectady-Troy	0.0	0.0	0.0	19.5
Norfolk-Portsmouth. VA	0.0	0.0	0.0	0.0
Akron. OH	0.0	0.0	0.0	0.0
East Chicago	0.0	0.0	0.0	0.0
Greensboro	0.0	0.0	0.0	0.0