INSTITUTE for REPLICATION

No. 4
I4R DISCUSSION PAPER SERIES

Replication of “How Much Does Immigration Boost Innovation?”

Taylor Wright

October 2022

I4R Discussion Paper Series

I4R DP No. 4

Replication of "How Much Does Immigration Boost Innovation?"

Taylor Wright

Brock University, St. Katherines, ON/Canada

October 2022

E-Mail: joerg.peters@rwi-essen.de
RWI – Leibniz Institute for Economic Research
Hohenzollernstraße 1-3
45128 Essen/Germany
www.i4replication.org

ISSN: 2752-1931

Replication of "How Much Does Immigration Boost Innovation?"

Taylor Wright*
Brock University

October 12, 2022

Abstract

A common approach to identifying the causal impact of immigration on outcomes involves using a "shift-share" or Bartik instrument exploiting country-specific immigration inflows (shifts) and location specific prior shares for the same countries. New econometric findings suggest this instrumental variables approach uses identifying variation not from the shifts, as previously believed, but rather from the shares and suggest a battery of checks to explore the sensitivity of estimates. In this note, I first replicate Hunt and Gauthier-Loiselle (2010) which estimates the effects of immigration on innovation via patenting, and second deploy these new checks from the econometric literature on shift-share instruments. I find that the results of Hunt and Gauthier-Loiselle (2010) (skilled immigration increases innovation and has positive spillovers on the innovation of others) replicate and hold up well to these new tests.

1 Introduction

There is an extensive literature analyzing the impacts of immigration not only on the outcomes of immigrants themselves but also in analyzing any potential spillover effects on the native born population. This literature has spent extensive time on labor market impacts (e.g. Borjas et al. (1997), Card (2001) and Edo (2019) for a survey) but has also extended to the role of immigration on innovation (e.g. Blit et al. (2018), Bound et al. (2017), Peri et al. (2016)). The study of the impacts of immigration on innovation are also important because many countries, including the United States make decisions not only about the overall target of immigrants but also about which types of immigrants they would like to welcome. Immigrants who foster innovation are not only valuable in and of their own research, inventions, and discoveries but they can also provide positive spillovers for those already living in the host country. These contributions to innovation could push economic output up and contribute to economic growth Hunt and Gauthier-Loiselle (2010) examine the impacts of immigration on innovation which they measure through US patenting behaviour.

The authors document that immigrants in the United States tend to pursue backgrounds in STEM fields while the native born population are less likely to pursue these careers. They examine the impacts of patenting behaviour of the immigrant and native

*I am grateful for financial support from Open Philanthropy.

born populations split by education status: college graduates, post-college graduates, and graduates in science and engineering.

When examining the impacts of immigration, concerns about endogenity of immigration location decisions has lead researchers to search for a quasi-experimental approach. Building on early work analyzing immigration and labour market outcomes (e.g. Altonji and Card (1991), Card (2001)) much of the literature has used shift-share or Bartik instruments. These instrumental variables combine aggregate flows of immigrants (the "shifts") with previous periods' distribution of immigrants across space (the "shares"). The intuition behind these instruments is that the due to the observation that immigrants from the same country tend to co-locate, lagged shares will have predictive value of immigration flows while having the benefit of being divorced from current economic phenomenon that may influence immigration location decisions. The implication is that identification comes from the shift portion of the instrument.

Hunt and Gauthier-Loiselle (2010) use this type of instrument with state-level data finding that for immigrant college graduates and 1% increase in the share increases patenting by 12-18% while for post-college graduates the corresponding figure is 19-27% depending on the exact specification. As these numbers are larger that the patenting rates for immigrant college graduates at the individual level, the authors conclude there is evidence of positive spillover effects on innovation from skilled immigration.

Recently, however, there has been work that calls into question this identification assumption (see for example Adao et al. (2019), Borusyak et al. (2020)). Goldsmith-Pinkham et al. (2020) decompose these shift-share instruments into their component pieces and find that they are equivalent to GMM estimation with the shares as instruments. This equivalence results implies that identification is actually coming from the lagged "shares" rather than the inflow "shifts". The authors provide tools to examine the sensitivity of estimates using these shares. Simultaneously, Jaeger et al. (2018) have argued that these instruments are vulnerable to the conflation of short and long run responses to immigration. In particular, they argue that when inflows ("shifts") generate general equilibrium adjustments that take time to dissipate and there is a strong degree of serial correlation in the lagged "shares" and inflow "shifts" may conflate the short-run impacts with the longer term adjustment process to previous inflows.

The goal of this note is twofold: First I examine the reproducibility of the results of Hunt and Gauthier-Loiselle (2010) both by using the supplied scripts and data, and by attempting to reconstruct their instrument myself. Second, I examine the implications of this new Bartik instruments literature for the estimates, conclusions, and robustness of Hunt and Gauthier-Loiselle (2010).

2 Reproducibility

Using the data and `Stata` code made available on Open ICPSR by the authors, I was able to reproduce the tables contained in the manuscript exactly. I have uploaded a condensed version of the code (a single script) for the paper to recreate the tables. Additionally, I will note that estimation using `R 4.2.0` (R Core Team (2022)) and `fixest` (Bergé (2018)) also reproduces the authors OLS and WLS estimates exactly.

The authors data and code provide the final datasets and scripts for analysis but not the raw data or scripts used to construct the final datasets. I attempted to reconstruct the instrumental variable based on the authors' description in section I.B. and re-estimate the main IV results. I was unable to match the values of the authors' instrument exactly but the values of my instrument are quite close to those of the original paper. Table 1 recreates the instrumental variable estimates from Table 7 in Hunt and Gauthier-Loiselle's work using my recreated instrument, while Table 2 does the same for Table 8. In both cases, the first column presents the original estimates while the second presents those using the recreated instrument. The values using my instrument are similar in magnitude (they are now actually slightly larger) and remain statistically significant.

In sum, I consider this a successful exercise in reproducibility as I am able to directly recreate the authors findings and while I cannot recreate the instrument used exactly based on my understanding of their approach, the instrument I do create produces nearly identical estimates that offer no change in conclusions.

Table 1: Recreation of Table 7 (IV only)

	HGL (1)	Constructed (2)
Panel A. Base Specification	30.270***	30.518***
	(7.156)	(7.387)
Panel B. Base Specification without California	26.253***	25.938***
	(7.085)	(6.910)
Panel C. Base specifications without year 2000	18.921*	19.291**
	(7.111)	(7.138)
Panel D. Include BEA region dummies	23.398***	23.705***
	(5.363)	(5.631)
Panel E. Include state dummies	24.485***	24.694***
	(6.250)	(6.544)
Panel F. Include BEA region dummies and percent electrical workers 1980 × year dummies	23.088**	23.268**
	(6.964)	(7.053)
Panel G. Include BEA region dummies and percent electrical workers 1980 × year dummies and 1940 immigrant shares (λ)	24.463**	24.364**
	(7.766)	(7.878)
Panel H. Include BEA region dummies and percent electrical workers 1980 × year dummies; exclude share college natives	17.633**	17.837**
	(5.641)	(5.771)
Panel I. Include BEA region dummies × post-1980	18.552*	18.530*
	(7.647)	(7.629)
Panel J. Include BEA region dummies × post-1980; exclude share college natives	12.284*	12.129*
	(6.082)	(6.016)

Results are weighted by $1/(1/pop_{t+1} + 1/pop_t)$ as in Hunt and Gauthier-Loiselle (2010). Standard errors in parentheses. Statistical significance indicators in table correspond to * =.05, **=.01, and ***=0.001

Table 2: Recreation of Table 8 (IV only)

	College graduates		Post-college	
	HGL (1)	Constructed (2)	HGL (3)	Constructed (4)
Panel A. Include BEA region dummies percent electrical workers 1980 × year dummies; IV excludes share college natives	17.633**	17.837**	18.913	20.198
	(5.641)	(5.771)	(14.272)	(15.017)
Panel B. Include BEA region dummies × post-1980; IV excludes share college natives	12.284*	12.129*	26.960*	26.568+
	(6.082)	(6.016)	(13.356)	(13.944)

Results are weighted by $1/(1/pop_{t+1} + 1/pop_t)$ as in Hunt and Gauthier-Loiselle (2010). Standard errors in parentheses. Statistical signifiance indicators in table correspond to + = .1, *=.05, and **=.01

3 Advances in our understanding of Bartik instruments

In recent years there have been several prominent papers that suggest we should reconsider the validity of Bartik (shift-share) instruments. Goldsmith-Pinkham et al. (2020) decompose Bartik instruments to demonstrate that identification is actually an argument about the shares (in this case the 1940 population shares for each country of origin) rather than the shifts (here, the annual flows of skilled—college, post college, science and engineering)—immigrants from each country of origin). More specifically, Bartik instruments are equivalent to a GMM strategy with original shares as the instruments. This pooled-exposure design requires identification assumption that there are no other shocks correlated with the differential exposure.

Focusing on the IV estimates from Table 8 (columns 2 and 4), which are the preferred specifications noted in Hunt and Gauthier-Loiselle (2010), I recreate the coefficients using the `bartik.weights` package in `R` (Chen (2018)).[1] While column 2 of Table 8 provides an estimated effect of 17.6, the coefficient I recover is 20.3. For column 4, HGL find 18.9 and I find 36.6.

Table 3: Negative and positive weights

	Sum	Mean	Share
Panel A. College graduates			
Negative	-0.099	-0.002	0.454
Positive	1.099	0.019	0.546
Panel B. Post-college			
Negative	-0.229	-0.005	0.444
Positive	1.229	0.020	0.556

Table 3 presents the sum, mean, and share of negative and positive weights. For college graduates (Panel A), 45.4% of weights are negative and for post-graduates (Panel B) 44.4% of weights are negative. Negative weights don't pose a problem under the assumption of constant effects, but could under heterogeneous treatment effects.

[1] The `bartik.weights` package is slightly outdated, so I have updated the package to correspond to the authors' `Stata` version in order to compute some elements for this analysis. The codes will be available at https://github.com/taylorjwright/hgl_2010_replication

Figure 1: Top 5 Weighted Countries by Year, College graduates

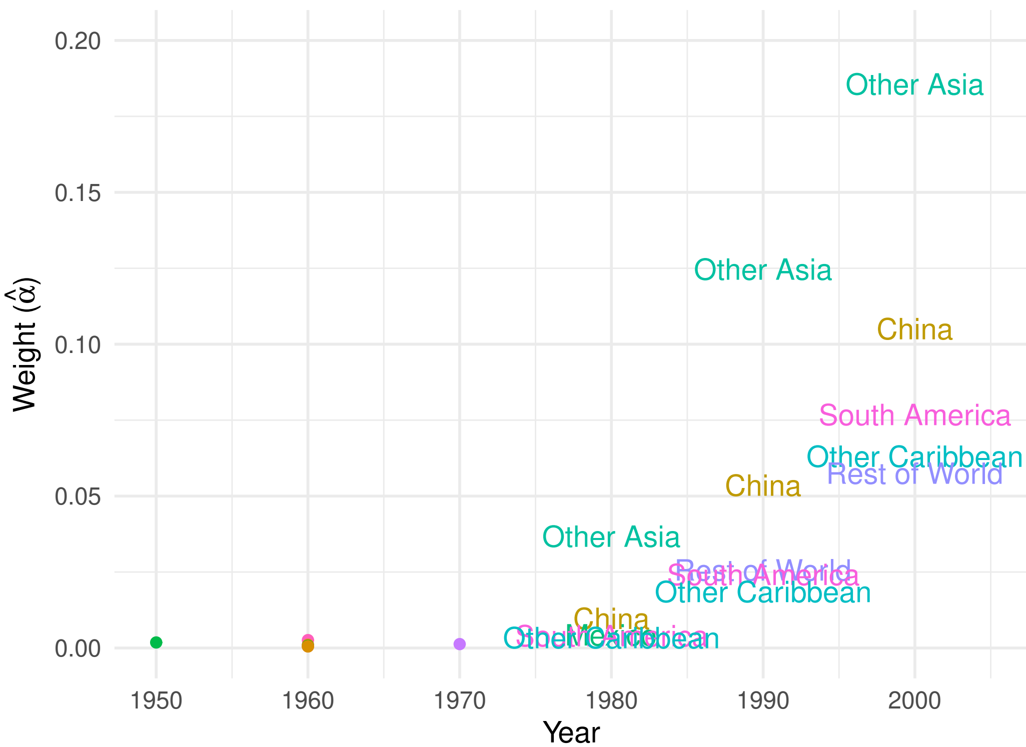

Because shift-share instruments combine many sources of variation, it is not clear which variation is driving the estimates. Additionally, these weights, (α_k) can be interpreted as reflecting the sensitivity of the estimates to the misspecification of the k^{th} instrument. That is, the larger the α_k the larger the bias if that instrument is misspecified. To help characterize the sources of variation and better understand how exposed a research design is to the misspecification sensitivity, Goldsmith-Pinkham et al. (2020) suggest the calculation of the Rotemberg weights assigned to each country of origin. In this setting there are 18 countries/regions that are used as instruments: United Kingdom, Ireland, Italy, Germany, Poland, Russia, Other Europe, Canada, Mexico, Puerto Rico, Cuba, Other Caribbean, Central America, South America, China, India, Other Asia, and Rest of World.

Figures 1 and 2 plots the weights on the y-axis with the year the x-axis for college graduates and post-college graduates, respectively. Pre-1980 all countries received similarly small weights and for legibility I omit naming them. Post-1970 both college and post-college grads see immigrants from "Other Asia" and China receiving relatively large weights.

Figure 2: Top 5 Weighted Countries by Year, Post-college

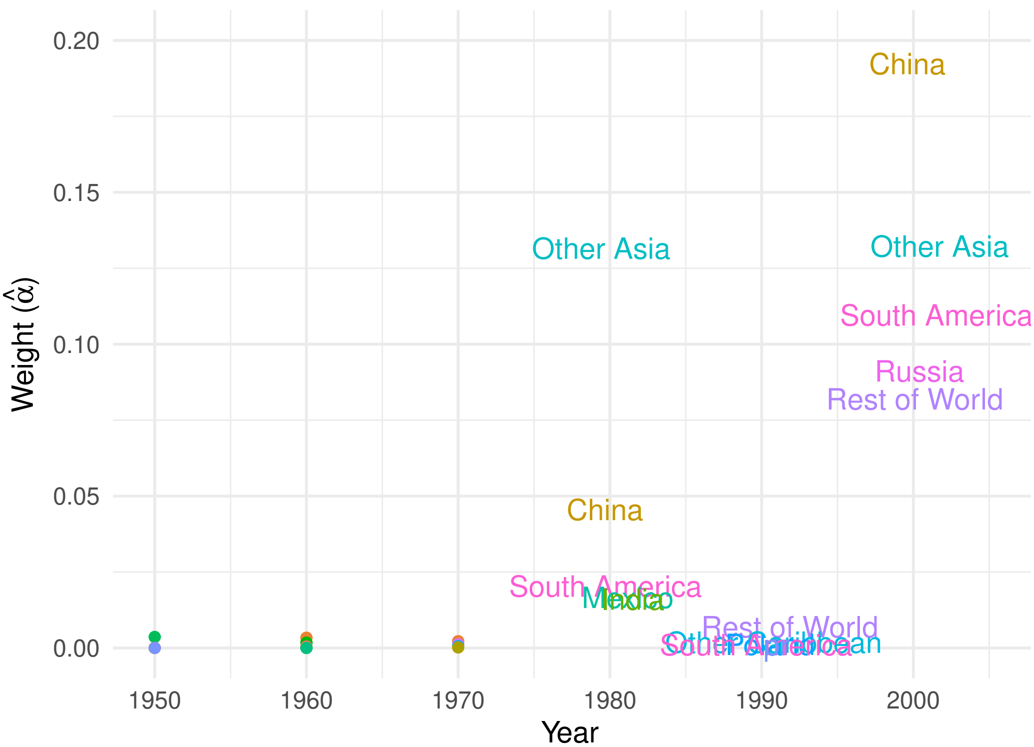

Table 4: Negative and positive weights

	$\hat{\alpha}_k$	g_k	$\hat{\beta}_k$
Panel A. College graduates			
Other Asia, 2000	0.186	598762.06	92.499
Other Asia, 1990	0.125	453060.84	-23.403
China, 2000	0.105	300815.38	42.655
South America, 2000	0.077	196964.48	3.978
Other Caribbean, 2000	0.063	124223.39	-6.924
Panel B. Post-college			
China, 2000	0.192	176500.41	101.565
Other Asia, 2000	0.132	149934.17	241.476
Other Asia, 1980	0.132	121234.09	-92.183
South America, 2000	0.110	85350.87	8.974
Russia, 2000	0.091	95231.81	-18.372

Additionally, Table 4 provides the top 5 country-years for college graduates in Panel A, while Panel B does the same for post-grads. For college graduates the top 5 countries received over 45% of the weight and for post-college graduates the top 5 made up over 65%. Furthermore, four of the top five country-years were from 2000. The idea here is that the weights shed some light on where the variation used in estimation is coming from and highlights which units would have an out-sized impact on the estimates in the face of violations of identification assumptions. In this setting, following the logic of Goldsmith-Pinkham et al. (2020) we should mainly be thinking about the comparisons between places with more and fewer immigrants from "Other Asia" countries (i.e. Asian countries outside China and India) and China and whether or not these locations also have other characteristics that might predict changes in patenting through non-immigration channels and in particular for the year 2000.

To provide a little additional context, the largest country weights that Goldsmith-Pinkham et al. (2020) find in their replication of Card (2009) are 48% for Mexico (high school equivalent workers) and 15% for Philippines (college equivalent workers). So, while these countries are not individually as highly leveraged as Mexico in Card (2009), a small selection of country-years are receiving an out-sized share of the weight.

While this replication exercise is focused on Hunt and Gauthier-Loiselle (2010)... The more common identification assumption that is invoked is that the shares are uncorrelated with changes in the error term after conditioning on observables, which Goldsmith-Pinkham et al. (2020) point out is the same notion that is applied in difference-in-differences frameworks. However, if there are other characteristics that predict changes in the outcome of interest that might operate through a non-immigration channel (as described above) then this assumption no longer holds. It is still possible to have a consistent estimator however, by appealing to the presence of many, exogenous, independent shocks. However, it is not clear that this second approach would be satisfied in this setting as there are relatively few shocks (roughly 100 country-years drawn from 18 countries) used as instruments compared to the hundreds of industries used in Borusyak et al. (2020) where this result is derived) and as Jaeger et al. (2018) note, the countries these shocks come from have become highly correlated after 1970.

Table 5 present the correlations between α_k, β_k, and g_k. These statistics are calculated on the aggregate measures across years for each country. In particular we are interested in the share of variance in the weights, α_k that can be explained by the g_k—that is, how much of the weights are explained by the shocks (immigrant inflows). Panel A presents these correlations for college graduates and I find a correlation of .869. In Panel B, I find a correlation of .616 for post-college graduates. For context, Goldsmith-Pinkham et al. (2020) find correlations of 0.991, and 0.766 between α_k and g_k for high school equivalent and college equivalent workers in their replication of Card (2009). The takeaway here is that the shocks explain a great deal of the variance in weights for college graduates ($\sim$ 75%) and somewhat less for post-college graduates ($\sim$ 38%). It is also worth pointing out that though these correlations are lower than those for the Card (2009) replication, they are still a great deal larger than the canonical setting or even in a China shock example present in the working paper version of Goldsmith-Pinkham et al. (2020). This helps provide some support for the common intuition of the shift-share instruments in this setting.

Table 5: Correlations

	$\hat{\alpha_k}$	$\hat{\beta}_k$	g_k
Panel A. College graduates			
$\hat{\alpha_k}$	1.000		
$\hat{\beta}_k$	0.012	1.000	
g_k	0.869	0.241	1.000
Panel B. Post-college			
$\hat{\alpha_k}$	1.000		
$\hat{\beta}_k$	-0.164	1.000	
g_k	0.616	-0.170	1.000

Despite the apparent support for the intuition that identification comes from the shocks, Goldsmith-Pinkham et al. (2020) suggest that in what they refer to as the "immigrant enclave" literature it is natural to think about identification as coming from the shares. In particular, these immigrant shocks from various countries are affecting locations differentially and that this might be related to the "pull" factors of the existing shares (the immigrant enclaves). In contrast, as noted above, the shocks identification requires that there exist some random "push" factors and notably that there are sufficiently many of them to wash out the endogeneity of the shares. Given that there tend to be relatively few instruments in these immigration settings, it seems unlikely that this holds.

Regardless, we can examine certain aspects of these identifying assumptions. Goldsmith-Pinkham et al. (2020) suggest several exercises including examining the relationship between location characteristics and origin country shares; and examination of pretrends; and the use of alternative estimators and over-identification tests.

Table 6: Relationship between Origin Country Shares and Characteristics

	Other Asia	China	South America	Other Caribbean	Russia	Rest of World	Bartik College	Bartik Post-college
Population 1940 (log)	0.027	0.047**	0.066+	0.071	0.057*	0.041*	0.001	0.001
	(0.027)	(0.016)	(0.035)	(0.054)	(0.027)	(0.016)	(0.001)	(0.000)
State personal income	0.181	0.149*	0.120*	0.108	0.086*	0.085**	0.014***	0.005***
per capita 1940 (log)	(0.129)	(0.068)	(0.047)	(0.077)	(0.042)	(0.024)	(0.003)	(0.001)
Num.Obs.	294	294	294	294	294	294	294	294
R^2	0.180	0.408	0.431	0.247	0.444	0.567	0.223	0.241

Each column reports results of a single regression of a 1940 origin country share on 1940 characteristics. Results are weighted by $1/(1/pop_{t+1} + 1/pop_t)$ as in Hunt and Gauthier-Loiselle (2010). Standard errors in parentheses. Statistical signifiance indicators correspond to + = .1, *=.05, and **=.01

Table 6 presents the results from the location characteristics and origin country shares exercise. Each column is the result of running a regression of a 1940 origin country share on the 1940 characteristics used in the analysis by Hunt and Gauthier-Loiselle (2010) (state personal income per capita and population). The authors included a similar table (Table 5), but this was not broken down by origin country and given the out-sized weights on several countries, we would like to explore a the validity of the identification for those countries specifically. The first thing to note is that these characteristics explain a great deal of variation in the shares, ranging from 18% to over 50%. I find that for these top origin countries there is a positive relationship between state personal income per capita and the shares that is statistically significant for several of the top origin countries. Population is also positively correlated with the shares and statistically significant for several of the origin countries. Additionally, there is a positive correlation between the instruments and the characteristics.

Figure 3: Top 5 Weighted Countries by Year, Post-college

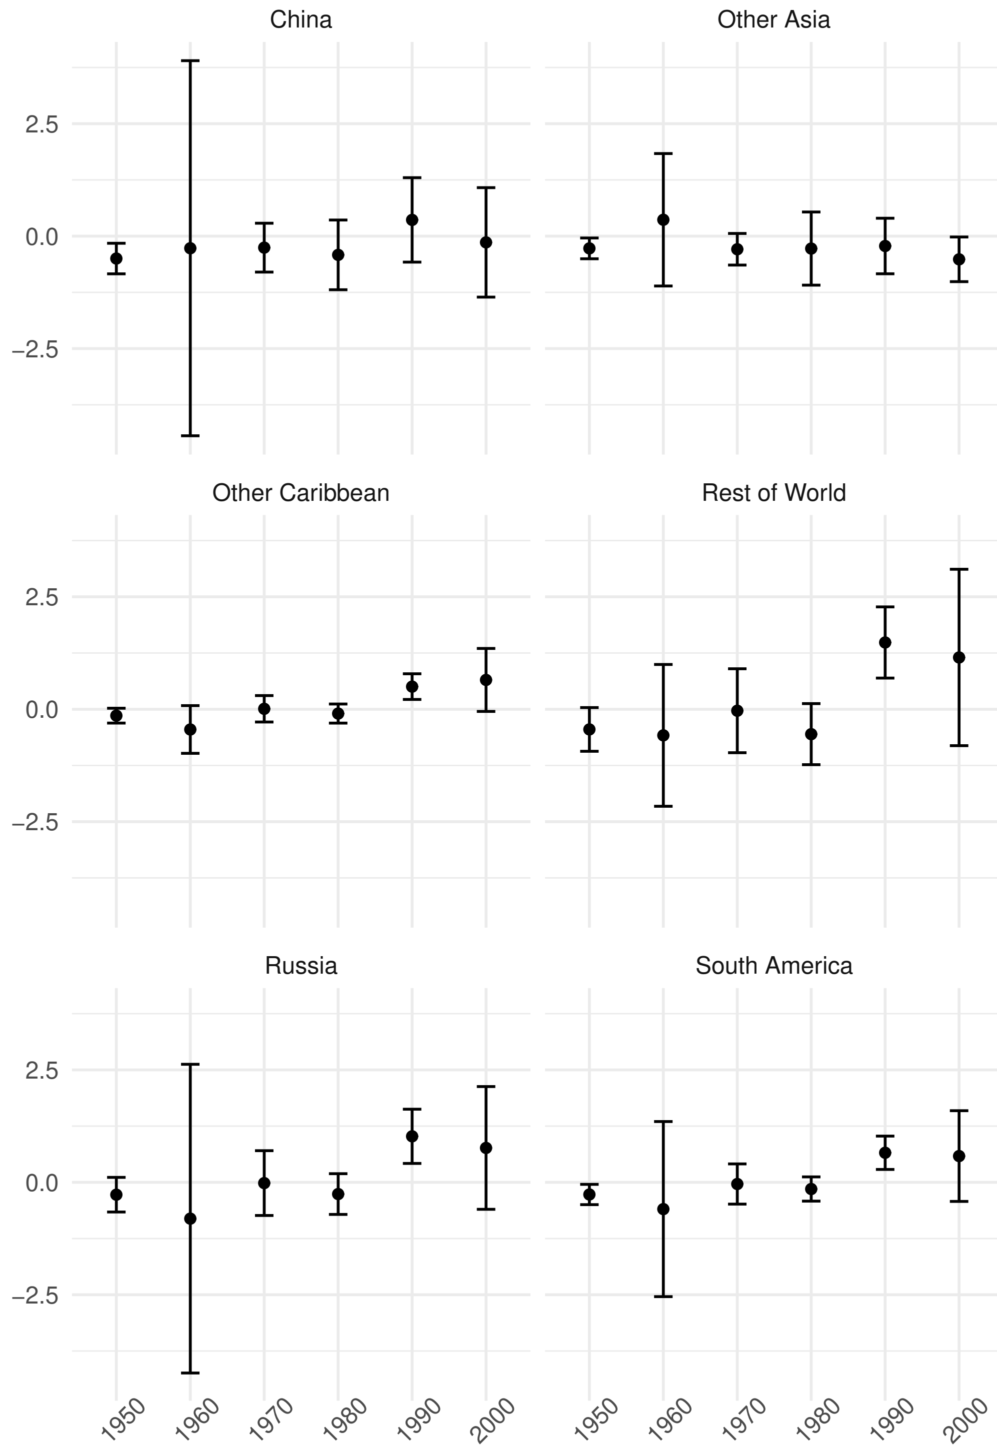

Figure 4: These figures report pretrends for the overall instrument and the top-5 Rotemberg weight origin countries as reported in panel B of Table 4. The coefficients are estimated using the reduced-form regression of equations (11) and (12) with their 1980, 1990, and 2000 values (that is, we include all the controls in Card 2009 in Table 6, columns 3 and 7, and re-estimate year-by-year). Hence, the 2000 coefficient corresponds to the reduced-form coefficient estimated in Table 6. The Others are Cyprus, New Zealand, Israel, and Australia.

Figure 5: Top 5 Weighted Countries by Year, Post-college

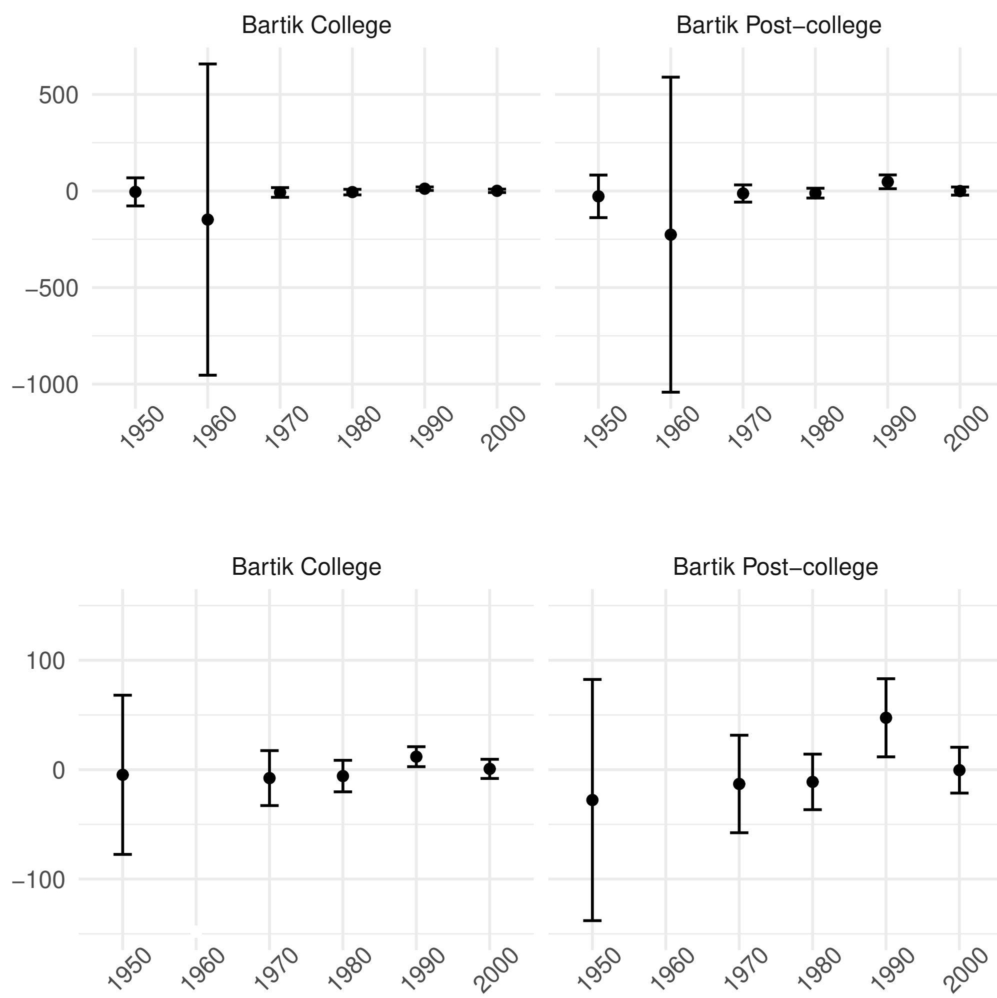

Figure 6: These figures report pretrends for the overall instrument and the top-5 Rotemberg weight origin countries as reported in panel B of Table 4. The coefficients are estimated using the reduced-form regression of equations (11) and (12) with their 1980, 1990, and 2000 values (that is, we include all the controls in Card 2009 in Table 6, columns 3 and 7, and re-estimate year-by-year). Hence, the 2000 coefficient corresponds to the reduced-form coefficient estimated in Table 6. The Others are Cyprus, New Zealand, Israel, and Australia.

Figures 3 and 5 present the "pretrends" which are the reduced form estimates of the main specification (the dependent variable is the difference in log patents across ten years, with a lead of one year compared to the control variables) estimated separately for the 1940 shares of each of the top weighted origin countries (and the instruments themselves). Each point on these figures is a regression restricting the sample to that year. One of the most striking elements of these figures is the enormous standard errors for 1960. These large standard errors are present in the instruments (both post-college graduates and college graduates) but also for each of the top weighted origin countries and especially China and Russia. This aligns with Hunt and Gauthier-Loiselle (2010) who indicate the the instrument is weak in the lower immigration decades (1940-1950 and 1950-1960). Due to the large standard errors for the instruments overshadowing any trends in the other years, I also present the estimates omitting 1960 which highlights the larger standard errors in 1950 for the instruments. While there is no policy change in the traditional difference-in-differences sense that would naturally lend itself to examining pretrends, Goldsmith-Pinkham et al. (2020) write "We suspect that researchers will be more comfortable with the plausibility of their empirical design if parallel pretrends are satisfied for the instruments to which their estimates are most sensitive to misspecification." (p. 2606). What we are then examining is whether 1940 shares from these origin countries predict systematically larger patenting changes and the existence of trends (especially leading up to year 2000 whose shocks we saw had out-sized weight in the instrument). While there does not appear to be any statistically significant pretrends, we may be concerned the serial correlation in inflows, which the elevated point estimates in 1990 and 2000 for several of the countries could characterize, could mean that the shocks from the inflows are not cleanly tied to a specific period (this is a point made by Jaeger et al. (2018)).[2] However, it does not appear that the 1940 shares for the two origin locations that received high weight in 1990 and 2000 (China and Other Asian countries) predict larger patenting changes in those years, which is perhaps reassuring on this point.

Table 7: Alternative Estimators

	(1)
OLS	14.80974
TSLS	15.97702
Bartik TSLS	20.30000
LIML	16.39724
MBTSLS	16.20218

This table reports a variety of estimates of the effects of immigration on patenting rate. The regressions are at the state level and include a time periods 1950-2000. The TSLS row is my replication of column 2 Hunt and Gauthier-Loiselle (2010). The TSLS (Bartik) row uses the Bartik instrument. The TSLS row uses each origin country share interacted with the immigration inflows separately as instruments. The MBTSLS row uses the estimator of Kolesár et al. (2015) with the same set of instruments. The LIML row shows estimates using the limited information maximum likelihood estimator with the same set of instruments. Results are weighted by $1/(1/pop_{t+1} + 1/pop_t)$.

Lastly, I estimate alternative estimators as suggested by Goldsmith-Pinkham et al. (2020) using the `ManyIV` package introduced alongside Kolesár (2018). Table 7 presents the results comparing the standard two-stage least squares estimator (TSLS), the Limited Information Maximum Likelihood (LIML), the Modified Bias-corrected TSLS (MBTSLS), and Bartik TSLS estimators. There is a great deal of correspondence among these estimators, which is suggestive evidence against misspecification (in the homogeneous effect interpretation). However, the Sargan over-identification test rejects (though the Modified Cragg-Donald does not) for TSLS.

[2] Jaeger et al. (2018) specifically examine the validity of long-term estimates in the context of immigration. They point out that the long-term estimates of immigration effects are not well identified in the United States because of the concentration of immigrants coming from only a few countries, with little variation post-1970. If researchers would like to use a 1 period lag (here, and often elsewhere, 10 years because data comes from the decennial census) they argue that we can therefore only provide estimates of effects in the 1980s. As the purpose of this replication is to examine the sensitivity of estimates to this new understanding of the source of variation used in the estimation, I do not engage with this critique of the shift-share instruments in the immigration enclave setting, though I do acknowledge its importance for this literature.

Figure 7: Top 5 Weighted Countries by Year, College graduates

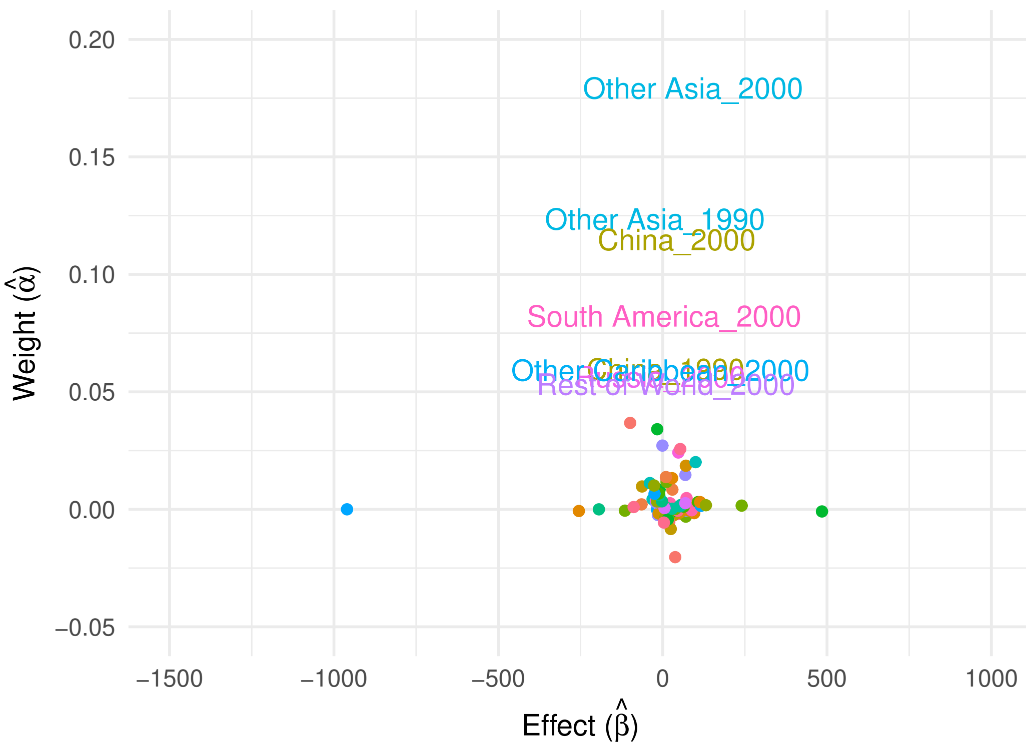

Figure 8: Top 5 Weighted Countries by Year, Post-college

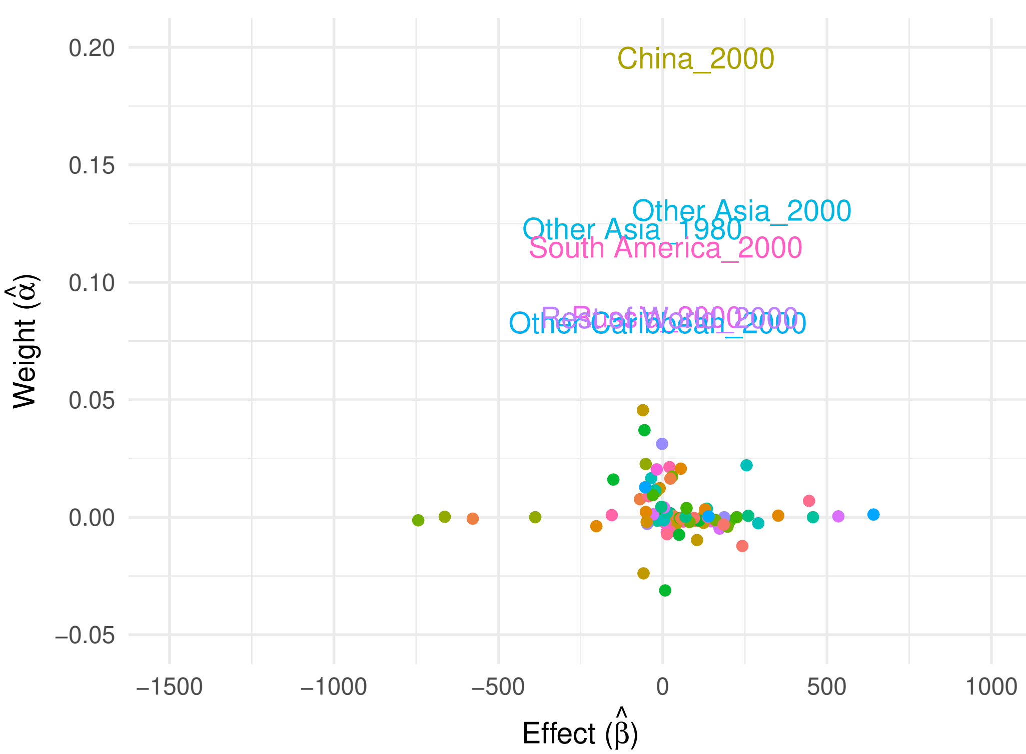

Additionally, Figures 7 and 8 plot the weights on the y-axis with the estimated coefficients on the x-axis for college graduates and post-college graduates, respectively. One conclusion from this is that the higher weight origin countries do not have dramatically different β_k than other countries. There are also a few country-year observations that have very large (in absolute terms) β_k but these receive essentially 0 weight.

4 Conclusion

How sensitive are the results from shift-share instruments to the recent literature that finds identification most likely comes from the shares, especially in the immigration enclave setting? Following the guidance for empirical researchers in Goldsmith-Pinkham et al. (2020), I revisit Hunt and Gauthier-Loiselle (2010) where the shares are the 1940 shares of migrants from a particular origin country living in a particular location and the shocks/shifts are the national immigrant inflows.

I first am able to exactly reproduce the original analysis and further reconstruct the instrument used to produce very similar results. Next I find that for college graduates the Rotemberg weights are very heavily explained by immigrant inflows while this is less true for the post-college graduates. In this case, the explanatory power of shocks (immigrant inflows) actually provides a fair characterization of the variation used in the estimation. Third, I find that a very few origin country-years (mainly from the year 2000) account for the bulk of the weight in the estimator. This suggests that the comparison we should have in mind is between places with greater and fewer immigrants form China or Other Asian countries in 2000. Fourth, I do see some patterns in the correlations between 1940 characteristics and immigrant shares from top weighted origin countries and instruments, especially per capita income. Fifth, I do not find statistically significant pretrends for the top weighted origin countries though there may be some small concern hinted at by the elevated point estimates in 1990 and 2000 (and consistent with Jaeger et al. (2018)). Sixth, I do not find large differences among the alternative estimators.

Overall, it appears Hunt and Gauthier-Loiselle (2010) holds up quite well under the new scrutiny brought by the recent shift-share instrument literature.

References

Adao, R., Kolesár, M. and Morales, E.: 2019, Shift-share designs: Theory and inference, *The Quarterly Journal of Economics* **134**(4), 1949–2010.

Altonji, J. G. and Card, D.: 1991, The effects of immigration on the labor market outcomes of less-skilled natives, *Immigration, trade, and the labor market*, University of Chicago Press, pp. 201–234.

Bergé, L.: 2018, Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm, *CREA Discussion Papers* (13).

Blit, J., Skuterud, M. and Zhang, J.: 2018, An Analysis of the Patenting Rates of Canada's Ethnic Populations, *Canadian Public Policy-analyse De Politiques* **44**(1), 1–20.

Borjas, G. J., Freeman, R. B., Katz, L. F., DiNardo, J. and Abowd, J. M.: 1997, How much do immigration and trade affect labor market outcomes?, *Brookings papers on economic activity* **1997**(1), 1–90.

Borusyak, K., Hull, P. and Jaravel, X.: 2020, Quasi-Experimental Shift-Share Research Designs.

Bound, J., Khanna, G. and Morales, N.: 2017, Understanding the Economic Impact of the H-1B Program on the U.S., *National Bureau of Economic Research* .

Card, D.: 2001, Immigrant inflows, native outflows, and the local labor market impacts of higher immigration, *Journal of Labor Economics* **19**(1), 22–64.

Card, D.: 2009, Immigration and inequality, *American Economic Review: Papers and Proceedings* **99**(2), 1–21.

Chen, J.: 2018, *bartik.weight: Estimating Rotemberg Weights for a Bartik "Instrument"*. R package version 0.1.0.
URL: *https://github.com/jjchern/bartik.weight*

Edo, A.: 2019, THE IMPACT OF IMMIGRATION ON THE LABOR MARKET, *Journal of Economic Surveys* **33**(3), 922–948.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H.: 2020, Bartik Instruments: What, When, Why, and How, *American Economic Review* **110**(8), 2586–2624.

Hunt, J. and Gauthier-Loiselle, M.: 2010, How Much Does Immigration Boost Innovation?, *American Economic Journal: Macroeconomics* **2**(2), 31–56.

Jaeger, D. A., Ruist, J. and Stuhler, J.: 2018, Shift-Share Instruments and the Impact of Immigration.

Kolesár, M.: 2018, Minimum distance approach to inference with many instruments, *Journal of Econometrics* **204**(1), 86–100.

Peri, G., Shih, K. and Sparber, C.: 2016, Stem workers h 1b visas and productivity in us cities, *Journal of Labor Economics* .

R Core Team: 2022, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: *https://www.R-project.org/*