Economics Bulletin

Volume 41, Issue 4

Immigration and housing rents in Switzerland: Identification in a shift-share research design

Dario Fauceglia
Zurich University of Applied Sciences and University of St.Gallen

Tilman Slembeck
Zurich University of Applied Sciences and University of St.Gallen

Abstract

Drawing on geo-coded property data, we exploit aggregate immigration shocks in a shift-share instrument as the identifying quasi-random variation - conditional on control variables and fixed effects - to estimate the elasticity of residential housing rents to local immigration in Switzerland. In our balanced sample of municipalities the estimated average elasticity ranges between 2.40 and 2.98. Interestingly, exposure-robust inference, which avoids potentially downward biased standard errors in shift-share regressions as shown by Adao et al. (2019) (AKM), does not systematically decrease the precision of the estimated immigration effect in our application. However, these exposure-robust standard errors reveal that the elasticity of housing rents to immigration does not statistically differ from unity across specifications at the 5% significance level.

We thank Daniel Sager for providing data and helpful discussions and comments and Kirill Borusyak for helpful comments and providing stata codes for data transformation (ssaggregate) and null-imposed inference.

Citation: Dario Fauceglia and Tilman Slembeck, (2021) "Immigration and housing rents in Switzerland: Identification in a shift-share research design", *Economics Bulletin*, Vol. 41 No. 4 pp.2561-2573 .
Contact: Dario Fauceglia - fauc@zhaw.ch, Tilman Slembeck - stil@zhaw.ch.
Submitted: April 12, 2021. **Published:** December 29, 2021.

1 Introduction

To identify the elasticity of housing rents to local immigration, we use a shift-share (or "Bartik", 1991) instrumental variable (SSIV) that allocates inflows of immigrants from different nationalities to municipalities exposed differentially to them, with exposure depending on immigrants' past nationality shares in municipalities' populations. This SSIV is relevant because immigrants tend to move to local enclaves of people who share their nationality. It is intended to address unobservables correlated with immigration and rents, such as improved local amenities, or reverse causality in the sense that immigrants might respond differently to current and future expected housing rents than natives. As shown by recent econometric advances, the identifying variation in this shift-share research design may stem from local nationality shares that drive spatial settlement patterns of immigrants (Goldsmith-Pinkham et al., 2020) (GPSS) or immigration shocks (or "shifters") at the aggregate level. (Borusyak et al., 2021) (BHJ) (and AKM). In contrast, previous applied literature on housing argued that shares and shocks must simultaneously meet the exclusion restriction (e.g. Saiz, 2007 or Accetturo et al., 2014). As opposed to existing housing studies examining the immigration effect with SSIVs, such as Saiz (2007), Accetturo et al. (2014), Sá (2015), Degen and Fischer (2017), Sharpe (2019) or Moallemi and Melser (2020), we take into account the results of BHJ in the specification of our empirical model in order to be transparent about exploiting aggregate immigration shocks as our source of identifying variation within the shift-share instrument. Since the source of exogenous variation used for identification remains unclear in these previous shift-share studies, the credibility of their identification strategies is difficult to assess. Importantly, as required in the framework of BHJ, we construct our instrument with a sufficient number of immigration shocks from different countries and years and isolate the random variation in these shocks with exposure-weighted control variables and fixed effects. Again in contrast to the housing literature cited above, we also recover and report standard errors that are robust to residual correlation across municipalities exposed to similar nationality shares, potentially leading to a false rejection of the null of no immigration effect according to AKM. This exposure-robust inference does not systematically reduce the significance of the positive immigration effect on housing rents and this finding, which mitigates the concerns expressed in AKM, is another contribution of this paper. In addition, it suggests uncorrelated residuals along exposure shares so that we arguably also identify the immigration effect in the alternative econometric framework of GPSS, which relies on the (conditional) exogeneity of shares for identification. Finally, there is little or no empirical evidence on the nexus between immigration and housing rents from outside the US (see Saiz, 2003, 2007 and Sharpe, 2019 for the US). We fill part of this gap with this paper on the effect of immigration on local housing rents in Switzerland.

2 Data

We employ proprietary geo-coded data on offered rents and housing characteristics provided by Meta-Sys. These information have been collected from postings on advertising websites active in Switzerland. Basten et al. (2017) show for data drawn from a popular online platform that aggregate (cantonal) indices of offered rental prices follow mostly aggregate indices of agreed rents that are based on realized transactions. We exclude postings on commercial properties. The remaining data is sourced from the Swiss Federal Statistical Office (BFS) and the World's Bank WDI database. Importantly, the immigration data taken from the annual population statistics of the BFS has been consolidated by the BFS ensuring its consistency over the entire sample period.

3 A hedonic price regression

We estimate the following hedonic price model (De Haan and Diewert, 2013) with around 1.082 million observed properties i over 2004 -2018 ($= t$) by pooled OLS

$$log(r_{ilt}) = Q_{ilt}^{'}\Omega + \mu_{lt} + \epsilon_{ilt}, \qquad \mathbb{E}[\epsilon_{ilt}|Q_{ilt}, \mu_{lt}] = \mathbb{E}[\epsilon_{ilt}] = 0 \tag{1}$$

in order to separate the municipality-level determinants of log rents $log(r_{ilt})$, as captured by the estimated fixed effects $\hat{\mu}_{lt}$, from rent effects due to changes in the quality composition of properties within a municipality. Specifically, the vector of quality variables Q_{ilt} controls among others for (log) construction year, living space surface, additional property area, number of rooms, scenic views (e.g. lake) and for being a single-family house or the first party moving in after building or renovation. To construct our main estimation sample, we keep the predicted annual log housing rents $\widehat{log(r_{lt})} \equiv \hat{\mu}_{lt}$ of municipalities that we observe in every year. This selection rule results in a balanced sample of 10094 municipality-year pairs. Our dependent variable in equation (4) is the annual change in log rents $\Delta\widehat{log(r_{lt})} \equiv \hat{\mu}_{lt} - \hat{\mu}_{lt-1}$ that has a mean of 0.01 and a first and third quartile corresponding to -0.015 and 0.034 respectively.

4 Theoretical considerations

We understand immigration into municipalities as a shifter of local housing demand. As in Saiz (2007), we therefore employ gross immigrant inflows to construct our treatment variable as opposed to net migration inflows that also reflect movements along housing demand. Since immigration creates an upward pressure on rental prices, it could lead to out-migration of local natives and previously settled migrants. Immigration may also result in native flight unrelated to rent changes but driven by preferences for segregation (Saiz and Wachter, 2011). This potential displacement of local residents dampens the required increase in rents in order to restore a new housing equilibrium and should be captured by our estimates as part of the reaction to the initial immigration shock.

5 A SSIV estimator

We study the change in annual log housing rents $\widehat{log(r_{lt})}$ in municipality l associated with a 1% increase in the lagged immigration rate $\widetilde{IM}_{lt-1} \equiv \frac{IM_{lt-1}}{POP_{lt-2}}$, which is defined as the local annual inflow of immigrants in $t-1$ (IM_{lt-1}) as a share of the municipality's population in $t-2$ (POP_{lt-2}). Our coefficient of interest is β, which measures the β% change in average housing rents across municipalities in response to a 1% increase in the explanatory variable $\widetilde{IM}_{lt-1}$. This elasticitiy β may be recovered from

$$\widehat{log(r_{lt})} = \beta \widetilde{IM}_{lt-1} + Z'_{lt-1}\gamma + \varepsilon_{lt}, \tag{2}$$

where Z_{lt-1} is a vector of controls, while $\widetilde{IM}_{lt-1}$ is instrumented by the "shift-share"-variable I_{lt-1} that has the following form:

$$I_{lt-1} \equiv \sum_{c=1}^{54} s_{lct_0} \cdot IMS_{ct-1}, \qquad s_{lct_0} \equiv \frac{POP_{lct_0}}{POP_{lt_0}}, \qquad IMS_{ct-1} \equiv \frac{IM_{ct-1}}{POP_{ct_0}}, \tag{3}$$
$$s_{lct_0} \geq 0 \quad \forall \ (l,c), \qquad S_{lt_0} \equiv \sum_{c=1}^{54} s_{lct_0} < 1 \quad \forall \ l,$$

with s_{lct_0} denoting the number of nationals from origin country c in municipality l (POP_{lct_0}) as a share of the municipality's total population (POP_{lt_0}) in the initial year 2001 ($= t_0$) that predates the sample period 2004-2018; IMS_{ct-1} is the aggregate number of immigrants from c entering Switzerland (IM_{ct-1}) in $t-1$ divided by the aggregate number of nationals from c that lived in Switzerland in 2001 (POP_{ct_0}); S_{lt_0} is the initial share (in t_0) of foreign nationals in municipality l. We predict local immigration I_{lt-1} by a weighted average of IMS_{ct-1}, with shares s_{lct_0} measuring the differential exposure of municipalities to aggregate immigration IMS_{ct-1} from specific origin countries c. We take first differences Δ of equation (2) to obtain:

$$\Delta\widehat{log(r_{lt})} = \beta \Delta\widetilde{IM}_{lt-1} + \Delta Z'_{lt-1}\gamma + \Delta\varepsilon_{lt}, \tag{4}$$

instrumenting $\Delta\widetilde{IM}_{lt-1}$ with $\Delta I_{lt-1} \equiv \sum_{c=1}^{54} s_{lct_0} \Delta IMS_{ct-1}$. Before we discuss how to identify β armed with this specification in Section 7, we present in Section 6 a specification equivalent to (4) that we employ to recover exposure-robust standard errors, as suggested in BHJ and AKM, and a moment condition that must be satisfied for identification in the framework of BHJ.

6 An equivalent country-level IV estimator

BHJ show that $\hat{\beta}$ from estimating (4) is equivalently obtained by a IV regression that uses the average country exposure $s_c \equiv \frac{1}{721}\sum_{l=1}^{721} s_{lct_0}$ as weights, and in which the aggregate immigration shocks by countries ΔIMS_{ct-1} serve directly as instruments in estimating

$$\overline{\Delta \widehat{log}(r_{ct})}^{\perp} = \beta \overline{\Delta \widetilde{IM}}_{ct-1}^{\perp} + \overline{\Delta \varepsilon}_{ct}^{\perp}, \tag{5}$$

where $\overline{\quad}$ indicates a weighting of the variables by exposure (e.g. $\overline{\Delta \widetilde{IM}}_{ct-1}^{\perp} \equiv \frac{1}{721}\sum_{l=1}^{721} \frac{s_{lct_0}}{s_c} \Delta \widetilde{IM}_{lt-1}^{\perp}$) so that specification (5) as opposed to (4) varies over (c,t) instead of (l,t). $\perp$ denotes residualization on ΔZ_{lt-1}. Given instrument relevance, β is identified if the following moment condition in the considered non-iid setting is satisfied: [1] [2]

$$\mathbb{E}\left[\frac{1}{14}\sum_{t=1}^{14}\sum_{c=1}^{54} s_c \Delta IMS_{ct-1}\overline{\Delta\varepsilon}_{ct}\right] = \mathbb{E}\left[\frac{1}{14\times 721}\sum_{t=1}^{14}\sum_{l=1}^{721}\underbrace{\sum_{c=1}^{54} s_{lct_0}\Delta IMS_{ct-1}}_{\equiv \Delta I_{lt-1}}\Delta\varepsilon_{lt}\right] = 0, \tag{6}$$

The left-hand side of equation (6) expresses this condition within the framework of specification (5), while the right-hand side is the more conventional equivalent formulation based on specification (4). More fundamentally, equation (6) implies that β is identified under the exactly same set of conditions in both specifications. Specifically, the left-hand side of equation (6) shows that the orthogonality of ΔIMS_{ct-1} with $\overline{\Delta\varepsilon}_{ct} \equiv \frac{1}{721}\sum_{l=1}^{721}\frac{s_{lct_0}}{s_c}\Delta\varepsilon_{lt}$ when weighted by s_c guarantees the orthogonality between the shift-share instrument ΔI_{lt-1} and municipality-level confounders $\Delta\varepsilon_{lt}$ on the right-hand side of equation (6) (and vice versa). To achieve identification according to the left-hand side of equation (6), lagged changes in country-level immigration ΔIMS_{ct-1} must be uncorrelated (s_c-weighted) in expectation with $\overline{\Delta\varepsilon}_{ct}$ that reflects the average changes in current unobserved determinants of rent growth in municipalities disproportionately affected by immigration from country c in terms of relative exposure $\frac{s_{lct_0}}{s_c}$.

[1] To derive the left-hand side of equation (6), a consistent estimate $\hat{\gamma}$ of γ in equation (4) is assumed so that $\Delta\varepsilon_{lt} = \Delta\varepsilon_{lt}^{\perp}$ and $\overline{\Delta\varepsilon}_{ct} = \overline{\Delta\varepsilon}_{ct}^{\perp}$ holds asymptotically.

[2] Equation (6) is a generalization of $E[\Delta I_{lt-1}\Delta\varepsilon_{lt}] = 0$ for a non-iid setup that is required since exposure shares create instrument correlation by treating ΔIMS_{ct-1} as random and potentially also cross-residual correlation by mediating shocks other than immigration.

7 Identification from quasi-randomized shocks

As proposed by BHJ, the moment condition (6) is fulfilled if the aggregate immigration shocks are as-good-as-randomly assigned, implying each ΔIMS_{ct-1} has the same mean regardless of realized s_c and $\overline{\Delta\varepsilon}_{ct}$. Since immigration from different countries and periods differ systematically, random shock assignment may solely hold after conditioning on controls. We thus rely on the first-difference model (4) in which municipality fixed effects have been differenced out.[3] These fixed effects not only absorb time-invariant municipality-level confounders included in $\frac{1}{14}\sum_{t=1}^{14}\varepsilon_{lt}$ from ε_{lt} but also time-invariant differences in country-specific immigration $\frac{1}{14}\sum_{t=1}^{14}IMS_{ct-1}$ from IMS_{ct-1}. Similarly, the year fixed effects included in ΔZ_{lt-1} also remove year-specific differences in average immigration changes $\frac{1}{54}\sum_{c=1}^{54}\Delta IMS_{ct-1}$ from ΔIMS_{ct-1} apart from taking out average changes in year-specific unobservables $\frac{1}{721}\sum_{l=1}^{721}\Delta\varepsilon_{lt}$ from $\Delta\varepsilon_{lt}$. Alternatively, we employ distinct sets of year fixed effects for each of the eight municipality types within the seven major Swiss regions. For example, this prevents a bias whenever aggregate immigration is driven by diverging rental price dynamics for urban as opposed to rural municipalities within a region. In some specifications we include unit fixed effects in equation (4) that capture separate linear trends $\frac{1}{14}\sum_{t=1}^{14}\Delta\varepsilon_{lt}$ for each municipality while they also purge $\frac{1}{14}\sum_{t=1}^{14}\Delta IMS_{ct-1}$ from ΔIMS_{ct-1}. These unit fixed effects take into account that both, past and present, aggregate immigration may be primarily attracted to a few fast growing "superstar" cities with steadily rising rents (Gyourko et al., 2013), as documented in Albert and Monras (2021) for the US. This would invalidate the use of the immigration shocks ΔIMS_{ct-1} (and of the shares s_{lct0}) as a source of identifying variation in our shift-share specification unless such long-term trends are controlled for (see also Sharpe, 2019 on this). ΔZ_{lt-1} also contains the change in the number of vacant rental properties and the unemployment rate at municipality-level. In sensitivity checks, we account for origin countries' GDP changes and include an indicator for whether citizens of a EU/EFTA country were granted full settlement rights in a given year according to *Free Movement of Persons* agreement. Because S_{lt0} varies across municipalities, all controls that have country-level counterparts must be included in exposure-weighted form in ΔZ_{lt-1}, as described in Table 1. Otherwise one would leverage non-random differences in the share of Swiss nationals $1 - S_{lt_0}$ between municipalities.

[3]Estimation in first-differences also addresses possible non-stationarities of immigration and prices in equation (2).

8 Estimates and inference

Columns 1 to 7 of Table 1 display $\hat{\beta}$ when we pursue identification through shocks. The narrow range of estimates between 2.397 and 2.983 suggests that the specification in column 1 including municipality-level controls as well as year and (differenced out) municipality fixed effects may suffice to isolate the random component of ΔIMS_{ct-1}. These estimates imply that a 1% increase in the number of immigrants in relation to a municipality's population raises average rents by 2.4% to 2.98%. Our average rent to immigration elasticity appears to be somewhat larger than in the US, in particular with respect to the results presented in Sharpe (2019). More specifically, Saiz (2007) concludes that the rent to immigration elasticity is around unity in the US (but rising up to 1.73 in a fixed effects specification), which is statistically not significantly different from our set of results as discussed below, while Sharpe (2019) obtains an elasticity significantly below unity in his shift-share specification that in addition to Saiz (2007) also controls for initial city characteristics. Apart from differences related to identification strategy already discussed, the differing size magnitude of the immigration effect is most likely related to the more inelastic housing supply to prices in Switzerland compared to the US, as confirmed by Caldera and Johansson (2013) that reports a more than ten times higher responsiveness of housing supply to prices in the US relative to Switzerland.[4] This indicates more pervasive geographical and regulatory constraints in Switzerland compared to the US on average that result in lower housing supply elasticities to prices and arguably also to rents for which there are, however, no comparing numbers for the US (von Ehrlich et al., 2018). Moreover, our rent-to-immigration elasticities are in line with the estimates of price-to-immigration elasticities obtained by Degen and Fischer (2017) for the Swiss housing market. We report conventional heteroskedasticity-robust standard errors (and CI's), which allow residuals to be spatially correlated across and serially within groups of neighboring municipalities. AKM show formally and in simulations that a positive correlation of residuals across municipalities with similar shares (that are identical to or correlated with s_{lct0}) results in overrejection of the null of no effect ($\beta_0 = 0$) with conventional standard errors. Therefore, as suggested by BHJ, we also obtain standard errors valid under arbitrary spatial dependence in the residuals, such as along exposure shares, by estimating β from equation (5) at the level of identifying shocks, including relevant country-level controls.[5] These standard errors that are robust to exposure and country clustering of residuals in addition to heteroskedasticity confirm a significant positive effect of immigration on rents around the 5% significance level (see columns 1 to 7

[4] See also Saiz (2010) for a less pronounced difference between Switzerland and the US in the housing supply elasticities.

[5] We also recover exposure-robust first-stage F-statistics as in BHJ.

Table 1: Shift-share IV estimates of the effect of immigration on rents (2004-2018)

	1	2	3	4	5	6	7	8
$\hat{\beta}$	2.760	2.397	2.718	2.983	2.541	2.423	2.658	2.854
Standard error (SE) and confidence interval (CI)								
Municipality group clustered SE	(1.241)	(1.367)	(1.606)	(1.845)	(1.219)	(1.431)	(1.605)	(1.069)
and 95% CI (conventional)	[0.296, 5.224]	[-0.318, 5.111]	[-0.470, 5.907]	[-0.680, 6.645]	[0.120, 4.961]	[-0.418, 5.264]	[-0.529, 5.844]	[0.732,4.975]
Country clustered SE	(1.413)	(1.316)	(1.040)	(1.179)	(1.198)	(1.229)	(1.010)	
and 95% CI (exposure-robust)	[-0.010,5.530]	[-0.182,4.975]	[0.680,4.757]	[0.673,5.293]	[0.192,4.890]	[0.014,4.831]	[0.678,4.637]	
Null imposed country clustered 90% CI	[0.973, 10.11]	[0.600, 7.711]	[0.584, 8.001]	[0.807, 8.073]	[1.008, 8.591]	[0.809, 7.717]	[0.417, 7.107]	
Null imposed country clustered 95% CI	[0.630, 24.69]	[0.220, 13.43]	[$-\infty, +\infty$]	[-0.331, 28.96]	[0.706, 19.73]	[0.489, 14.00]	[$-\infty, +\infty$]	
Municipality-level controls (in equation 4)								
Δ(# of vacant properties)$_{lt-1}$	✓	✓	✓	✓	✓	✓	✓	✓
Δ(Unemployment rate)$_{lt-1}$	✓	✓	✓	✓	✓	✓	✓	✓
S_{lt_0} x Year FE	✓	✓			✓	✓		
S_{lt_0} x Region-MunicipalityType-Year FE			✓	✓			✓	
$\sum_{c=1}^{54}(s_{lct_0}\Delta\mathbf{1}(\text{Full movement treaty}_{ct-1}))$		✓		✓		✓		
$\sum_{c=1}^{54}(s_{lct_0}\Delta(\text{Log real GDP}_{ct-1}))$		✓		✓		✓		
S_{lt_0} x Municipality FE (linear trend)					✓	✓	✓	
Municipality & Region-MunicipalityType-Year FE								✓
Country-level controls (in equation 5)								
Year FE	✓	✓	✓	✓	✓	✓	✓	
Δ(Full movement treaty$_{ct-1}$ FE)		✓		✓		✓		
Δ(Log real GDP$_{ct-1}$))		✓		✓		✓		
Country FE (linear trend)					✓	✓	✓	
# of municipality-year pairs	10094	10094	10094	10094	10094	10094	10094	10094
# of country-year pairs	756	756	756	756	756	756	756	756
First-stage F-stats (exposure-robust, except (8))	64.52	45.06	85.45	53.95	64.00	46.38	81.32	77.87

Note: Regressions include a constant (columns 1-4) or unit fixed effects (columns 5-8).

of Table 1). We also construct and report exposure-robust confidence intervals that impose a null hypothesis $\beta = \beta_0$ and may have better finite-sample properties, in particular in the presence of few or concentrated clusters (see AKM and BHJ for details). These null-imposed confidence intervals become wider asymmetrically around $\hat{\beta}$, such that $\hat{\beta}$ remains statistically significant in most instances, but substantially larger positive effects are not rejected anymore.[6] Interestingly, the exposure-robust confidence intervals – with and without a null imposed – and the conventional standard errors reveal that the estimated elasticities are never significantly above unity at the 5% significance level.

9 Identification via shares?

According to Table 1, considering exposure-robust standard errors (with no null imposed) does not decrease the precision of $\hat{\beta}$. This is in contrast to Autor et al. (2013) in which exposure-robust standard errors increase substantially compared to conventional standard errors (see Table C2 in BHJ). This suggests the absence of unobserved trends

[6] This pattern also appears in AKM and BHJ and is related to low power of null-imposed inference, as simulations in Borusyak and Hull (2020) show.

in $\Delta\varepsilon_{lt}$ that are mediated by past settlement shares s_{lct0} in our application, while the more generic industry employment shares used in Autor et al. (2013) may also capture other trends than increased imports from China to the US, such as industry-specific technology or taste shocks. Put differently, it appears that in our setting the differential exposure of municipalities to immigrants' nationalities affects rents through immigration only, which is a requirement to identify β through the shares s_{lct0}. Akin to a "parallel trends" assumption in difference-in-differences, the exclusion restriction in the shares view is therefore that each exposure share is not related to unobserved trends in the error term conditional on controls, that is $\mathbb{E}[\Delta\varepsilon_{lt} s_{lct0} | \Delta Z_{lt-1}] = 0 \quad \forall\, c$ (see GPSS for details). This identifying assumption does not preclude an heterogeneous effect of immigration on rents, depending, for example, on the skill intensity of immigrants or the level of housing supply elasticities in the immigrants' locations.[7] In column 8 of Table 1, we relax the assumption of "parallel trends" across shares by including unit and time fixed effects (not exposure-weighted) that permit diverging linear municipality trends and municipality-type-specific rent changes within regions in the first-difference equation (4). $\hat{\beta}$ in column 8 of Table 1 is not statistically different from previous estimates displayed in columns 1 to 7, further supporting the idea that identification via shares is feasible as well.

10 Conclusion

Our empirical analysis conducted within a shift-share research design shows that there is a significantly positive effect of immigration on housing rents in Switzerland. Although we rely on aggregate immigration shocks as our (conditionally) exogenous source of variation for identification, we reach the same conclusion when we estimate the immigration effect in the "shares" econometric framework.[8] This positive immigration effect on housing rents also persists when we control for preexisting trends in municipalities that capture, for instance, steady rent appreciations in "superstar" cities with thriving economies over a longer period. Moreover, unreported results suggest that the empirical findings are also robust to estimations in samples that are unbalanced but include a more comprehensive set of municipalities instead.[9] Importantly, although the estimated rent to immigration elasticities range between 2.397 and 2.983, exposure-robust inference indicates that these elasticities do not significantly differ from unity. This analysis can be extended in various directions. For

[7] In practice, $\hat{\beta}$ corresponds to a weighted average of heterogeneous treatment effects β_{lct} that vary potentially along municipalities l, origin countries c and years t (see BHJ Appendix A.1).

[8] In our panel SSIV setting, it may also be important to address a potential bias that conflates the short- and long-run response to immigration (Jaeger et al., 2018).

[9] These results are available upon request.

instance, we have not yet considered the channels through which immigration affects rents nor studied effect heterogeneity or geographical spillover effects of immigration across municipalities that are spatially connected within regional housing and labor markets.[10]

[10]Relatedly, a quantitative spatial model would also help unveil the general-equilibrium mechanisms influencing rents but buried by the fixed effects in the present analysis (Redding, 2020).

References

Accetturo, A., Manaresi, F., Mocetti, S., and Olivieri, E. (2014). Don't stand so close to me: the urban impact of immigration. *Regional Science and Urban Economics*, 45:45–56.

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

Albert, C. and Monras, J. (2021). Immigration and spatial equilibrium: the role of expenditures in the country of origin. *Working Paper.*

Autor, David, H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68.

Bartik, T. J. (1991). Who benefits from state and local economic development policies? *Kalamazoo ME: W.E. Upjohn Institute for Employment Research.*

Basten, C., von Ehrlich, M., and Lassmann, A. (2017). Income taxes, sorting and the costs of housing: Evidence from municipal boundaries in switzerland. *Economic Journal*, 127(601):653–687.

Borusyak, K. and Hull, P. (2020). Non-random exposure to exogenous shocks: Theory and applications. *NBER Working Paper 27845.*

Borusyak, K., Hull, P., and Jaravel, X. (2021). Quasi-experimental shift-share research designs. *Review of Economic Studies (Forthcoming).*

Caldera, A. and Johansson, Å. (2013). The price responsiveness of housing supply in oecd countries. *Journal of Housing Economics*, 22(3):231–249.

De Haan, J. and Diewert, E. (2013). Hedonic regression methods. *Handbook on Residential Property Price Indices (OECD).*

Degen, K. and Fischer, A. M. (2017). Immigration and swiss house prices. *Swiss Journal of Economics and Statistics*, 153(1):15–36.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Gyourko, J., Mayer, C., and Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–99.

Jaeger, D. A., Ruist, J., and Stuhler, J. (2018). Shift-share instruments and the impact of immigration. *NBER WP 24285.*

Moallemi, M. and Melser, D. (2020). The impact of immigration on housing prices in australia. *Papers in Regional Science*, 99(3):773–786.

Redding, S. J. (2020). Trade and geography. *NBER WP 27821*.

Sá, F. (2015). Immigration and house prices in the uk. *The Economic Journal*, 125(587):1393–1424.

Saiz, A. (2003). Room in the kitchen for the melting pot: Immigration and rental prices. *Review of Economics and Statistics*, 85(3):502–521.

Saiz, A. (2007). Immigration and housing rents in american cities. *Journal of Urban Economics*, 61(2):345–371.

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3):1253–1296.

Saiz, A. and Wachter, S. (2011). Immigration and the neighborhood. *American Economic Journal: Economic Policy*, 3(2):169–88.

Sharpe, J. (2019). Re-evaluating the impact of immigration on the us rental housing market. *Journal of Urban Economics*, 111:14–34.

von Ehrlich, M., Schöni, O., and Büchler, S. (2018). On the responsiveness of housing development to rent and price changes: Evidence from switzerland. *Study on behalf of State Secretariat for Economic Affairs (SECO).*

Appendix

Consistency

In the framework of BHJ, consistency requires sufficiently independent shocks, each with relatively small average exposure s_c in order for a law of large numbers to apply. Figure A1 reveals that some nationalities have relatively high (normalized) average exposure shares across municipalities in the cross-section, which could potentially prevent us from obtaining consistent estimates. This preoccupation, however, is substantially mitigated by correctly evaluating the concentration of shares in relation to the entire sample, including the time-series dimension of the data. This results in the following Herfindahl index of share concentration that should converge to zero to ensure consistency: $\sum_{c=1}^{54} \sum_{t=1}^{14} \widetilde{s}_{ct}^2 \to 0$, with $\widetilde{s}_{ct} \propto s_c$ normalized to add up to one over all panel observations. In our application this index, which is also an inverse measure of the effective sample size, equals an arguably sufficiently low 0.0076. Although serial and cross-sectional shock correlation is not considered in this variant of the Herfindahl index in the sense that it reduces the effective sample size, we presume that independent shock variation is large enough for the application of a law of large numbers. This may particularly hold after considering first differences, conditioning extensively on fixed effects and country-level controls and by appealing to the idea that immigration waves from specific countries or regions tend to die out relatively quickly over time.

Figure A1: Country exposure shares as a fraction of local foreign populations in 2001 averaged over municipalities $\left(=\frac{s_c}{\sum_{c=1}^{54} s_c}\right)$