Shift-Share Instruments and the Impact of Immigration

– preliminary –

Joakim Ruist

Gothenburg University

Jan Stuhler

Universidad Carlos III de Madrid,

SOFI, CReaM, and IZA

David A. Jaeger

CUNY Graduate Center, Universität zu Köln,

CReAM, CESifo, IZA, and NBER

April 2017

Acknowledgements: Jan Stuhler acknowledges funding from the Spanish Ministry of Economy and Competitiveness (MDM2014-0431 and ECO2014-55858-P), and the Comunidad de Madrid (MadEco-CM S2015/HUM-3444). We thank Michael Amior, George Borjas, Christian Dustmann, Tim Hatton, Joan Llull, Marco Manacorda, Simen Markussen, Joan Monras, Elie Murard, Barbara Petrongolo, Uta Schoenberg, JC Suarez Serrato and seminar participants at the Universidad Autonoma de Barcelona, London School of Economics, Colegio Carlo Alberto, Duke University, Queen Mary University, Royal Holloway University, Gothenburg University, the Norwegian School of Economics in Bergen, the Helsinki Center of Economic Research, the Frisch Centre in Oslo, the University of Navarra, the Institute for the Study of Labor in Bonn, the 2017 PSE-CEPII Workshop on the Migration, and the Milan Labor Lunch Series for comments.

Shift-Share Instruments and the Impact of Immigration on Wages

Abstract

Many studies in the immigration literature rely on geographic variation in the concentration of immigrants to identify their impact on the labor market. National inflows of immigrants are interacted with their past geographic distribution to create an instrument, in the hopes of breaking the endogeneity between labor market conditions and the location choice of immigrants. We present evidence that estimates based on this *shift-share* instrument are subject to bias from the conflation of short- and long-run responses to local shocks. The bias stems from the interplay of two factors. First, local shocks may trigger adjustment processes that gradually offset their initial impact. Second, the spatial distribution of immigrant inflows typically changes little over time. In the U.S., both the country-of-origin composition and spatial distribution of immigrant arrivals have been almost perfectly serially correlated in recent decades, with the same cities repeatedly receiving large immigrant inflows. Estimates based on the conventional shift-share instrument are therefore unlikely to identify a causal effect. We propose a "double instrumentation" solution to the problem that — by isolating spatial variation that stems from *changes* in the country-of-origin composition on the national level — produces estimates that are likely to be less biased than those in the previous literature. Our results are a cautionary tale for a large body of empirical work, not just on immigration, that rely on shift-share instruments for causal identification.

Studies on the labor market impact of immigration are often based on spatial variation of immigrant inflows across areas. Typically, inflows at the aggregate level are combined with the lagged geographic distribution of immigrants to create an instrument, in the hopes of addressing the endogeneity of their location choices with respect to local labor demand (Altonji and Card 1991, Card 2001). With dozens of publications in leading journals in the last decade, this "past-settlement" instrument is a crucial element in the "spatial correlation" literature on immigration, and has been used to identify supposedly exogenous labor supply shocks also for other questions of interest. Moreover, it is a prominent example for a category of instrumental variables that share the same underlying rationale – combining *local* economic compositions with shifts on the *aggregate* level to predict spatial variation in a variable of interest. These "shift-share" instruments have become popular in a wide range of literatures and, in a quest for better identification, have introduced spatial variation also in settings that traditionally relied on time-series analysis.[1]

Despite a proliferation of studies, the past settlement instrument has not resolved a long-standing dispute regarding the labor market effects of immigration or, more generally, how local labor markets adjust to supply shocks (see, for example, Borjas 2014 and Card and Peri forthcoming). Estimates of the wage impact that rely only on the past settlement instrument tend to be less negative than those from the factor proportions approach, or those that rely on natural experiments that produce exogenous inflows of immigrants (see, for

[1] For example, Bartik (1991) combines the local industry composition with national changes in employment across industries to isolate local labor demand shock. Kovak (2013) interacts the local industry composition with tariff changes to examine the impact of trade reform, while Autor, Dorn, and Hanson (2013) interact consider aggregate trade flows to examine the impact of Chinese imports on labor markets in the US. Nunn and Qian (2014) interact the probability that a country receives aid with time variation in US food shipments to study their effect on civil conflicts. Shift-share instruments are also central in a surging literature on local fiscal multipliers (e.g. Nakamura and Steinsson 2012, Wilson 2012).

example, Aydemir and Kirdar 2014; Llull 2014; Dustmann, Schoenberg, and Stuhler forthcoming; and Monras 2015). Moreover, the magnitude and even sign of estimates from the spatial correlation approach appear more variable than estimates from alternative empirical approaches (Dustmann, Schoenberg and Stuhler 2016), and may change sign even when applied to different time periods within the same country (Borjas 1999).

We suggest that these inconsistencies in the literature arise partly from the conflation of short- and long-run responses to local supply shocks. The problem stems from the interplay of two factors. First, local supply shocks may trigger prolonged general equilibrium adjustments that gradually offset their initial local impact. A region hit by a local shock may eventually experience positive wage growth, and such regional adjustments may take a decade or more (Blanchard and Katz 1992, Eberts and Stone 1992, Greenaway-McGrevy and Hood 2017). Second, the origin-composition and settlement patterns of immigrants are correlated over time. This applies in particular to the U.S., which due to its large area appears as an attractive setting for the spatial correlation approach. But the origin-composition and settlement patterns have here been almost perfectly serially correlated in recent decades, with the same cities repeatedly receiving large and predictable immigrant inflows. Because of these two factors, the spatial correlation approach tends to conflate the (presumably negative) short-run wage impact of recent immigrant inflows with the (presumably positive) movement towards equilibrium in response to previous immigrant supply shocks.

The spatial correlation literature using the past settlement instrument relies on partial equilibrium adjustments that can be neither too fast, leading to no observed differences across markets (Borjas 1999), or too slow, which can lead to violation of the instrument exogeneity. Interestingly, the latter problem can be worse than the former – it introduces biases that can dominate the short-term impact of current immigration, resulting in a sign reversal and a *positive* estimated effect of immigration on wages. We therefore maintain that the existence of an equilibrium adjustment process poses a problem for estimation of the labor market effect of immigration, regardless of its speed. By placing the past settlement instrument in a theoretical framework, this and other potential violations of the exogeneity of the instrument become clearer than in the "ad-hoc" implementations that are common in the applied literature.

Using data from the U.S. Census and American Community Survey from 1960 to 2011, we illustrate how use of the past settlement instrument exacerbates these biases. Because the country of origin mix of the inflow of immigrants is so similar over time, the correlation between the predicted decadal immigrant inflow rate across metropolitan areas and its lag is consistently high – since the 1980s, the correlation has been 0.96-0.99. As a consequence, the standard instrumental variable approach captures not only the short-term impact, but also the longer term adjustment process to previous inflows. The resulting estimates have no clear interpretation, because the respective weights on the short- and longer term adjustment vary across applications, and because the latter are likely to also affect labor market outcomes in "control" areas. The greatest strength of the past settlement instrument, its impressive ability to predict current flows, can thus turn into a weakness. In some sense, if the instrument is "too strong", it is difficult to believe that it constitutes a shock that is unrelated to the dynamics of the local labor market.

Our results suggest, however, that periods with substantial changes in the country of origin composition provide variation that can be exploited with a variant of the past settlement strategy. We show that a "double instrumentation" procedure, in which both current and past immigrant inflows are instrumented by versions of the past settlement instrument that vary in their national but not local components, isolates an exogenous component of observed inflows that is uncorrelated to local demand *and* past supply shocks. The procedure is demanding, as the consequences of current and past immigrant inflows on the local level can be distinguished only if there is sufficient innovation in their composition on the national level. We show that in the U.S. the enactment of the Immigration and Nationality Act of 1965, which led to a large break in the composition of immigrants (Hatton 2015), provides sufficient variation for its application.

Using this procedure, we estimate that the wage impact of immigration in the 1970s was more negative than estimates based on the conventional shift-share instrument would suggest. However, the estimated impact of the 1960s immigrant inflow on wage growth in the 1970s is positive, and in some specifications of similar magnitude as the negative impact of the 1970s inflow, suggesting that immigration may not have a persistent negative effect on the *relative* local wage level. Innovations in the composition of migrants in the U.S. make the 1970s therefore a particularly interesting case, and similar compositional breaks are observed in other countries. In contrast, U.S. immigrant inflows after 1980, with their persistent country-of-origin composition, are not conducive for such analysis.

The issue that we emphasize is particularly salient for the past settlement instrument and the spatial correlations immigration literature, but in principle extends to many other types of shift-share instrument. Shift-share instruments combine *local* “shares” and *aggregate* “shifts” to generate spatial variation in a variable of interest. An intrinsic issue that we illustrate here is that the local shares are always highly serially correlated, whether constructed from the composition of demographic groups, industries or other characteristics. For shift-share instruments to be valid we thus require one of two conditions to hold: either the national “shifts” are not serially correlated, or the variable of interest does not trigger dynamic equilibrium adjustments in local outcomes. In contexts where there are sudden shocks on the national level, shift-share instruments may meet the required exogeneity assumptions. In others, like the immigration literature, care must be taken to insure that there is sufficient variation over time to plausible interpret the results as causal effects. Variants of the shift-share methodology, such as the one proposed here, can be used to isolate spatial variation that is uncorrelated with the spatial distribution of past shocks.

I. Spatial Correlations and the Past Settlement Instrument

By number of publications, the spatial correlation approach is the dominant identification strategy in the economic immigration literature, and its central identification issue is the selection problem.[2] Immigrants do not randomly sort into labor markets, but rather are attracted to areas with favorable demand conditions (Jaeger 2007). A simple comparison between high- and low-immigration areas may therefore yield an upward-biased estimate of the impact of immigration. The problem is notoriously difficult to solve and arises even in those cases in which natural experiments generate exogenous variation in immigrant inflows at the national level.

To address the selection problem, a large number of studies exploit the observation that immigrants tend to settle into existing cities with large immigrant populations. This tendency, noted in Bartel (1989) and Lalonde and Topel (1991), was first exploited by Altonji and Card (1991) to try to identify the causal impact of immigration on natives' labor market outcomes. Altonji and Card use only the geographic distribution of all immigrants. Card (2001) refined this instrument by noting Bartel's observation that immigrants locate near previous immigrants *from the same country of origin*. For each labor market, he created a predicted inflow based on the *previous* share of the immigrant population from each country of origin combined with the *current* inflow of immigrants from those countries of origin at the

---

[2] See Peri 2016, Dustmann, Schoenberg and Stuhler 2016, or the report by the National Academy of Science 2016, for recent reviews. The main alternative is to exploit differences in the concentration of immigrants across across skill (e.g. education-experience) groups (Borjas, 2003). The skill-cell approach identifies only relative effects and can be sensitive to the definition of skill groups and other assumptions (see Dustmann and Preston 2012, Borjas 2014; Dustmann, Schoenberg and Stuhler 2016).

national level. The potential advantage of this specification arises from the considerable variation in the geographic clustering of immigrants from different countries of origin. Card's shift-share instrument then is, specifically

$$\widetilde{m}_{jt} = \sum_{o} \frac{M_{ojt^0}}{M_{ot^0}} \frac{\Delta M_{ot}}{L_{ojt-1}}, \tag{1}$$

where M_{ojt^0}/M_{ot^0} is the share of immigrants from country of origin o in location j at reference date t^0, ΔM_{ot} is the number of new arrivals from that country at time t at the national level, and L_{ojt-1} is the local population in the previous period. The expected inflow rate $\widetilde{m}_{jt}$ is therefore a weighted average of the national inflow rates from each country of origin, with weights that depend on the distribution of earlier immigrants at time t^0. We refer to this as the "past settlement instrument", but other terms are used in the literature (e.g. "network," "supply-push," or "enclave instrument"). Like other shift-share instruments the past settlement instrument has intuitive appeal because it generates variation at the *local* level by exploiting variation in *national* inflows, which are arguably less endogenous with regard to the local labor market.[3]

It is difficult to overstate the importance of this instrument for research on the impact of immigration on labor markets. Few literatures rely so heavily on a single instrument or variants thereof. Appendix Table 1 presents a list of articles published in top general and field journals in economics, plus a number of recent papers that perhaps better reflect current usage of the instrument.[4] With around 30 publications in the last decade alone, it is one of the most

[3] Studies vary in their choice of t^0 and how temporally distant it is from *t*. Saiz (2007) predicts national immigrant inflows using characteristics from each origin country to address the potential endogeneity of national inflows to local conditions. Hunt (2012) and Wozniak et al. (2012) remove the area's own inflows from the national inflow rate to reduce the endogeneity to local conditions.

[4] Most studies listed in Appendix Table 1 use a version of the Card (2001) instrument as their main strategy to address the selection bias, although some use the simpler Altonji and Card (1991) variant. Others combine the past settlement instrument with other (mostly distance-

popular instrumental variables in labor economics. While most applications focus on questions related to immigration, authors have begun to use the instrument as a convenient way to generate (potentially exogenous) variation in labor market conditions to examine outcomes like fertility (Furtado and Hock, 2010) or parental time investment (Amuedo-Dorantes and Sevilla, 2014).

The arguments offered in support of the validity of the instrument vary somewhat across studies. A typical motivation is given by Card (2009):

> *"If the national inflow rates from each source country are exogenous to conditions in a specific city, then the predicted inflow based on [Card's] equation (6) will be exogenous."*

Although this statement captures the instrument's intuitive appeal, the term "exogenous" can be misunderstood.[5] The instrument is a function of national inflow rates and local immigrant shares. It may therefore not be exogenous in the sense of satisfying the exclusion restriction required for the instrument to be valid if the shares are correlated with unobserved local conditions, even if the national inflow rates are unrelated to those conditions.

To the best of our knowledge, ours is the first attempt to evaluate the validity of the instrument within a simple model of labor market adjustment, although various concerns have been expressed previously. Borjas (1999) notes that the exclusion restriction necessary for the validity of the instrument may be violated if local demand shocks are serially correlated, leading to correlation between the immigrants shares used in the construction of the instrument and subsequent demand shocks. Pischke and Velling (1997) note that mean

---

based instruments) to increase strength of the first-stage or use the instrument for robustness tests or as a reference point for other identification strategies.

[5] Deaton (2010) argues that a lack of distinction between "externality" (i.e. the instrument is not caused by variables in the outcome equation) and "exogeneity" (validity of the IV exclusion restriction) causes confusion in applied literatures. Such distinction might indeed be useful for the discussion around the past settlement and other shift-share instruments.

revision in local unemployment rates may introduce bias if immigrant shares are correlated with the unemployment rate, and Amior (2016) notes that immigrant shares tend to be correlated with area-specific demand shocks related to the local industry structure.

None of these concerns appear problematic enough, however, to explain the surprisingly varying and sometimes positive estimates produced by using the past settlement instrument to identify the impact of immigration on labor market outcomes. In particular, serial correlation in local labor demand should be addressed if the instrument is constructed using settlement patterns that are sufficiently lagged (e.g. Dustmann, Fabbri, and Preston 2005; Dustmann, Frattini, and Preston 2013; Wozniak and Murray 2012; Orrenius and Zavodny 2015). We argue instead that the past settlement instrument almost surely violates the exogeneity assumption by conflating short- and long-run responses to local shocks. As we show, the common strategy of choosing t^0 to be at a substantially earlier point in time offers no protection because the violation arises not from serial correlation in outcomes, or correlates of the initial immigrant distribution, but from the endogenous response to immigrant inflows themselves.

II. The Past Settlement Instrument and Local Labor Market Adjustments

We examine the validity of the past settlement instrument in a simple model of local labor markets. The core issue can be described in a simple dynamic setting, in which local labor markets adjust in response to spatial differentials in current economic conditions. We first study concerns raised in the previous literature, and proposed solutions, and then turn towards problems that stem from the prolonged response of labor markets to local demand and immigration-induced supply shocks.

Output in labor market j at time t is given by

$$Y_{jt} = \theta_{jt} K_{jt}^{\alpha} L_{jt}^{1-\alpha}, \tag{2}$$

where L_{jt} is labor, K_{jt} capital, θ_{jt} is local total factor productivity and α is capital's share of output. Labor is paid its marginal product such that

$$log w_{jt} = \log(1-\alpha) + log\theta_{jt} + \alpha log k_{jt}, \tag{3}$$

with $k_{jt} = K_{jt}/L_{jt}$ denoting the capital-labor ratio. If in the long run capital is perfectly elastically supplied at price r, the optimal capital-labor ratio will be

$$log k_{jt}^{*} = \frac{1}{1-\alpha}\log\left(\frac{\alpha}{r}\right) + \frac{1}{1-\alpha} log\theta_{jt}. \tag{4}$$

It will be affected by the local productivity level θ_{jt} but, because of the constant returns to scale assumption inherent in the production technology, not by the local labor aggregate L_{jt}.

The local labor aggregate consists of natives, N_{jt}, and immigrants, M_{jt}. The inflow of newly-arrived immigrants as a share of overall employment in the local labor market is therefore

$$m_{jt} = \Delta M_{jt}/L_{jt-1}. \tag{5}$$

Assuming that the spatial distribution of immigrant arrivals is partly determined by the distribution of previous immigrants and partly by currently local demand conditions, we can decompose this flow as

$$m_{jt} = \lambda \underbrace{\sum_{o} \frac{M_{ojt-1}}{M_{ot-1}} \frac{\Delta M_{ot}}{L_{jt-1}}}_{past\ settlement\ pull} + (1-\lambda) \underbrace{\frac{f(log w_{jt})}{\sum_j f(log w_{jt})} \frac{\Delta M_t}{L_{jt-1}}}_{economic\ pull}$$

where $0 \leq \lambda \leq 1$ measures the importance of existing enclaves relative to local economic conditions, as captured by $f(log w_{jt})$ with $f' > 0$. If $\lambda < 1$ we are therefore faced with the selection problem – immigrants prefer to locate in areas with with favorable demand condition. Our formulation implies that immigrants may be responsive to (relative) wage growth, such that OLS estimates of their wage impact will be biased upward even when the dependent variable is wage growth instead of wage levels. Adding a noise term to allow for unobserved heterogeneity across cities would not affect our argument.

The Local Adjustment

A key, but often not explicitly discussed, issue for the spatial correlation literature is the local adjustment process – in particular the response of other factors of production – triggered by immigrant-induced local labor supply shocks.[6] The main concern in the literature is that if other factors adjust quickly, the observed impact of immigration at the local level may not represent the overall impact at the national level. The longer the time elapsed between the supply shock and measurement, the less likely the data will uncover any impact of immigrants on local wages (Borjas 1999). Researchers therefore assume that estimates exploiting the spatial distribution of immigrants are biased towards zero (e.g. Borjas 2006, Cortes 2008), or argue that only limited spatial adjustments occur in their period of study.

However, research on regional evolutions in the U.S. concludes that spatial adjustments may take around a decade or more (e.g. Blanchard and Katz 1992, Ebert and Stone, 1992, Greenaway-McGrevy and Hood, 2016). Recent evidence from the migration literature points likewise to a prolonged adjustment period (e.g. Monras 2015, Borjas 2015, Amior and Manning 2015, Braun and Weber 2016, Edo 2017), and it has been observed that local wages remain depressed long after other types of shocks (e.g. Autor, Dorn, Hanson 2016).

We therefore show that even if adjustments to local shocks occur slowly, the assumptions necessary for the past settlement instrument to identify the causal effect of

[6] Labor supply shocks may affect capital flows (Borjas, 1999) and internal migration (Card, 2001; Dustmann *et al.,* 2015; Amior and Manning, 2015), but may also affect human capital accumulation (Smith, 2012; Hunt, 2012), the production technology of firms (Lewis, 2011; Dustmann and Glitz, 2015), or occupational choice (Peri and Sparber, 2009).

immigration will not be met if the instrument is serially correlated. Adjustment processes could take different forms (e.g. Greenaway-McGrevy and Hood, 2016) and the relative importance and speed of individual channels, such as internal migration, is disputed (e.g. Card 2001, Borjas 2014). To illustrate our point it however suffices to consider a single response function or *error correction model (ECM)* that abstracts from the channel of adjustment. Specifically, assume that the local capital-labor ratio does not equilibrate immediately in period *t*, but rather adjusts sluggishly according to

$$logk_{jt} = logk_{jt-1} - m_{jt} + \gamma\left(logk^{*}_{jt-1} - logk_{jt-1}\right). \tag{7}$$

The capital-labor ratio declines in response to immigrant inflows but, barring any subsequent shocks, returns to the optimal level over subsequent periods. The coefficient γ measures the speed of this convergence. As we use decadal data (i.e. the average migrant has entered five years before measurement), the assumption $\gamma \approx 1$ might not be implausible, but our argument also holds if the convergence process is slow ($0 < \gamma \ll 1$), if it begins immediately in period t, if is triggered already by the expectation of immigrant inflows, or if the recovery is only partial (e.g. Bartik 1991 notes that local shocks may have long run effects by affecting human capital accumulation).

We therefore explicitly allow for a local labor market to be in disequilibrium. The error correction model given by Equation (7) allows simultaneously for wages to respond to a contemporaneous labor supply shock and for labor market dynamics in form of a lagged disequilibrium term. A similar error correction model is described and motivated by Amior and Manning (2015) for the case of population dynamics in response to labor demand shocks.

While the specific mechanisms or timing are less important, the degree to which the adjustment process in area j affects wages in *other* areas will affect the interpretation of our empirical results. For example, the capital-labor ratio may adjust either because of capital inflows or native internal migration ($logk_{jt} = logK_{jt} - logL_{jt}$ and thus $\Delta logk_{jt} =$

$\Delta logK_{jt} - \Delta logL_{jt}$). An important distinction between the two channels is that internal migration – population movements from one area to another – is necessarily spatial, while it is less obvious if the accumulation of capital in one area affects its supply in others. To fix ideas we thus decompose the overall adjustment coefficient γ into

$$\gamma = \gamma_K + \gamma_L \tag{8}$$

where γ_K captures the importance of *internal* adjustment processes (such as local savings and investment) while γ_L represents the importance of spatial or *external* adjustment processes (such as migration between areas).

The Selection Problem

In this model the past settlement instrument addresses the selection problem, if combined with a first-differenced specification in wages.[7] To illustrate, assume that the capital-labor ratio is at its optimum for all areas in period 0 and in period 1 there are different immigrant inflows to each area. From equations (3) and (7), the wage level in labor market *j* equals

$$logw_{j1} = \log(1-\alpha) + log\theta_{j1} + \alpha\left(logk_{j0}^{*} - m_{j1}\right) \tag{8}$$

and a regression of first-differenced wages $\Delta logw_{j1}$ on immigrant inflows m_{j1} instrumented by the past settlement instrument $\widetilde{m}_{j1}$ has

[7] The past settlement instrument is unlikely to address selection in wage levels. OLS estimates are biased by non-random sorting of recent *arrivals* with respect to wage levels, but IV estimates would suffer from non-random sorting of immigrant *stocks*. There is little reason to expect that the latter is much less of a concern, in particular since the past settlement instrument suggests a close relationship between stocks and new arrivals, and spatial differences in wage levels are persistent (Moretti 2011). Most of the literature uses first-differenced or fixed-effect specifications (e.g. Dustmann et al. 2005).

$$plim\ \tilde{\beta}^{IV}_{t=1} = \frac{Cov(\tilde{m}_{j1}, \Delta log w_{j1})}{Cov(\tilde{m}_{j1}, m_{j1})} = -\alpha + \underbrace{\frac{Cov(\tilde{m}_{j1}, \Delta log\theta_{j1})}{Cov(\tilde{m}_{j1}, m_{j1})}}_{demand\ shocks} \tag{9}$$

where the covariance terms represent their population values. The asymptotic bias term in equation (9) illustrates a key concern about the past settlement instrument (e.g. Borjas 1999, Hunt and Gauthier-Loiselle 2010, Aydemir and Borjas 2011, Dustmann and Glitz 2015). If productivity or other labor demand shifts are serially correlated (Amior and Manning 2015), then past immigrant inflows and thus the instrument might be correlated with demand shifts in the current period. Common solutions in the literature are to test for serial correlation in the residuals of the wage regression (Dustmann, Frattini and Preston 2013) or to lag the base period t^0 sufficiently aback, as to minimize the potential that the instrument is correlated with current demand shifts. If, in addition, the flow of immigrants by country of origin at the *national* level are unaffected by current local labor demand conditions, the instrument will be uncorrelated with current demand shifts. Since our concern is not about time dependence in *external* processes we abstract from this issue by assuming that $log\theta_{jt}$ follows a random walk.

The Overlapping Response Problem

Our fundamental concern is that even in the absence of serial correlation in external processes, serial dependence is generated *endogenously* by immigration inflows. The past settlement instrument violates the exogeneity condition because of the interplay of two factors.

First, local shocks lead to general equilibrium adjustment processes that may gradually offset their initial local impact, such that a negative wage response is succeeded by recovery and positive wage growth. As described above, such adjustments can plausibly extend over more than one decade. Variables constructed from the U.S. census data commonly capture arrivals in the preceding decade, such that the average migrant has entered the U.S. about five years before the measurement of wages. Part of the local adjustment, in particular the recovery of wages, may plausibly occur after five years and thus in the *next* period. Second, the spatial distribution of immigrant inflows in the U.S. is highly serially correlated. The past settlement instrument aggravates this issue, as it is motivated by the very idea of serial correlation in immigrant inflows. The instrument isolates that part of the variation in current inflows that is predictable by past stocks and thus past cumulative inflows up to time t^0.

Together, these observations imply that the short-term response to new immigrant arrivals overlaps with the lagged response to past immigrant inflows – and that the standard IV estimator used in the literature conflates these short- and long-term responses. This "*overlapping response hypothesis*" mirrors arguments from the recent literature on labor demand shocks, which argues that persistent trends in labor demand can trigger important population dynamics on the local level, and that this persistence needs to be accommodated for if one wishes to estimate the response of labor markets to local demand shocks (Amior and Manning 2015, Greenaway-McGrevy and Hood 2016). Trends in immigration-induced labor supply can be even more persistent, suggesting that such arguments are important also for the related literature.

We can use our model to illustrate the resulting bias and its properties. Equation (9) showed a special case that abstracted from the problem, as the local market was assumed to be in steady state when an unexpected immigration inflow occurred in $t = 1$. This assumption is implicitly made also in previous studies. But in the next period, a regression of first-differenced wages on instrumented immigrant inflows would yield

$$plim\tilde{\beta}^{IV}_{t=2} = -\alpha + \underbrace{\frac{\alpha\gamma}{1-\alpha}\frac{Cov(\tilde{m}_{j2}, \Delta log\theta_{j1})}{Cov(\tilde{m}_{j2}, m_{j2})}}_{lagged\ demand\ shocks} + \underbrace{\alpha\gamma\frac{Cov(\tilde{m}_{j2}, m_{j1})}{Cov(\tilde{m}_{j2}, m_{j2})}}_{lagged\ supply\ shocks} \tag{10}$$

The two new bias components arise from the endogenous response of the capital-labor ratio to local shocks in the previous period. First, it responds to past local *demand* shocks that occurred in $t = 1$. These are potentially correlated with the instrument in the current period ($Cov(\widehat{m}_{j2}, \Delta log\theta_{j1}) > 0$) if immigrants are attracted to areas with growing labor demand. Second, the capital-labor ratio responds to the immigration-induced *supply* shock that occurred in the previous period. Either response raises the marginal productivity of labor, and therefore wages, leading to an upward bias in our estimates.

The two bias terms are endogenously generated and arise independently from the assumed time series properties of the local demand shocks θ_{jt}. The first term illustrates that demand shocks can generate bias even if they are not serially correlated. Intuitively, if local shocks trigger local adjustments, immigrant shares must not only be uncorrelated with *current* but also with *past* demand shocks. Choosing t^0 to be temporally distant may therefore be advantageous even if demand shocks are not serially correlated. As this is a common strategy in the literature, we assume below that the instrument $\widehat{m}_{jt}$ is sufficiently lagged and uncorrelated to (the current adjustment to) past demand shocks.

The supply-side bias is harder to address. Its size in $t = 2$ depends on the ratio $Cov(\widehat{m}_{j2}, m_{j1})/Cov(\widehat{m}_{j2}, m_{j2})$, which is the slope coefficient in a regression of past on current immigrant inflows, using past settlement shares to instrument current inflows. This coefficient will be small if the past settlement instrument is a substantially better predictor for current immigrant inflows in area j than inflows in the previous period. As we will show, this is unfortunately often not the case in the U.S. context. Instead, the coefficient fluctuates around and is sometimes larger than 1: while the past settlement instrument is a good predictor for immigrant inflows in the intended period, it is also a similarly good predictor for immigrant inflows in previous periods. Importantly, choosing t^0 to be temporally distant does not address this bias.[8]

The size of the supply-side bias in $\tilde{\beta}_t^{IV}$ equation (10) is proportional to the speed of convergence γ. However, in a more general setting with repeated immigrant inflows, this speed may have little influence on the size of the bias. The regression of first-differenced wages on instrumented immigrant inflows in period t has (see Appendix A.x for derivation)

$$plim\tilde{\beta}_t^{IV} = -\alpha + \underbrace{\alpha\gamma\sum\nolimits_{s=0}^{t}(1-\gamma)^s\frac{Cov(\widetilde{m}_{jt}, m_{jt-1-s})}{Cov(\widetilde{m}_{jt}, m_{jt})}}_{lagged\ supply\ shocks}, \tag{11}$$

such that the size of γ will matter little if the predictable component of immigrant inflows is highly serially correlated. In the extreme case, if the past settlement instrument predicts immigrant inflows in all past periods equally well, expression (11) simplifies (because $lim_{t\to\infty}\gamma\sum_{s=1}^{t}(1-\gamma)^s = 1$) to

$$plim\tilde{\beta}_t^{IV} = -\alpha + \alpha\underbrace{\frac{Cov(\widetilde{m}_{jt}, m_{jt-1})}{Cov(\widetilde{m}_{jt}, m_{jt})}}_{lagged\ supply\ shocks}, \tag{12}$$

which does not depend on the speed of local convergence γ. Intuitively, it does not matter if an ongoing local adjustment process has been triggered by immigrant inflows in the previous or an earlier period if both are equally correlated with our instrument. With few exceptions, the serial correlation in immigrant inflows is so extraordinarily high in the U.S., even using changes over a decade, that the speed local of convergence may matter little in practice.[9]

---

[8] Lagging the instrument further aback may reduce the numerator in the ratio $Cov(\widehat{m}_{j2}, m_{j1})/Cov(\widehat{m}_{j2}, m_{j2})$ but, by reducing its ability to predict inflows in the intended period, also the denominator. In principle, the bias may intensify if the denominator shrinks more strongly than the numerator. In the U.S. Census, the ratio is insensitive to the choice of base period t^0.

[9] What does however matter is the assumption that in the long run, immigrant inflows have no persistent effect on local relative wages. If the local recovery is only partial, the size of the supply-side bias in equation (12) would shrink proportionally. If immigration has instead a positive long-run effect on local wages (e.g. via agglomeration and density externalities, Peri

The supply-side bias alone can thus turn the IV estimate of the impact of immigration from negative to positive. As the bias is proportional to the true wage impact of immigration (in our model given by $-\alpha$), this conclusion holds even when the true wage impact is strongly negative. OLS estimates suffer from selection bias, but are less affected by the overlapping response bias if the actual inflows m_{jt} vary more than their predictable component $\widetilde{m}_{jt}$ across decades (as they do in the U.S. Census), as this would reduce the final term in the expression corresponding to equation (10). A priori it is therefore not clear if IV estimates are more accurate than OLS estimates.

The Overlapping Response Problem with Anticipation

We so far assumed that immigrant inflows occur as a "shock", to which local markets respond only in hindsight. However, if these inflows occur repeatedly, and repeatedly in the same areas, their arrival might be anticipated. For example, firms or workers in Los Angeles experiencing steady inflows of Mexicans during the 1970s may have expected further Mexican inflows in the 1980s.

The idea that labor markets adjust in anticipation, and thus concurrently or even before a demand or supply shift actually occurs, is explored already in Topel (1986). But the role of expectations has received less attention in the spatial correlation literature, and it is hard to judge how sophisticated expectations are, or how strongly households and firms will respond. Immigrant arrival rates across cities in the U.S. are so stable and thus so predictable some degree of anticipation seems likely, but that firms and workers may not necessarily respond to anticipated arrivals. For example, Eberts et al. (1992) argue that the assumption that households move years in advance of an anticipated demand shocks (as made in Topel 1986) is not realistic.

2016), the bias increases accordingly.

We will consider two cases here that, together with our baseline case in which anticipation plays no role, may plausibly bound the truth. In the first version the expected inflow of migrants in the next period is equal to the current rate, i.e. $E[m_{jt+1}] = m_{jt}$. In the second version agents combine the observed composition of immigrants in the city with a correct forecast of the national inflow in the next period, i.e. $E[m_{jt+1}] \cong \tilde{m}_{jt+1}$. In the first model agents are naive, simply extrapolating from the current to the next period. In the second they predict as well as an econometrician armed with census data. The truth is plausibly in between.

If the capital-to-labor ratio responds similarly to anticipated and realized shocks, then the error correction model changes from equation (7) to

$$logk_{jt} = logk_{jt-1} - m_{jt} + \gamma\left(logk^{*}_{jt-1} - logk_{jt-1} - E[m_{jt}]\right). \qquad (7')$$

The "naive" expectation $E[m_{jt+1}] = m_{jt}$ would not affect the probability limit given in equation (9), but equation (10) would change to

$$plim\tilde{\beta}^{IV}_{t=2} = -\alpha + \ \ldots \ + 2\alpha\gamma\frac{Cov(\tilde{m}_{j2}, m_{j1})}{Cov(\tilde{m}_{j2}, m_{j2})} \qquad (10')$$

The bias from a response to the supply shock is now twice as large, because the capital-labor ratio responds both to the immigrant inflow in *t=1* as well as to the expected inflow in *t=2*, and the latter is equal to the former. With the "sophisticated" expectation $E[m_{jt+1}] = \tilde{m}_{jt+1}$, already the estimates in *t=1* would be affected, and equation (10) would instead change to

$$plim\tilde{\beta}^{IV}_{t=2} = -\alpha + \ \ldots \ + \alpha\gamma\frac{Cov(\tilde{m}_{j2}, m_{j1})}{Cov(\tilde{m}_{j2}, m_{j2})} + \alpha\gamma \qquad (10'')$$

The bias is similar in both anticipation models if $Cov(\tilde{m}_{j2}, m_{j1}) \approx Cov(\tilde{m}_{j2}, m_{j2})$.

Extending these arguments to a generic period *t* shows that under either anticipation model, the bias term is largest in the period after a structural break in the distribution of immigrants occurs, when the lagged response to the unexpected immigrant inflow in the previous period coincides with the anticipatory response to updated beliefs about their distribution in the future.

Interpretation of Conventional IV Estimator

How should estimates from the conventional IV estimator then be interpreted? According to equation (11), they capture a weighted average of the short- and long-run response of local relative wages to immigration, which depends on two sets of weights. The first set depends on the degree to which the instrument predicts current vs. past immigrant inflows. This is context-specific, so the estimator $\tilde{\beta}_t^{IV}$ will weight the short- and long-term response differently in different applications. The second set of weights depends on the degree to which local wage recovery ($\gamma = \gamma_L + \gamma_K$) stems from internal adjustment processes ($\gamma_K > 0$) or spatial spillovers such as internal migration that affect wages also in other areas ($\gamma_L > 0$). If part of the adjustment is spatial, then the long-run wage impact of immigration on area j as partially captured by $\tilde{\beta}_t^{IV}$ represents only a *relative* effect in relation to other areas which themselves are indirectly affected by immigration, not the long-run effect of immigration on the overall economy. In other words, while the long-run effect of immigration on the host economy is of prime interest, conventional spatial correlation estimates are unlikely to be informative about it.

For both these reasons, the estimator $\tilde{\beta}_t^{IV}$ is hard to interpret. The aim of spatial correlation studies is typically to estimate the short-run local wage impact of immigration before spatial adjustments occur, such that the local reflects the national impact. From this perspective, the conventional estimator $\tilde{\beta}_t^{IV}$ is biased. Even if our aim is to estimate only the impact on immigration on local *relative* wages, the estimator has the undesirable property that it weights the short- and long-run impact differently across applications.

III. Revising the Past Settlement Instrument

Our model illustrates the difficulty of consistently estimating the labor market impact of immigration using the past settlement instrument. In the presence of prolonged spatial adjustment following local labor market shocks, we require an instrument that

- does not correlate with contemporaneous and past demand shocks,
- explains the locational choices of immigrants, and
- is uncorrelated to their choices in the previous period.

The last two conditions are testable, while in the absence of information on demand shifts the first requires a theoretical argument. The past settlement instrument potentially satisfies the first condition if we choose t^0 to be sufficiently in the past and quite clearly satisfies the second condition, so the crucial problem is its correlation to past supply shocks.

This issue can be addressed in various ways. First, in periods in which the country of origin composition of migrants changes strongly, the past settlement instrument will be less correlated with past supply shocks, and estimates based on the past settlement instrument should be less biased. We explore this hypothesis in our empirical analysis. Second, the bias from overlapping responses is also reduced in settings in which the national inflow rate is temporarily increased (as in Gonzalez and Ortega, 2011). Third, one can exploit origin-specific "push factors" that led to a change in national inflows of a particular origin group, as recently done by Aydemir and Kirdar (2013), Llull (2014), Monras (2015), Chalfin (2015), and Carpio and Wagner (2015). While the use of push factors is motivated by the desire to break the potential endogeneity of national inflows to local conditions – for example, more Mexicans may enter the United States if the California labor market is strong – they may under some conditions, also address the overlapping response problem. Specifically, the overlapping response bias in equation (11) can be eliminated if the push factor triggers immigrant flows that are uncorrelated to previous inflows.

But such exogenous push factors are unfortunately rare. We propose therefore to consider all arrivals, but to isolate innovations in local immigrant inflows that are uncorrelated with past inflows. Intuitively, this can be accomplished by first regressing the past settlement instrument $\widetilde{m}_{jt}$ on its lag $\widetilde{m}_{jt-1}$, and then using the residual from this regression to instrument current immigrant inflows. In practice it is more useful to directly use both current and lagged instrument in our wage regression (which yields the same coefficient on current inflows). Specifically, we regress local wage growth on both current and past immigrant inflows,

$$\Delta log w_{jt} = \beta_0 + \beta_1 m_{jt} + \beta_2 m_{jt-1} + \varepsilon_{jt}\ , \tag{13}$$

and instrument the two endogenous variables by

$$\widetilde{m}_{jt} = \sum_o \frac{M_{ojt^0}}{M_{ot^0}} \frac{\Delta M_{ot}}{L_{ojt-1}} \text{ and } \widetilde{m}_{jt-1} = \sum_o \frac{M_{ojt^0}}{M_{ot^0}} \frac{\Delta M_{ot-1}}{L_{ojt-2}}\,. \tag{14}$$

where immigrant stocks by country of origin are measured at the same reference date t^0 for both instruments.

This "double instrumentation" addresses two distinct problems. The instrumentation of m_{jt} by $\widetilde{m}_{jt}$ addresses the selection problem. The inclusion of m_{jt-1} and its instrumentation by $\widetilde{m}_{jt-1}$ addresses the overlapping response problem. Other, seemingly more direct strategies to control for past economic conditions do not suffice. Controlling directly for actual immigrant inflows m_{jt-1} without instrumentation by $\widetilde{m}_{jt-1}$ would introduce a mechanical relationship to local demand shocks.[10] And lagging the instrument further aback, a common

[10]Note that the residual from a regression of the past settlement instrument on past immigrant inflows is equal to

$$\hat{\epsilon}_{jt} = \widetilde{m}_{jt} - a - b m_{jt-1},$$

where $b > 0$ is the slope coefficient in a regression of $\widehat{m}_{jt}$ on m_{jt-1}. However, m_{jt-1} depends strategy for other reasons, does not address the overlapping response problem. The validity check recently proposed by Peri (2016) – to test if the past settlement instrument is correlated with lagged wage growth – while otherwise useful, would not reliably detect the overlapping response problem. The absence of such correlation is precisely one of the possible consequences when the short-run wage impact and longer-term wage recovery to immigrant inflows overlap.[11] Controlling for past wage growth in the wage regression does not suffice for the same reason.

Our model provides predictions on the signs and relative magnitudes of coefficients in the estimating equation (13). The coefficient β_1 captures the wage impact of immigration in the short run (what is normally the coefficient of interest in the literature), and is likely negative, while the coefficient β_2 captures the longer term reaction to past supply shocks and expected to be positive.[12] By summing over both we may thus in principle hope to capture the longer-term effect of immigration on local relative wages. But its interpretation is not straightforward; the coefficient β_2 captures the lagged response of local wages in areas that experienced immigrant inflows *relative* to wages in other areas. However, in the long run, immigrant inflows in one area are likely to affect economic conditions in other areas ($\gamma_L > 0$ in our model), such that area comparisons do not capture the overall effect of immigration on the economy.

---

positively on local demand shocks in that period, introducing bias (see equation (10)).

[11] In our model, a regression of lagged wage growth on the past settlement instrument $\widetilde{m}_{jt}$ estimates $\alpha(\gamma \sum_{s=0}^{t} (1-\gamma)^s \, Cov(\widetilde{m}_{jt}, m_{jt-2-s}) - Cov(\widetilde{m}_{jt}, m_{jt-1}))/Var(\widetilde{m}_{jt})$, and the term in brackets can be approximately zero if immigrant inflows are highly serially correlated.

[12] Specifically, in our model β_1 should be equal to $-\alpha$, while β_2 should be positive and – if lagged adjustments are completed within about one decade *or* if immigrant inflows are highly serially correlated – of similar magnitude. However, other frameworks (e.g. with frictions, as in Chassambouli and Peri 2015, or Amior 2016) would predict other magnitudes.

When estimated by two-stage least squares (2SLS), the corresponding first-stage equations are

$$m_{jt} = \pi_{10} + \pi_{11}\widetilde{m}_{jt} + \pi_{12}\widetilde{m}_{jt-1} + u_{jt} \quad (15)$$

$$m_{jt-1} = \pi_{20} + \pi_{21}\widetilde{m}_{jt} + \pi_{22}\widetilde{m}_{jt-1} + v_{jt} \quad (16)$$

Intuitively, the "right" instrument should predict each of the endogenous variables – for example, the immigrant selection equation of our model suggests $\pi_{12} = \pi_{21} = 0$. If we are willing to impose such restrictions we can estimate equation (15) using a systems estimator, with potential efficiency gains compared to the 2SLS procedure. However, a systems estimator would require a structural interpretation of our first stage equation. As immigrant selection may be more complicated than assumed in our model, we present 2SLS estimates as our baseline specification.

If the stock variables at t^0 used for construction of $\widetilde{m}_{jt}$ and $\widetilde{m}_{jt-1}$ are the same, the difference between the two instruments comes only from time variation in the composition of national inflows. Card's (2001) decomposition into country of origin groups is therefore essential, while the simpler variant of the instrument used by Altonji and Card (1991) would not isolate innovations in supply at the local level. However, the instruments will still be highly correlated if the composition of national inflows change little from one period to the next.[13] While the "double instrumentation" procedure in equations (13) through (16) addresses both the selection and the overlapping-response bias in theory, it may not work in finite samples. Whether the procedure is feasible in practice must therefore be demonstrated in each context.

IV. Data and Descriptive Statistics

[13] If the national percentage changes in the population from each country of the origin are the same from one period to the next, Card's (2001) instrument reduces to Altonji and Card's (1991) instrument.

We use data from the 1960-2000 U.S. Censuses and the merged 2007-2011 American Community Surveys (ACS), all obtained through IPUMS (Ruggles, *et al.* 2015). For convenience we will refer to the merged ACSs as the year 2010.[14] We define an immigrant as a person born in a country other than the U.S. (excluding outlying U.S. territories) and a newly-arrived immigrant as a foreign-born person that immigrated during the last decade. We divide immigrants into 39 countries and regions of origin.[15] In descriptive results that use data that goes back to the 1940 Census, we use the same 17 countries and regions that were used by Card (2001) because of the limited information on countries of origin in those data.

The entire immigrant populations by origin and local area are used in the construction of the past settlement instrument, which is used to instrument immigration rates in the labor force. We conduct our analysis across both metropolitan statistical areas (MSAs) and across commuting zones (CZs). MSAs are the standard unit of analysis in the existing literature and, because of their better comparability, also the baseline unit in our analysis. We include in the analysis all MSAs that can be identified in all Censuses, use data on finer spatial units to make their boundaries as consistent over time as possible, and finally exclude three MSAs in which boundary changes were particularly large between the 1960, 1970, and 1980 Censuses, and for which finer information cannot be used to make them more consistent.[16] This leaves

[14] We use 2007-2011 rather than, for example, 2008-2012, because the MSA definitions changed with the 2012 ACS.

[15] We separately include each country of origin with at least 5,000 observations in the 1990 census, except Cambodia, Iran, Laos, Thailand, and Vietnam, which were not separately coded in all Censuses. All remaining countries of origin are merged into the regions Latin America, Western Europe, Eastern Europe, Asia, Africa, Australia and New Zealand, and Others. Countries that split or merged after 1970 (the USSR, Yugoslavia, Czechoslovakia, and Germany) are coded as the merged unit throughout (e.g. the separate states of the Russian Federation continue to be coded as one unit after the breakup as the USSR, and West and East Germany are merged prior to 1990). Hong Kong and Taiwan are coded as part of China.

[16] These are Bridgeport and New-Haven-Meriden, CT, and Worcester, MA. For all three, their total recorded populations more than triple between the 1960 and 1970 Censuses, and then shrink again by more than two-thirds in the 1980 Census. No other MSA comes close to an

us with a sample of 109 MSAs. The definition of commuting zones is based on Tolbert and Sizer (1996), and applied to Censuses using codes provided by Autor and Dorn (2013).

Our outcome variable is the average log weekly wage among the native labor force in an area. We restrict our wage sample to those who are 18-64 years of age and have 1-40 years of potential experience (age minus expected age at completion of school), and drop who currently attend school, who live in group quarters, or who are self-employed. To reduce the influence of outliers – some wages are as low as, or below, one dollar per week – we drop the bottom and top percentile of wages in each census year. Dropping the top percentile matters little while the choice of cut-off point at the bottom has a non-negligible but, as we show, limited effect on our estimates. To address selectivity bias from changes in the composition of workers we residualize wages using separate national-level regressions for each census year that control for six education levels (high school dropout, high school degree, some college but no degree, bachelor degree, master degree, and professional or doctoral degree), 40 potential experience levels, gender interacted with marital status, three races (white, black, and other), and nine U.S. Census divisions.

We show the characteristics of immigrant inflows by decade in Table 1. Immigration has been high and the immigrant share of the population has risen steadily from 5.2 percent in 1970 to 13.6 percent in 2010. The coefficient of variation of the share of recent arrivals by MSA shrunk by a half over the same period, indicating that immigrants were more geographically dispersed in the earlier decades. We present evidence on the formal and effective skills of immigrant arrivals in Section V.3.

With the Immigration and Nationality Act of 1965, enacted in June 1968, the composition of immigrants changed considerably (Hatton, 2015). Among new arrivals in the 1970 census (i.e. those who arrived in the 1960s, only a minority of which arrived after the

equally problematic pattern in the data.

change in admissions policy was implemented), 41 percent were of Canadian or European origin, whereas in 1980 (those arriving in the 1970s, after the policy change) the corresponding share was only 17 percent. At the same time, the share of Latin Americans and Asians among the newly-arrived rose from 54 percent for those arriving in the 1960s to 75 percent for those arriving in the 1970s. Over the following three decades, there are no correspondingly large compositional changes.

This compositional change is further illustrated in the table's final block, the first row of which reports the correlation from one decade to the next in the shares of all 38 origins (excluding "Other") in the national inflows. The correlation in country of origin shares between those arriving in the 1960s and those arriving in the 1970s is 0.59 while the correlation is between 0.96 and 0.99 in subsequent decades. As shown in the following row, the pattern is the same if we exclude Mexicans, although the correlations are smaller. The last row confirms that there was indeed a unique structural break in immigrant composition in the 1970s. We cannot identify newly-arrived immigrants in the data prior to 1970, so here we show the cross-decade correlations for all immigrants instead. The immigrant stocks in 1970 and 1980 have a correlation coefficient of 0.65, whereas the three earlier pairwise correlations are all above 0.94 and those afterwards are at least 0.90.

In Figure 1 we show the correlations between the recent arrival shares of immigrants arriving in the 1960s and 1970s (top half) and 1970s and 1980s (bottom half). The graphs to the left show all 39 country-of-origin groups and those on the right exclude the outliers Mexico and Cuba. The correlation between the 1960s and 1970s shares is lower without Mexico, as apart from the Mexican share there is large variation in origin shares between the two decades.

V. Estimating the Impact of Immigration

Our data allow us to estimate the wage impact of recent immigrant arrivals in the U.S. for four different decades. As a benchmark, Panel A of Table 2 presents OLS estimates from a regression of the decadal growth in (residualized) log wages of all workers on the immigrant inflow rates across MSAs. Parts of the literature focuses on wage growth among men, for which the results are similar (not reported). Panel B presents the corresponding IV estimates, together with the first-stage coefficient on the instrument. The instrument is the conventional shift-share variable defined in equation (1), where the reference period is the beginning of the decade. The instrument is a strong predictor of immigrant inflows in all decades.

Two observations stand out. First, the estimates are positive for some decades. Selection may generate an upward bias in the OLS estimates and, once we instrument the immigrant inflow rate using the past settlement instrument, the estimates become indeed more negative. However, the differences are modest; the IV estimate for the 1980s is still positive and large (as well as statistically significant). Second, the point estimates differ substantially across the decades. Estimates on the Commuting Zone level, shown in Appendix Table A.2, are similar. Borjas, Freeman and Katz (1997) and Borjas (1999) note that the spatial correlation approach leads to different estimates for the 1970s and 80s, and we find that this variability extends to instrumental variable estimates based on the origin-version of the past settlement instrument, to more recent periods, and to different spatial definitions.

It is only in the 1980 census (i.e. for inflows during the 1970s) that we find a more than marginally negative IV estimate of the effect on wages. As shown in the previous section, this is the period in which a change in the U.S. admission policy created a substantial shift in the composition of immigrant arrivals across origin groups, plausibly making their distribution across MSAs less related to their spatial distribution in the previous decade. We report the correlations between actual immigrant inflows and the past settlement instrument and their respective lags in Panel A of Table 3. As expected, this correlation is lower for immigrant inflows in the 1970s than in the later decades: 0.82 compared to 0.92-0.96. This gap becomes even larger when considering the instrument instead of actual inflows: 0.70 compared to 0.96-0.99.

Serial correlation will therefore be an important issue when estimating the wage impact no matter what decade one considers. Yet in the 1970s at least there is some variation between the decades, whereas the serial correlation in the instrument is nearly perfect in the later decades – immigrant arrivals are predicted to enter again and again the same cities.

Our theoretical argument implies that all the IV estimates in Table 2 are upward-biased. However, together with the observed break in the spatial distribution of immigrant arrivals, it also suggests that this bias should be smallest in the 1980 Census – in which we indeed find the most negative coefficient estimate. As it was caused by the Immigration Act of 1965, the break itself was likely not anticipated. However, workers and firms may have expected that it had a persistent effect on the distribution of immigrant arrivals, and that arrivals as observed during the 1970s were informative about the likely distribution of arrivals in subsequent decades. In this case, the Immigration Act may also explain why the spatial correlation estimates are most positive in the 1990 Census (see Section II).[17]

Based on Equation (10), we can estimate the key bias components. In particular, the bias is proportional to the ratio between the two pair-wise correlations of the instrument and past and current inflows. Since the past settlement instrument draws from the national composition in the period for which it is constructed, we may hope that the numerator of this ratio is substantially smaller than its denominator. As shown in Panel B of Table 3, this is unfortunately not the case: in the later decades, the past settlement instrument is more strongly

[17] The question if workers and firms respond to expectations plays a more important role in this argument than the question how expectations are exactly formed. The spatial distribution of inflows in the 1970s were so similar to the inflows in the 1980s that even a naive extrapolation of the former would accurately predict the latter.

correlated to past inflows than to the current inflows it is supposed to predict. Notably, this is a natural pattern when the national composition changes very little, since past inflows are closer in time to the reference period t^0 used in the construction of the instrument. Lagging this reference period further aback, while weakening the predictive power of the instrument, does not change this pattern, as can be seen from the rows using t-2 as the base period (i.e. constructing the instrument from the immigrant distribution two decades prior to the year of observation).

Some studies in the literature combine spatial variation in immigrant inflows across areas with their density across skill groups.[18] Depending on the outcome variable of interest, the explanatory variable may be the rate of immigration in a particular education group (Cortes, 2008; Hunt, 2012), or the relative skill content of immigration (Card, 2009; Lewis, 2011) in an area. Panel C of Table 3 shows that such measures – we consider the immigration rates of high skilled (at least bachelor degree) and low skilled (less than a bachelor degree) workers, as well as the logarithm of the ratio of high skilled to low skilled immigrants – are likewise highly serially correlated. The serial correlation in the skill-specific inflow rates and instruments is consistently close to the corresponding values of the total rate: modest in the 1970s and high in all later decades. The serial correlation in the log skill ratio is high in all periods. The overlapping response problem affects therefore also those empirical strategies that exploit both area and skill-cell variation.[19]

[18] See Peri (2016) or Dustmann, Schoenberg and Stuhler (2016) for an overview. By using both spatial and skill-cell variation, one can difference out unobserved factors that lead to higher or lower wages of all workers in a city (see Card, 2007). However, only relative wage effects of immigration across skill groups are identified.

[19] However, the magnitude of the problem may be different. The assumption that average wages are mean reverting because labor demand is perfectly elastic in the long run is standard in the literature (even though wage differences between cities are persistent, see Moretti 2010), but differences in local skill-specific wages may be more persistent.

V.1 First-stage Results

To isolate innovations in the immigration inflow rate, we implement the "double instrumentation" procedure proposed in Section III. We illustrated that even in the 1970s, the past settlement instrument is highly correlated with its lag. The key question is therefore if there is sufficient variation to distinguish the two, and we begin by presenting the first-stage results from the 2SLS estimation of equation (13) in Table 4.[20] An important issue is the choice of reference period for the construction of the two instruments. For comparison, we start by using immigrant stocks one decade prior to the year of observation (as in Table 2) for the construction of both instruments.

By and large, the pattern for the 1970s is as hoped (column (1)): the instrument for the 1960s is the main predictor of inflows in that decade, while the 1970s instrument has the largest coefficient estimate for the 1970s inflow – the "right" instrument predicts the right endogenous variable. In contrast, there is little hope to study the wage impact of immigration in the later decades, as illustrated in columns (2) to (4). Because of the high serial correlation in national inflow shares, the two instruments carry almost the same information. This issue is reflected in the Sanderson-Windmeijer statistic, which indicates that the coefficients on the two endogenous variables cannot be separately identified. As a consequence, the coefficient estimates jump from one decade to the next, with little apparent sense in the relative sign of the coefficients in one versus the other first stage equation.

Even for the 1970s some questions remain, as the coefficient for the instrument and its lag in the first stage of the 1970s inflow have nearly the same size. Such pattern is not unreasonable – perhaps new arrivals are attracted to areas that were popular destinations already in previous decades – but it illustrates that parameter restrictions on the first stage

[20] To simplify comparison between first-stage coefficients, we rescale the lagged instrument so that both instruments have the same mean. This has no effect on coefficients in the second stage.

coefficient imposed from a theoretical model might be invalid. As such, we rely on 2SLS instead of systems equation estimation for our baseline results. Moreover, the choice of 1970 as reference period for both instruments is not ideal as it is mechanically related to immigrant inflows during the 1960s, and lagging it aback decreases the coefficient on the current instrument further. This is not an issue for estimation of the second stage, as its slope coefficients depend on the (weighted) difference of the two respective first-stage coefficients, which is large and positive in any specification. But it illustrates that our proposed "double" procedure puts far higher requirements on the data than the conventional procedure with a single instrument, leading to the question if its predictive power could be improved further.

We therefore test a modification of the past settlement instrument that is motivated by the observations that (i) new arrivals from different origin groups cluster to very different degrees, and that (ii) these differences are persistent over time. This pattern is illustrated in Figure A.1, which plots a measure of the clustering of new arrivals in each of our four decades against the same measure in the previous decade. The clustering of an origin group is measured by the slope coefficient from a regression of its actual inflow against its predicted inflow rates across MSAs, i.e. the origin-specific past settlement instrument. This coefficient varies substantially across origin groups, with the spatial distribution of arrivals from some being highly predictable, while others are less likely to settle into existing clusters.

We can use these differences in the propensity to which new arrivals cluster into existing enclaves to further improve the predictive power of the past settlement instrument. Specifically, we construct the "cluster-propensity" version of the instrument

$$\widetilde{m}_{jt}^{p} = \sum_{o} \hat{p}_{ot} \frac{M_{ojt^0}}{M_{ot^0}} \frac{\Delta M_{ot}}{L_{ojt-1}}, \tag{17}$$

where $\hat{p}_{ot}$ is the predicted origin-specific propensity to cluster in period t. Of course, the propensity p_{ot} can be directly estimated, but this propensity could be endogenous to local economic conditions, in particular for origin groups that are concentrated in only few cities.

However, Figure A.1 shows that the origin-specific propensity to cluster is quite stable over time – some groups cluster continually more than others – suggesting that factors other than local demand conditions are also at work. To address the endogeneity concern further we construct two versions of $\widetilde{m}_{jt}^{p}$ that use indirect information to predict p_{ot}. For the first version ("GIV-A"), $\hat{p}_o$ is defined as the average of the estimated propensities to cluster in both periods *t* and *t-1*, thereby reducing any direct relationship with the current period. For the second version ("GIV-B"), $\hat{p}_o$ is defined as the estimated propensity to cluster in the future period *t+1*. The first version of the instrument has the advantage that it treats the two regressors in equation (13) symmetrically. The second version has the advantage that by not drawing from the current period *t* it addresses the endogeneity concern more comprehensively, but it introduces an asymmetry in that $\hat{p}_o$ is constructed from a period that is closer in time to the other components of $\widetilde{m}_{jt}^{p}$ than $\widetilde{m}_{jt-1}^{p}$. Yet another option would be to predict p_{ot} from past variables. In our data and time period of interest, we cannot observe the propensity to cluster in past periods, but we do observe the spatial concentration of past arrivals across MSAs – which is indirectly related to the origin-specific propensity to cluster. For example, we can construct the *dissimilarity index* of each origin group relative to the total population, and Appendix Figure A.2 shows that an origin's group dissimilarity index in 1960 is indeed a strong predictor for the propensity of its 1970s arrivals to cluster. We prefer estimates from the GIV-A and GIV-B specification, as they use more direct information on an origin group's propensity to cluster, but the second-stage estimates remain similar if we construct $\hat{p}_{ot}$ in equation (17) indirectly via this dissimilarity index instead.

We use these generated instruments in columns (5) to (7) of Table 4. The modification does not address the severe lack of variation in later decades, but compared to column (1) the Sanderson-Windmeijer F-statistic increases markedly, and the 1970s inflows are now predominantly predicted by the 1970s instrument. As expected, this improvement in the "look" of the first stages has however only a negligible effect on second-stage estimates, as we show below. Changing the base year for construction of the instrument has only a modest effect on the first-stage estimates either: while conceptually the earlier 1960 base year is clearly preferable (see Section II), variable and area definitions from the 1970 Census are more consistent with the data underlying our outcome variable, so apriori it is not clear which is the better choice. We report both specifications below, but the choice has little effect on estimates in the second stage.

V.2 Second-stage Results: Baseline and Robustness Tests

We thus estimate the impact of immigrant inflows on wage growth in the 1970s, reporting our results in Table 5. We again report estimates from different specifications, varying the construction of the instrument, the definition of the outcome variable, the weighting scheme, or the inclusion of control variables in columns (1) to (9). For comparison we report the conventional IV estimate first (Panel A), before showing the 2SLS second-stage (Panel B) and reduced form estimates (Panel C) of our proposed double instrumentation procedure. Our model provides clear predictions on its signs: the coefficient on the 1970s inflows captures the wage impact of recent arrivals in the short run, which in a factor proportions model is negative. The coefficient on 1960s inflows captures the longer term reaction to local shocks, and is predicted to be positive.

The coefficient on recent immigrant arrivals is indeed significantly negative. In our baseline specification, the impact of a one-percent (as a share of the local labor force) immigrant inflow is estimated to decrease average wages by about 0.7 log points. This estimate is substantially more negative than the corresponding conventional IV estimates in Panel A of Table 5 (or Table 2), consistent with our expectation that these estimates are upward-biased. Moreover, we find a positive and statistically significant positive coefficient on the predicted inflow in the past decade, consistent with our expectation that this coefficient measures the longer term adjustment of local labor market to local supply shocks. In absolute terms, this coefficient is not that much smaller than the negative coefficient on current inflows, suggesting that local wages largely recover from an immigration-induced supply shock within one decade. As discussed, these estimates capture however only the impact on local *relative* wages – with spatial equilibrium adjustments over time, immigration may have still have a positive or negative longer term effect on the national labor market.

The conventional spatial correlation estimates are sensitive to certain specification choices (Panel A), as are the double IV second-stage estimates (Panel B). The choice of base year for the instrument (column 2) has limited consequences, as have the trimming of the wage variable (column 5) or many other choices related to the construction of our variables that we do not report here (such as the use of current or lagged population as denominator when measuring the immigrant inflow rate). Taking the difference in an origin group's propensity to cluster, while improving the appearance of the first stages (see previous section), has only a limited effect on the second-stage estimates either (columns 3 and 4). As the scaling is not comparable to the standard version of the instruments we do not report the reduced form coefficients here. However, the choice of weighting scheme or control variables does matter. Studies in the spatial correlation literature sometimes weight MSAs by population levels, which in our case reduces both the conventional and double IV second-stage, but not the reduced-form estimates (column 6). Solon, Wooldridge and Haider (2015) note that the motivation for weighting by absolute populations is not clear, as it may neither help in the estimation of population-average causal effects nor increase efficiency.[21] As the

[21] Since all but three MSAs have populations above 100,000, individual-level uncertainty is unlikely to be an important factor in our sample, and we found only limited evidence for heteroscedasticity of the error term with respect to population size. We do use population weights on the commuting zone level, as many commuting zones have quite small populations.

variance of the dependent variable declines approximately linearly in log population, weighting by log population appears more reasonable (column 7).[22] These estimates remain quite similar to our unweighted baseline estimates.

MSAs differ substantially in their industry structure, so an obvious concern is a potential correlation between the past settlement instrument and changes in local labor demand from industry-specific or sectoral demand shifts. Estimates are however quite robust to the inclusion of a Bartik instrument for wage changes explained by lagged industrial compositions (column 8), or to the inclusion of local manufacturing or other industry shares (not shown).[23] On the other hand the inclusion of Census division fixed effects does indeed alter significantly the second-stage coefficients of both the conventional and double IV estimates (column 9). In contrast, the reduced-form estimates remain quite stable across all specifications, indicating that much of the observed changes in second-stage coefficients stems from a re-scaling of the first stages.

Appendix Table A.3 shows the corresponding estimates on the Commuting Zone level. While the Commuting Zone definition covers the whole country and comprises more observations, it appears less suitable for our purpose as the quality and comparability of those observations are lower than on the standard MSA level.[24] Estimates on the Commuting Zone

---

[22] We report heteroscedasticity-robust standard errors in Table 5. However, these may be downward biased in Panel B and C because of small sample bias (more specifically, the high correlation between the two instruments). Conventional estimates of the standard error are larger, but the coefficient estimate on recent arrivals remains significant at the 1 or 5 percent level in all specifications.

[23] A particular concern could be the large swings of prices and wages in the oil industry. While its local employment share is a highly significant predictor, it does not affect much the coefficients on immigrant inflows.

[24] Many commuting zones are small, so the measures of immigrant inflow rates suffer from sampling error (see Aydemir and Borjas, 2011). Moreover, many small commuting zones receive hardly any immigrants, but are unlikely to be good controls as they differ from more popular destinations also in other respects. To partially address the differences in size we include the lag of log population in all specifications reported in Table A.3.

level follow the same broad pattern as on the MSA level, but the point coefficients are smaller in absolute value, and the coefficient on the lag if often quite close to zero.

These sensitivity tests lead to two conclusions. On the one hand, the coefficient on recent immigrant arrivals becomes consistently more negative once we switch from the conventional to our double-IV procedure, across specifications, and different spatial aggregation schemes. Our empirical results therefore support our core argument that conventional 2SLS estimates based on the shift-share instrument are upward-biased due to the high correlation between current and past immigrant inflows. On the other hand, both the conventional and double-IV 2SLS estimates are quite sensitive to certain specification choices, and the latter are strongly driven by those few cities that around 1970 experienced a substantial break in the predicted immigrant inflow rate.[25] We therefore do not put much trust on any specific point estimate for the wage impact of immigrant arrivals; ultimately the data seems still ill-suited to give a precise answer to this question.

With these caveats in mind, our results suggest that local immigration-induced supply shocks have a non-negligible, negative effect on local wages in the short run. While the overall patterns is consistent with standard theories of factor demand, our estimates imply a more negative effect than the standard competitive model, and are at the lower end of the spectrum of the existing literature.[26] Of course, our point is precisely that our estimates are not

[25] The MSAs with the largest influence on the slope coefficients in equation (13) are Miami, which – because of the change in the country of origin-composition – was predicted to receive substantially less immigrants during the 1970s than in the previous decade, and San Antonio and El Paso in Texas, which were predicted to receive substantially more.

[26] Altonji and Card (2001), Monras (2015), Lllull (2015), Borjas and Monras (2016) find similarly large effects on average or relative wages. In the competitive model, the average wage effect is bounded by the capital share of production (see Borjas, 2015). A promising avenue for future research is the observation that even conditional on a wide range of observables, immigrant arrivals tend to earn lower wages than natives. A gap between marginal productivity and wages can have important consequences in a search model with frictions (see Chassamboulli and Palivos 2014; Chassamboulli and Peri 2015; Amior 2016).

directly comparable: while some previous studies compound the (presumably negative) short- and (presumably positive) longer-run wage response to immigration, we aim to isolate the former. And while our estimates are more negative, they need to be interpreted against the backdrop of the high serial correlation in local immigrant inflows that we documented. Even though several cities received large immigrant inflows of 10 or even 20 percent during the 1970s, our estimates do *not* suggest that relative wage levels changed much in response: these cities also received large inflows during the 1960s, and much of the implied respective wage effects cancel out. Instead, immigration appears to be a comparatively unimportant determinant of *local* wage levels, dominated by other factors such as the broad regional trends on the Census division level, or by industry-specific trends in wages.

V.3 Second-stage Results: Heterogeneity Across Subgroups

After estimating the response in average wages, we next study the impact of immigrant arrivals on subgroups defined by demographic characteristics and skills. The distributionary consequences of immigration are a common concern (e.g. Card, 2009), but another important motivation for us here is to provide additional support of the validity of our empirical strategy. While there is little consensus about the impact immigration on average wages, there is agreement on its distributionary consequences: workers who in terms of skills are more similar to immigrant workers, and thus more substitutable in production, are expected to experience a more adverse wage impact, and workers who are very dissimilar may instead benefit.[27] It would thus be a warning sign if we find the estimated adverse wage impact to be clustered among workers who are unlikely to directly compete with immigrant arrivals on the labor market.

[27] The mean and distributionary consequences of immigration are determined by different sets of structural parameters in the standard model (Borjas, 2013), which may help to explain why there is more agreement on one than the other.

In fact, the distributionary consequences of immigration seem theoretically so clear-cut that they are used as identifying assumptions in strands of the literature that are focused on the relative wage effects of immigration (Katz, Borjas and Freeman 1996, Borjas 2003, Card 2009, Ottaviano and Peri 2012). These studies typically cut the labor market into skill groups defined by education and age or experience, ore more recently along the distribution of wages (Dustmann, Frattini and Preston, 2013). A common concern in this *skill-cell approach* is that the observed education and age of immigrant arrivals may not be a good proxy for their effective skill in the destination country. Conditional on observable characteristics, immigrant arrivals earn significantly lower wages than native workers (e.g. Bratsberg and Ragan 2002, Borjas 2003, Dustmann and Preston 2012, Dustmann, Schönberg and Stuhler 2016), perhaps because experience and education gained in the origin country are less well rewarded as skills obtained in the destination country.

This "*downgrading*" of immigrant arrivals needs to be accounted for in order to determine which native groups they are likeliest to be closest substitutes. We here follow a method described in Dustmann, Schönberg and Stuhler 2016 to impute the *effective* education and experience of immigrant workers, based on their observed density across occupation-wage cells relative to native workers of different education-experience types (for alternative imputation methods see Borjas 2003 and Docquier, Ozden, and Peri 2014). We implement this imputation procedure in each Census for immigrant arrivals in the preceding decade, interacting 3-digit occupations and ten wage deciles, and distinguishing between two experience (*inexperienced* and *experienced*) and two education groups (*unskilled* with at most high school degree vs. *skilled* with at least some college attendance).

The results for the 1970, 1980 and 1990 Census are reported in Table 6. The first and second block of columns compare the observed education-experience distribution of natives and immigrant arrivals, while the third block reports the imputed effective skill of these arrivals. Immigrant arrivals are overrepresented among young workers, and the degree of overrepresentation is similar in each of the three Censuses (around 10-15%). In contrast, immigrant arrivals are observed to be slightly *more* educated than natives in the 1970 Census, but less educated in subsequent decades. This contrast is partly explained by the change in origin-composition, as arrivals from countries whose share among all immigrants rose after the Immigration Act of 1965 were on average less educated than those with declining shares.

However, application of our imputation procedure suggests – based on the observation that even arrivals with high education often work in occupation-wage cells in which native workers with low education are overrepresented – that the *effective* education of immigrant arrivals is substantially lower than their observed education in each of the censuses. Taking this downgrading into account, they are therefore strongly underrepresented among skilled workers even in the 1970 Census (21.8 vs 30.4%), and the gap increases further in the following decades. The implications of our imputation procedure regarding experience groups are more mixed: in the 1970 Census, immigrant arrivals are estimated to have more effective than observed experience, while the pattern is reversed in the following decades.[28] In the 1980 Census, immigrant arrivals are most overrepresented (by about 50%) among inexperienced and unskilled workers.

To illustrate against which workers immigrant arrivals are likely to compete it is also informative to study their density within the native wage distribution (Dustmann, Frattini and Preston 2013, and Dustmann, Schönberg and Stuhler 2016). In Figure 3 we show where immigrant arrivals are actually situated in the native wage distribution (dashed line), and where we would assign them if they received the same return to their experience and

[28] The distribution of inexperienced and experienced workers across occupation-wage cells is quite similar among unskilled natives (but very different for skilled native workers). The distinction between experience groups seems therefore less consequential, consistent also with the observation that the returns to experience are lower among the unskilled than the skilled.

education as natives (the solid line). The x-axis measures the percentiles of the wage distribution. The y-axis is the density of a particular group relative to natives (horizontal line at 1). For instance, a point (2, 20) means that members of the group are twice as likely as natives to be located at the 20th percentile of the native wage distribution.

The figure illustrates that based on their observed characteristics, immigrants should be slightly overrepresented at the bottom and the very top of the wage distribution. However, they are actually strongly overrepresented at the bottom, in particular around the 10^{th} to 20^{th} wage percentile, and strongly underrepresented at the top. To summarize, our evidence suggests that immigration should have the most adverse wage impact on native workers with low education and in the lower part of the wage distribution. What pattern we should expect across experience groups is less clear-cut, but it is reasonable to expect a more adverse effect among young, especially young and unskilled workers.

We report 2SLS estimates by subgroup in Table 7. Column (1) reports our preferred specification from Table 5, in which we use 1960 shares to construct both instruments. Restricting the sample to male workers yields similar estimates (column 2). But the estimated wage impact is quite different for unskilled than skilled workers (columns 3 and 4): unskilled workers experienced a larger wage drop upon arrival of new immigrants, but also a more substantial wage recovery in the next decade. Our estimates are largest in absolute size for young, unskilled workers (column 8), and in the bottom centiles of the wage distribution (columns 9 and 10). In contrast, our estimates suggest that immigration had substantially less (and possibly no) effect on local wages in the top of the wage distribution.

Again, we do not wish to put much weight on any specific point estimate. Ultimately, the available variation is too limited: by isolating recent immigrant arrivals from previous inflows we use a substantially narrower source of variation than previous studies in the literature. As a consequence, the estimates for some subgroups are relatively imprecise.

However, the overall pattern lines up well with theoretical expectations: the wage estimates appear most negative in those groups in which immigrant arrivals were strongly overrepresented, and more negligible in others. This result supports the interpretation that our empirical strategy is indeed capturing the wage impact of immigration, and not any other local shocks – while it is not unlikely that the spatial distribution of immigrant arrivals during the 1970s may have correlated with other types of shocks, it seems less plausible that both the spatial and subgroup pattern overlapped so closely.

VII. Conclusions

To study the impact of immigrant arrivals on labor markets is notoriously difficult. Their locational choices are not random, and the economy may adjust in many different ways to a change in local factor supplies. To establish causal identification in spite of these issues, much of the existing studies the short-term wage response using the past settlement instrument – a shift-share instrument that combines national inflows with the locational patterns of immigrants in a previous period.

We showed that this approach is unlikely to identify a well-defined causal effect of interest when there is only limited change in the composition of immigrant inflows on the national level. In such setting, the inflow rates of immigrants across cities will tend to be highly serially correlated. In recent decades in the U.S., the rates have been nearly perfectly correlated, with the same cities receiving again and again large inflows. As a consequence, the shift-share instrument predicts not only recent arrivals, but is also a great (and often better) predictor for arrivals in a previous decade.

The conventional IV estimator does then not only capture the short-term response to recent immigrant arrivals, but also the longer term adjustment processes that such arrivals may trigger. This compound effect is hard to interpret; the weights that the estimator puts on the short- and longer-term wage response will differ across applications, as the correlation of the instrument with recent and past immigrant inflows is context-specific. Moreover, the longer-term response of local relative wages itself is hard to interpret, as it may reflect spatial adjustment processes that eventually affect also "control" areas that were not directly exposed to immigrant inflows.

The greatest strength of the past settlement instrument, its ability to predict current flows to local labor markets, can thus turn into a weakness. In some sense, if the instrument is "too strong", it is difficult to believe that it constitutes a shock to local labor markets, or that it can plausibly separate the exogenous from the endogenous in the actual immigrant flows. The flipside of this argument is that the instrument exogeneity assumption may be more plausible in settings in which it is more difficult "to get the instrument to work" – in which the first-stage link between past settlements and inflows is weaker because the source of these inflows has been less stable over time, as is for example the case in many European countries (which are however also smaller, providing less spatial variation). In short, there may be better prospects to satisfy the exclusion restriction in settings in which the rank condition is harder to fulfil.

To address these issues systematically we proposed a revised estimation procedure, which isolates the variation in local immigrant inflows that is uncorrelated to inflows in the previous period. We implemented this idea by including the previous inflow in our estimating equation, and instrumenting it with a version of the past settlement instrument – a "*double instrumentation*" procedure that captures and separates both the short-term response, and the longer term adjustment of local relative wages to immigrant inflows. However, the procedure is data demanding, as the two instruments will typically be highly collinear – in the U.S. Census, there are not sufficient innovations in the location choices of immigrants in the later decades to distinguish the short and long-term response.

However, due to a structural break in the composition of immigrants after the Immigration and Nationality Act, we do observe some innovations in the spatial distribution of immigrant arrivals between the 1960s and 1970s – allowing us to apply our revised estimator in that period. We find estimates that are at the lower end of the previous literature, suggesting that the initial wage impact of immigration on natives can be large. However, we also find that this decline of local wages is reversed in the next period; cities that received large (predicted) immigrant inflows in the 1960s, but smaller inflows during the 1970s, tend to experience a relative wage increase. Immigration may thus have little, if any, adverse effect on local wages in the longer run. These findings are consistent with a simple competitive model; a shock in the supply of one factor depresses the returns to that input temporarily, but factor adjustments wash out that effect over time.

The idea to decompose immigrant inflows by origin groups (Card, 2001) is crucial for our estimation procedure. While this decomposition has – in our data – little effect on the conventional IV estimator, it allows us to isolate innovations in local immigrant inflows that are caused by changes on the national level. To increase the predictive power of the instrument further, one may exploit that new arrivals from different origin groups have a different propensity to cluster into past settlements.

Still, there are a number of important problems. We think our findings do demonstrate that the high serial correlation in immigrant inflows is problematic for reduced-form identification strategies from spatial data. But it remains to be seen if our more specific hypotheses – that the short-term wage impact is more negative than the conventional IV estimator suggests, and the longer term response positive – can be confirmed in data other than the U.S. Census. Even if we accept these findings, their interpretation is not straightforward. Their general pattern is consistent with the standard competitive model, but the magnitude of the short-term wage impact is not. Moreover, we only identify the long-term effect of immigration on *local relative* wages, while the long-run effect on the national economy is the more important question (e.g. Peri, 2016). However, we think the apparent inconsistencies in the spatial correlation literature, and the lack of consensus about the short-term effect of immigration – which in principle should be much easier to identify – are also a major stumbling block for research on this and other interesting questions in the literature.

Finally, our findings illustrate an intrinsic property of *shift-share* instruments that can be quite problematic. Shift-share instruments impute local shocks by combining local "shares" of industry, demographic or other compositions and aggregate "shifts", but the local share are always highly serially correlated. For such instruments to be valid even in the presence of dynamic adjustment processes, we thus require that their national components are not serially correlated. In contexts where there are frequent changes or a sudden shock on the national level, shift-share instruments may meet this assumption. In others, like the immigration literature, care must be taken to insure that there is sufficient variation over time to plausible interpret the results as causal effects. The variant of the shift-share methodology that we propose here can then be used to generate spatial variation that is uncorrelated with the spatial distribution of past shocks.

References:

ALTONJI, J. G., AND D. CARD (1991): "The Effects of Immigration on the Labor Market Outcomes of Less-skilled Natives," in Immigration, Trade and the Labor Market, NBER Chapters, pp. 201–234. National Bureau of Economic Research, Inc.

AMIOR, M. (2016): "The Labor Impact of Migration Under Monopsony," mimeo, Cambridge University.

AMIOR, M., AND A. MANNING (2015): "The Persistence of Local Joblessness," CEP Discussion Papers dp1357, Centre for Economic Performance, LSE.

AMUEDO-DORANTES, C., AND A. SEVILLA (2014): "Low-Skilled Immigration and Parenting Investments of College-Educated Mothers in the United States: Evidence from Time-Use Data," Journal of Human Resources, 49(3), 509–539.

AUTOR, D. H., AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103(5), 1553–97.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103(6), 2121–68.

AUTOR, D., D. DORN, AND G. H. HANSON (2016): "The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade," *Annual Review of Economics*, 8, 205–240.

AYDEMIR, A., AND G. BORJAS (2011): "Attenuation Bias in Measuring the Wage Impact of Immigration," Journal of Labor Economics, 29(1), 69-112

AYDEMIR, A., AND M. G. KIRDAR (2013): Quasi-Experimental Impact Estimates of Immigrant Labor Supply Shocks: The Role of Treatment and Comparison Group Matching and Relative Skill Composition, IZA Discussion Paper No. 7161.

BARTEL, A. P. (1989): “Where Do the New U.S. Immigrants Live?,” Journal of Labor Economics, 7(4), 371–91.

BARTIK, T. J. (1991): Who Benefits from State and Local Economic Development Policies?, in Books from Upjohn Press. W.E. Upjohn Institute for Employment Research.

BLANCHARD, O. J., AND L. F. KATZ (1992): “Regional Evolutions,” Brookings Papers on Economic Activity, 23(1), 1–76.

BOUND, J., D. A. JAEGER, AND R. M. BAKER (1995): “Problems with Instrumental Variables Estimation when the Correlation between the Instruments and the Endogenous Explanatory Variable is Weak,” Journal of the American Statistical Association, 90(430), 443-450.

BORJAS, G. J., R. B. FREEMAN, AND L. F. KATZ (1997): “How Much Do Immigration and Trade Affect Labor Market Outcomes?,” *Brookings Papers on Economic Activity*, 28(1), 1–90.

BORJAS, G. (1999): “The Economic Analysis of Immigration,” in Handbook of Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 3 of Handbook of Labor Economics, chap. 28, pp. 1697–1760. Elsevier.

BORJAS, G. (2003): “The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market*,” Quarterly journal of Economics, 118(4), 1335–1374.

BORJAS, G. J. (2006): “Native Internal Migration and the Labor Market Impact of Immigration,” *The Journal of Human Resources*, 41(2), pp. 221–258.

BORJAS, G. (2013): “The Analytics of the Wage Effect of Immigration,” *IZA Journal of Migration*, 2(1), 22.

BORJAS, G. (2014): Immigration Economics, Harvard University Press

BORJAS, G. (2015): The Wage Impact of the Marielitos: A Reappraisal, NBER Working Paper, No. 21588, NBER.

BORJAS, G. J., AND J. MONRAS (2016): "The Labor Market Consequences of Refugee Supply Shocks," NBER Working Papers 22656, National Bureau of Economic Research, Inc.

BRAUN, S. T., AND WEBER, H. (2016), "How Do Regional Labor Markets Adjust to Immigration? A Dynamic Analysis for Post-war Germany", Kiel Institute for the World Economy Discussion Paper No. 2025.

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. (2007): "How Immigration Affects U.S. Cities," CReAM Discussion Paper Series 0711, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London.

CARD, D. (2009): "Immigration and Inequality," American Economic Review, 99(2), 1-21.

CARD, D:, AND G. PERI (forthcoming): "Immigration Economics: A Review," Journal of Economic Literature, forthcoming.

CARPIO, X. V. D., AND M. WAGNER (2015): "The Impact of Syrians Refugees on the Turkish Labor Market*," Discussion paper, mimeo.

CHALFIN, A. (2015): "The Long-Run Effect of Mexican Immigration on Crime in US Cities: Evidence from Variation in Mexican Fertility Rates," *American Economic Review*, 105(5), 220–25

CHASSAMBOULI, A., AND G. PERI (2015): "The Labor Market Effects of Reducing the Number of Illegal Immigrants," *Review of Economic Dynamics*, 18(4), 792–821.

CHASSAMBOULLI, A., AND T. PALIVOS (2014): "A Search-Equilibrium Approach To The Effects Of Immigration On Labor Market Outcomes," *International Economic Review*, 55, 111–129.

CORTES, P. (2008): "The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data," Journal of Political Economy, 116(3), 381–422.

DUSTMANN, C., F. FABBRI, AND I. PRESTON (2005): "The Impact of Immigration on the British Labour Market," Economic Journal, 115(507), F324–F341.

DUSTMANN, C., T. FRATTINI, AND I. PRESTON (2013): "The Effect of Immigration along the Distribution of Wages," The Review of Economic Studies, 80(1). 2

DUSTMANN, C., AND A. GLITZ (2015): "How Do Industries and Firms Respond to Changes in Local Labor Supply?," Journal of Labor Economics, 33(3), 711 – 750.

DUSTMANN, C., AND I. PRESTON (2012): "Comment: Estimating the Effect of Immigration on Wages," *Journal of the European Economic Association*, 10(1), 216–223.

DUSTMANN, C., U. SCHOENBERG, AND J. STUHLER (2016): "The Impact of Immigration: Why Do Studies Reach Such Different Results?," Journal of Economic Perspectives, 30(4), 31–56.

DUSTMANN, C., U. SCHOENBERG, AND J. STUHLER (forthcoming): "Labor Supply Shocks, Native Wages, and the Adjustment of Local Employment", Quarterly Journal of Economics, forthcoming.

EBERTS, R. W., AND J. A. STONE (1992): Wage and Employment Adjustment in Local Labor Markets, in Books from Upjohn Press. W.E. Upjohn Institute for Employment Research.

FURTADO, D., AND H. HOCK (2010): "Low Skilled Immigration and Work-Fertility Tradeoffs among High Skilled US Natives," American Economic Review, 100(2), 224–28.

GONZÁLEZ, L., AND F. ORTEGA (2011): “How do very open economies adjust to large immigration flows? Evidence from Spanish regions,” *Labour Economics*, 18(1), 57–70.

GREENAWAY-MCGREVY, R., AND K. K. HOOD (2016): “Worker migration or job creation? Persistent shocks and regional recoveries,” *Journal of Urban Economics*, 96(C), 1–16.

HATTON, T. J. (2015): “United States Immigration Policy: The 1965 Act and its Consequences,” Scandinavian Journal of Economics, 117(2), 347–368.

HUNT, J. (2012): “The Impact of Immigration on the Educational Attainment of Natives,” IZA Discussion Papers 6904, Institute for the Study of Labor (IZA).

HUNT, J., AND M. GAUTHIER-LOISELLE (2010): “How Much Does Immigration Boost Innovation?,” American Economic Journal: Macroeconomics, 2(2), 31–56.

JAEGER, D. A. (2007): “Skill Differences and the Effect of Immigrants on the Wages of Natives,” Discussion paper, mimeo.

KOVAK, B. K. (2013): “Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?,” American Economic Review, 103(5), 1960–76.

LALONDE, R. J., AND R. H. TOPEL (1991): “Labor Market Adjustments to Increased Immigration,” in Immigration, Trade and the Labor Market, NBER Chapters, pp. 167–199. National Bureau of Economic Research, Inc.

LEWIS, E. (2011): “Immigration, Skill Mix, and Capital Skill Complementarity,” Quarterly Journal of Economics, 126(2), 1029-1069

LLULL, J. (2014): “The Effect of Immigration on Wages: Exploiting Exogenous Variation at the National Level,” CReAM Discussion Paper Series 1436, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London.

MONRAS, J. (2015): “Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis,” Sciences Po Economics Discussion Papers 2015-04, Sciences Po Departement of Economics.

MORETTI, E. (2011): “Local Labor Markets,” in Handbook of Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 4 of Handbook of Labor Economics, chap. 14, pp. 1237–1313. Elsevier.

NAKAMURA, E., AND J. STEINSSON (2014): “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” American Economic Review, 104(3), 753–92.

NUNN, N., AND N. QIAN (2014): “US Food Aid and Civil Conflict,” American Economic Review, 104(6), 1630–66.

ORRENIUS, P. M., AND M. ZAVODNY (2007): “Does immigration affect wages? A look at occupation-level evidence,” Labour Economics, 14(5), 757 – 773.

PERI, G. (2016): “Immigrants, Productivity, and Labor Markets,” Journal of Economic Perspectives, 30(4), 3-30

PISCHKE, J.-S., AND J. VELLING (1997): “Employment Effects of Immigration to Germany: An Analysis Based on Local Labor Markets,” The Review of Economics and Statistics, 79(4), pp. 594–604.

RUGGLES. S., K. GENADEK, R. GOEKEN, J. GROVER, AND M. SOBEK (2015), Integrated Public Use Microdata Series: Version 6.0 [Machine-readable database], University of Minnesota

SAIZ, A. (2007): “Immigration and housing rents in American cities,” Journal of Urban Economics, 61(2), 345–371.

SOLON, G., S. HAIDER, AND J. M. WOOLDRIDGE (2015): “What Are We Weighting For?,” *Journal of Human Resources*, 50(2), 301–316.

TOLBERT, C. M., AND M. SIZER (1996): "U.S. commuting zones and labor market areas: A 1990 update," Discussion paper, Staff report No. AGES-9614, US Department of Agriculture, Washington, DC.

TOPEL, R. H. (1986): "Local Labor Markets," Journal of Political Economy, 94(3), pp. S111–S143.

THE NATIONAL ACADEMY OF SCIENCE (2016), "The Economic and Fiscal Consequences of Immigration", Panel on the Economic and Fiscal Consequences of Immigration, by Francine D. Blau and Christopher Mackie, Eds.

WILSON, D. J. (2012): "Fiscal Spending Jobs Multipliers: Evidence from the 2009 American Recovery and Reinvestment Act," American Economic Journal: Economic Policy, 4(3), 251–282.

WOZNIAK, A., AND T. J. MURRAY (2012): "Timing is everything: Short-run population impacts of immigration in US cities," Journal of Urban Economics, 72(1), 60–78.

Table 1: Characteristics of immigrant inflows

Variable	1950	1960	1970	1980	1990	2000	2010
National immigrant share	0.076	0.056	0.052	0.067	0.087	0.117	0.136
Share recent arrivals							
Nation			0.016	0.025	0.037	0.044	0.032
Average MSA			0.014	0.020	0.029	0.037	0.028
Standard deviation across MSAs			0.018	0.022	0.034	0.030	0.019
Coefficient of variation across MSAs			1.31	1.11	1.17	0.81	0.66
Share of recent arrivals from							
Canada + Europe			0.41	0.17	0.13	0.16	0.12
Mexico			0.11	0.23	0.24	0.33	0.28
Other Latin America			0.26	0.20	0.24	0.21	0.23
Asia			0.17	0.32	0.32	0.26	0.31
Africa/Other			0.05	0.08	0.08	0.04	0.06
Serial correlation in national composition							
Recent arrivals, 38 origins (excl. Other)				0.59	0.99	0.96	0.98
Recent arrivals, excluding Mexico				0.37	0.95	0.90	0.95
Immigrant stocks, 16 origins (excl. Other)	0.99	0.99	0.94	0.65	0.90	0.97	1.00

Note: Based on U.S. census data and 109 MSAs (see text). Recent arrivals are immigrants who arrived during the last decade.

Figure 1: The composition of immigrant arrivals in the U.S.

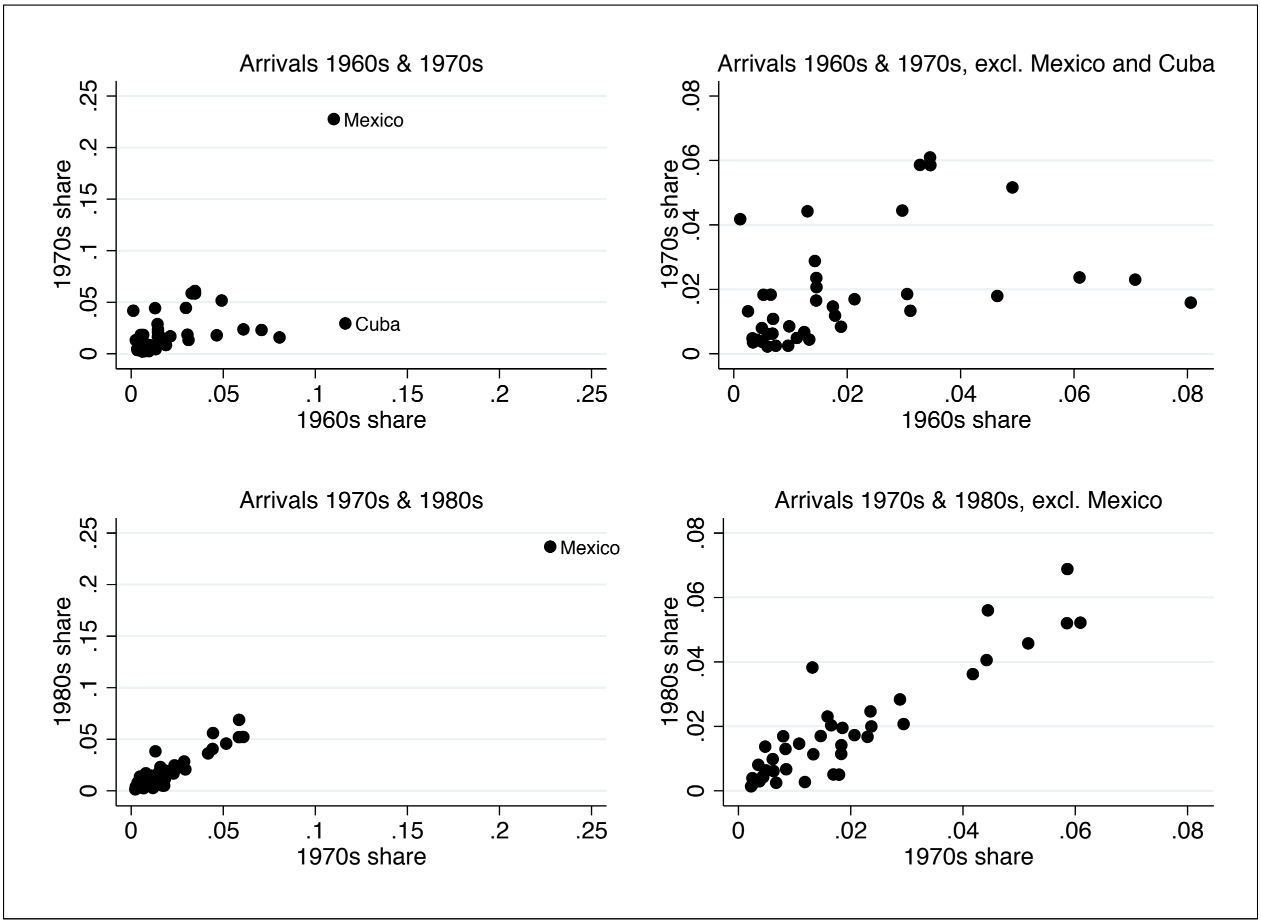

Note: Each observation is the share of all newly-arrived immigrants that were born in a specific country. N=39.

Table 2: Estimated Wage Impact of Immigration

Panel A: OLS	1980	1990	2000	2010
Imm. inflow rate	-0.167	0.449**	0.120	-0.041
	(0.138)	(0.127)	(0.118)	(0.144)

Panel B: 2SLS	1980	1990	2000	2010
Imm. inflow rate	-0.352*	0.382**	-0.092	-0.047
	(0.155)	(0.102)	(0.112)	(0.136)
First stage	0.686**	0.976**	0.629**	0.749**
	(0.132)	(0.175)	(0.114)	(0.058)
R-squared	0.674	0.775	0.655	0.832
F-statistic	26.97	31.11	30.53	170

Note: Based on U.S. census data and 109 MSAs. The table reports the slope coefficient in a regression of the change in residual log wage on the immigrant inflow rate in the decade preceding each census year. Reference year for past settlement instrument is beginning of decade. Robust standard errors in parentheses, ** $p<0.01$, * $p<0.05$.

Table 3: Correlation in Local Immigrant Inflows

Panel A: Serial Correlation	1980	1990	2000	2010
Actual inflows	0.82	0.96	0.92	0.96
Past settlement IV	0.70	0.99	0.96	0.99

Panel B: Predicted Inflows	1980	1990	2000	2010
Correlation(inflow,IV)				
IV baseperiod t-1	0.82	0.88	0.81	0.91
IV baseperiod t-2	0.73	0.69	0.68	0.78
Correlation (past inflow,IV)				
IV baseperiod t-1	0.62	0.96	0.93	0.95
IV baseperiod t-2	0.51	0.81	0.81	0.83

Panel C: Serial Correlation by Skill Group	1980	1990	2000	2010
Actual inflows				
high skilled	0.79	0.95	0.94	0.97
low skilled	0.81	0.95	0.88	0.93
log skill ratio	0.62	0.80	0.76	0.73
Past settlement IV				
high skilled	0.70	0.97	0.98	0.99
low skilled	0.72	0.98	0.98	0.99
log skill ratio	0.88	0.95	0.99	0.99

Note: Each entry is a pairwise correlation across 109 MSAs. Panels A and C report the serial correlations in actual inflows and in the past settlement IV: over all immigrants in Panel A, and over subgroups or ratios between those in Panel C. Panel B compares the correlation between the IV and the inflow it is supposed to predict, with that between the IV and the previous inflow. Baseperiod t-1 and t-2 mean that the instrument is constructed using the immigrant distribution in the observation year minus one and two decades respectively.

Table 4: Double Instrumentation: First Stage

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1980	1990	2000	2010	1980	1980	1980
IV base period	1970	1980	1990	2000	1970	1960	1960
Generated IVs (with clustering)	-	-	-	-	Version A	Version A	Version B
(i) First stage for X=Imm(t)							
Z(t)	0.415**	-0.190	-0.557	1.049*	0.916**	0.941**	1.303**
	(0.115)	(0.456)	(0.288)	(0.400)	(0.158)	(0.165)	(0.192)
Z(t-1)	0.325**	1.221*	1.211**	-0.303	0.208**	0.247**	-0.0577
	(0.0478)	(0.510)	(0.265)	(0.415)	(0.0775)	(0.0777)	(0.0723)
F-statistic	166.2	15.15	107.9	103.7	310.9	364.4	667.9
Sanderson-Windmeijer F-stat.	47.68	6.05	1.11	0.85	94.3	62.96	73.9
(ii) First stage for X=Imm(t-1)							
Z(t)	-0.0980*	0.268	-0.502	1.339*	0.176*	0.180	0.468**
	(0.0469)	(0.163)	(0.279)	(0.513)	(0.0696)	(0.0946)	(0.121)
Z(t-1)	0.719**	0.376*	1.430**	-0.157	0.875**	0.907**	0.472**
	(0.0171)	(0.173)	(0.241)	(0.534)	(0.0334)	(0.0808)	(0.0659)
F-statistic	5261	53.32	532.5	175.7	7306	268	273.2
Sanderson-Windmeijer F-stat.	153.46	6.35	1.1	0.87	378.58	107.66	118.3

Note: N=109. X denotes actual immigrant inflows and Z denotes the past settlement instrument for the decade preceding each census year. For comparability the two instruments are scaled to the same mean, i.e. the mean of the instrument in period t. Using standard (decomposed by origin group) version of past settlement instrument in columns (1)-(4), and generated instruments in columns (5)-(7) that incorporate the estimated clustering of origin groups in the 1970 and 1980 Census (Version A) or 1990 Census (Version B). Robust standard errors in parentheses. ** p<0.01, * p<0.05

Table 5: Double Instrumentation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
IV base period	1970	1960	1960	1960	1960	1960	1960	1960	1960
			GIV-A	GIV-B	trim 5%	abs weights	log weights	Bartik	region FE
Panel A: 2SLS									
Imm(t)	-0.352*	-0.428**	-0.321	-0.244	-0.408**	-0.284*	-0.411**	-0.463**	-0.596*
	(0.155)	(0.158)	(0.190)	(0.207)	(0.149)	(0.113)	(0.156)	(0.162)	(0.235)
Panel B: 2SLS w/ Double IV									
Imm(t)	-0.773**	-0.883**	-0.879**	-0.915**	-0.805**	-0.532*	-0.863**	-0.944**	-1.317**
	(0.214)	(0.235)	(0.211)	(0.225)	(0.209)	(0.250)	(0.235)	(0.248)	(0.437)
Imm(t-1)	0.575**	0.666**	0.649**	0.677**	0.582**	0.359	0.655**	0.705**	0.893**
	(0.156)	(0.171)	(0.142)	(0.148)	(0.151)	(0.251)	(0.171)	(0.184)	(0.305)
Panel C: Reduced Form									
Imm(t)	-0.378**	-0.373**	-	-	-0.332**	-0.407**	-0.375**	-0.398**	-0.416*
	(0.0921)	(0.0924)	-	-	(0.0769)	(0.123)	(0.0934)	(0.0964)	(0.160)
Imm(t-1)	0.245*	0.304*	-	-	0.239*	0.245	0.308*	0.315**	0.263
	(0.0945)	(0.125)	-	-	(0.117)	(0.220)	(0.125)	(0.0928)	(0.185)

Note: Based on U.S. census data and 109 MSAs. The dependent variable is the change in residual log wages by MSA between the 1970 and 1980 Census. Generated instruments in columns (3) and (4) incorporate the estimated clustering of origin groups in the 1970 and 1980 Census (Version A) or 1990 Census (Version B), standard version of past settlemnt instrument in all other columns. Bottom 5% of wages trimmed in column (5), observations weighted by lagged total or log population in columns (6) and (7). Columns (8) and (9) include a Bartik IV or Census Division fixed effects as control variables. Robust standard errors in parentheses, ** p<0.01, * p<0.05.

Table 6: The Observed and Effective Skills of Immigrant Arrivals

		Natives			Immigrants: Observed			Immigrants: Effective		
		Potential experience			Potential experience			Potential experience		
Year	Education	1-20 yrs	21-40 yrs	Total	1-20 yrs	21-40 yrs	Total	1-20 yrs	21-40 yrs	Total
1970	Low	*34.9%*	*34.7%*	*69.6%*	*37.8%*	*28.4%*	*66.2%*	*33.8%*	*44.4%*	*78.2%*
	High	*20.4%*	*10.0%*	*30.4%*	*28.2%*	*5.6%*	*33.8%*	*19.4%*	*2.4%*	*21.8%*
	Total	*55.4%*	*44.7%*		*66.0%*	*34.0%*		*53.2%*	*46.8%*	
1980	Low	*33.7%*	*23.3%*	*56.9%*	*40.8%*	*20.2%*	*61.0%*	*62.0%*	*13.0%*	*75.0%*
	High	*32.2%*	*10.8%*	*43.1%*	*34.0%*	*5.0%*	*39.0%*	*22.3%*	*2.7%*	*25.0%*
	Total	*65.9%*	*34.1%*		*74.8%*	*25.2%*		*84.4%*	*15.6%*	
1990	Low	*27.1%*	*18.6%*	*45.6%*	*41.9%*	*17.1%*	*59.1%*	*61.9%*	*9.4%*	*71.3%*
	High	*37.2%*	*17.2%*	*54.4%*	*34.5%*	*6.4%*	*41.9%*	*24.8%*	*3.9%*	*28.7%*
	Total	*64.2%*	*35.8%*		*76.5%*	*23.5%*		*86.7%*	*13.3%*	

Note: Based on U.S. census data. The table reports the distribution of natives and recent immigrants (those who arrived within the past decade) across observed or imputed education-experience cells. The imputation of effective skills for immigrants is based on the distribution of workers across wage centiles and 2-digit occupations and described in Dustmann, Schönberg and Stuhler (2016).

Figure 2: Position of Immigrants in 1980 Native Wage Distribution

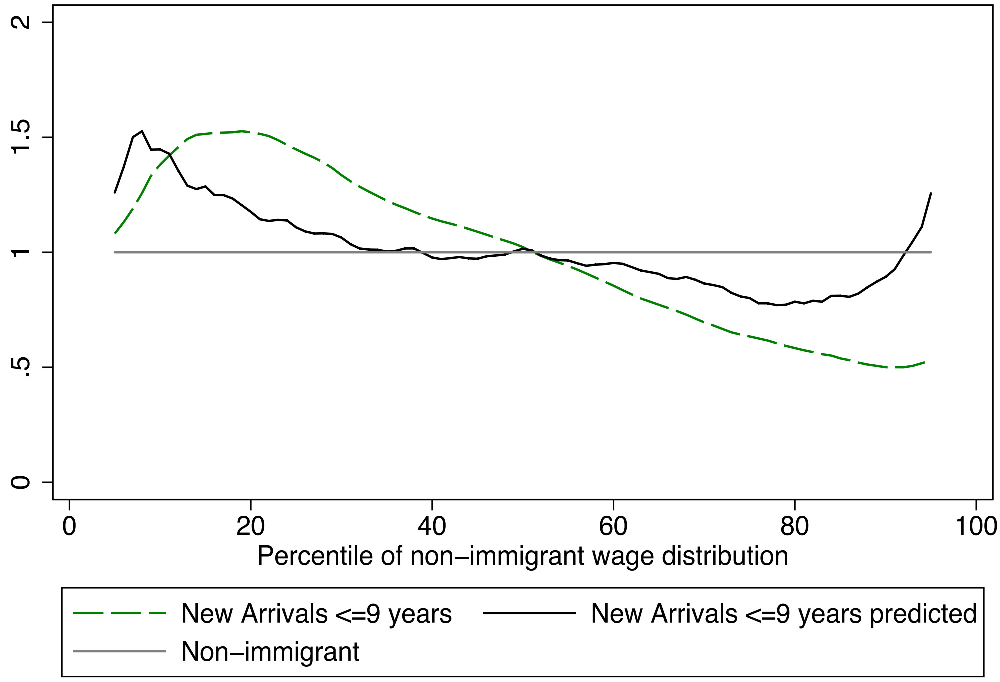

Note: The figures plots kernel estimates of where recent immigrants (who arrived over the past decade) are actually situated in the native wage distribution (the dashed lines), and where we would assign them if they received the same return to their experience and education as natives (the solid lines). The horizontal line shows as a reference the native wage distribution. The kernel estimates are above the horizontal line at wages where immigrants are more concentrated than natives, and below the horizontal line at wages where immigrants are less concentrated than natives.

Table 7: Second Stage by Subgroup

	(1)	(2)	(3)	(4)
	all	male	low skilled	high skilled
Imm(t)	-0.883**	-0.741**	-0.971**	-0.224
	(0.235)	(0.281)	(0.289)	(0.401)
Imm(t-1)	0.666**	0.497*	0.714**	0.301
	(0.171)	(0.214)	(0.207)	(0.402)
	(5)	(6)	(7)	(8)
	young (age <=30)	mid-age (31-50)	old (51-64)	young, low skilled
Imm(t)	-1.076**	-0.643*	-0.746	-1.371**
	(0.334)	(0.267)	(0.438)	(0.462)
Imm(t-1)	0.930**	0.457*	0.541	1.074**
	(0.240)	(0.194)	(0.335)	(0.326)
	(9)	(10)	(11)	(12)
	p10	p25	p75	p90
Imm(t)	-0.728	-1.021	-0.239	-0.328
	(0.757)	(0.555)	(0.253)	(0.292)
Imm(t-1)	0.481	0.578	-0.183	-0.269
	(0.757)	(0.479)	(0.350)	(0.394)

Note: Based on U.S. census data and 109 MSAs. The dependent variable is the change in residualized mean (columns 1-8) or percentile (columns 9-12) of log wages in the decade preceding the census year. Low (high) skilled are workers with at most a high school degree (5+ years of college). Estimation by 2SLS. Instrument constructed using dissimilarity index, base period is 1960 for both instruments. Robust standard errors in parentheses. ** p<0.01, * p<0.05.

Figure A.1: The Clustering of Origin Groups

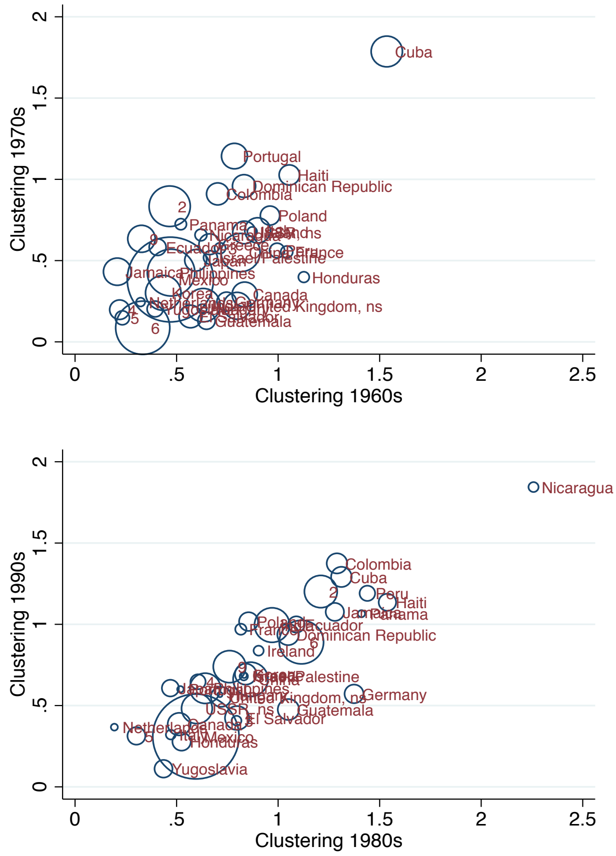

Note: The figure plots a measure of the clustering of new arrivals in an origin group in one against the previous decade. The clustering of an origin group is measured by the slope coefficient from a regression of its actual inflow rate against its predicted inflow rate (i.e. the origin-specific past settlement instrument) across MSAs. Base period for the construction of the instrument is the beginning of the respective decades, i.e. 1960 for the top left panel. Observations scaled by the origin share in national inflows.

Figure A.2: The Spatial Concentration and Clustering of Origin Groups

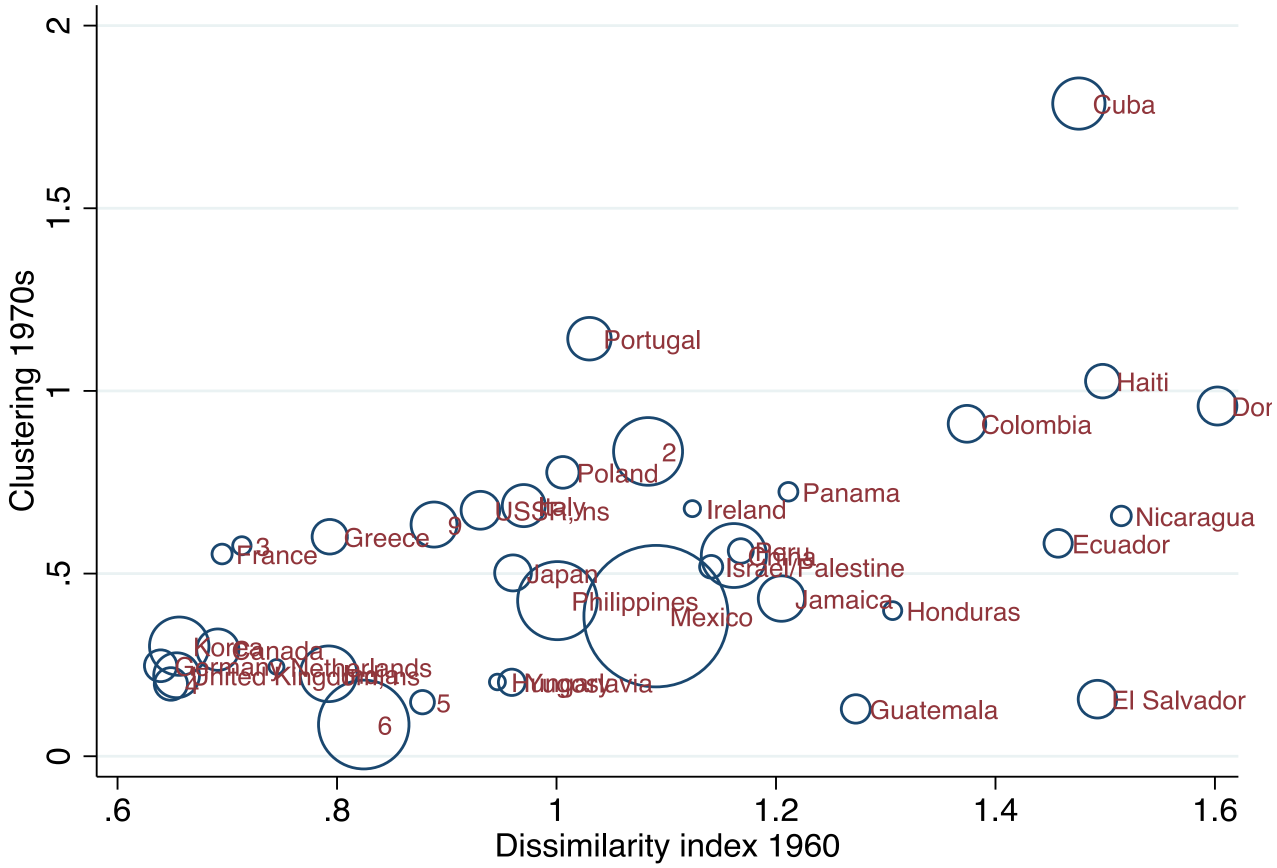

Note: The figure plots a measure of the clustering of new arrivals in an origin group during the 1970s against their spatial concentration in 1960. The clustering of an origin group is measured by the slope coefficient from a regression of its actual inflow rate against its predicted inflow rate (i.e. the origin-specific past settlement instrument) across MSAs. The spatial concentration is measured by the group's dissimilarity index compared to the total population in 1960. Scaled by the origin's share in national inflows in the 1970s.

Table A.1: Publications using the Past Settlement Instrument

Authors	Year	Journal	Outcome
Altonji and Card	1991	Book chapter	Native labor market outcomes
Card and DiNardo	2000	AER: P&P	Internal migration
Card	2001	JOLE	Internal migration, labor market outcomes
Fairlie and Meyer	2003	JOLE	Native self-employment
Dustmann, Fabbri and Preston	2005	Economic Journal	Native labor market outcomes
Hatton and Tani	2005	Economic Journal	Internal migration
Ottaviano and Peri	2005	JoUE	Native wages and employment
Ottaviano and Peri	2006	J. of Econ. Geography	Native wages and housing market
Reed and Danziger	2007	Am. Econ. Review	Native labor market outcomes
Saiz	2007	JoUE	Housing market
Cortes	2008	J. Political Econ.	Prices (goods and services)
Frattini	2008	mimeo	Prices (goods and services)
Kugler and Yuksel	2008	mimeo	Native labor market outcomes
Peri and Sparber	2009	AEJ: Applied	Task specialization
Card	2009	AER: P&P	Native labor market outcomes
Iranzo and Peri	2009	ReStat	Schooling externalities and productivity (TFP)
Hunt and Gauthier-Loiselle	2010	AEJ: Macro	Innovation
Furtado and Hock	2010	AER: P&P	Fertility
Boustan	2010	Quarterly J. of Econ.	Residential segregation
Kerr and Lincoln	2010	JOLE	Science and engineering employment, patenting
Cortes and Tessada	2011	AEJ: Applied Econ.	Labor supply, household work and services
Lewis	2011	Quarterly J. of Econ.	Investment in automation
Gonzalez and Ortega	2011	Labour Econ.	Labor market outcomes
Farré, Libertad and Francesc	2011	B.E. J. Econ. A&P	Female labor supply
Cortes and Tessada	2011	AEJ: Applied	Female labor supply
Cascio and Lewis	2012	AEJ: Policy	Residential and school segregation
Beaudry and Green	2012	Econometrica	Wage determination
Bianchi, Buonanno, and Pinotti	2012	J. Eur. Econ. Ass.	Crime
Smith	2012	JOLE	Youth employment
Wozniak and Murray	2012	JoUE	Population, internal migration
Hunt	2012	Working Paper	Educational attainment
Peri	2012	ReStat	Productivity (TFP)
Malcho-Moller, Munch and Skaksen	2012	Scan. J. Econ.	Firm-level wages
Dustmann, Frattini and Preston	2013	ReStud	Native labor market outcomes
Lafortune	2013	AEJ: Applied	Marriage market
Ottaviano, Peri and Wright	2013	Am. Econ. Review	Native labor market outcomes
Monras	2013	Working Paper	Native labor market outcomes
Bell, Fasani and Machin	2013	ReStat	Crime
Facchini, Mayda and Mendola	2013	Working Paper	Native labor market outcomes
Amuedo-Dorantes, Sevilla	2014	J. Human Res.	Parental time investment
Cortes and Pan	2014	J. Health Econ.	Supply of native nurses
Aydemir and Kirdar	2014	Working Paper	Native labor market outcomes
Llull	2014	Working Paper	Native labor market outcomes
Piyapromdee	2014	Working Paper	Native labor market outcomes and welfare
Ganguli	2015	JOLE	Knowledge diffusion
Orrenius and Zavodny	2015	JOLE	Educational choices
Amior	2015	Working Paper	Native labor market outcomes
Del Carpio, Özden, Testaverde, and Wagner	2015	Scan. J. Econ.	Native labor market outcomes
Dustmann and Glitz	2015	JOLE	Firm adjustment
Özden and Wagner	2015	Working Paper	Native labor market outcomes
Machin and Muprhy	2015	Working Paper	Higher education
Chalfin	2015	AER: P&P	Crime
Ottaviano, Peri and Wright	2015	Working Paper	Firm-level trade of services
Forlani, Lodigiani and Mendolicchio	2015	Scan. J. Econ.	Female labor supply
Cattaneo, Fiori and Peri	2015	J. Human Res.	Native labor market outcomes
Kasy	2015	JoUE	Location choices with social externalities
Sharpe	2015	PhD Thesis	Housing market
Ransom and Winters	2016	Working Paper	STEM education and employment
Fernandez-Huertas, Ferrer and Saiz	2016	Working Paper	Residential segregation
Fassio, Kalantaryan and Venturini	2016	Working Paper	Productivity (TFP)
Foged and Peri	2016	AEJ: Applied	Native labor market outcomes

Note: The table lists publications that use a version of the past settlement instrument and their outcome of interest. JOLE=Journal of Labor Economics, JoUE=Journal of Urban Economics, AEJ=American Economic Journal, Restat=Review of Economics and Statistics, ReStud=Review of Economic Studies.

Table A.2: Estimated Wage Impact of Immigration (Commuting Zones)

Panel A: OLS	1980	1990	2000	2010
Imm. inflow rate	-0.300**	0.581**	-0.048	-0.034
	(0.0837)	(0.0878)	(0.109)	(0.116)

Panel B: 2SLS	1980	1990	2000	2010
Imm. inflow rate	-0.381**	0.574**	-0.261	-0.005
	(0.0972)	(0.0833)	(0.170)	(0.0819)
First stage	0.782**	1.017**	0.602**	0.678**
	(0.166)	(0.0570)	(0.115)	(0.0819)
R-squared	0.663	0.891	0.714	0.823
F-statistic	22.28	318.4	27.28	68.53

Note: Based on U.S. census data and 741 Commuting Zones. The table reports the slope coefficient in a regression of the change in residual log wage on the immigrant inflow rate in the decade preceding each census year. Reference year for past settlement instrument is beginning of decade. Observations weighted by population in reference period. Robust standard errors in parentheses, ** p<0.01, * p<0.05.

Table A.3: Double Instrumentation: Second Stage (Commuting Zones)

	(1)	(3)	(6)	(7)
IV base period	1970	1970	1970	1970
		trim 5%	Bartik	region FE
Panel A: 2SLS				
Imm(t)	-0.271	-0.279*	-0.301*	-0.516*
	(0.140)	(0.124)	(0.148)	(0.206)
Panel B: 2SLS w/ Double IV				
Imm(t)	-0.359	-0.321	-0.398	-0.836*
	(0.219)	(0.186)	(0.247)	(0.390)
Imm(t-1)	0.141	0.0676	0.156	0.409
	(0.162)	(0.131)	(0.201)	(0.268)
Panel C: Reduced Form				
Imm(t)	-0.169**	-0.151**	-0.189**	-0.273**
	(0.0635)	(0.0550)	(0.0679)	(0.0745)
Imm(t-1)	0.00353	-0.0573	0.00323	0.0256
	(0.0706)	(0.0735)	(0.0744)	(0.107)

Note: Based on U.S. Census data and 741 Commuting Zones. The dependent variable is the change in residual log wages by commuting zone between the 1970 and 1980 Census. All regressions include lag log population as control variable and are weighted by lagged total population size. Bottom 5% of wages trimmed in column (2). Columns (3) and (4) include a Bartik IV or Census Division fixed effects as control variables. Robust standard errors in parentheses, ** p<0.01, * p<0.05.