Foundational Essays in Immigration Economics

George J. Borjas

Introduction

This volume collects the essays that represent what I think are my most important publications on the economics of immigration. I have been working on immigration-related issues since the early 1980s. I never imagined at the time that it would become the research topic that would occupy most of my thinking time and define my career. Nevertheless, as time went by, I realized that the economic approach had never been applied to the study of international immigration, so each question I asked led to another and each new paper raised new puzzles to think about.

Perhaps a word of warning is in order, as this introduction partly describes my personal journey and participation in building the subfield that we now call "immigration economics." To a surprising extent, that journey reflects the environment where I matured intellectually and my reaction to random personal contacts that led to research projects many years later. I hope this description of how my research evolved over the decades gives a sense of why certain questions sparked my interest (while others did not) and what I was trying to learn in this life-long investigation.

At the time I began to think about the economics of immigration, immigration was not a policy concern in any of the industrialized countries that would soon see a surge in immigrant flows. Even in the United States, which is typically thought of as a "nation of immigrants," the foreign-born share of the population had reached a historic low of 4.7 percent in 1970. Given the unimportance of immigration in the debate over social policy, it is not too surprising that academic economists paid little attention to it.

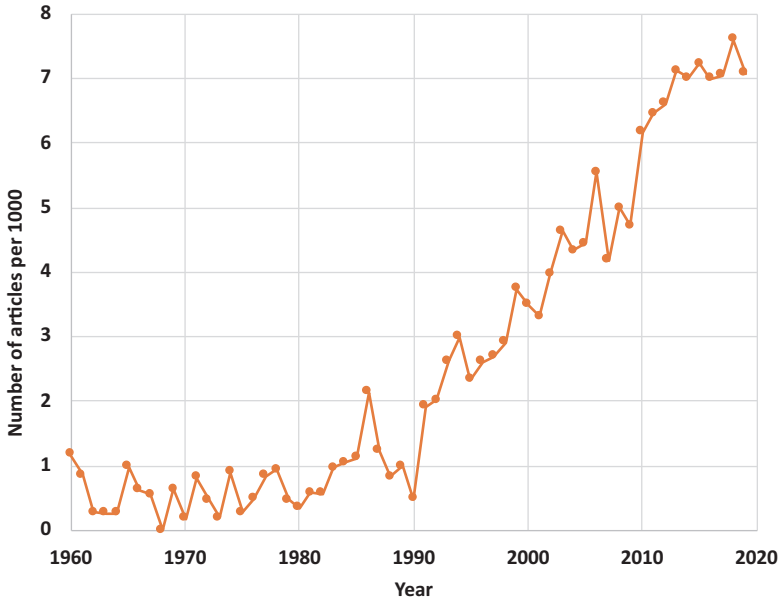

Figure 1: Immigration articles published in economics, 1960–2019.
Source: Web of Science, Social Science Citation Index. The figure gives the number of articles published in economics (per 1000 articles published) where the topic of the study is either immigration or international migration, or the article has the words "immigration" or "immigrant" or "foreign born" in the title.

Figure 1 uses data from the Web of Science to illustrate the trend in the number of immigration-related papers published in economics journals since 1960. Fewer than one paper per year (per 1000 papers published in economics) dealt with immigration issues through the mid-1980s. There was clearly some self-selection in the type of economist attracted to immigration issues at the time. It could not have been because they thought their paper would make a policy splash and end up being discussed in the front page of the *New York Times*. In my case, the attraction was that the study of immigration issues could increase our understanding of the role of human capital in the labor market and why some people chose one course of economic action (for example, move) while other people did not. This all began to change in the 1990s, as immigration increasingly became a central and contentious policy issue in many countries.

By the 2010s, the (relative) number of papers published dealing with immigration issues had increased seven-fold from what it was in the 1970s and early 1980s.

The papers collected in this volume are divided into five general topics. Part I examines the issue of economic assimilation, how the human capital stock of immigrants (typically proxied by earnings) changes over time, and how the assimilation process leads to a narrowing of the earnings gap between immigrants and natives.

Part II deals with the self-selection of immigrants. It is an obvious fact, but it needs restating: Not everyone becomes an immigrant. And, equally important, immigrants are not randomly selected from the population and, as a result, differ in fundamental ways from those who choose to stay.

Part III considers the intergenerational mobility of immigrants. The effects of current immigration can persist for several generations, depending on how quickly the children and grandchildren of immigrants become like natives.

Part IV investigates the labor market impact of immigration, perhaps the most contentious issue in the debate over immigration policy. It is not surprising that large supply shocks are likely to change the income distribution. Immigration is no exception to the general rule that policies typically leave winners and losers in their wake, and the labor market impact of immigration often defines who the winners and losers are.

Finally, Part V addresses the benefits from immigration. The same forces that lead to lower wages after a large supply shock also imply that some persons will benefit from the changed wage structure. As with the calculation of gains from international trade, these benefits will outweigh the lost wages suffered by the losers, generating an immigration surplus.

I. Economic Assimilation

Soon after entering the doctoral program in economics at Columbia University in 1971, James Heckman offered me the life-changing opportunity of becoming his research assistant at the New York office

of the National Bureau of Economic Research. Although I nominally got my graduate education at Columbia, it was at the 261 Madison Avenue headquarters of the Bureau that I spent much of my graduate student years and it was there where I learned how to do research and think like an economist.

Heckman was a new faculty hire at Columbia, having just gotten his doctorate from Princeton. He was beginning to work on the issues that would eventually lead to the Nobel Prize-winning selection bias correction (Heckman, 1979). After he left for the greener pastures at the University of Chicago in 1973, I became Jacob Mincer's research assistant at the Bureau. Mincer had already derived what would eventually be known as the Mincer Earnings Function (Mincer, 1974). Not surprisingly, I developed an interest in human capital and earnings determination, an interest that would eventually become my doctoral thesis: incorporating job separations (i.e., quits and layoffs) into an analysis of how earnings evolved over the life cycle (Borjas, 1981).

It was at the Bureau where I got a sense of how paradigm-changing research was done by having the unique opportunity to observe and assist (in my very limited way) the day-to-day progress of the research agendas of Heckman and Mincer. I spent countless hours interacting with a superb group of graduate students who went on to have distinguished academic careers (including Ann Bartel and Kenneth Wolpin). And I watched in awe as the frontiers of labor economics and related fields were literally being pushed out paper after paper by Michael Grossman, James Smith, Finis Welch, and the many other applied microeconomists who populated the offices up and down the hall from mine.

Sometime in the mid-1970s, Barry Chiswick, who himself had been a product of the Columbia "human capital factory" a decade earlier, came to present his latest paper at Columbia's Labor Workshop. In hindsight, this study, which examined the economic assimilation of immigrants in the United States, jump-started the modern analysis of the economics of immigration (Chiswick, 1978). It was immediately obvious that Chiswick's study was an important extension of the Mincer framework. Mincer had related earnings to the

number of years of labor market experience. Chiswick argued that applying the framework to foreign-born men required we distinguish between United States and foreign labor market experience, as part of the human capital acquired in other countries was probably specific and not transferable to the United States labor market.

Chiswick's analysis of the 1970 Census cross-section revealed three intriguing empirical regularities. First, the age-earnings profile of immigrants was below that of natives at the time of arrival, so that immigrants suffered an initial wage disadvantage. Second, the slope of the immigrant age-earnings profile was steeper, so that their earnings grew faster than those of comparable natives. Chiswick argued that the steeper slope reflected a faster rate of acquisition of human capital in the post-migration period, such as learning the English language. Third, the two age-earnings profiles eventually crossed, so that immigrants actually earned *more* than natives in the long run. This overtaking point took place at around 10 to 15 years after arrival. Chiswick (1978, p. 900) attributed the overtaking to the selection mechanism that produced the immigrant population: immigrants were "more able and more highly motivated."

Although I am an immigrant, having left Cuba as a kid in the initial wave of refugees after Castro's 1959 takeover, I had no intellectual interests in any immigration-related issues during my graduate student years. Nevertheless, given the nature of my thesis (extending the human capital model to a context where people changed jobs and acquired general and specific training in various "locations"), I saw the conceptual resemblance with Chiswick's work on immigration. Inevitably, I began to wonder about how the difficult endogeneity issues that plagued my thesis would show up in the study of immigrant earnings.

At the time, I knew little about the history of immigration to the United States, but I was vaguely familiar with recent trends in Cuban migration and wondered whether the strong positive correlation between earnings and years-since-migration reported in Chiswick's study might be capturing something other than assimilation. I had heard that the first wave of Cuban immigrants who arrived immediately after the communist takeover was composed of

the entrepreneurs and intellectual elites who lost substantially from the regime change, while the later waves of immigrants were more representative of the Cuban population. This information formed the basis of a question that I asked at the seminar: What does the correlation between earnings and years-since-migration really mean if there are differences in the human capital stock that different waves of immigrants bring to the United States?

I have no recollection of how the question was answered. In fact, I could never have predicted that the question would linger in my mind for years before I would do anything about it. Nevertheless, I found the topic intriguing enough that I would teach the Chiswick paper in my labor economics classes after I left Columbia, and sometimes re-ask my question in midterm or final exams. It is again worth noting there was little policy-related reason to have spent any time discussing this topic in class in the late 1970s and early 1980s. The paper showed up in my lectures and exams because it was intellectually close to my interests: how does one think about the evolution of earnings over the life cycle when workers are in different markets at different points in time?

The class lectures and exam questions—plus the fact that I was then living in Southern California, where large waves of immigrants were beginning to settle—kept the topic fresh in my mind and I realized that the solution to my seminar query required that we track specific waves of immigrants over time. That tracking would help determine if the earnings of specific immigrants indeed grew faster than the earnings of natives. I decided to wait until the 1980 census data file was publicly released and resolved that I would tackle the question immediately thereafter. I would select specific cohorts of immigrants, find them in both the 1970 and 1980 censuses, and track them to see what wage growth the cohorts had actually experienced.

Once I got the 1980 census data all set up, it took all of one morning to discover that the implications of the Chiswick cross-section regression model and the implications of the tracking exercise I had long envisioned were very different. I vividly remember looking at the initial regression output on a computer screen in my office, seeing how the two sets of regressions led to very different results, and thinking

"What in the world!" (though those were not quite the words that came to mind). Little did I know that the regressions I ran that morning would shape my career and cement my life-long interests in investigating what economics has to say about immigration.

I wrote up my empirical results in the paper "Assimilation, Changes in Cohort Quality, and the Earnings of Immigrants," which was published in the *Journal of Labor Economics* in 1985 (see Chapter 1). That study used the 1970 and 1980 census cross-sections to make the point that because different cohorts of immigrants differed in their "cohort quality," a cross-section correlation between earnings and years-since-migration provided no information whatsoever about the assimilation process. In fact, I discovered that part of the reason for the fact that earlier cohorts earned more was that the earlier cohorts had more human capital at the time of arrival. Instead of rapid assimilation, the panel analysis revealed a decline in the skills of immigrants across successive waves from about 1950 through 1980.

My interest in this topic has never gone away. I made a point of returning to it after every new decennial census was released, simply because I was curious to find out what the latest data showed. Borjas (1995), for example, added the 1990 census cross-section to the panel, allowing the tracking of specific immigrant cohorts over a 20-year period (from 1970 to 1990). The availability of three cross-sections gave me the opportunity to build a more formal model of the assimilation and cohort effects. This model has now been estimated in many other contexts: the United States in the early 20th century (Abramitzky *et al.*, 2014); Germany (Dustmann, 1993); Israel (Eckstein and Weiss, 2004); Sweden (Edin *et al.*, 2000); and a comparison of Australia, Canada, and the United States (Antecol *et al.*, 2006).

I returned to this topic once again in the paper titled "The Slowdown in the Economic Assimilation of Immigrants: Aging and Cohort Effects Revisited Again," published in the *Journal of Human Capital* in 2015 (see Chapter 2). By then, I had access to several more decennial censuses, and I could track specific immigrant waves over 30 or

40 years. The long length of the panel produced a new empirical finding. My early work focused on cohort differences in the *level* of entry wages. The longer panel revealed that not only were there cohort differences in entry wages, but that there were also cohort differences in the rate of growth of earnings. Specifically, the immigrants who entered the United States in the 1960s or 1970s experienced much faster wage growth than the immigrants who entered the country in the 1980s and 1990s.

Despite the large number of studies on economic assimilation published since the 1980s, the research continues apace, and new hypotheses keep being introduced and explored in a continuing attempt to understand what the tracking of immigrant earnings reveals about the assimilation process. Albert *et al.* (2020), for example, examine the hypothesis that the shape of an immigrant's age-earnings profile depends on the volume of immigration at the time the immigrants entered the country as well as on the size of the immigrant flow in the years following their entry. In short, the analysis of immigrant assimilation must also incorporate the issues that come up in the study of the labor market impact of immigration (a subject that will be discussed at length in what follows).

II. Self-Selection of Immigrants

After I finished my 1985 "Assimilation and Changes in Cohort Quality" paper, I realized that the study had raised and left unanswered a major question: *Why* did the skills of immigrants decline across successive waves? I immediately turned my attention to answering that question. It became obvious to me that the answer must rely on some variation of "different types of immigrants are choosing to come to the United States."

And this is where my Columbia background came in quite handy. Being Heckman's research assistant familiarized me with the self-selection problem. Moreover, during one of my years at Columbia, Sherwin Rosen visited the department and taught a graduate-level labor economics class that focused exclusively on the theoretical and conceptual underpinnings of the field. Rosen, too, was interested in

self-selection issues and devoted at least one entire lecture to the Roy model (Roy, 1951) that was then being slowly recognized as a landmark work in labor economics (although it had been totally ignored at the time of publication). In its original iteration, the Roy model addressed a simple and profound question: why do some people choose to become hunters and other people choose to become fishermen?

Despite my early introduction to these models and methods, I have to admit that the application of the Roy model to immigration was far from obvious as I set out to tackle the unanswered question from my 1985 paper. In fact, I detoured into a bunch of dead ends for about a year as I tried model after model and regression after regression trying to understand which factors drove the decline in immigrant skills across cohorts.

I was driving from Santa Barbara to Santa Monica on the Pacific Coast Highway one morning and totally out of the blue the Roy model comes to mind and I immediately see its relevance for immigrant self-selection. In the context of immigration, the basic insight is that the rate of return to skills is the key variable that determines the location where people would rather work. If the source country's labor market offers a very low rate of return to skills, highly skilled workers will want to move to a place where those skills are better rewarded, and the migrant flow will be positively selected. In contrast, if the source country's labor market offers a very high rate of return to skills, highly skilled workers will have little incentive to move elsewhere and the migrant flow will be negatively selected. By the time I got to Santa Monica, I knew that I had finally gotten a handle on the problem that had been pestering me and that I had a model to derive (which I did by the end of the day!), some regressions to run, and a paper to write. The paper, titled "Self-Selection and the Earnings of Immigrants," was published in the *American Economic Review* in 1987 (see Chapter 3).

Ideally, the empirical analysis would have compared the skills of the workers who chose to stay in a particular country of origin with the skills of the compatriots who chose to move to the United States. Depending on how the rate of return to skills in the source country

compared to the rate of return in the United States, the immigrant flow would be either positively or negatively selected. The data needs for this kind of transparent exercise, however, are stringent, requiring information on the skills of both movers and stayers for a particular country of origin. In the mid-1980s, such data did not exist.

I instead pursued a second-best strategy for the empirical analysis. Using the data reported in the United States census, I estimated the correlation between the earnings of immigrants in the United States and the rate of return to skills in the source country. The Roy model predicts that the correlation should be negative: more skilled workers should originate in countries where the rate of return to skills is low. The data confirmed the presence of this negative correlation (although it was weak). That correlation helps give a partial answer to the question raised by my 1985 paper. The earnings of successive immigrant waves declined because of a change in the national origin mix of the immigrant population between 1950 and 1980, away from the industrialized countries of Western Europe where the rates of return to skill tend to be low towards the developing countries where the rates of return to skill tend to be high.

My 1987 self-selection paper has influenced many subsequent studies in the literature, mainly because it provides an intuitive and empirically falsifiable hypothesis. Armed with the Roy model, we should be able to predict which subset of persons will move from country A to country B. The most influential of these studies (in fact, the first to conduct a proper empirical test of the model) was the paper by Chiquiar and Hanson (2005), which focused on the selection characterizing the large flow of Mexican immigrants to the United States.

The Chiquiar–Hanson study compared the education of Mexicans who move and are observed in the United States Census with the education of Mexicans who stay and are observed in the Mexican census. The rate of return to education in Mexico is relatively high, so that the Roy model predicts that the Mexicans who move would be relatively less educated than the Mexicans who stay. The Chiquiar–Hanson analysis revealed that the immigrant flow from Mexico was neither positively nor negatively selected. The Mexican immigrants

seem to lie in the middle of the Mexican education distribution. The authors conjecture that perhaps migration costs are sufficiently high that the least-educated Mexicans are liquidity-constrained and simply cannot afford the move even though they might have an incentive to do so.

One problem with the data used in the Chiquiar–Hanson study, which contaminates the evidence, is the presence of large numbers of Mexican undocumented immigrants in the United States, many of whom do not appear in the United States Census data (so that the education distribution of the movers is measured with error). Several papers have revisited the question using data from exclusively Mexican sources, which observe movers *before* they move. The empirical analysis in some of these studies tends to be more supportive of the Roy model, with the evidence indicating that the Mexicans who choose to move earn substantially less than those who stay (see Ibarraran and Lubotzky, 2007; Fernández-Huertas Moraga, 2011; and Kaestner and Malamud, 2014).

The literature is now bursting with papers that test the predictions of the Roy model in various contexts. The selection of migrants from Puerto Rico to the United States seems consistent with the model's predictions (Ramos, 1992), as is the migration of German university graduates to various countries (Parey *et al.*, 2017), the migration of Italian workers from southern to northern Italy (Bartolucci *et al.*, 2018), and the migration of urban Norwegians to the United States in the early 20$^{\text{th}}$ century (Abramitzky *et al.*, 2012). In contrast, the emigration of workers from Israel to the United States (Gould and Moav, 2016) and the sorting of foreign-born workers across Organisation for Economic Co-operation and Development (OECD) countries (Grogger and Hanson, 2011) does not fully support the model's predictions.

I have returned to the analysis of self-selection issues several times during my career. It is well known that many immigrants eventually decide to return to their source country after "testing the waters" elsewhere and finding that things did not work out as they expected, or perhaps they intended to migrate temporarily for a few years with the full intention of returning after accumulating sufficient wealth.

Bernt Bratsberg, who was my doctoral student, and I began to think about the self-selection of the return migrants in the late 1980s. That collaboration led to the paper titled "Who Leaves? The Outmigration of the Foreign-Born," published in the *Review of Economics and Statistics* in 1996 (see Chapter 4).

The Roy model has clear implications about which types of immigrants choose to return home. The return migrants should be the "marginal" immigrants, those who had the least to gain from the initial emigration. As a result, any unexpected negative shock in their new homes encourages them to reconsider their decision and drives some of them back. The return migrants, therefore, would be the least skilled if the initial flow was positively selected, or the most skilled if the initial flow was negatively selected. In short, return migration accentuates the initial selection. The Borjas–Bratsberg study documented that this type of selection neatly characterized return migrants from the United States. Subsequent work shows that the model's predictions are also consistent with the return migration flows to Egypt (Wahba, 2015) or from Finland to Sweden (Rooth and Saarela, 2007).

More recently, Ilpo Kauppinen, Panu Poutvaara, and I revisited both the theoretical underpinning and empirical testing of the Roy model (Borjas *et al.*, 2019). We had access to administrative data from Denmark that reported the history of earnings for the entire Danish population and specifically identified the subsample of persons who eventually emigrated out of Denmark. The richness of the data motivated us to extend the Roy model by introducing the statistical concept of stochastic dominance, which allows the ranking of two alternative distributions, into the framework. The Roy model predicts a specific ranking of the income distributions of movers and stayers. In the context of Denmark, where the rate of return to skills is low relative to potential destinations, the movers are positively selected and the cumulative distribution of the (pre-migration) earnings of the movers is to the right of (or stochastically dominates) the cumulative distribution of the earnings of stayers.

III. Ethnic Capital

During the 1980s, as I continued my study of immigration issues within the context of an economic approach, I also began to get a broader perspective about what "really matters." It became evident that the focus of the assimilation literature on the evolution of earnings during an immigrant's life cycle was partly missing the point. What was really important was the *long-term* assimilation process, as defined by the outcomes experienced by the children and grandchildren of immigrants.

My first investigation of the intergenerational mobility process used various censuses to track particular national origin groups across generations (Borjas, 1993a). For example, we could use the 1980 Census to calculate the average earnings of Mexican immigrants, and then jump forward to the 2010 Census to calculate the average earnings of United States-born workers who had at least one Mexican-born parent. Not surprisingly, those immigrant groups with the highest earnings in 1980 had the highest-earning descendants 30 years later. This "intergenerational correlation" was very strong, about 0.6, implying that about 60 percent of the earnings variation across immigrant groups in the first generation persists into the second.

The numerical value of the intergenerational correlation between the average earnings of first- and second-generation ethnic groups was higher than expected. Typically, the correlation between the earnings of parents and children lies somewhere between 0.3 and 0.4 (Solon, 1999), so that the intergenerational correlation at the ethnic group level was noticeably stronger.

To explain this discrepancy, I introduced the hypothesis that ethnicity *per se* may play an independent role in intergenerational mobility, acting as an externality in the production of human capital. This hypothesis was developed in a paper titled "Ethnic Capital and Intergenerational Mobility," published in the *Quarterly Journal of Economics* in 1992 (see Chapter 5). In particular, I argued that the skills of the next generation depend not only on parental inputs (which is the typical explanation for the correlation in earnings between

parents and children), but also on the overall quality of the ethnic environment in which the children grow up, or ethnic capital.

The introduction of ethnic capital into an economic model of intergenerational mobility has one important implication: if the external effect of ethnicity is sufficiently strong, ethnic differences in skills observed in this generation may persist for several generations. This persistence would then explain why the intergenerational correlation in the skills of immigrant groups might be larger than the simple correlation between individual parents and children.

The empirical analysis reported in the paper related the skills of a particular child, measured in terms of educational attainment or earnings, to both the skills of the parent and to the average educational attainment or earnings of the ethnic group in the parental generation (which was the empirical proxy for ethnic capital). The regressions suggested that about half of the 0.6 intergenerational correlation is attributable to the link between parents and children, while the other half is attributable to the ethnic capital effect.

The ethnic capital hypothesis has been applied to explain a number of empirical regularities that produce ethnic stickiness in various characteristics across generations. These applications include the role of ghettos in the transmission of ethnic capital (Cutler *et al.*, 2005); the intergenerational transmission of education in ethnic settings (Chakraborty *et al.*, 2019); and the link between ethnicity and self-employment (Wang and Maani, 2014).

As soon as my initial ethnic capital study was completed, I began to worry about what exactly the concept of ethnic capital was capturing. In other words, what is it about ethnicity that leads to stickiness in human capital accumulation across generations and produces such a high value of the intergenerational correlation in the average earnings of ethnic groups? It seemed obvious, for instance, that the ethnic capital externality might work through the ethnic neighborhood, which has long been a dominant feature of cities. The neighborhoods cluster people of similar backgrounds, and foster a set of cultural attitudes, social contacts, and economic opportunities that affect workers throughout their lives. The ethnic capital effect then

arises because children are frequently exposed to a specific set of ethnic characteristics as they grow up.

In "Ethnicity, Neighborhoods, and Human Capital Externalities", published in the *American Economic Review* in 1995 (see Chapter 6), I extended the analysis of ethnic capital to directly link neighborhoods, ethnic segregation, and intergenerational mobility. I obtained data that provided detailed information on the zip code of residence for a sample of young workers, making it possible to determine the extent to which ethnic groups segregate in particular neighborhoods and to measure the distinct impact of this segregation on the process of human capital accumulation. The main finding of the neighborhood study was that residential segregation and the influence of ethnic capital are intimately linked. In particular, the impact of ethnic capital on the skills of the next generation arises partly because the ethnic capital variable proxied for the socioeconomic background of the neighborhood where the children were raised, and these neighborhood characteristics influenced intergenerational mobility.

A natural experiment conducted in Sweden confirms that the skill composition of the ethnic enclave influences the economic outcomes of immigrant children (Åslund *et al.*, 2011). Between 1985 and 1994, Sweden randomly placed new refugees across a large number of cities, rather than letting them settle in the traditional gateways, such as Stockholm and Malmö. By chance, some refugees ended up in neighborhoods where their ethnic compatriots were a bit more skilled. It turns out that the children that were randomly assigned to neighborhoods where a higher fraction of the compatriots were better educated ended up doing far better in school.

IV. Labor Market Impact

As far as I know, Grossman (1982) was the first paper to empirically measure the impact of immigration on the receiving country's labor market. This pioneering study conjectured that it should be possible to measure the wage impact of immigration by comparing the evolution of wages in labor markets that are affected

differentially by immigrant flows. The earnings of substitutable workers, for instance, should decline more in those markets that had a larger immigrant influx. Grossman operationalized her insight by introducing an assumption that greatly influenced subsequent work. She assumed that a labor market was delineated by the boundaries of a metropolitan area. The empirical analysis, therefore, consisted of estimating a "spatial correlation," which compared wage outcomes in cities that received many immigrants to the outcomes in cities that received few. Grossman's estimate of the spatial correlation suggested that the labor market effects of immigration were relatively small.

I began to work on this issue shortly after Grossman's paper was published, and the results from those initial forays are summarized in Borjas (1987). Like Grossman, I also found a weak correlation between immigration and wages across cities (in the context of a Generalized Leontief production function framework), but I was not convinced by the results. I suspected that immigrants are rational workers who will most likely settle in high-wage cities, and this built-in positive correlation would make it difficult to isolate whatever negative impact the supply shock may have had on native wages. Moreover, natives, too, are rational workers. They will respond to the wage cut following an immigrant supply shock by moving elsewhere, effectively diffusing the impact of immigration to other localities and suggesting that the local labor market may be the wrong place to look for the impact of immigration.

My dissatisfaction with the framework, and the fact that I could not figure out how to get rid of the built-in endogeneity in a credible way, led me to shelve that part of my research program for a few years. During the hiatus, Altonji and Card (1991) proposed addressing the endogeneity issue by using lagged immigration as an instrument for current immigration. Although this instrument (or some variation thereof) is commonly used in the literature and now goes under the name of a "shift-share" instrument, I have always had my doubts. The justification for the lagged instrument is that new immigrants tend to settle in the same places where earlier waves settled, perhaps because of lower information costs associated with ethnic networks. However, the settlement decision of the earlier immigrant

waves was also endogenous—perhaps those pioneering waves also wanted to settle in high-wage places. Any serial correlation in wages at the local level would invalidate the instrument.

This concern over the validity of the shift-share instrument continues to inspire a lot of research. Jaeger *et al.* (2018) argue that because local wages are correlated over time, the shift-share instrument captures both the impact of the current immigrant wave (which is likely to be negative) as well as the impact of earlier immigrant waves (which is likely to be positive because the local economy is adjusting to earlier flows).

A conversation with my Harvard colleague Richard Freeman at the annual American Economic Association meetings sometime in the early 1990s rekindled my interests. We were organizing an immigration conference for the National Bureau of Economic Research and he proposed that we do some work on this issue along with Larry Katz. The collaboration culminated in the paper titled "How Much Do Immigration and Trade Affect Labor Market Outcomes," published in the *Brookings Papers on Economic Activity* in 1997 (see Chapter 7). That paper provides a frontal critique of the spatial correlation approach that had, by then, become entrenched in the literature. Building on the concurrent work on changes in the United States income distribution and wage structure (Katz and Murphy, 1992), the paper used a simple supply–demand framework to simulate the impact of immigration in the *national* labor market. Using estimates from the wage structure literature for the elasticity of substitution between different types of skill groups, we concluded that, despite what the existing estimates of the spatial correlation suggested, immigration may have had a substantial impact on the wage of low-skill native workers.

There was one element in the Borjas–Freeman–Katz approach that left me unsatisfied. Our empirical evidence resulted from a *simulation* of a supply–demand model for workers classified into two skill groups (high and low). I would have felt more comfortable if the evidence had been based on an actual analysis of data rather than on feeding "off-the-shelf" elasticity estimates to a mathematical model of the labor market. This concern stayed with me for

a few years, as I kept trying to find a way to retain the national labor market setting but produce empirical results using a regression framework.

Once again, my Columbia training paid off. Jacob Mincer's emphasis on labor market experience as an important component of human capital suggests that we define skill groups not only in terms of education, but also in terms of labor market experience. Young immigrants would more likely compete with young natives, and older immigrants would likely compete with older natives. The education–experience breakdown allowed for the demarcation of a large number of distinct skill-based labor markets. That insight led to the paper "The Labor Demand Curve *Is* Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," published in the *Quarterly Journal of Economics* in 2003 (see Chapter 8).

The 2003 paper has two distinct sections. The first half is a purely descriptive analysis, analogous to the empirical spatial correlation literature, that correlates wages and immigration across labor markets. Instead of defining labor markets in geographic terms, I exploited wage differences across national labor markets differentiated by skills. The descriptive analysis in my 2003 paper revealed a strong negative correlation between wages and immigration across these skill-based labor markets.

The second half of the paper introduced a structural model that allows for the presence of cross-effects when estimating the wage impact of immigration. Much of the literature has (wrongly) focused on how the entry of a particular type of immigrant, say low-educated workers, affects the wage of natives in the same skill group. But there are cross-effects as well: The immigration of low-educated workers also affects the wage of high-educated workers.

Because of the large numbers of potential skill groups defined by education and experience (there were 32 such groups in my 2003 paper), it is impossible to estimate all the cross-effects unless the analysis imposes some structure on how different types of workers interact. This is where the assumption of an aggregate production function comes in. I assumed that the technology could be

represented by a nested CES framework. This technical assumption allowed for easy measurement of the elasticity of substitution across education groups and across experience groups. The paper then showed how these elasticities of substitution could be used to simulate the impact of any potential supply shock on the wage structure.

One aspect of the 2003 paper that some readers have noted is that the title emphasizes that "The Labor Demand Curve *Is* Downward Sloping," rather than the catchier, policy-relevant, and more provocative result that immigrants have a negative impact on native wages. The title choice was endogenous and directly reflects what drove my research interests. What I always found interesting about studying immigration is what it teaches about the way labor markets work, rather than the immediate policy concern of whether immigrants are assimilating quickly or slowly, or whether the wage effect is small or large. The evidence may have policy implications and will surely be praised or belittled by advocates (typically out of context) to weave a particular narrative. But the policy concern was not the inspiration for the paper nor what motivated me to spend countless hours working out various models and processing large quantities of data. The 2003 paper is nothing more than the product of my years-long attempt to understand where the spatial correlation literature was going wrong and what could be done to set it right.

The analysis in the paper has been replicated in many other contexts, including estimating the descriptive skill-cell regression using data for other countries such as Germany (Steinhardt, 2011), Mexico (Mishra, 2007), Norway (Bratsberg *et al.*, 2014), and Spain (Carrasco *et al.*, 2008). The estimates of the wage elasticities for the United States have also been refined by taking into account the endogeneity of particular types of immigrant flows, using instruments that exploit international variation in distance from the United States, as well as such push factors as wars, natural disasters, and political regime changes (Llull, 2018a). These instruments, which seem far more exogenous than the shift-share instrument, tend to generate substantially more negative wage effects.

There has also been work extending the nested CES framework. Ottaviano and Peri (2012) and Manacorda *et al.* (2012) argue that the structural framework in my 2003 paper ignored the possibility that there could be complementarity between observationally equivalent immigrants and natives. In other words, immigrants with the same education and experience as natives need not be perfect substitutes because they work in different occupations (Peri and Sparber, 2009; Llull, 2018b), or because employers downgrade the skills immigrants acquired prior to migration (Dustmann *et al.*, 2013).

Card (2009) focuses on a different elasticity of substitution, arguing that high school graduates and high school dropouts are perfect substitutes. This particular assumption greatly affects the results of a simulation exercise that predicts the wage impact of a supply shock. Even though low-skill immigration led to a very large percent increase in the number of workers without a high school degree, the relative impact of low-skill immigration on supply is far smaller if the low-skill workforce is artificially expanded by including the very large number of native high school graduates.

These studies show that the simulation of wage impacts in a nested CES framework is sensitive to the specific assumptions used to restrict the set of potential production interactions among groups. Moreover, the estimates of the various elasticities of substitution are sensitive to sample definitions and regression specifications. As an example, the Ottaviano–Peri (2012) evidence on complementarity between observationally equivalent immigrants and natives would vanish if the authors had followed the convention in the literature and used the mean log wage of a skill cell as the measure of productivity rather than the unusual log mean wage (Borjas *et al.*, 2012). The sensitivity has inspired a search for alternative ways of allowing for cross-effects without imposing the restrictions of the nested CES framework. Dustmann *et al.* (2013), for example, estimate the differential wage impact of a particular supply shock at every point of the wage distribution, while Llull (2018b) explicitly models how immigration influences the occupational choice of natives.

My 2003 labor market impact paper raised a new puzzle: The spatial correlation and national labor market approaches led to very different findings about how native wages respond to immigration. As noted earlier, even putting aside the fact that immigrants do not settle randomly across cities, the spatial correlation approach still faced the problem that local labor markets are not closed islands. An influx of (substitutable) immigrants into a city will lower that city's wage and encourage native workers to settle elsewhere. This diffusion of the local impact of immigration throughout the national labor market could produce the finding that even though immigration and wages are not strongly correlated across cities, immigration still has a substantial wage effect at the national level.

The natural next step, therefore, would be a synthesis that would jointly examine the wage impact of immigration and the internal migration of native workers. My attempt at such a synthesis led to the paper titled "Native Internal Migration and the Labor Market Impact of Immigration," published in the *Journal of Human Resources* in 2006 (see Chapter 9). Not surprisingly, the spatial correlation turns out to be an amalgam of two different parameters: the wage elasticity that captures the short-run wage impact and the crowd-out effect that gives the number of natives who avoid a particular labor market for every immigrant who moves there. Moreover, the theoretical framework made it possible to convert the measured spatial correlation into the wage elasticity by using information on the size of the crowd-out effect.

The joint analysis of immigration and native internal migration has guided recent research that attempts to further elucidate what spatial correlations truly measure. Amior (2020), for example, examines local labor market dynamics after immigrant supply shocks, and cannot reject the hypothesis that immigration crowds out native population growth on a one-to-one basis. Dustmann *et al.* (2017) examine labor market flows from Czechoslovakia to Germany after the fall of the Iron Curtain and find that much of the local wage impact was attenuated by a large internal migration response, a response that was mostly due to a reduction of native inflows into immigrant-targeted areas rather than larger outflows.

I also became interested in the labor market impact of high-skill immigration, particularly after a conversation with a friend I made at the many Little League and scouting activities we both attended in the late 1990s. My friend had a Ph.D. in mathematics from MIT and had found it difficult to get an academic job because the collapse of the Soviet Union had sent a large number of first-rate Soviet mathematicians into the American labor market.

About a decade after that conversation, I met Kirk Doran, an economist at the University of Notre Dame, who personally knew people with similar experiences. We soon discovered that the American Mathematical Society maintains a data archive that records every single research paper published by every single mathematician in the world, so that we could track the careers of American mathematicians before and after the Soviet supply shock. Due to the separation imposed by the Cold War, Soviet and American mathematicians had gone in different directions for several decades. We then measured the impact of the Soviet supply shock on two different groups of American mathematicians: those who specialized in "Soviet fields," like differential equations, and those who specialized in "American fields," like statistics.

It was hard to predict how things would turn out. On the one hand, the presence of the Soviet mathematicians could increase the productivity of the Americans working on differential equations—the Americans were then exposed to theorems they did not know before. On the other hand, those same American mathematicians had to compete with the Soviet émigrés for jobs and publication space.

The resulting paper, titled "The Productivity of American Mathematicians and the Collapse of the Soviet Union," was published in the *Quarterly Journal of Economics* in 2012 (see Chapter 10). It turns out that the publication rate of the group of American mathematicians whose research overlapped most with that of the Soviets declined precipitously after 1990. Put simply, the American mathematicians who competed most directly with the Soviet émigrés lost out. They found it harder to hold onto their academic jobs and to produce publishable research.

It is crucial to determine the net impact of high-skill immigration because this particular type of immigration represents a potential source of substantial economic benefits to natives. High-skill immigrants might move the knowledge frontier outwards, increasing the productivity of the entire workforce. A large literature attempts to document the existence of this spillover effect, measuring the impact of the "high-tech" H1-B program (Kerr and Lincoln, 2010; Doran *et al.*, 2016); the impact of the migration of Jewish scientists from Nazi Germany on the productivity of American scientists (Moser *et al.*, 2014), and how the exodus of those Jewish scientists affected the productivity of the students and colleagues left behind (Waldinger, 2010, 2012). The evidence is far from clear-cut. High-skill migration can sometimes be highly beneficial. At the same time, it can also impose a substantial cost on the high-skill workers who most directly compete with the immigrants.

My most recent work on the labor market impact of immigration has focused on the consequences of the Mariel boatlift. On April 20, 1980, Fidel Castro declared that Cuban nationals wishing to move to the United States could leave freely from the port of Mariel. About 125,000 Cubans left before an agreement between the Carter administration and the Castro regime closed the escape valve in October 1980. The flow was composed of relatively low-skill workers. Nearly 60 percent of the adult refugees lacked a high school diploma. Although the Mariel supply shock increased the total number of workers in Miami by about 8 percent in about two or three months, it increased the number of the least-educated workers by almost 20 percent.

The labor market impact of the Mariel boatlift was first studied in Card's (1990) classic paper. Card's comparison of the labor market in pre- and post-Mariel Miami relative to the change observed in a set of cities that served as a control group indicated that little changed in Miami despite the large number of new workers. The average wage in Miami did not fall, and the unemployment rate remained unchanged relative to what was happening elsewhere.

My interest in revisiting the Mariel episode was sparked when a close (non-economist) acquaintance read an early draft of my book,

We Wanted Workers: Unraveling the Immigration Narrative (Borjas, 2016), told me that all the models in the world were no match for the transparency of the Mariel episode, and wondered if the evidence in the Card paper was as clear-cut as everyone seemed to believe. I started looking more closely at the original study and quickly saw that Card's empirical analysis had not examined the earnings of high school dropouts, the group most likely to be affected by the supply shock. To make matters worse, Card had assembled the control group of cities based partly on the fact that they resembled Miami *after* Mariel. I quickly decided to reanalyze the data, leading to the paper titled "The Wage Impact of the *Marielitos*: A Reappraisal," which was published in the *Industrial and Labor Relations Review* in 2017 (see Chapter 11).

The reappraisal almost instantly (literally within an hour after I started looking at the data) produced a result that seems to have been overlooked all these years. The average wage of the least educated workers in Miami (many of whom were African American) fell sharply after 1980 and did not recover fully until about 1990. The magnitude of the adverse wage impact was substantial. The wage of workers without a high school diploma in Miami (relative to the control cities) dropped by as much as 30 percent by 1985.

Not surprisingly, the very different set of results reported in my 2017 paper produced a surge of academic interest in the labor market consequences of the Mariel episode. The flurry of subsequent papers includes Peri and Yasenov (2019) and Clemens and Hunt (2019). These studies show that the evidence is sensitive to the definition of a "low-skill" worker and to how the analyst adjusts for changes in the racial composition of the sample. Monras (2021) extends the analysis in an important way by showing that about half of the recovery in the wage of high school dropouts in Miami by the end of the 1980s can be attributed to an internal migration response by low-skill native workers, who chose to relocate in other cities.

An important lesson from the conflicting evidence is that the labor market data available to study the Mariel episode is quite limited—in particular, the number of workers sampled in Miami by the Current Population Surveys is small. The data constraint has encouraged a

search for alternative sources of information about Miami's labor market at the time. Some of my recent work (Anastasopoulos *et al.*, 2021) illustrates the promise of these alternative data sources by looking at the number and textual content of help-wanted ads placed in the *Miami Herald* before and after Mariel, which reveal a substantial post-Mariel drop in the number of help-wanted ads for low-skill positions.

V. Economic Benefits

The last aspect of the economics of immigration that interested me through the years was the calculation of the benefits from immigration. As an increasing number of economists began producing estimates of the spatial correlation that presumably measured the adverse impact of immigration on wages, it occurred to me that those estimates (even if they were correct) were only telling half the story. Economic theory suggests that the wage drop will produce some benefits as well. The intuition is similar to that employed for justifying the gains from international trade. The gains from the lower labor costs produced by immigration accrue to everyone who uses those workers, from employers who hire the cheaper workers, to the families looking for help to run the household, to consumers buying the cheaper goods and services. As with trade, the dollar gains accruing to the natives whose gain *must be* numerically larger than the dollar losses suffered by the natives who lose. On net, therefore, immigration increases the wealth of natives, and this gain is called the "immigration surplus."

I began to teach the basic model in my classes and would sometimes mention it to colleagues as I traveled giving lectures at various universities. The idea (and a very preliminary calculation) even makes a one-paragraph appearance in the "Immigration" entry I wrote for the *Fortune Encyclopedia of Economics* (Borjas, 1993b). I was always struck by the fact that everyone seemed intrigued. The motivation for actually sitting down and deriving the formal framework came after Alan Krueger asked me (at a dinner after I had just presented a seminar at Princeton) if I would organize an immigration

symposium for a journal he edited at the time. The derivation of the immigration surplus (within the framework of a competitive labor market) and my attempt at producing numerical estimates of the gains appeared in a paper titled "The Economic Benefits from Immigration," published in the *Journal of Economic Perspectives* in 1995 (see Chapter 12).

As I was working out the model and plugging numbers into formulas, I was surprised that no matter how hard I tried, the estimated surplus was small. The latest estimates (as of 2015) suggest that current immigration, which increased labor supply by over 16 percent, raised the short-run wealth of natives by about 0.3 percent, or $50 billion annually (Borjas, 2016, p. 155). The explicit calculation of the gains from immigration reveals that the standard competitive framework cannot yield large estimates of the immigration surplus without a major change in assumptions, such as the presence of human capital spillovers produced by the entry of high-skill immigrants.

The concept and measurement of the immigration surplus has received substantial attention in the subsequent literature. Ben-Gad (2008) derives the surplus using a growth model; Amuedo-Dorantes and de la Rica (2013) estimate the surplus for the Spanish labor market; and Amior and Manning (2020) discuss the surplus in the context of a monopsony model.

Shortly after I worked out the mechanics of the immigration surplus, I was appointed to a National Academy of Sciences panel that would evaluate the costs and benefits from immigration. One specific charge to the panel was the calculation of the net *fiscal* impact. Superficially, the exercise seems trivially simple: add up the taxes immigrants pay and subtract the cost of the public services they receive. I quickly realized, however, that the devil is in the details. The National Academy has produced two sets of estimates of the fiscal impact since the 1990s (Smith and Edmonston, 1997; and Blau and Mackie, 2017), and both times the *sign* of the fiscal impact depends entirely on assumptions. How long is the long run (for example, 25 years, 50 years, or 300 years)? What are the assumed paths of future federal expenditures and revenues? How much should immigrants be charged for their consumption of public goods, such as police protection?

I dislike "playing with assumptions" games, so I was never tempted to waste my time doing calculations along the lines suggested by the National Academy exercise. But I became interested in the interaction between the welfare state and immigration. My attention instead turned to an issue that emphasized the economics of the problem: Does the huge interstate dispersion in welfare benefits affect the location choices of new immigrants?

The economic approach has a clear-cut prediction about how immigrants should behave. Persons born in the United States often find it expensive to move across states. Suppose that migration costs are mainly fixed costs—and that these fixed costs are high. The existing differences in welfare benefits across states may not produce a large number of native movers because the interstate benefit differentials might be swamped by the migration costs. In contrast, low-income immigrants are a self-selected sample of persons who chose to bear the fixed costs of the move. Suppose that once the costs of moving to the United States are incurred, it costs little to choose one particular state over another. The sample of newly arrived immigrants will then tend to live in the "right" state. Borjas (1999) examined this hypothesis empirically and documented that newly arrived immigrants who were likely to qualify for public assistance indeed clustered in those states that offered the most generous welfare benefits.

As I was working out the details of the welfare magnet hypothesis, I realized that the idea also had labor market implications, perhaps generating a new source of economic benefits. Labor market efficiency requires that the wage (for a given type of worker) be equalized across labor markets, such as cities or states.

Suppose again that migration costs are fixed costs. The existing wage differentials across states may not induce large numbers of native workers to move because the (high) migration costs swamp interstate differences in income opportunities. As a result, native internal migration will not arbitrage interstate wage differentials. However, new immigrants have already decided to incur the fixed cost of the move to the United States. If it costs little to choose one particular state of destination over another, income-maximizing immigrants will choose the destination that offers the best opportunities.

The new immigrants will then settle in the state that offers the highest wage. As a result, the immigrant population may play a disproportionately large role in helping the national labor market attain an efficient allocation of resources.

The paper that explores this hypothesis, titled "Does Immigration Grease the Wheels of the Labor Market?", was published in the *Brookings Papers on Economic Activity* in 2001 (see Chapter 13). The study not only documented that newly arrived immigrant workers settled in those local labor markets where their skills were best rewarded (akin to a "labor market magnet"), but also developed a model to calculate the benefits accruing to natives as a result of this optimal allocation of immigrant workers across markets. It turned out, however, that the magnitude of the gains from the "greasing the wheels" effect was relatively small, about the same magnitude as the immigration surplus.

The hypothesis that immigrants endogenously settle in those labor markets where they are most productive is now part of the "tool kit" in immigration economics and has been revisited in recent research. For example, Cadena and Kovak (2016) document that Mexican immigrants responded more strongly to the labor demand shock of the Great Recession than low-skilled natives. In particular, low-skilled Mexicans were much more likely to move away from the cities hardest hit by the negative demand shock. This "excess mobility" produced a benefit as it helped attenuate the negative impact of the local shock on the native workforce. Albert and Monras (2019) extend the model by providing another reason that produces immigrant clustering into particular labor markets. Part of the income accruing to immigrants will be remitted to their families in the source countries, and the amount of remittance can be higher when the immigrant population in the United States settles in high-wage cities.

VI. Conclusion

The study of the economics of immigration has matured in the past four decades. We now have a set of models, econometric methods, and datasets that provide a solid foundation for future work.

This introduction described the intellectual journey I traveled as my curiosity and academic interests led me to pursue a specific avenue of immigration-related research which, more often than not, raised questions that led into another avenue of immigration-related research. The studies presented in this volume represent the product of this life-long journey.

I was very fortunate to begin that journey at a time when few people cared about the policy content of the answer my research was giving. In other words, the perceived value of the work did not depend on whether the paper reported that immigration assimilation was fast or slow, whether immigrants were positively or negatively selected, or whether the negative wage impact was sizable or trivial. Instead, it was the discovery that the economic approach produced valuable insights about immigration.

Unfortunately, somewhere along the way, the study of immigration became radioactively entwined with political considerations. This unhealthy coupling made it increasingly difficult for many to separate the intellectual contribution of any particular study from the awareness that the study provided an answer that could be used to advance a particular political agenda. The study of immigration issues and its future direction, however, are far too important to be left to policy advocates.

References

Abramitzky, R., L. P. Boustan, and K. Eriksson (2012). "Europe's Tired, Poor, Huddled Masses: Self-Selection and Economic Outcomes in the Age of Mass Migration," *American Economic Review* 102: 1832–1856.

Abramitzky, R., L. P. Boustan, and K. Eriksson (2014). "A Nation of Immigrants: Assimilation and Economic Outcomes in the Age of Mass Migration," *Journal of Political Economy* 122: 467–506.

Albert, C. and J. Monras (2019). *Immigration and Spatial Equilibrium: The Role of Expenditures in the Country of Origin*, Working Paper, Universitat Pompeu Fabra.

Albert, C., A. Glitz, and J. Llull (2020). *Labor Market Competition and the Assimilation of Immigrants*, Working Paper, Barcelona GSE.

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,* J. M. Abowd and R. B. Freeman (eds.). Chicago: University of Chicago Press, pp. 201–234.

Amior, M. (2020). *The Contribution of Immigration to Local Labor Market Adjustment*, Working Paper, Hebrew University of Jerusalem.

Amior, M. and A. Manning (2020). *Monopsony and the Wage Effects of Migration*, Working Paper, Hebrew University of Jerusalem.

Amuedo-Dorantes, C. and S. de la Rica (2013). "The Immigration Surplus and the Substitutability of Immigrant and Native Labor: Evidence from Spain," *Empirical Economics* 44: 945–958.

Anastasopoulos, L. J., G. J. Borjas, G. G. Cook, and M. Lachanski (2021). "Job Vacancies and Immigration: Evidence from the Mariel Supply Shock," *Journal of Human Capital* 15: 1–33.

Antecol, H., P. Kuhn, and S. J. Trejo (2006). "Assimilation via Prices or Quantities? Sources of Immigrant Earnings Growth in Australia, Canada, and the United States," *Journal of Human Resources* 41: 821–840.

Åslund, O., P.-A. Edin, P. Fredriksson, and H. Grönqvist (2011). "Peers, Neighborhoods, and Immigrant Student Achievement: Evidence from a Placement Policy," *American Economic Journal: Applied Economics* 3: 67–95.

Bartolucci, C., C. Villosio, and M. Wagner (2018). "Who Migrates and Why? Evidence from Italian Administrative Data," *Journal of Labor Economics* 36: 551–588.

Ben-Gad, M. (2008). "Capital-Skill Complementarity and the Immigration Surplus," *Review of Economic Dynamics* 11: 335–365.

Blau, F. D. and C. Mackie (2017). *The Economic and Fiscal Consequences of Immigration*, Washington, DC: National Academy Press.

Borjas, G. J. (1981). "Job Mobility and Earnings over the Life Cycle," *Industrial and Labor Relations Review* 34: 365–376.

Borjas, G. J. (1987). "Immigrants, Minorities, and Labor Market Competition," *Industrial and Labor Relations Review* 40: 382–392.

Borjas, G. J. (1993a). "The Intergenerational Mobility of Immigrants," *Journal of Labor Economics* 11: 113–135.

Borjas, G. J. (1993b). "Immigration." In, *The Fortune Encyclopedia of Economics*. D. Henderson (ed.) New York: Time-Warner Books, pp. 484–490.

Borjas, G. J. (1995). "Assimilation and Changes in Cohort Quality Revisited: What Happened to Immigrant Earnings in the 1980s?" *Journal of Labor Economics* 13: 201–245.

Borjas, G. J. (1999). "Immigration and Welfare Magnets," *Journal of Labor Economics* 17: 607–637.

Borjas, G. J. (2016). *We Wanted Workers: Unraveling the Immigration Narrative*, New York: Norton.

Borjas, G. J., I. Kauppinen, and P. Poutvaraa (2019). "Self-Selection of Emigrants: Theory and Evidence on Stochastic Dominance in Observable and Unobservable Characteristics," *Economic Journal* 129: 143–171.

Borjas, G. J., J. Grogger, and G. H. Hanson (2012). "On Estimating Elasticities of Substitution," *Journal of the European Economic Association* 10: 198–210.

Bratsberg, B., O. Raaum, M. Røed, and P. Schøne (2014). "Immigration Wage Impacts by Origin," *Scandinavian Journal of Economics* 116: 356–393.

Cadena, B. and B. Kovak (2016). "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession," *American Economic Journal: Applied Economics* 8: 257–290.

Card, D. (1990). "The Impact of the Mariel Boatlift on the Miami Labor Market," *Industrial and Labor Relations Review* 43: 245–257.

Card, D. (2009). "Immigration and Inequality," *American Economic Review* 99: 1–21.

Carrasco, R., J. F. Jimeno, and A. C. Ortega (2008). "The Effect of Immigration on the Labor Market Performance of Native-Born Workers: Some Evidence from Spain," *Journal of Population Economics* 21: 627–648.

Chakraborty, T., S. Schüller, and K. F. Zimmermann (2019). "Beyond the Average: Ethnic Capital Heterogeneity and Intergenerational Transmission of Education," *Journal of Economic Behavior and Organization* 163: 551–569.

Chiquiar, D. and G.H. Hanson (2005). "International Migration, Self-Selection, and the Distribution of Wages: Evidence from Mexico and the United States," *Journal of Political Economy* 113: 239–281.

Chiswick, B. (1978). "The Effect of Americanization on the Earnings of Foreign-Born Men," *Journal of Political Economy* 86: 897–921.

Clemens, M. A. and J. Hunt (2019). "The Labor Market Effects of Refugee Waves: Reconciling Conflicting Results," *Industrial and Labor Relations Review* 72: 818–857.

Cutler, D. M., E. L. Glaeser, and J. Vigdor (2005). "Ghettos and the Transmission of Ethnic Capital." In *Ethnicity, Social Mobility, and Public Policy Comparing the USA and UK*, G. C. Loury, T. Modood, and S. M. Teles (eds.), Cambridge: Cambridge University Press, pp. 204–221.

Doran, K., A. Gelber, and A. Isen (2016). *The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries*, NBER Working Paper No. 20668.

Dustmann, C. (1993). "Earnings Adjustment of Temporary Migrants," *Journal of Population Economics* 6: 153–168.

Dustmann, C., T. Frattini, and I. P. Preston (2013). "The Effect of Immigration along the Distribution of Wages," *Review of Economic Studies* 80: 145–173.

Dustmann, C., U. Schönberg, and J. Stuhler (2017). "Labor Supply Shocks, Native Wages, and the Adjustment of Local Employment," *Quarterly Journal of Economics* 132: 435–583.

Eckstein, Z. and Y. Weiss (2004). "On the Wage Growth of Immigrants: Israel 1990–2000," *Journal of the European Economic Association* 2: 665–695.

Edin, P.-A., R. J. LaLonde, and O. Åslund (2000). "Emigration of Immigrants and Measures of Immigrant Assimilation: Evidence from Sweden," *Swedish Economic Policy Review* 7: 163–204.

Fernández-Huertas M. J. (2011). "New Evidence on Emigrant Selection," *Review of Economics and Statistics* 93: 72–96.

Gould, E. D. and O. Moav (2016). "Does High Inequality Attract High Skilled Immigrants?" *Economic Journal* 126: 1055–1091.

Grogger, J. and G. H. Hanson (2011). "Income Maximization and the Selection and Sorting of International Migrants," *Journal of Development Economics* 95: 42–57.

Grossman, J. B. (1982). "The Substitutability of Natives and Immigrants in Production," *Review of Economics and Statistics* 54: 596–603.

Heckman, J. J. (1979). "Sample Selection Bias as a Specification Error," *Econometrica* 47: 153–161.

Ibarraran, P. and D. Lubotsky (2007). "Mexican Immigration and Self-Selection: New Evidence from the 2000 Mexican Census." In *Mexican Immigration to the United States,* G. J. Borjas (ed.), Chicago: University of Chicago Press, pp. 159–192.

Jaeger, D., J. Ruist, and J. Stuhler (2018). *Shift-share Instruments and the Impact of Immigration*, NBER Working Paper No. 24285.

Katz, L. and K. J. Murphy (1992). "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics* 107: 35–78.

Kaestner, R. and O. Malamud (2014). "Self-Selection and International Migration: New Evidence from Mexico," *Review of Economics and Statistics* 96: 78–91.

Kerr, W. R. and W. F. Lincoln (2010). "The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention," *Journal of Labor Economics* 28: 473–508.

Llull, J. (2018a). "The Effect of Immigration on Wages: Exploiting Exogenous Variation at the National Level," *Journal of Human Resources* 53: 608–622.

Llull, J. (2018b). "Immigration, Wages, and Education: A Labor Market Equilibrium Structural Model," *Review of Economics Studies* 85: 1852–1896.

Manacorda, M., A. Manning, and J. Wadsworth (2012). "The Impact of Immigration on the Structure of Wages: Theory and Evidence from Britain," *Journal of the European Economic Association* 10: 120–151.

Mincer, J. (1974). *Schooling, Experience, and Earnings*. New York: National Bureau of Economic Research.

Mishra, P. (2007). "Emigration and Wages in Source Countries: Evidence from Mexico," *Journal of Development Economics* 82: 180–199.

Moser, P., A. Voena, and F. Waldinger (2014). "German Jewish Émigrés and U.S. Invention," *American Economic Review* 104: 3222–3255.

Monras, J. (2021). "Local Adjustments to Immigrant-Driven Labor Supply Shocks," *Journal of Human Capital* 15: 204–235.

Ottaviano, G. and G. Peri (2012). "Rethinking the Effect of Immigration on Wages," *Journal of the European Economic Association* 10: 152–197.

Parey, M., J. Ruhose, F. Waldinger, and N. Netz (2017). "The Selection of High-Skilled Migrants," *Review of Economics and Statistics* 99: 776–792.

Peri, G. and C. Sparber (2009). "Task Specialization, Immigration, and Wages," *American Economic Journal: Applied Economics* 1: 135–169.

Peri, G. and V. Yasenov (2019). "The Labor Market Effects of a Refugee Wave: Synthetic Control Method Meets the Mariel Boatlift," *Journal of Human Resources* 54: 267–309.

Ramos, F. A. (1992). "Outmigration and Return Migration of Puerto Ricans," In *Immigration and the Workforce: Economic Consequences, for the United States and Source Areas,* G. J. Borjas and R. B. Freeman (eds.).Chicago: University of Chicago Press, pp. 49–66.

Rooth, D.-O. and J. Saarela (2007). "Selection in Migration and Return Migration: Evidence from Micro Data," *Economics Letters* 94: 90–95.

Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers* 3: 135–146.

Smith, J. P. and B. Edmonston (1997). *The New Americans: Economic, Demographic, and Fiscal Effects of Immigration.* Washington, DC: National Academy Press.

Solon, G. R. (1999). "Intergenerational Mobility in the Labor Market." In *Handbook of labor economics*, Volume 3A, O. Ashenfelter and D. Card (eds.). Amsterdam: Elsevier, pp. 1761–1800.

Steinhardt, M. F. (2011). "The Wage Impact of Immigration in Germany: New Evidence for Skill Groups and Occupations," *The B.E. Journal of Economic Analysis and Policy* 11.

Wahba, J. (2015). "Selection, Selection, Selection: The Impact of Return Migration," *Journal of Population Economics* 28: 1432–1475.

Waldinger, F. (2010). "Quality Matters: The Expulsion of Professors and the Consequences for Ph.D. Student Outcomes in Nazi Germany," *Journal of Political Economy* 118: 787–831.

Waldinger, F. (2012). "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany. *Review of Economic Studies* 79: 838–861.

Wang, X. and S. A. Maani (2014). "Ethnic Capital and Self-Employment: A Spatially Autoregressive Network Approach," *IZA Journal of Migration* 3.