Forced Displacement and Education

The Impact of Refugees on the Education of Host Country Students

Prepared for MIT Inspire
Submitted on 13 January, 2016
EC223

There was no formal mentor in the completion of this research project. However, a special thanks is owed to Eric A. Hanushek of Stanford University's Hoover Institute for his consultation on computational methods regarding international assessments. Additionally, I am grateful to thank my economics teacher Patrick Romans for his feedback on an earlier draft of this project, and my high school peers for numerous engaging discussions of my research.

Abstract

In this paper, I incorporate refugee population dynamics into an international education production function. In other words, this paper concerns the effect of refugees on the educational outcomes of native students in host countries. To conduct this analysis, I first adjust the PISA scores from the five most recent years of testing, such that they are comparable across time. Then, I use a fixed-effects regression to find the "average" effects of the changes in the refugee population. Surprisingly, in a regression of adjusted international test scores of native students on the change in the number of school-age refugees and individuals in refugee-like situations over a three-year period per thousand school-age inhabitants, controlling for institutional factors, family inputs, student characteristics, and serial correlation of test scores, the coefficient is 0.198. Similarly, when the lagged number of refugees and individuals in refugee-like situations is included in the regression, the coefficient is 0.069. The positive values of the coefficients suggest that the size of the population of refugees and asylum seekers is positively correlated with better educational outcomes for host country students. However, as is the case in many multivariate regressions, these results are not conclusive because the coefficients are not statistically significant. Additionally, the size of the coefficients suggests a greater impact per refugee in the case of sudden population changes, confirming the refugee shock hypothesis.

Introduction

Globally, the population of concern to the UN Refugee Agency (UNHCR) is comprised of 54.9 million individuals. Over 25% of that population are refugees, many of whom are increasingly seeking refuge in neighboring, less-developed countries (Ruiz and Vargas-Silva, 2013). For this reason, minimizing the negative impacts of refugee populations is a focus of refugee assistance and understanding the educational consequences for host students at an aggregate level can inform relevant policy decisions.

I first present a fixed effects model to estimate the impact of refugee population dynamics on the performance of host country students in the Program for International Student Assessment (PISA), an international test in mathematics, science, and reading. I then analyze the coefficients of the refugee population data in order to test the ‘forced migration shock’ hypothesis. In refugee studies, ‘forced migration shock’ is used to refer to sudden influxes of refugees, asylum seekers, and internally displaced persons. It is thought to have a disproportionately large impact on host communities, compared to steady population fluctuations.

The two main bodies of literature that are relevant to this research are the economics of education and the economics of forced migration. Due to a previous lack of data availability, research in the economics of education has only emerged over the past two decades. In the literature, test scores are often used as valuable proxies for educational quality, as well as indicators of human capital. The determinants of such test scores fall into four main categories; family inputs, school inputs, institutional variables, and student ability. Changes in the refugee population principally affect school and student variables, in terms of refugees as peers. Less directly, it is likely that refugees affect institutional and family variables as well, though less directly.

Like the literature on the economics of education, the empirical research on the economics of forced displacement is mostly limited to country-level or regional analysis. Much of this local research is concerned with the effects of refugee populations on prices and wages in their host communities. However, after an extensive literature search, only one piece of research on the impact of refugees on host children has been published. Javier E. Baez (2011) finds that an influx of refugees into northern Tanzania resulted in a 15 to 20% increase in the rate of infection in host children, and a 5% increase in the infant mortality rate of natives. More importantly to this analysis, he found that the influx of refugees resulted in a 7.1% reduction in host student schooling duration and an 8.5% reduction in literacy.

Given Baez's research in Tanzania, I anticipate that the size of the refugee population will negatively impact the educational outcomes of host students. Additionally, considering the emphasis on refugee shock as the principal mechanism of the effect of refugees on host communities, I anticipate that the change in refugee population will have a more significant impact on host student test scores than the lagged size of the population.

To conduct this analysis, I first adjust the average PISA scores of native students using the National Assessment of Educational Progress (NAEP), such that they are comparable across time. I proceed by using the accompanying PISA survey data, in addition to World Bank development indicators, in order to control for the variation in the test scores. Finally, I use UNHCR demographic data to perform a fixed-effects regression of the adjusted PISA scores.

The results of the regression have notably high standard errors. That said, the estimated coefficient of the lagged number of school-age UNHCR persons of concern per 1,000 school-age inhabitants is slightly positive. Similarly, the estimated coefficient of the change in school-age UNHCR persons of concern per 1,000 school-age inhabitants over the three year period between PISA tests is slightly positive.

These results seem to directly contrast with the evidence available from Tanzania, in terms of correlation. However, the refugee shock hypothesis, echoed by Baez (2011), is supported by the scale of the coefficients. Additionally, the high variation in the "average" positive correlation can be explained by the classical observation of winners and losers in refugee host countries.

The immediate policy implications of these findings relate directly to GDP. Though a causal link has not been established, Countries with average PISA scores that are one standard deviation higher than the average score experience 2% more GDP growth. Furthermore, despite the question of causality in education, it is theoretically apparent that the skills of workers affects GDP. As a consequence, if a greater number of school-age UNHCR persons of concern usually improves host student human capital, then the hosting of the refugees that are most in need of asylum, from a humanitarian perspective, is thoroughly incentivised.

Literature Review

The available literature on the economics of forced displacement is notably limited. Of the literature available, most is conducted at the local or national level. In terms of the observed effects on host communities, the impacts of refugees are varied. Essentially, among communities that host refugees, there are winners and losers. That said, the only published paper on the effects of refugees on host children has found drastically negative impacts.

To start, a distinction must be made between economic migrants and forced migrants. The literature on economic migration is quite well-developed. However, economic migrants and forced migrants vary at a level as fundamental as location selection. While economic migrants often migrate to locations that reward their specific characteristics, forced migrants often face limited decisions that are based primarily on security, the cost of relocation, and government policies (Lehrer, 2012; Falek et al., 2011; Bubb et. al. 2011). Refugees are also likely to be different from economic migrants in terms of unobserved factors related to their displacement, like exposure to violence and loss of illiquid assets (Ruiz and Vargas-Silva, 2013). In addition to those differences, refugees are widely assumed to be more economically isolated (e.g. in camps or centers) than economic migrants (Betts el al., 2014). These distinctions require that the economic and social impacts of refugees on their host communities are not assumed to be the same as those of economic migrants.

One extensively researched topic is the effect of forced migration on host labor markets. Kondylis (2010) found that refugee schock negatively impacted the labor market outcomes, in terms of employment, of both men and women in Bosnia. Alternatively, Maystadt and Werwinp identify both winners and losers in in Tanzania, where household level welfare and wages for the formal sector rose with an influx of refugees.

<u>Data</u>

Most datasets are from Organization for Economic Cooperation and Development (OECD), Institute of Education Sciences, UNHCR, and World Bank databases. Average scores of host country students and the accompanying standard errors and standard deviations are available for all five test years (2000, 2003, 2006, 2009, and 2012) and all three subjects (math, reading, and science). For the purposes of this analysis, I define "host country student" as an individual born in the country in which the PISA test was administered. Thus, the scores are obtained by filtering the average scores of each country by the student's country of birth.

Country data is available for most OECD countries for all years, and some non-OECD countries, mostly for 2009 and 2012. The scores are comparatively adjusted such that the mean score among OECD countries is 500 and the standard deviation is 100 for every test. However, because the countries going into the mean and the real performance of those countries may change from test to test, it is necessary to adjust the test scores.

To adjust the scores such that they are comparable across time, US-specific NAEP scores are used. The NAEP is a test of 4th, 8th, and 12th graders in the same three core subjects and numerous peripheral topics. Despite the focus of the NAEP on abstract knowledge and the focus of the PISA on applied knowledge, the scores are highly correlated, making their combined use reasonable. Because the NAEP testing isn't conducted on a cycle, linear interpolation is used to approximate US scores on non-NAEP years.

The accompanying PISA surveys were also used in controlling for other factors that have been observed to influence student scores. The survey data is used as percent values, based on student responses and school information. The sample is of all students, not host country students, and there is some concern over non-response bias. That said, using data at this aggregate level helps to avoid selection problems implicit in observing variables at the student level. Data that did not meet the OECD's reporting standards was omitted.

To include country-level institutional variables in the model, World Bank Development Indicators such as per-capita GDP[1] and expenditure on education as a percent of the overall GDP are used. Here, there is debate in the literature over whether there is truly a causal link between GDP and PISA scores. However, this distinction is unimportant in controlling for other country-level factors that account for variation in PISA scores.

The last key data source was the UNHCR statistical handbook, which was used in conjunction with World Bank population data to calculate the number of school-age members of

[1] PPP was used for GDP values.

the UNHCR's population of interest[2] per 1,000 school-age inhabitants. For countries in which this demographic data was not available, the number of refugees per 1,000 inhabitants of all ages is used. Due to this inconsistency, these two sets of population data are never used in conjunction when the econometric method incorporates lagging or differencing.

The data, overall, is similar to unbalanced panel data, meaning that it contains multiple observations over multiple individuals, where the observation period is not always consistent. The data is not a perfect panel, but a pooled time series cross-sectional data set, because the observations are actually cross-sectional, and not longitudinal. However, considering that the PISA testing is conducted used a clustered random sample with over 1,000 tests, it is reasonable to assume that the PISA sample is roughly representative of that country's students.

Panel Data Regression

Objectives

The goal of the econometric analysis is to identify the causal effect of refugees on the educational outcomes of the native host student. In order to do this, the econometric method must identify the coefficient in a regression of adjusted test scores of host country students on the number of UNHCR persons of concern. To perform a regression of that nature, the international test scores must first be adjusted to be intertemporally comparable. Then, the regression coefficients must be calculated in a way that accounts for unobserved heterogeneity within individuals. Finally, using the coefficients for both the change in the number of persons of concern over the preceding three-year period and the lagged number of persons of concern, the hypothesis of forced migration shock can be tested using the values of both coefficients.

[2] The population includes refugees, stateless persons, asylum seekers, and internally displaced persons.

Adjusted Scores

In order to regress PISA scores across time, it is necessary to adjust the scores. Each year, the mean PISA score of OECD countries is 500, and the standard deviation is 100. However, the countries going into the mean vary, as it can be expected that the performance of those countries does. In order to adjust the PISA scores, I employ a modification of the method proposed by Eric A. Hanushek and Ludger Woessmann (2012). The method takes advantage of the NAEP test scores, which are intertemporally comparable.

To begin, I express the difference between US scores at any given time and the US scores in 2012 in terms of the standard deviation of native US students in the 2012 PISA test:

$$U_{g,s,t} = (NAEP_{g,s,t} - NAEP_{g,s,2012})\frac{SD_{s,2012}{}^{PISA}}{SD_{g,s,2012}{}^{NAEP}} \qquad \text{(eq. 1)}$$

where U is the score difference for a given grade (4th, 8th, or 12th), subject (math, reading, or science), and time (2000, 2003, 2006, 2009, 2012), NAEP is the grade- time- subject- specific NAEP score, and SD is the standard deviation of the 2012 PISA test scores among native US students. Data for non-PISA years was imputed using a linear interpolation.

I then develop a value X to compare the test score variation across countries. Assuming that variation in real performance between the economically developed and relatively stable OECD countries is relatively constant over time, the average OECD PISA score can be used to standardize the variation by using the average OECD score for native students:

$$X^{i}{}_{s,t} = PISA^{i}{}_{s.t} - PISA^{OECD}{}_{s,t} \qquad \text{(eq. 2)}$$

where X is the score variation for the PISA test of country i in subject s at time and $PISA^{OECD}$ is the average score among all countries for a given subject and time.

Lastly, the previously defined values are used to place the scores on a common metric. Using a given country's reported PISA score, the adjusted score can be calculated:

$$I^{i}_{s,t} = X^{i}_{s,t} - X^{US}_{s,t} + PISA^{US}_{s} + \sum_{a} U^{US}_{a,s,t} \qquad \text{(eq. 3)}$$

where I is the adjusted score for country i, using the previously defined variables. Notably, because PISA tests only sample 15-year-olds, while the NAEP samples three grade groups, the value of U is summed across all ages. Because aggregation is a widely accepted method, this aggregation-based modification is appropriate (Hanushek and Wößmann, 2011).

Fixed Effects Regression

I then conduct OLS regression on the data using a fixed-effects. Using an F-test to compare pooling to fixed-effects methods, significant effects were observed, meaning that some variation in the data is accounted for by unobserved, constant differences between the countries. For example, variables that can't be observed -- like cultural values concerning education -- can be accounted for in each country using a fixed-effects method. Additionally, fixed effects modeling allows any of the dependant variables to be statistically endogenous (Econometric Analysis of Panel Data 3e., p. 19). Because endogeneity is a leading concern in the analysis of international education data, a fixed effects method is the optimal approach to this regression.

Fixed-effects methods account for unobserved homogeneity within individuals through a slight modification of the basic OLS model:

$$y = \alpha_{NT} + X\beta + Z_{\mu}\mu + v \qquad \text{(eq. 4)}$$

where y is a vector of dependant variables, X is a matrix of independent variables, Z_{μ} is a matrix of individual dummy variables, and v is a stochastic disturbance that is ~ iid.

An OLS regression solves for an intercept α, a vector of coefficients β, and the set of fixed parameters μ. For the purposes of the fixed effects model, homoskedasticity is assumed.

The UNHCR data also mandates some alterations to the standard regression method. For nearly every year and nearly every country, the UN has published the number of refugees that each country hosts per 1,000 inhabitants. However, the UNHCR has not published demographic data for the refugee populations of many countries. As a consequence, the vectors for lagged refugees and the change in refugees includes some values that are based on the number of refugees per 1,000 inhabitants of all ages, which can be viewed as a proxy variable, and other values that are based on the number of school-age persons of concern per 1,000 school-age inhabitants.

To ensure that this inclusion does not drive the regression coefficients, a separate vector of dummy variables was established. The values of the vector were 1 for data that described all ages, and 0 for the data of interest. In the vector X, the vector of dummy variables was included to allow these observations to have their own slope. Two interaction terms, between the dummy variables and each of the refugee variables, were also added to the matrix of independent variables to allow the all-ages observations to have their own slopes as well. Although this inclusion is expected to both bias the estimators and decrease their precision, it is necessary in order to maintain a reasonable number of observations.

Results

The results of the fixed-effects analysis are presented below in Table 1.

Table 1: Fixed-Effects OLS Estimates of the Effect of Refugees on Educational Outcomes
Dependant variable: Adjusted international PISA average scores (2000-2012)

	Coefficient	Standard Error

Change in school-age population of concern (per 1,000 school-age inhabitants)	0.198	(0.156)
Lagged school-age population of concern (per 1,000 school-age inhabitants)	0.069	(0.068)
Change in refugees of all ages interaction term	-0.244	(0.401)
Lagged refugees of all ages interaction term	-0.233	(0.980)
Refugees of all ages dummy	31.212	(25.837)
No. of Observations	158	
No. of Countries	68	
R-squared	0.866	
Adjusted R-squared	0.302	

TABLE 1. Notes to Table 1: Data is from the PISA, NAEP, World Bank Development Indicators, and the UNHCR statistical yearbooks, and covers four testing years (2003, 2006, 2009, and 2012). The "all ages" dummy is a dummy variable that takes on a binary value of one for data that was imputed from UNHCR data for refugees per 1,000 inhabitants, and 0 if the data is for the number of school-age persons of concern per 1,000 school-age inhabitants. The dummy variable serves to give the imputed data its own intercept, and the interaction terms allow the imputed data for each variable to have its own slope. The controls used in the regression are based on those identified in the literature on education production functions. All coefficients presented above were not statistically significant.

Table one shows that a greater number of lagged, school-age UNHCR person of concern per 1,000 native host country students slightly increases the expected average PISA score. The same effect is even more dramatically evident, nearing 1 point. However, even that value places the impact of each school-age person of concern per 1,000 native host country students at half of one percent of the standard deviation in the adjusted test scores of the sample. Additionally, considering that each additional refugee has more than two times the impact of each lagged refugee, the refugee shock hypothesis appears to be confirmed by the regression results.

Despite this evidence, and as a matter of econometric theory, the direction of the causal relationship is not clear. In fact, of those forced migrants who have a choice of where to migrate, they are likely to select a location that is suitable to their individual characteristics. This has only been observed in the contest of labor and skills, but it is possible that refugees seek out educational institutions of higher quality. The migration of refugees and asylum seekers may not be perfectly exogenous.

In addition, the interpretation of the refugee data does not lead to a very clear conclusion. The dummy variable indicated an intercept for the refugee-only data that is an average of 31.2 points higher than the data for school-age persons of interest. However, this could be a matter of the dummy variable being a proxy for unobserved factors, such as development. In the population data set, most countries without published demographic data are well-developed OECD members. The UNHCR reports demographic data much more often for non-OECD PISA countries. In this way, the coefficient could be biased up.

Similarly, the interaction term which allows the data for refugees of all ages to have its own slope is likely to be biased in the same way. The standard errors of the coefficients of these interaction terms are also much higher than those of the combined data. With that in mind, when the average slope of the interaction term and the school-age term is considered, the estimated effect of the change in refugees is approximately zero and the estimated effect of lagged refugees is negative, but with an extremely large margin of error. On the other hand, the coefficient for lagged refugees of all ages may truly be representative of a slightly negative effect. Essentially, the interaction terms for the refugee data do not suggest a clear conclusion.

For all values, the data reflects a large margin of error. First of all, this is quite common among panel data sets. However, there are more possible interpretations. For instance, the R-squared value is reasonable, while the adjusted R-squared value is extremely low. This can be traced back to co-linearity in the regressors. In other words, it's likely that the inclusion of variables such as the all-ages interaction term reduced the statistical significance of related regressors.

Beyond an interpretation just within the data, the large margins of error make sense in the context of the literature. In regard to host communities, winners and losers have been observed, as discussed in the literature review. Perhaps this variation in outcomes is the root cause of the lack of statistical significance in the data.

Whatever the causes of the variation in the observed effects are, it is still true that they represent an average effect. Regardless of their statistical significance of the coefficients, and even considering the the large standard errors, it would still appear that, on average, the number of school-age UNHCR persons of concern per 1,000 school-age inhabitants appears to be positively correlated with adjusted PISA scores.

Conclusion

In essence, I found that while data at the less accurate and relevant level of refugees per 1,000 inhabitants doesn't lead to a clear conclusion, the average effect of school-age persons of concern on the educational outcomes of host country students is slightly positive. This has immediate policy implications at a level as economically fundamental as GDP. A strong correlation has been observed between the adjusted average PISA test scores of a country and that country's GDP. In fact, a PISA score that is one standard deviation higher than average is correlated with a 2% higher relative GDP growth rate (Hanushek and Wößmann, 2012). Given this strong association, and given that an additional UNHCR person of concern per 1,000 school-age inhabitants accounts for one half of a percent of one standard deviation from the mean expected test score, host countries appear to have a strong incentive to take humanitarian action.

Bibliography

Hanushek, E.A. and Wößmann, L. ”Education and Economic Growth.” International Encyclopedia of Education, 2012, vol. 2, pp. 245–252.

------------. “Do better schools lead to more growth? Cognitive skills, economic outcomes, and causation.” The Journal of Economic Growth, July 2012, vol. 17, pp. 267–321.

------------. “The Economics of International Differences in Educational Achievement.” Handbook on the Economics of Education, 2011, vol. 3.

Bubb, Ryan; Kremer, Michael and Levine, David I., 2011. "The Economics of International Refugee Law." The Journal of Legal Studies, 2011, vol. 40 (2), pp. 367–404.

Ruiz, Isabel and Vargas-Silva, Carlos. “The Economics of Forced Migration.” The Journal of Development Studies, April 2013, vol. 49 (6), pp. 772–784

Maystadt, Jean-François and Verwimp, Philip. "Winners and Losers Among a Refugee-Hosting Population," Households in Conflict Network Working Paper No. 60, 2009.

Betts, Alexander et al. “Refugee Economics” Humanitarian Innovation Project, University of Oxford, June 2014.