Dynamic Amenities and Path Dependence in Location Choice: Evidence from Danish Refugee Placement

Farid Farrokhi
Purdue

David Jinkins*
Copenhagen Business School

January 5, 2021

Abstract

We examine whether spending time in a location causes a person to stay there longer. We exploit a 1999 change in Danish refugee settlement policy, which required refugees to stay in their conditionally randomly assigned settlement location for at least three years. In a difference-in-difference design, we compare changes in the length of time refugees chose to remain in their settlement location relative to non-refugee immigrants. In our preferred specification, we find the required stay caused refugees to be 7.1 percentage points more likely to remain in their settlement city thirteen years after their initial placement. We also calibrate a simple framework of location choice, and find that an additional year of staying in a location creates an opportunity cost that equals 15% of the one-time static cost of moving.

Keywords: Location choice, Amenities, Moving costs, Immigrants, Conditionally-random allocation
JEL Classification: I38, J61, R23

1 Introduction

A long-standing tradition in economic geography has emphasized the importance of path dependence in the location of factors of production (Krugman, 1991; Bleakley and Lin, 2012). While it is well-documented that cities are persistent as the locations of economic activity, little is known about the specific mechanisms that cause the persistence. We take a step toward identifying such mechanisms by studying the effect of staying in a location on an individual's likelihood of staying there longer. We call this phenomenon "dynamic

*We are grateful for funding from the Independent Research Fund Denmark grant number 8019-00031B which made this project possible. We would also like to thank for helpful comments from Niels Henning Bjørn, David Hummels, Birthe Larson, Rie Martens, Herdis Steingrimsdottir, Chong Xiang, and participants at a NOITS meeting, a Kraks Fund seminar, and seminars at Copenhagen Business School and Purdue. All remaining errors are our own.

amenities" in contrast to static amenities that are independent of a person's history of residential location. Dynamic amenities arise from growing roots in a location. The benefit that a person derives from living in a city rises over time as the person learns about local amenities, gets attached to them, and as her social network grows in the city.

In addition to examining the existence of dynamic amenities as a mechanism behind path dependence in economic geography, we measure their contribution to moving costs. Across the literature, moving costs are estimated as structural residuals inferred from migration flows. Kennan and Walker (2011) estimate that moving costs for an average mover across the American states amount to $312,000 (in 2010 dollars). Tombe and Zhu (2019) find that migrating between China's provinces reduces lifetime utility on average by nearly a factor of three. These large estimates point to a range of migration frictions above and beyond direct, monetary costs. Here, we explore the extent to which dynamic amenities contribute to these moving costs.

In this paper, we exploit a natural experiment to test the causal proposition that staying in a location makes people less likely to leave. It is straightforward to show the simple correlation, that the longer someone is in a location the less likely she is to leave. In the Danish data we have examined, this relationship is strong and robust to controlling for marital status, age, number of children, and home ownership status. However, this simple correlation does not lend much support to our hypothesis that amenities are dynamic. The identification challenge is that people who remain in a city longer might simply prefer that location for time-invariant reasons, not because they have grown more attached over time. The same static reasons that make a person select a location in which to reside also lead her to stay there longer, creating a negative correlation between the probability of moving and the amount of time spent in a location. Separating dynamic amenities from this underlying correlation is inherently challenging.

To address this identification challenge, we exploit variations from a policy change in the placement of refugees in Denmark. Prior to 1999, refugees arriving in Denmark were allocated subsidized housing in different Danish municipalities. They were encouraged by social workers to stay in their placement municipality for at least eighteen months. After a law change in 1999, refugees were legally required to remain in their placement location for three years.[1] A literature on Danish refugee settlement has argued that refugee placements were random conditional on several factors such as family status and country of origin Damm (2005); Azlor et al. (2018). Since most of these factors are observable, we follow the literature in treating placements as conditionally random.

To test our hypothesis, we apply a difference-in-difference regression in which the outcome variable is

[1]There were some exceptions to this policy, for example if a refugee found regular full-time work in another municipality. Details of the policy both before and after 1999 are discussed in Section 2.

whether an immigrant is still in her first location thirteen years after being granted a visa in Denmark. The longest required stay is three years, so this is ten years after the end of the required stay. The treatment is the policy change which required refugees to stay longer in their allocated location. If dynamic amenities are important, we expect that, relative to non-refugee immigrants, treated refugees should be more likely to still be in their allocated location after thirteen years. In our preferred specification, we find that thirteen years after the initial allocation, treated refugees were 7.1 percentage points more likely to still be in their allocated location relative to non-treated refugees. The unconditional probability of a refugee in our sample staying in his initial location after thirteen years is 55.3%, meaning that the dynamic amenity effect accounts for 12.8% of the baseline likelihood of staying. We take our difference-in-difference result as evidence for a causal effect of dynamic amenities.

Several features of the refugee settlement policy make the implementation of our exercise challenging. The nature of this challenge is in defining the sample and control variables to make the exercise as close as possible to an ideal randomized experiment. For example, immigrant students tend to leave their first observed location after a fixed number of years with a high probability, therefore they are difficult to compare with refugees. Therefore it is important to either control for student visas or exclude them from our regressions. In addition, refugees with special needs were often settled in Copenhagen. For this and other reasons we discuss below, refugees' settlement to Copenhagen might be non-random. Moreover, given that Denmark does not have a particularly large population, as we add controls or restrict our sample size we tend to lose power in our estimation. For all these reasons, we conduct a meta-exercise in which we run our difference-in-difference regression on a wide range of specifications. The sign and statistical significance of our estimates are for the most part robust to these different specifications. The point estimates range from two to nine percentage points.

To provide an additional robustness check, we run our analysis using a regression discontinuity design. Identification in our difference-in-difference design relies on a common-trend assumption between refugee and non-refugee immigrants. Rather than a common-trend assumption, our regression discontinuity design relies on the assumption that refugees who were granted visas just before the implementation of the policy are randomly drawn from the same population as those who were granted visas just after the policy implementation. Our baseline regression discontinuity point estimates imply that refugees arriving just after the policy implementation are 16.3 percentage points more likely to be in their allocated city after thirteen years in Denmark. The point estimates in our regression discontinuity design are larger compared to those of our difference-in-difference but not as accurately estimated. We take the sign and magnitude of the estimates in our regression discontinuity exercise as a confirmation of our difference-in-difference results.

We explore potential mechanisms for the policy effect by replacing the outcome variable in our difference-in-difference specification with other characteristics of refugees. We find that after the policy refugee household heads were more likely to have children in their first years in Denmark and were more likely to be working. The policy had no effect on marriage status or on refugees receiving government support. These results suggest that being forced to stay in a location made refugees more likely to exert effort to look for a local job, and have children earlier. Social connections to coworkers, other parents, and other children may have tied refugees more strongly to their settlement location.

Finally, we calibrate the implied costs that dynamic amenities generate for leaving a location relative to static costs of moving. Using a simple model of location choice, we derive two equations that can be brought to our data for a calibration. One of the equations is a structural counterpart of our difference-in-difference regression meant to be informative about the effects of dynamic amenities. The other equation contains information about both dynamic amenities and static costs by exploiting bilateral flows of immigrants moving across regions within Denmark. We find that staying in a location for an additional year creates an opportunity cost that is equal to 15.6% of one-time static moving costs. To put this number in perspective, suppose the accumulation of the implied dynamic opportunity cost over time was linear and the discount rate was 5%. Then dynamic amenities would create a location-specific asset that after seven years is equal in value to the static cost of moving.

This paper contributes to the urban economics and economic geography literature in several ways. On the empirical side, this literature has shown in a number of different contexts the importance of history in the current location of economic activity (Bleakley and Lin, 2012; Kline and Moretti, 2013; Dalgaard et al., 2018). To give an example inspired by Bleakley and Lin (2012), many cities in the United States are built around navigable rivers or portage sites. These were important trade and transport hubs at one time, but now technology has made them obsolete. And yet, the cities persist.[2] In particular, we provide causal evidence in support of studies that emphasize the importance of tenure in a location through local ties or home attachment as reflected in people's tendency to reside near their birthplace (Zabek, 2019; Coate and Mangum, 2019). On the theoretical side, Allen and Donaldson (2018) study the path of spatial equilibrium and its potential dependence on history. Our project contributes to the empirical part of this literature by showing that part of the persistence in spatial distribution of economic activity is driven by the dynamic amenity effect.

In addition, our paper contributes to the literature examining the costs of relocation (Kaplan and

[2]The literature suggests that the importance of history in the current location of cities may depend on the empirical context. Davis and Weinstein (2002) examine the persistence of the distribution of population in Japan, finding that the long-run city size in Japan was robust to shocks as big as the destruction of Japanese cities in WWII. Miguel and Roland (2011) also find support for persistence when considering the US bombing of Vietnam as a shock. Bosker et al. (2007), however, find support for multiple equilibria in German city growth given the bombing of Germany in WWII. Likewise, Michaels and Rauch (2018) find support for the importance of history by examining the effect of the collapse of the Western Roman Empire on city locations in Britain and France.

Schulhofer-Wohl, 2017). Using migration flows to infer the magnitude of moving costs, this literature has found huge costs. For example, as we mentioned earlier, Kennan and Walker (2011) estimate that the inferred cost per move between US states in 2010 was on average \$312,000 (in 2010 dollars) which amounts to 42% of median lifetime income calculated at age 30. In addition, a number of studies have shown that home ownership reduces mobility (Munch et al., 2006; Mian and Sufi, 2014; Yagan, 2014), and that people tend to remain in their locations even as the location is hit by a negative productivity shock (Autor et al., 2015). We complement these studies by finding that an important part of these moving costs stem from dynamic amenities. A parallel to moving costs in economic geography is trade costs in international trade. A number of studies in the trade literature have aimed at uncovering frictions behind the large estimates of trade costs, when inferred as structural residuals of trade flows, highlighting the role of inventories (Alessandria et al., 2010), timeliness (Hummels and Schaur, 2013), information (Allen, 2014), and creating and maintaining international relationships (Eaton et al., 2020).

Our empirical study is inspired by an extensive literature using natural experiments in urban and labor economics (Baum-Snow and Ferreira, 2015). Several studies have taken advantage of forced migration to study the benefits of moving on future human capital and earning (Nakamura et al., 2016; Voigtlaender et al., 2020). More closely, we build on a literature that investigates the conditionally random allocation of refugees to cities in Nordic countries (Edin et al., 2003; Damm, 2005; Damm and Dustmann, 2014; Foged and Peri, 2016; Dustmann et al., 2018; Azlor et al., 2020). This literature often uses data from refugees or immigrants as instruments to study economic or social outcomes for the entire sample of society (e.g. How large are agglomeration economies? How does the number of immigrants in a location affect voter sentiment?) Focusing on refugees themselves, Eckert et al. (2020) study the effect of big city experience on wages within the refugee community.

2 Policy Background

In this section, we describe the Danish policy for settling refugees in the period up to 1998, and changes made in 1999.

2.1 Danish Refugee Policy 1986-1998

Starting in the 1980's and up until 1999, the practicalities of assigning refugees to permanent housing in particular locations was handled by the non-profit Danish Refugee Council (DRC). As in other Scandinavian countries, low-cost housing is offered to asylum seekers who are granted refugee status in Denmark.[3] Our description of the policy in this period summarizes Damm (2005), and the interested reader should read that paper for more detail. In the early 1980's, refugees were assigned housing in their preferred location. Due to difficulty finding housing, and also because of political pressure to disperse refugees out of major urban centers, in 1986 the DRC started a policy of conditionally random allocation of refugees to counties.[4]

The goal of the policy was to distribute refugees proportionally to the population of the 15 counties in Denmark. Only 182 of the 275 Danish municipalities received refugees in this period. The DRC only assigned refugees to municipalities which had necessary facilities for integration, which in practice meant that refugees were not usually placed in the most rural municipalities. The DRC had temporary offices in municipalities. The location of the offices would rotate over time across municipalities. Refugees were placed in the municipalities which had temporary offices when their refugee visas were granted. This mechanism led to clustering of nationalities in particular municipalities. Refugees did not need to accept the DRC housing offer if they could find their own permanent housing. In practice, around 90% of refugees took the housing offer from the DRC in the period 1986-1998. Because of a large inflow of refugees from the wars in the former Yugoslavia in the early 1990's, housing availability became the primary driver of placement and refugees were placed in rural municipalities as well.

An asylum seeker granted refugee status in this period was asked whether he would like help from the DRC in finding housing. If he agreed, he filled out a form including his preference of county and the reasons for that preference. Around ten days later, he would be provided temporary housing in the county, and the local office of the DRC would begin the search for permanent housing in the county. A refugee could ask to be reassigned if he was unhappy with the county he was allocated, and the DRC would reallocate him to another county. In practice this did not happen often. The average wait for permanent housing in the assigned county was 6-7 months. If the number of refugees requesting a particular county exceeded the quota for that county, priority was given to those with close family in the county or with special medical or educational needs available there. As refugee inflows to Denmark increased markedly in the 1990's, quotas in desirable counties filled up quickly in that period.

The literature that has studied Danish refugees' settlement policy indicates that the refugees' allocation across counties was random conditional on several controls. In particular, Damm (2005) cites seven factors as influencing the initial assignment of refugees to locations: marital/family status, nationality, year of refugee status, location of close family or friends, special health needs, special educational needs, and finally

[3]Sometimes in this literature, the housing will be described as permanent. In Denmark, this often means a rental contract without a time-limit. Due to renter protection laws, it is difficult to evict renters as long as they pay their rent. A typical rental contract in the spot market has a fixed term, allowing landlords to evict the tenant at the end of the contract.

[4]Prior to a reform in 2007, there were 14 Danish counties, and two "municipal" counties Copenhagen and Frederiksberg, both part of the Copenhagen metropolitan area. At the time these counties were divided up further into 275 municipalities.

a strong preference for a particular location. Of these, the first three are recorded in our data. Damm (2005) suggests proxies for the next three factors, and argues that the final factor is not too important since so few refugees asked to be reassigned. In an empirical analysis, she finds that all of the listed factors are statistically significant in influencing initial location assignment in the Danish register data, but combined their explanatory power is modest with $R^2 \approx 0.13$ which she takes as suggestive evidence that refugee placement was indeed conditionally random.

After being placed in a county, the DRC provided instruction in Danish language, culture, and job training, and the refugees received a means-tested subsidy. Refugees were encouraged by social workers to stay in their assigned county for at least 18 months to complete the coursework, but they were legally free to leave as long as they could find housing elsewhere. Refugees could still receive the subsidy if they chose to leave.

2.2 Major policy change 1999

Partly in response to the large inflow of refugees from the former Yugoslavia in the mid 1990's, the Danish Integration Act was passed by the Danish legislature in 1998, coming into effect in 1999 (Udlaendinge- og Integrationsministeriet, 1998). Our understanding of the effects of this law change draws on the summary in Azlor et al. (2020), as well as our own reading of the law and interviews with Danish government bureaucrats who were involved with implementing the law.

Below we discuss three changes to the law that bear relevance for our study. The first effect of the 1999 Danish Integration Act was that refugees became more tied to their assigned location. The second effect was that the length of integration coursework was lengthened from 18 months to three years. Finally, the method for dispersing refugees across Danish counties changed.

After 1999, it was more difficult for refugees to move away from their placement location. As an important background aside, in Denmark one must register with the government when moving to a new home. Under the new law, until the end of the training period refugees were not allowed to register in homes outside of their assigned municipality without either approval of the destination municipality, or the demonstration that they had a work position with a contract in the destination municipality.[5] Furthermore, Azlor et al. (2018) report that the subsidy from the government was tied to residing in the assigned location for the training period.[6]

[5] In our interviews, one bureaucrat said that it was always possible to register in a destination municipality, but that refugees were under the impression that it was illegal to move.

[6] We have obtained a sample of the letter sent to refugees when they received their visa from the Danish Immigration Service. This letter includes the following paragraph, loosely translated from Danish:

> You cannot choose where you will live yourself, because it is a large–and expensive–task for the municipality to

With the passage of the new law, much of the responsibility for refugee settlement formerly delegated to the DRC was brought into the Danish government. The choice of municipalities to settle refugees was taken over by the Danish Immigration Service (DIS). Each year, DIS was tasked with estimating the number of refugees which would arrive the next year. These expected refugees were then allocated annually to municipalities based on the number of foreigners currently residing in the municipality and the total population of the municipality. If more refugees arrived than expected, DIS would adjust the quotas as needed. Locations with fewer foreigners relative to population were assigned higher refugee quota. Refugees would only rarely in special circumstances be allocated to a location with its quota already filled.[7]

The placement of a particular refugee was decided by DIS based in part on an interview with the refugee himself. In making the allocation decision, DIS considered factors such as whether the refugee had family already in Denmark, the nationality of the refugee, and any special educational or medical needs. We will explain the ways we control for these factors in Sections 3 and 4.

3 Data and Descriptive Statistics

3.1 Data Sources

Our empirical analysis relies on Danish register data. Among these, we primarily use two registers. The first is the Residence Register. This register is created using information from the Danish Immigration Service on all first-time residence permits granted from 1997 until the present.[8] Important variables include the type of permit, the date the permit was granted, and a personal identification number of the recipient linkable across data sets. The second register we use is the Population Register. This register contains annual demographic information on all Danish residents beginning in 1986. Important variables for us are the municipality of residence, family/marital status, age, gender, and country of origin. We link the Population Register with the Residence Register using the personal identification number. In addition to these primary registers, we

organize an integration program. As a starting point, for you to follow an integration program and ultimately receive a permanent residence, it is a requirement that you live for 3 years in the municipality which the Ministry for Foreigners has decided you should live in. There can also be consequences for the payment of your subsidy if you insist on moving. You can find another municipality to live in if you have found permanent full-time work there.

[7]These parts of the Danish Integration Act of 1999 were largely unchanged until after 2012, when major changes to the law were passed by the Danish legislature particularly in 2016. These changes were partly a response to the then ongoing refugee crisis in Europe induced by the Arab Spring and Syrian Civil War (Dustmann et al., 2017). In addition to generally becoming more strict on immigration, the rules about how long refugees need to remain in the allocated municipality and what types of training programs they must attend have been overhauled several times since 2012.

[8]We obtain information on the first time someone is granted a residence permit within a permit type. For example, we observe the date at which someone is granted a family reunification visa, but not the date of renewal if that visa is renewed. If the same person later gets a work visa, we will observe the date that the work visa is granted.

supplement our analysis with information from several other Danish registers, including information about family relationships from the Family Register and about income from the Income Register.[9] Because our primary empirical exercise focuses on an immigration policy change in 1999, we will mostly be interested in behavior of people who arrived in Denmark around the time of the policy change. Concretely, our data is composed of all household heads who were granted a residence permit in Denmark from 1986 to 2016. Our sample contains only immigrants who were working age adults when the residence permit was granted, which we define as older than 18 and younger than 60. Household head refers to a single adult not married to an existing Danish resident, not the child of an existing Danish resident, and not the parent of an existing Danish resident.[10] In other words, we only analyze immigrants who had no close family connections in Denmark prior to their arrival.[11]

We break this sample of immigrants into two groups. The first is household heads granted refugee permits ("refugees"), and the second is household heads granted non-refugee residence permits ("non-refugee immigrants").[12] An additional challenge with our data set is that we only observe first visa information for immigrants arriving in or after 1997. In order to identify refugees in the period from 1986 to 1996, we follow the strategy used in much of the literature, and use immigrant origins to differentiate refugees from non-refugees (Damm, 2005; Damm and Dustmann, 2014; Dustmann et al., 2018). Each of these studies uses a slightly different set of origins. We use the list of refugee-sending countries from Damm (2009), which is based on the top 11 sources in the period 1986-1993, and add the former Yugoslavian republics since these were important sources in the mid-1990's. In practice, we classify any immigrant arriving from the following countries as a refugee before 1997: Lebanon, Iran, Iraq, Somalia, Sri Lanka, Vietnam, Poland (before 1990), Afghanistan, Ethiopia, Romania, Chile, Yugoslavia Federal Republic, Yugoslavia, Serbia & Montenegro, and Bosnia.

[9]The Danish register acronyms are: Residence Register (OPHG), Population Register (BEF), Income Register (IND), and Family Register (FAM). Except for the exact date a visa was granted in the Residence Register, most other variables in our data set (e.g. location, marital status, number of children) are measured on January 1st. This means that we observe demographic characteristics of immigrants for the first time on January 1st in the calendar year following the year in which the visa was granted.

[10]If a married heterosexual couple is granted residence in the same year, we follow the Denmark Statistics protocol that the female is designated household head. If a married homosexual couple is granted residence, the protocol is that the older spouse is the household head.

[11]We do not observe more distant family relationships, such as whether an immigrant is the cousin or uncle of an existing Danish resident.

[12]Due to technical issues at DIS, data on immigration cases is not always linked to the ultimate personal identification numbers of residence permit recipients. In the Residence Register data, Denmark Statistics has made an effort to match missing personal identification numbers using nearest neighbor matching on attributes like country of origin, gender, age, etc. These methods work well when immigration flows are small and diverse, but not as well when immigrants are mostly homogeneous and from large source countries. In our analysis we drop these imputed observations, which represent 15.6% of the observations in the Residence Register. We also drop immigrants who we observe being granted a visa in a particular year, but who do not appear in our other registers until after the end of the following year.

Since our paper is about the decision to stay in a location, the definition of location will be important. Ideally we want to consider integrated local labor markets or commuting zones. To this end, the municipality level in Denmark is too fine. The ideal unit would be what used to be known as the county. Unfortunately this exercise is complicated by the fact that in 2007 there was a major Danish administrative reform that eliminated the county level of government (and incidentally also eliminated 173 municipalities). Thus for the purposes of this paper, we combine the county prior to 2007 with the statistical unit "province" reported after 2007 as our unit of analysis. The province unit is similar, but not exactly the same as the former Danish administrative county. In particular, to compare before and after 2007, we must combine the two counties and two provinces for Central Jutland and South Jutland, because an important regional Danish city (Horsens) changed from the county South Jutland to the province Central Jutland. We also consolidate the counties/province in the Copenhagen metropolitan area. We end up with the following seven locations, which we call "cities": Copenhagen, Bornholm, West and South Zealand, Fyn, Central and South Jutland (Aarhus), West Jutland, and North Jutland. We put these cities on a map in Figure 1. In most of our analysis below we omit Bornholm, as this small island had very little immigration of any kind in the period we study.

3.2 Descriptive Statistics

We tabulate our descriptive statistics for all adult household heads separately for our two groups: Refugees and Non-refugee immigrants. Except where obvious, we limit our descriptive statistics to immigrants arriving in the years 1986-2003, because this corresponds to the sample in our empirical exercises below. We report statistics separately for immigrants arriving before and after the 1999 refugee policy change. All statistics we report correspond to the immigrants' first observed year in our data.

Figure 2 contains the number of immigrants in our sample population by year. The number of non-refugee immigrants before 1997 should be viewed cautiously, because the method we use classifies people less likely to be refugees (say from less common refugee sending countries) as a non-refugee immigrant, even if these people are in fact refugees. In particular, the drop in the level of non-refugee immigration from 1996 to 1997 is likely to simply be that we capture non-refugee immigrants more accurately in 1997. The rapid increase in immigration after 2003 was due to the EU expansion, which allowed free immigration from Eastern Europe. Growth in overall immigration stalled following the financial crisis of 2008, and then resumed with the recovery beginning in 2012. One somewhat startling observation to come out of this table is that while immigration has been rapidly growing over the period, Denmark granted markedly less refugee visas in the mid 2000's relative to the 1980's and 1990's. This trend was reversed during the Syrian refugee crisis in the 2010's.

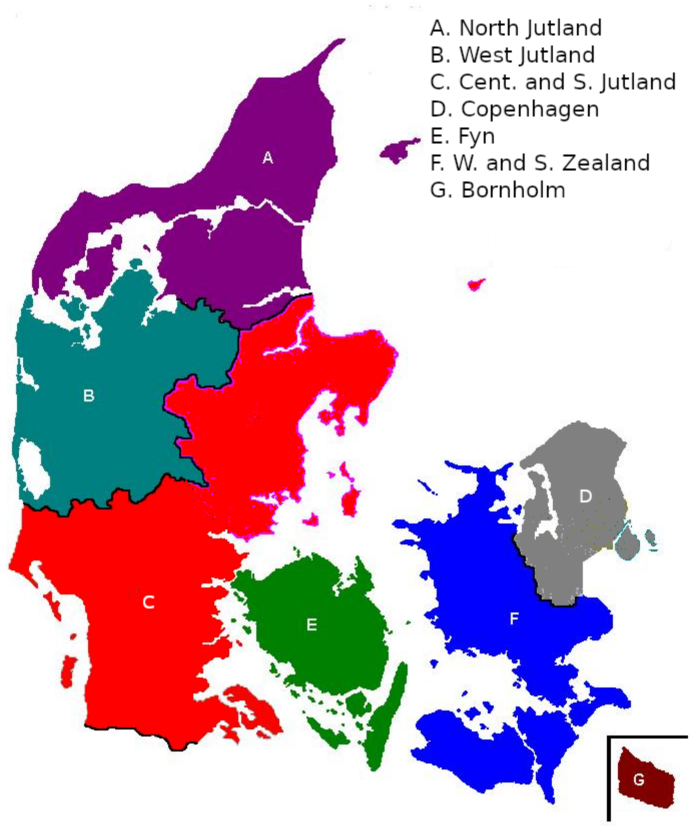

Figure 1: Danish geography

Notes: This figure shows the cities we will use as units of analysis. This map is adapted from a map by Wikipedia user Boeing720, and is released under the Creative Commons Attribution-Share Alike 3.0 Unported license.

Figure 3 shows the major countries of origin for refugees receiving visas from 1986-2003. We split the sample up into refugees coming before and after 1999. The composition of refugee countries of origin changed over this period, depending on the location of wars and natural disasters in the world. Syrians and Afghanis were more likely to arrive after 1999, while Bosnians and Sri Lankans were more likely to arrive before 1999. Iran and Iraq were major refugee sending countries during both periods.

Table 1 contains descriptives for immigrant household heads. Refugee household heads are more likely to be male, while non-refugee household heads are slightly more likely to be female.[13] Refugees are more likely to be married on arrival, and have both more children and larger family sizes. Non-refugees are younger on average than refugees particularly after 1999, because many of them arrive as students.

Highest education levels on immigration are presented in Table 2. This information is only for immigrants who have a recorded education level, which make up, for example, only 12% of non-refugee immigrants whose first visa was granted between 1999 and 2003. With that caveat in mind, recorded education levels of refugees are stochastically dominated by education levels of non-refugees. Both groups are more educated in the period after 1999 than in the period before 1999.

Table 3 presents the residence location of immigrant household heads in the first year after receiving their visas. The most striking pattern visible in the table is that more than half of non-refugee immigrants live in Copenhagen when their visas are granted, but only a quarter of refugees live there. Refugees are also significantly more likely to live in the relatively rural and remote parts of Denmark outside of Central and South Jutland (Aarhus) and Fyn when their visas are granted.

Table 4 contains counts of visa types for non-refugee household heads after 1997 when our visa information starts. The two most popular forms of visa are for students, and for general EU citizens. If we focused on non-household heads, family reunification would probably be more important. As it stands, this type of visa is rather unimportant, which is reassuring since we are worried about family relations of refugees falling into our control group. Student visas are a rather large category, and students often stay only for a few years and then leave their location within Denmark, or Denmark altogether. Since this profile is quite different from the standard refugee profile, we sometimes exclude student visas in our analysis below.

Finally, Table 5 reports unconditional staying probabilities for the sample of immigrant household heads. Refugee household heads are much more likely to remain in their placement cities than non-refugee immigrants in all years after their initial placement.[14] Anticipating our estimation method below, we

[13] Intuition that single refugees are likely to be male is not wrong. Part of the reason we register so many female household heads among refugees is that if a married couple arrives together, the female is designated the household head. The reason that non-refugees have so many male household heads is that as we will see, many of the non-refugee immigrants arrive as students, who are usually single.

[14] A number of factors play a role in the staying rate of non-refugee immigrants versus that of refugees. In contrast to refugees, non-refugee immigrants have the option of returning to their home country. In addition, compared to refugees, non-refugee

also include the difference-in-difference for staying rates in the final column. While the numbers here are consistent with our dynamic amenities hypothesis, we should be careful not to interpret this table causally. As argued above, the assignment of refugees to locations can only be regarded as random when we control for, at a minimum, the origin country and placement city of refugees.

Figure 2: Refugee and non-refugee immigrants by year

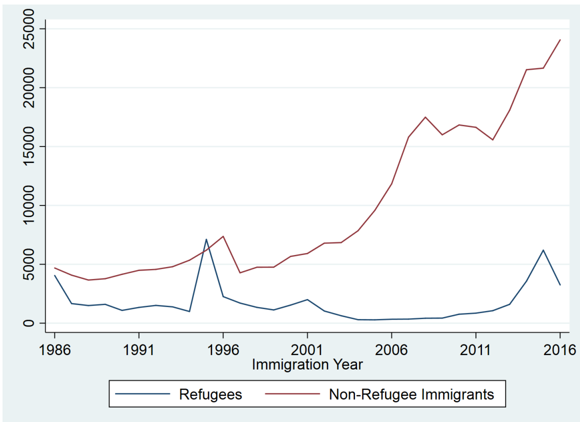

Notes: This figure shows the number of refugees and non-refugee immigrants in our sample granted visas between 1986 and 2016

Figure 3: Refugee countries of origin

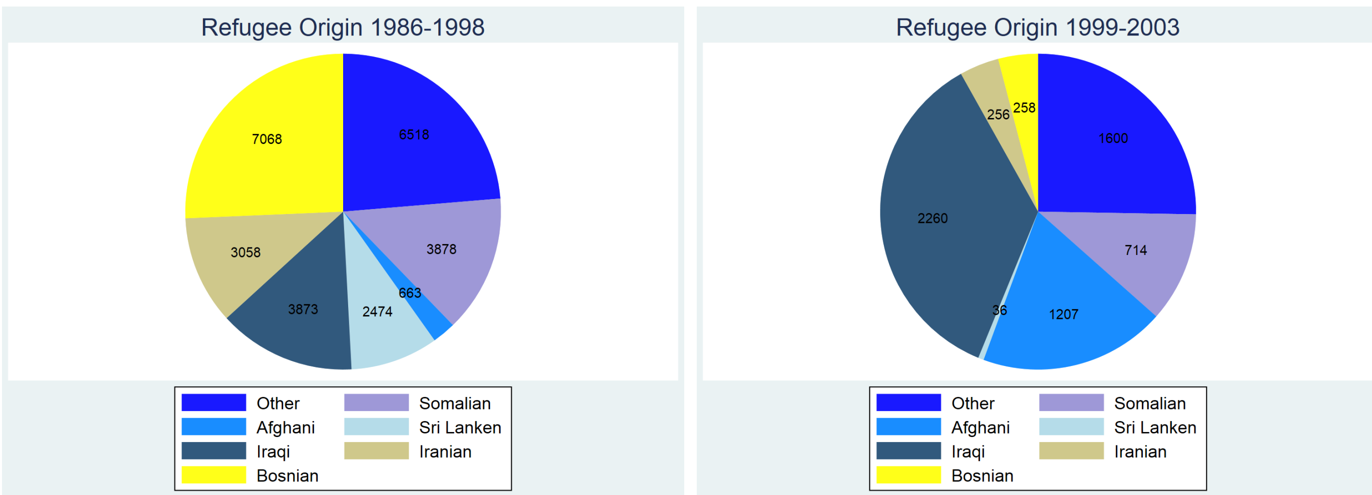

(a) Visa granted pre-1999

(b) Visa granted post-1999 inclusive

Notes: This figure shows major countries of origin for refugee household heads who first received their visas from 1986-2016.

immigrants may benefit from more moving opportunities across cities within Denmark. On the other hand, since a refugee's preferred residence location is likely different from the location he is assigned by the Danish Immigration Service, a refugee may have more incentive to ultimately leave than a non-refugee immigrant who chooses his first residence location himself. Overall the data indicate that the net effect of these forces is that refugees stay longer than non-refugee immigrants.

Table 1: Descriptive Statistics

	Before 1999		After 1999	
	Refugees	Non-Refugees	Refugees	Non-Refugees
Female	0.49	0.54	0.39	0.53
Children	0.62	0.25	0.58	0.13
Family size	1.87	1.46	1.80	1.27
Married	0.47	0.24	0.52	0.16
Age	30.39	27.84	32.63	26.05
Count	27532	62091	6331	29965

Notes: This table shows descriptive statistics for immigrant household heads granted their first Danish visa from 1986-2003. Statistics refer to the first year in Denmark.

Table 2: Education Levels

	Before 1999		After 1999	
	Refugees	Non-Refugees	Refugees	Non-Refugees
Primary school	44%	25%	15%	12%
General secondary	18%	16%	28%	16%
Vocational internships	18%	34%	24%	31%
Short higher education	4%	8%	5%	10%
Medium long higher education	10%	11%	16%	17%
Long higher education	7%	8%	11%	14%

Notes: This table shows education level upon arrival in Denmark for immigrant household heads first granted visas 1986-2003, if recorded. Short and medium long higher education are professional post-secondary training programs. Long higher education includes university bachelor and masters programs.

Table 3: Initial Location

	Before 1999		After 1999	
	Refugees	Non-Refugees	Refugees	Non-Refugees
Copenhagen	26%	53%	23%	52%
Bornholm	0%	0%	1%	0%
West and South Zealand	11%	5%	14%	4%
Fyn	11%	7%	9%	6%
Central and South Jutland	31%	23%	27%	25%
West Jutland	11%	5%	13%	5%
North Jutland	11%	6%	12%	7%

Notes: This table shows the location where immigrant household heads first granted visas in the period 1986-2003 resided in their first year in Denmark.

Table 4: Visa Types Percentages

Type of visa	Before 1999	After 1999
Work	9%	7%
EU	48%	32%
Family reunification	7%	2%
Student	30%	50%
Other types	6%	4%

Notes: This table shows the number of first visa types granted to non-refugee immigrant household heads granted visas in the period 1986-2003.

Table 5: Unconditional staying probabilities

	Before 1999		After 1999		
	Refugees	Non-Refugees	Refugees	Non-Refugees	Diff-in-Diff
After 3 years	80.9%	40.7%	95.4%	35.2%	20.0%
After 8 years	68.4%	23.9%	70.3%	20.1%	5.7%
After 13 years	61.0%	20.1%	62.9%	16.0%	6.0%

Notes: This table shows the unconditional fraction of refugees and non-refugee household heads granted visas in the period 1986-2003 who remain in their first observed city after three, eight, and thirteen years. The last column is difference-in-difference in probabilities of staying after 1999 relative to before, between refugees and non-refugee immigrants.

4 Empirical Strategy and Findings

As we discussed in Section 2, the assignment of refugees to locations was random only conditional on a number of factors such as country of origin, marital status, and the year of refugee status. In addition, refugees were asked in placement interviews about their educational and medical needs as well as location preferences. Before 1999, refugees assigned to unattractive areas were even able to refuse the assignment and be reassigned, although this rarely happened in practice.

We mean to choose a baseline sample intended to make refugee allocations as close as possible to a random assignment. First, we include immigrants granted first visas from 1986 to 2003.[15] This allows us to look at outcomes up to thirteen years after the refugee visa was granted, that is, outcomes ten years after the longest years of stay as required by the policy change in 1999.

[15]There is no attrition in our main empirical exercise by construction. All controls are measured on arrival, so we observe these for all immigrants. The main outcome variable is a dummy which takes the value of one if an immigrant is still in her placement location. We record a zero for this outcome variable if the immigrant either resides in another city in Denmark or if she drops out of the data for any reason. This is important to avoid selection, as non-refugee immigrants are more likely than refugees to leave Denmark rather than move to another Danish city.

In our baseline sample, we do not include immigrants with student or family visas because students have different profiles than refugees and people with family connections in Denmark are strongly tied to particular locations. We omit Copenhagen, because while many non-refugee immigrants choose to first live there, only refugees with special needs are placed there. Finally, we only include refugees placed in the last six months of the calendar year after the policy change in order to avoid those that were able to choose their settlement location when the quotas were not yet filled. Below we will examine how the results are affected by cutting the data in different ways.

We turn into our main exercise in Section 4.1 where we run difference-in-difference regressions to identify dynamic amenities —that staying in a location causes a person to stay there longer. Then, in Section 4.2 we run two additional exercises as robustness checks. In particular, we examine (1) alternative specifications in which the sample and controls are different but the source of identification is the same, and (2) a regression discontinuity design in which the source of identification is not the same. In Section 4.3, we discuss the potential mechanisms behind our results, such as labor force participation and additional fertility.

4.1 Difference-in-Difference Regressions

We use a "difference-in-difference" design to establish causal evidence for dynamic amenities. Specifically, we consider the following regression:

$$Y_{i,\tau} = \beta_{dd}\mathbb{1}_{\{i=\text{refugee}\}} \times \mathbb{1}_{\{t_i^{visa}>=1999\}} + \gamma_r \mathbb{1}_{\{i=\text{refugee}\}} + \gamma_p \mathbb{1}_{\{t_i^{visa}>=1999\}} + X_i\delta + \epsilon_i \tag{1}$$

The dependent variable $Y_{i,\tau}$ is an indicator that equals one if person i was still in her first observed city τ years after her visa was granted in year t_i^{visa}.[16] $\mathbb{1}_{\{i=\text{refugee}\}}$ is an indicator that equals one if person i was a refugee, and zero otherwise (i.e. if i was a non-refugee immigrant). $\mathbb{1}_{\{t_i^{visa}>=1999\}}$ is an indicator that equals one if person i was granted a visa after the policy change in 1999, and zero otherwise (i.e. if the visa was issued before 1999).

The coefficient of interest is β_{dd}. If the common-trend assumption is valid, β_{dd} is the causal effect of the 1999 policy change on the probability that a refugee is still in her assigned city τ years after her visa was granted. We argue that the relevant aspect of the policy change for this outcome is that refugees were required to stay in their placement location for three years, rather than merely encouraged to stay for eighteen months.[17] Dynamic amenities, in contrast to static ones, depend on a person's duration of stay

[16] The exact amount of time since the granting of the immigrant's visa depends on the calendar date that the visa was granted. Our first observation of an immigrant's location is on January 1st in the calendar year after the visa was granted. The fraction of immigrants in their first observed location after three years then means the fraction of immigrants who are still in the first observed location when we observe them on January 1st between three and four years after their visa was granted.

[17] We describe other aspects of the policy change in Section 2.2, and discuss implications of these other aspects of the policy

(tenure) in a location. While the incentives to move that stem from static amenities are not affected by the policy change, we expect the policy change to have an effect on moving through lengthening refugees' location tenure in their initial placement. Consequently, comparing long-run outcomes for the cohorts that arrived before and after the policy change can be used to identify dynamic amenities.

In our case, the common-trend assumption means that in the absence of the policy shock, refugee and non-refugee immigrants would have had common time trends in terms of our dependent variable. While this assumption is impossible to directly check, we can examine refugee and non-refugee time trends in the period before our policy experiment to provide suggestive evidence. We do this graphically in Figure 4. We plot the mean residuals ϵ_i of the following regression, separately for refugees and non-refugee immigrants by immigration year,

$$Y_{i,\tau} = X_i\delta + \epsilon_i \tag{2}$$

In Subplots 4a, 4b, and 4c we include as controls X_i only first city by origin country fixed effects. These fixed effects are important to use in all specifications, because refugee allocations to cities depended on country of origin both before and after the policy change in 1999. We examine outcomes three, eight, and thirteen years after the granting of visas. That is, since the policy change required three years of forced stay, we add two increments of five years.[18]. In Subplots 4d, 4e, and 4f, we include what we will refer to as full controls: fixed effects for age on arrival, gender, marital/family status on arrival, and first city by year of immigration linear time trends.[19] The first vertical line on the left is for 1997, the year in which we can identify refugees based on visa information. The second vertical line is for 1999, the year of the policy change. The dotted lines are 95% confidence intervals.

change for the interpretation of our results in Section 4.3.

[18]We cannot look farther without severely reducing our post policy sample size, since our final year of data is 2016. In any case, we will show results below which suggest that the estimated effect of the policy appears to be constant after around seven years in Denmark.

[19]We also use fixed effects for number of children and family size when we truncate the sample at 1997, because that is when this information becomes available in our data.

Figure 4: Mean Residual Rate of Staying in First City by Year of Visa

First City by Origin FE as Controls

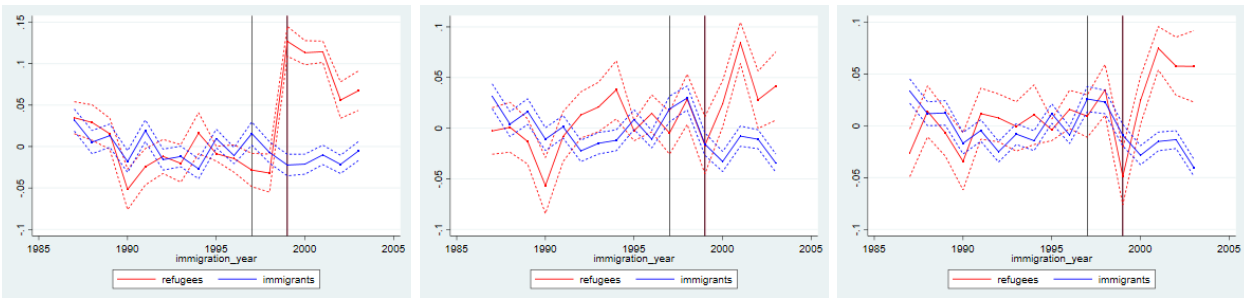

(a) After three years (b) After eight years (c) After thirteen years

Full Controls

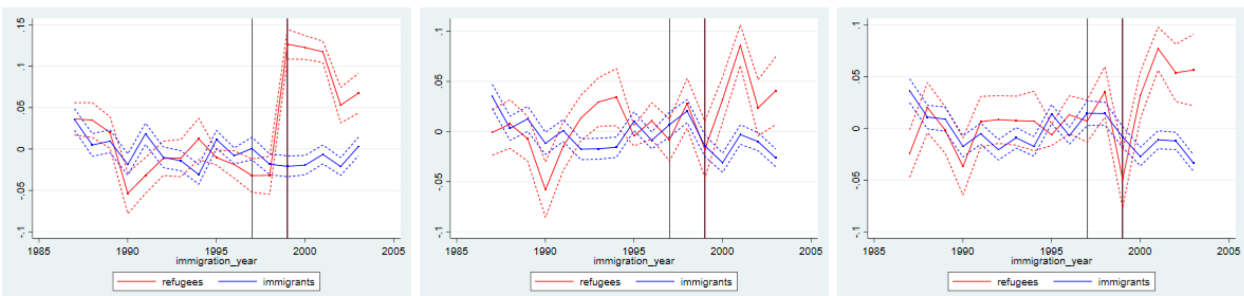

(d) After three years (e) After eight years (f) After thirteen years

Notes: This figure shows the mean rate of staying in a location in 3, 8, and 13 years after arrival, conditional on controls, by year of immigration, separately for refugees and non-refugee immigrants. In Subplots (a)-(b)-(c), the controls are first city interacted with origin country fixed effects. In Subplots (d)-(e)-(f), we allow for our full set of controls consisting of fixed effects for age, gender, marital/family status (all upon arrival) and first city by year of immigration linear time trends.

The pre-policy trends for refugees and non-refugee immigrants appear to be similar, particularly in regressions on outcomes for three and thirteen years after the granting of immigrant visas. For the outcomes eight years after the visa was granted, the pre-policy trends in residuals are more noisy, but do not diverge in any consistent way. After the policy change, there is a clear divergence in the expected direction. In particular, refugees who arrive after 1999 are more likely to still be in their first city in all plots. That the policy was immediately effective is suggested in Subplots 4a and 4d. After the policy change, all cohorts of refugees are much more likely to remain in their placement city three years out relative to immigrants. In contrast, the relative outcomes of the two groups appears to be stable before the policy change. Outcomes after eight and thirteen years are similar, although the cohort granted visas in exactly 1999 at the onset of the policy change saw a decrease in their likelihood of remaining in their placement cities relative to non-refugee immigrants.[20] All other cohorts of refugees react to the policy consistently with the hypothesis that refugee were more likely to remain in their placement city after the policy change.

We report in Table 6 the estimation results of our difference-in-difference equation (1) for $\tau = 13$ years after the visa was granted. The first two columns of the table contain results for our entire sample period from 1986-2003, while the second two columns contain 1997-2003, the period for which we can more confidently identify both refugees and non-refugee immigrants. The coefficient on the refugee indicator, γ_r, is positive, sizable, and statistically significant in all specifications. For those who arrived before the policy change, refugees are 17-21 percentage points more likely to be in the placement city than non-refugees thirteen years after the visa was granted. On the one hand, we might have expected this coefficient to be negative, since non-refugee immigrants could select their preferred settlement location while refugees were initially conditionally randomly allocated. On the other hand, non-refugee immigrants had more moving options than refugees. In particular, they could move to a third country, or return to their home country. Since there are more moving options available to non-refugee immigrants, they are more likely to leave. It is an empirical question whether the former (selection) or the latter (moving options) is dominant. Here, we find that overall the latter is the dominant force.

The coefficient on Post-1999, γ_p, is negative, and statistically significant in all specifications. The estimates imply that non-refugee immigrants were three to five percentage points less likely to remain in their first city for thirteen years if they arrived after 1999. To interpret this effect, recall that in this specification we omit immigrants whose first city was Copenhagen. The trend of leaving the first city may reflect the increasing attractiveness of the Copenhagen area relative to the rest of Denmark in recent years as reflected by the growing share of overall Danish population living in Copenhagen.[21]

Our coefficient of interest, β_{dd}, the interaction between the refugee and post-policy indicators, is positive, statistically significant, and relatively large. We estimate this coefficient to be 7.1 percentage points in our preferred specification under Column (2). We obtain similar estimates in the other specifications, except for the specification truncating the data from 1997 with controls where our estimate is slightly lower at 4.4 percentage points. In words, the 1999 policy change caused refugees to be four to seven percentage points more likely to be observed in their settlement location thirteen years after receiving their visas. To put the

[20] Experimenting with the data, we have found that this result is caused by refugees from Iraq. In Figure A.3 in the appendix we omit Iraqi refugees and re-plot Figure 4. In these plots, the decrease in the relative likelihood of staying in the placement city for thirteen years disappears. We have spoken with an Iraqi family who arrived in Denmark around this time period, but we could not identify any particular reason why Iraqis should be different than other refugees. We continue to include Iraqis in the main analysis, and found that dropping them has little effect on our main estimates.

[21] Our own calculations from Denmark Statistics data show that 41% of the Danish population lived in the Greater Copenhagen Area in 2008, and 44% in 2020.

magnitude of our difference-in-difference estimate in perspective, we note that the unconditional probability that a refugee is observed in his initial location after thirteen years is 55.3%. This means that, based on our preferred specification, dynamic amenities account for 12.8% of the baseline likelihood of staying.

In addition, we run our difference-in-difference regression for each year at which we observe the outcome. In Figure 5, we plot the the coefficient β_{dd} for analogues of Column (2) in Table 6, for τ ranging from 2 to 13 years. As can be seen in the figure, the maximum effect is three years after the visa is granted. This isn't surprising, since the policy essentially forces refugees to remain in their allocated city for three years.[22] After the third year the effect falls in magnitude. The size of the effect is stable from year seven until year thirteen, suggesting that the long-run effect of the policy change was to cause refugees to be around seven percentage points more likely to settle in their placement city.

The takeaway from these exercises is that the policy clearly was implemented as planned and had a strong effect on the fraction of refugees staying in their placement city after three years. More importantly, even thirteen years after the visa was granted, the policy had a statistically and economically significant effect on the fraction of refugees who stayed in the placement city. We interpret our difference-in-difference regression result as evidence in favor of the hypothesis that relocation costs grow over time as a person gets more attached to a location, and that the resulting path-dependence through dynamic amenities can create a long-lasting effect on the location choice of workers.

[22]There were some conditions under which refugees were allowed to change municipalities before three years had passed, which we discuss in Section 2.2 above.

Table 6: Estimation Results —Main Difference-in-Difference Specification

Dependent variable: Still in initial placement after thirteen years in Denmark				
	(1)	(2)	(3)	(4)
Refugee	0.175***	0.170***	0.184***	0.205***
	(0.0158)	(0.0156)	(0.0326)	(0.0321)
Post-1999	-0.0350***	-0.0306***	-0.0504***	-0.0437***
	(0.00846)	(0.00839)	(0.0159)	(0.0156)
Refugee X Post-1999	0.0723***	0.0706***	0.0717***	0.0437**
	(0.0127)	(0.0126)	(0.0177)	(0.0175)
Observations	54,835	54,835	12,408	12,408
R-squared	0.273	0.288	0.258	0.295
Truncate at 97	NO	NO	YES	YES
Controls	NO	YES	NO	YES
First City Trend	NO	YES	NO	YES

Notes: Standard errors are in parentheses. Significance levels are indicated by *** as p<0.01, ** as p<0.05, * as p<0.1. All specifications include country of origin by first city fixed effects. Controls for full data include fixed effects for age, gender and family type (marital status) measured at arrival. If truncated at 1997, controls include family size and number of children measured at arrival as well. We omit immigrants whose first city is Copenhagen and those with student or family reunification visas.

Figure 5: Difference-in-Difference Estimated Coefficient by Years from Arrival

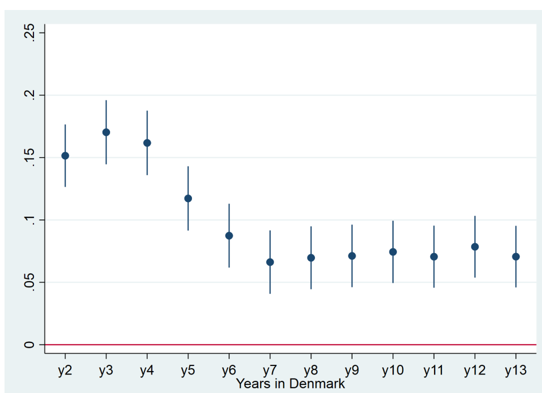

Notes: This figure shows the difference-in-difference coefficient, i.e. the interaction between refugee and policy indicators, from separate regressions for analogues of Column (2) in Table 6 with the outcome of interest (i.e. whether a person is still in her initial location after τ years) ranging for τ between 2 and 13.

4.2 Robustness Analysis

We run two distinct exercises as robustness checks on our main difference-in-difference results which we reported in Section 4.1. First, we re-examine our difference-in-difference regressions using a wide range of alternative specifications. Second, we design a regression discontinuity in which the identification relies on variations within the group of refugees only. These robustness checks give a range of estimates that complement our preferred difference-in-difference point estimates, and reinforce our main finding in supporting dynamic amenities as a factor that causes path dependence in the location of workers.

4.2.1 Alternative Specifications

As with most observational studies, the number of specifications we could choose from is voluminous, both in terms of which controls we include, and exactly with which sample we run our analysis. As we argued in Section 2, refugees' assignment to locations was random only conditional on a number of control variables. In particular, across all of our specifications, we allow for country of origin by first city fixed effect. These fixed effects are important, because both before and after the policy change, refugees were more likely to be assigned to municipalities where a community from their country of origin already existed. That is, refugees might be assigned to cities they prefer, invalidating the conditional random allocation assumption important for removing selection on location choice.[23]

As a robustness analysis, we run our difference-in-difference regression for a wide range of alternative specifications. In Figure 6, we plot the resulting coefficient, β_{dd} from regressions in which we set $\tau = 13$ and consider these alternative specifications:

- In regressions labeled as "Controls" we include age, gender, marital/family status, and if truncated at 1997, also the number of children, and family size fixed effects, all measured the first time they are observed, whereas in "No controls" we do not include them.

- Regressions labeled as "City trend" include a linear time trend in first city, and if "No City trend", then there is no trend included.

- In regressions labeled as "All mnths", the sample includes post-1999 refugees arriving in all calendar months, whereas in "Last 6 mnths" the sample includes only post-1999 refugees arriving in the last six months of the year. As discussed above, the distinction could matter since refugees might be able to influence their settlement locations only toward the first calendar months of a year when quotas were not yet filled.

[23]With our current fixed effects strategy, we exploit changes in behavior of all refugees relative to all non-refugee immigrants in response to the policy change, allowing for differences across mean staying rates for different countries of origin in different first cities. We have experimented with including a three-way fixed effect of country of origin by first city by post-1999 dummy. With this three-way fixed effect, we would be identifying based on differences across mean staying rates of refugee and non-refugee immigrants from a country before and after the policy change. None of our results with that three-way fixed effect are statistically significant. The reason is simply that some countries (Afghanistan, say) overwhelmingly send refugees, and the other countries (Germany, for example) overwhelmingly send non-refugee immigrants. Hence, we have very little remaining variation within countries.

- In regressions labeled as “All visas” the sample includes immigrants with all visa types, whereas in “No Stu/fam” we drop students and family reunification visas. We do so because students have very different moving pattern relative to other immigrants (such as a higher probability of moving to a new city after graduation), and family reunification visa receivers may be more tied to a location than refugees due to family already living there.
- In regressions labeled as “All cities” our sample includes refugees first observed in all cities, whereas in “No Cph”, we drop those whose we first observe in Copenhagen. As discussed earlier, Copenhagen was only available to refugees with special needs, for example special educational needs or certain medical conditions, none of which we observe in the data.
- In regressions labeled as “All prev yrs” the sample includes all years before 1997, while in “After 97 inc”, we drop immigrants arriving before 1997. Identification of refugees is difficult before 1997 because we do not have direct information on visa type prior to 1997. We are more confident in our identification of refugees and non-refugee immigrants after 1997.

In each plot, we order the coefficients by point estimates. The coefficient of interest ranges from around two to nine percentage points, and is statistically significant in all but seven of the 64 specifications. There are not many obvious patterns among specifications which deliver low or high point estimates. The strongest pattern is that truncating the data to use only immigrants granted visas in 1997 and afterward tends to reduce the size of the coefficient.

In the appendix we include similar plots for outcomes of three and eight years after visas are granted (Figures A.1 and A.2). Three years after visas are granted we see a large policy effect (7 to 19 percentage points) across specifications. The results for eight years after visas are granted are roughly in line with those for thirteen years after visas are granted, although the variance of estimates is higher: around thirty percent of the specifications are not statistically significant, and the lowest point estimates are slightly below zero.

Figure 6: Alternative Difference-in-Difference Specifications

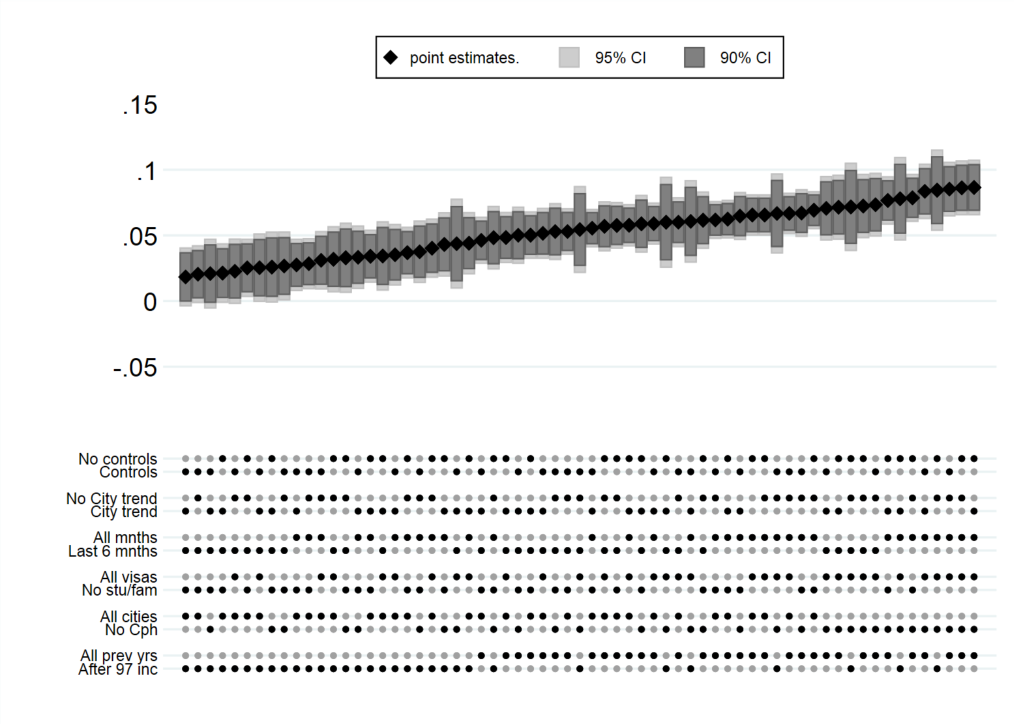

Notes: This figure shows the difference-in-difference coefficient for a wide range of specifications with different controls and sample cuts. Refer to the text in Section 4.2.1 for detailed description of these specifications. In front of each row below the graph, a black filled circle determines which feature is enabled in the corresponding specification. For example, in the first regression on the far left side, we include all controls, no city trend, refugee and non-refugee immigrants who were granted visa in all months, we exclude student and family reunification visas, we include all cities as the first city of placement, and the sample contains only those who came after 1997.

4.2.2 Regression Discontinuity

As another robustness check, we consider a regression discontinuity design centered around the policy change on January 1st, 1999. While we are using the same data and natural experiment, the identification assumption in this exercise is different from that in our baseline difference-in-difference regression. A sufficient identification assumption in this section is that refugees who are granted visas just after the policy change are drawn from the same population as those who were granted visas just beforehand. Therefore comparing outcomes for these two groups is analogous to running a randomized controlled trial.

The object we are trying to estimate in our regression discontinuity exercise is:

$$\beta_{rd} = \mathbb{E}\left[Y^1_{i,\tau} - Y^0_{i,\tau} | t_i^{visa} = t_0\right] \tag{3}$$

Here $Y^1_{i,\tau}$ is an indicator variable which equals one if individual i would have still been in her placement city after τ years in Denmark if she were treated by the post-1999 policy. Similarly, $Y^0_{i,\tau}$ is an indicator variable which equals one if individual i would have still been in her placement city after τ years in Denmark if she were *not* treated by the post-1999 policy. t_i^{visa} refers to the *day* on which the visa was granted, and t_0 is January 1st, 1999.[24] As is standard in the quasi-experimental literature, the challenge is that we only observe people in either the treated state or the untreated state.

Under the identification strategy described above, we can estimate β_{rd} with the following estimator:

$$\begin{aligned}
\hat{\beta}^{+} &= \arg\min_{\beta_0,\beta_1} \sum_{i=1}^{n} \mathbb{1}\left(t_i^{visa} >= t_0\right)\left(Y_i - \beta_0 - \beta_1\left(t_i^{visa} - t_0\right)\right)^2 K_{bw}\left(t_i^{visa} - t_0\right) \\
\hat{\beta}^{-} &= \arg\min_{\beta_0,\beta_1} \sum_{i=1}^{n} \mathbb{1}\left(t_i^{visa} < t_0\right)\left(Y_i - \beta_0 - \beta_1\left(t_i^{visa} - t_0\right)\right)^2 K_{bw}\left(t_i^{visa} - t_0\right) \\
\hat{\beta}_{rd} &= \hat{\beta}_0^{+} - \hat{\beta}_0^{-}
\end{aligned} \tag{4}$$

Intuitively, we estimate a weighted linear regression on distance from the time of policy implementation, with weights given by K_{bw}, which is a kernel function with bandwidth bw. We use a triangular kernel, and bandwidth chosen by the method suggested in Calonico et al. (2014), and report their suggested robust confidence intervals in our results.[25] As before, we run regressions for the dependent variable for tenure τ ranging from two years to thirteen.

Focusing on tenure $\tau = 13$, Figure 7 contains a bin plot showing how mean staying probabilities vary around the threshold, with each of the bins corresponding to approximately one month. Our full estimates are presented in Figure 8, including 95% robust confidence intervals. The bandwidths calculated using the Calonico et al. (2014) method vary from year to year, ranging from 151 to 221 days around the cutoff. The estimates are reassuringly consistent with our baseline difference-in-difference estimates. The point estimates are if anything a bit higher, implying that those who arrived after the policy change were 15 percentage points more likely than those who arrived before the policy change to still be in their placement city after thirteen years in Denmark. The regression discontinuity point estimates are however estimated with more noise. All in all, the resulting estimates from this robustness check reinforces our finding in support of dynamic amenities.

[24] In the difference-in-difference exercise, the date on which visa was granted referred to the corresponding year rather than day.

[25] We implement this method using the Stata package rdrobust on our entire sample. Obviously we cannot look only at refugees arriving in the final months of the calendar year as we did in our difference-and-difference exercise, since our regression discontinuity is calculated on the days surrounding January 1st. While the most recent version of rdrobust allows for the inclusion of controls, we found that including a large number of fixed effects is infeasible. Dropping those that were placed in Copenhagen and including only small dimensional controls has little effect on the estimates.

Figure 7: Regression Discontinuity, Binned Staying Probability 13 Years After Residence Granted

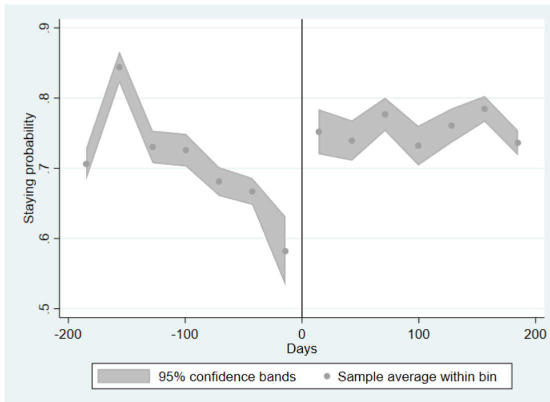

Notes: This figure shows mean staying probabilities for refugees granted residence in Denmark around January 1st, 1999. Each bin represents one month (29 days) for 200 days before and after the policy change. The shaded area represents 95% confidence intervals for each bin mean.

Figure 8: Regression Discontinuity Estimates by Years After Residence Granted

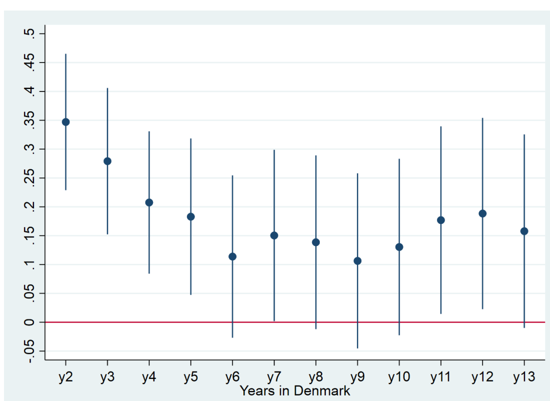

Notes: This figure shows our regression discontinuity results from separate regressions in which the outcome of interest (i.e. whether a refugee is still in her initial location after τ years) ranges for τ between 2 and 13.

4.3 Discussion and Mechanisms

The main objective of our study is to test for a causal relationship from staying in a location on the likelihood of staying there longer, the phenomenon which we call dynamic amenities. A caveat to our empirical strategy is that, as discussed in Section 2, the 1999 change in policy involved lengthening the required stay, but also lengthening the integration programming the municipalities set up for the refugees, and a change in the bureaucracy handling refugee affairs. The effect we find above is the combined effect of all policy changes. We argue that the most salient part of the policy for location choice is the lengthening of the required stay, because training in Danish work culture and language could be used in any part of Denmark.

Once this causal relationship is established, two subsequent questions come to mind. First, what are the mechanisms through which staying in a location creates an attachment to that location? Second, how large are the opportunity costs of this attachment if one leaves a location? We briefly discuss the first question here using extensions of our baseline regressions, and make an effort to address the second one using a simple model of location choice in the next section.

To study the mechanisms behind our result, we replace the dependent variable in our preferred difference-in-difference specification with a number of other individual characteristics. In a word, we are testing for the effect of being forced to remain in a location for three years on dependent variables such as work outcome or having more children.[26] If the policy had an effect on these variables, it may shed light on what mechanisms are behind the persistence in refugees location choice. We consider the following outcomes: (a) whether a refugee is married if he was single upon arrival, (b) whether a refugee is still married if he was married upon arrival, (c) additional number of children relative to the first year of stay in Denmark, (d) whether a refugee is working if he was not working in his first observed year, (e) whether a refugee is working if he was working in his first observed year, (f) whether a refugee receives government support.[27]

[26] An important caveat on these regressions is that we must drop observations for any immigrant who leaves Denmark, since we only observe individual characteristics such as employment as long as an immigrant remains in Denmark. This may introduce selection into our estimates if the choice to leave Denmark is correlated with our outcomes of interest. We avoided this selection in our main exercises above by defining the outcome variable of interest as staying in the first observed city. If someone leaves Denmark, we know that they did not stay.

[27] We say a refugee is working if she has non-zero labor income. We say a refugee receives government support if he has non-zero income classified as either unemployment payments or welfare payments.

Figure 9: Mechanism regressions

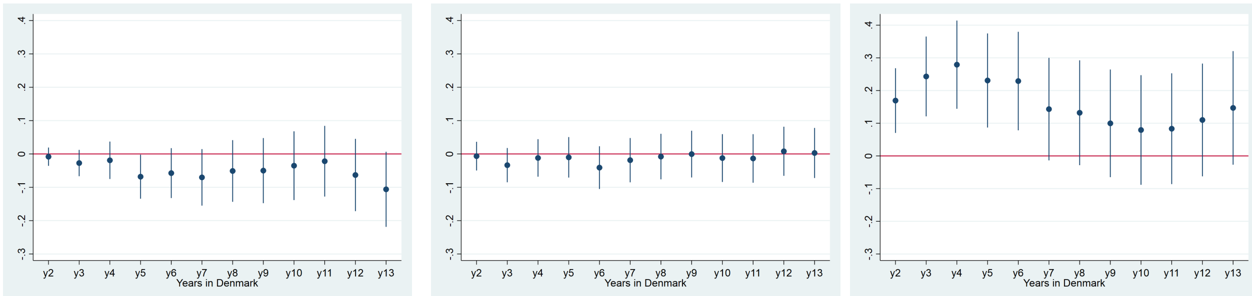

(a) Married (single on arrival) (b) Married (married on arrival) (c) Additional Children

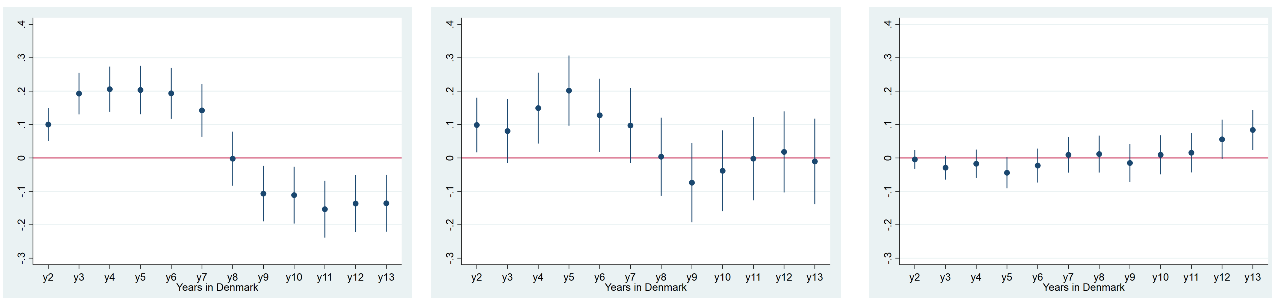

(d) Working (unemp. on arrival) (e) Working (working on arrival) (f) Any government support

Notes: This figure shows the difference-in-difference coefficient, i.e. the interaction between refugee and policy indicators, with various outcome variables in separate regressions for years in Denmark ranging from one year to thirteen years. The outcome variables are: (a) whether the refugee is married for the sample who were single upon arrival, (b) whether the refugee is still married for the sample who were married upon arrival, (c) additional number of children relative to the first year in Denmark, (d) whether the refugee is working for the sample who were not working in the first year, (d) whether the refugee is working for the sample who were working in the first year, (f) government support. Except for dropping observations where a refugee leaves Denmark and the outcome variables, the regression sample and specification is identical to that in Column 4 of Table 6.

Figure 9 presents the results from this exercise in a form analogous to the way we presented the main results in Figure 5. In Panels (a) and (b), we see that the policy had no significant effect on marriage, neither for those who were single on arrival, nor for those married on arrival. Panel (c) shows the effect of the policy on children. Since we have fixed effects for number of children on arrival, we are examining additional children as an outcome. The estimates appear to be positive and statistically significant in the first few years, and not statistically different from zero after a decade or so. Since children are a stock rather than a flow, this panel points to change in the timing of fertility. Refugees forced to stay for three years had ultimately the same number of children as migrants who were not treated, but treated refugees had their children earlier.

Panel (d) and (e) show that treated refugees were more likely to find a job whether or not they were observed working in their first year. This effect is significant for the first few years, and fades away by the eighth year in Denmark, turning negative for those who were not initially observed working. We split the sample into those working in the first observed year and those not working in the first observed year because the control group of non-refugee immigrants may be more comparable if they have the same work status as refugees in their first year. The graph for the pooled sample looks similar. Panel (f) shows that the policy had little effect on government support take up.

Our results indicate that refugees who were forced to stay for three years were more likely to immediately have children and find work in their placement city. It is plausible that there is a sunk cost to searching for work, and it makes sense to do so only if one plans on remaining in a location for a long enough period to cover the cost. Once one has a job, the job itself keeps one in a location. Moreover, coworkers naturally become friends.[28] Having children, by the same token, may create a natural social network of other parents for the refugees themselves, and the children may also develop friends and connections of their own as they get older. These social bonds may be costly to break by leaving the placement city. This interpretation is speculative, and we hope that future research can further illuminate the mechanism through which dynamic amenities operate.

5 Calibration Exercises

In this section, we measure the magnitude of moving costs that arise from dynamic amenities relative to total moving costs. To do so, we write down a simple model of location choice with which we can calibrate the effect from dynamic amenities compared to static costs of moving.

5.1 A Simple Model of Location Choice

To help with the exposition, let i refer to individual persons, j and k to locations, τ to the amount of time spent in a location which we refer to as "tenure", and t to a particular time period. In the first period which we call $t-1$, person i is characterized by her initial location, $j \in J$, as well as the number of periods she has been there, $\tau = 1, 2, ...$ We call (j, τ) the state of person i. We denote the second period by t, which we think of as a shorthand way of modeling preferences for the rest of lifetime. Consider person i in the state

[28]The policy change in 1999 which required refugees to remain in the settlement location for three years also extended integration coursework from one and a half to three years, and changed the bureaucracy surrounding refugee placement. We cannot rule out that these facets of the reform may have had a direct effect on job finding. For example, another possible reason that refugees are more often observed in the labor force after 1999 is that, as described in Section 2.2, refugees were allowed to move to a different municipality only if they had a job contract there. This may have given refugees more incentive to work in the formal sector rather than the informal sector, where they would not be recorded as working in our data.

(j,τ). We denote the logarithm of her utility in period t if she selects to reside in location $k \in J$ by $v^i_{kt}(j,\tau)$,

$$v^i_{kt}(j,\tau) = \begin{cases} \omega_{kt} + \eta^i_k - c + z^i_{kt} & k \neq j \\ \omega_{kt} + \eta^i_k + \phi\tau + z^i_{kt} & k = j \end{cases}$$

Here, ω_{kt} is city-specific amenities for the average person in period t, η^i_k is time-invariant person-specific valuation of amenities of location k, c is the static cost of moving, ϕ is the utility from one additional year of growing roots in a location, and z^i_{kt} is person-location specific preference shock in period t.

We assume that $\exp(z^i_{it})$ is drawn independently across persons and over time from a Fréchet distribution with dispersion parameter θ and a location parameter normalized to one.[29] Then, the probability that person i in the state (j,τ) moves to city k compared to staying in city j is:

$$\frac{\pi^i_{kt}(j,\tau)}{\pi^i_{jt}(j,\tau)} = \left[\frac{\exp(\omega_{kt} + \eta^i_k - c)}{\exp(\omega_{jt} + \eta^i_j + \phi\tau)}\right]^\theta$$

which can be written in log terms as:

$$\ln \pi^i_{kt}(j,\tau) - \ln \pi^i_{jt}(j,\tau) = (\tilde{\omega}_{kt} - \tilde{\omega}_{jt}) + (\tilde{\eta}^i_k - \tilde{\eta}^i_j) - \tilde{c} - \tilde{\phi}\tau \tag{5}$$

where $\tilde{\omega}_{jt} \equiv \theta\omega_{jt}$, $\tilde{\eta}^i_j \equiv \theta\eta^i_j$, $\tilde{c} \equiv \theta c$, $\tilde{\phi} \equiv \theta\phi$. Equation (5) is the probability of moving to city k relative to staying in city j for individual i in the state of (j,τ). This relative probability is larger if average utility received from city k relative to j is larger due to higher wages, lower rents, better climate, etc. captured by time-city-specific variable ω. It is larger if person i attaches a higher value to city k relative to j captured by person-city variable η. It is also larger if static moving costs c are smaller, and if the dynamic amenity value due to staying in city j for τ periods, $\phi\tau$, is lower. The effects from all these channels will be higher if θ as the elasticity of moving with respect to changes in the net non-stochastic benefit of moving is larger.

Without loss of generality, we set the unconditional mean of person-specific amenities to zero $\mathbb{E}[\tilde{\eta}^i_k] = 0$. Let $G_t(j,\tau)$ refer to any generic group of people in year t who are in the state of (j,τ). Due to selection, the conditional expectation of person-specific preferences for a location given current location and tenure $\mathbb{E}\left[\tilde{\eta}^i_k | i \in G_t(j,\tau)\right]$ is generically nonzero. Define $P^G_{kt}(j,\tau)$ as the relative average log probabilities in period

[29] This normalization is without loss of generality since we include η^k_i as a person-city-specific term. In addition, an equivalent formulation is to have z being drawn from Extreme Value Type I with scale parameter $1/\theta$ and location parameter zero. Let $\exp(Z) = Y$ be a random variable that has a Fréchet distribution with dispersion parameter θ and location parameter one, that is $\Pr(Y \leq y) = \exp(-y^{-\theta})$ then: $\Pr\Big(Z \leq z\Big) = \Pr\Big(Y \leq \exp(z)\Big) = \exp\Big(-(\exp(z))^{-\theta}\Big) = \exp\Big(-\exp(-\theta z)\Big)$.

t of all persons in group $G(j,\tau)$ moving to k compared to staying in j,

$$\begin{aligned} P_{kt}^{G}(j,\tau) &\equiv \mathbb{E}[\ln \pi_{kt}^{i}(j,\tau) - \ln \pi_{jt}^{i}(j,\tau) | i \in G_t(j,\tau)] \\ &= (\tilde{\omega}_{kt} - \tilde{\omega}_{jt}) + \mathbb{E}[\tilde{\eta}_k^i - \tilde{\eta}_j^i | i \in G_t(j,\tau)] - \tilde{c} - \tilde{\phi}\tau \end{aligned} \tag{6}$$

We define the expression in equation (6) for two groups. First, let $R_t^{(\tau_0)}(j,\tau)$ be a group of people who were *randomly* assigned to location j with a required stay of duration of τ_0 periods (refugees). Second, let N be a *non-randomly* assigned group (non-refugee immigrants). A key feature of group R is that since their assignment to j was random, for any year within the forced stay $\tau = 1, 2, ..., \tau_0$,

$$\mathbb{E}[\tilde{\eta}_k^i | i \in R_t^{(\tau_0)}(j,\tau)] = 0, \quad \text{for all } k \in J \tag{7}$$

We use our structure to derive equations that describe refugees' moving differentials relative to non-refugee immigrants. Consider refugees as a randomly assigned group $R^{(\tau_0)}$ who were required to stay in j for a duration of τ_0. By structure of the policy, refugees have to stay in their initial location until the end of the required stay. Right at the end of staying in j for τ_0 periods, in some year t_0, refugees in group $R^{(\tau_0)}$ can choose to move according to their will. Consider a change in policy, such as the one in Danish immigration policy explained in Section 2, that required a lengthening of the forced stay from τ_0 to τ_1. Our difference-in-difference design compares moving probability differentials of group $R^{(\tau_0)}$ and $R^{(\tau_1)}$ relative to non-refugee immigrants. Let the years that correspond to the end of forced stay before and after the policy change be t_0 and t_1. For people in $R^{(\tau_0)}$, choices begin after τ_0 periods in year t_0, and for people in $R^{(\tau_1)}$, choices begin after τ_1 periods in year t_1. Specifically,

$$\left[P_{kt_1}^{R^{(\tau_1)}}(j,\tau_1) - P_{kt_1}^{N}(j,\tau)\right] - \left[P_{kt_0}^{R^{(\tau_0)}}(j,\tau_0) - P_{kt_0}^{N}(j,\tau)\right] = \underbrace{-\tilde{\phi}(\tau_1 - \tau_0)}_{\lambda^{dd}} + \epsilon_{kjt_0t_1}^{dd} \tag{8}$$

Here, our identification assumption requires that the difference between unobserved preferences of the non-refugee group for cities j and k does not systematically change with the calendar year.[30] In addition, using the in-flow and out-flow migration rates of refugees between any pair of cities,

$$-0.5\left[P_{kt}^{R^{(\tau_1)}}(j,\tau_1) + P_{jt}^{R^{(\tau_1)}}(k,\tau_1)\right] = \underbrace{\tilde{c} + \tilde{\phi}\tau_1}_{\lambda^{cc}} + \epsilon_{kjt}^{cc} \tag{9}$$

Here, the residual can be thought of as the variation in static moving costs around the mean, and $\tilde{c}$ is meant

[30] Specifically, the error term is given by $\mathbb{E}[\tilde{\eta}_j^i - \tilde{\eta}_k^i | i \in N_{t_1}(j,\tau)] - \mathbb{E}[\tilde{\eta}_j^i - \tilde{\eta}_k^i | i \in N_{t_0}(j,\tau)]$ where every term including tenure τ is the same but the sample of non-refugee immigrants change between year t_0 and t_1.

to measure the symmetric component of static costs.[31]

We can recover ϕ relative to static cost c by dividing the two terms that are confounded by the effect of migration elasticity. This is an advantage of our approach despite the challenge of identifying dynamic amenities ϕ or static migration costs c separately from the migration elasticity θ. Specifically, obtaining λ^{dd} and λ^{cc} from equations (8) and (9),

$$\frac{\phi}{c} = \frac{-\lambda^{dd}}{(\tau_1 - \tau_0)\lambda^{cc} + \tau_1 \lambda^{dd}} \tag{10}$$

We now turn to bringing equations (8) and (9) to our Danish data on immigration and using the results to measure ϕ/c based on equation (10).

5.2 Bringing the Structure to Data

We face two empirical challenges in order to calibrate the ratio of ϕ to c. First, flows of migration in a given origin-destination-year for a given group of immigrants with a given length of tenure is too sparse to be informative. To deal with this issue, we aggregate moving flows along the dimension of cities, and along the dimension of calendar years. We aggregate cities based on geographical proximity into three regions consisting of: 1. Copenhagen 2. Western and Southern Zealand and Fyn, and 3. Jutland including Aarhus. We aggregate calendar years into four-year periods, consisting of three periods for those who arrived before the policy change (1989-92, 1993-96, 1997-2000), and three periods for those who arrived after the policy change (2003-06, 2007-10, 2011-14).[32]

Putting all this together, the data which we use for our structural exercise consists of 3 by 3 matrices of across-region moving flows for the two groups of refugees and non-refugee immigrants, with the outcomes of interest, including length of tenure, to be observed in three periods for the sample who arrived before the policy change and three periods for those who arrived after the policy change.

The second challenge we face is that variable $P^G_{kt}(j,\tau)$, defined in equation (6), is the mean of the log probability of migration, not the log of the mean probability. Using the log of mean moving flows can create a bias that gets larger for smaller moving flows. An additional advantage of aggregating our data, as explained above, is to mitigate concerns about such small flows. We further address the potential for such a bias by using a second order Taylor approximation of the mean log function to minimize the error in our

[31] Consider a more flexible specification in which static costs vary across pairs of origin-destination (j,k), that is to replace $\tilde{c}$ by $\tilde{c}_{jk}$. Then, the RHS of equation (9) would equal $\tilde{c}^S_{jk} + \tilde{\phi}\tau_1 + \epsilon^{cc}_{kjt}$, where $\tilde{c}^S_{jk}$ is the symmetric component between $\tilde{c}_{jk}$ and $\tilde{c}_{kj}$.

[32] These choices are governed by the structure of our data and exercise. The most recent year we observe a person who arrived before 1999 and has a tenure of one year is 2000, and the earliest year we observe a person who arrived in or after 1999 and has a tenure of three years is 2003.

variable construction.[33]

In addition, we emphasize that our calibration is contingent on the policy requirement for years of stay. In particular, given the structure of Danish immigration policy, equations (7) and (8) hold only at $\tau_0 = 1$ and $\tau_1 = 3$. This is because we can control for the selection margin in refugees' residential location only within the period of their required stay —in which they do not select their region of residence. Hence, we can calibrate ϕ based on moving decisions in the first and third years of tenure in a location. The benefits from staying for another year in a location might be non-linear in the years of stay. For example, a person may grow roots in a place in early years of stay with these benefits to flatten as she stays there longer. Or, such benefits may be modest in the first few years and peak only after several years. We will then require a linearity assumption to attribute our calibrated ϕ to other years of stay after the third year.

With these considerations, we bring our structure to data. In line with the specifics of Danish immigration policy, we set $\tau_0 = 1$ corresponding to the before-policy periods referred to as t_0, and $\tau_1 = 3$ corresponding to the after-policy periods referred to as t_1. We assign weights to observations based on the population of immigrants in the origin region in pre-policy period.[34]

Our calibration delivers $\lambda^{dd} = -13.5$ and $\lambda^{cc} = 63.5$. Plugging these into equation (10) then sets ϕ/c at 0.156. This means that the opportunity cost of moving that is generated by staying in a location for an additional year is 15.6% of the static cost of moving. To put this calibrated value in perspective, suppose the dynamic amenity is a location-specific asset whose value accumulates at a linear rate over time, and the annual discount rate is 5% in line with the estimates of annual cost of capital. Then, a back-of-the-envelope calculation suggests that seven years of dynamic amenities generate an opportunity cost of moving that equals current one-time static moving costs.

Discussion. Our goal in the above calibration is only to give the reader a rough idea of the quantitative significance of dynamic amenities. Since we aggregated our data along several dimensions to deal with the sparsity of migration flows, we could not include important controls as in our reduced form estimations. Our approach could allow for such controls only if we had a larger sample of refugees. That being said, we have

[33]Specifically, $\mathbb{E}[\ln(X)] \approxeq \ln(\mathbb{E}[X]) - \frac{1}{2\ln(\mathbb{E}[X])^2}\mathbb{V}[X]$. Neither the problem we face nor the solution we adopt is rare. For example, in the context of trade rather than migration, Kleinman et al. (2020) consider a Taylor-series expansion, and show that the bias from using the log of the mean rather than the mean of the log is small for small policy changes.

[34] We have 54 observations for equation (8) in which λ_{dd} is the intercept. These observation are from six inter-region flows (all pairs of j and $k \neq j$) for the nine combinations of periods t_0 and t_1. The 25th, 50th, and 75th percentile among these observations are -29.64, -10.47, and -2.07. To run this regression we must choose a tenure τ for non-refugee immigrants. Theoretically any value would work. In practice, we choose $\tau = 1$ because moving frequency is higher in the first year, and this property helps us minimize the approximation error in calculating the average log values. Setting τ for non-refugee immigrants at higher values increases our calibrated ϕ/c. In this sense, our estimates of λ_{dd} are conservative. In addition, we have 18 observations for equation (9) for which λ_{cc} is the intercept. These observations are from 6 inter-region flows for all three periods after the policy change. The 25th, 50th, and 75th percentile among these observations are 42.11, 46.38, and 57.44.

several remarks to make in support the robustness of our calibration even though our data on flows between locations is sparse.

First, in our specification, we make a few conservative assumptions such as the one mentioned in Footnote 34. Also, without using weights in our regressions our calibrated ratio would be 24.7% (with $\lambda^{dd} = -16.8$, $\lambda^{cc} = 59.2$), which is larger than the 15.6% we find in our preferred specification. Second, in our reduced-form difference-in-difference regression with tenure $\tau = 3$, including controls only alters the difference-in-difference coefficient a modest amount. This suggests that including controls in our structural difference-in-difference equation (8) may not have a drastic effect on our calibration results. Finally, we have tried to partially control for refugee origin and age by restricting our sample to immigrants between 25 and 49 years from the top Middle Eastern refugee-sending countries, Iraq, Afghanistan, Syria, and Iran, which together account for more than half of refugees in the post-policy period. Even with this generous cut, the number of non-zero flow observations that we can use to estimate equation (8) falls from 54 in our full sample to only 16 in the restricted sample. Repeating our structural difference-in-difference exercise on this restricted sample, we calibrate $\lambda^{dd} = -7.0$. The resulting ratio of ϕ/c is around half of the ratio we calibrate with the full sample.

6 Concluding Remarks

In this paper, we examine the hypothesis that people become more attached to a location the longer they live there. We are motivated to evaluate this hypothesis because in the literature little is known about the mechanisms behind the persistence of cities and what causes path dependence in workers' locations, about forces behind moving costs, and why available estimates of moving costs are so large.

The quantitative trade and economic geography literature has grown quickly with the recent development of proper tools and the availability of appropriate data. We have also benefited from these tools in other areas of our research. However, in this paper we are shifting away from that approach. Studies in this literature have incorporated an increasingly complicated apparatus to conduct policy analysis in general equilibrium. The key benefit from this literature is arguably putting numbers on variables, such as wages, prices, and welfare indices, in response to counterfactual policy. This benefit comes with its own costs, however, since one has to make strong assumptions such as restrictive functional forms to bring these models to data. All those restrictions not only call into question the validity of the numbers these studies generate, but also limit the usefulness of these models in the identification of new mechanisms.

For these reasons, we have intentionally took a more reduced-form approach in order to impose less restrictive assumptions for identification. To identify a relationship that points to a causal effect beyond correlations, we utilize a unique opportunity created by a policy change in refugees' settlement in Denmark. The resulting identification in our difference-in-difference study can be thought of as an average treatment effect on a particular group of people in a certain social context. Without more structure we cannot turn these estimates into implied measures of moving costs. As such, we provide an example of such structure in our simple calibration.

References

Alessandria, G., Kaboski, J. P., and Midrigan, V. (2010). Inventories, lumpy trade, and large devaluations. *American Economic Review*, 100(5):2304–39.

Allen, T. (2014). Information frictions in trade. *Econometrica*, 82(6):2041–2083.

Allen, T. and Donaldson, D. (2018). The geography of path dependence. *Unpublished manuscript*.

Autor, D. H., Dorn, D., and Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584):621–646.

Azlor, L., Damm, A. P., and SCHULTZ-NIELSEN, M. L. (2018). Local labour demand and immigrant employment. *Rockwell Foundation Working Paper*.

Azlor, L., Damm, A. P., and Schultz-Nielsen, M. L. (2020). Local labour demand and immigrant employment. *Labour Economics*, page 101808.

Baum-Snow, N. and Ferreira, F. (2015). Causal inference in urban and regional economics. In *Handbook of regional and urban economics*, volume 5, pages 3–68. Elsevier.

Bleakley, H. and Lin, J. (2012). Portage and path dependence. *The quarterly journal of economics*, 127(2):587–644.

Bosker, M., Brakman, S., Garretsen, H., and Schramm, M. (2007). Looking for multiple equilibria when geography matters: German city growth and the wwii shock. *Journal of Urban Economics*, 61(1):152–169.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Coate, P. and Mangum, K. (2019). Fast locations and slowing labor mobility.

Dalgaard, C.-J., Kaarsen, N., Olsson, O., and Selaya, P. (2018). Roman roads to prosperity: Persistence and non-persistence of public goods provision.

Damm, A. P. (2005). *The Danish Dispersal Policy on Refugee Immigrants 1986-1998: A Natural Experiment?* Department of Economics, Aarhus School of Business Aarhus.

Damm, A. P. (2009). Ethnic enclaves and immigrant labor market outcomes: Quasi-experimental evidence. *Journal of Labor Economics*, 27(2):281–314.

Damm, A. P. and Dustmann, C. (2014). Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review*, 104(6):1806–32.

Davis, D. R. and Weinstein, D. E. (2002). Bones, bombs, and break points: the geography of economic activity. *American Economic Review*, 92(5):1269–1289.

Dustmann, C., Fasani, F., Frattini, T., Minale, L., and Schönberg, U. (2017). On the economics and politics of refugee migration. *Economic policy*, 32(91):497–550.

Dustmann, C., Vasiljeva, K., and Piil Damm, A. (2018). Refugee migration and electoral outcomes. *The Review of Economic Studies*.

Eaton, J., Eslava, M., Jinkins, D., Krizan, C., and Tybout, J. (2020). A search and learning model of export dynamics. Technical report, Working paper.

Eckert, F., Hejlesen, M., and Walsh, C. (2020). The return to big city experience: Evidence from refugees in denmark.

Edin, P.-A., Fredriksson, P., and Åslund, O. (2003). Ethnic enclaves and the economic success of immigrants—evidence from a natural experiment. *The quarterly journal of economics*, 118(1):329–357.

Foged, M. and Peri, G. (2016). Immigrants' effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics*, 8(2):1–34.

Hummels, D. L. and Schaur, G. (2013). Time as a trade barrier. *American Economic Review*, 103(7):2935–59.

Kaplan, G. and Schulhofer-Wohl, S. (2017). Understanding the long-run decline in interstate migration. *International Economic Review*, 58(1):57–94.

Kennan, J. and Walker, J. R. (2011). The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251.

Kleinman, B., Liu, E., and Redding, S. J. (2020). International friends and enemies.

Kline, P. and Moretti, E. (2013). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly Journal of Economics*, 129(1):275–331.

Krugman, P. (1991). History and industry location: the case of the manufacturing belt. *The American Economic Review*, 81(2):80–83.

Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Michaels, G. and Rauch, F. (2018). Resetting the urban network: 117–2012. *The Economic Journal*, 128(608):378–412.

Miguel, E. and Roland, G. (2011). The long-run impact of bombing vietnam. *Journal of development Economics*, 96(1):1–15.

Munch, J. R., Rosholm, M., and Svarer, M. (2006). Are homeowners really more unemployed? *The Economic Journal*, 116(514):991–1013.

Nakamura, E., Sigurdsson, J., and Steinsson, J. (2016). The gift of moving: Intergenerational consequences of a mobility shock. Technical report, National Bureau of Economic Research.

Tombe, T. and Zhu, X. (2019). Trade, migration, and productivity: A quantitative analysis of china. *American Economic Review*, 109(5):1843–72.

Udlaendinge- og Integrationsministeriet (1998). Integrationslov, lov nr. 474.

Voigtlaender, N., Becker, S., Grosfeld, I., Grosjean, P., and Zhuravskaya, E. (2020). Forced migration and human capital: Evidence from post-wwii population transfers. *The American Economic Review*.

Yagan, D. (2014). Moving to opportunity? migratory insurance over the great recession. *Job Market Paper*.

Zabek, M. A. (2019). Local ties in spatial equilibrium.

A Additional Figures and Tables

Figure A.1: Diff-in-diff specifications, 3 years after visa was granted

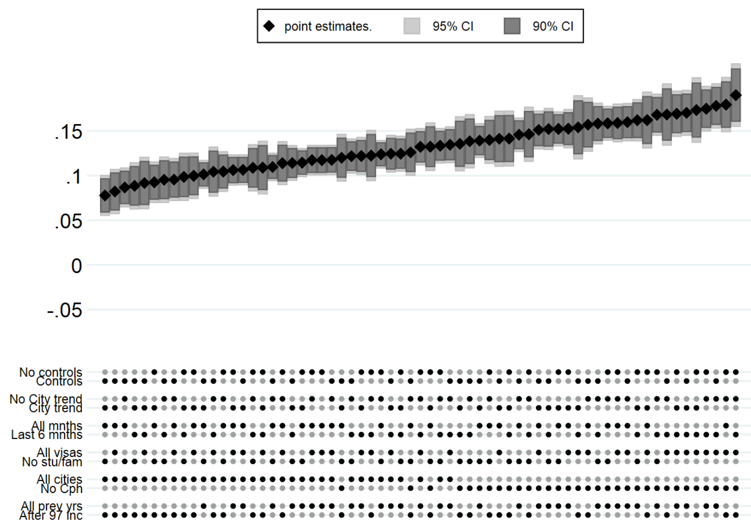

Figure A.2: Diff-in-diff specifications, 8 years after visa was granted

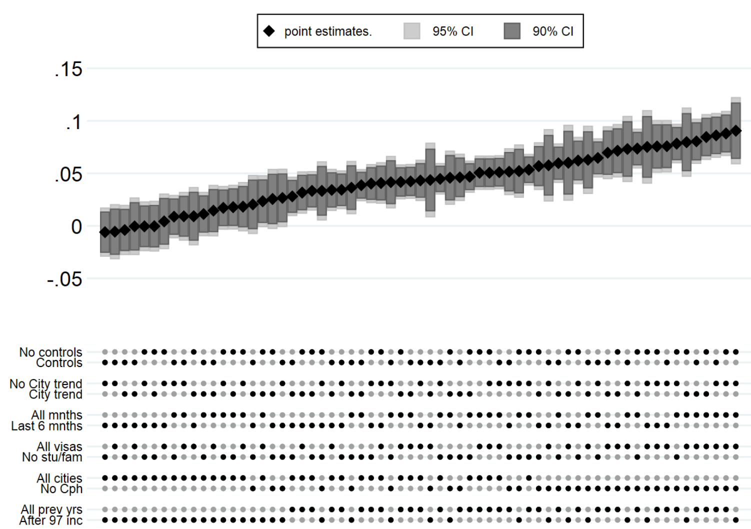

Figure A.3: Difference-in-Difference Plots: No Iraq

Residual Fraction of Refugees Still in First City by Year of Visa, First City by Origin FE

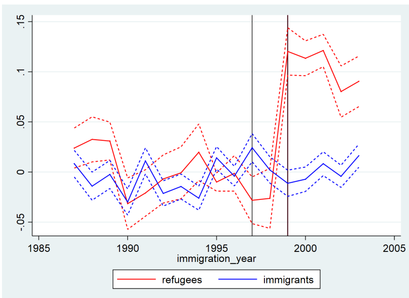

(a) After three years

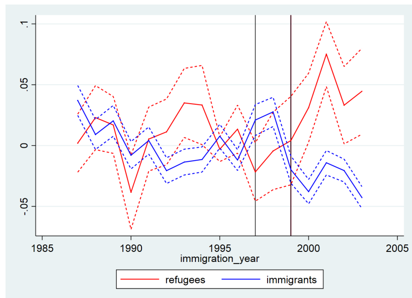

(b) After eight years

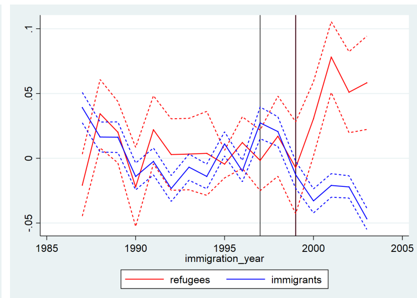

(c) After thirteen years

Residual Fraction of Refugees Still in First City by Year of Visa, Full Controls

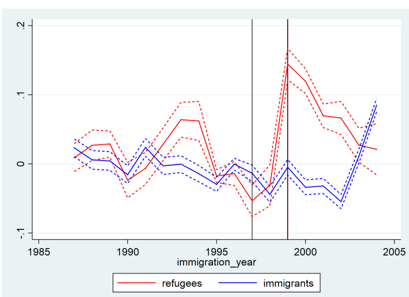

(d) After three years

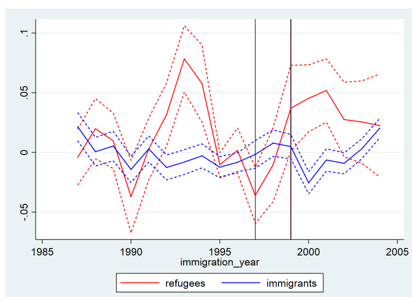

(e) After eight years

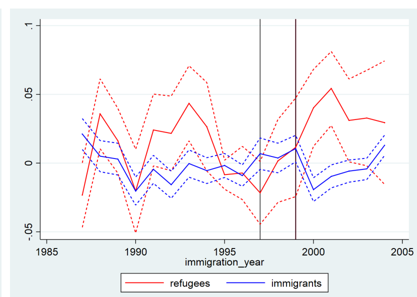

(f) After thirteen years