Department of Economics

Working Paper Series

Homeowner Preferences after September 11th, a Microdata Approach

Adam Nowak

Juan Sayago-Gomez

Working Paper No. 17-17

This paper can be found at the College of Business and Economics Working Paper Series homepage:

http://business.wvu.edu/graduate-degrees/phd-economics/working-papers

Homeowner Preferences after September 11th, a Microdata Approach

Adam Nowak*
West Virginia University

Juan Sayago-Gomez †
Icesi University

July 17, 2017

*CORRESPONDING AUTHOR, Department of Economics, College of Business and Economics, PO Box 6025, Morgantown WV 26506-6025; email: adam.d.nowak@gmail.com;

†Department of Economics, School of Business and Economic Studies; email: jtsayago@icesi.edu.co

Abstract

The existence of homeowner preferences - specifically homeowner preferences for neighbors - is fundamental to economic models of sorting. This paper investigates whether or not the terrorist attacks of September 11, 2001 (9/11) impacted local preferences for Arab neighbors. We test for changes in preferences using a differences-in-differences approach in a hedonic pricing model. Relative to sales before 9/11, we find properties within 0.1 miles of an Arab homeowner sold at a 1.4% discount in the 180 days after 9/11. The results are robust to a number of specifications including time horizon, event date, distance, time, alternative ethnic groups, and the presence of nearby mosques. Previous research has shown price effects at neighborhood levels but has not identified effects at the micro or individual property level, and for good reason: most transaction level data sets do not include ethnic identifiers. Applying methods from the machine learning and biostatistics literature, we develop a binomial classifier using a supervised learning algorithm and identify Arab homeowners based on the name of the buyer. We train the binomial classifier using names from Summer Olympic Rosters for 221 countries during the years 1948-2012. We demonstrate the flexibility of our methodology and perform an interesting counterfactual by identifying Hispanic and Asian homeowners in the data; unlike the statistically significant results for Arab homeowners, we find no meaningful results for Hispanic and Asian homeowners following 9/11.

Key Words: house prices, ethnicity, homeowner preferences, terrorism, September 11th
JEL Codes: R21, R23, R31, J15

We are grateful to Omar Alothimeen, Dan Grossman, Brad Humphreys, Crocker Liu, Amir Neto, Patrick Smith, and participants in the "Local Externalities" panel at ASSA - AREUEA. We are incredibly grateful to Vikram Maheshri for his discussion of the paper. All errors are our own.

1 Introduction

Preferences over neighborhood and neighbor characteristics are fundamental to models of economic sorting. On September 11, 2001 (9/11), 19 terrorists from 4 Arab countries - Egypt, Lebanon, Saudi Arabia and the United Arab Emirates - attacked targets in New York City and Washington DC. The How Americans Responded (HAR) survey carried out immediately after 9/11 found that the attacks affected individual preferences for both Arabs and Muslims.[1] A majority of respondents in HAR (70%) viewed African, Hispanic, and Asian Americans as favorable, but less than half of the respondents (46%) viewed Arabs or Muslims as favorable. Furthermore, in the 30 days after 9/11, the Anti Defamation League recorded 12 instances of either anti-Arab or anti-Muslim violence.[2] This study investigates whether or not the events of 9/11 caused preferences for Arab neighbors to change, and if these changes were reflected in residential real estate prices.

Using a differences-in-differences approach in a hedonic pricing model, we test this hypothesis using transactions for single-family homes obtained from the King County Assessor's Office. King County is in Washington state and includes the Seattle metropolitan area. We exploit cross-sectional and temporal variation in sale prices in order to isolate price effects attributable to 9/11. Our results are both plausible and statistically significant: 1) relative to sales before 9/11, houses sold 180 days after 9/11 and within 0.1 miles of an Arab neighbor sold at a 1.4% discount, 2) this effect is temporary as the discount is not statistically significant 180 days after 9/11, and 3) the effect is not attributable to the presence of nearby mosques. These conclusions are robust across a wide number of model specifications and identifying alternatives.

Economic research has focused on 2 channels by which economic variables, such as property prices, can be affected by terrorism: expectations and preferences. In the context of real estate, the expectations channel is straightforward. Property valuations will decrease following a terrorist attack if 1) the likelihood of a future terrorist attack increases and 2) a particular piece of real estate is a realistic target of a terrorist attack. Although HAR indicates that 79% of respondendents became more concerned of another terrorist attack following 9/11, it is unlikely that individuals living in single-family homes believed their homes were potential targets of a terrorist attack.[3] By

[1] http://www.isr.umich.edu/cps/har/ How Americans Responded is a survey project at the Institute for Social Research, the University of Michigan. The 613 panel participants in the project were first contacted October 17, 2001 and then re-surveyed April, 16 2002.

[2] http://archive.adl.org/terrorism_america/adl_responds.html

[3] This is not to say that all properties are unlikely targets of terrorism as Abadie and Dermisi (2008) find evidence

using single-family homes in a location more than 3,000 miles from New York City, we preclude any price effects resulting from the expectations channel.

Static preferences towards Arab neighbors are in line with correspondence experiments that find landlords discriminate against non-Arabs or non-Muslims, Ahmed and Hammarstedt (2008) and Bosch et al. (2010). Similar to Gautier et al. (2009), we use an event study where non-Arab homeowners initially have arbitrary preferences for their current Arab neighbors, but the events of 9/11 change these preferences. Because we examine transactions within 0.1 miles of Arab homeowners, we use *neighbor* in the geographic sense. We note it would be interesting to perform an analysis similar to Linden and Rockoff (2008) and Pope (2008) where the arrival of an Arab neighbor post 9/11 impacts property prices. However, given the short-term nature of the price effects and the small number of Arab homeowners in the data, such an analysis is not possible.[4]

Of course, a change in preferences towards Arab neighbors does not necessarilly mean that current homeowners themselves form unfavorable views of Arab neighbors post 9/11. Rather, the value of the property is also based on expected preferences of potential buyers. Given that less than half of the respondents in HAR viewed Arabs as favorable, it is not implausible that shortly after 9/11 transaction prices would reflect the probability that a non-negligible number of potential buyers would view Arab neighbors unfavorably.[5] A limitation of this study is that we cannot distinguish between the current homeowner and potential buyer channels. In what follows, we remain agnostic as to the true source of the change in preferences.

Unfortunately, our data set is comparable to many transaction level data sets that do not indicate the ethnicity of the buyer or seller. However, our data set does include the full name (first and last name) of both the buyer and seller. As such, we identify ethnicity based on an individual's full name. Related approaches have been used extensively in the biostatistics literature and are known as *name-ethnicity matching*. Gautier et al. (2009) use an informal but pragmatic approach where

of the expectations channel at work in high-rise office buildings. Similarly, homeowners living in high-rise residential buildings might might form similar expectations. Although Seattle is a major metropolitan area and the assessor's data includes condominium transactions in multi-floor buildings, there are too few units and transactions following 9/11 than can be used for this analysis. Furthermore, unlike Abadie and Dermisi (2008), it is not clear what are the landmark buildings in downtown Seattle. In short, high-rise residential buildings are not the focus in this paper but a possible avenue for future research.

[4]The small number of Arabs purchasing home in the 180 days after 9/11 in the assessor data is the primary reason for this insignificance.

[5]This can be done using a search model similar to Krainer (2001) where 1) current homeowners receive non-negative utility from living near a nearby Arab neighbor (Type I), 2) some potential buyers are also Type I, 3) the remaining potential buyers receive disutility from an Arab neighbor (Type II), and 4) the transaction price is decreasing in the probability that potential buyers are Type II.

research assistants from Turkey and Morocco manually identify *Turkish* and *Moroccan* names in 20,148 transactions. Given our data set includes 302,065 transactions and we do not have any Arab research assistants, manually identifying Arab names in the assessor data is neither practical not feasible. Based on anecdotal evidence, similar situations are common to many economic researchers.

In order to identify Arab homeowners, we use a supervised learning algorithm common in the machine learning literature. The idea is to use a labeled data set in order to predict labeling on a different, unlabeled data set. For the application at hand, we use a binary labeling where countries are either members of the Arab League or not. We then apply this labeling to the set of Summer Olympic rosters for all countries 1948-2012. Next, we estimate a binomial classifier (a regularized or penalized logit) using indicator variables for names as the explanatory variables. Finally, we use the estimated binomial classifier to label buyer and seller names in the assessor data. In short, names in the assessor data are labeled based on the conditional likelihood that a given name would be found on the Olympic roster of an Arab League country.

The supervised learning algorithm we describe provides two advantages over manual classification. First, unlike manual classification, our procedure can be scaled to large data sets. In real estate settings, large data sets are commonplace. Second, our procedure is quite flexible and can be used to classify individuals by name into any number of groups. As a demonstration of this flexibility, we perform a falsification test using Asian and Hispanic homeowners and, in contrast to Arab homeowners, find no significant price effects post 9/11 for either of these two groups.

We focus on preferences for Arab neighbors as a matter of practicality despite results in the HAR indicating borh Arabs and Muslims were viewed as unfavorable. Because the Arab World is not the same as the Islamic World, and vice-versa, we compare alternative identification schemes based on religious and geographic considerations. We also examine countries with large Muslim populations as well as the countries in the recently proposed US Census racial category: *Middle East and North African* (MENA). Including non-Arab countries with a large Muslim population intensifies the price effect; using MENA countries mitigates the price effect as by definition MENA is a purely geographic definition that includes Israel but excludes Sudan. As a whole, we interpret these findings as evidence that homeowner preferences also respond to non-Arab Muslim neighbors.

Of course, automated classification schemes are neither perfect nor perfect substitutes for hu-

man classifiers. Regardless, in the spirit of Gautier et al. (2009), an independent undergraduate economics student confirmed the probable ethnicity in a small sample of buyers classified as Arab in the assessor data. As a whole, the machine learning approach developed here is both a practical and flexible way to identify unobserved ethnicity using observed names in large data sets.

2 Literature Review

2.1 Terrorism

This study is part of a growing theoretical and empirical literature investigating the relationship between terrorism and economic variables. Theoretical results focus on how terrorism is different from other risks. Lakdawalla and Zanjani (2005) demonstrate how terrorism insurance differs from catastrophe insurance. Becker et al. (2004) examine terrorism in a behavioral model that includes both fear and risk aversion. Glaeser and Shapiro (2002) compare the impacts of war and terrorism on the formation and dissolution of urban centers and find war has had a larger impact on urban formation than terrorism.

Despite the low probability of a terrorist event, there does appear to be significant empirical evidence that terrorism has a large impact on macroeconomic variables. In one of the earliest studies, Abadie and Gardeazabal (2003) find terrorism in the Basque region of Spain decreased regional per capita GDP by a non-trivial 10%. Examples of other economic variables affected by terrorism include stock markets Zussman and Zussman (2006) and Abadie and Gardeazabal (2003), foreign direct investment Abadie and Gardeazabal (2008), consumption Eckstein and Tsiddon (2004), industrial organization Berrebi and Klor (2010), birth weight Camacho (2008), vacancy rates Abadie and Dermisi (2008), and house prices Besley and Mueller (2012) and Elster et al. (2017). Of course, these economic costs are in addition to any non-economic costs associated with a decrease in quality of life, Frey et al. (2007).

Implicit in these studies are the beliefs that economic agents have about the probability of future terrorism. These beliefs are updated after the realization of relevant events. For instance, Abadie and Gardeazabal (2003) find violence that the stock market reacts positively when a credible truce is reached. Zussman and Zussman (2006) and Zussman et al. (2008) find that the Israeli stock market reacts in response to key events in the Israeli-Palestinian conflict. Using a regime-switching model,

Besley and Mueller (2012) find there is a difference between short-term violence and long-term perceptions of persistent violence in Northern Ireland; only when homeowners perceive that the world is in a non-violent state do regional house price indexes trend upwards. Abadie and Dermisi (2008) find vacancy rates in notable Chicago office buildings increased in the months following 9/11.

2.2 Preferences and the Real Estate Market

This study is also related to other studies that find preferences for neighbors are correlated with socioeconomic variables. Using US Census microdata, Bayer et al. (2004) find race is a fundamental cause of sorting in the San Francisco Bay area. Other studies suggest preferences of natives for non-immigrant neighbors can lead to segregation, Cutler et al. (2008) and Benabou (1993). Saiz and Wachter (2011) find house prices metropolitan areas with more immigration grow faster, but house prices neighborhoods with a higher concentration of immigrants grow more slowly relative to the metropolitan area at large. Saiz and Wachter (2011) suggest one interpretation is that natives prefer native neighbors. That being said, Bayer et al. (2007) find price effects attributable to race can be misleading when neighborhood quality is not observed. As a whole, these studies find evidence that homeowners prefer neighbors with common social, racial, linguistic, and ethnic identities. We build on these studies and ask if preferences for ethnicity are static, or if preferences respond to acts of terrorism.

Several studies have used correspondence experiments to test for discrimination against Arabs and Muslims in the real estate market. Ahmed and Hammarstedt (2008) find landlords in Sweden are less willing to respond to rental applications when the applicant's name is strongly associated with Islam or Arab ethnicity. Using applicant names *Erik Johansson, Maria Andersson* and *Mohammed Rashid*, Ahmed and Hammarstedt (2008) find emails sent using the name Mohammed Rashid received fewer callbacks and invitations to view the property. In a similar experiment in Spain, Bosch et al. (2010) find the low response rate still exists after controlling for socioeconomic factors. Using hand-written applications, Carpusor and Loges (2006) find similar evidence of discrimination in the American rental market. Overall, these findings are not unlike the experiences of other minorities in rental markets, Hanson and Hawley (2011). By construction, in all of these correspondence experiments, the ethnicity of the applicant is signaled by name and name alone.

In addition to cross-sectional studies, event studies have also been used to test for changing

attitudes towards Arabs and Muslims at the neighborhood level. Gautier et al. (2009) investigate property prices following the murder of Dutch filmmaker Theo van Gogh by a recent convert to radical Islam. Following van Gogh's murder, listing prices in nearby Muslim neighborhoods decreased by 3%. Further, Gautier et al. (2009) find evidence of increasing segregation after the murder as Muslims became more likely to purchase homes in areas with an already large concentration of Muslims. Similar results at the neighborhood level are found following London subway bombings, Ratcliffe and von Hinke Kessler Scholder (2015). The results in (Gautier et al., 2009) and Ratcliffe and von Hinke Kessler Scholder (2015) are best interpreted as the effect of terrorism on home prices in nearby Muslim neighborhoods. In contrast, we are interested in local impacts at the property level in neighborhoods that do not necessarily have a clustering of Arab homeowners. Furthermore, Seattle is more than 3,000 miles from New York and Washington DC and was not a target of any terrorism on 9/11.

2.3 Name-Ethnicity Matching

In order to estimate our binomial classifier, we use a hierarchical data generating process common in the text modeling literature. In these models, words in a body of text are assumed to be drawn from a multinomial distribution where the probability of each word is drawn from one or more topics, Hofmann (1999), Blei et al. (2003). For example, in the economics literature, texts with *urban* topics are more likely to include the words *residential*, and *rental*, whereas texts with *international* topics are more likely to include *tariff* and *trade*. In this study, the first and last name of each Olympian (words) are viewed as realizations from a multinomial distribution where the probability of each name is a function of the Olympian's ethnicity (topic).

Automated approaches for matching name to race or ethnicity, name-ethnicity matching, have been well studied in the biostatistics literature. Examples include Coldman et al. (1988), Burchard et al. (2003), Fiscella and Fremont (2006). Moreover, examples where the researcher generates names from a given ethnicity are common in correspondence experiments in economics, Bertrand and Mullainathan (2004), Ahmed and Hammarstedt (2008), Bosch et al. (2010), Hanson and Hawley (2011), Hanson et al. (2016). Applications where real names are used to infer ethnicity include Humphreys et al. (2016) and Gautier et al. (2009).

A standard econometric approach to creating a binomial classifier based on observables is to use

a logit or similar binary choice model. In the name-ethnicity approach, the number of observables (names) can be large. In this high-dimensional covariate setting, maximum likelihood estimation will overfit the data and out-of-sample predictions are at best misleading, Hastie et al. (2015). Out-of-sample performance is of critical importance to us as our primary goal is predicting ethnicity in the assessor data. Noting this, we use an ℓ_1 regularized estimator. This regularization yields an estimator that has been shown to have superior out-of-sample performance in high-dimensional logistic models relative to both un-regularized estimators and ℓ_2 regularized estimators, Ng (2004). Furthermore, unlike manual classification, the classifier can be scaled to large data sets.

We estimate the binomial classifier by first labeling the Olympians as either representing an Arab League country or not. By using labeled data, we use a supervised learning algorithm. An alternative when the research has no ex-ante knowledge of groups is to identify latent ethnic groups by applying an unsupervised learning algorithm on an unlabeled set of names. One of the more popular unsupervised learning algorithms in textual analysis is the Latent Dirichlet Allocation (LDA), Blei et al. (2003). Similar to principal components, it is necessary for the researcher to label or identify the ex-post estimated LDA topics. In unreported results, we find that the LDA resulted in a group that could most likely be identified as an Arab League group. In any event, because we are interested in classifying based on ex-ante, specified groups, we leverage this specificity and use the supervised learning algorithm described below.

3 Identifying Ethnicity and the Binomial Classifier

In what follows we use the following definitions: a *full name* is a first name and last name pair, and a *name* is either a first name or a last name. We treat each full name as an exchangeable set of names: one or more first and last names.[6] For example, American Olympian *carl lewis* is expressed as the 2-element set $\{carl, lewis\}$. In our analysis, we treat hyphenated names as two names and retain the hyphen in order to acknowledge the split. For example, French Olympian *jean-claude killy* becomes the 3 element set $\{jean\text{-}, claude, killy\}$, and Syrian Olympian *nasser al-shami* is the 3-element set $\{nasser, al\text{-}, shami\}$. We also remove diacritics from the full names in a practical manner as *josé* becomes *jose*. In the Olympic rosters, we identify 69,648 unique names. In order to

[6] Viewing the names in this was is common in the textual analysis literature and is known as a *tokenization* approach where each name is a *token*. Examples of the tokenization approach to text data include Gentzkow and Shapiro (2010), Taddy (2013), and Nowak and Smith (2017).

focus on the more frequent names, we drop all names that occur fewer than 10 times in the rosters. Doing so leaves $P = 3,212$ unique names.[7]

As mentioned above, we use indicator variables for names as the explanatory variables in a logit model. That being said, it is instructive to view each full name as a $P \times 1$ vector, x_n, of 1s and 0s where the p element is equal to 1 if name p is in the full name of Olympian n and 0 otherwise. For *carl lewis*, x_n has 2 elements with a 1, and the remaining elements are 0: the pth element of x_n, x_{np}, corresponding to *carl* will be equal to 1 and likewise for *lewis*. Because we view first and last names as exchangeable variables, our estimator does not distinguish *carl lewis* from *lewis carl*. However, such distinctions are unlikely to have a material impact on our estimator. Moreover, the exchangeability assumption is also practical as first and last names are not explicitly identified in the assessor data.

Next, we set the indicator variable $y_n = 1$ if an Olympian comes from an Arab country and 0 otherwise. In order to identify Arab countries, we use the list of countries in the Arab League in addition to other lists described in more detail. Using y_n and x_n, one could estimate a logit model where the explanatory variables are indicator variables for each of the P names. The probability that $y_n = 1$ is then given by

$$\Pr(y_n = 1|X_n, \phi) = \frac{e^{\phi_0 + \sum_p X_{np}\phi_p}}{1 + e^{\phi_0 + \sum_p X_{np}\phi_p}} \tag{1}$$

In Equation (1), when $0 < \phi_p$, the presence of name p increases the likelihood that Olympian n represents a country that is a member of the Arab League, and vice-versa for $\phi < 0$. When $\phi_p = 0$, name p does not help to predict y_n. The parameter ϕ_0 controls the unconditional $\Pr(y_n = 1)$.

For fixed P, ϕ_p can be consistently estimated by maximum likelihood. In our application, the assumption of fixed P is difficult to defend as P is large by conventional standards even after filtering out the least common names in the Olympic rosters. As such, a maximum likelihood estimation of ϕ is at worst infeasible when the data is seperable or at best prone to overfit the data, Hastie et al. (2015).[8] Of course, we could increase the cutoff, exclude more names, and reduce the number of explanatory variables. However, such an approach is imprudent as names that are removed might be relevant for classification.

[7] In unreported results, we found that a cutoff of 5 produced similar results to a cutoff of 10.

[8] In this setting, the data set is separable if any name is only present a single group, i.e. *jose* is only found in the non-Arab League rosters. Hastie et al. (2015) provides further details on separable data sets.

Therefore, we retain a large set of names - large P - and use a regularized likelihood estimator. Specifically, we place an ℓ_1 penalty on the individual ϕ_p parameters. We then choose ϕ to minimize the following

$$-\sum_n \Pr(y_n = 1|X_n, \phi)^{y_n} \left[1 - \Pr(y_n = 1|X_n, \phi)\right]^{1-y_n} + \lambda \sum_p |\phi_p| \tag{2}$$

Define the solution to Equation 2 as ϕ^*. The first term in Equation 2 is the negative likelihood of the sample using the individual likelihood given in Equation 1. The second term is an ℓ_1 penalty term that regularized the magnitudes of the elements in ϕ. The parameter λ is a tuning parameter.[9][10]

When $\lambda = 0$, there is no penalty on ϕ, and ϕ^* is the maximum likelihood estimator for the logit. When $0 < \lambda$, ϕ^* is well-defined even for separable data sets. Because of the shape of the ℓ_1 penalty, some entries of ϕ^* can be set equal to 0 when $0 < \lambda$. As mentioned above, when $\phi_p^* = 0$, name p cannot be used to classify y_n. With this interpretation, minimizing Equation 2 simultaneously performs variable selection and coefficient estimation. In any event, by including the penalty term $\lambda \sum_p |\phi_p|$, ϕ^* can be used to classify names out-of-sample as, unlike the maximum likelihood estimator, a regularized estimator is less likely to overfit the data in-sample, Ng (2004). Furthermore, this functional form for the penalty term has been shown to yield superior out-of-sample prediction compared to an ℓ_2 penalty, Ng (2004). We emphasize, out-of-sample performance is fundamental to our results.

4 Model Specification

In order to test for changes in the local valuation of Arab neighbors, we use a difference-in-differences approach and compare a control and treatment group before and after 9/11. The two groups are defined using distance to the nearest Arab homeowner. As such, properties located 0-0.1mi (0.1-0.3mi) from an Arab neighbor are in the treatment (control) group; these cutoffs are validated, below. Properties further than 0.3mi from an Arab homeowner are neither in the treatment nor control group. To the extent that the events of 9/11 were unexpected, we interpret the results as causal: if not for 9/11, there would be no local price effects attributable to Arab neighbors. Of

[9] In our analysis, we select λ using 5 fold cross-validation. The results are robust to λ near the cross-validated choice of λ

[10] We use the `glmnet` package in `R` to solve Equation 2. The solution is found by using a quadratic approximation to the true regularized likelihood.

course, given the singular nature of these events, it is not fair to generalize the results to generic terrorist events.

In order to identify relevant distance thresholds, we first estimate a base hedonic model for property n in census tract c sold in quarter t

$$p_{nct} = x_{nct}\beta + \delta_{ct} + u_{nct} \tag{3}$$

Here, p_{nct} is the log price of the house, x_{nct} is a vector of house attributes including log square footage, bedrooms, bathrooms, and age of the property, β is a vector of implicit prices, δ_{ct} is a census tract specific price at time t, and u_{nct} is an error term. We begin each quarter on the 11th of March, June, September, and December. In doing so, the immediate effects of 9/11 are captured by only a single δ_{ct} for a given c.

We determine the distance cutoff for the treatment group using a method similar to Linden and Rockoff (2008). After estimating Equation 3, we collect the residuals for all transactions located within 0.3mi of an Arab homeowner. We then estimate a local polynomial regression in order to estimate the price gradient as a function of distance from the nearest Arab neighbor.[11] The 95% confidence interval for the local polynomial using transactions 180 days before 9/11 is presented in green in Figure 1. Point estimates for the same local polynomial regression 180 days after 9/11 are presented as a red line in Figure 1.[12] If the events of 9/11 caused local preferences for Arab neighbors to change, the changes are extremely local as the point estimates are within the pre-9/11 95% confidence interval beyond 0.1mi. In any event, Figure 1 suggests a cutoff of 0.1mi for the treatment group.

We create the indicator variable $D^{0.1}_{nct} = 1$ ($D^{0.3}_{nct} = 1$) if any Arab neighbor is currently living within 0.1mi (0.3mi) of the property at time t and 0 otherwise. As further described below, we use buyer and seller names for all transactions between the years 1982-2016 in order to identify the current homeowner. In the absence of any time-varying treatment effects, cross-sectional differences between the control and treatment groups can be estimated using

$$p_{nct} = x_{nct}\beta + D^{0.1}_{nct}\psi^{0.1} + D^{0.3}_{nct}\psi^{0.3} + \delta_{ct} + u_{nct} \tag{4}$$

[11] Similar to Linden and Rockoff (2008), we estimate a local polynomial of order 3 and use a bandwidth of 0.1 nearest-neighbors using the `locfit` package in R.

[12] A similar plot is produced when using all sales before 9/11 and sales 365 days before 9/11.

In Equation 4, $\psi^{0.3}$ captures local cross-sectional price effects common to all properties within 0.3 miles of an Arab homeowner. Cross-sectional differences include both arbitrary preferences for race and ethnicity as well as local amenities and neighborhood quality. The coefficient $\psi^{0.1}$ captures cross-sectional differences in price between the control and treatment groups. When $\psi^{0.1} > 0$ ($\psi^{0.1} < 0$), properties in the treatment group sell at a premium (discount) to properties in the control group. Figure 2 displays the treatment and control group for a randomly selected house.

In the absence of time-varying effects, boundary effects, or detailed data on race and ethnicity, it is difficult if not impossible to disentangle race and ethnicity preferences from local amenity effects, Bayer et al. (2007). We follow Pope (2008) and Linden and Rockoff (2008) and test for time-varying price effects attributable to Arab neighbors by comparing prices between the control and treatment groups before and after 9/11. As mentioned above, we are interested in investigating whether or not 9/11 caused perceptions of Arab neighbors to change in a negative way. In order to test this, we create the indicator variable $Post_{nct} = 1$ if the transaction occurred after 9/11 but before March 10, 2002, a time period of 180 days. We later investigate 180-365 days following and 0-180 days before 9/11. For all of these event windows, we estimate

$$p_{nct} = x_{nct}\beta + D^{0.1}_{nct}\psi^{0.1} + D^{0.3}_{nct}\psi^{0.3} + Post_{nct} \times D^{0.1}_{nct}\tau^{0.1} + Post_{nct} \times D^{0.3}_{nct}\tau^{0.3} + \delta_t + \mu_c + u_{nct} \quad (5)$$

Here, $\tau^{0.1}$ and $\tau^{0.3}$ capture time-varying price effects relative to the location of Arab neighbors following 9/11. If $\tau^{0.3} < 0$, properties within 0.3 miles of Arab neighbors experienced price declines relative to the rest of the market. If the events of 9/11 caused homeowners' preferences towards Arab neighbors to change, price effects should be stronger the closer the property is to an Arab neighbor, $\tau^{0.1} < 0$.

5 Data

5.1 Name Ethnicity Data

Data used in this paper comes from two sources. The first source is the set of Summer Olympic rosters from 1948 to 2012. These rosters were downloaded from the Olympic Reference website.[13]

[13] http://www.sports-reference.com/olympics/

Each roster includes the full name of the Olympian, age, gender and nationality. As mentioned above, we label countries as Arab if the country is a member of the Arab League. Six member states formed the Arab League in 1945: Jordan, Syria, Saudi Arabia, Lebanon, Egypt, Iraq. However, the Arab League has expanded over the years and now includes 22 member states that cover Northern Africa and the Middle East. We use the 22 member roster of Arab League nations. A complete list of these countries and other country lists is presented in the sppendix. Table 1 shows the 20 most frequent names from the Arab and non Arab League countries. Figure 3 presents the names in a word cloud where the size of the font is associated with a greater frequency of the name within group.

We also investigate alternative lists of countries in order to better interpret and validate our results. First, we include countries with a large number and percentage of Muslim citizens (Bangladesh, Indonesia, Iran, Pakistan, and Turkey) alongside Arab League countries. Second, we include countries in the proposed MENA racial category from US Census. The MENA country list notably includes Israel but excludes Arab League members Somalia and Sudan. Therefore, using the MENA country list places more of an emphasis on geography than ethnicity. It should be pointed out that the US Census also considers an ethnic-based definition of MENA using self-reported ethnic identifiers including Arab, Berber, Kurdish, Middle Eastern, and others.

Other publicly available data sets have been used by resarchers to identify ethnicity including Wikipedia Treeratpituk and Giles (2012) and IMDB Rachevsky and Pu (2011).[14] We use the Summer Olympic data set in lieu of these other data sets as the Summer Olympic data set 1) provides a sufficient number of observations for the training set and 2) the number of individuals from Arab League countries is significantly larger in the Summer Olympic data than in the Wikipedia or IMDB data. In total, there are $N = 90,636$ unique Olympians from 221 unique countries.

5.2 Pricing Data

Transaction data for single-family homes comes from the King County Assessor's Office that includes the Seattle metro area.[15] The data set is publicly available and includes information on property attributes, buyer and seller names, transaction price, and other relevant information. We

[14]The authors are not aware of any publicly available databases provided by Wikipedia. Databases on actors, directors, etc. provided by IMDB and are available at http://www.imdb.com/interfaces

[15]http://www.kingcounty.gov/depts/assessor.aspx

filter out outlying observations using reasonable filters described in the appendix. In order to exclude any effects attributable to the volatility of housing prices in the mid 2000s, we limit our data set to transactions between January 1, 1982 and December 31, 2002. This leaves 302,065 total transactions in the study.

Using December 31, 2002 as a cutoff provides us with more than 15 months of sales post 9/11 that we can use to identify time-varying price effects. Summary statistics for the data are provided in Table 2. The average transaction price is $217,418. The average house has 1,960 square feet, 3.3 bedrooms, 1.5 bathrooms, and was built in 1965. Although using data back to 1982 might seem excessive, our results are not sensitive to this starting date. In the appendix, we show that our results do not change substantially when using the sub-periods 1990-2002 or 2000-2002.

It is important to note that we first identify the locations of current homeowners at any point in time using the entire set of buyer and seller names for all transactions between the years 1982-2016. For example, an Arab homeowner who buys in 1990 and sells in 1995 would be the current owner between these years. As a result, all sales between 1990-1995 for any properties located within 0.1mi of this Arab homeowner would have $D^{0.1} = 1$.

In addition to using buyer names, we also use seller names to identify Arab homeowners. For instance, an Arab homeowner who buys a home in 1980 and sells in 2010 would never appear as a buyer in the set of transactions 1982-2016. However, this individual would appear as a seller in 2010. Absent this considerations, sub-period analysis using the years 2000-2002 would only include the locations of Arab homeowners who purchased a property between January 1, 2000 and December 31, 2002. In short, we identify the locations of all Arab homeowners regardless of whether or not the transaction occurred during the relevant sub-period. Because our data covers more than 30 years of transaction data, we consider any bias attributable to unobserved Arab homeowners negligible.

Using the binomial classifier, we identify the locations of 494 Arab homeowners between 1982 and 2002. This implies 0.17% of homeowners are of Arab ancestry and is comparable to the 2000 US Census where 0.4% of the King County population has Arab ancestry.[16] A sample of the locations of homeowners identified as Arab on 9/11 are displayed in Figure 4. Unlike Gautier et al. (2009), the locations of Arab homeowners are fairly dispersed throughout Seattle. Although the data is publicly available, for privacy purposes, we do not disclose the actual names of the individuals who

[16] https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?src=CF

are identified as Arab. A list of the homeowners identified as Arab is available from the authors upon request. Table 3 displays the total number of observations in the control and treatment groups before and after 9/11. There are 7,320 transactions in the treatment group with 445 transactions occurring between 0 and 180 days after 9/11. Likewise, there are 35,768 transactions in the control group with 2,488 transactions between 0 and 180 days after 9/11.

Table 4 compares control variables in the control group to control variables in treatment group. Mean comparison tests indicate that the treatment group differs from the control group in a statistically significant but economically insignificant manner. Figure 5 displays the distribution of the control variables for the 2 groups. Similar to the evidence in Table 4, the distributions for the control and treatment groups are similar, although a greater portion of properties in the treatment group are built after 1985.

6 Results

6.1 Arab Name Identifiers

Results for the regularized logit model are presented in Table 5. As expected, the results in Table 5 indicate that relative frequency in Table 1 is a strong indicator of nationality. Names that are strong indicators of being from a country in the Arab League or not are as expected. The strongest predictor of non-Arab status is *jose*.

It is interesting to note that the most common Arabic name, *mohamed*, and its variants are not strong predictors of an individual being from an Arab League country. This should not be surprising, as *mohamed* is a both an Islamic and Arab name found throughout many non-Arab countries. Results in Table 5 do not imply that *mohamed* should not be used to signify an Arab or Mulsim applicant as in Ahmed and Hammarstedt (2008) or Bosch et al. (2010). Rather, Table 5 indicates which names are the strongest predictors but does not display which names would have the most influence on landlords.

We also find that predictions from our method would be comparable to other classifiers used in the literature. Bertrand and Mullainathan (2004) identify distinctively black names using the relative frequency of names between racial groups (black or white); distinctively black names are names with the largest ratio of relative frequencies. Figure 6 plots ϕ^* against the log of the ratio

of relative frequencies for Arab and non-Arab countries for all names that occur at least once in both sets. The positive relationship between ϕ^* and the log ratio confirms the notion that a ratio-based method similar to Bertrand and Mullainathan (2004) would yield similar classifications to the binomial classifier we describe above.

Using ϕ^* and Equation 1, we can calculate $\Pr(y_n = 1|X_n, \phi^*)$ for buyer and seller names in the assessor data. Unlike Gautier et al. (2009), we do not have access to an Arab research assistant. However, a manual inspection of the names by an Arab undergraduate economics student confirmed probable ethnicity for names above this cutoff. Based on manual inspection, we create the indicator variable $Arab = 1$ for Arab homeowners if $0.35 < \Pr(y_n = 1|X_n, \phi^*)$ and $Arab = 0$ otherwise. In the Appendix, we present similar results when using a cutoff of 0.5.

6.2 Common Trends

Figure 7 presents the quarterly price index for the control, treatment, and all sales 12 quarters before 9/11 and 4 quarters after 9/11. The price index uses the convention that quarters being on the 11th of March, June September, and December, and uses $t = 0$ for the period June 12, 2001 - September 10, 2001; $t = 1$ corresponds to the period September 11, 2001 - December 10, 2001.

Figure 7 presents graphical evidence that property prices in the treatment and control groups followed a common trend. In unreported results, we can not reject the null hypothesis that the price indexes for the control and treatment groups follow a different linear time trend for $t \leq 0$. In $t = 1$, there appears to be a notable decrease in the treatment group price index relative to the control group price index. For $2 \leq t$, the price indexes for the control and treatment groups appear to return to their pre 9/11 relationship.

6.3 Price Effects

Table 6 presents our first set of results. All standard errors are clustered two ways at the quarter and census tract levels. The first column estimates Equation 5 using additively separable census tract and quarter fixed effects. Additively separable fixed effects preclude heterogeneous price trends across census tracts. In the cross-section, properties in the treatment group sell at a 1.1% discount relative to properties in the control group. In the 180 days after 9/11, properties in the treatment group sell at a 1.8% discount relative to the control group. These results indicate a short-term,

local, negative price effect attributable to Arab neighbors caused by the events of 9/11 possibly caused by a change in underlying preferences.

Results in Column 2 include interaction terms between the census tract and quarter fixed effects and control for heterogeneous price trends across census tracts. As expected, cross sectional effects attributable to Arab neighbors do not change much but the time-varying effects do change. Unlike the results in Column 1, after controlling for heterogeneous census tract price trends, there is no significant difference in price between properties in the control and treatment groups pre 9/11. However, similar to results in Column 1, there is significant statistical evidence of local price effects related to Arab neighbors post 9/11. Relative to the control group, properties in the treatment group sold at a 1.4% discount. The result is significant at the 1% level.

In order to estimate a price effect absent any comparison to properties in the control group but instead relative to all properties in the market, Column 3 uses $D^{03}(1 - D^{0.1})$ instead of $D^{0.3}$ in Equation 5. In contrast to the other columns in the table, the coefficient on $D^{0.1} \times Post$ is now interpreted as the price effect relative to all properties in the market. The coefficient for $D^{0.1}$ indicates properties in the treatment groups decreases 2.7% relative to the rest of the market in the 180 days after 9/11.

Column 4 tests for longer-term price effects by defining $Post = 1$ if the transaction is 180-365 days after 9/11 and $Post = 0$ otherwise. In contrast to the results in Column 2, there is no significant price difference in the control and treatment groups. Comparing the results in Columns 2 and 4, if 9/11 changed preferences for Arab neighbors, these changes were short-term. We further investigate the time decay, below.

In our discussion of common trends, we did not reject the null hypothesis that the control and treatment groups have a common linear price trend. However, it is possible that the results in Columns 1-3 reflect short-term momentum from unobserved price trends specific to properties in the treatment group. Alternatively, price effects in Columns 1-3 might not be caused by 9/11 but reflect existing, short-term price trends immediately before 9/11. In order to rule out this possibility, we use $Post = 1$ if the transaction is 0-180 days before 9/11 and $Post = 0$ otherwise.[17] Doing so, we test for abnormal price effects between March 12, 2001 and September 10, 2001. Results indicate no evidence that the results in Columns 1-3 are driven by any short-term momentum in the market.

[17] We avoid using the term $Pre = 1$ in order to keep the tables concise.

More importantly, there is no significant difference between the treatment and control groups 180 days before 9/11. In conclusion, the results in Column 5 rule out any momentum effects present in the treatment group.

6.4 Alternative Model Specifications

Table 7 presents results for alternative model specifications. For reference, we reproduce the results in Column 2 of Table 6 in Column 1 of Table 7. Column 2 of Table 7 eliminates all sales greater than 0.5 miles from an Arab homeowner. We do this in order to remove any possible systematic difference between properties near Arab homeowners and properties much further away. The results in Column 2 for the 89,049 transactions less than 0.5 miles from an Arab homeowner are comparable to results using all 302,065 transactions.

Previous results assume that the control variables enter linearly into the regression function. Column 3 provides a robustness check of this assumption and allows for a flexible relationship between sale price and square footage, bedrooms, and bathrooms. In Column 3, we sort properties into bins based on square footage where the bins are 0-500sqft, 500-1,000sqft, etc. up to 4,500-5,000sqft. We then include indicator variables for the bins as explanatory variables in lieu of a continuous measure of log square footage. We also include indicator variables for bedrooms and bathrooms. Results in Column 3 are comparable to the results when using a linear specification for the control variables.

Column 4 investigates if the price effect attributable to Arab neighbors post 9/11 is more intense at smaller distances within the treatment group. In order to do so, we create the variable $\frac{0.1-Distance}{0.1}$ where $Distance$ is the distance to the closest Arab homeowner in miles. The product $D^{0.1} \times \frac{0.1-Distance}{0.1}$ is constrained to be between 0 and 1 and assumes a linear relationship between proximity to Arab neighbors and price. $D^{0.1} \times \frac{0.1-Distance}{0.1} = 1$ when $Distance = 0$ and $D^{0.1} \times \frac{0.1-Distance}{0.1} = 0$ when $0.1 \leq Distance$. Results in Column 4 of Table 7 indicate that none of the coefficients pre 9/11 are statistically different from 0 for the control and treatment groups.

The coefficient on $D^{0.1} \times \frac{0.1-Distance}{0.1} \times Post$ is statistically different from 0 and equal to -0.047; thus, estimated price effects attributable to Arab neighbors in the treatment group appear to be entirely captured by distance. This is not too surprising given Figure 1. Taken in conjunction with the null result for the coefficient on $D^{0.1} \times Post$, properties located immediately next to an

Arab neighbor, $Distance = 0$, sold at a 4.7% discount relative to the control group. A property 0.05mi away (approximately 264ft) is expected to sell at a 2.35% price discount. Of course, it is not possible for single-family, detached properties to be immediately next to each other. Regardless, Column 4 provides evidence that the price effects for Arab neighbors post 9/11 are decreasing in distance and supports the notion that non-Arab homeowners are reacting to Arab neighbors.

Price trends in Figure 7 presents evidence that the price effect is transitory. Column 5 investigates the speed at which prices return to pre 9/11 relationships. In order to do so, we include the number of days singe 9/11 divided by 180. This ratio is equal to 0 on 9/11 and increases to a value of 1 180 days after 9/11. Results in Column 5 indicate that post 9/11, the treatment group saw a price decline of 2.1% relative to the control group. However, this price effect decays over time. 180 days after 9/11, price effects in the treatment group have increased by 1.4% compared to their initial 2.1% levels. Alternatively, approximately 2/3 of the initial price effects have disappeared 180 days after 9/11. Of course, these time effects do not take into account any lags associated with listing the property. Absent any listing data from Multiple Listing Services or similar databases, it is not possible to determine the exact date a property was listed. In any event, the results in Column 5 present suggestive evidence that price effects attributable to Arab neighbors attenuated over time.

Because the events of 9/11 occur at the beginning of Q4, it is possible that the results above are possible are merely capturing seasonal effects. In order to rule this out, Column 6 of Table 7 includes seasonal fixed effects for the quarter of sale. The results are comparable to Columns 1-3. Thus, it does not appear that our results are driven by a seasonal effect.

6.5 Alternative Ethnic Groups

We can use the binomial classifier and the Olympic data to classify individuals based on any list of countries. In order to demonstrate this flexibility, and perform an interesting counterfactual, we identify homeowners from East Asian and Hispanic countries and perform the same analysis. For the East Asian list we use: China, Japan, Mongolia, South Korea, and North Korea. For the Hispanic group, we use all countries in Central and South America. Preliminary results for these two groups is presented in Figure 8. As expected, and unlike Figure 1, there does not appear to be any local price effects for either East Asian or Hispanic neighbors.

Formal results are presented in Table 8 for East Asian and Hispanic neighbors; complete results for both East Asian and Hispanic neighbors are presented in the appendix. In contrast to the price effects for Arab neighbors reported in Column 1, Columns 2 and 3 indicate that there does not appear to be any price effects associated with either Eat Asian or Hispanic neighbor types in the 180 days post 9/11. Alternatively, changing preferences post 9/11 were very acute and limited to Arab neighbors. These results are in line with survey results from the HAR where views of Asian and Hispanic groups were not affected by the events of 9/11.

Results for Arab neighbors rely on using countries in the Arab League as a means to identify Arab homeowners. Of course, this identification scheme is not without its flaws and limitations. First and foremost, we emphasize that we estimate effects with respect to Arab neighbors and not Muslim neighbors. The Islamic World is not homogeneous and its 1.6 billion members come from various denominations, ethnic groups, and more than 200 countries, Miller (2009). However, Islam is the dominant religion in Arab League countries with a majority of Muslims in these countries being either Sunni or Shia. Noting this, we expand our set of countries to include the Arab League countries as well as countries with a significantly large number and percentage of Muslim citizens: Indonesia, Pakistan, Bangladesh, Turkey, and Iran.[18] We exclude India as the number of Indian Muslims is large but constitutes only 13.4% of the total Indian population.

Results when using this expanded list of countries are presented in Column 4 of Table 8. Compared to the results using only the Arab League countries, results when including non-Arab, Muslim countries indicate a much larger price effect. Compared to the control group, the treatment group experienced price declines of 3.2%. This price effect is more than twice than the 1.4% price effect when using the Arab League countries, alone. Of course, this expanded list of countries is by no means a definitive list of Muslim countries in much the same way the resulting list of Olympians is by no means a definitive list of Muslim names. Rather, this expanded list of countries is designed to incorporate additional names that are predominantly Muslim thereby expanding our results to both Arab and non-Arab Muslims. Regardless, results in Column 4 of Table 8 suggest price effects post 9/11 are applicable to many members of the Muslim World and not only Arab Muslims.

Indeed, we acknowledge the concept of ethnicity is by no means restricted to geographic defitions. Previous research by the US Census has used both country of origin and self-reported

[18] $http://www.pewforum.org/files/2009/10/Muslimpopulation.pdf$

ethnicity (Arabic, Kurdish, or Berber) in forming a definition of Arab, Asi and Beaulieu (2013). Presumably for practical purposes, other studies rely on geography as a basis for identification, Gautier et al. (2009), Bosch et al. (2010). We demonstrate the appropriateness of our Arab League choice by comparing our results to a strictly geographic definition: the MENA category proposed by the US Census.[19]

Results using countries in the MENA category are presented in Column 5 of Table 8. The price effect is mildly significant and less than both the the Arab League only results and the Arab League with additional Muslim country results.[20] One possible interpretation of these weaker results is that the notion of geography is explicit in the MENA category. As such, MENA countries include various ethnic and religious groups. Notably, the MENA category includes Israel. For the study at hand, the strict geographic requirements present in the list of MENA countries erroneously includes a number of Jewish individuals in both the treatment and control groups. As such, results in Column 5 are understandable as the HAR indicates Americans did not for unfavorable views of Jewish Americans post 9/11.

6.6 Mosques

Unlike Gautier et al. (2009), Figure 4, indicates that there does not appear to be an area of Seattle where Arab homeowners are concentrated. We present additional evidence using methods in Duranton and Overman (2005) for purchases by Arab homeowners before and after 9/11.[21] Duranton and Overman (2005) use the distribution of pairwise distance as a measure of agglomeration; here, we use the distribution as a measure of clustering by Arab homeowners. Figure 9 presents the distribution of pairwise distances between purchases by Arab homeowners before and after 9/11. As seen in Figure 9, we find no evidence of clustering as the pairwise distances are within the 95% confidence interval in both periods.

Although there is no significant evidence that Arab homeowners live in concentrated areas, it is possible that some Arab homeowners choose to live near mosques or Islamic centers. If that is true, it is possible that estimated price effects reflect a mosques or Islamic centers effect and not an

[19] $https://www.census.gov/content/dam/Census/newsroom/press-kits/2017/2015nct_presentation_jones.pdf$

[20] p-value=0.056

[21] The Duranton and Overman (2005) statistic is calculated as the empirical probability mass function for all pairwise distances, d_{ij}, for all $i, j = 1, ..., I$ pairs of members in a subgroup (here, Arab homeowners) of I individuals. We use a probability mass function as we bin all pairwise distances at 0.25mi for smoothing purposes. The 95% confidence interval is calculated using 2,000 random draws of properties.

Arab neighbor effect. This proposition is not outlandish, as Bogin (2012) find prices near mosques in Baltimore, MD decline 17% after 9/11. In order to rule out the effects of mosques and Islamic centers, we include similar cross-sectional and time-varying indicators in the estimating equation. We identify 23 mosques or Islamic centers in the Assessor data that we believe to be in existence before 9/11.[22]

Table 9 presents the results when including the mosques and Islamic center indicators. The price effects attributable to Arab neighbors in Table 9 are comparable to our previously reported results. In contrast to Bogin (2012), we find no significant mosques or Islamic centers effect when allowing for heterogeneous price trends across census tracts. However, we do find large mosques and Islamic center effects when using additively separable fixed tract and quarter fixed effects.

7 Limitations

Because we identify Arabs based on name, our identification method would identify converts to Islam who adopt a traditional Arab-Islamic name as Arabs, i.e. American Olympian *Cassius Clay* changed his name to *Muhammad Ali* after converting to Islam. Although Pew Research finds that 21% of American-born Muslims are converts, it is unclear what percentage of converts change their name as this is not required by Islam.[23] To the extent homeowners react to Arab neighbors and not Muslim neighbors, this measurement error in the identifier would certainly bias our results towards 0. However, to the extent that homeowners react to Muslim neighbors and not only Arab neighbors, the results we report could be interpreted as a response to the presence of Muslim neighbors.

In addition, we identify Arabs using transaction data and can not identify Arab renters. To the best of our knowledge, the only way to remedy this is to obtain information on the identity of individual renters in the market. In order to mitigate the effect of potential Arab renters in nearby multifamily properties, we present results in the appendix after removing properties less than 0.3mi from a multifamily property that are similar to results in Table 6. Unfortunately, King County does not include the the mailing address of the owner of record. Were this data available going back to 1982, we could do a similar procedure where we removed all properties less than 0.3mi from a non

[22] We were not able to identify any mosque or Islamic center closures post 9/11 in a web search and cannot rule out survivorship bias. The addresses and locations of the mosques and Islamic centers are available from the researchers upon request.

[23] $http://www.pewresearch.org/2007/05/22/muslim-americans-middle-class-and-mostly-mainstream/$

owner-occupied property. If we assume that all non owner-occupied properties are rented out, this would effectively remove all rental properties from the analysis. In any event, we fully acknowledge this limitation as a limitation of all studies that use names in order to identify relevant ethnic or religious types.

8 Conclusion

This study uses transaction data and presents evidence that 9/11 caused a change in homeowner perceptions of Arab neighbors. Specifically, homeowners with Arab neighbors within 0.1mi sold their properties at a 1.4% discount within 180 days of 9/11. A notable limitation of this study and other comparable studies is that we cannot identify the source of the price declines. Specifically, we cannot disentangle homeowner preferences for Arab neighbors from homeowner expectations of the preferences of potential buyers.

In order to identify ethnicity, we use a supervised learning algorithm trained using Summer Olympic rosters from 1948 to 2012. The algorithm can be used to classify buyer and seller names for various groups in large data sets where manual classification is not possible. In future work, we plan to compare our classifier to manual, crowd-sourced classifiers including the Amazon Mechanical Turk.[24] However, given the non-negligible costs associated with Amazon Mechanical Turk, we reserve this for future work.

Our identification scheme is comparable to identification in other studies that use well-chosen geographic identifiers in the absence of explicit ethnic or religious data. Not surprisingly, ethnic groups with no apparent ties to the events of 9/11 do not appear to be the source of any significant price effects. Using Arab League countries alone or alongside additional predominantly Muslim countries yields more intense price effects. We also demonstrate that relying on geography alone can yield misleading classifications that reduce both the estimated price effect and its significance. That being said, we conclude that researchers must make judicious choices when creating relevant country lists.

In any event, results indicate that preferences for specific ethnic groups can be changed by significant events. Although we document a negative change in preferences, we are hopeful that positive changes in preferences for various ethnic groups can be obtained.

[24] https://www.mturk.com/mturk/

References

Abadie, A. and Dermisi, S. (2008). Is terrorism eroding agglomeration economies in central business districts? lessons from the office real estate market in downtown chicago. *Journal of Urban Economics*, 64(2):451–463.

Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, pages 113–132.

Abadie, A. and Gardeazabal, J. (2008). Terrorism and the world economy. *European Economic Review*, 52(1):1–27.

Ahmed, A. M. and Hammarstedt, M. (2008). Discrimination in the rental housing market: A field experiment on the internet. *Journal of Urban Economics*, 64(2):362–372.

Asi, M. and Beaulieu, D. (2013). Arab households in the united states: 2006–2010. *American Community Survey Briefs*, 10(20).

Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of political economy*, 115(4):588–638.

Bayer, P., McMillan, R., and Rueben, K. S. (2004). What drives racial segregation? new evidence using census microdata. *Journal of Urban Economics*, 56(3):514–535.

Becker, G. S., Rubinstein, Y., et al. (2004). Fear and the response to terrorism: an economic analysis. *University of Chicago mimeo*.

Benabou, R. (1993). Workings of a city: location, education, and production. *The Quarterly Journal of Economics*, 108(3):619–652.

Berrebi, C. and Klor, E. F. (2010). The impact of terrorism on the defence industry. 77(307):518–543.

Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *The American Economic Review*, 94(4):991–1013.

Besley, T. and Mueller, H. (2012). Estimating the peace dividend: The impact of violence on house prices in northern ireland. *The American Economic Review*, 102(2):810–833.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.

Bogin, A. N. (2012). *Three Essays on the Property Value Impact of Neighborhood Disamenities*. PhD thesis, Syracuse University.

Bosch, M., Carnero, M. A., and Farre, L. (2010). Information and discrimination in the rental housing market: Evidence from a field experiment. *Regional science and urban Economics*, 40(1):11–19.

Burchard, E. G., Ziv, E., Coyle, N., Gomez, S. L., Tang, H., Karter, A. J., Mountain, J. L., Pérez-Stable, E. J., Sheppard, D., and Risch, N. (2003). The importance of race and ethnic background in biomedical research and clinical practice. *New England Journal of Medicine*, 348(12):1170–1175.

Camacho, A. (2008). Stress and birth weight: evidence from terrorist attacks. *The American Economic Review*, 98(2):511–515.

Carpusor, A. G. and Loges, W. E. (2006). Rental discrimination and ethnicity in names1. *Journal of Applied Social Psychology*, 36(4):934–952.

Coldman, A. J., Braun, T., and Gallagher, R. P. (1988). The classification of ethnic status using name information. *Journal of Epidemiology and Community Health*, 42(4):390–395.

Cutler, D. M., Glaeser, E. L., and Vigdor, J. L. (2008). Is the melting pot still hot? explaining the resurgence of immigrant segregation. *The Review of Economics and Statistics*, 90(3):478–497.

Duranton, G. and Overman, H. G. (2005). Testing for localization using micro-geographic data. *The Review of Economic Studies*, 72(4):1077–1106.

Eckstein, Z. and Tsiddon, D. (2004). Macroeconomic consequences of terror: theory and the case of israel. *Journal of Monetary Economics*, 51(5):971–1002.

Elster, Y., Zussman, A., and Zussman, N. (2017). Rockets: The housing market effects of a credible terrorist threat. *Journal of Urban Economics*, 99:136–147.

Fiscella, K. and Fremont, A. M. (2006). Use of geocoding and surname analysis to estimate race and ethnicity. *Health Services Research*, 41(4p1):1482–1500.

Frey, B. S., Luechinger, S., and Stutzer, A. (2007). Calculating tragedy: Assessing the costs of terrorism. *Journal of Economic Surveys*, 21(1):1–24.

Gautier, P. A., Siegmann, A., and Van Vuuren, A. (2009). Terrorism and attitudes towards minorities: The effect of the theo van gogh murder on house prices in amsterdam. *Journal of Urban Economics*, 65(2):113–126.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71.

Glaeser, E. L. and Shapiro, J. M. (2002). Cities and warfare: The impact of terrorism on urban form. *Journal of Urban Economics*, 51(2):205–224.

Hanson, A. and Hawley, Z. (2011). Do landlords discriminate in the rental housing market? evidence from an internet field experiment in us cities. *Journal of Urban Economics*, 70(2):99–114.

Hanson, A., Hawley, Z., Martin, H., and Liu, B. (2016). Discrimination in mortgage lending: Evidence from a correspondence experiment. *Journal of Urban Economics*, 92(March):48–65.

Hastie, T., Tibshirani, R., and Wainwright, M. (2015). *Statistical learning with sparsity: the lasso and generalizations*. CRC Press.

Hofmann, T. (1999). Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57. ACM.

Humphreys, B., Nowak, A., and Zhou, Y. (2016). Cultural superstitions and residential real estate prices: Transaction-level evidence from the us housing market. West Virginia University Working paper.

Krainer, J. (2001). A theory of liquidity in residential real estate markets. *Journal of urban Economics*, 49(1):32–53.

Lakdawalla, D. and Zanjani, G. (2005). Insurance, self-protection, and the economics of terrorism. *Journal of Public Economics*, 89(9):1891–1905.

Linden, L. and Rockoff, J. E. (2008). Estimates of the impact of crime risk on property values from megan's laws. *The American Economic Review*, 98(3):1103–1127.

Miller, T. (2009). Mapping the global muslim population: a report on the size and distribution of the world's muslim population. *Washington, DC: Pew Research Center*.

Ng, A. Y. (2004). Feature selection, l 1 vs. l 2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78. ACM.

Nowak, A. and Smith, P. (2017). Textual analysis in real estate. *Journal of Applied Econometrics*, 32(4):896–918.

Pope, J. C. (2008). Fear of crime and housing prices: Household reactions to sex offender registries. *Journal of Urban Economics*, 64(3):601–614.

Rachevsky, L. and Pu, K. Q. (2011). Selection of features for surname classification. In *Information Reuse and Integration (IRI), 2011 IEEE International Conference on*, pages 15–20. IEEE.

Ratcliffe, A. and von Hinke Kessler Scholder, S. (2015). The london bombings and racial prejudice: Evidence from the housing and labor market. *Economic Inquiry*, 53(1):276–293.

Saiz, A. and Wachter, S. (2011). Immigration and the neighborhood. *American Economic Journal: Economic Policy*, 3(2):169–188.

Taddy, M. (2013). Multinomial inverse regression for text analysis. *Journal of the American Statistical Association*, 108(503):755–770.

Treeratpituk, P. and Giles, C. L. (2012). Name-ethnicity classification and ethnicity-sensitive name matching. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*.

Zussman, A. and Zussman, N. (2006). Assassinations: Evaluating the effectiveness of an israeli counterterrorism policy using stock market data. *The Journal of Economic Perspectives*, 20(2):193–206.

Zussman, A., Zussman, N., and Nielsen, M. Ø. (2008). Asset market perspectives on the israeli–palestinian conflict. *Economica*, 75(297):84–115.

Table 1: 20 Most Frequent Arab League and Non-Arab League Names

Arab League		Non-Arab League	
Name	Count	Name	Count
mohamed	557	peter	920
al-	520	jose	864
el-	452	john	859
ahmed	230	de	817
abdel	186	kim	795
ali	180	van	748
ibrahim	102	david	620
ben	101	juan	581
hassan	81	maria	580
sayed	79	carlos	542
abdul	76	lee	540
mahmoud	75	paul	533
abdullah	71	michael	507
khaled	60	robert	492
moustafa	59	luis	479
youssef	55	martin	451
omar	49	jan	438
hussain	48	daniel	407
saleh	47	aleksandr	395
said	42	jean-	393

Table 1 displays the total counts for each name for Arab League and non Arab League countries. Names are taken from Summer Olympic rosters 1948-2102. There are 90,636 Summer Olympians from 221 countries.

Table 2: Summary Statistics for Transaction Data

Statistic	Min	Mean	Median	Max	St. Dev.
Sale Price in $1,000s	45.000	217.418	182.000	1,700.000	136.840
Square Footage	480	1,960.371	1,860	4,850	756.511
Construction Year	1900	1965.219	1969	2002	26.076
Age (years)	0	30.195	25	102	26.137
Bedrooms	1	3.316	3	6	0.840
Bathrooms	1	1.450	1	3	0.573
Sale Year	1982	1995.414	1996	2002	4.785

Table 2 displays summary statistics for the 265,255 transactions in the King County Assessor's data.

Table 3: Control and Treatment Counts Pre and Post 9/11

	$Post = 0$	$Post = 1$	Row Sum
$D^{0.1} = 1$	6,875	445	7,320
$D^{0.3}(1 - D^{0.1}) = 1$	33,725	2,043	35,768
Column Sum	40,600	2,488	43,088

Table 3 displays the total number of transactions near Arab neighbors, before and after 9/11. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days after 9/11.

Table 4: Control Variables in the Treatment and Control Group

Variable	Mean Treatment	Mean Control	t-statistic	p-value
log(sqft)	7.495	7.484	2.23	0.026
Bedrooms	3.332	3.309	2.125	0.034
Bathrooms	1.546	1.478	-8.727	0.00
Construction Year	1969.769	1966.047	10.327	0 .00

Table 4 presents t-statistics for the differnece in means of the control variables for transactions in the control group to transactions in the treatment group. The control group are transaction 0.1-0.3 miles from an Arab homeowner. The treatment group are transactions within 0.1mi of an Arab homeowner. There are 33,014 transactions in the control group and 6,910 transactions in the treatment group. Standard errors are heteroskedasticity robust.

Table 5: Strong Predictors from the Penalized Logistic Model

Arab League			Non-Arab League		
Name	ϕ^*	Count	Name	ϕ^*	Count
fouad	9.838	18	diouf	-8.301	13
salem	9.091	34	akhtar	-8.262	14
khalifa	8.963	27	carolyn	-7.325	20
khaled	8.961	61	teodoro	-6.416	13
mourad	8.713	14	kerstin	-6.311	48
nabil	8.696	21	patrick	-6.124	196
hicham	8.462	17	hernan	-6.058	31
kamel	8.397	27	ud-	-6.008	22
riadh	8.077	11	cedric	-5.806	31
fawzi	8.027	12	diop	-5.604	19
yahia	7.905	11	khan	-5.54	57
jamal	7.894	16	singh	-5.287	359
ramadan	7.878	18	abdoulaye	-5.175	18
khamis	7.862	24	nunez	-5.16	32
tarek	7.73	17	keita	-5.14	24
younes	7.671	12	reza	-5.137	47
abou	7.599	20	eddie	-5.075	53
adel	7.481	29	larry	-4.613	62
alaa	7.321	12	filho	-4.506	40
gamal	7.242	14	syed	-4.301	20

Table 5 displays the 20 names that are the strongest predictors of being from an Arab League and the 20 names that are the strongest predictors of not being from an Arab League country. ϕ^* minimizes the ℓ_1 penalized likelihood model in Equation 2. The probability of being from an Arab League country is given by $\Pr(y_n = 1|X_n, \phi) = \frac{e^{\phi_0 + \sum_p X_{np}\phi_p}}{1+e^{\phi_0 + \sum_p X_{np}\phi_p}}$. Counts indicate the total number of times a given name is found in the Olympic rosters from 1948-2012.

Table 6: Price Effects for Arab Neighbors post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.011^{**}	-0.006	-0.010	-0.006^{**}	-0.007^{**}
	(0.005)	(0.005)	(0.008)	(0.003)	(0.003)
$D^{0.3}$	-0.000	-0.005		-0.005^{***}	-0.005^{***}
	(0.005)	(0.005)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			-0.005		
			(0.005)		
$D^{0.1} \times Post$	-0.018^{***}	-0.014^{***}	-0.027^{***}	-0.005	0.006
	(0.005)	(0.005)	(0.005)	(0.009)	(0.009)
$D^{0.3} \times Post$	0.009^{***}	-0.013^{***}		-0.008	-0.002
	(0.003)	(0.002)		(0.006)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.013^{***}		
			(0.002)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.852	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 6 displays results for the price effects attributable to Arab neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table 7: Price Effects for Arab Neighbors post 9/11, Alternative Specifications

	Base	Distance ≤ 0.5	Bins	Distance	Time	Seasonal
$D^{0.1}$	−0.006	−0.008	−0.011	−0.006	−0.006	−0.006
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
$D^{0.3}$	−0.005	0.003	−0.000	−0.005	−0.0005	−0.005
	(0.005)	(0.004)	(0.005)	(0.007)	(0.005)	(0.005)
$D^{0.1} \times \frac{0.1-Distance}{0.1}$				0.000		
				(0.003)		
$D^{0.1} \times Post$	−0.014***	−0.016***	−0.018***	0.003	−0.021***	−0.014***
	(0.005)	(0.004)	(0.005)	(0.011)	(0.007)	(0.004)
$D^{0.3} \times Post$	−0.013***	−0.008	0.009***	−0.013***	−0.013***	−0.013***
	(0.002)	(0.006)	(0.003)	(0.002)	(0.005)	(0.005)
$D^{0.1} \times Post \times \frac{0.1-Distance}{0.1}$				−0.047**		
				(0.022)		
$D^{0.1} \times Post \times \frac{Days}{180}$					0.014***	
					(0.001)	
Num. obs.	302,065	89,049	302,065	302,065	302,065	302,065
R^2	0.879	0.862	0.883	0.879	0.879	0.879
Tract × Quarter FE	✓	✓	✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	0-180 Days	0-180 Days	0-180 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 7 displays results for the price effects attributable to Arab neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. The variable *Distance* is the distance between the property and the nearest Arab homeowner. The variable *Days* is the number of days since September 11, 2001. The *Base* model uses all transactions since 1980. The *Distance* model uses all transactions within 0.5mi of any Arab homeowner at any point in time. The *Bins* model uses indicator variables for square footage binned every 500 square feet as well as indicator variables for bedrooms, bathrooms. The *Distance* model allows the price effect to vary with *Distance*. The *Time* model allows the price effect to vary with *Days*. The *Seasonal* model includes quarter fixed effects for Q2, Q3, and Q4. All regressions except the *Bins* model include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table 8: Price Effects for Alternative Neighbor Types post 9/11

	Arab League	East Asian	Hispanic	Expanded	MENA
$D^{0.1}$	-0.006	-0.004	-0.016^{***}	0.011	-0.005
	(0.005)	(0.003)	(0.003)	(0.007)	(0.004)
$D^{0.3}$	-0.0005	-0.010^{***}	-0.018^{***}	-0.005	-0.012^{***}
	(0.005)	(0.004)	(0.004)	(0.007)	(0.005)
$D^{0.1} \times Post$	-0.014^{***}	0.001	-0.001	-0.032^{***}	-0.009^{*}
	(0.005)	(0.010)	(0.005)	(0.004)	(0.005)
$D^{0.3} \times Post$	-0.013^{***}	-0.020^{***}	-0.022^{***}	-0.008^{***}	-0.012^{***}
	(0.002)	(0.005)	(0.003)	(0.003)	(0.004)
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.879	0.879	0.879	0.879	0.879
Tract × Quarter FE	✓	✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	0-180 Days	0-180 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 8 displays results for the price effects attributable to neighbor types following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of a specific homeowner type. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. *Base* uses the Arab League as the identifying countries. *EastAsian* uses China, Korea, and Japan as the identifying countries. *Hispanic* uses Mexico, Spain, and all Latin and South American countries as the identifying countries. *Expanded* uses Arab League countries, Bangladesh, Indonesia, Iran, Pakistan, and Turkey as the identifying countries. $MENA$ uses the Middle Eastern and North African countries as defined in the US Census. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table 9: Price Effects for Arab Neighbors and Mosques post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.011^{***}	-0.008	-0.016^{*}	-0.008^{***}	-0.009^{***}
	(0.003)	(0.005)	(0.009)	(0.003)	(0.003)
$D^{0.3}$	-0.003^{**}	-0.009		-0.008^{***}	-0.009^{***}
	(0.001)	(0.006)		(0.002)	(0.002)
$D^{0.3} \times (1 - D^{0.1})$			-0.009		
			(0.006)		
$M^{0.1}$	0.006	0.001	-0.032^{***}	-0.002	-0.002
	(0.009)	(0.016)	(0.012)	(0.009)	(0.009)
$M^{0.3}$	-0.036^{***}	-0.033		-0.031^{***}	-0.031^{***}
	(0.004)	(0.025)		(0.006)	(0.006)
$M^{0.1} \times (1 - M^{0.3})$			-0.033		
			(0.025)		
$D^{0.1} \times Post$	-0.021^{*}	-0.018^{**}	-0.027^{***}	-0.004	0.002
	(0.011)	(0.007)	(0.006)	(0.009)	(0.011)
$D^{0.3} \times Post$	0.012^{**}	-0.009^{***}		-0.010	0.003
	(0.006)	(0.003)		(0.007)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.009^{***}		
			(0.003)		
$M^{0.1} \times Post$	-0.114^{**}	-0.013	-0.011	0.047	0.050
	(0.052)	(0.017)	(0.009)	(0.040)	(0.040)
$M^{0.3} \times Post$	0.038^{**}	0.003		-0.037	-0.038
	(0.017)	(0.023)		(0.033)	(0.033)
$M^{0.1} \times (1 - M^{0.3}) \times Post$			0.003		
			(0.023)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.852	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 9 displays results for the time-varying effect of Arab neighbors and mosques following September 11, 2000. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $M^{0.1} = 1$ ($M^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of a Mosque or Islamic Center. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Figure 1: Price Gradient Near Arab Homeowners

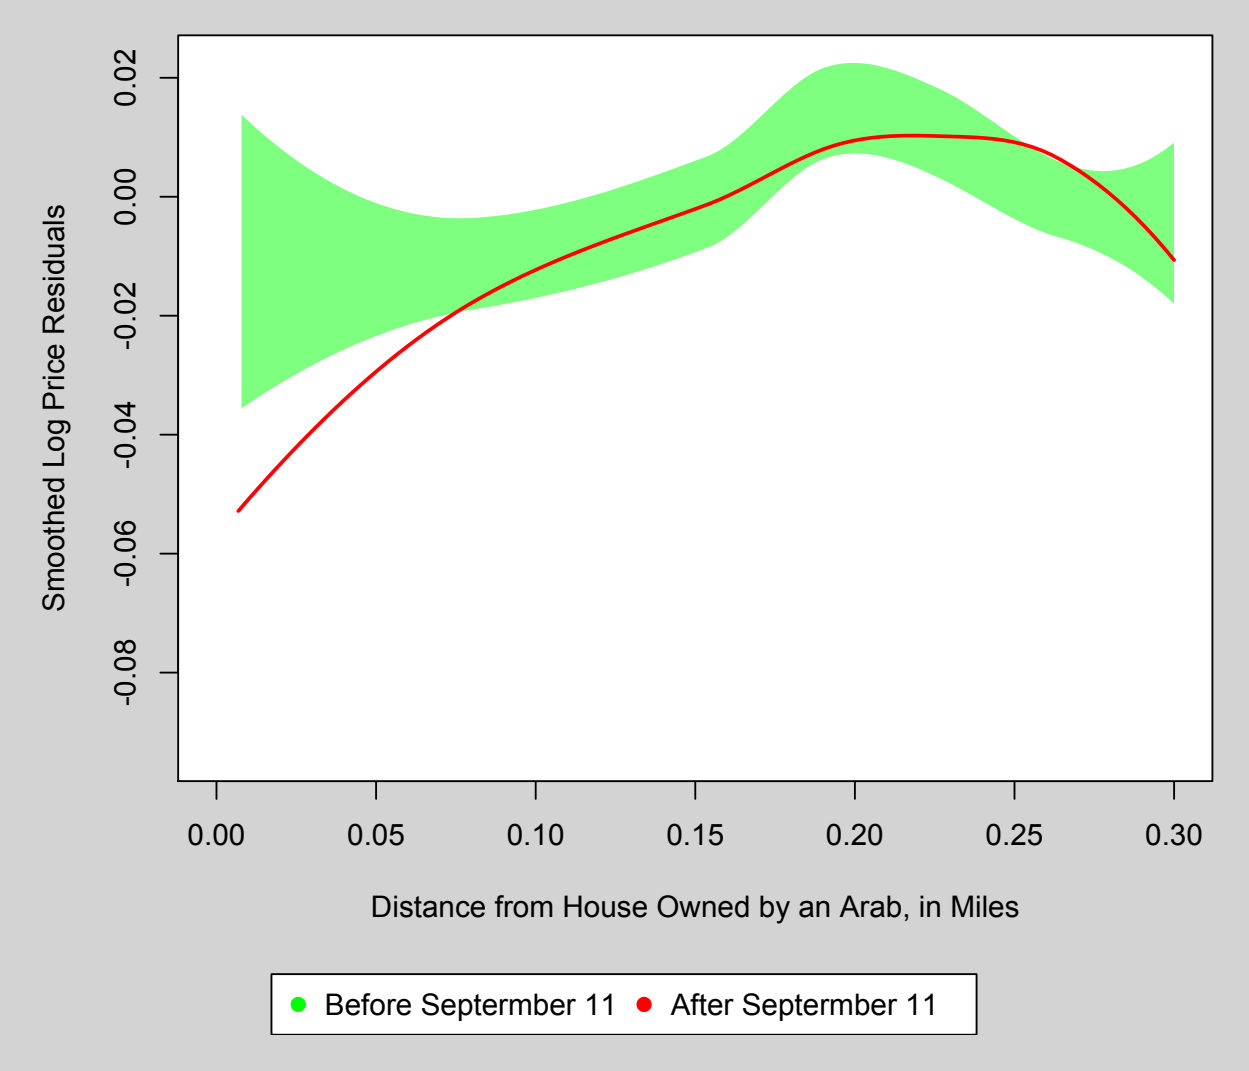

Figure 1 displays a local polynomial estimation of hedonic residuals less than 0.3 miles from an Arab neighbor as a function of distance from the nearest Arab neighbor. The 95% confidence interval for transactions 0-180 days before 9/11 is presented in green, and the point estimates for transactions 0-180 days after 9/11 are displayed in red.

Figure 2: Treatment Area and Control Area

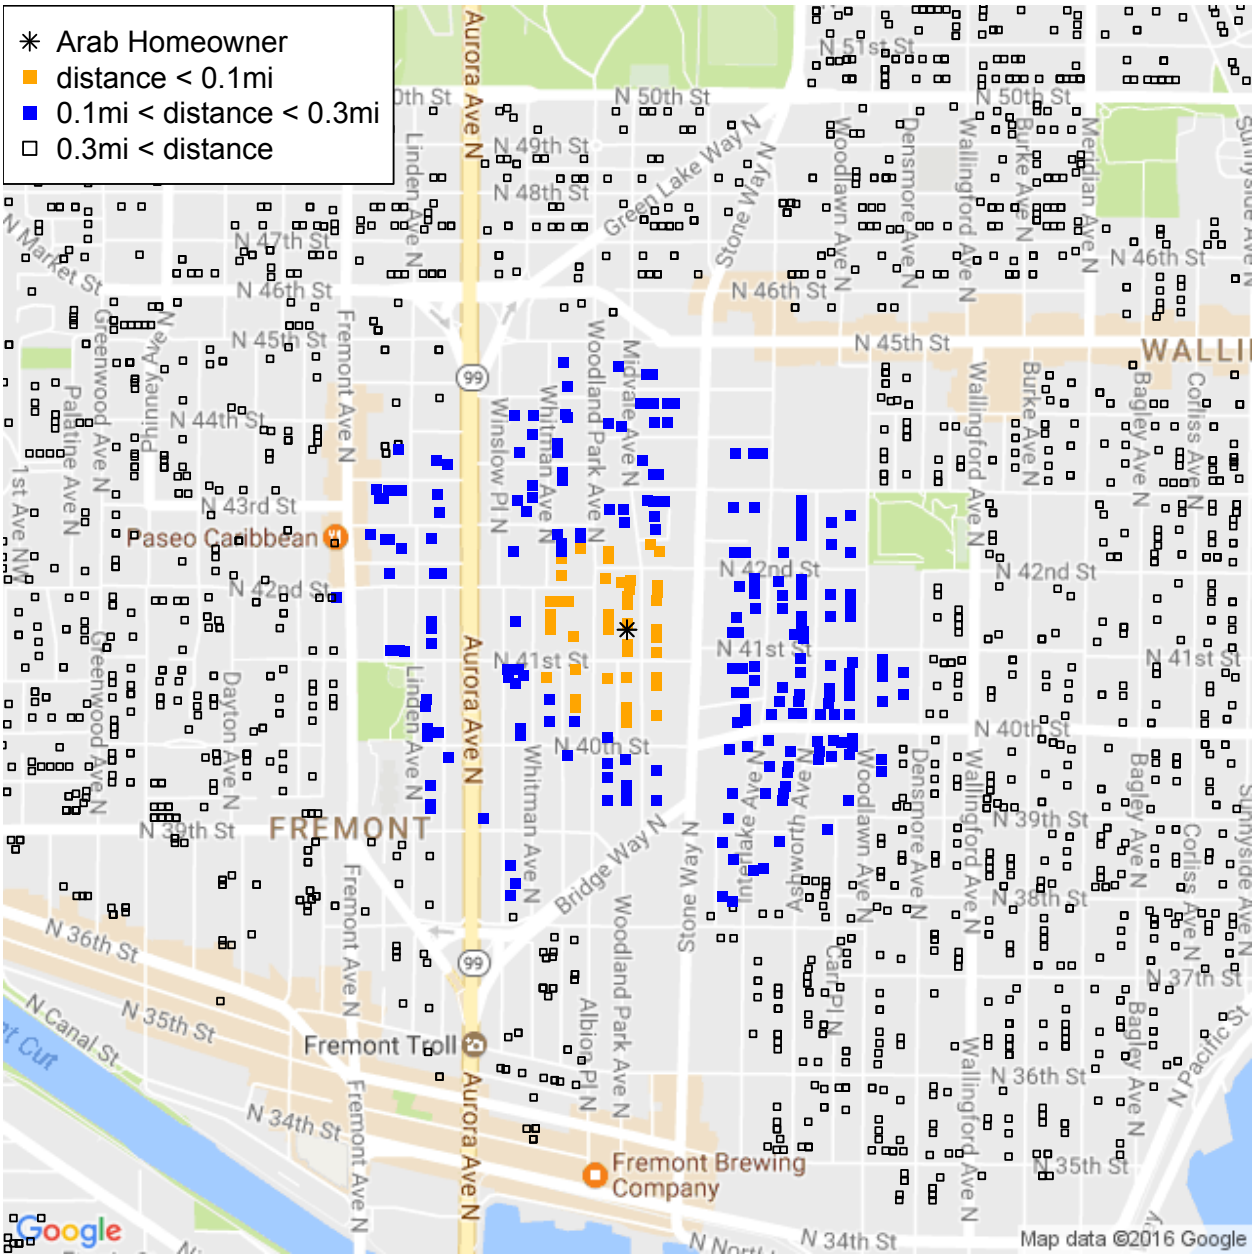

Figure 2 displays an example of a random property, properties less than 0.1 miles, properties 0.1-0.3 miles, and properties more than 0.3 miles away.

Figure 3: Olympic Roster Names

Arab League Names

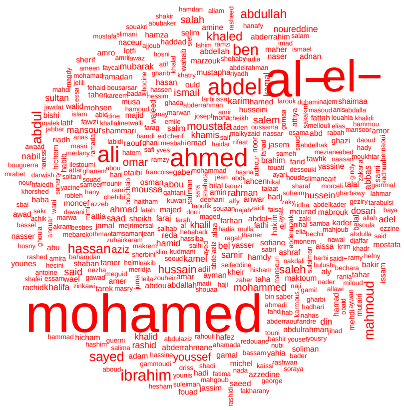

Non-Arab League Names

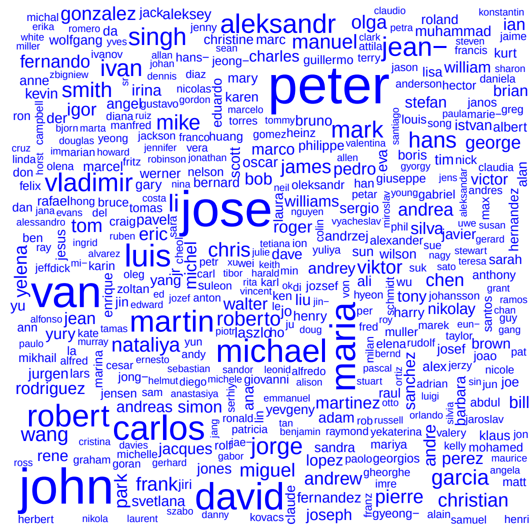

Figure 3 displays the 300 most frequent names on the Olympic rosters for each country. More frequent names are indicated with a larger font.

Figure 4: Location of Arab Homeowners

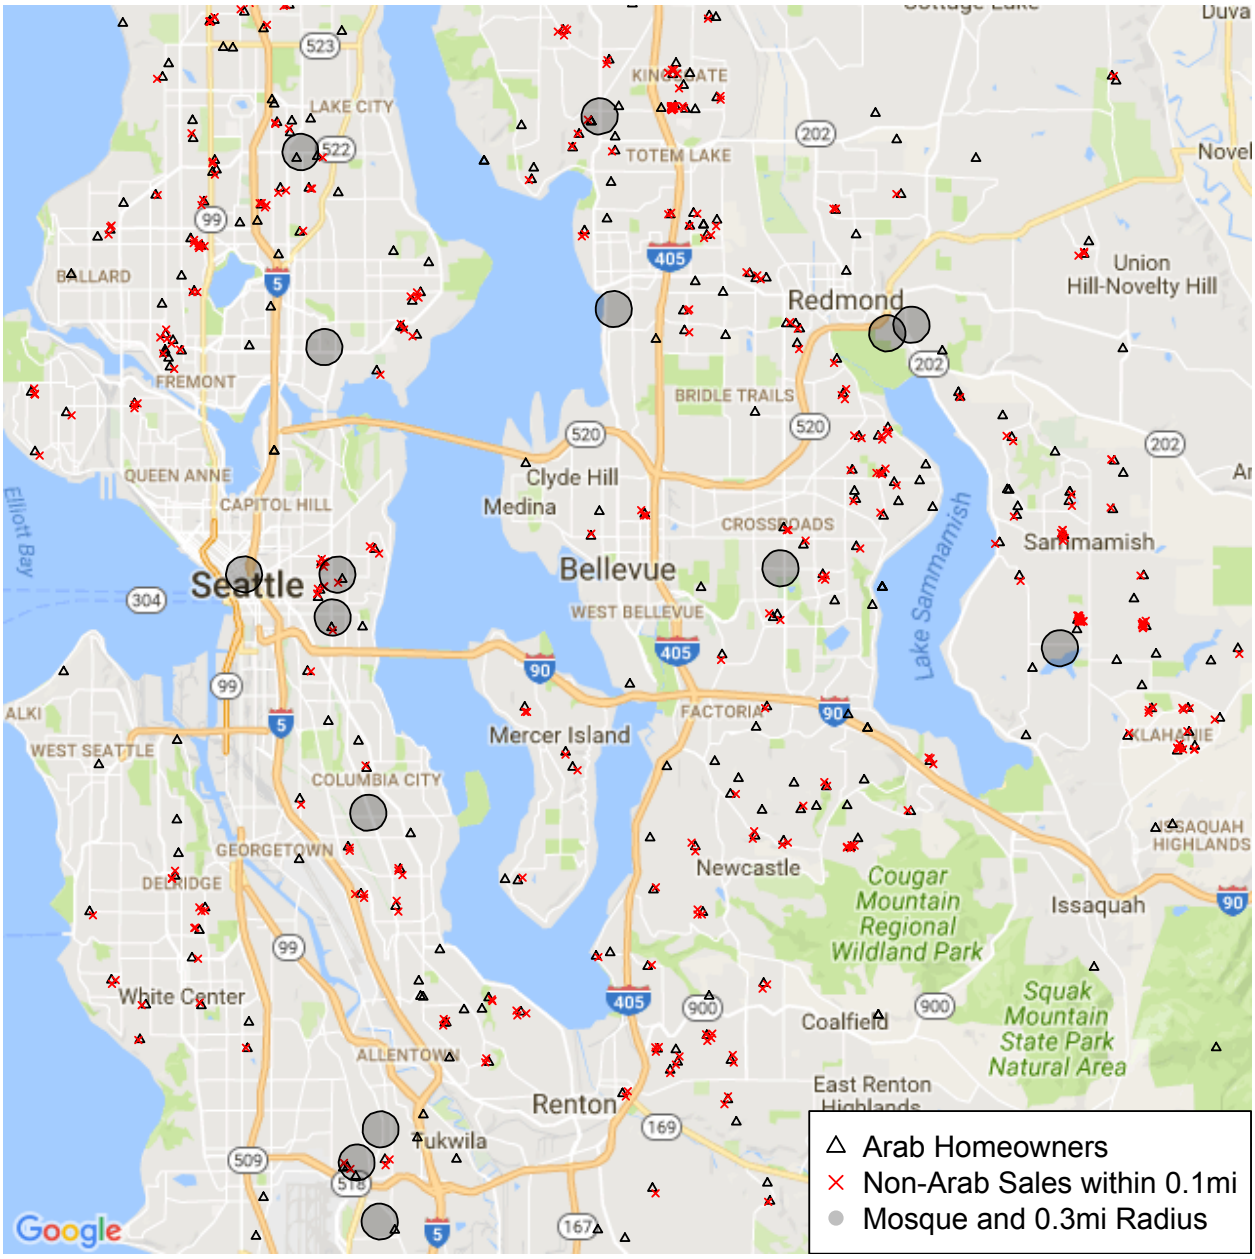

Figure 4 displays a significant area of the study area in the vicinity of Seattle, Washington. Figure 4 displays Arab homeowners on 9/11, transactions by Non-Arab homeowners within 0.1 miles of an Arab neighbor 0-180 days after 9/11, and mosques with a 0.3 mile radius.

Figure 5: Treatment Area and Control Control Variables

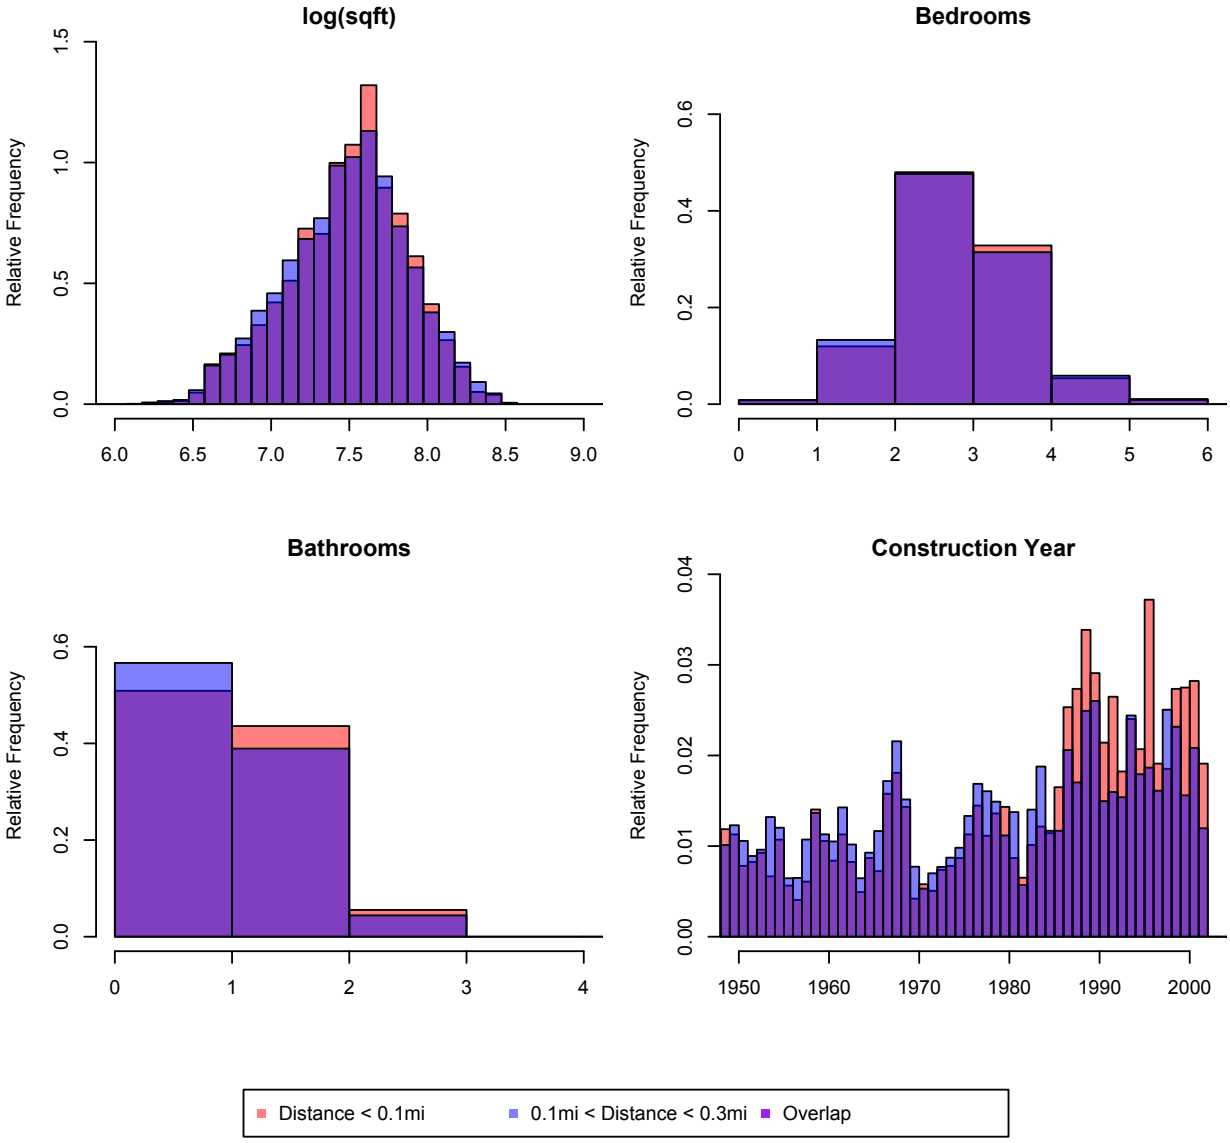

Figure 5 displays histograms of the control variables used in the hedonic regressions for transactions in the control and treatment group. The control group includes all transactions 0.1-0.3 miles from an Arab neighbor, and the treatment group includes all transactions less than 0.1 miles from an Arab neighbor.

Figure 6: Logit Coefficients and Relative Frequency

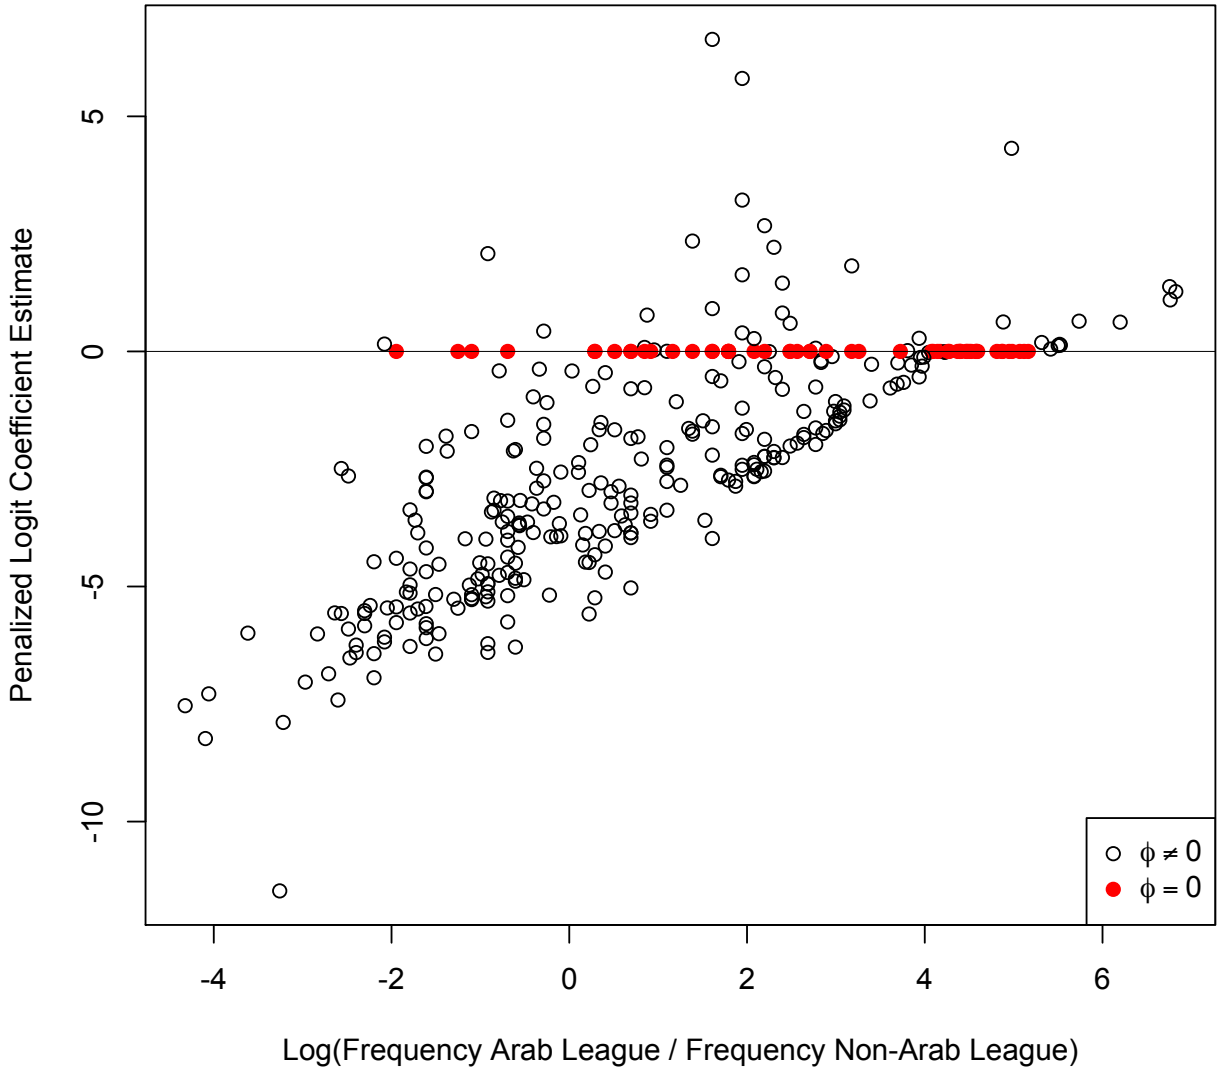

Figure 6 displays the coefficients of ϕ^* relative to the log of the relative frequencies of names between the Arab Leage and non Arab League countries. ϕ^* is the set of coefficients that minimize the penalized liklelihood in Equation 2. The number of times name p is found in the full names of Olympians from group g, $N_g(p)$, divided by the total number of Olympians in group g, N_g. The relative frequency for name p in group g is equal to $f_g(p) = \frac{N_g(p)}{N_g}$. The log relative frequency is equal to $\log(f_{ArabLeague}(p)) - \log(f_{NonArabLeague}(p))$.

Figure 7: Common Trends in the Price Index

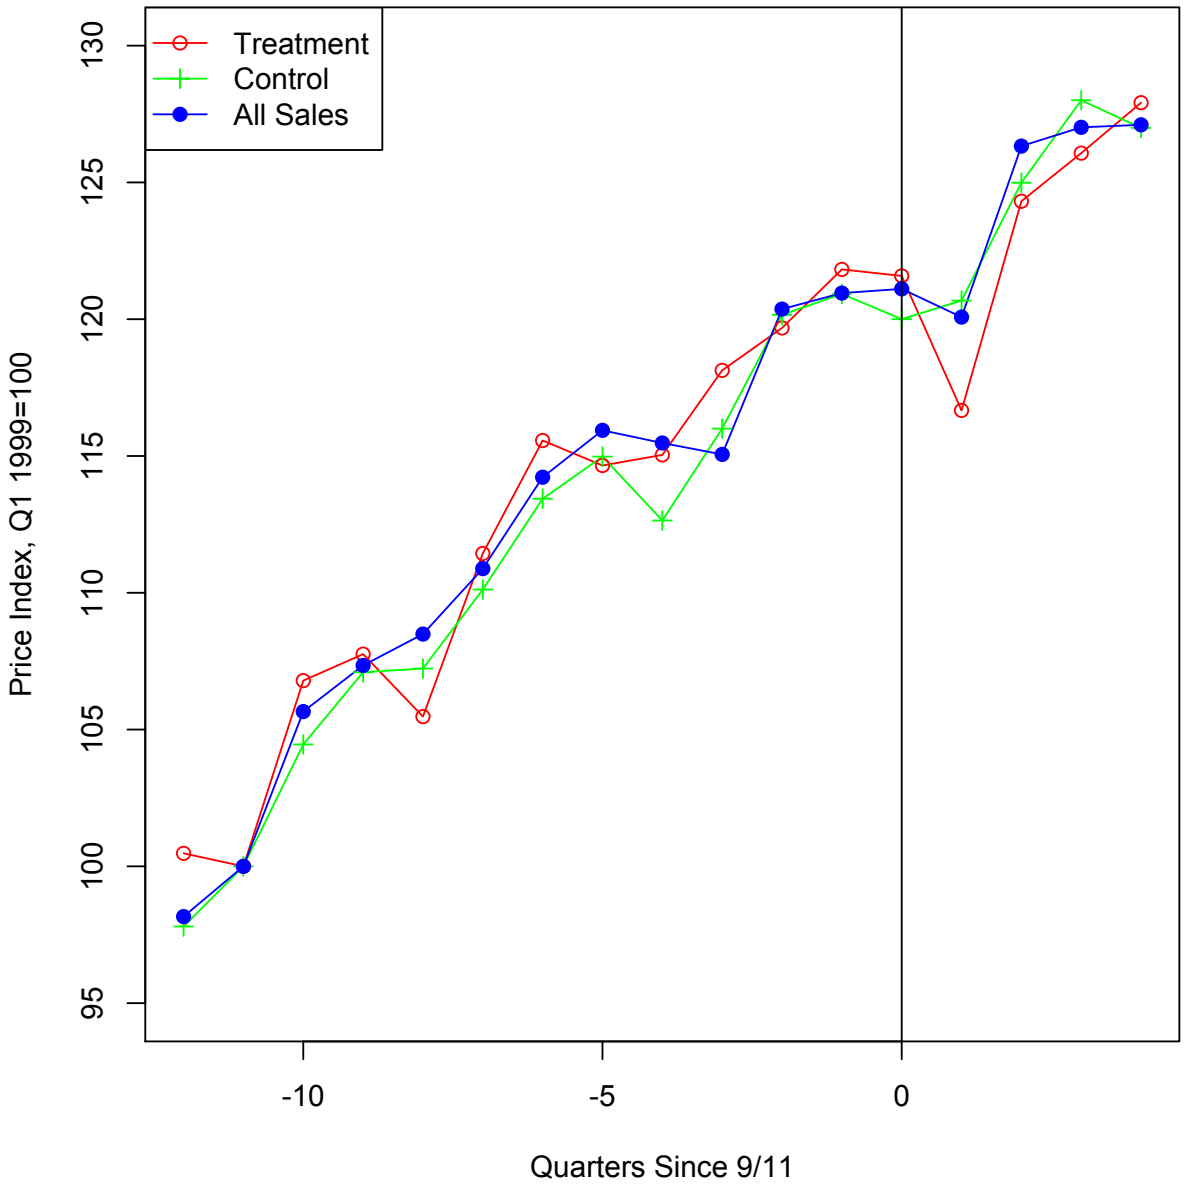

Figure 7 presents the quarterly hedonic price index for the *treatment*, *control*, and *allsales* sample of sales. The *treatment* group is the set of all sales within 0.1mi of an Arab neighbor. The *control* group is the set of all sales 0.1-0.3mi from an Arab neighbor. The *allsales* sample is the set of transactions in the data. Quarters begin on the 11th of March, June, September, and December. The period $t = 0$ corresponds to the period June 12, 2001 - September 10, 2001; $t = 1$ corresponds to the period September 11, 2001 - December 10, 2001. The index is normalized to 100 in the period 11 quarters before September 11, 2001.

Figure 8: Alternative Ethnic Group Price Gradients

East Asian Homeowners

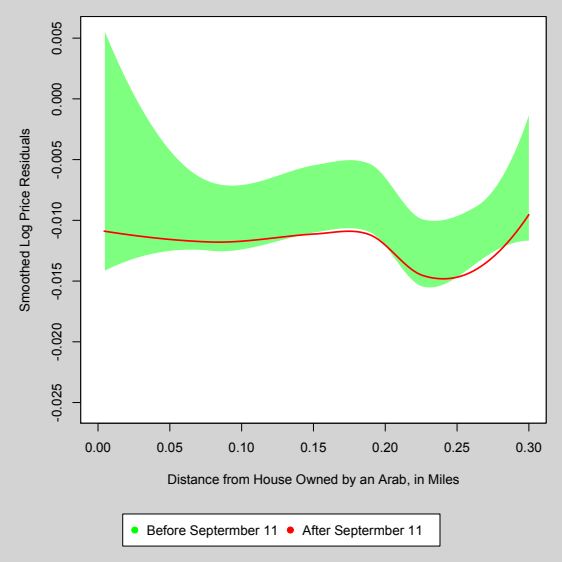

Hispanic Homeowners

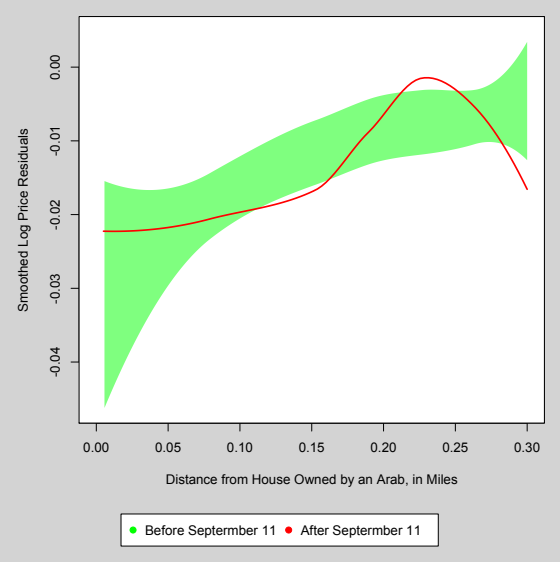

Figure 8 displays a local polynomial estimation of hedonic residuals from transactions less than 0.1 miles from East Asian and Hispanic neighbor as a function of distance from the nearest East Asian or Hispanic neighbor. The 95% confidence interval for transactions 0-180 days before 9/11 is presented in green, and the point estimates for transactions 0-180 days after 9/11 are displayed in red.

Figure 9: Concentration of Arab Homeowners

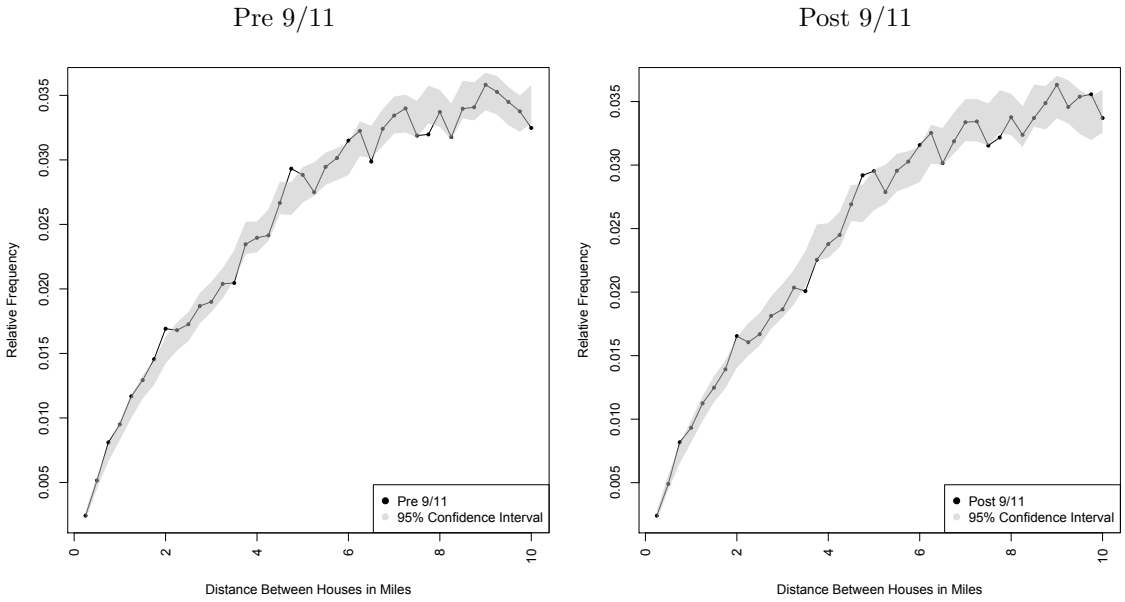

Figure 9 presents the probability mass function of pairwise distances for each Arab homeowner in the data before and after 9/11. Distances are grouped into 0.25mi bins. Confidence intervals are calculated using 2,000 random samples drawn from all transactions in the respective time periods.

Appendix

Data filters

1. Remove all transactions with sale prices less than the 0.5% quantile and sale prices greater than the 99.5% quantile of sale prices
2. Remove all transactions with square footage less than the 0.5% quantile of square footage and sale prices greater than the 99.5% quantile of square footage
3. Remove all transactions with more than 6 bedrooms
4. Remove all transactions with more than 3 bathrooms
5. Remove all transactions for properties constructed before 1900

Table A1: Countries and Types

Country	Arab League	Expanded	MENA	East Asian	Hispanic
algeria	✓	✓	✓		
argentina					✓
bahrain	✓	✓	✓		
bangladesh		✓			
bolivia					✓
brazil					✓
chile					✓
china				✓	
colombia					✓
comoros	✓	✓			
costa rica					✓
cuba					✓
djibouti	✓	✓			
dominican republic					✓
ecuador					✓
egypt	✓	✓	✓		
el salvador					✓
guatemala					✓
honduras					✓
hong kong				✓	
indonesia		✓			
iran		✓	✓		
iraq	✓	✓	✓		
israel			✓		
japan				✓	
jordan	✓	✓	✓		
kuwait	✓	✓	✓		
lebanon	✓	✓	✓		
libya	✓	✓	✓		
mauritania	✓	✓			
mexico					✓
morocco	✓	✓	✓		
nicaragua					✓
north korea				✓	
oman	✓	✓	✓		
pakistan		✓			
palestine	✓	✓	✓		
panama					✓
paraguay					✓
peru					✓
qatar	✓	✓	✓		
saudi arabia	✓	✓	✓		
somalia	✓	✓			
south korea				✓	
spain					✓
sudan	✓	✓			
syria	✓	✓	✓		
tunisia	✓	✓	✓		
turkey		✓			
united arab emirates	✓	✓	✓		
united arab republic	✓	✓	✓		
uruguay					✓
venezuela					✓
yemen	✓	✓	✓		

Table A1 lists the countries in the Arab League, the Arab League as well as countries with a significant number and percentage of Muslim citizens, Middle East and North Africa countries as defined by the US Census. We also include countries we define as East Asian and Hispanic.

Table A2: Price Effects for Arab Neighbors post 9/11, 1990-2002

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.011^{**}	-0.006	-0.011	-0.007^{**}	-0.007^{***}
	(0.005)	(0.005)	(0.008)	(0.003)	(0.003)
$D^{0.3}$	-0.004	-0.005		-0.005^{***}	-0.006^{***}
	(0.005)	(0.005)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			-0.005		
			(0.005)		
$D^{0.1} \times Post$	-0.018^{***}	-0.013^{***}	-0.026^{***}	-0.002	0.007
	(0.005)	(0.005)	(0.005)	(0.008)	(0.009)
$D^{0.3} \times Post$	0.009^{***}	-0.012^{***}		-0.007	-0.002
	(0.003)	(0.002)		(0.006)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.012^{***}		
			(0.002)		
Num. obs.	262,601	262,601	262,601	262,601	262,601
R^2	0.827	0.851	0.851	0.851	0.851
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A2 displays results for the price effects attributable to Arab neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1990 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A3: Price Effects for Arab Neighbors post 9/11, 2000-2002

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	−0.003	−0.002	−0.011	−0.004	−0.005
	(0.005)	(0.005)	(0.008)	(0.004)	(0.004)
$D^{0.3}$	−0.012**	−0.009*		−0.010***	−0.011***
	(0.005)	(0.005)		(0.003)	(0.003)
$D^{0.1} \times (1 - D^{0.3})$			−0.009*		
			(0.005)		
$D^{0.1} \times Post$	−0.015***	−0.017***	−0.024***	−0.003	0.005
	(0.006)	(0.006)	(0.005)	(0.009)	(0.010)
$D^{0.3} \times Post$	0.004	−0.006***		−0.003	0.003
	(0.003)	(0.002)		(0.007)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			−0.006***		
			(0.002)		
Num. obs.	69,431	69,431	69,431	69,431	69,431
R^2	0.803	0.820	0.820	0.820	0.820
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A3 displays results for the price effects attributable to Arab neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 2000 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A4: Price Effects for Arab Neighbors post 9/11, Cutoff 0.5

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	−0.008	−0.004	−0.009	−0.005	−0.005*
	(0.006)	(0.006)	(0.008)	(0.003)	(0.003)
$D^{0.3}$	−0.003	−0.005		−0.005***	−0.006***
	(0.005)	(0.006)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			−0.005		
			(0.006)		
$D^{0.1} \times Post$	−0.021***	−0.018***	−0.027***	−0.002	0.005
	(0.005)	(0.006)	(0.005)	(0.010)	(0.010)
$D^{0.3} \times Post$	0.012***	−0.009***		−0.008	−0.003
	(0.002)	(0.002)		(0.006)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			−0.009***		
			(0.002)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.852	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A4 displays results for the price effects attributable to Arab neighbors following September 11, 2001. Arab neighbors are identified using a 0.5 probability cutoff. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A5: Price Effects for East Asian Neighbors post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.006^{**}	-0.004	-0.014^{***}	-0.004^{***}	-0.004^{***}
	(0.003)	(0.003)	(0.005)	(0.001)	(0.001)
$D^{0.3}$	-0.008^{**}	-0.010^{***}		-0.010^{***}	-0.010^{***}
	(0.004)	(0.004)		(0.001)	(0.001)
$D^{0.1} \times (1 - D^{0.3})$			-0.010^{***}		
			(0.004)		
$D^{0.1} \times Post$	-0.004	0.001	-0.019^{**}	-0.001	-0.001
	(0.011)	(0.010)	(0.009)	(0.005)	(0.005)
$D^{0.3} \times Post$	0.017^{***}	-0.020^{***}		-0.016^{**}	-0.016^{**}
	(0.004)	(0.005)		(0.008)	(0.008)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.020^{***}		
			(0.005)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.852	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A5 displays results for the price effects attributable to East Asian neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an East Asian homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A6: Price Effects for Hispanic Neighbors post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.019^{***}	-0.016^{***}	-0.033^{***}	-0.015^{***}	-0.016^{***}
	(0.003)	(0.003)	(0.005)	(0.001)	(0.001)
$D^{0.3}$	-0.018^{***}	-0.018^{***}		-0.018^{***}	-0.018^{***}
	(0.004)	(0.004)		(0.001)	(0.001)
$D^{0.1} \times (1 - D^{0.3})$			-0.018^{***}		
			(0.004)		
$D^{0.1} \times Post$	-0.001	0.001	-0.022^{***}	-0.005	0.000
	(0.004)	(0.005)	(0.003)	(0.006)	(0.005)
$D^{0.3} \times Post$	-0.000	-0.022^{***}		-0.007	-0.009
	(0.007)	(0.003)		(0.006)	(0.006)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.022^{***}		
			(0.003)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.853	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A6 displays results for the price effects attributable to Hispanic neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Hispanic homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A7: Price Effects for Arab and Expanded Neighbors post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	0.007	0.011	0.006	0.009**	0.009**
	(0.007)	(0.007)	(0.010)	(0.004)	(0.004)
$D^{0.3}$	0.003	−0.005		−0.006***	−0.005**
	(0.006)	(0.007)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			−0.005		
			(0.007)		
$D^{0.1} \times Post$	−0.041***	−0.032***	−0.039***	0.005	0.002
	(0.004)	(0.004)	(0.006)	(0.012)	(0.013)
$D^{0.3} \times Post$	0.002	−0.008***		0.006	−0.001
	(0.004)	(0.003)		(0.007)	(0.008)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			−0.008***		
			(0.002)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.853	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A7 displays results for the price effects attributable to Arab, Indonesian, Pakistani, Bangladeshi, Turkish, or Iranian neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab, Indonesian, Pakistani, Bangladeshi, Turkish, or Iranian homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A8: Price Effects for Middle East and North African Neighbors post 9/11

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.008^{**}	-0.005	-0.018^{***}	-0.005^{**}	-0.006^{***}
	(0.004)	(0.004)	(0.006)	(0.002)	(0.002)
$D^{0.3}$	-0.003	-0.012^{***}		-0.013^{***}	-0.013^{***}
	(0.004)	(0.005)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			-0.012^{***}		
			(0.005)		
$D^{0.1} \times Post$	-0.014^{***}	-0.009^{*}	-0.021^{***}	-0.011	0.004
	(0.005)	(0.005)	(0.003)	(0.007)	(0.008)
$D^{0.3} \times Post$	0.012^{***}	-0.012^{***}		-0.003	0.001
	(0.004)	(0.004)		(0.006)	(0.006)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.012^{***}		
			(0.004)		
Num. obs.	302,065	302,065	302,065	302,065	302,065
R^2	0.853	0.879	0.879	0.879	0.879
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A8 displays results for the price effects attributable to Middle East and North African neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Middle East and North African homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A9: Price Effects for Arab Neighbors post 9/11, No Nearby Multifamily Properties

	Model 1	Model 2	Model 3	Model 4	Model 5
$D^{0.1}$	-0.013^{***}	-0.008	-0.015^{**}	-0.008^{***}	-0.008^{***}
	(0.005)	(0.005)	(0.007)	(0.003)	(0.003)
$D^{0.3}$	-0.003	-0.007		-0.007^{***}	-0.008^{***}
	(0.005)	(0.005)		(0.002)	(0.002)
$D^{0.1} \times (1 - D^{0.3})$			-0.007		
			(0.005)		
$D^{0.1} \times Post$	-0.012^{**}	-0.009^{***}	-0.024^{***}	-0.001	0.001
	(0.006)	(0.003)	(0.002)	(0.009)	(0.010)
$D^{0.3} \times Post$	0.005^{***}	-0.015^{***}		-0.014^{**}	-0.001
	(0.002)	(0.002)		(0.006)	(0.007)
$D^{0.1} \times (1 - D^{0.3}) \times Post$			-0.015^{***}		
			(0.002)		
Num. obs.	250636	250636	250636	250636	250636
R^2 (full model)	0.860	0.888	0.888	0.888	0.888
Tract + Quarter FE	✓				
Tract × Quarter FE		✓	✓	✓	✓
Post 9/11 Window	0-180 Days	0-180 Days	0-180 Days	180-365 Days	-180-0 Days

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A9 displays results for the price effects attributable to Arab neighbors following September 11, 2001. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of an Arab homeowner. The variable $Post = 1$ if the transaction is 0-180 days, 0-90 days, or 180-365 days after September 11, 2001. All single-family property less than 0.1mi from a multifamily property are removed. All regressions include log square footage, bedrooms, bathrooms, and construction year as control variables. All regressions use all transactions January 1, 1982 to December 31, 2002. All regressions use standard errors clustered at the census tract and quarter levels.

Table A10: Control and Treatment Counts Pre and Post 9/11, all Types

Panel A: East Asian			
	$Post = 0$	$Post = 1$	Row Sum
$D^{0.1} = 1$	59,772	2,934	62,706
$D^{0.3}(1 - D^{0.1}) = 1$	99,767	4,246	104,013
Column Sum	159,539	7,180	166,719

Panel B: Hispanic			
	$Post = 0$	$Post = 1$	Row Sum
$D^{0.1} = 1$	78,510	4,071	82,581
$D^{0.3}(1 - D^{0.1}) = 1$	111,962	3,861	115,82
Column Sum	190,472	7,932	198,404

Panel C: Expanded Arab League			
	$Post = 0$	$Post = 1$	Row Sum
$D^{0.1} = 1$	13,240	828	14,068
$D^{0.3}(1 - D^{0.1}) = 1$	54,003	3,186	57,179
Column Sum	67,243	4,014	71,257

Panel D: Middle East and North Africa			
	$Post = 0$	$Post = 1$	Row Sum
$D^{0.1} = 1$	12,042	741	12,783
$D^{0.3}(1 - D^{0.1}) = 1$	49,900	2,955	52,855
Column Sum	61,942	3,696	65,638

Table A10 displays the total number of transactions near various neighbor types, before and after 9/11. The variable $D^{0.1} = 1$ ($D^{0.3}$=1) if the transaction is within 0.1mi (0.3mi) of a given neighbor type. The variable $Post = 1$ if the transaction is 0-180 days after 9/11.

Figure A1: Price Index by Census Tract

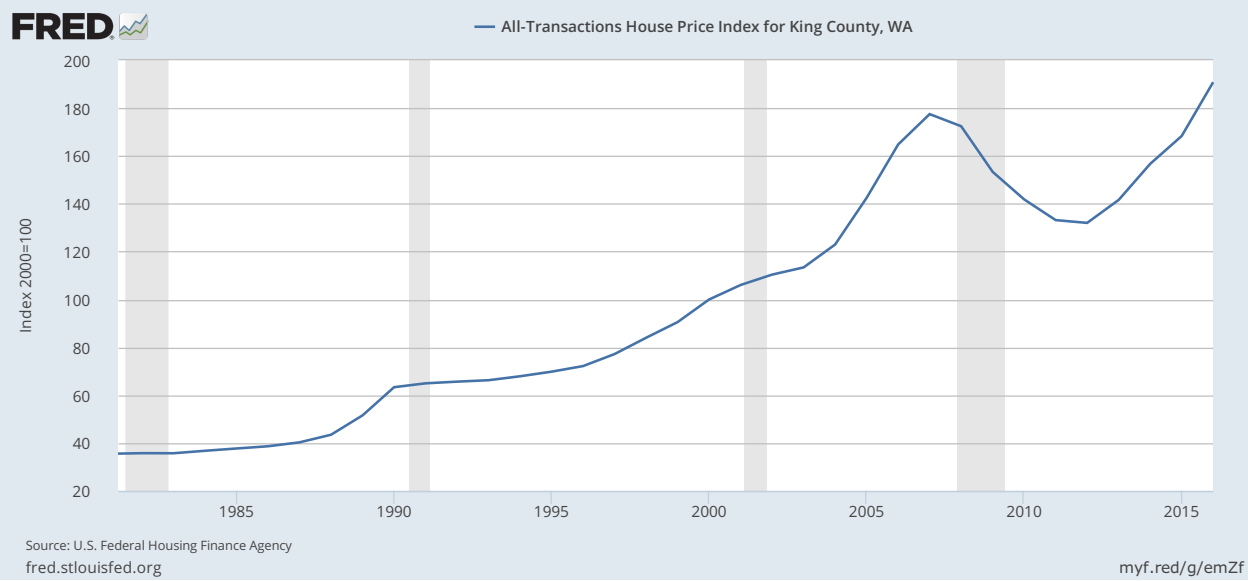

Figure A1 presents the Federal Housing Finance Administration house price index for King County, Washington. Source: Federal Reserve Bank of St. Louis and U.S. Federal Housing Finance Administration. *https://fred.stlouisfed.org/series/ATNHPIUS53033A*