U.S. Housing as a Global Safe Asset: Evidence from China Shocks*

[PRELIMINARY; PLEASE DO NOT CITE WITHOUT AUTHORS' PERMISSION]

William Barcelona Nathan Converse Anna Wong

January 1, 2020

Abstract

This paper examines the causal impact of international capital flows from China on one of the least accurately measured and understood global safe haven assets: U.S. residential real estate. We demonstrate using aggregate capital flows data that the recent rise in unrecorded inflows in the U.S. balance of payments is likely attributable to inflows being used to purchase U.S. residential properties, mainly to inflows originating from China. We then exploit a novel, direct measure of Chinese demand for U.S. residential properties at the local level by using a web traffic dataset from a real estate listing website that specializes in marketing foreign residential properties to users based in Mainland China. With a difference-in-difference matching framework, we find that house prices in China-exposed areas of major U.S. cities have on average grown seven percentage points faster than in similar neighborhoods with low exposure to Chinese buyers over the period 2010-2016, when the U.S. experienced two episodes of surges in safe haven flows from China. We then show that the time variation in the average excess price growth in China exposed-areas comoves closely with macro-level measures of U.S. capital inflows from China. Following periods of economic stress in China, inflows from China to the U.S. jump and the price growth gap widens, suggesting that Chinese households view U.S. housing as a safe haven asset.

JEL classification: F3, F6

Keywords: China, real estate, capital flows, safe assets

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the view of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. We are grateful for comments from Carol Bertaut, Steve Kamin, Tyler Goodspeed, and seminar participants at the Board of Governors of the Federal Reserve for their insightful comments at the early stages of this project. The authors are employed in the Division of International Finance at the Federal Reserve Board. Correspondence: nathan.l.converse@frb.gov or anna.wong@frb.gov, Federal Reserve Board, Washington DC 20551.

1 Introduction

This paper uses a unique dataset to examine the causal impact of international capital flows on one of the least accurately measured and understood global safe haven assets: U.S. residential real estate. Flows of foreign capital into the U.S. housing market are not measured in the U.S. balance of payment data, but rather relegated to the inconspicuous "statistical discrepancy" line in official statistics.[1] Historically, the unmeasured net capital inflows subsumed into the statistical discrepancy have tended to rise during times of global financial stress, such as during the Asian Financial Crisis and the Global Financial Crisis (Flatness et al., 2009). However, as we will show, in the past ten years unmeasured capital inflows have particularly surged during periods of economic distress in China, for example rising to 1.5 percent of U.S. GDP in 2015, a period when the scale of capital outflows from China dwarfed that of other countries. Moreover, the measurement of private international capital flows presents many challenges, but tracking Chinese households' purchases of foreign assets is particularly difficult due to the widespread use of informal channels in order to circumvent Chinese capital controls.

The unknown scale of capital inflows into the residential real estate market in the U.S. and indeed most advanced economies has posed significant challenges not only for researchers but also for policymakers. An increasing number of national and local authorities in countries such as Canada, New Zealand, Australia, and Hong Kong have imposed restrictions on foreign buyers, basing their policy action on anecdotal reports of substantial purchases by foreign Chinese buyers. Some countries, such as Canada, have left it to individual municipal governments to adopt policies vis a vis foreign buyers as they see fit. Other countries, such as New Zealand, have implemented policies at the national level.

This paper tackles these measurement problems by first closely examining relevant macroeconomic data from China and the U.S. and second by using a novel micro-level dataset that allows us to directly measure cross sectional variation in Chinese demand for U.S. residential real estate. In particular, we highlight the unusually close comovement between the unmeasured capital in-

[1] While the most recent IMF Balance of Payments Manual (BPM6) specifies that "real estate investment, including investment properties and vacation homes" should be included in foreign direct investment (FDI), the U.S. and indeed most countries do not measure such flows. As a result, these flows are captured only in the residual line of the balance of payments. Some countries, notably China, refer to this line in the balance of payments as "errors and omissions." For this reason, we will use the term "statistical discrepancy" when referring to U.S. official statistics and "errors and omissions" when referring to Chinese data.

flows captured flows in the U.S. statistical discrepancy and inflows of money and deposits from China recorded by the U.S. Treasury International Capital (TIC) System. Then we use a unique micro dataset to measure variation in Chinese demand for residential real estate across U.S. cities and estimates the causal impact of Chinese households' purchases of housing in cities all over the United States. Our micro data come from a rich web traffic dataset obtained from a popular Chinese website that specializes in listing foreign residential properties, broadly similar to Zillow but tailored for Chinese users. The dataset provides us with counts of Chinese users' views of U.S. properties in over 9,000 individual U.S. ZIP codes, thus giving us a direct measure of variation in Chinese residents' demand for local U.S. residential housing across the country. We then adopt a difference-in-difference matching framework to compare house price growth in areas heavily exposed to Chinese foreign buyers with price changes in areas that have low exposure to Chinese capital inflows, but which are otherwise similar. With this empirical framework, we disentangle the impact of the foreign demand shock from that of national and local factors, such as changes in mortgage rates or local economic activity, as well as unobserved local factors.

We find evidence of an impact of foreign capital inflows on the evolution of house prices in major U.S. cities. The magnitude of the premium in price growth in China-exposed over non-exposed areas averages about 2 percent in the two years following each of the two episodes of China shocks in the last decade. We then use local projections to assess the relationship between the price effects we uncover and capital inflows from China, and find that the latter variable explains the majority of the widening in price growth differentials after Chinese stress episodes. Overall then, the evidence we present is consistent with inflows from China picking up after periods of economic stress in China as foreign Chinese households purchase U.S. residential real estate as a safe haven asset.

We conduct various checks to confirm that the patterns we uncover reflect the effect of capital inflows from China on the U.S. housing market rather than variation in local, national, or global economic conditions. We perform a placebo matching exercise in which we identify as the treatment group areas with high house price growth and show that the price gap between these hot markets and matched controls evolves completely differently than the gap for the areas we identified China-exposed. This provides further confirmation that the treatment effect we uncover is not merely a result of Chinese households buying into hot markets. When relating the price growth gap we uncover to macro-level variables, we show that the gap is not explained by the state of the U.S. business cycle or the level of mortgage rates. We also verify that the price growth gap is unrelated to deposit inflows from countries other than China. Thus the data are not consistent with alternative stories relating the price growth gap to local, national, or global financial conditions.

This paper contributes to three areas of literature. First, we add to the literature on how international capital flows affect housing prices by offering a fresh interpretation of existing macro data and a novel measure of foreign demand on U.S. housing at the national level. Recent work has documented that real estate purchases by foreign residents affects house prices in individual cities in the U.S. (Li et al., 2019), the United Kingdom (Badarinza and Ramadorai, 2018), and Germany (Bednarek et al., 2019). Due to a lack of data on foreign ownership of housing, these studies have proxied for foreign demand through indirect inferences. For example, Li et al. (2019) proxy for Chinese purchases of residential real estate in three cities in California by using the share of ethnic Chinese residents. Badarinza and Ramadorai (2018) also derives the foreign demand shock by linking the ethnicity of neighborhoods in London to related-country shocks. Our paper differs from both studies, not only in that that our analysis is national in scope, but also in our measure of foreign demand, page views data from a popular real estate listing site catered to users in China, which gives a more direct measure of demand from foreign buyers physically located in China. In addition, we control for other factors affecting prices by matching exposed and unexposed U.S. ZIP codes on relevant observable characteristics.

Other work has used VAR analysis to show that capital inflows put upward pressure on housing prices by loosening financial conditions in the recipient economy (Sa et al., 2014; Cesa-Bianchi et al., 2015), or employed regression analysis proxying foreign inflows with the recipient's country's current account (Aizenman and Jinjarak, 2014) , we provide evidence that flows going directly to the purchase of residential real estate also contribute to price increases.

Second, we contribute to the literature on out-of-town buyers. Previous work has shown that out of town buyers have significantly increased housing prices in specific U.S. cities (Chinco and Mayer, 2016; Favilukis and Van Nieuwerburgh, 2018). However, previous research has not separated identified the effect of foreign buyers on prices.

Third and finally, we add to the literature on measuring the foreign assets and liabilities of countries. The existence of a global statistical discrepancy between estimated foreign assets and liabilities has been well-known (Lane and Milesi-Ferretti, 2007), and existing work has found that the main source of the discrepancy had been the underreporting of assets held in portfolio equity and debt in 2004, the last year of data in Lane and Milesi-Ferretti (2007). Particularly, the underreporting of U.S. assets held by foreign countries accounted for about half of the global statistical discrepancy. Studies that focus on the U.S. statistical discrepancy have found that there is substantial net unrecorded inflows from both financial derivatives and residential real estate (Curcuru et al., 2009), and that unmeasured capital flows tend to rise during times of global stress events (Flatness et al., 2009). In particular, Curcuru et al. (2009) find that the inclusion of residential real estate, which should be classified as a FDI flow, would substantially increase the net U.S. international investment liability position. Our findings suggest that cross-border purchases of U.S. residential real estate, which had previously been of a modest size, has grown rapidly in the past decade because of shocks from China and the global safe asset attribute of this asset class, and it is now a substantial source of the U.S. statistical discrepancy.

This paper proceeds as follows. In the next section, we carefully examine data on capital flows between China and the U.S. and find patterns consistent with purchases of U.S. residential real estate by residents in China. In Section 3, we describe the data we use to identify cross-sectional variation in demand for U.S. housing from China and outline our matching methodology, then present our micro-level results. In Section 4, we assess how our micro-level estimates relate to macro-level measures of capital inflows. Section 5 concludes.

2 Motivating Facts

A variety of circumstantial evidence points to substantial capital flows originating in China to the U.S. residential real estate market, particularly during episodes of deteriorating macroeconomic conditions in China over the past decade. In this section, we describe this evidence in detail.

A commonly referenced source for data on foreign purchases of residential U.S. real estate is the U.S. National Association of Realtors (NAR), which publishes an annual estimate of foreign purchase of U.S. residential real estate based on a voluntary survey of realtors. The NAR estimates that purchases by Chinese nationals increased from $11 billion in 2010 (17 percent of all foreign purchases) to $30 billion in 2018 (25 percent). A drawback to this data is that the voluntary survey typically has a low response rate (3 percent in 2016), and the survey does not guarantee that respondents use a uniform definition of what constitutes a foreign buyer or a uniform methodology for assessing what country he or she is from (e.g. country of residence versus nationality). Thus, while the NAR data are consistent with media coverage in indicating a substantial increase in purchases of residential real estate by Chinese buyers, they are at best an imprecise measure—albeit the only available measure, hence its popular use.

Aggregate capital flows data from balance of payment (BOP) accounts offer additional clues to the size of foreign purchases of U.S. residential real estate. We begin by noting a striking increase since 2010 in the comovement of private capital outflows from China and missing net capital inflows to the United States as captured by the statistical discrepancy line in the U.S. BOP. Figure 1 shows that the 12-quarter rolling correlation between the two variables went from being zero or negative to above 0.8 and remained elevated from 2012 to 2017.[2] A positive statistical discrepancy indicates that there are capital inflows that are not being captured in official statistics. While the most recent IMF Balance of Payments Manual (BPM6) specifies that foreign purchases of residential real estate should be included in foreign direct investment, the U.S. and indeed many countries do not measure such flows. As a result, these flows are captured only in the residual statistical discrepancy line of the balance of payments. In the case of the U.S. BOP, the statistical discrepancy primarily reflect two missing assets: financial derivatives, and the object of interest for our study—foreign purchases of U.S. residential real estate assets. We next explore bilateral capital flows data between the U.S. and China at a more granular level.

Large cross-border transactions, particularly for real estate purchases, oftentimes involve transactions between a foreign bank and a U.S. bank. So naturally, one might ask: how do recorded banking outflows from China comove with U.S. data on flows into U.S. banks from China? The left panel Figure 2 plots gross outflows via money and deposits as reported in Chinese balance of payment statistics, along with recorded flows into deposits at U.S. financial institutions from China and Hong Kong, obtained from the U.S. Treasury International Capital (TIC) System.[3] We include flows from Hong Kong because of the widely documented practice of Chinese households'

[2] The statistical discrepancy is calculated as the difference between the U.S. measured current account deficit and the measured financial inflows from abroad that finance that deficit.

[3] Like many countries, China does not publish bilateral capital flows data. But the the U.S. does, via the TIC System.

Figure 1: Rolling Correlation: Net private capital outflows from China and the U.S. statistical discrepancy

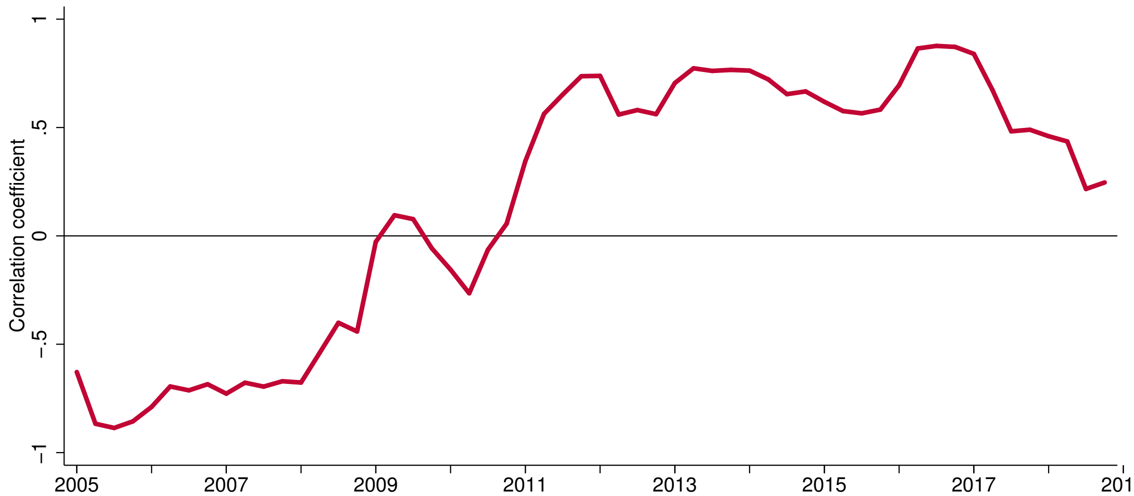

Source: Haver, authors' calculations. Graph plots correlation between 4−quarters rolling sum of net private outflows from China and US statistical discrepancy.

using banks in Hong Kong as a conduit when moving funds abroad.

Panel (a) of Figure 2 shows a striking degree of comovement between total Chinese deposit outflows and the pattern of Chinese deposit inflows to the United States. Most notably, outflows from China via banks and bank inflows to the United States from China both spiked during the two recent periods of concern about a so-called hard landing for the Chinese economy, first in 2011-2013 and again in 2014-2016. This comovement suggests that residents in China shifting money abroad place a substantial share into the U.S. banking system. At the peak of the first episode, in the fourth quarter of 2011, inflows to the United States accounted for 29 percent of total Chinese money and deposit outflows. And at the height of the second episode, when China unexpectedly devalued its currency in the third quarter of 2015, flows from China and Hong Kong into U.S. deposits accounted for 48 percent of total Chinese money and deposit outflows.

To verify that the comovement observed in panel (a) of Figure 2 is not simply a reflection of a high degree of banking integration between the U.S. and China, we examined the correlation between foreign banking outflow from other countries and bilateral banking inflows to the U.S. from those same countries. The results of this exercise, which can be found in the Appendix (Figure A1 and also Table A1) confirm that the comovement we observe for Chinese flows is not the norm. Rather, from 2010 onward, we observe an unusually close relationship between bank flows out of

Figure 2: Outflows from China and U.S. inflows

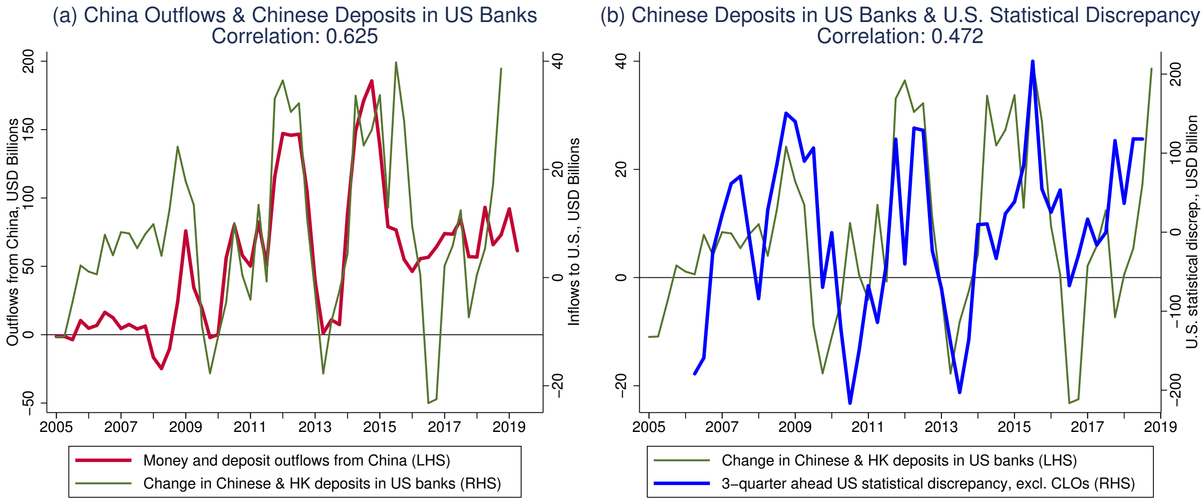

Source: Haver, TIC System, and authors' calculations. Series are 4−quarter sums.

China to the rest of the world and bank flows from China into the U.S.

We next examine the relationship between banking inflows to the U.S. from China and the U.S. statistical discrepancy, and find further evidence indicating that inflows from China are being used to purchase residential real estate in the U.S. As we discussed previously, foreign purchases of residential real estate are one factor known to contribute to the statistical discrepancy. Another factor known to have systematically contributed to the U.S. statistical discrepancy over the past fifteen years is the purchase of U.S. loans by offshore entities set up to issue collateralized loan obligations (CLOs).[4] Using methodology detailed in Liu and Schmidt-Eisenlohr (2019), we construct estimates of the contribution of CLOs to the statistical discrepancy and subtract them out. The resulting adjusted series is shown by the blue line in panel (b) of Figure 2, and makes clear that CLOs were not the major driver of the large positive values of the statistical discrepancy observed in recent years. The size of the statistical discrepancy remaining after stripping out the effects of offshore issuance of CLOs suggests that the NAR estimates for the size of foreign purchases of U.S. residential real estate are likely on the low side.[5]

[4]For more details on this see BEA (2019), Guse et al. (2019), and Liu and Schmidt-Eisenlohr (2019).

[5]The statistical discrepancy measures residual net capital inflows, not gross flows, so it cannot be directly interpreted as foreign purchases of U.S. real estate. However, Curcuru et al. (2009) estimates that U.S. has consistently experienced net capital inflows in residential real estate over the period 1989-2007, with inflows in this asset class much larger and rising much more rapidly than outflows. As a result, the variation in net real estate flows is almost entirely driven by changes in the gross real estate inflows. As of 2007, U.S. liabilities in residential real estate was $798 billion, compared with the $198 billion in U.S. claims on residential real estate abroad. If the remainder net sta-

In normal times, the U.S. statistical discrepancy is small in size and has an average value of zero.[6] But the discrepancy has historically registered sizeable positive values during bouts of international financial turmoil such as the Asian Financial Crisis and (as seen in panel (b) of Figure 2) the Global Financial Crisis, due to unrecorded safe haven flows into the United States (Flatness et al., 2009). It is therefore notable that not only do banking inflows from China (the green line in Figure 2) peak during the 2011-13 and 2014-16 periods of heightened concern about a hard landing in China, but the U.S. statistical discrepancy became large and positive as well, despite those not being periods of *global* financial stress.

Panel (b) of Figure 2 show inflows from China to the U.S. Banking system along with the the three-quarter-ahead value of the U.S. statistical discrepancy. Figure 2 makes clear that the U.S. statistical discrepancy peaked three quarters after bank inflows from China and Hong Kong during the two episodes of economic distress in China during the period we are studying. Conversely, the statistical discrepancy dropped to zero three quarters after banking inflows from China and Hong Kong dropped to their lowest level ever, in the third quarter of 2016. In Appendix Figure A2, we show that the correlation between the two series demonstrate a strong positive value starting at a lag of two quarters, peaking at a lag of three quarters.

What is the significance of the three quarter lag in the strong relationship between banking inflows from China and the U.S. statistical discrepancy? In fact, it is further suggestive of substantial inflows of Chinese capital to the U.S. residential real estate market. This is because bank transactions involving foreigners are recorded in the U.S. balance of payments while real estate transactions are not. Consider an example in which a foreign resident moves money into a U.S. bank to purchase a house in the United States. When the foreign resident deposits money in a U.S. bank, the bank reports an increase in its liabilities to foreigners, which shows up as capital inflow from abroad. Six to nine months later, when the same foreign resident takes the money out to purchase a house, the bank reports a drop in its liabilities to foreigners, generating a capital outflow in official statistics. The foreigner has purchased a claim on a U.S. asset (the house), which is technically an inflow of direct investment from abroad, but which in practice is not recorded

tistical discrepancy after exclusion of CLOs does purely reflect net residential real estate flows, then the size of gross inflows, or gross foreign purchases of U.S. residential real estate purchases, would be larger than the net statistical discrepancy ex-CLOs shown in panel (b) of Figure 2.

[6] Although these flows do imply a large net position, as demonstrated by Curcuru et al. (2009)

in the balance of payments. This unrecorded FDI inflow adds to the U.S. statistical discrepancy, pushing it upwards. The three quarter lag in the relationship is consistent with foreign residents depositing funds in U.S. banks and then taking between six and nine months to find a house to buy and settle the resulting real estate transaction, a very plausible timeframe.

To more formally establish the connection between U.S. missing inflows and Chinese capital outflows, we regress the four quarter moving average of U.S. statistical discrepancy and three China-specific variables that proxy for shocks: Chinese FX reserve sales, net Chinese money and deposits outflows, and changes in the Chinese macro climate, measured by the coincident macro climate index published by the Chinese National Bureau of Statistics (NBS). Taking into account the lagged relationship evident in Figures 1 and 2, we lag these explanatory variables three quarters.

Because Figure 1 indicated that the relationship between Chinese flows and U.S. inflows changed dramatically in 2010, in the regressions we allow the coefficient on the Chinese variables to vary over time. Specifically, we create dummy variables for pre- and post- 2010Q2 periods and interact them with each China variable. Additionally, in the post-2010Q2 period, we allow the coefficient on the China variables to vary depending on whether it represents a positive or negative signal regarding the outlook for the Chinese economy. Net foreign exchange reserve sales, for example, would be a negative signal, indicating that the authorities are intervening against currency depreciation pressure emanating from private market participants. Conversely, net foreign exchange reserve purchases would suggest intervention to dampen appreciation due to net capital inflows to China. Finally, we include the year-on-year log change in the VIX to control for global financial conditions more generally.[7]

The results, shown in Table 1, further confirm that negative (positive) shocks from China are associated with an increase (decrease) in U.S. statistical discrepancy net inflows since 2010Q2, with a three quarter lag. The pre-2010 China shocks are not important in explaining the safe haven flows, as none of these Chinese variables are significant when interacted with the pre-2010Q2 dummy. The VIX, our measure of global financial conditions, is insignificant across all specifications, suggesting that unrecorded capital inflows to the U.S. are better explained by Chinese factors than by global financial conditions more generally. Strikingly, the regressions using Chinese capital

[7] The VIX measures implied volatility on options to buy the S&P500 U.S. equity index and is widely viewed as an indicator of global risk sentiment.

Table 1: Relationship Between U.S. Statistical Discrepancy and Economic Conditions in China

	US_stat_discrep		
	(1)	(2)	(3)
China FX Rsv sales x post-2010Q2(-)	0.264		
	(0.134)		
China FX Rsv sales x post-2010Q2(+)	-0.212**		
	(0.0725)		
China FX Rsv sales x pre-2010Q2	-0.0524		
	(0.0851)		
China deps outflows x post-2010Q2(-)		0.554**	
		(0.205)	
China deps outflows x post-2010Q2(+)		-2.474*	
		(0.961)	
China deps outflows x pre-2010Q2		0.358	
		(0.808)	
China macro chg x post-2010Q2(-)			6.852**
			(2.006)
China macro chg x post-2010Q2(+)			-4.143*
			(1.592)
China macro chg x pre-2010Q2			1.818
			(1.302)
dlnVIX	-8.020	-0.165	-5.246
	(8.128)	(9.806)	(7.733)
Observations	66	66	65
R^2	0.274	0.222	0.288

Standard errors in parentheses
Note: Positive coefficients correspond to net positive inflows in U.S. statistical discrepancy. All flows variables are four-quarter rolling sums. Macroeconomic conditions index is measured as the year-on-year percent change. VIX is measured as the year-on-year change in the log. All Chinese variables are lagged three quarters.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

outflows variables have substantial explanatory power: the R-square is 0.42 for the net foreign reserves sales regressions, and 0.28 for the regressions with Chinese money and deposit outflows.

A back of the envelope calculation based on these results suggests that each $1 billion in net reserves sales by the Chinese authorities since 2010Q2 was associated with a $0.3 billion increase in unrecorded capital inflows to U.S. For each $1 billion in net money and deposit outflows from China, there was a $0.7 billion increase.

In this section, we have presented a variety of indirect evidence that since 2010, Chinese residents have shifted money into the U.S. which has been used to purchase residential real estate. We had shown that capital outflows from China have made their way into the U.S. banking system and also—with a lag—contributed to the large and positive unrecorded inflows seen over the period. The time structure of the relationships we have uncovered is consistent with capital inflows being used to purchase real estate. In the next section, we lay out empirical methodology for rigorously testing the impact of these flows on the U.S. residential real estate market at the micro level.

3 The Effect of Chinese Demand on U.S. House Prices

Our goal is to estimate the causal impact of Chinese inflows on the house price growth in the U.S. residential real estate market. The identification challenge is to be able to disentangle the impact of a foreign demand shock from other factors affecting house prices, such as U.S. domestic economic conditions like mortgage rates, as well as unobserved local factors. To overcome these challenges, we adopt a difference-in-difference matching framework. The general strategy is to compare house price growth in areas heavily exposed to foreign Chinese buyers to that in geographically proximate areas that have low exposure to Chinese demand, but which are otherwise similar in other attributes that matter for house price growth. In the language of the matching literature our outcome variable is house price growth, the areas exposed to Chinese capital inflows are the treatment group, and areas not of interest to Chinese buyers make up the control group.

Three assumptions are central to identification in the matching methodology we adopt: the conditional independence assumption, the overlap assumption, and the independent and identically distributed assumption (Caliendo and Kopeinig, 2008). By using house price *growth* as the outcome variable, we difference out the time-invariant unobservable factors that affect the *level* of house prices of the neighborhoods we study. And by both conditioning the matches on observables that explain Chinese demand and restricting the geographic proximity of the control areas, we minimize selection bias, including bias generated by time-variant unobserved common local factors, which could jointly affect the treatment and outcome variables. The resulting difference in the house price growth of the treated and the matched control neighborhoods, we argue, produces an unbiased estimate of the causal effect of Chinese demand on local house prices in the U.S.

This section first describes in detail the dataset we use to measure exposure to Chinese demand. After outlining our matching design, we compute the gap in house price growth between the treated and control ZIP codes.

3.1 Data

To measure the exposure of an area's residential real estate market to Chinese capital inflows, we make use of a unique dataset that allows us to directly measure Chinese demand for residential real estate at the micro level. We obtained web traffic data from an real estate listing website called Juwai, which is based in China and caters to individuals resident in China who are looking to purchase residential property abroad (the name of the company translates to "living abroad"). The Juwai dataset provides us with the number of views of properties in each U.S. ZIP code, each month, from each Chinese city. In other words, Juwai provides us with data at the city-pair month level. We have three months of property views data from November 2016 to January 2017. Over these three months, there were 670,000 total views originating from China of U.S. real estate properties across more than 7000 ZIP codes, or 917 core-based statistical areas (CBSAs).[8] The geographic dispersion of these views can be seen in Figure 3, with 70 percent of the views concentrated in just 20 U.S. cities, of which about a third are in California and almost 20 percent in the Greater Los Angeles area. The top 20 U.S. cities that received the most real estate property views from China are listed in Appendix Table A2 .

To our knowledge, no other researcher has compiled a comparable dataset which *directly* measures Chinese demand at the local level for the entire United States. Nonetheless, we validate our data by checking the relationship between Juwai property views and two variables that ex ante

[8] A CBSA is defined by the U.S. Office of Management and Budget as a geographic area that consists of one or more counties containing an urban center of at least 10,000 people and adjacent countries that are socioeconomically tied to the the urban center via commuting. In the rest of the paper, we will use the terms CBSA and city interchangeably.

Figure 3: Web Traffic Hits from China of U.S. Listed Properties, by CBSA

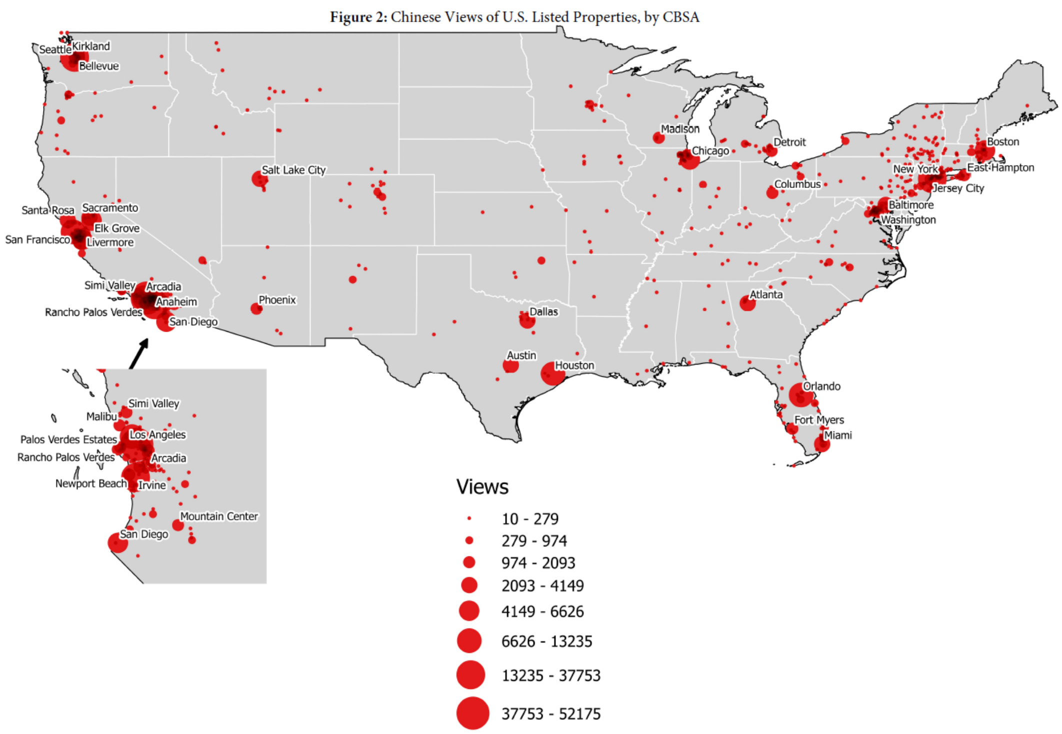

we expect to correlate with Chinese purchases of residential real estate: Airport passenger arrivals from China and the share of real estate transactions done in cash. While it is certainly possible to purchase a property from abroad without ever visiting it in person, discussions with representatives of Juwai indicated that most buyers do indeed come to the U.S. In addition to providing information about an overseas real estate property on their website, Juwai also offers consulting services for the potential buyer and helps refer the potential buyer to a real estate brokerage firm abroad. According to the Juwai CEO, potential buyers often make one visit to the city in which their property of interest is located and make a purchase within six months. We validate the Juwai data by pairing the passenger arrival data from a Chinese city to a U.S. city. Figure 4 shows that the city-pair arrival data has a 0.54 correlation with the Chinese city-U.S. city pair Juwai data.

To further confirm that the Juwai views data do meaningfully capture Chinese demand for residential real estate in individual cities and ZIP codes, we also checked the relationship between number of Juwai views in each ZIP code and the share of residential real estate transactions done in cash for the same ZIP code. Given that foreign buyers often prefer the more expedient option of settling real estate transactions with cash, we would expect that areas where Chinese residents

Figure 4: Juwai views vs. airline passenger arrivals

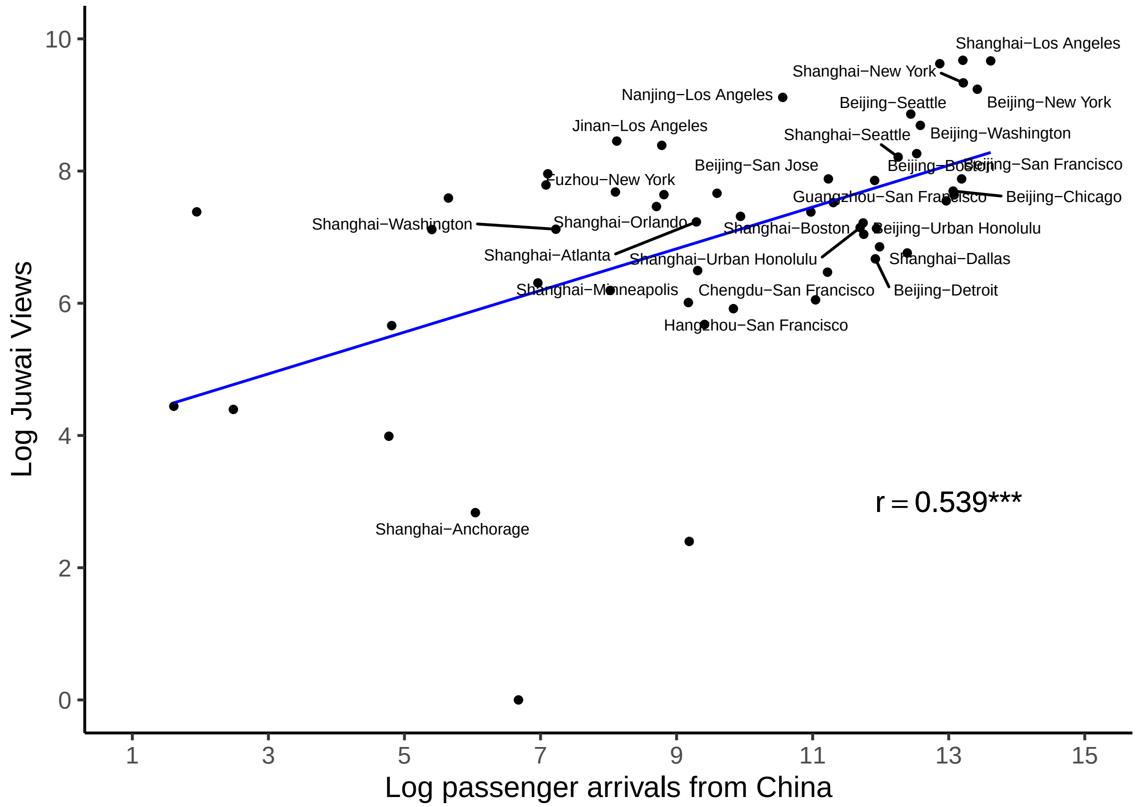

Source: Juwai and FAA.

purchase more properties would also have a higher share of sales in cash. And indeed, the correlation between views and the share (in value terms) of purchases done in cash is 0.157 (statistically significant at the one percent level). Figure 5 plots the significant and positive relationship between views and the cash purchases share. In Appendix Figure A3, we show this relationship holds even more strongly for major U.S. cities, including Seattle (correlation 0.34), Washington DC (0.34), Los Angeles (correlation 0.324), and New York (0.24).

3.2 Matching Design and the Drivers of Chinese Demand

A necessary condition for unbiased estimates of matching is that the potential outcome variable be conditionally mean independent of the treatment. Put differently, unbiasedness requires that we control for factors that affect both house price growth and foreign buyer demand. The most likely confounding factor in this context is past price growth. Given that many Chinese buyers are purchasing houses in the U.S. as an investment, it seems reasonable to expect that they would

Figure 5: Juwai views vs. cash sales share

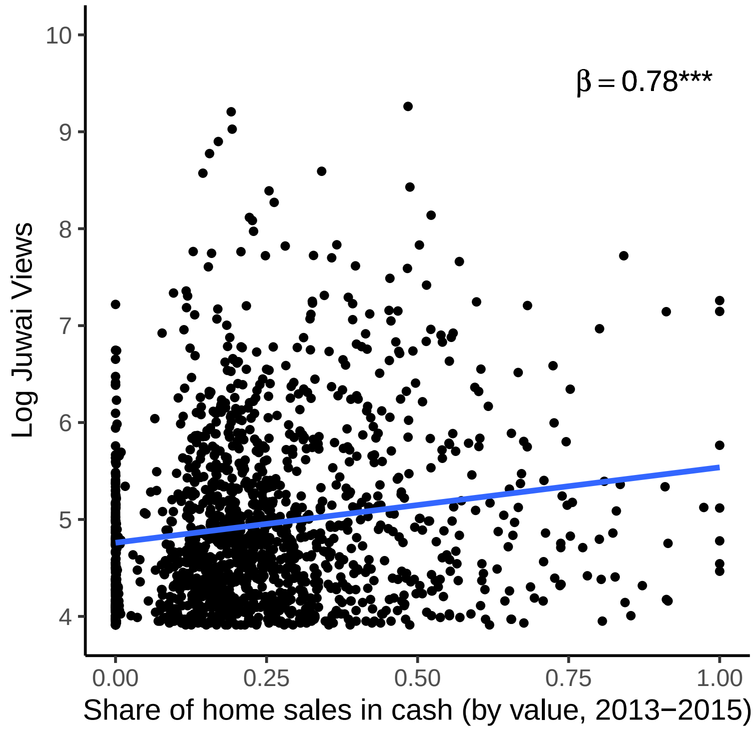

Sources: Juwai and CoreLogic.

self-select into local housing markets that tend to experience greater house price appreciation. Such a selection effect would upwardly bias any estimate of the price effect of foreign Chinese buyers. In fact, we will show in this section that once we use fixed effects control for unobserved factors at the city level, house prices in the zip codes that attract the attention of Chinese buyers do not tend to historically appreciate faster than those that do not.

To gauge the severity of the selection bias, we ask the question: Do foreign Chinese buyers tend to systematically buy from places that have experienced more rapid price appreciation? To answer the question, we estimate a cross-sectional regression model of determinants of foreign Chinese buyer demand, both at the CBSA and ZIP code level. The dependent variable is the log number of Juwai views in the geographic area. Our choice of explanatory variables is guided by a survey of potential Chinese buyers carried out by Juwai, the company from which we obtained our property views data. Users cited the following motivations for overseas property purchases: 46 percent said "lifestyle," 42 percent investment, 64 percent emigration, and 83 percent education. Accordingly, we included explanatory variables that reflect these objectives, and added variables that capture the accessibility of the location to a foreigner arriving from China. The following is the specification for the cross section regression at the CBSA level, where i denotes a CBSA:

$$\begin{aligned}\Delta ln\_views_i = \alpha &+ \beta_1 chinese\_share\_init_i + \beta_2 dist\_to\_china_i + \beta_3 univ_i \\ &+ \beta_4 ln\_pop\_init_i + \beta_5 ln\_med\_price\_init_i \\ &+ \gamma_1 temp_i + \gamma_2 unemp_i + \gamma_3 commute\_time_i + \delta hist\_apprec_i + u_i\end{aligned} \tag{1}$$

The initial population size ($ln\_pop\_init_i$) is included because foreign buyers tend to be drawn to larger cities, with the wealth of amenities that a large city can offer. The initial median house price ($ln\_med\_price\_init_i$) is included because various studies hypothesized that Chinese buyers are more drawn to the luxury segment of the local housing market. We measure the accessibility of an area to Chinese buyers with two variables: the initial share of Chinese in local population ($chinese\_Share\_init_i$), and the (log) distance from the nearest airport with top passengers arrivals from China ($dist\_to\_china_i$). We include a variable that captures the average historical annual house price appreciation to reflect the investment motive of foreign buyers (the average price appreciation between 2001 to 2006, a period when Chinese or foreign capital inflows play little role in the U.S. housing boom; we also separately try the average appreciation for 2001 to 2010. We also include a variable that captures access to universities (number of universities or distance to the closest university), to reflect the commonly cited reason of purchasing houses close to the university attended by their offspring. Not only might obtaining housing for offspring studying at a U.S. university motivate Chinese parents to purchase property, but having a child studying in the U.S. university makes Chinese households eligible for a much greater annual quota of U.S. dollar purchases than is otherwise the case. In addition, we include other variables that commonly motivate local buyers to purchase a house, such as mean commute time, temperature, and unemployment rate (only available at the CBSA level).

We also regress Chinese property views on our set of determinants at the ZIP code level. For comparison with our CBSA-level regressions, we first include the three variables which are only available at the CBSA level (temperature, unemployment rate, and commute time):

$$\begin{aligned}\Delta ln\_views_z = \alpha &+ \beta_1 chinese\_share\_init_z + \beta_2 dist\_to\_china_z + \beta_3 univ_z \\ &+ \beta_4 ln\_pop\_init_z + \beta_5 ln\_med\_price\_init_z \\ &+ \gamma_1 temp_i + \gamma_2 unemp_i + \gamma_3 commute\_time_i + \delta hist\_apprec_i + u_i\end{aligned} \tag{2}$$

However, we then run a specification which drops variables available only at the CBSA label (temperature, unemployment) and instead include CBSA fixed effects (θ_i).

$$\begin{aligned}\Delta ln\_views_z = \alpha &+ \beta_1 chinese\_share\_init_z + \beta_2 dist\_to\_china_z + \beta_3 univ_z \\ &+ \beta_4 ln\_pop\_init_z + \beta_5 ln\_med\_price\_init_z \\ &+ \delta hist\_apprec_s + \theta_i + u_i\end{aligned} \tag{3}$$

The CBSA-specific fixed effects should control for other determinants of house prices found in the literature, such as changes in construction or wage costs or supply elasticities (Saiz, 2010), both of which have been measured at the CBSA-level in existing literature. We separately estimate the regression for up to 624 CBSAs (in all 50 states and the District of Columbia) where data were available, as well as for a trimmed sample of 224 CBSAs which had at least 50 views (located in 40 states).

Table 2 displays the results for the CBSA-level regressions. The variables included explain the pattern of Chinese housing search pattern well, with an adjusted R-squared of 0.7 to 0.8, and confirm a number of existing hypotheses about the motivation of foreign Chinese buyers. Large U.S. cities, cities with high existing share of ethnic Chinese population, and markets with higher median housing prices tend to draw foreign Chinese buyers. These findings are consistent with those in existing literature (e.g. Badarinza and Ramadorai, 2018, IMF GFSR, 2019). We also find that places with more universities and those that are located closer to airports with direct flights to China also have more Juwai hits, though these relationships are not consistently significant. Factors important to domestic buyers—commute time, unemployment—are not important in motivating potential Chinese buyers.

Most notably, average house appreciation during the 2001-2006 housing boom period is consis-

Table 2: Determinants of Demand from Foreign Chinese Buyers – CBSA-Level Regressions

	Full Sample				CBSAs with > 50 views			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Initial Chinese share	0.293***	0.305***	0.291***	0.258***	0.229***	0.225***	0.244***	0.213***
	(0.0599)	(0.0601)	(0.0620)	(0.0623)	(0.0592)	(0.0605)	(0.0625)	(0.0627)
Distance to China	-0.00604	0.0209	-0.0193	-0.00689	-0.119	-0.130*	-0.146*	-0.120
	(0.0493)	(0.0524)	(0.0527)	(0.0538)	(0.0615)	(0.0654)	(0.0648)	(0.0656)
Number of colleges	0.00804	0.00433	0.00101	-0.000144	0.0145*	0.0146*	0.00792	0.00749
	(0.00633)	(0.00642)	(0.00633)	(0.00642)	(0.00642)	(0.00682)	(0.00681)	(0.00701)
Population	1.023***	1.050***	1.107***	1.109***	0.711***	0.727***	0.846***	0.840***
	(0.0350)	(0.0366)	(0.0378)	(0.0385)	(0.0604)	(0.0664)	(0.0699)	(0.0728)
Initital median home price	0.807***	0.762***	0.469***	0.710***	0.430***	0.452***	0.208	0.431**
	(0.0908)	(0.0930)	(0.118)	(0.102)	(0.125)	(0.130)	(0.157)	(0.138)
Average temperature		-0.0132**	-0.0174***	-0.0148**		-0.00299	-0.0114	-0.00584
		(0.00462)	(0.00478)	(0.00479)		(0.00666)	(0.00685)	(0.00668)
Initial unemployment rate		0.00300	-0.0214	0.00201		0.0112	-0.0213	0.0197
		(0.0138)	(0.0159)	(0.0154)		(0.0272)	(0.0296)	(0.0274)
Initial average commute		0.0149	0.0155	0.0177		-0.0117	-0.00925	-0.00112
		(0.0102)	(0.0106)	(0.0108)		(0.0181)	(0.0186)	(0.0187)
Ave. Δ home price, pre-crisis			0.0404***				0.0452**	
			(0.00958)				(0.0154)	
Ave. Δ home price, pre-2010				0.0271				0.0276
				(0.0164)				(0.0292)
Observations	624	624	556	556	231	231	224	224
R^2	0.803	0.806	0.824	0.819	0.711	0.712	0.735	0.725
Adjusted R^2	0.802	0.804	0.821	0.816	0.705	0.701	0.724	0.714

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

tently positive and significant (columns 3 and 7). That the coefficient on this variable is positive and significant confirms that omitted variable bias and reverse causality are potential concerns: Even when we control for other relevant covariates, house prices in the places that foreign Chinese buyers desire tend to historically appreciate faster than those that do not, at least at the city level. At the same time, when we include as a regressor average house price appreciation from 2001 to 2010, the variable is not significant. For this reason, in what follows we control only for pre-crisis price changes, rather than price changes up to the start of the period when Chinese capital flows to the U.S. picked up.

The results of our ZIP code-level regressions suggest that the unobserved factors that drive house price appreciation are mostly at the broader city level. As Table 3 shows, the signs on the coefficients of the same variables that explain the city level regressions well are similar. The coefficient on the historical house price appreciation in the period 2001-2006 is also positive and significant. However, the significance disappears once we included CBSA-specific fixed effects.

Table 3: Determinants of Demand from Foreign Chinese Buyers – ZIP-Level Regressions

	Dependent Variable: Log Juwai views			
	(1)	(2)	(3)	(4)
Initial Chinese share	0.0884***	0.0903***	0.0655***	0.0698***
	(0.0211)	(0.0214)	(0.0149)	(0.0161)
Distance to China	-0.0856*	-0.0792*	-2.832***	-3.039***
	(0.0359)	(0.0381)	(0.188)	(0.248)
Distance to nearest college	-0.227***	-0.215***	-0.279***	-0.268***
	(0.0336)	(0.0308)	(0.0292)	(0.0278)
Population	0.353***	0.427***	0.380***	0.425***
	(0.0262)	(0.0292)	(0.0229)	(0.0261)
Inital Median HH Income	0.531***	0.400**	0.194*	0.221*
	(0.106)	(0.128)	(0.0915)	(0.0916)
Average temperature	0.00835	0.000181		
	(0.00441)	(0.00453)		
Initial unemployment rate	0.0260*	0.00985		
	(0.0112)	(0.0121)		
Initial average commute	-0.00350	-0.0124		
	(0.00558)	(0.00637)		
Ave. Δ home price, pre-crisis		0.0433***		0.0146
		(0.00695)		(0.0144)
Observations	8139	6571	8142	6572
R^2	0.224	0.232	0.383	0.378
Adjusted R^2	0.223	0.231	0.327	0.323
CBSA FE	No	No	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Another key observation from the ZIP code level regressions is that there is much greater randomization of foreign Chinese demand *within* a CBSA, which is favorable for satisfying the overlapping assumption critical for a well-designed matching framework. The adjusted R-squared for the same group of explanatory variables in the CBSA regressions decreased substantially to 0.2-0.3 in the ZIP code level regressions. The lower R-square allows for a larger pool of untreated ZIP codes that share overlapping characteristics with the treated ZIP codes.

Taken altogether, the city-level and ZIP code-level regressions appear to suggest that foreign Chinese buyers tend to select into a city where they would buy a house based on city-specific factors that could be drive up house prices, but once the city of their desire is chosen, they tend to randomize selection into a ZIP code. City-level unobserved effects likely present a positive selection bias on the treatment.

For these reasons, we propose two solutions to alleviate the potential bias. First, we conduct the matching at the ZIP code level. Second, the matching should be conducted *within* the same CBSA, such that the treatment effect on the treated would difference out the unobserved city-level factor. In the next section, we describe in details the implementation of our matching framework.

3.3 Match Results

We now formally describe our matching design. We conduct the matching in two ways. The first, informed by the discussion in the previous section, matched a treatment ZIP code with a control ZIP code within the same CBSA. The treatment group (henceforth referred to as "Treatment Indicator 1") are those ZIP codes which were in the top decile for Juwai views within their respective CBSAs, and the control are drawn from the ZIP codes below the 50th percentile for Juwai views within the same CBSA. But a problem with this definition is that some CBSAs contain only a limited number of ZIP codes (i.e. Santa Rosa in California), and therefore have limited overlapping support and was unable to produce any matches, even if we restricted the matches to cities that contain more than 30 ZIP codes. Furthermore, in small cities, the untreated would presumably be more likely to experience spillovers effects from the treated ZIPs, thus biasing our estimate of the outcome variable to the downside.

We thus devised an alternative treatment group "Treatment Indicator 2") that somewhat relaxes the constraints of the match. This second treatment group is defined as the top 5 percent of

Juwai views nationally, and the control group defined as those in the bottom 30 percent nationally. Instead of restricting the matches to be within the same CBSAs, we allowed the treated ZIPs to be matched with those that are located outside of the treated ZIP codes city, but with the restriction that the city must be of similar attractiveness to foreign Chinese buyers – in practice, these are cities whose aggregate Juwai views percentile are similarly ranked as those CBSAs in which the treated ZIP codes are located. This last detail is important because we are not allowing a match with just any ZIP codes in the country, but only relaxing the match criteria so as to cities that draw a similar level of interest from foreign Chinese buyers. In this way, we expand the population of potential matched control areas while still minimizing the city-specific unobserved heterogeneity that could bias our estimates of the treatment effect.[9]

Both definitions of treatment yielded about 370 treatment ZIP codes, across 20 CBSAs in the first definition and 33 CBSAs in the second definition. The first group of cities accounts for 43 percent of total U.S. employment and similarly by as much in terms of population, and the second group of cities accounts for 52 percent of total U.S. employment and of U.S. population. The average median home price, as of December 2016, is $550,000 for the first treatment group, and $680,000 for the second. In comparison, the National Association of Realtors (NAR) reported an average purchase price of $940,000 by foreign Chinese buyer in 2016. Our much lower median purchase price of Chinese buyers may suggest that the survey based collection method employed by NAR may have over-sampled the top-end buyers, which is not surprisingly given that those purchases are more high-profiled and garnered more attention from realtors.

The covariates on which the treated ZIP codes are matched to untreated ZIP codes within the same CBSA are: (1) population size in 2010, (2) percent of ethnic Chinese population in 2010, (3) log median house price in 2010, (4) distance from the nearest college, (5) average commute time in 2010, and (6) historical average house price appreciation over the period 2001-2006. For Treatment Indicator 2, the additional covariate is the percentile of Juwai views of the CBSA to which a ZIP code belongs to.

Each treated ZIP code is matched to five untreated ZIP codes, with replacement and using the nearest-neighbor algorithm, which minimizes the weighted sum across the differences of each of the

[9]With the second treatment definition, about 30 percent of the matched control ZIP codes belong to the same city as the treatment ZIP code.

Figure 6: Matched Co-variates for Treatment vs. Average Control Group

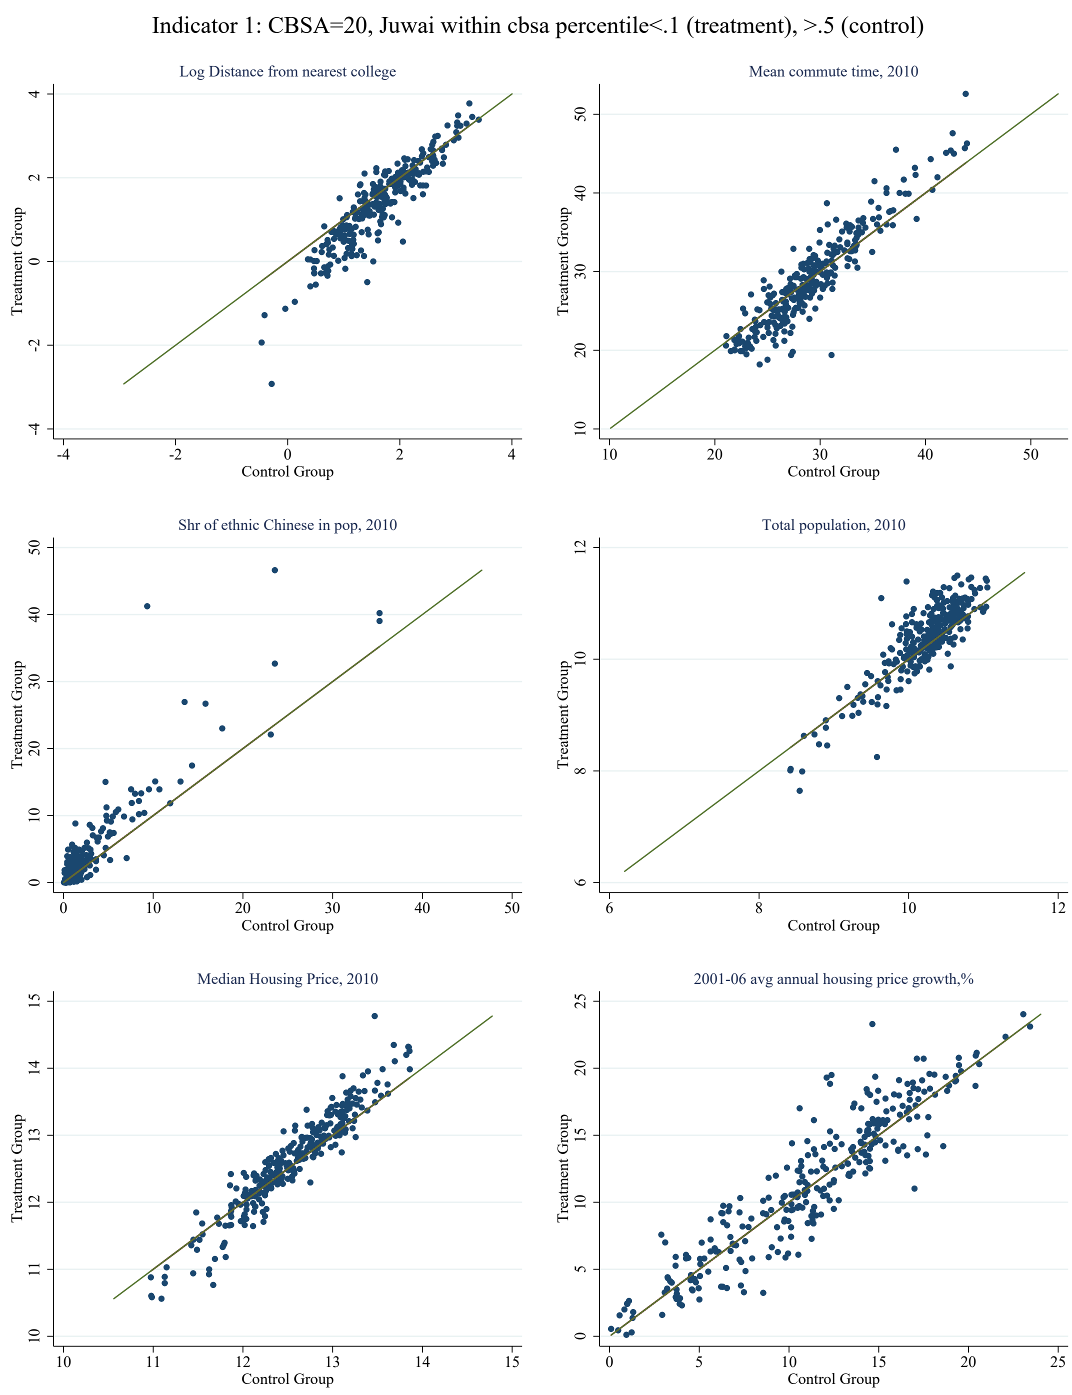

Note: The line represents a 45 degree line, i.e. points on the line have the same value for control and treatment group.

Figure 7: Matched Co-variates vs. Average Control Group

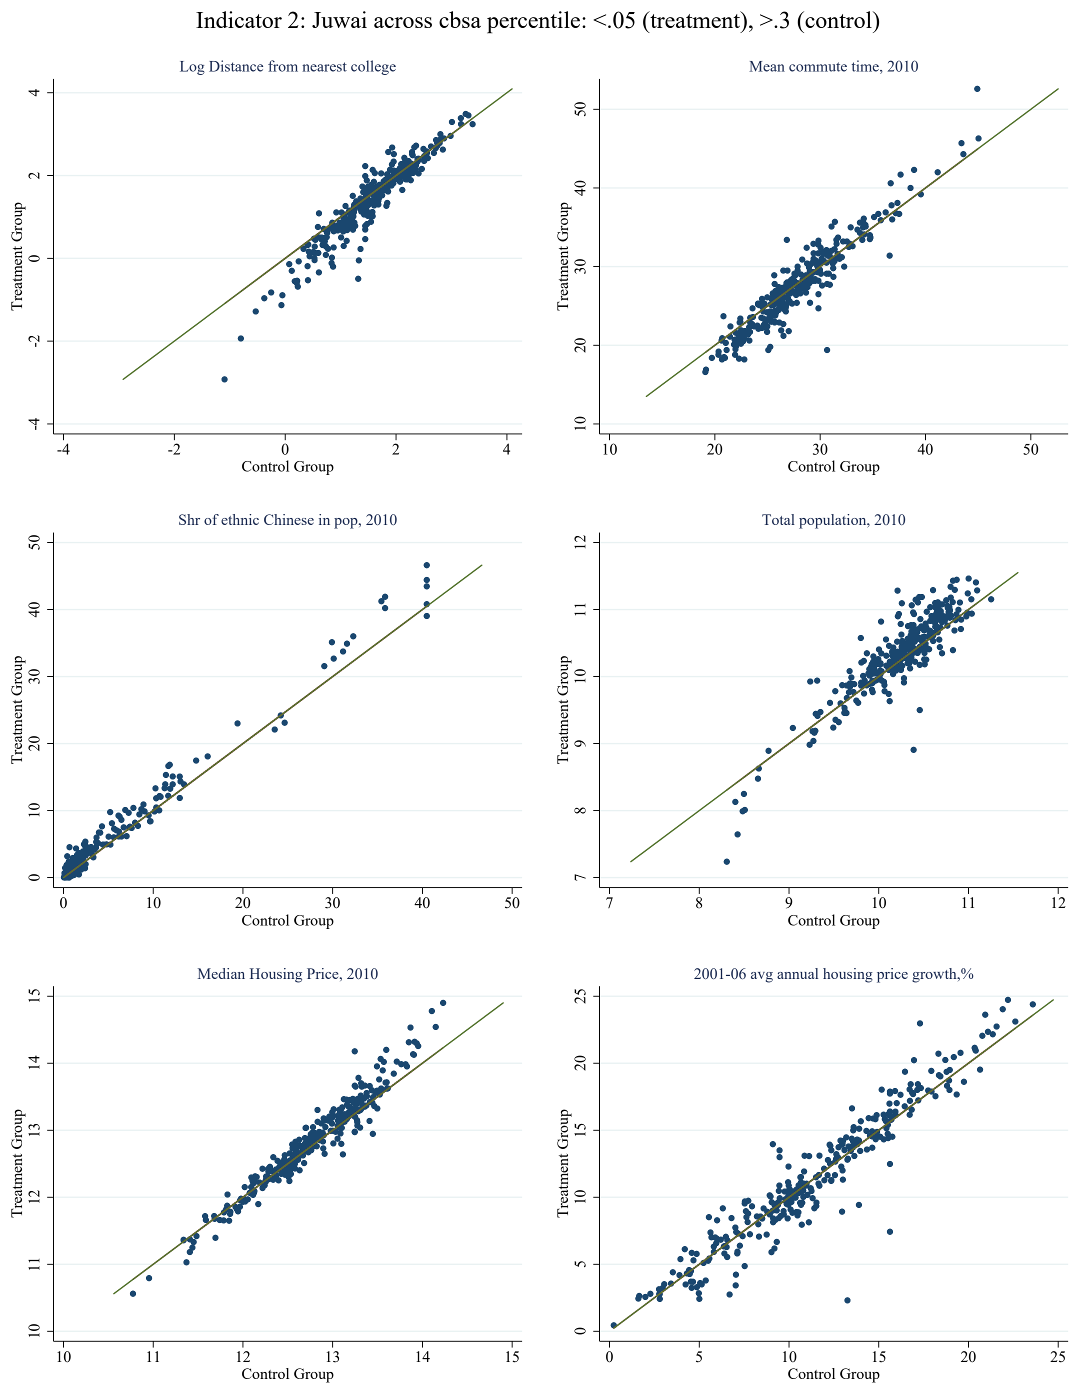

Note: The line represents a 45 degree line, i.e. points on the line have the same value for control and treatment group.

matching covariates. Figure 6 and Figure 7 compare each of the matched covariates of the treated ZIP code to the average of their matched control for Treatment Indicators 1 and 2, respectively. It shows that the values are similar between the treated and control group across each of these matched characteristics. The slope of these scatter plots are mostly statistically insignificant from one (except for share of ethnic Chinese and median home prices for both indicators, and the average 2001-2006 house price growth rate for indicator 1). From these plots, one can also see that there is a substantial overlapping support for the matches between the treated and the control, satisfying a crucial condition for the matching design. We also illustrate our matching process by providing the example of Seattle. (see Figures A4 in Appendix A)

3.4 Cumulative Impact: Average Treatment on the Treated

The outcome variable, or the object of our interest, is difference in house price growth between treated and untreated areas. Figure 8 shows the fitted kernel distribution of 6-year cumulative house price growth in the 2010-2016 period of the treated ZIPs and the matched control ZIPs. For comparison purpose, we also show the distributions of the house price growth for the 2000-2006 period, when the United States experienced a broad housing market boom. By design, the distribution of house price growth of the treated over this earlier period was to be similar to the matched ZIPs, confirmed by these plots. Indeed, for both definitions of treatment, a two-sample Kolmogoros-Smirnov test for equality of distribution found that they are statistically indistinguishable in this period. However, the distribution of house price growth in the period 2010-2016, when the China shock was present, had shifted right of the control ZIP codes. The rightward shift of the treated ZIPs in this later period is even more apparent using the second treatment definition. A test for equality of distribution confirms with high significance that the housing price appreciation of the treated group is larger than for the control group for both definitions of treatment.

The difference in the mean growth rates of the two distributions provides the average treatment effect on the treated (ATET) of exposure to foreign Chinese capital over the 6-year window. For treatment definition 1, the mean house prices over this period for the treatment grew 7 percent faster than for the control group, or 1.1 percent faster per year. For treatment definition 2, mean house price growth for the treatment group was 14 percent faster over the 6 year period, or 2.2 percent per year, compared to that of the control.

Figure 8: Distribution of 6-yr House Price Growth

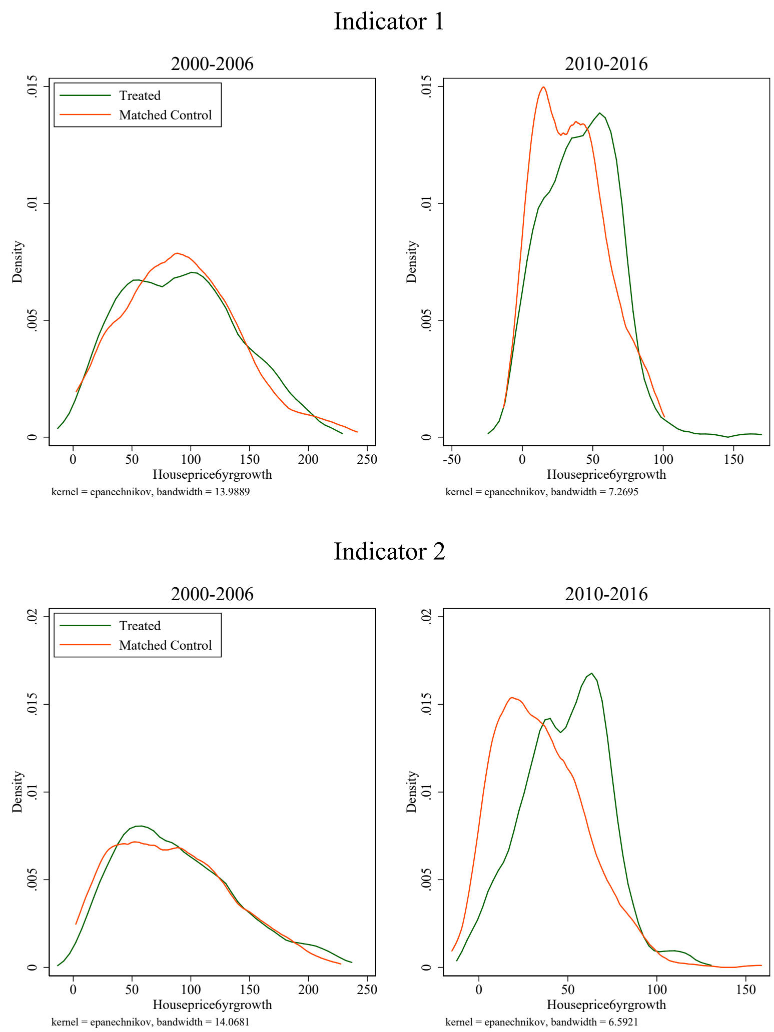

Another way to visualize the cumulative impact is look at the evolution of the house price index of the treatment and control group (Figure 9). Because of their differential growth rates, the house price level of the treatment group has diverged significantly from the control group in recent years. For the 20 cities included in the treatment definition 1, the divergence picked up after 2008, and for the 34 cities in the treatment group 2, the divergence picked up since 2010.

Figure 9: House Price Level (Weighted Avg, Index 2006m12=100)

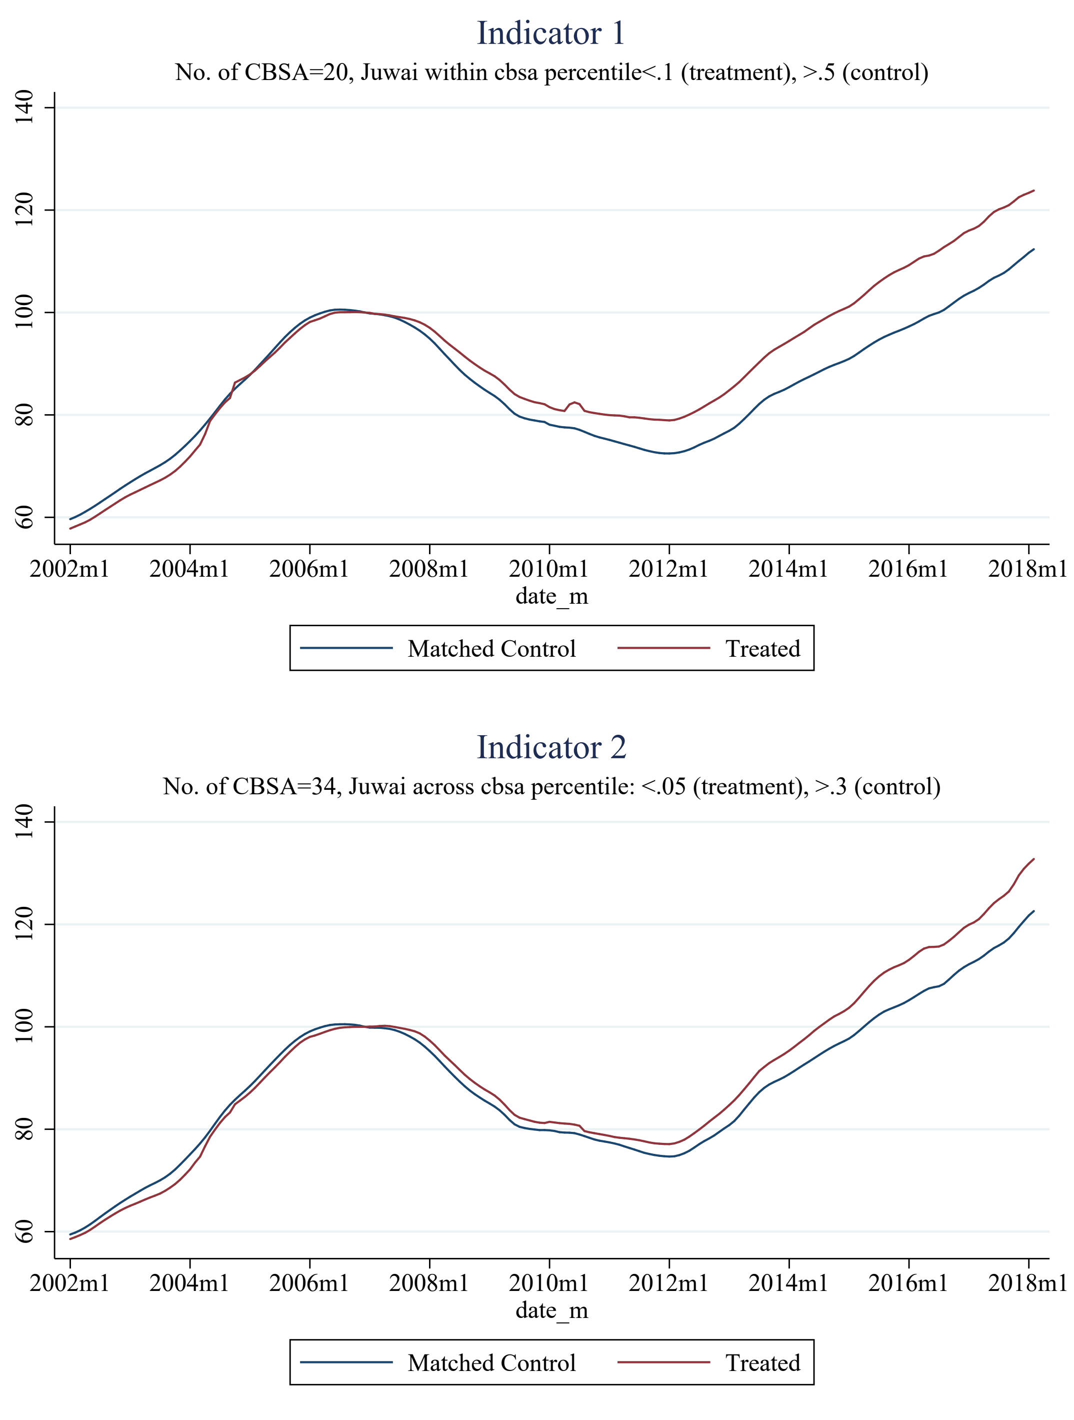

3.5 Evolution of the Average Treatment Effect on the Treated Over Time

Because there were multiple periods of China shocks since the Global Financial Crisis, we now explore the time evolution of the average treatment effect on the treated. We calculate the ATET as the weighted average of the difference in the 2-year house price growth between the treatment and control group. We calculated this on a rolling monthly basis. Figure 10 shows the time evolution of this variable for both Treatment indicator 1 and 2, with a 95 percent confidence interval around using the standard errors estimated by the nearest-neighbor matching procedure. Both treatment definitions produced three local peaks in the premium of house price growth of the treated over the control ZIPs: in 2009m1, 2012m12, and 2016m1. For Treatment indicator 1, the two-year house growth premium of the treated was 3 percent at its 2009 peak, 2 percent at its 2013 peak, and close to 2 percent at its 2016 peak. For indicator 2, the premium on house price growth was 4 percent in 2009, 3 percent in 2013, and also around 3 percent in 2016.

There is considerable heterogeneity across cities in the gap between exposed and non-exposed ZIP codes. Figure 11 plots the evolution of the estimated treatment effect for the 10 cities which are most exposed to Chinese demand. Specifically, we selected the 10 CBSAs that had the largest share of their component ZIP codes exposed to Chinese demand, according to our Treatment Indicator 2.[10] The differences that we see across cities is reassuring, in that it provides further evidence that our national-level estimates are not simply picking up national or international macroeconomic trends that happen to affect the areas where Chinese residents search for houses.

[10] Recall that our Treatment Indicator 1 specifically defines treated ZIP codes as being in the top decile of views within each CBSA, so it would not make sense to rank cities this way.

Figure 10: Average Treatment Effect on the Treated

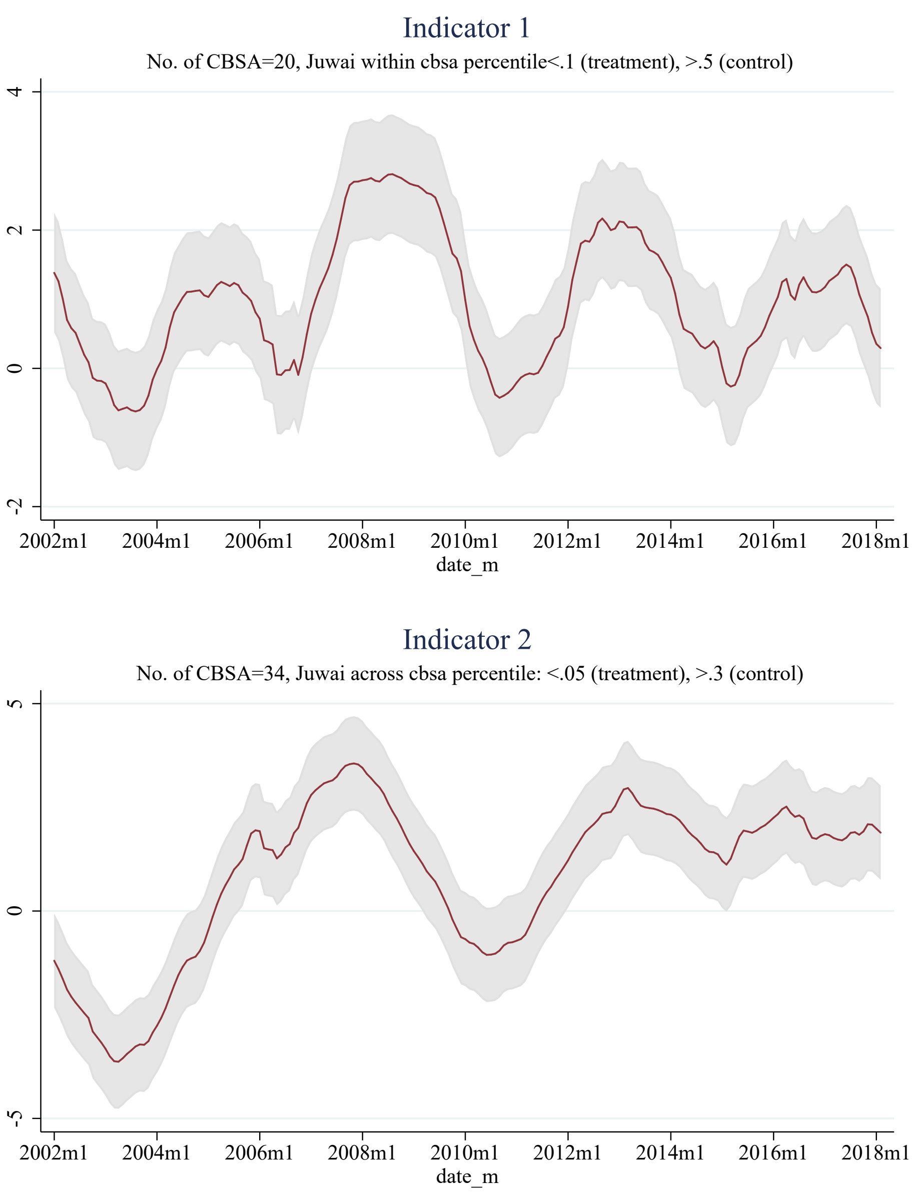

Figure 11: Average Treatment Effect on the Treated, Top 10 Cities

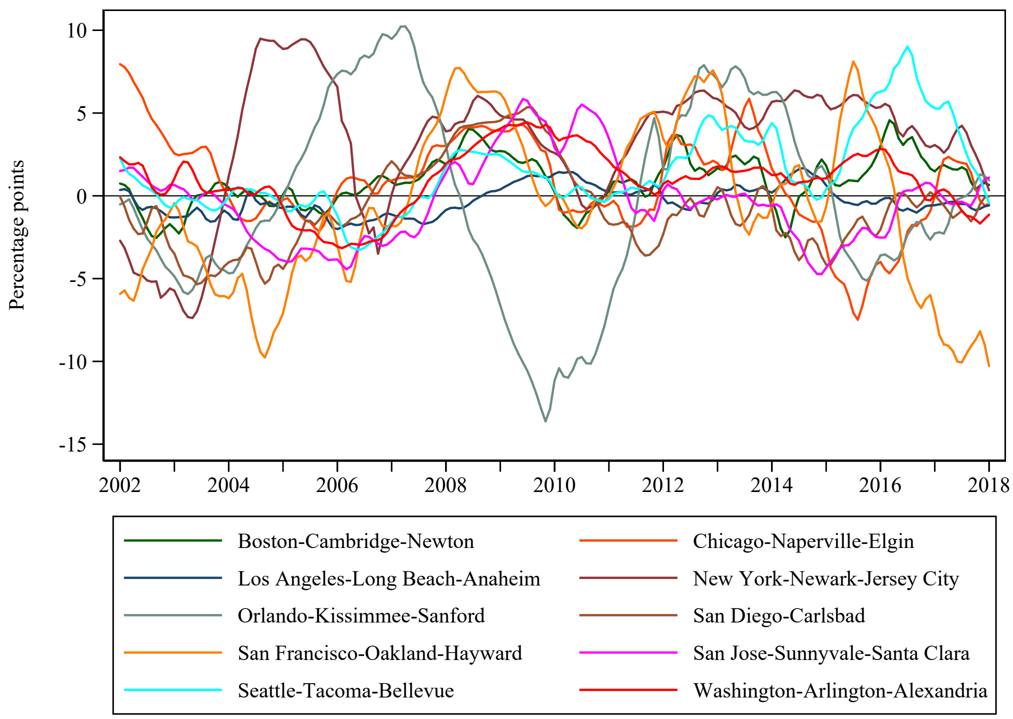

3.6 Robustness Tests: Placebo Matching

Despite our matching exercise, the possibility remains that the observables on which we have matched China-exposed ZIP codes to unexposed control ZIP codes do not fully capture pre-existing differences in price growth. If that is the case, then the price growth divergence we have found would merely reflect the self-selection of Chinese demand in real estate markets with rapid price growth, rather than any causal effect of Chinese housing purchases. To further allay this concern, we now conduct a placebo exercise.

Rather than analyzing the gap between China-exposed areas in the U.S. and matched control areas, we instead select as a placebo treatment group the ZIP codes that scored in the top 5 percentile of house price growth in the period 2010-2016 but are not defined as treatment group in our earlier exercise. We match them with similar ZIPs using the same procedure described in Section 5.1. The results are shown in Figure 12 .

The time evolution of the house price growth of this placebo treatment group is drastically different that the pattern evinced in the actual treatment group. By design, this group of ZIP codes experienced strong price appreciation over the period 2010-2016, hence Figure 12 does show an increase in house price growth rate over this period. However, unlike the results with the actual treatment group, the price growth gap only turns positive once in 2015, and the price growth gap is significantly negative for an extended period. It appears that the placebo ZIP codes suffered sharper housing price declines during the global financial crisis and also and saw a sharper rebound since 2012. Broadly speaking, the drivers of the price growth in the placebo ZIPs that we selected based on their having hot housing markets appear to be different than the drivers of price growth in the actual treatment group. And more importantly, the trajectory of the gap between the placebo group and their matched controls bears no obvious relationship to the measures of capital inflows from China that we analyzed in Section 2.

Figure 12: Placebo Difference in House Price Growth and Level

4 Connecting the Macro and the Micro: Capital Inflows from China and U.S. House Prices

In this section, we link the macroeconomic variables we examined in Section 2 with the micro-level effects that we identified in the previous section. We show that the price gap between China-exposed and non-exposed areas is significantly related to deposit inflows to the U.S. from China, and that this relationship is strongest after three quarters, consistent with capital flows from China entering the U.S. housing market. We then rule out alternative explanations for this comovement. First, we control for U.S. economic conditions in order to rule out the possibility that the dynamics of the price gap and capital inflows simply reflect the state of the U.S. economy. Second, we show that deposit inflows from other countries are unrelated to the price gap, demonstrating that the China inflow-price gap comovement is not generated by global economic conditions, but rather from China-specific shocks. Thus the analysis in this section provides further evidence that Chinese households have moved money into the U.S. via the banking system, money then used to purchase residential property. This then generates excess price growth in China-exposed areas as well as increasing the size of the U.S. statistical discrepancy when the funds leave the banking system.

Figure 13 plots the relationship between deposit inflows from China and Hong Kong along with the evolution of the gap between house price growth China-exposed ZIP codes and the matched matched controls.[11] Recall that in Section 2 we presented evidence that funds brought to the U.S. via the banking system were being used for house purchases with a lag. The relationship between inflows and the estimated treatment effect also exhibits this behavior: the contemporaneous correlation is only 0.06, but rises to 0.36 with a three-quarter lag. This is why Figure 13 plots the treatment effect three quarters ahead. The degree of comovement between the two series is striking, and we see that peaks in capital inflows from China and Hong Kong coincide with peaks in the treatment effect three quarters ahead. It is also notable that there was essentially no relationship beteen the two series prior to 2010, the year in which China liberalized some controls on capital outflows.

In the remainder of this section, we will examine the relationship between the divergence in price

[11] To conserve space, in this section we work exclusively with estimates constructed using Treatment Definition 1 described in the previous section.

Figure 13: Average Treatment Effect on the Treated and Capital Inflows from China

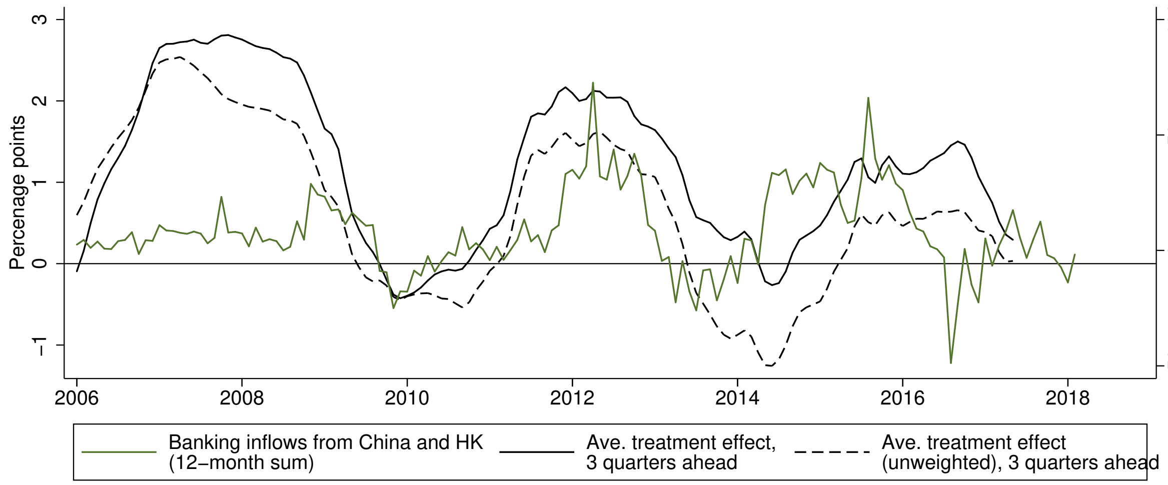

Treatment effect is the difference in 2−year house price growth between China−exposed ZIP codes and matched controls, calculated using Treatment Definition 1.
Sources: TIC system, authors' calculations.

growth and Chinese inflows to the U.S. more formally. In particular, we estimate the cumulative response of the gap between China-exposed and non-China exposed U.S. ZIP codes to capital inflows by estimating the following local projection, following Jordá (2005):

$$ATET_{t+h} = \alpha^h + \beta^h \text{China\_Deposit\_Inflows}_t + \gamma_1^h \Delta NFP_t + \gamma_2^h r_t^{mort} + \sum_{j=1}^{9} X_{t-j}\Lambda_{\text{j}}^h + \varepsilon_t$$

Where $ATET_{t+h}$ is the treatment effect plotted in Figure 10: the difference in two-year house price growth between China-exposed ZIP codes and matched non-exposed ZIP codes. We measure capital inflows from China (China_Deposits_Inflows$_t$) using the one-month change in deposits held by residents of China and Hong Kong in U.S. financial institutions, measured as a percentage of the level of Chinese and Hong Kong deposits at the end of the previous period. As we discussed in Section 2, deposit inflows is the capital flow appearing in official statistics most likely to reflect households moving money into the U.S. to purchase residential properties. Again, we include flows from Hong Kong based on media reports and mainland Chinese policies that suggest that Hong Kong acts as a major conduit for capital outflows from the Chinese household sector. To control for the state of the U.S. economy as it relates to the housing market, we include as controls the month-on-month growth in seasonally adjusted U.S. non-farm payrolls (ΔNFP_t) as well as the average 30-year mortgage rate in the U.S. (r_t^{mort}). The matrix X_{t-j} contains lagged values of the dependent variable, the shock, and the controls, with our specification containing nine lags. We experimented with alternative lag structures and differencing (e.g.year-on-year rather than month-on-month changes); the results were qualitatively similar to those presented below.

The results of the estimation are presented in Figure 14. We find a significant and positive relationship between deposit flows from China to the U.S. and the gap between house price growth between exposed and non-exposed ZIP codes. The effect peaks at around eight months. Recall that in Section 2 we showed that deposit outflows from China showed a strong correlation with the U.S. statistical discrepancy with a lag of three quarters and noted that this pattern is consistent with Chinese residents moving money into U.S. banks and using it to purchase real estate on average three quarters later. It is therefore striking that the impulse response in Figure 14 is also similarly consistent with such timing. As the China deposit inflows variable enters our specification in logs, the estimates in Figure 14 imply that a one percentage point increase in flows generates an 0.008 percentage point widening in the gap in price growth between China-exposed and non-exposed areas in the U.S. This effect may seem small at first glance, but recall from Figure 13 above that inflows reached roughly $70 billion during periods of concern about a China hard landing in 2012 and 2015. Thus our coefficient estimates explain the vast majority of the increase in the gap between China-exposed and non-exposed areas of the U.S. during those periods.

In contrast to the significant effect evident in Figure 14, the relationship between the treatment effect and the domestic U.S. variables relevant to the housing market that we include in our specification is either negative or not statistically significant, as shown in Figure 15. This increases our confidence that the effect we price uncover and its relationship with inflows from China is not merely an effect of the areas we classify as China-exposed being more sensitive to national or global shocks which also happen to attract inflows form China.

To demonstrate that our results reflect a causal relationship between capital flows from China and the price of U.S. real estate, we conduct another placebo test. The primary concern with our analysis relating inflows to the U.S. from China to U.S. house prices is omitted variable bias—perhaps these two variables are both driven by some third factor, such as the state of the global economy. To verify that this is not the case, we test whether there is a relationship between inflows from other countries and gap in house price appreciation between the China-exposed and non-exposed areas we are studying.

Figure 14: Treatment Effect and Deposit Inflows from China

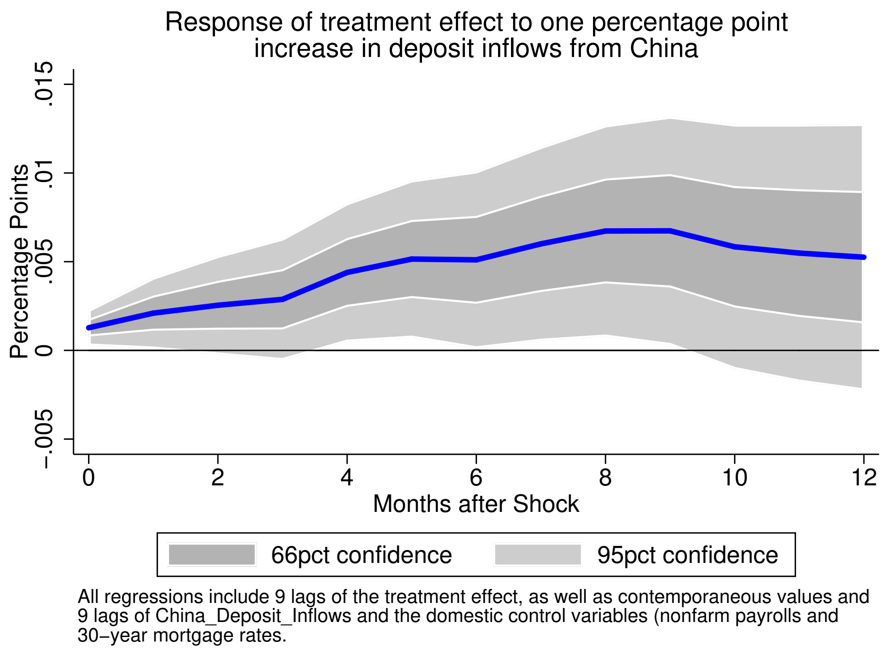

Figure 15: Treatment Effect and U.S. Variables

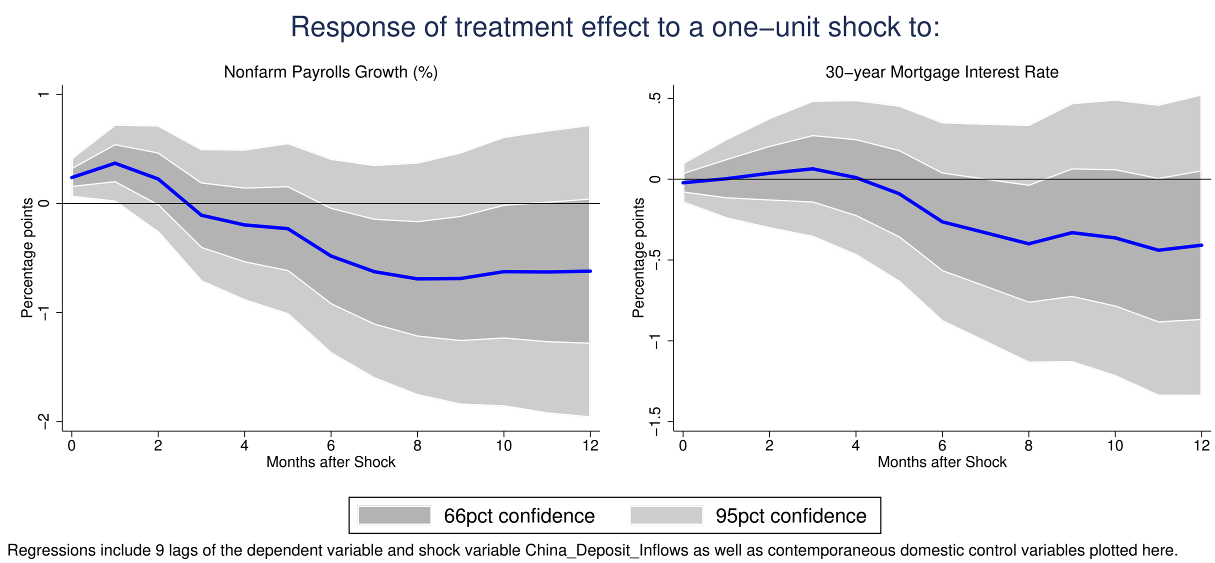

In Figure 16 we repeat the original exercise of estimating the relationship between deposit inflows to the U.S. and the average treatment effect estimated above, but instead of using inflows from China and Hong Kong as the shock variable we use capital flows from several other countries. In particular, we analyze the relationship for the ten countries other than China that are the top sources of foreign deposits in U.S. financial institutions and which are not global financial centers.[12] Figure 16 makes clear that inflows from these countries generally have no significant relationship with the relative growth of house prices in the China-exposed ZIP code we study. Importantly, China is not consistently the largest source of deposit inflows to the U.S. nor is the variance of flows from China systematically higher than flows from the other countries we focus on in this placebo exercise. Thus the lack of significance for these other countries is not simply due to flows from China being larger or more variable. Moreover, where the relationship is significant, it is usually negative. In addition, the point estimates of the effects for these countries are orders of magnitude smaller than what we find for China. Overall, this exercise confirms that our results on the relationship between capital inflows from China and the trajectory of house prices in China-exposed areas is not the result of both variables moving in response to general global financial conditions.

[12]We exclude from this exercise financial centers including the the Cayman Islands, the Bahamas, and the British Virgin Islands.

Figure 16: Flows from other countries and U.S. real estate prices

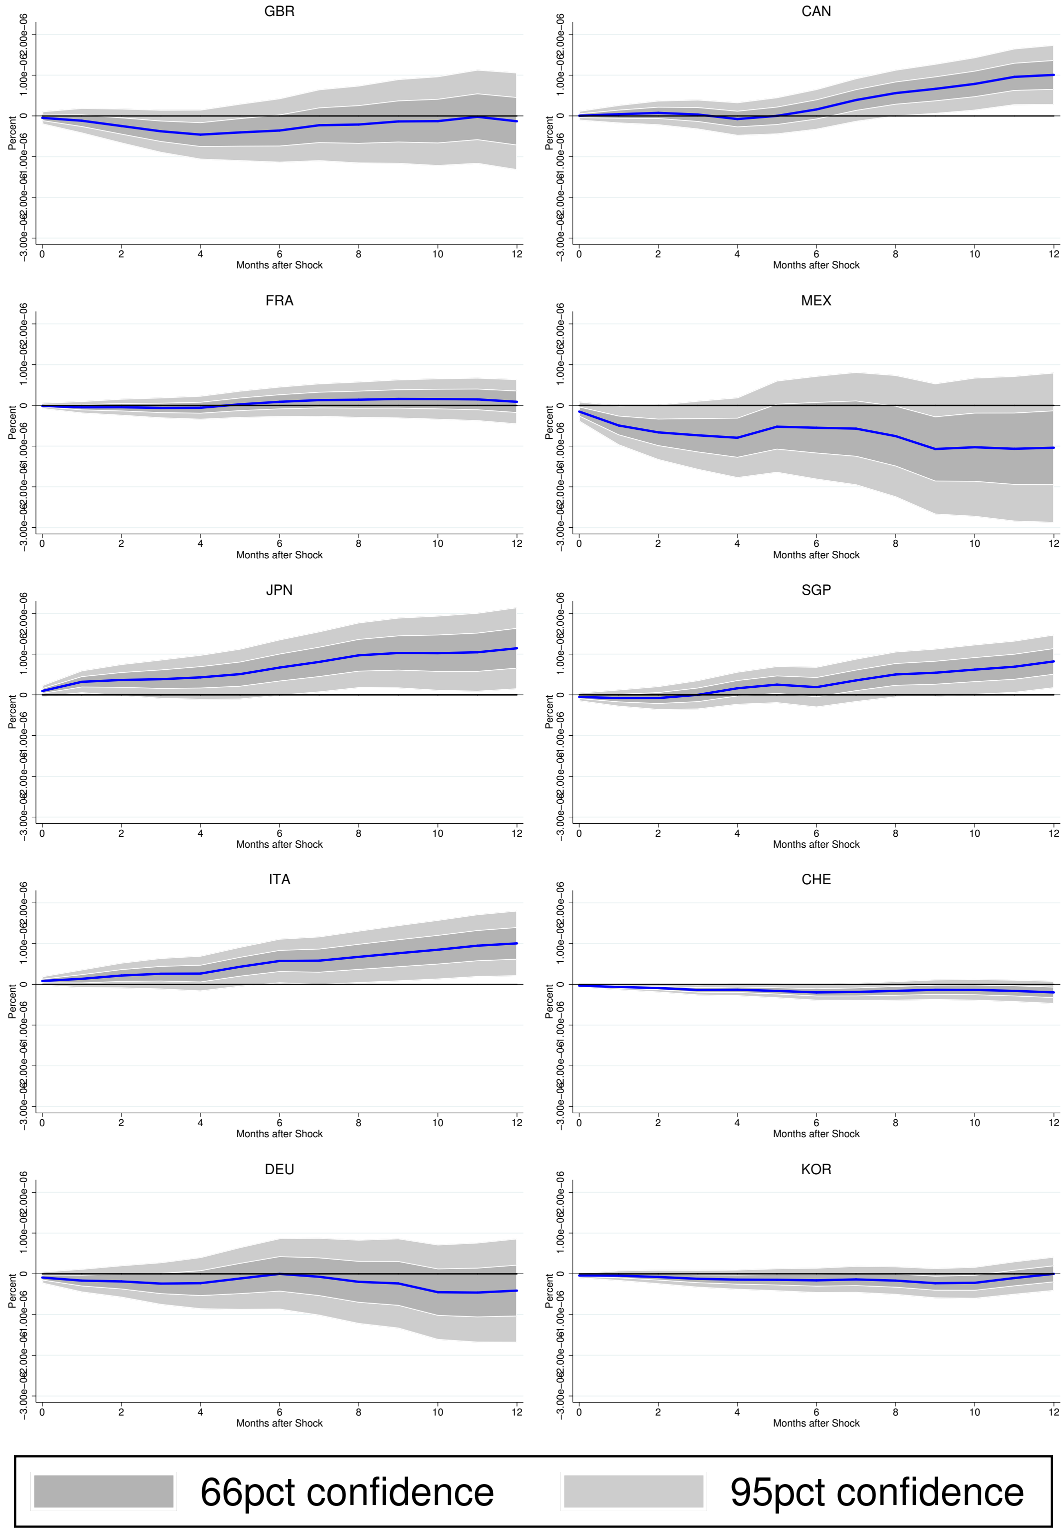

All regressions include 9 lags of the treatment effect as well as contemporaneous values and 9 lags
of deposit inflows from China and the two U.S. controls (nonfarm payrolls and the 30−yr mortgage rate.

5 Conclusion

In this paper we have presented a broad range of evidence that suggest substantial purchases of U.S. residential real estate by safe haven inflows from China following periods of economic stress in China since 2010. We have shown that this novel type of safe haven capital flow has generated multiple China shocks in the U.S. housing market, with house prices in China-exposed areas have rising significantly faster than those in areas not exposed to Chinese demand.

At the macro-level, we have shown that measures of macroeconomic and financial stress in China, as well as inflows of deposits from China and Hong Kong into the U.S., strongly commove with the unrecorded capital inflows captured by the U.S. statistical discrepancy, with a lag of three quarters. And we have discussed in detail why, due to the way U.S. balance of payments statistics are collected, this relationship is consistent with Chinese households moving money into U.S. and subsequently using the funds to purchase residential real estate.

Our micro-level analysis made use of a novel dataset to directly measure variation in Chinese demand for residential real estate across U.S. ZIP codes, showing that house prices have increased on average faster in China-exposed areas than those in otherwise similar areas which have not attracted interest from potential Chinese buyers. After exploring in detail the drivers of the foreign interest at the city and ZIP code level, we matched ZIP codes in our dataset that attract a relatively high level of Chinese interest with observationally similar control ZIP codes receiving relatively low interest. The resulting estimate of the average treatment effect on the treated ZIP codes indicated that exposure to Chinese demand accelerated price growth by on average one or two percent per year after 2010. Looking at the dynamics of this treatment effect over time, we found that the price growth gap widened markedly following periods of economic stress in China, an indication that Chinese households have purchased U.S. residential real estate as a safe haven asset.

Finally, we linked the time varying average treatment effect that we estimated using micro data with the aggregate measures of capital inflows from China that we initially analyzed. In local projections we found that the two series were significantly related, with the timing of the peak response consistent with the findings of our macro-level analysis.

Throughout the paper, we have conducted robustness tests to rule out alternative explanations for our findings. A placebo exercise defining hot housing markets as the treated group generated an estimated treatment effect that was qualitatively different from ours. Our estimated treatment effect was not significantly related to domestic U.S. variables. And deposit inflows from countries other than China which are important sources of capital flows to the U.S. behaved differently and were not significantly related to the treatment effect we estimated.

Overall our findings suggest that housing markets in major U.S. cities have been subject to global safe haven flows from China following periods of economic stress in China. The fact that authorities in most countries do not collect data on foreigners purchases of residential real estate makes clear the novelty of type of capital flow. The Chinese governments purchases of U.S. Treasury debt has previously attracted attention, as have Chinese firms direct investment in the United States. Our findings show that another, very different type of capital inflow has significant economic effects in the U.S., both at the local and the macroeconomic levels.

References

Aizenman, J. and Jinjarak, Y. (2014). Real estate valuation, current account and credit growth patterns, before and after the 20089 crisis. *Journal of International Money and Finance*, 48(Part B):249 – 270. Current account imbalances and international financial integration.

Badarinza, C. and Ramadorai, T. (2018). Home away from home? foreign demand and london house prices. *Journal of Financial Economics*, 130(3):532 – 555.

BEA (2019). U.s. net international investment position first quarter 2019, year 2018, and annual update. Technical report, Bureau of Economic Analysis.

Bednarek, P., te Kaat, D., , Ma, C., and Rebucci, A. (2019). Capital flows, real estate, and city business cycles: Micro evidence from the german boom. Mimeo.

Caliendo, M. and Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72.

Cesa-Bianchi, A., Cespedes, L. F., and Rebucci, A. (2015). Global liquidity, house prices, and the macroeconomy: Evidence from advanced and emerging economies. *Journal of Money, Credit and Banking*, 47(S1):301–335.

Chinco, A. and Mayer, C. (2016). Misinformed speculators and mispricing in the housing market. *The Review of Financial Studies*, 29(2):486–522.

Curcuru, S. E., Thomas, C. P., and Warnock, F. E. (2009). *Current Account Sustainability and Relative Reliability*, pages 67–109. University of Chicago Press.

Favilukis, J. Y. and Van Nieuwerburgh, S. (2018). Out-of-Town Home Buyers and City Welfare.

Flatness, A., Whitaker, E., and Yuskavage, R. (2009). Annual revision of the u.s. international accounts. *Survey of Current Business*, 89(7):35– 47.

Guse, M., Park, W., Saravay, Z., and Yook, Y. (2019). Collateralized loan obligations in the financial accounts of the united states. Technical report, Federal Reserve Board.

Jordá, . (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Lane, P. R. and Milesi-Ferretti, G. M. (2007). The external wealth of nations mark ii: Revised and extended estimates of foreign assets and liabilities, 1970-2004. *Journal of International Economics*, 73(2):223 – 250.

Li, Z., Shen, L. S., and Zhang, C. (2019). Capital flows, asset prices, and the real economy: A 'china shock' in the us real estate market. Unpublished.

Liu, E. and Schmidt-Eisenlohr, T. (2019). Who owns u.s. clo securities? Technical report, Federal Reserve Board.

Sa, F., Towbin, P., and Wieladek, T. (2014). Capital inflows, financial structure and housing booms. *Journal of the European Economic Association*, 12(2):522–546.

Saiz, A. (2010). The Geographic Determinants of Housing Supply*. *The Quarterly Journal of Economics*, 125(3):1253–1296.

Appendix

Figure A1: Correlations: Foreign deposit outflows
vs U.S. foreign deposit inflows

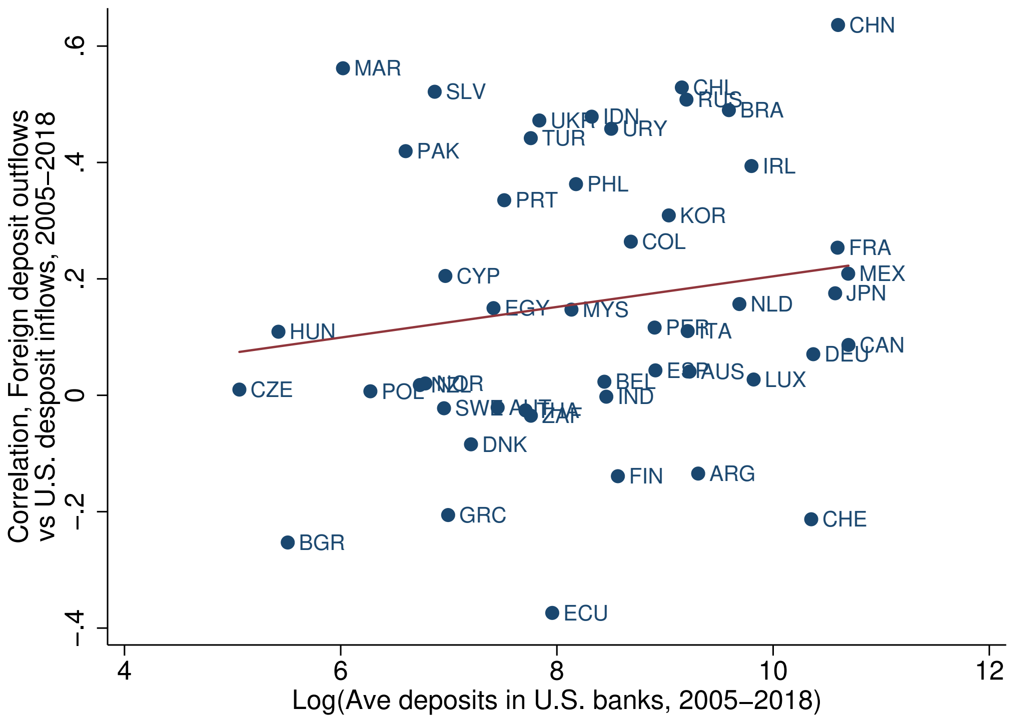

Table A1: Correlations: Foreign deposit outflows vs U.S. foreign deposit inflows

Country	Correlation between foreign deposit outflows & U.S.deposit Inflows	Ave. deposits in banks (USD Billions)
CAN	0.086	44220.02
MEX	0.209	44115.93
CHN	0.636	40169.52
FRA	0.254	39986.32
JPN	0.175	39096.38
DEU	0.071	31945.02
CHE	-0.213	31336.30
IRL	0.394	18025.86
NLD	0.157	16123.29
BRA	0.490	14628.95
ARG	-0.135	11012.84
AUS	0.041	10169.16
ITA	0.110	9993.05
RUS	0.508	9881.36
CHL	0.529	9468.34
KOR	0.309	8390.27
ESP	0.043	7414.02
PER	0.116	7360.84
COL	0.264	5906.38

Source: TIC data and IMF BoPS. Excludes financial centers.

Figure A2: Correlation Structure—Bank inflows to the U.S. from China and Hong Kong vs. the U.S. statistical discrepancy

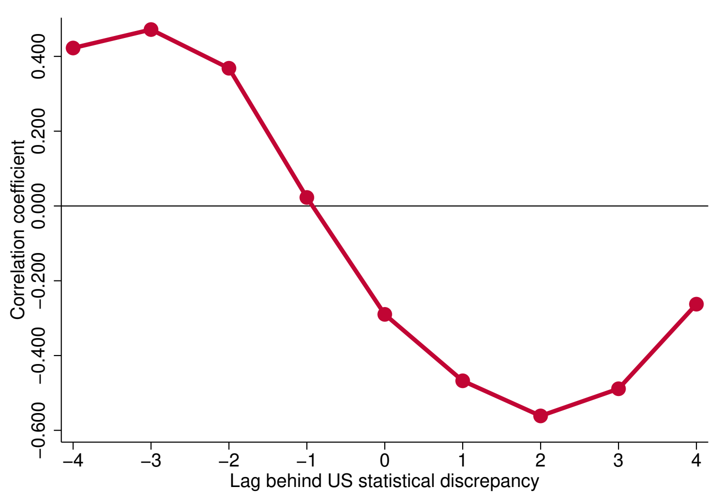

Source: Haver, TIC System, authors' calculations.

Juwei Veiews Data, Additional Information

In additional to providing information of an overseas real estate property on their website, Juwai also offers consulting services for the potential buyer and helps refer the potential buyer to a real estate brokerage firm abroad. As reported by the Juwai CEO, potential buyers would oftentimes make one visit to the city which their property of interest is located and make a purchase within 6 months. We validate the Juwai data by pairing the passenger arrival data from a Chinese city to a U.S. city. Figure 4 shows that the city-pair arrival data has a 0.54 correlation with the Chinese city-U.S. city pair Juwai data. We also cross-checked the cities which have high Juwai views with the percent of houses purchased wit

Table A2: Share of Juwai U.S. Listing Views, by CBSA

Rank	CBSA	State	Share of Juwai U.S. Views
1	Los Angeles-Long Beach-Anaheim	CA	18.9%
2	New York-Newark-Jersey City	NY-NJ-PA	12.3%
3	Seattle-Tacoma-Bellevue	WA	5.5%
4	Riverside-San Bernardino-Ontario	CA	4.3%
5	San Jose-Sunnyvale-Santa Clara	CA	3.0%
6	Houston-The Woodlands-Sugar Land	TX	2.8%
7	San Francisco-Oakland-Hayward	CA	2.8%
8	Orlando-Kissimmee-Sanford	FL	2.6%
9	Chicago-Naperville-Elgin	IL-IN-WI	2.2%
10	Miami-Fort Lauderdale-West Palm Beach	FL	2.2%
11	Boston-Cambridge-Newton	MA-NH	2.0%
12	San Diego-Carlsbad	CA	2.0%
13	Washington-Arlington-Alexandria	DC-VA-MD-WV	2.0%
14	Sacramento–Roseville–Arden-Arcade	CA	1.9%
15	Philadelphia-Camden-Wilmington	PA-NJ-DE-MD	1.4%
16	Urban Honolulu	HI	1.4%
17	Atlanta-Sandy Springs-Roswell	GA	1.4%
18	Oxnard-Thousand Oaks-Ventura	CA	1.2%
19	Dallas-Fort Worth-Arlington	TX	0.9%
20	Detroit-Warren-Dearborn	MI	0.9%

Figure A3: Juwai views vs. cash sales share, by city

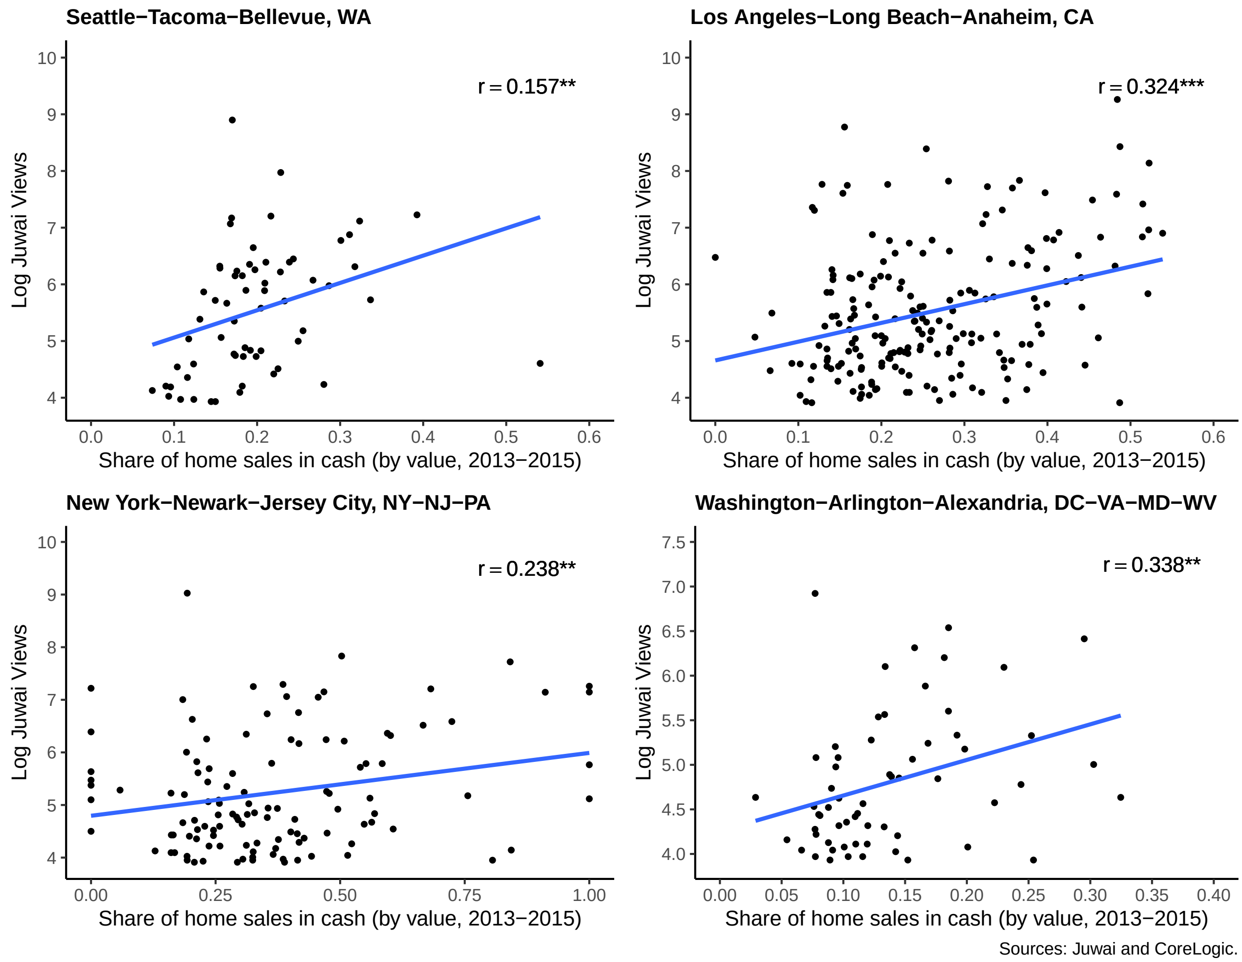

Figure A4: An Example of Matching (Indicator 1): Seattle