

The Financial Consequences of Wanting to Own a Home*

Caitlin S. Gorback, McCombs School of Business, UT-Austin[†]

and

Gregor Schubert, UCLA Anderson School of Management[‡]

February 2025

Abstract

We study the causal effects of homeownership affinities on tenure choice, household sensitivity to credit shocks, and retirement portfolios. Exploiting exogenous variation in affinities across U.S. immigrants' countries of origin, we find that a 10pp higher affinity causes a 1.5pp higher homeownership rate in the U.S. Using exogenous credit-supply shocks, we show that high-affinity households are more responsive to credit availability, and less likely to default on mortgage payments. By retirement, high-affinity households realize higher homeownership, greater total wealth, and larger real estate shares in their portfolios. These effects are largely driven by appreciation.

*A previous version of this paper was circulated under the title “The Cultural Origins of Immigrant Homeownership”. We thank discussants and participants at 2023 ASSA-AREUEA, UCLA Anderson Finance Seminar, the Tel Aviv University Conference on Low Income Housing Supply and Housing Affordability, the AREUEA Meetings in Dublin and Tokyo, the Southern Economic Association Meeting, the Harvard Development and Urban Workshop, the Federal Housing Finance Agency, the UT Austin Economics and UT Austin McCombs faculty seminars, the 2022 Pre-WFA Real Estate Symposium and NBER Real Estate Summer Institute, the University of Houston Empirical Microeconomics Seminar, the 2023 BFI Women in Empirical Microeconomics Conference, the 2023 Utah Eccles Innovations in Housing Affordability Summit, the University of Nevada in Las Vegas, and the 2024 UEA meeting for valuable comments and suggestions. We also thank the UCLA Ziman Center for Real Estate’s Rosalinde and Arthur Gilbert Program in Real Estate, Finance and Urban Economics for generous funding.

[†]Email: caitlin.gorback@mcombs.utexas.edu

[‡]Email: gregor.schubert@anderson.ucla.edu

I. Introduction

In the United States, homeownership has long been viewed by policymakers as a bellwether for quality of life and a core tenet of the American Dream (Layton, 2021; Goodman and Mayer, 2018). As such, many policies aim to support high homeownership rates, including, among others, the mortgage interest rate deduction, securitization of mortgage products, and low-down payment programs for first-time homebuyers, though none have pushed homeownership sustainably above 65% (Layton, 2022). These policies are touted as supporting household wealth formation. However, while some of these policies have been shown to be effective in increasing homeownership, they may also reduce household mobility (Fonseca and Liu, 2024a) and have been suggested to substitute for other forms of savings (Mian and Sufi, 2009; Bhutta, 2015; Cocco, 2005). It is therefore important to understand the impact of homeownership on housing returns and portfolio choice, and what makes policy targeted at homeownership more or less effective.

In this paper, we use variation in immigrants' affinities for homeownership as a laboratory to study how homeownership impacts tenure choice, household wealth, and portfolio choices. We show that a higher affinity for homeownership, proxied by origin-country homeownership rates by marital status, provides plausibly exogenous variation in the decision of whether or not to own a home, even conditional on a comprehensive set of socioeconomic and demographic characteristics. These affinities are robust to a variety of country-of-origin controls, and attenuate with measures of assimilation in the U.S., suggesting a cultural origin unique to each country's institutional and financial history. This variation then enables us to show that a greater affinity for being a homeowner has a causal effect on average homeownership among foreign-born U.S. residents. Moreover, it changes housing cycle exposure for the high affinity groups by causing a greater sensitivity to exogenous credit supply shocks. These households also default less often than other households that also bought at the peak. Finally, we show that a higher desire to own a home results in a greater likelihood of owning a home at retirement age, a greater share of retirement assets coming from a household's primary residence, and greater overall wealth at retirement.

In order to show how a greater preference for owning a home can affect the eventual portfolio decision of households, we first augment a simple tenure choice model, based on Brueckner (1986), by allowing for heterogeneity in affinities for owning a home. We use this conceptual framework to show how these affinities impact households' decision to own vs. rent and their sensitivity to credit market shocks. We first show that, in a model of housing consumption in which cultural affinities lower the perceived user cost of owning a home, these consumption benefits lead households to switch from renting to owning at higher rates. Specifically, the propensity to own increases with a household's homeownership affinity at all tax rates above a minimum threshold. Next, we add borrowing constraints to the model and expand the setting to two time periods, requiring households to save for downpayments should they wish to purchase in the second period. We show that higher downpayments erode the benefits of owning relative to renting, and that the entry into homeownership for those with a high affinity is more sensitive to these changes in credit constraints. Moreover, those with high affinities are less likely to default in times of hardship.

Second, we test these predictions empirically. Taking the model’s implications to the data requires us to develop a measure of homeownership affinity. Building on the labor literature that leverages cultural norms in the country of origin to provide variation in U.S. labor force participation (Blau, Kahn, and Papps, 2011), we study the cultural norms around homeownership for the set of immigrants to the U.S. The U.S. foreign-born population is large and has been growing, representing about 10% of residents in 2000, and rising to 14% by 2019. Enabling us to study the impact of cultural norms on homeownership, these immigrants hail from over 150 countries, with large differences in their expectations around homeownership. Building on this observation, we collect new data on cultural norms around homeownership among foreign-born U.S. residents, with variation driven by households’ differential birthplaces. Specifically, we define “homeownership affinity” as the rate of homeownership in a household head’s country of origin (“HOCO”). Using international census microdata data we compute recent homeownership rates by marital status across 85 origin countries; these cover the majority of countries from which migrants to the U.S. originate and nearly 80% of the immigrant population.

We combine this data with household microdata from the American Community Survey (ACS) data to document the passthrough of cultural norms of homeownership among immigrants in the U.S. We show that the probability of owning a home is higher for all federal marginal tax rates above 10% among immigrant households with above-median homeownership in their country of origin. Using household-level regressions, we find that these country-of-origin-related norms significantly affect household tenure choice in the U.S., with immigrants from countries with strong cultural norms around housing more likely to own homes in the U.S., while those with weaker preferences opt instead to rent. The size of the effect suggests that a 10 pp increase in homeownership in a householder’s country of origin is associated with a 1.5 pp increase in the likelihood of being a homeowner as a U.S. resident. The interquartile range of differences in *HOCO* among foreign-born residents causes a 2.8 pp difference in homeownership, for a relative magnitude, this equals 44% of the standard deviation in homeownership across U.S. states.

This effect survives rigorous controls for households’ financial standings, demographics, household size and composition. We also control for a variety of origin-country characteristics, such as inflation experience, credit supply, property rights, and measures of cultural closeness to the U.S. To address concerns that immigrants are an assortatively selected group, meaning they may come from wealthier backgrounds in their origin country, we control for mean characteristics among immigrants from each country. We address sample selection concerns by showing that the results are robust to using alternative measures of origin preferences, removing particularly large immigrant population groups, and replicating the results in the American Housing Survey (AHS).

We explore different mechanisms for *why* origin country homeownership passes through to tenure choice in the U.S. by estimating how this effect varies with the characteristics of the householder or their spouse. We find that the effect is only mildly attenuated for those who immigrated as children, and is stronger for female householders, those without a college degree, and those who have been in the U.S. for longer. Moreover, the effect is much smaller for householders who marry

a U.S.-born spouse. For those with a foreign-born spouse, the spouse’s *HOCO* has a similar effect on tenure choice as the householder’s *HOCO*. These analyses suggest that this effect likely reflects a latent difference in the desire to own a home that a household realizes over time, and which weakens with greater assimilation.

This first set of empirical results suggests, on the one hand, that cultural norms evolve in each country according to its own unique history of institutions, and cannot be explained by any particular set of controls. On the other hand, the surviving variation likely reflects cultural preferences transferred from immigrants’ origin country to the U.S., and has important effects on financial decisions, consistent with the growing literature on the effects of cultural and social backgrounds and past experiences on financial decision-making (e.g. Malmendier and Wellsjo (2024)).

Next, we test the model’s second implication that higher affinity households—in our context those with a high desire to own a home because of their origin—should be more responsive to mortgage credit supply shocks. We first provide descriptive evidence that households with a higher affinity for homeownership that were renters in the year 2000 were more likely to see their homeownership increase by 2005 even without a change in household characteristics. Put differently, high *HOCO* renters were more likely to be “marginal” with regard to the expansion in homeownership during the 2000-2005 boom. Moreover, we find that this effect is driven to a large degree by households that were more constrained in the year 2000—as measured by having household characteristics that were associated with low homeownership rates—but had a high affinity for owning a home. This example suggests that credit booms may see greater take-up among households that are primed to want to buy a home, leading to a divergence in homeownership.

In order to test this idea more formally, we move from cross-sectional household-level analysis to group-level analysis over time to study the causal impact of mortgage credit supply shocks on groups with different homeownership norms. If the sensitivity to changes in credit supply varies with homeownership norms, this represents a novel factor that might be influencing the transmission of credit policy and housing market dynamics in response to shocks. We aggregate households into panels of groups by homeownership norms within each commuting zone for 2000-2015 and compare how their homeownership changes in response to exogenous shocks to mortgage credit supply in a stacked panel of long differences over the 5-year periods that correspond to the boom (2000-2005), bust (2005-2010), and recovery (2010-2015) periods of the housing cycle.

We construct exogenous local credit shocks as “shift-share” instruments based on a county’s exposure to lenders that are expanding their mortgage portfolio for reasons unrelated to local credit demand. This approach follows a literature using non-local lender-level shocks to build exogenous shocks to the availability of credit to borrowers (e.g. Khwaja and Mian (2008); Greenstone, Mas, and Nguyen (2020)). We follow Gilchrist, Siemer, and Zakrajsek (2018) and Garcia (2020) and estimate the non-local component of the national change in mortgage lending by each lender as time-varying lender fixed effects when controlling for county-level mortgage lending activity. This isolates the degree to which a lender increases its lending activity in a way that is unrelated to the changes in overall credit demand in the counties where the lender has a presence. Counties that have

higher exposure to lenders that are expanding their lending activity nationally (e.g. due to relative changes in their ability to fund mortgages) should see higher *relative* mortgage lending activity - but by construction this shock is uncorrelated with national trends and regional differences in average credit demand changes. We show that these exogenous shocks strongly predict time variation in local loan originations at the county level.

We create pseudo-panels for different groups within counties for 5-year periods, using the ACS microdata, and estimate the effect of these exogenous credit shocks on changes in homeownership over time. We compare the relative homeownership effects of the credit shocks between groups that differ in their affinity for homeownership but otherwise are exposed to the same local economic trends and have similar observable characteristics. We find that homeownership responds significantly more to increases in credit availability for high affinity immigrant groups, relative to low affinity immigrants. A 1 SD higher shock to loan originations increases the gap in homeownership between high and low affinity immigrants groups by an additional 1.06 pp over a five-year period, relative to a baseline gap of 0.86pp. This effect is robust to controlling for group-by-county and commuting-zone-by-year fixed effects. To the degree that credit shocks drive housing market cyclicality, foreign-born residents from high-HOCO countries are thus more likely to respond to credit easing as a result of policy or regulatory changes.

Consistent with their greater sensitivity to mortgage credit, we show that high affinity households saw much more growth in homeownership during the 2000-2005 credit boom than low-affinity groups; moreover, those with high affinities also seem to hold on to their homes more strongly during downturns, experiencing a smaller decline in homeownership during the 2005-2015 period, relative to other population groups. We document that high *HOCO* homeowners are less likely to miss mortgage payments in later years, conditional on their financial circumstances, as proxied by debt-to-income and cost-to-income ratios. Focusing on the housing boom period, we find that households that purchased a home at the end of the housing boom in 2006-2007 on average experienced twice the delinquency rates of households that purchased homes in other years. However, having an affinity for homeownership lowers the likelihood of mortgage delinquency: households with a one SD higher *HOCO* who bought at the peak saw 38% lower delinquency rates than other households in their same purchase year cohort.

What does a greater tendency to own a home mean for long-run household financial portfolios and wealth formation? While this is an important question for policy purposes, it is usually difficult to identify causal variation in tenure choice. The first part of our paper establishes that an affinity for homeownership robustly induces differences in tenure choice that are plausibly exogenous with regard to other financial windfalls experienced by a household. We utilize this variation to study portfolio allocations later in life. Under the assumption that cultural norms around homeownership impact a household's portfolio allocation at retirement only through the household's decision to purchase a home or not, we can therefore use *HOCO* differences to study the causal effects of tenure choice on household retirement portfolios. This design is akin to a reduced-form instrumental variables estimation.

For this analysis, we use restricted-access data from the Health and Retirement Study (HRS), which, to our awareness, is unique in combining information on households' financial positions, housing consumption, and detailed country-of-origin information.¹ We construct a sample of foreign-born survey respondents' household balance sheets at (or near) age 60, matched to their *HOCO* information, and estimate the effect of the latter on wealth, portfolio composition, and real estate holdings, while controlling for a rich set of other household characteristics.

We find that a greater affinity for homeownership leads to a higher likelihood of owning one's home at retirement, even conditional on total non-housing wealth. As a result, for respondents with 1 pp higher *HOCO*, the net value of their primary residence on average represents a 0.9 pp greater share of their retirement portfolio. Moreover, we estimate that a higher affinity for homeownership causes households to retire with higher total wealth, consisting of \$3.5K higher wealth coming from their primary residence, and a \$4.6K increase in other assets, including other real estate. A standard deviation increase in *HOCO* (12 pp) is associated with increases in primary residence value by 38% and other asset wealth by 43% relative to their sample averages. This wealth effect also exists at the intensive margin: even within homeowners, those with greater *HOCO* end up with higher wealth at retirement.

What can explain this overall increase in wealth as a result of homeownership? We find no evidence that households with greater affinity for homeownership are more likely to own stocks or invest a greater share of their portfolio in stocks, making it unlikely that shifts in preferences for risky assets other than real estate explain this effect. Instead, we find evidence that the wealth effect of homeownership may be the consequence of being in the right place (and home) at the right time: *HOCO* has a greater effect on the real estate component of household wealth in states that experienced large run-ups in house prices. Combined with the overall increase in house prices across the country in recent decades, this suggests that homeownership effects on wealth may not be automatic, i.e. occurring through forced savings, but may rather depend on an exceptional ability or luck in timing housing market cycles.

These findings have important policy implications. While we document a new channel for wealth inequality by showing that a greater affinity for homeownership *does* drive wealth at retirement, we also find that much of this effect is driven by windfalls, leaving little scope for deliberate policy action. In our setting, we exploit variation that nudges one's preferences towards housing on the extensive margin, and find that this does not yield higher savings without good luck in becoming a homeowner in the right location at the right time. This suggests that policies targeting homeownership alone are not enough to reliably improve retirement wealth. Instead, policy also needs to target the intensive margin of savings, conditional on homeownership (see, e.g. Bernstein and Koudijs (2024)). Furthermore, the modest benefits of housing wealth at retirement must be weighed against the risks of promoting homeownership to marginal populations as documented in Keys, Mukherjee, Seru, and Vig (2010); Adelino, Schoar, and Severino (2016); Justiniano, Primiceri, and Tambalotti (2019); Gupta (2019). A better alternative would be policies that encourage savings directly, such as tax-advantaged retirement savings accounts or financial literacy programs.

Related literature & contributions. We contribute to a growing literature showing how experience impacts households’ financial choices: Botsch and Malmendier (2020) and Malmendier and Wellsjo (2024) show that consumers who experienced high inflation are averse to adjustable rate mortgages, and also opt into higher rates of homeownership. Moreover, there is a strong correlation between where people grow up and their tenure choice as adults, potentially as parents subsidize their children’s purchases through equity extraction (Benetton, Kudlyak, and Mondragon, 2024), or through children adopting their parents’ asset preferences (Charles and Hurst, 2003). Ringo (2020) finds that homeownership rates in one’s county of birth in the U.S. predicts adult homeownership, and Marcén and Morales (2020) document that the U.S. homeownership rate among immigrants is correlated with homeownership rates in their country of origin in a 2016 cross-section of U.S. households. Happel, Karabulut, Schäfer, and Tuzel (2022) show that differences in attitudes to homeownership in Germany as a result of exposure to residential housing destruction during World War II persist even in the long run.

We add to the growing literature studying the take-up of credit among marginal homebuyers. This literature began with the acknowledgment that many of the policies expanding access to credit in the mid-2000’s, such as private label securitization and exotic mortgage structures, led to increased homeownership transitions among marginal homebuyers. These subprime homebuyers defaulted on their mortgages at historic rates, driving much of the housing bust during the Great Recession (Mian and Sufi, 2009; Bhutta, 2015). Bayer, Ferreira, and Ross (2016) found that many of these marginal homebuyers tended to be racial minorities, with large populations of black and Hispanic borrowers newly taking up credit and subsequently defaulting. We document that a subset of immigrants look to be similarly marginal in their responses to mortgage credit expansion, potentially exposing them to default risk, or on the upside, large increases in housing wealth, depending on their market timing.

Finally, we add to the growing understanding of how homeownership interacts with households’ broader wealth formation and portfolio choice (e.g. Cocco (2005); Yao and Zhang (2005); Chetty, Sándor, and Szeidl (2017)). This is an old literature, yet difficult to improve on using microdata methods, as very few datasets offer comprehensive surveys of households’ financial *and* housing portfolios. Most recently, Sodini, Van Nieuwerburgh, Vestman, and von Lilienfeld-Toal (2023) use rich datasets from Sweden on demographics, housing, finances and debt in conjunction with a natural experiment in which municipalities decided whether or not to privatize public housing. The authors find overwhelming benefits to homeownership in their setting: homeowners accumulate more wealth and allocate a higher share to risky assets. Complementing this literature, we exploit plausibly exogenous variation in the ownership decision to study U.S. households’ sensitivities to credit cycles and the subsequent retirement portfolio consequences.

We proceed with Section II, which grounds the proposed effects of an affinity for homeownership in a simple model of tenure choice under borrowing constraints. Section III follows, describing how we collect and build our data. Section IV outlines our estimation strategy establishing immigrant norms’ impacts on household-level homeownership, and reviews our results. Section V outlines

how we measure immigrants' homeownership responses to credit supply shocks and discusses our findings. Section VI examines how an affinity for homeownership impacts portfolio allocations at retirement. Section VII concludes.

II. Theoretical Framework

We first develop a stylized model of tenure choice incorporating an affinity for housing into a user cost model, and later add borrowing constraints. We consider housing consumption to be exogenous in order to keep the model tractable and extend the model in Brueckner (1986).

To provide the basic intuition, first consider a static tenure choice problem: a household decides whether to rent, in which case its consumption x^R , net of rent payments Q and taxes t is

$$x^R = (1 - t)y - Q,$$

where y is the household income. If the household decides to become a homeowner instead, it has to make a down-payment αP , where P is the house price, at the beginning of the period, which has an opportunity cost r .² The household finances the remainder $(1 - \alpha)P$ of the house price with mortgage payments at interest rate r , which are tax-deductible, sells the house again at the end of the period for price P and receives back its equity αP . Moreover, the household experiences a cultural cost of being a homeowner ϕ , which represents the degree to which cultural benefits of homeownership (e.g. pride, sense of security, control premium) offset the additional costs of homeownership (e.g. mental burden, lower mobility, time cost of maintenance). Households from groups with a relatively high cultural affinity for homeownership will experience a lower cost of homeownership. Consumption for homeowners is therefore

$$\begin{aligned} x^H &= (1 - t)(y - \alpha r P - r(1 - \alpha)P) - \alpha P + \alpha P - \phi \\ &= (1 - t)y - (1 - t)rP - \phi \end{aligned}$$

Note that the initial equity payment and later return of principal net out and that the need to finance the equity payment means that the household effectively pays interest on the full purchase price. For now, we are implicitly assuming that the household has no wealth or financing constraint in coming up with the downpayment - and we will modify this assumption later. Imposing a zero profit condition for landlords, such that $Q = rP$, this means that households are indifferent between owning or renting if

$$trP = \phi,$$

and households with relatively high tax advantages (high t) or high affinity for homeownership (low ϕ) will choose to own rather than rent. A common way of visualizing this “user cost” approach to tenure choice is to plot user costs for owners and renters over tax rates. Figure 2 illustrates the trade-off in this static version of the model: the owner-occupied user cost declines with the

household's marginal tax rate, while the rental cost does not vary with tax rates. Therefore, there is some tax rate t^* above which households choose to own and below which they rent. However, the location of this indifference point is shifted by the cultural affinity for homeownership. A household with a lower mental cost (higher affinity) for homeownership will have an owner-occupied user cost curve that is shifted down for all tax rates, leading to a lower indifference point t^{**} above which a household chooses to own rather than rent. That is, *higher affinity for homeownership results in higher homeownership rates.*

The model built thus far has two testable implications we take to the data in Section IV. First, those households with a high affinity for homeownership should have lower user costs than those with lower affinities; however, since ϕ is unobserved by the researcher, along with many other standard inputs to user costs not outlined above (such as the opportunity cost of downpayments, or the household's risk premia), we cannot credibly construct household level user costs without significant noise. Instead, we study how homeownership rates vary with affinities. Intuitively, households with lower user costs of owning will more often select into homeownership.

The model's second implication is that between t^{**} and t^* the probability that a household owns its home is increasing in ϕ ; below t^{**} everyone rents, and above t^* everyone owns. These thresholds are of course simplifications for two representative households, such that empirically we should observe, for a given tax rate, immigrant households with stronger cultural norms favoring homeownership will own at higher rates all along the tax schedule.

Saving for downpayments. In order to evaluate the role of changes in credit supply and how they interact with an affinity for homeownership, we need to make the model dynamic and introduce a role for the downpayment constraints α which did not matter in the static version. We do so by adding a period in which households need to save for the downpayment in the subsequent period, following Brueckner (1986). Now, households first live through a period 0 in which all households are renters, earn their income and can save for period 1. In period 1, they have the choice to become homeowners as before, but now have to finance the downpayment αP from their (weakly positive) savings s in the previous period. Moreover, any excess savings are invested with taxable returns r during period 1 and are available to be consumed at the end of the period.

Thus, in period 0, consumption of all households - whether future renters (R) or homeowners (H) - is given by

$$x_0 = (1 - t)y_0 - s - Q$$

where households may choose different savings rates s^R or s^H depending on their optimal tenure choice in the following period. The period 1 consumption amounts then become

$$\begin{aligned} x_1^R &= (1 - t)y_1 + (1 + (1 - t)r)s^R - Q \\ x_1^H &= (1 - t)y_1 + (1 + (1 - t)r)s^H - (1 - t)rP - \phi, \end{aligned}$$

where x_1^R and x_1^H denote the consumption of households that choose to rent and own, respectively,

and $s^H \geq \alpha P$ needs to hold for the homeowner.³ Each household maximises the objective function

$$u(x_0) + \theta u(x_1)$$

where $u(\cdot)$ is a strictly concave utility function and θ is the discount factor. We assume that owning always dominates renting for households for whom the downpayment constraint is not binding (i.e. for which $s^R \geq \alpha P$), and focus on the households for which the constraint is binding ($s^H = \alpha P$). The owner-renter utility differential for these constrained households is given by

$$\Omega = u(x_0^H) + \theta u(x_1^H) - u(x_0^R) - \theta u(x_1^R).$$

Credit supply shocks. If we think of a mortgage credit supply change as a change in the downpayment requirement α , we can derive the resulting change in the homeowner utility for constrained households (substituting $s^H = \alpha P$):

$$\frac{\partial \Omega}{\partial \alpha} = P(-u'(x_0^H) + \theta(1 + (1-t)r)u'(x_1^H)) < 0,$$

where the inequality follows from the fact that when the downpayment constraint is binding, s^H is higher than optimal and moves further away from its optimal level if the downpayment requirement increases. That is, higher downpayment requirements - a tightening of mortgage credit - reduce the utility from homeownership holding other parameters constant.

How does this effect of a tightening of mortgage credit supply vary with household affinity for homeownership? Differentiating with regard to the cultural cost of homeownership ϕ , we find

$$\frac{\partial \Omega}{\partial \alpha \partial \phi} = -P(1 + (1-t)r)u''(x_1^H) > 0,$$

so for households where the cost ϕ is larger, i.e. they have a *lower* affinity for homeownership, the negative impact of an increase in downpayments is mitigated, while higher affinity households experience a larger utility decline after mortgage credit tightens. Intuitively, households with a higher consumption benefit from being homeowners in the later period would benefit more from smoothing their consumption by shifting some consumption towards the earlier period when they are renters. However, a downpayment constraint requires them to lower consumption in the earlier period in order to save so they can afford to become homeowners. A higher downpayment constraint – and therefore higher required savings – lowers relative consumption in the earlier renter period even more for future homeowners. This leads some households to stay renters and forgo the painful period of low consumption and high savings that becoming a homeowner entails. This decline in the relative attractiveness of becoming a homeowner is stronger for high homeownership propensity households because they will already, on average, have chosen homeownership even at lower consumption levels in the first period, and pushing that level of early period consumption lower is therefore more painful for them.

Mortgage default for homeowners. How does behavior change *after* households enter homeownership? We are interested in how a higher affinity might impact the likelihood of defaulting on a mortgage among homeowners, when faced with similar financial circumstances. Assume that two otherwise identical homeowners (with the same observable income and tax rates) are faced with a circumstance where they can default on their mortgage. For simplicity, assume that the mortgage costs have increased, such that if the household does not default and become a renter again, it would reduce disposable income of the household by δ . We can think of these costs as being associated with mortgage structure, such as ARM resets, or variable HELOC rates (Bernanke, 2007; ?). Thus, the current homeowners face a choice between the following consumption bundles:

$$\begin{aligned} x^{\text{default}} &= (1 - t)y - Q \\ x^{\text{no default}} &= (1 - t)y - (1 - t)rP - \phi - \delta, \end{aligned}$$

where the utility gain from defaulting is given by

$$\Delta(y, \phi) = u(x^{\text{default}}(y, \phi)) - u(x^{\text{no default}}(y, \phi)).$$

Note that, with a strictly concave utility function $u(\cdot)$, a higher affinity (lower homeownership cost ϕ) household, always has a lower incentive to default, *ceteris paribus*, as they get more utility from being a homeowner for a given income, i.e.

$$\frac{\partial \Delta(y, \phi)}{\partial \phi} \Big|_y > 0.$$

Now, assume that default becomes more attractive (δ increases) because of a further increase in adjustable mortgage rates, or a decline in local income opportunities combined with an inability to move without defaulting (Fonseca and Liu, 2024b). Then,

$$\frac{\partial \Delta(y, \phi)}{\partial \delta \partial \phi} \Big|_y > 0.$$

That is, for a given level of income, a shock to the mortgage carrying cost that makes default more attractive has a bigger impact on low affinity households. The intuition is that high affinity households with the same income are experiencing higher levels of utility and thus lower changes in marginal utility from a given change in income while remaining a homeowner. This means that *low affinity homeowners are more likely to increase their defaults* when they encounter economic conditions that make it financially harder to pay their mortgages.

This simple model shows (1) how homeownership affinity can increase homeownership, (2) that households with high affinities respond more strongly to mortgage credit expansions than do their low-affinity peers, and (3) that high affinity households should have lower mortgage default rates, both on average, and in response to shocks that increase the incentive to default. We will test these hypotheses empirically in sections IV and V.

III. Data

A. *Measuring Affinities for Homeownership*

The concept of cultural norms or an “affinity” is not ex-ante well defined, but can be thought of as the mass lived experience or expectations around a given institution in a specific culture (or in this case, country). For example, Blau et al. (2011) use the relative female-to-male labor force participation rates in origin countries as a measure of cultural norms around women’s labor supply. Analogously, we define homeownership norms to be the rate of homeownership in one’s country of origin. We use a number of different approaches to try to measure the reference homeownership rate in the country of origin that is relevant for understanding tenure choices by foreign-born U.S. residents of different ancestry.

Homeownership rates in country of origin. One of the key concerns in measuring homeownership across countries is whether statistics provided by different national and international organizations measure the same concept and whether the sampled populations are representative of the underlying population. Moreover, to the degree that residents in different life stages will have different peers and reference groups in their home country, we need to take into account heterogeneity in homeownership within countries as well. To address these concerns, our main approach for constructing homeownership rates in the country of origin follows Marcén and Morales (2020) in using harmonized international census microdata from IPUMS (Center, 2020) to construct our own homeownership rates for comparable subgroups within each country. In particular, for each country we retain households where the head is 18-69 years old, and categorize each household further by the marital status of the household head.⁴ We compute separate homeownership rates for each marital status in the expectation that homeownership reference points related to these major life events most cleanly transfer to the cultural expectations of U.S. residents. Where countries have multiple census years available in IPUMS, we retain the sample that is closest to the year 2015. After dropping 5 censuses that had implausibly low average homeownership rates,⁵ we are left with homeownership data by marital status for 70 countries of origin in the US census data.

This baseline sample contains origin country data covering 72% of the birth places of all the foreign-born residents in the U.S. in the year 2000. Table I shows example data for the largest country of origin groups in the U.S. covered by this data. Note that variation between households in the mapping to origin-country homeownership rates can come either from variation in the homeownership for different marital statuses in their country of origin, or from variation in the marital status distribution among household heads from that country residing in the U.S.

An important advantage of using country origins of foreign-born residents to study the effects of preferences for homeownership is that there is large variation in the cultural backgrounds of foreign-born residents relative to observable differences between U.S.-born residents: Figure 1 highlights the large amount of variation in homeownership when comparing rates across countries, relative to comparing U.S. states. In the U.S., New York stands out as the state with the lowest homeownership rate, at an average of 50.86%, driven by the large population share in New York City, which has

a homeownership rate of around 30%.⁶ Eight states have homeownership rates below 60%, twelve have rates above 70%, but for the most part, states hover in the 60% region, consistent with a national homeownership rate in the mid-60’s. In contrast, there is much more cross-country variation in homeownership. Switzerland takes the bottom spot, with a homeownership rate of 31.5%. An additional 12 countries have mean homeownership rates below 60%. 57 countries have homeownership rates above 70%, and among those, 18 have rates higher than 90%.

Alternative homeownership statistics. As a robustness check, we also hand-collected data on average national homeownership rates (not disaggregated by marital status or age) from various international organizations and government statistics agencies. We were able to find data for 62 countries through this method.⁷ In the 50 countries where these two samples overlap, the correlation between the national average homeownership rates obtained through both methods is 87%.

B. Other Origin-Country Characteristics

Country economic characteristics. In order to analyze the drivers of homeownership in origin countries, we also collect a cross-sectional data set of country characteristics from official data sources. In particular, we collect data on GDP per capita, inflation rates, urbanization rates, and domestic credit for 2000-2020 from the World Bank’s World Development Indicators. In addition, we also use indices of property rights protection and investment freedom from the Heritage Foundation for the years 2000-2008.

Major religion by country. To assess the role of particular determinants of culture, we also collect a data set of the major religion of each country in our sample. We manually access data from the United Nations Statistics Division, Wikipedia, and the CIA World Factbook (in this order) on a country’s population by religion. We combine different religious denominations into the categories Christianity, Islam, Buddhism, Hinduism, and other (which includes Atheists and Agnostics), and assign to each country the category with the largest population share. For the few cases (e.g. China) where official sources are not clear on which of the major religions best describes the majority practice in the country, we assigned “Other”.

C. Tenure Choice and Household Finances

American Community Survey (ACS). In order to measure household-level tenure choices for foreign-born U.S. residents, we use microdata from the American Community Survey for the years 2005-2019.⁸ This is a repeated cross-section of a representative sample of U.S. households, which contains ~3.4 M households with foreign-born heads who are 18-69 years old whom we are able to match to homeownership rates in their countries of origin. The key variable in this data set for our analysis is the measure of whether the household owns the home that it lives in. In addition, this data allows us to measure the characteristics of the head’s spouse, including their country of origin, which enables analyses that measure the effect of within-family differences in origin countries on household tenure choices. Moreover, we construct the following control variables at the household and individual level from this data: household income, education of the household head and their

spouse, years since immigrating of the head and their spouse, age and gender of the head and their spouse, number of children living in the household, race of the household head, and marriage status. **American Housing Survey (AHS)**. We augment the ACS data with data from the American Housing Survey. The AHS is also a repeated cross-section, but instead of surveying a representative sample of households, its focus is on a sample of housing units. We use data from 2001-2019, covering ten waves of the biannual survey, which yields 175k households over 585k observations. We have HOCO for about 30k of the households with foreign-born heads. Like in the ACS, we observe many socioeconomic and demographic characteristics of the households, including household income, education of the household head, years since immigrating of the head, age and gender of the head, number of children living in the household, and marriage status. While the ACS microdata provides us with a public use microdata area (PUMA) which we can map to a commuting zone, the AHS data only provides 15 core-based statistical areas (CBSA), or a catch-all non-CBSA code. The AHS also provides detailed data on mortgage origination and other financial characteristics of the household.

Panel Study of income Dynamics (PSID). One confounder of concern could be that immigrants from high homeownership countries come to the U.S. with higher levels of wealth than those from low homeownership origins. Wealth is not observable in either the AHS or ACS data. As such, we utilize a set of shared socioeconomic, demographic, location, and origin variables to impute wealth, using loadings calculated from the PSID. We use the 1997 and 2017 waves of the PSID to construct household wealth, as these two waves refreshed the PSID household sample, adding new immigrant households to the panel. We then regress household non-housing net worth on a suite of socioeconomic and demographic characteristics, as well as origin and destination fixed effects, to estimate their expected contributions to net worth. Finally, we use these coefficients to construct non-housing net worth estimates in either the AHS or ACS data. See Appendix A for details.

D. Measuring Credit Supply

In order to measure local mortgage market outcomes, mortgage lending by lender, and to construct mortgage market liquidity shocks we use data collected under the **Home Mortgage Disclosure Act (HMDA)**. We obtain the loan-level data for 2007-2017 from the Consumer Financial Protection Bureau’s website, which provides information on the lender and location of the property as well as property and loan characteristics, and supplement it with loan-level data for 1998-2006 available from the National Archives. We match lenders to their parent company using the crosswalk maintained by Robert Avery⁹ and aggregate the data at the level of total loans and loan volume by lender-county-year. We retain only first-lien purchase loans that were originated for single-family (1-4 units) site-built properties that are owner-occupied, in order to focus our analysis on the type of lending that is most likely associated with a household’s transition into homeownership. We construct a panel of lender-county-year data for 1998-2017.

E. Portfolio Allocations

In order to study the long-run household portfolio consequences of an affinity for homeownership, we use data from the Health and Retirement Study (HRS), which is provided by the University of Michigan (2024). This data set is uniquely useful for our question of how tenure choice preferences affect household portfolios in the long run, as it (1) is a nationally representative survey with detailed information about household finances, (2) captures household portfolios at the age of retirement, and (3) contains information about the country of origin of foreign-born participants. We obtain access to the restricted version of the HRS data, allowing us to link current participant location and national origins to their financial information. This data allows us to provide some of the first information about differences in the end-of-working-life portfolios of immigrants from different origin countries.

Within the HRS, we use the RAND HRS Longitudinal data (which combines and standardizes information from different waves) and, in order to make the included households comparable, focus on married / partnered households, for which we retain one person (the financial respondent where applicable). Among persons representing a household, we keep only respondents that are 55-65 years old at the time of at least one survey wave and we find the survey wave that comes closest to them being 60 years old. We construct a number of different variables characterizing respondents' retirement portfolio size and composition as well as real estate ownership, then merge this information with their country of origin information, including HOCO. See Appendix B for further detail on variable construction and data transformation in the HRS.

The key outcome variables that we focus on are whether households own real estate in their retirement portfolio, the value of that real estate, the value of different components of their household wealth and the share allocated to risky assets in the form of stocks.

IV. Preferences for Homeownership and Tenure Choice

In order to establish the baseline effect of homeownership in the country of origin on tenure choice in the U.S., we use household-level data to estimate specifications of the following form, among foreign-born U.S. residents:

$$HO_{ht} = \gamma HOCO_{ht}^{head} + \beta' X_{ht} + \zeta_{CZ,t} + \varepsilon_{ht} \quad (1)$$

Here, $HOCO_{ht}^{head}$ is the homeownership for householders, h , with the same marital status in the country of origin of the household head; X_{ht} is a vector of control variables consisting of household and individual characteristics; $\zeta_{CZ,t}$ represents fixed effects controlling for common variation at the commuting zone or year level, t . In fact, the full specification for most of our analyses flexibly includes commuting-zone-by-year level fixed effects, which control for the fact that immigrants might sort into cities with trends in housing markets or local labor markets that also affect homeownership rates. In this section, we establish that there is indeed an effect of homeownership in the country of

origin (HOCO) on immigrants’ tenure choice in the U.S., and that this effect is robust to plausible confounding effects.

A. *Heterogeneity in Ownership Along the Tax Schedule*

In order to motivate the more formal regression estimates that follow and provide some intuition for our result in line with the user cost model that is commonly used to discuss tenure choice, we first consider how a higher affinity for homeownership affects actual homeownership among immigrants as a function of household tax rates.

As noted in Section II, the static user cost model implies immigrant households with stronger cultural norms favoring homeownership will own at higher rates all along the tax schedule. To test this implication of the model, we use the NBER’s TAXSIM model (v35) to compute federal marginal tax rates for all immigrant households in our sample. Figure 3 plots a binned scatter-plot of homeownership rates by federal marginal income tax rate, splitting our immigrant sample population into those with above-median HOCOs and below-median HOCOs. Importantly, the tenure choice model outlined above has a household deciding whether to own or rent the *same* unit, therefore we include CBSA-by-year fixed effects to proxy for location and labor market, as well as a suite of property characteristics including bedrooms, bathrooms, acreage, square footage, the home’s age, type (i.e. single- or multi-family), and the presence of garage or car-port. These controls adjust for the fact that the researcher cannot observe the counterfactual rental (purchase) price a homeowner (renter) considered before buying (renting).

The figure shows that immigrants with stronger cultural norms for housing do select into homeownership at higher rates along the entire tax schedule. This gap between above- and below-median HOCO marginally attenuates at the top marginal income tax rates as predicted by the model, since the benefits of owning are highest here for all groups. These findings suggest that high-HOCO households may end up being observationally similar to marginal homebuyers, as they are selecting into homeownership at lower marginal tax rates, and thus lower incomes due to tax rate progressivity. Moreover, these results support the idea that differences in homeownership associated with affinities stemming from different origins are likely not explained by differences in financial circumstances, which we would expect to be strongly correlated with imputed tax rates.

B. *Results: Baseline household level effects of HOCO on homeownership*

To determine the size of the effect of country-of-origin homeownership on tenure choices among foreign-born U.S. residents, we estimate equation 1 in the household level ACS data. The results are shown in Table II. Each column shows the estimated effect of homeownership in the country of origin of the household head, matched by marital status, on whether the household owns their home in the U.S. The first column shows the raw association between the two variables, while the second column flexibly controls for a large number of characteristics of the household, such as income, household size, children, years since immigration, and age and education of the household head. These control variables capture other factors that are known to drive tenure choices. Additionally,

they address the concern that immigrants from countries with high homeownership rates may be incidentally selected with regard to characteristics that drive higher homeownership rates in the U.S, such as being high-income. That the relationship survives these controls and remains statistically significant, shown in column 2, suggests a direct link between the homeownership experiences in immigrants' origin country and their choices in the U.S.

One important concern is that immigrants from different countries arrive in the U.S. with different levels of (non-housing) wealth that are systematically related to homeownership in their country of origin, and which affect their ability to acquire a home. To control for this possibility, we use PSID data to estimate the relationship between household wealth and a rich set of household characteristics and then apply those estimates in the larger ACS sample to impute household net worth. While this is likely not a perfect proxy for wealth, when we include this control in the estimation in column 3 of Table II, the coefficient is only modestly reduced relative to the one estimated without this wealth proxy. Moreover, the analysis in Section IV.A showed that the gap in homeownership between low and high HOCO immigrants is quite stable along the marginal tax rate that households face, suggesting that household finances are likely not a key driver of the gap. Similarly, we show in Section IV.D below that the effect is larger, the *longer* a migrant has been in the U.S., which makes wealth upon arrival an unlikely explanation.

Additionally, we want to control for the possibility that immigrants from higher homeownership countries might be sorting into labor markets or housing markets (here captured by commuting zone boundaries) that are more or less conducive to homeownership on average or which experience rising homeownership rates during the years in the sample. Column 4 therefore includes commuting zone-by-year interacted fixed effects. That is, our preferred estimate in column 4 identifies only off within-labor-market variation in tenure choices within each year between immigrant households that are identical with regard to their observable characteristics noted above, except in the household head's country-of-origin. The estimated elasticity of homeownership with regard to *HOCO* in this specification is 15%, and is significant at the 1% level.

To put this number in perspective, note that it means that the interquartile range of *HOCO* across households in our sample, which is between 58 ppt and 76 ppt, is associated with a 2.8 ppt difference in homeownership rates. For comparison, this represents 44% of the standard deviation in homeownership across U.S. states, or 39% of the gap in homeownership rates between native and foreign-born households, conditioning on all the controls in Table II. This number is also quite similar to findings from Ringo (2020), who, using U.S. *county*-of-origin homeownership rates at age 14 as an instrument for an individual's likelihood of owning a home later in life, finds that a one-standard deviation increase in county-of-childhood homeownership rates is associated with a 2pp increase in homeownership.¹⁰

To visualize how different countries contribute to this effect, we can aggregate the residualized household homeownership rates and marriage status *HOCOs* to the origin country group level. We plot the relationship between the two variables in Figure 4. The top graph shows the raw averages of homeownership in the U.S. among foreign-born residents, plotted over the averages of the marital

status matched *HOCOs*. The lower graph residualizes these variables with regard to the full set of control variables in column 4 of Table II before aggregating. While the marker sizes are not weighted by number of households in the U.S., both graphs use a weighted-linear fit overlaid in the dashed red line, which accounts for differential numbers of immigrants from origin countries. These graphs make it clear that homeownership rates vary widely between countries, with average rates conditional on marital status of more than 90% in Romania, Hungary, and Cambodia at the upper end, and rates below 50% in Turkey, Ghana, and Switzerland at the lower end.

It is reassuring that no obvious grouping of countries by geography or wealth emerges from these graphs. For instance, several low-income countries can be found both among the highest and lowest *HOCO* and U.S. homeownership countries. Moreover, the most visible outliers from the fitted linear relationship in the residualized graph are Switzerland and South Sudan, which only constitute 0.14% and 0.0009% of the U.S. foreign-born population.

C. Robustness checks

In this section, we explore different specifications that control for potential omitted variables and alternative explanations *other than country of origin* that could explain this association between *HOCO* and homeownership among immigrants in the U.S.

Origin country characteristics. An important concern with regard to the mechanism for our findings in Table II is that they may reflect omitted variable bias arising from immigrant households bringing other aspects of the lifestyle or economic situation of their country of origin with them, which then incidentally affect their homeownership in the U.S., without a particular preference over the latter. We consider a number of plausible home country characteristics that might be driving homeownership both in the U.S. and the origin country – their unconditional correlations with origin country average homeownership rates are shown in Appendix Figure BI: On the one hand, we consider economic characteristics, such as GDP per capita, urbanization rate, annual inflation rate, and domestic credit to private borrowers as a share of GDP. These capture the idea that low wealth, exposure to economic risk, urban lifestyles, or a lack of access to credit may all keep homeownership in the country of origin artificially low, and this lack of resources and lifestyle choices may be replicated in the immigrants’ experiences in the U.S. Note, for example, that homeownership is unconditionally negatively correlated with urbanization at a country level as shown in Appendix Figure BI.

On the other hand, the legal system and property rights may make homeownership a more or less attractive investment relative to the security of, and access to, other assets. The persistence of the resulting attitudes towards homes as part of a household portfolio may be part of the cultural differences that explain immigrants’ homeownership in the U.S. To capture these dimensions, we consider indices of property rights and investment freedom created by the Heritage Foundation. As can be seen in Appendix Figure BI, these indices are highly positively correlated with GDP per capita and urbanization, and associated with lower inflation. Moreover, some studies suggest that religion may play a role in tenure choices, for example through their positions on usury (e.g.

Das, Coulson, and Ziobrowski (2019)), so we also test how these results change when including fixed effects for which of the major religions (Buddhism, Christianity, Islam, Hinduism, Other) predominates in the country of origin.

We test the importance of these confounding country-of-origin characteristics by including all of them as control variables in the estimation. The results are shown in column 1 of Table III: the *HOCO* effect on U.S. homeownership is consistently positive and significant in all specifications, and the magnitude of the coefficient is slightly bigger and more precisely estimated even when all country characteristics are controlled for. Among the country characteristics ((see Appendix Table AI for the full set of coefficients), only urbanization has a significant (and positive) effect on U.S. homeownership when all of our baseline controls are already included. These results suggest that each country’s homeownership norms evolve through a complex system of institutions, culture, etc., none of which completely explain the variation in norms.

Immigrant group selection. Another potential concern with our estimates above is that immigrants from particular countries are *selected* with regard to their tendency to buy a house when in the U.S. For example, if migrants from high homeownership origin countries tend to be selected to be particularly wealthy or more educated than average, e.g. because only the elites are able to pay the costs associated with migrating, then these selected characteristics might drive homeownership in the U.S. This concern should be mitigated by the fact that our baseline regressions already directly control for foreign-born households’ income, education, and household size, such that any such selection would have to be orthogonal to these observables. However, to additionally ensure that our results are robust to the *average* characteristics of migrants from that origin, column 2 of Table III shows a version of our regression where we also control for the average share of college-educated household heads, mean household income, mean number of children and mean household size among migrants from that origin country. In addition, we control for the GDP per capita of the origin country, in case immigrant selection is driven by the economic opportunities in the country of origin. Given that some of these migrant characteristics may actually be affected by homeownership, they may represent “bad controls” (Angrist and Pischke, 2008) that shut down the causal channel of interest – and the estimates including these controls should therefore be interpreted with caution. As the results in column 2 show, while the estimated effect size is reduced when including these additional controls, it is nonetheless statistically significant and implies an elasticity of 14% of U.S. tenure choice with regard to the household *HOCO*.

Largest origin groups. As Table I shows, the share of U.S. migrants from different origin countries is highly concentrated. As a result, it is important to establish whether our *HOCO* effect estimates are driven by idiosyncrasies of the biggest origin countries or also hold for the smaller foreign-born origin groups in the U.S. In column 3 of Table III, we therefore exclude any households with heads from one of the five largest immigrant groups in the U.S. as of the year 2000 (Mexico, Philippines, India, China, Canada) from the sample. The results show the estimated *HOCO* elasticity for this smaller sample is of a similar size, at 15%, and significant at a 1% level, which suggests that the large origin countries are not driving the baseline results.

HOCO measurement. Our measure of homeownership preferences assigns *HOCO* measured in international census microdata by household head marital status. In column 4, we show that our results are not due to mismeasurement of *HOCO* in this international census microdata. We hand-collect average national homeownership rates from official government sources or international organizations like the OECD. This sample covers a smaller number of countries (62, compared to 74 in our microdata sample). However, the *HOCO* constructed from the national average homeownership rates in this alternative data set still has an effect on U.S. tenure choice that is significant at the 10% level, and the estimated effect size of 15% is again very similar to our baseline estimate using the microdata *HOCO* by marital status.

Alternative household sample. The other concern around sample selection may be that the American Community Survey household data has a sampling methodology or differential response rates that correlate both with *HOCO* and U.S. tenure choices. In column 5, we instead use data from the American Housing Survey, which collects a much smaller sample using a different methodology focused on being representative of housing units in the U.S. and tracking them over time. When we estimate our baseline effect in this separate data set of only $\sim 46\text{K}$ foreign-born householders, we obtain a significant estimated *HOCO* elasticity of 13%, so our baseline effect does not seem to be driven by methodological idiosyncrasies of the ACS sample.

D. Mechanisms

To better understand what mediates the effect of country-of-origin homeownership on the choices of U.S. immigrants, we can explore which characteristics of the household head affect the strength of the effect.

Second-generation immigrants. One subgroup of foreign-born immigrants that is unlikely to be selected to immigrate based on their own unobservable characteristics, but may nonetheless be culturally influenced by the attitude towards homeownership in their country of origin are second-generation immigrants. While the ACS data does not allow us to identify countries of origin for the parents of residents born in the U.S.,¹¹ we can identify immigrants who came to the U.S. as children (18 years old or younger), and consider whether our results hold for this subgroup. Table IV, column 1, shows that even though the sample size is substantially reduced for this analysis, the pass-through of *HOCO* to the children’s decision to own a home as adults later in life is only modestly attenuated.¹² We estimate that immigrants who arrived as children have a 12% passthrough of homeownership attitudes from their country of origin. This is consistent with a related literature that has shown that an affinity for homeownership is transmitted from parents to children, which in this case ends up propagating the influence of the country of origin.

Gender. An additional dimension of household characteristics that might matter for the transmission of *HOCO* into tenure choices is the gender of the household head. To explore this possibility, we estimate a specification that interacts the *HOCO* variable with indicators of the householder’s gender. The results are shown in column 2 of Table IV. We find that the effects of *HOCO* on households headed by men are smaller by about a quarter of the original effect size. While the

gender of the householder may be associated with other differences in household structure, the fact that the transmission is significant and large for both households headed by both genders suggests that the effects are likely not driven by differences across immigrant groups in the prevalence of female-headed households.

Education. As higher education enables immigrants to attain more knowledge about their host country and greater exposure to other perspectives, it is commonly associated with higher rates of assimilation (Dustmann, 1996). In column 3 of Table IV we explore whether this holds for homeownership effects: we find that the effect of *HOCO* on tenure choice is about a third smaller for college-educated householders.

Time since arrival. If the *HOCO* effects represent something like the expression of a cultural reference point at different life stages, rather than an endowment when arriving in the U.S., immigrants would be expected to take a couple of years to save up and establish themselves before taking the decision to settle somewhere and buy a home in line with their *HOCO*. In contrast, if the mechanism is a more short-lived attitude or social pressure to own a home that was prevalent in the country of origin, we would expect the effects to fade over time in the new country. We test these ideas in column 4 of Table IV by comparing the *HOCO* effect among immigrants who arrived in the U.S. less than 10 years before the survey, and long-term residents who have been in the U.S. for a period of more than 10 years. The results show that the *HOCO* effect is significantly larger among long-term foreign-born residents, and smaller and not statistically significant among recent immigrants. In addition to showing that the effect lasts even years after arrival, this finding also allays some potential concerns with our *HOCO* effect estimation: on the one hand, this finding is not compatible with reverse causality, where immigrants move to the U.S. *because* they acquire a home, for instance as an investment, which may be driven by real estate dynamics in the country of origin. On the other hand, it makes it less likely that unobserved differences in the wealth of new arrivals from different countries— those not captured by the financial control variables – drive the effect.

Origin of the spouse. If culture matters, we would expect married householders' tenure choices to also be influenced by the origin of their spouse. Marrying a U.S.-born spouse might indicate a greater desire to assimilate, which should be associated with a lower importance of cultural factors for the householder. In contrast, if the mechanism for *HOCO* effects operates through a greater level of sophistication with regard to the U.S. financial system, marrying a U.S.-born spouse would, if anything, indicate a greater ability to exploit this advantage. Column 5 of Table IV includes only married householders and interacts the *HOCO* variable with whether the spouse is U.S.-born: we find that the *HOCO* effect is almost entirely undone for immigrants that marry a U.S.-born spouse, which seems aligned with a cultural explanation for the effect. If the spouse is not U.S.-born, the householder may also consider the *HOCO*-based preferences of their spouse when making tenure choices: Column 6 shows the independent effect of the spouse's affinity for homeownership, as well as an interaction for whether both spouses are from the same country. We find that the effect of a foreign-born spouse's *HOCO* is also significant and statistically

indistinguishable from that of the householder.

To summarize, this section shows that the effects of origin country homeownership vary in a way that is consistent with a “cultural” explanation for this financial behavior. For example, foreign-born householders who are more likely to assimilate in the U.S. because they went to college or married a U.S.-born spouse are less likely to act in line with their HOCO. At the same time, the effects persist long after the householder immigrated, or if they arrived as a child, and the spouse’s origin matters to a similar degree. All of these factors point to an effect that is more latent than the economic circumstances of the householder when they first arrive in the U.S.

V. Credit Cycles and Affinities for Homeownership

The previous section established that the homeownership in foreign-born U.S. residents’ countries of origin passes through into actual tenure choices when living in the U.S.—and that this effect is robust to different specifications and control variables. One of the key novel predictions from the conceptual framework in Section II is that these differences in affinity for homeownership have important consequences for the effects of financial shocks and policy choices on homeownership. In particular, we propose that a higher desire to own a home makes households more responsive to financial shocks, such as changes in the availability of credit, which then changes their housing market outcomes.

Testing this hypothesis is one of the key goals of this paper because how different demographic groups respond to credit expansions can have important impacts on their financial situation. For instance, other researchers have shown that racial minority households in the U.S. saw large increases in homeownership during the peak of the housing boom of the mid-2000s, which then resulted in an increased risk of mortgage delinquency and foreclosure during the subsequent bust (Bayer et al., 2016). Moreover, homeownership is associated with differences in child outcomes, mobility, and many other aspects of household behavior (Dietz and Haurin, 2003). As a result, differential responsiveness to mortgage credit supply shocks can lead to differences in a number of group outcomes, which has important implications for the effectiveness and welfare consequences of policies that try to promote homeownership.

In this section, we show that an affinity for homeownership resulting from foreign-born residents’ country of origin substantially affects the impact of mortgage credit supply shocks on housing market outcomes

A. *Who benefited from the 2000-2005 Credit Expansion?*

The conceptual framework in Section II suggests that, among renters, high HOCO households should be more sensitive to credit supply expansions in whether they eventually transition into homeownership. Moreover, the model implies that this effect should be stronger among households that are more financially constrained in their decision to own a home.

However, when considering aggregate changes in homeownership across HOCO groups in response to a credit expansion, it will also matter whether the *share* of households in each group that is “marginal,” those previously indifferent households for whom credit expansion changes their tenure choice, is large. For instance, in the cases where most high HOCO households that are renters are so financially constrained that they could never own a home, or if most are so rich that they always own a home, even a high sensitivity to credit expansions would not translate into a large change in homeownership for the high HOCO group.

To provide intuition for the regression analysis that follows, we therefore first examine the question of whether the increase in homeownership during the 2000-2005 boom in housing markets—which is often attributed to an expansion in credit supply (e.g. Mian and Sufi (2009); Justiniano et al. (2019); Mian and Sufi (2022))—affected high HOCO households differently.

To capture the degree to which homeownership became more attainable for households with particular demographics, we use linear models to predict the propensity for homeownership in the cross-section in either 2000 or 2005 based on household characteristics, excluding HOCO itself.¹³ Then, applying the coefficients for predicting homeownership in 2005 to a household’s year 2000 characteristics and comparing the results to the homeownership predicted using year 2000 coefficients, tells us how much easier it got to become a homeowner from 2000 to 2005 for households with a particular suite of characteristics. This compares predicted homeownership status under two credit regimes, holding the household fixed.

Figure 5, panel A, shows how the characteristics of households with different HOCOs aligned with these changes in the propensity to attain homeownership. The positive slopes in the graph show that households with higher HOCOs were on average more likely to have characteristics associated with increased homeownership during 2000-2005 boom. That is, high HOCO foreign-born households were more “marginal”, on average, with regard to the changes in how easy it became to buy a home. Comparing the lines for renters and homeowners in the graph, we find that this was particularly true for renters with high HOCOs: these households, which were culturally inclined to *want* to become homeowners, but had yet to achieve that goal, saw a much larger increase in their propensity to own a home.

Were these larger increases in homeownership due to high HOCO renters being less financially constrained to begin with? Panel B of the same figure shows the relationship between the pre-boom predicted propensity to own a home—as a proxy for a lack of fundamental financial constraints—and the change in the homeownership during the boom. The graph compares this relation for low and high HOCO groups. We can think of the left side of panel B as comparing two foreign-born household heads in the year 2000 that were both unlikely to own a home based on their demographics, i.e. they were both constrained, financially or otherwise, in their ability to become a homeowner—but one household is low HOCO and the other is high HOCO. Then, the constrained high HOCO household’s particular mix of demographics was associated with a large increase in homeownership from 2000 to 2005, while a similarly constrained low HOCO household did not see a comparable increase in homeownership. This relationship reverts for less constrained households

on the right side of the graph: unconstrained low HOCO households were more likely to increase their homeownership between 2000 and 2005 than high HOCO households.¹⁴

Together, these analyses show that, during 2000-2005, high HOCO renters were particularly likely to increase their homeownership, and this effect was larger for households that were more constrained in their ability to buy a home before the boom. This pattern aligns with credit booms having particularly large effects among the households that are most eager to own a home for reasons related to their cultural origins.

Building on this intuition that high HOCO households are more marginal during credit-induced housing booms, in the next section, we show that exogenous increases in credit supply cause systematically larger changes in homeownership among high HOCO households than among low HOCO households.

B. Identification of mortgage credit supply shocks

In order to estimate the impact of mortgage credit on homeownership, we need to address the fact that a household’s access to credit is likely to be endogenous with regard to its income and wealth, which are affected by a group’s homeownership rates, and may also be affected by racial discrimination (Ambrose, Conklin, and Lopez, 2021). To estimate the causal effect of increased access to mortgage credit on homeownership rates among groups with different countries of origin, we therefore need to identify changes in credit supply that are plausibly exogenous with regard to group characteristics.

Our baseline analysis uses mortgage credit supply shocks based on county exposure to heterogeneous shocks to aggregate mortgage lending by different banks, following Gilchrist et al. (2018) and Garcia (2020). The intuition for the approach is that local mortgage lending may change either due to local changes in credit demand or due to a change in lenders’ willingness to originate new mortgages for idiosyncratic or common reasons - shifts in “credit supply” - that are unrelated to local market dynamics. To the degree that the idiosyncratic shifts in bank mortgage credit supply have a greater impact on counties where the affected banks have a higher market share, they can generate variation in local lending that avoids concerns over reverse causality from local housing market dynamics affecting lending.

To estimate lender fixed effects that are orthogonal to local mortgage demand shocks we run regressions of the following form: for any county c and lender j , we proceed in two steps, adapting the method in Gilchrist et al. (2018). First, we estimate

$$\Delta \ln L_{c,j,t} = \alpha_{c,t} + \eta_{j,t} + \varepsilon_{c,j,t}, \tag{2}$$

where $\Delta L_{c,j,t}$ is the change in the number of loans originated by lender j in county c over period t ; $\alpha_{c,t}$ are locality-by-year fixed effects, and $\eta_{j,t}$ are lender-by-period fixed effects. We keep only lenders in the sample in each period that have lending activity in at least 3 counties. We follow Gilchrist et al. (2018) and weight the data points by the geometric mean of the lender’s market

share in the county and the county’s share in the lender’s activity at the beginning of each period. Moreover, as we are interested in exposure to *relative* differences in lender activity, we re-center the estimated $\hat{\eta}_{j,t}$ terms for each period such that the loan-origination-weighted mean of lender shocks is zero.

In order to gain more power in estimating the effects of lending on homeownership—which varies only slowly and cyclically—and align the magnitude of our results with the housing cycles of interest, we estimate the effects of credit supply shocks on homeownership in long difference panels. We take long differences for the periods 2000-2005, 2005-2010, and 2010-2015 in lending when constructing the credit supply shocks (and also for the corresponding homeownership at the local group level). That is, the left-hand side of equation 2 is the change in lending from 2000 to 2005 by lender j in county c , the change in lending from 2005-2010... etc. Note that if a county’s lending declined because of local economic shocks, this will be captured by the county-by-period fixed effects, while a change in lending due to a lender’s change in ability to fund mortgages, independent of the dynamics of the markets in which the lender is present, will be captured by the lender-by-period fixed effects.

The shock to local mortgage lending in a county that is driven by exposure to lenders that are experiencing aggregate changes in their ability to fund mortgages is then constructed as the market-share weighted average in each county of the lender fixed effects:

$$S_{c,t} = \sum_{j \in C} \underbrace{\frac{L_{c,j,t}}{\sum_{j \in C} L_{c,j,t}}}_{\substack{\% \text{ of local lending} \\ \text{by lender } j}} \times \underbrace{\hat{\eta}_{j,t}}_{\text{Re-centered lender FE}}$$

Counties will experience below-average access to mortgage lending if they have greater exposure to lenders who curtail more national lending. To further purge this raw credit supply shock of county demand characteristics, e.g. the concern that over the period of study banks with credit supply shocks in a particular direction were associated with counties that are systematically different - beyond the level of credit demand in each period that is controlled for in equation 2. We estimate

$$S_{c,t} = \gamma_1 \underbrace{\hat{\alpha}_{j,t}^{\text{Loan}}}_{\substack{\text{Re-centered County} \\ \text{Loan count FE}}} + \gamma_2 \underbrace{\hat{\alpha}_{j,t}^{\text{Lend}}}_{\substack{\text{Re-centered County} \\ \text{Lending vol. FE}}} + \lambda_c + \lambda_t + \chi_{ct}, \quad (3)$$

weighting each county-period by its IRS population count. Here, $\hat{\alpha}_{j,t}^{\text{Loan}}$ and $\hat{\alpha}_{j,t}^{\text{Lend}}$ are estimates of the county-period fixed effect from equation 2 when the dependent variable is either the log change in loans or the log change in dollar lending volume, so they capture the extensive and intensive margin of changes in local credit demand. Moreover, we also control for county fixed effects and period fixed effects. We use the residual from this regression as our estimate of the credit supply shock in each county that originates from exposure to particular banks that is unrelated to local credit demand shocks, i.e. our main credit supply shock is $\tilde{S}_{ct} = \hat{\chi}_{ct}$.

The identifying assumption for using this shock to study the effects of exogenous shifts in credit

supply on homeownership is that exposure to lender fixed effects is not systematically correlated with differential within-county shocks to the demand for homeownership by low- and high-HOCO residents. This assumption could be violated, for example, if banks that lend relatively more in some geographic areas experience mortgage funding constraints at the same time as the demand for homeownership shifts in relative terms between low- and high-HOCO groups in the same areas in a way that is not correlated with general shifts in extensive and intensive margin lending in the county. Given the construction of the instrument above, we believe that this particular violation of the identifying assumption through correlated local exposure to particular banks and demand shifts is both unlikely and, if it occurs, correlated with permanent differences in county demographics, such as income or family structure. To address that possibility, we therefore also control for HOCO-group-by-county fixed effects in our estimations below.

What is the geographic variation in credit supply shocks underlying these results? Appendix figure BV shows the pattern of these credit supply shocks for the years 2000-2005 and 2010-2015, which are the beginning and end of our credit shock panel respectively. Note that there is both substantial variation in the shock size across counties, even within the same region or state, and also over time, as the regions with the largest credit supply shocks vary substantially across these two periods.

C. Credit Shock Sensitivity Estimation Approach

We use the credit supply shocks estimated off nonlocal-lending behavior by banks to identify the differential transition into homeownership across HOCO groups. We utilize a stacked 3-period panel over the periods 2000-2005, 2005-2010, and 2010-2015, which roughly correspond to the boom, bust, and rebound cycle in U.S. housing.

County-level effects. To verify that the credit supply shock constructed above is not just exogenous but also relevant, we estimate the effect of the shock on county-level lending, akin to the “pseudo-first stage” for the reduced-form estimation of within-county group level differences resulting from these shocks. At the county-by-period level, we estimate regressions of the form

$$\Delta \ln \text{Loans}_{ct} = \alpha_c + \alpha_t + \beta \tilde{S}_{ct} + \gamma' X_{ct} + \epsilon_{it},$$

which include both county and period fixed effects, as well as commuting-zone-by-period fixed effects that capture regional trends.

Within-county groups. One limitation of the household-level ACS data is that it represents a repeated cross-section of a representative sample of U.S. households, but individual households cannot be tracked over time. However, cross-sectional analysis of credit shock effects on homeownership is not suitable for our purposes: as we showed earlier there are large cross-sectional differences in *levels* of homeownership across different areas, which could be correlated with static geographic differences in exposure to lenders. To control for these time-invariant levels, we exploit variation in credit shocks *within* counties over time. Moreover, mortgage credit supply shocks al-

low marginal households to enter homeownership, and should therefore mainly be associated with *changes* in homeownership. Thus, the dependent variable should be in differences, which also requires variation over time.

Under the assumption that any sampling bias of different HOCO groups does not vary over time, we can aggregate the household-level ACS data used in earlier analyses into group level data for HOCO-county-year cells. We define 5 different groups g in the household data – US-born, high-HOCO foreign-born, low-HOCO foreign-born, Mexican-born, and no-HOCO data foreign-born – so that we can compare relative effects for foreign-born residents with different affinities for homeownership. High and low HOCO are defined as above and below the population-weighted median HOCO among non-Mexican foreign-born. We separate the Mexican-born cohort because they constitute the largest group of foreign-born residents and have HOCO values close to the median, such that minor differences in weighting can shift large foreign-born population shares between the low and high-HOCO categories, making the results highly sensitive to the idiosyncratic characteristics of this group. For each group-county-period cell we compute average homeownership rates, weighting each household by its ACS population weight.

Then, at the group-county-period level, we estimate regressions of the form

$$\Delta HO_{c,g,t} = \phi \tilde{S}_{c,t} + \beta_g \tilde{S}_{c,t-1} \times \mathbb{1}[HOCO_g] + \alpha_c + \alpha_t + \zeta_{c,g} + \psi_{CZ,t} + \varepsilon_{it}, \quad (4)$$

where $\Delta HO_{c,g,t}$ is the change in the local homeownership rate of group g in county ct over time; $\mathbb{1}[HOCO_g]$ is a categorical variable indicating the different groups, omitting the low-HOCO group as the reference group; α_c and α_t are county and period fixed effects; our most stringent specification also contains county-group level fixed effects $\zeta_{c,g}$ which control for the fact that some areas may generally make it easier or harder for groups with different affinities to move towards homeownership, leading them to have generally higher or lower growth in homeownership in our sample; and CZ-by-period fixed effects $\psi_{CZ,t}$ that account for regional trends in homeownership over time.

D. Affinity for homeownership and sensitivity to credit supply shocks

How does an affinity for homeownership effect the average response of households to mortgage credit shocks? The results of estimating equation 4 for the stacked long differences across periods 2000-2005, 2005-2010, and 2010-2015 are shown in Table V. The first two columns show the baseline effect of the mortgage credit supply shocks on the county-level change in new mortgage loans observed in HMDA. We see a positive increase in loan volume in response to credit shocks, confirming that these induced more potential homebuyers to submit loan applications.

Columns (3) and (4) then show the impact of these credit supply shocks on average homeownership changes, without distinguishing by HOCO group. The point estimates suggest that after a 1 pp increase in credit supply, homeownership grows by 1-1.7 pp. This confirms that the exogenous variation in mortgage credit supply in fact leads to significant changes in homeownership, such that this source of variation is plausibly large enough to pick up differential group-level responses.

Finally, columns (5) and (6) show the results where we allow the credit shock effects to vary across HOCO groups. The low HOCO foreign-born category is the reference group - such that the coefficient in the first row represents the baseline level of the credit shock effect for the low HOCO foreign-born—and all other coefficients are defined relative to it.¹⁵ Column (5) includes only period and county-group fixed effects, while column (6) also adds commuting-zone-by-period trends, that capture different trends in regional housing markets.

The results are similar across the two specifications: foreign-born residents with higher affinity for homeownership have a significantly higher responsiveness to credit shocks. While the effect of the credit supply shock on homeownership is also positive for U.S.-born locals and the low-HOCO foreign-born, these effects are too noisy to be statistically distinguishable from zero.

We see in comparing columns (5) and (6) to columns (3) and (4) that most of the impact of the lending shocks on homeownership is thus driven by the high-HOCO groups, suggesting that native-born residents and low-affinity residents were not very responsive to this particular source of credit supply variation. The units of the shock correspond to relative log point differences in national loan growth (controlling for local credit demand) among the banks that the average county in the commuting zone is exposed to. That is, the coefficient in row 3 of column (6) indicates that an exogenous 1 pp increase in national loan growth by the local banks leads to a nearly 4 pp increase in homeownership among the high HOCO cohort. In the context of homeownership gaps, per 5-year period, the change in the ΔHO gap between high and low HOCO households was 0.86pp. A 1-standard deviation credit shock of 0.28 implies that the ΔHO gap increases by 1.06 pp per five year period.

E. Housing boom and bust cycles, homeownership affinity, and delinquency

How does the effect of a desire to own a home on homeownership vary across the housing cycle? The results above may raise the question of whether the *HOCO* effects that we find are a phenomenon that is limited to housing booms. Table V shows that high-affinity household took up mortgage credit disproportionately over 2000-2015, when it was *relatively* more available in their county. Over this period, the national housing market also experienced severe boom-bust cycles in the *levels* of credit availability and in house prices that caused substantial dislocation in housing markets. As our framework in Section II suggested, and as we found in the case study on marginal demographics during the 2000-2005 boom in Section V.A, based on their higher affinity for homeownership we would expect high-*HOCO* households to have higher rates of *becoming* homeowners during boom periods. We would also expect them to have lower rates of *exiting* homeownership during downturns, as they are less likely to consider the carrying cost of a mortgage prohibitive relatively to groups that are less attached to homeownership.

In Figure 6 we show this differential responsiveness to the cycle visually in the raw homeownership changes by group. We break up the sample into the same periods that are also analyzed in the regression analysis: boom (2000-2005), bust (2005-2010) and rebound (2010-2015); and we additionally show data for the renewed boom period 2015-2019 that followed. The figure shows

that the desire to own a home systematically affected group level behaviors across different parts of the housing cycle. During the housing boom, a time of expanding credit, we see that homeownership changes the most among high-HOCO groups, increasing by almost 6%. During both the housing bust and initial rebound period, as credit markets recovered from the financial crisis, and homeownership was falling nationwide, it was falling *the least* among the high HOCO group, and then again rose fastest during the 2015-2019 period, albeit at only less than half the rate observed in the prior boom, growing at 2.6%. This is consistent with mortgage credit supply generally being less available after the Great Recession—but, nonetheless, high HOCO groups show a consistently higher tendency for moving into homeownership in all parts of the cycle.

Mortgage delinquency. *Why* do homeowners with a high *HOCO* experience smaller declines in homeownership during the housing bust? While we would expect homeownership to increase during periods of easing credit supply, we do not necessarily expect homeownership to fall when credit dries up. Only those with adjustable or floating rate mortgages will experience the immediate effect of a contraction in credit through rising mortgage rates, but as the overall economy deteriorates, households' financial circumstances may also take a hit. Whether households exit from homeownership as a result, depends on whether they are willing and able to endure the hardship associated with not defaulting on mortgage payments.

In Table VI, we explore how *HOCO* affects whether foreign-born households default on their mortgages over time, using data from the American Housing Survey. Using data from 2013-2017,¹⁶ we regress an indicator for having at least one missed mortgage payment in the past three months on *HOCO*, whether the household bought at the peak of the credit cycle in 2006-2007, and the interaction of the two. We also include CBSA-by-year fixed effects, detailed socioeconomic and demographic controls, controls for years the householder has lived in the US, as well as controls for debt-to-income and house cost-to-income ratios, which capture the degree of financial pressure that the household may be experiencing over the cycle. This analysis captures whether households with comparable finances behave differently in terms of mortgage delinquency depending on how attached they are to remaining a homeowner.

In column (1) of Table VI, we find that a 1 pp higher *HOCO* is associated with 8% fewer missed payments, suggesting that overall, higher affinity households seem to be more committed to paying their mortgage. In column (2), we split the sample by those who purchased at the peak of the cycle and explore heterogeneity by purchase date and *HOCO*. The raw delinquency rate in this sample of households that purchased at the peak of the boom is 11.6%, more than twice that of households that bought in other years, which is 5.6%.

We find that most of the variation in missed payments is driven by those households who purchased at the end of the housing boom; these households bought homes when they were most expensive and therefore often suffered large declines in their equity positions that may have resulted in their mortgages being underwater. We find that, for mortgages not originated at the peak of the boom, *HOCO* is still estimated to lower delinquency, but this estimate is no longer statistically significant. However, having a high HOCO significantly lowers delinquency rates among households

that bought at the peak, undoing a substantial share of the peak buying effect: a one standard deviation higher *HOCO* (which corresponds to a 12 pp difference in this AHS sample) is associated with a 4.4 pp lower delinquency rates for peak boom-originated mortgages,¹⁷ or a decline by 38% of the total delinquency rate for properties bought in those years. These results, in combination with Figure 6, suggest that high-affinity households have a high tendency to move into homeownership in all parts of the housing cycle and also hold on tighter to their homes during hard times than other households during market downturns. The same effects are shown visually in Appendix Figure BVI.

Housing Returns. While holding on more tightly to one's home may stabilize or improve the consumption flows that residents receive from owning their homes, they do not necessarily improve their financial returns. In Figure 7, we plot the annual housing return differentials by period of purchase, and *HOCO* group. We compute predicted returns to holding a home as the annual appreciation in the real value of the home in American Housing Survey data since the time that a household bought the home that they are living in. The coefficients shown in the figure show what share of these returns is predicted by the *HOCO* group of the current owner and the original purchase period. These return differentials control for the same household and demographic characteristics, commuting zone by year fixed effects, and wealth as in Table II. Additionally, to allow for differential returns to houses of varying qualities, we control for standard hedonic characteristics including bedrooms, bathrooms, acreage, square footage, the presence of a garage, the year the unit was built, and unit type.¹⁸

The omitted group is foreign-born households with below-median *HOCOs* who bought in the years before 2000. Historically, relative to low-*HOCO* households, high-*HOCO* households realized 27 basis points (bps) lower returns each year if they purchased prior to 2000. The second to the fifth bar in the figure show the baseline changes in returns for low-*HOCO* households that bought in later time periods. They clearly show the strong effect of differential timing relative to the real estate market cycle: homes bought during the boom years before the height of the recession in 2009 have low returns, whereas homes bought before and during the recovery in house prices during the 2010s experienced higher appreciation.

High-*HOCO* households who bought during the 2000s ended up realizing relatively higher annual returns than their low-*HOCO* peers, on the order of 42-53 basis points, but this advantage became smaller and not statistically significant for those buying during the 2010s. Importantly, these differential gains for high *HOCO* households are swamped by the negative returns realized for the cohort of buyers who purchased in the 2000s, which was on the order of 1.9-4.6 percentage points. Overall, a desire to own a home does not yield meaningfully different housing returns on average. Even though the greater willingness to hold on to one's home does seem to have modest benefits for the high-*HOCO* cohort that bought during the 2000s boom, these benefits do not outweigh buying at the wrong time. This importance of buying in the right market will resurface again in the next section when we analyze the implications for retirement portfolios of an affinity for homeownership.

VI. Retirement Investment Portfolio Choice

One potential issue with our analysis showing that an affinity for homeownership causes foreign migrants to be more likely to own their home is that it may take time for a desire for homeownership to translate into realized portfolio choices. Conversely, a greater desire to own a home might affect the *timing* of when a household acquires a home, but not the long run homeownership rate if we consider the full lifecycle. Moreover, acquiring a home may have consequences for overall household wealth and other financial portfolio choices. Understanding whether an affinity for homeownership may affect retirement portfolios is particularly important because of the large role that real estate plays in retirement portfolios.

Convergence of Homeownership. To see how a higher *HOCO* affects homeownership over an immigrant’s time in the U.S, we plot the estimated effect of being in different categories of years since immigration on homeownership, controlling for other household characteristics, by *HOCO* in Figure 8. The resulting homeownership curves reveal two insights: on the one hand, time since immigration matters substantially for the level of homeownership—for both low and high *HOCO* householders, homeownership varies by a factor of three between foreign-born householders that recently arrived in the U.S. and those who have been in the country for more than 40 years. Moreover, a higher *HOCO* affects the initial homeownership rates in the first years after arrival, but it has an even greater effect on differences in homeownership in the long run: after several decades in the U.S., high *HOCO* foreign-born households on average exceed the homeownership rates of U.S.-born households, while low *HOCO* households on average have homeownership rates that are 15 pp lower than those of U.S.-born households even in the long run. The graphs also show that the homeownership rate curves tend to flatten after several decades in the country. This analysis motivates the use of retirement portfolio data to get an estimate of the *long-run* financial consequences of wanting to own a home.

HRS data. To estimate the size of *HOCO* effects on household retirement portfolios, we use confidential microdata from the Health and Retirement Study (HRS) that links detailed information about household balance sheets for individuals retiring in the U.S. with their country of origin if they are foreign-born. While the sample size of this data is substantially smaller than the ACS sample, this data is uniquely suited to answering the questions of interest, as by the time that a household reaches retirement age, we can expect their preferences to have been mostly translated into portfolio choices, and we are not aware of any other data set that combines detailed household balance sheet information with country of origin information. By comparing holdings of real estate and other assets upon retirement, we can study the effects of *HOCO* on *long-run* investment choices, and the long-run consequences for the overall wealth and portfolio allocation choices of immigrants. The HRS data and construction of individual variables is described in detail in Appendix Section B.

Estimation specification. Using the HRS data, we estimate regressions of the form

$$Y_h^{age60} = \gamma HOCO_h + \beta' X_h^{age60} + \alpha_t + \varepsilon_h, \quad (5)$$

where Y_h^{age60} represents a household level financial characteristic measured in the survey wave where the financial respondent was closest to 60 years old; $HOCO_h$ represents the foreign-born financial respondent’s country of origin homeownership rate matched to their marriage status; X_h^{age60} are demographic controls that capture non-HOCO differences between respondents at age 60 and consist of categorical variable dummies for respondent marriage status, age, total household income, a dummy for gender, a dummy for whether they are at least a high school graduate, and a dummy for whether the respondent was born in Mexico; α_t is a dummy indicating the survey wave that the respondent data at age 60 is from. The baseline specification also controls for total non-housing wealth of the household. The country of origin is determined from the location where a respondent was born and we restrict the analysis to respondents who did not live in the U.S. when they were in school, in order to focus on individuals who grew up abroad.

A. Homeownership in the retirement portfolio

Many households seem to show a special preference for real estate when allocating their portfolios at retirement age: Poterba, Venti, and Wise (2011) document that over 80% of households own a home at retirement, and housing represents about half of non-annuitized wealth. We can see this by considering a household’s “discretionary” risky asset portfolio, i.e. funds that are not allocated through a retirement account, but rather can be moved freely by the household, and which are invested in either stocks, bonds, or real estate. Figure 9 shows that more than half of all retirees have *only* risky asset investments in the form of their primary residence—and do not hold any stocks or bonds outside of their retirement accounts. Thus, homeownership seems to occupy a special place in the set of asset classes that households could invest in—and cultural affinities for owning a home might be an important driver of this preference.

In order to quantify how much an affinity for homeownership affects realized tenure choices *at the end of the working life* of a household, we first estimate the effect of $HOCO$ on whether or not a retiring household owns their primary residence. Column (1) of Table VII shows this effect estimate: the estimated elasticity of realized retirement-age homeownership with regard to $HOCO$ is 0.77 and statistically significant in spite of the relatively small sample size of the HRS. This is larger than our preferred effect estimates from the ACS microdata, which might be due to retirement age immigrant households having had more time to establish themselves in the U.S. and realize their preferences for owning a home. This result confirms that an affinity for homeownership results in actual differences in tenure choices, even in the long run. Column (2) shows that this result is robust in magnitude, albeit weaker in statistical significance, when we exclude Mexican-born migrants who make up the majority of the sample. Moreover, note that one benefit of the HRS is that we are able to control not just for income but also for non-housing wealth of the households,

which supports our claim that this passthrough of country-of-origin characteristics is independent of wealth endowments of the migrants. Columns (3) and (4) confirm that this tenure choice ends up biasing the asset allocation in retirement portfolios: a 1 pp increase in a respondent’s *HOCO* is associated with a 0.9 pp higher share of their wealth portfolio consisting of the real estate value of their home.

To visualize this effect of *HOCO* on retirement age tenure choices, and check for non-linearities, consider Figure 10: If we split the respondents with *HOCO* data into terciles of affinity, homeownership at age 60 increases monotonically with being in a tercile with greater affinity for it.¹⁹

B. Retirement wealth and portfolio composition effects

Many policies supporting or subsidizing homeownership in the U.S. are motivated by the assumption that homeownership leads to greater wealth. Theoretically, this need not be the case. Housing is a risky and undiversified asset. As such, many models imply that households should compensate for this by concurrently investing in safe assets, and lowering their exposure to other risky assets, such as stocks (Cocco, 2005; Yao and Zhang, 2005). Additionally, because households must save for a downpayment, this can preclude investing in alternative assets early on in the life cycle, losing out on many years of compounding returns. On the other hand, the wealth effects from unexpected gains to housing equity can loosen budget constraints enabling investment in other risky assets (Chetty et al., 2017; Sodini et al., 2023). Even in the absence of a diversified portfolio, housing may contribute to wealth formation through forced savings, (Bernstein and Koudijs, 2024; Kovacs and Moran, 2023). The ultimate contribution of housing to one’s wealth at retirement is thus *ex ante* unclear, and an open empirical question.

How does this higher homeownership upon retirement change the retiree’s financial portfolio? To answer this question, we first consider whether high *HOCO* households end up with higher household wealth overall. In Table VIII, we find that households with higher *HOCO* indeed have significantly higher absolute wealth not just in the form of the net value of their primary residence (column (1)) but also in the rest of their retirement portfolio (column (2)). A 1 pp higher *HOCO* is associated with about a \$3.5K more valuable primary residence, and \$4.6K higher wealth in other assets. This effect is large, compared to a baseline average primary residence value of \$110K and wealth in other assets of \$129K, representing increase. It implies that a standard deviation increase in *HOCO* (12 pp) is associated with increases relative to these averages of 38% and 43%, respectively. Moreover, this effect does not just stem from the extensive margin effects of being more likely to choose homeownership – in columns (5) and (6) we show that differential affinity for homeownership also affects retirement portfolio wealth conditional on being a homeowner.

One potential explanation for this association between *HOCO* and retirement wealth could be differential preferences for risky assets that have higher returns – either associated with country of origin preferences, or as a result of real estate ownership affecting risk preferences – but in columns (3)-(4) and (7)-(8) we find no significant evidence that higher *HOCO* is associated with being more likely to own stocks nor that a higher share of the retirement portfolio (excl. real estate) is

being allocated to equities. Here, we condition on total non-housing wealth that might otherwise distort overall asset allocation choices. This stands in contrast to the theoretical findings in Cocco (2005) and Yao and Zhang (2005), who suggest housing investment substitutes for risky assets. Instead, these portfolio results are more consistent with Chetty et al. (2017), which notes that this substitution effect works in opposition to the potential wealth effect for households with growing equity in their homes.

Recent research by Sodini et al. (2023) on homeowners in Sweden seems to find the opposite result: being exogenously allocated homeownership after municipalities privatize existing units seems to increase not only housing wealth, but also stock market participation. What could explain our divergent findings? While our point estimates are positive, our smaller sample size may lack power for detecting small effects on stock market participation. However, different effects for U.S. residents may also be explained by differences in direct stock market participation across countries: according to the Survey of Health, Aging and Retirement in Europe (SHARE), 44% of married couples in Sweden directly participate in the stock market. On the other hand, according to the Survey of Consumer Finances, only 26% of couples in the U.S. do so. Owning a home in the U.S. context does not seem to significantly impact households' stock market participation, perhaps reflecting that for most households in the U.S. the effective barriers to participation outside of a retirement or pension plan seem to be quite high.

Our findings are instead consistent with concurrent work in the U.S. context studying the impact of losing one's home during the Great Recession, the flip side of the unexpected homeownership studied in Sodini et al. (2023). Artigue, Bayer, Ferreira, and Ross (2024) find that households that marginally kept their homes vs. marginally lost their homes between 2010 and 2013 have 19% higher homeownership by 2022. This persistent difference in homeownership rates does not seem to translate into differences in other long-term outcomes: credit scores, credit balances and auto loans across the two groups are statistically indistinguishable. These findings in combination with our own suggest that homeownership in the U.S. impacts few financial outcomes outside of those closely related to homeownership, namely housing wealth and overall wealth.

C. Regional house price dynamics and the HOCO effect on the retirement portfolio

Why does a higher *HOCO* also affect *overall* wealth? For many policymakers, the notion that owning a home is a key mechanism for building wealth is assumed to be obvious. After all, those who bought a home in earlier decades are now wealthy! For example, President Biden stated on May 31, 2024, that “owning a home is ... a pathway ... to pass down wealth to your children.”²⁰ Economists have questioned this assumption, highlighting that the returns to homeownership have been influenced by the good fortune of buying before a rise in house prices—which may be hard to replicate—and may look less favorable when considering the opportunity cost of not investing in other assets that have also seen high returns (Rappaport, 2010). As households endogenously select into homeownership, simple comparisons between homeowners and renters, while common in the policy discourse, are not convincing evidence that homeownership *causes* greater wealth. To

the degree that our *HOCO* measure is an exogenous driver of homeownership, conditional on other household characteristics, our results above provide a novel contribution by showing that a greater tendency to own a home caused higher wealth at retirement age.

However, it is important to understand the mechanism for this wealth effect of homeownership. While policy may make it easier or harder to buy a home, for example by lowering downpayments or subsidizing interest rate risk, these levers have limits. Housing policy as mediated through mortgages will only impact wealth creation if housing acts as a savings commitment device as in Kovacs and Moran (2023) or households adjust their consumption and leisure instead of substituting other forms of saving (Bernstein and Koudijs, 2024). If housing wealth accrual occurs instead through appreciation as in Sodini et al. (2023), homeownership policies that vary mortgage design will have less bite in terms of impacting household wealth, if—as is likely—they cannot also ensure favorable housing market dynamics after the purchase.

In section V, we showed that the desire to own a home increases households’ sensitivity to credit shocks, and also increases their tendency to acquire a home in all parts of the housing cycle, while reducing the tendency to default on a mortgage for homes bought at the peak of the boom. The effect of this behavior on the returns to buying a home is ambiguous: if it leads to buying a house at times when returns to buying a house are high, the household might experience a windfall from housing equity gains, but it can also affect household finances negatively if it results in buying a home even when prices are high and expected returns are low.

We can test whether the higher housing and overall wealth associated with a higher *HOCO* are due to higher saving, independent of the economic environment, or luck in having a fondness for an asset class that happens to have seen high appreciation rates in recent decades. While we do not have information on the exact timing and location of the purchase of the real estate that households own at the time of retirement, we can use the fact that most older households have low geographic mobility: we proxy for the likely housing capital gains that a household experienced with the past house price growth in the state where the household resides at age 60.

In Table IX, we explore whether the effect of *HOCO* on the value of retirement portfolios varies with price growth in households’ regional housing markets over the previous 20 years. We split respondents into terciles based on their state’s past house price growth. As the shown in Table IX column 1, there is strong evidence that the effect of *HOCO* on the value of the primary residence is only significant in states where housing markets experienced high price growth (top tercile). Additionally, column 2 shows that those are the only areas that saw a significant impact of higher *HOCO* on the value of households’ overall real estate portfolios (which include primary as well as secondary residences). The estimates suggest that high *HOCO* households have similar housing wealth upon retirement as low *HOCO* households if they live in states with unfavorable housing markets. In the highest house price growth states, a one standard-deviation increase in *HOCO* (roughly 12 pp) yields a substantial $\sim 82\text{K}$ increase in net primary residence value, or $\sim 199\text{K}$ in total net real estate value, or an 81% and 118% increase relative to the sample average. In contrast, column 3 shows that total wealth seems to be higher for high *HOCO* households in high

house price growth areas, but the coefficient is not statistically significant, and smaller than the real estate effects. In short, at least in the U.S. context, housing wealth owes more to price appreciation as in Sodini et al. (2023), and less to differential savings mechanisms. In column 3, we also show that there is no longer statistically significant evidence that *HOCO* affects other components of retirement wealth, once we allow HOCO to interact with regional housing market dynamics. In column 4, we show the differential effect of HOCO on stock ownership in locations with and without housing windfalls. We find weak evidence of higher affinity leading to higher stock ownership, and only in high house price growth states, consistent with Chetty et al. (2017).

D. Policy implications

Overall, the retirement portfolio effect findings are consistent with the wealth effects of owning a home and holding on to it impacting household portfolios mostly through unexpected house price gains. This is also consistent with our earlier finding that differential returns to *when* a household is buying a home over the housing cycle tend to be much larger than the difference in returns between buyers with different affinities who buy during the same period.

Note that this does *not* mean that the affinity for homeownership does not matter for policy. To the contrary, as our results show, high affinity households end up with higher housing wealth at retirement—an important policy goal—and are more likely to respond to credit expansions, which are a key tool of housing market policy. However, our results regarding the retirement portfolio and housing return effects of this higher homeownership show that it does not inevitably result in the desirable outcome of higher wealth for the affected households: the financial returns to having a higher affinity for homeownership that we find result mostly from getting “lucky” in either the location or timing of the purchase, not from a general effect on savings or asset accumulation that would benefit a homeownership household regardless of such windfalls.

This takeaway has important consequences for policy design: unless policymakers can guarantee that incentives to nudge households into becoming homeowners are timed favorably relative to the market cycle over time, or with regard to regional housing market dynamics, homeownership by itself is unlikely to result in substantial benefits in terms of wealth accumulation for the targeted households. Our findings show that, as it currently exists in the U.S., homeownership by itself may be a less effective policy lever than simply targeting savings directly through policy incentives or encouraging larger savings among homeowners (e.g. through stricter amortization policies as in Bernstein and Koudijs (2024)).

In fact, political pressures to “make homeownership affordable” might be stronger during high price periods at the peak of a boom and could thus push marginal households into homeownership at exactly the moments when this is least likely to benefit the affected households, as we showed in Section V.E. Broad incentives to enter homeownership, for instance through credit policy, have also been found to have negative unintended consequences for the targeted populations and the economy at large (Keys et al., 2010; Adelino et al., 2016; Justiniano et al., 2019; Gupta, 2019).

Finally, we note that while housing makes up a large share of households’ wealth after retirement,

it tends not to be used until very late in life. This could be to insure against the costs of very old age (Poterba et al., 2011), because of bequest motives (Bernheim, Shleifer, and Summers, 1985), or difficulty tapping into housing wealth (Mayer and Moulton, 2022). If all of a household’s wealth is tied up in housing, it may be harder to access than wealth in more liquid forms, such as wealth held in annuities or mutual funds. The measures of exogenous variation in homeownership driven by affinities that we construct may allow policy analysts to better quantify these costs and benefits of homeownership in retirement.

VII. Conclusion

Homeownership has long been a core tenet of the American Dream, often touted as the best way to build wealth over the life cycle. By utilizing cultural norms in homeownership to vary tenure choice in the U.S., we provide a new understanding on the impact of tenure choice on household wealth formation. We find that higher affinity households select into homeownership at higher rates, are especially sensitive to credit supply shocks, and thus exposed to increased house price volatility. At retirement, these households do end up realizing higher housing wealth and overall wealth, but this does not seem to substitute for other risky asset shares, nor does the wealth effect of unexpected housing returns yield differential stock market participation.

Our work shows that using origin-country preferences as a laboratory to study the causal effects of the desire to own a home provides new insights about why some households are more exposed to credit shocks, what drives retirement portfolio allocations, and that there is an affinity-related component driving homeownership and housing cycles. Moreover, our results have direct implications for understanding immigrants’ wealth positions in the U.S.

Given that one’s ancestry often intersects with more visible identifiers, such as race or ethnicity, policymakers should consider country-of-origin differences and other drivers of homeownership *affinity* when designing policies. In fact, we find that differences in preferences can amplify or dampen the impact of policies which hope to encourage homeownership. We highlight a particularly salient dimension of housing policy and the transmission of economic shocks to the housing market by studying the impact of changes in credit supply. Our results suggest that where someone was born before they moved to the U.S. can have long-lasting effects on how they interact with financial markets, as credit supply variations are more likely to lead to an investment in homeownership for groups that come from countries where homeownership is more common. However, there are many other dimensions of housing policy and housing market shocks that may show similar differences across foreign-born groups, but which we do not cover in this paper - and we hope that future research can fill that gap. Moreover, there is a large literature on the effects of homeownership on different social, political, and economic outcomes, and the importance of homeownership for financial crisis impacts. We hope that using exogenous variation in homeownership at the household level that is driven by country-of-origin characteristics can enable other researchers to better identify the causal effects of homeownership across these different domains.

REFERENCES

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2016, Loan originations and defaults in the mortgage crisis: The role of the middle class, *The Review of Financial Studies* 29, 1635–1670.
- Ambrose, Brent W, James N Conklin, and Luis A Lopez, 2021, Does borrower and broker race affect the cost of mortgage credit?, *The Review of Financial Studies* 34, 790–826.
- Angrist, Joshua D, and Jörn-Steffen Pischke, 2008, *Mostly harmless econometrics* (Princeton university press).
- Artigue, Heidi, Patrick Bayer, Fernando Ferreira, and Stephen Ross, 2024, The long-term consequences of losing a housing during the great recession, *Working Paper* .
- Bayer, Patrick, Fernando Ferreira, and Stephen L Ross, 2016, The vulnerability of minority homeowners in the housing boom and bust, *American Economic Journal: Economic Policy* 8, 1–27.
- Benetton, Matteo, Marianna Kudlyak, and John Mondragon, 2024, Dynastic home equity?, *Working Paper* .
- Bernanke, Ben S., 2007, The importance of financial education and the national jumptartcoalitionsurvey.
- Bernheim, B Douglas, Andrei Shleifer, and Lawrence H Summers, 1985, The strategic bequest motive, *Journal of Political economy* 93, 1045–1076.
- Bernstein, Asaf, and Peter Koudijs, 2024, The mortgage piggy bank: Building wealth through amortization, *The Quarterly Journal of Economics* 1767–1825.
- Bhutta, Neil, 2015, The ins and outs of mortgage debt during the housing boom and bust, *Journal of Monetary Economics* 76, 284–298.
- Blau, Francine D, Lawrence M Kahn, and Kerry L Papps, 2011, Gender, source country characteristics, and labor market assimilation among immigrants, *The Review of Economics and Statistics* 93, 43–58.

- Botsch, Matthew J., and Ulrike Malmendier, 2020, The long shadows of the great inflation: Evidence from residential mortgages, *Working Paper* .
- Brueckner, Jan K, 1986, The downpayment constraint and housing tenure choice: A simplified exposition, *Regional Science and Urban Economics* 16, 519–525.
- Center, Minnesota Population, 2020, Integrated public use microdata series, international: Version 7.3 .
- Charles, Kerwin Kofi, and Erik Hurst, 2003, The correlation of wealth across generations, *Journal of political Economy* 111, 1155–1182.
- Chetty, Raj, Lázló Sándor, and Adam Szeidl, 2017, The effect of housing on portfolio choice, *The Journal of Finance* 72, 1171–1211.
- Cocco, Joao F, 2005, Portfolio choice in the presence of housing, *The Review of Financial Studies* 18, 535–567.
- Das, Prashant, N Edward Coulson, and Alan Ziobrowski, 2019, Caste, faith, gender: Determinants of homeownership in urban india, *The Journal of Real Estate Finance and Economics* 59, 27–55.
- Dietz, Robert D, and Donald R Haurin, 2003, The social and private micro-level consequences of homeownership, *Journal of urban Economics* 54, 401–450.
- Dustmann, Christian, 1996, The social assimilation of immigrants, *Journal of population economics* 9, 37–54.
- Fonseca, Julia, and Lu Liu, 2024a, Mortgage lock-in, mobility, and labor reallocation, *The Journal of Finance* 79, 3729–3772.
- Fonseca, Julia, and Lu Liu, 2024b, Mortgage lock-in, mobility, and labor reallocation, *Real Estate Economics* 79, 3729–3772.
- Garcia, Daniel, 2020, Employment in the great recession: How important were household credit supply shocks?, *Real Estate Economics* .
- Gilchrist, Simon, Michael Siemer, and Egon Zakrajsek, 2018, The real effects of credit booms and busts, Technical report.

- Goodman, Laurie S, and Christopher Mayer, 2018, Homeownership and the american dream, *Journal of Economic Perspectives* 32, 31–58.
- Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen, 2020, Do credit market shocks affect the real economy? quasi-experimental evidence from the great recession and” normal” economic times, *American Economic Journal: Economic Policy* 12, 200–225.
- Gupta, Arpit, 2019, Foreclosure contagion and the neighborhood spillover effects of mortgage defaults, *The Journal of Finance* 74, 2249–2301.
- Happel, Jonas, Yigitcan Karabulut, Larissa Schäfer, and Selale Tuzel, 2022, Shattered housing, *Available at SSRN 4029064* .
- Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, 2019, Credit supply and the housing boom, *Journal of Political Economy* 127, 1317–1350.
- Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010, Did securitization lead to lax screening? evidence from subprime loans, *The Quarterly Journal of Economics* 125, 307–362.
- Khwaja, Asim Ijaz, and Atif Mian, 2008, Tracing the impact of bank liquidity shocks: Evidence from an emerging market, *American Economic Review* 98, 1413–42.
- Kovacs, Agnes, and Patrick Moran, 2023, Financial innovation, the decline in household savings, and the trade-off between flexibility and commitment, *Working Paper* .
- Layton, Don, 2021, The homeownership rate and housing finance policy– part 1: Learning from the rate’s history, Technical report, Joint Center For Housing Studies, Harvard University.
- Layton, Don, 2022, Homeownership rate and housing finance policy– part 2: Heavy lift to break the 65 percent barrier, Technical report, Joint Center For Housing Studies, Harvard University.
- Malmendier, Ulrike, and A. Steiny Wellsjo, 2024, Rent or buy? Inflation experiences and homeownership within and across countries, *Journal of Finance* 79, 1977–2023.
- Marcén, Miriam, and Marina Morales, 2020, The effect of culture on home-ownership, *Journal of Regional Science* 60, 56–87.

- Mayer, Christopher, and Stephanie Moulton, 2022, *The market for reverse mortgages among older Americans* (Oxford University Press).
- Mian, Atif, and Amir Sufi, 2009, The consequences of mortgage credit expansion: Evidence from the us mortgage default crisis, *The Quarterly journal of economics* 124, 1449–1496.
- Mian, Atif, and Amir Sufi, 2022, Credit supply and housing speculation, *The Review of Financial Studies* 35, 680–719.
- Pfeffer, Fabian, Bob Schoeni, Arthur Kennickell, and Patricia Andreski, 2016, Measuring wealth and wealth inequality: Comparing two u.s. surveys, *Journal of Economic and Social Measurement* 41, 103–120.
- Poterba, James, Steven Venti, and David Wise, 2011, The composition and drawdown of wealth in retirement, *Journal of Economic Perspectives* 25, 95–118.
- Rappaport, Jordan, 2010, The effectiveness of homeownership in building household wealth, *Economic Review-Federal Reserve Bank of Kansas City* 35.
- Ringo, Daniel R., 2020, Home ownership as a labor market friction, *Real Estate Economics* 49, 111–133.
- Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, 2020, Ipums usa: Version 10.0 .
- Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf. von Lilienfeld-Toal, 2023, Identifying the benefits from homeownership: A swedish experiment, *American Economic Review* 113, 3173–3212.
- University of Michigan, 2024, Health and retirement study.
- Yao, Rui, and Harold H. Zhang, 2005, Optimal consumption and portfolio choices with risky housing and borrowing constraints, *The Review of Financial Studies* 18, 197–29.

Notes

¹While the PSID does include an immigrant survey sample, the survey provides highly aggregated country-of-origin groups that significantly limit the variation in cultural norms relative to the confidential HRS data, which provides specific origin country information.

²For simplicity, we will assume that the household can borrow and invest at the same rate, but the conclusions do not qualitatively depend on this.

³Note that the homeowners here implicitly receive returns on their entire savings - even though αP is invested in the downpayment - because any downpayment reduces the mortgage cost at a rate r .

⁴We distinguish between 4 marital status categories: single or never-married, married, separated or divorced, and widowed.

⁵We dropped any sample with an average homeownership rate of less than 33% that could not be independently verified by an online search, which led us to drop: Argentina 1991, Benin 2002, Ethiopia 1994, Togo 2010, and Papua New Guinea 1990.

⁶According to a report from the New York City Comptroller’s office, “Spotlight: New York City’s Homeowner Housing Market,” March 2024. Accessed via <https://comptroller.nyc.gov/wp-content/uploads/documents/Spotlight-New-York-Citys-Homeowner-Housing-Market.pdf>

⁷This hand-collected comparison sample covers origin countries comprising 87% of all the foreign-born in the U.S. and 9.2% of the total U.S. population.

⁸This data was accessed through IPUMS (Ruggles, Flood, Goeken, Grover, Meyer, Pacas, and Sobek, 2020).

⁹We downloaded this crosswalk from Neil Bhutta’s website. The mapping to lender parent companies is only available through 2017 at which point the lender identification codes in HMDA change. Therefore, we limit the analysis to the years before for which we can consistently aggregate to parent companies over time.

¹⁰One may wonder why we did not use variation across counties within the US as our laboratory. First, our preferred dataset does not have data more granular than state for place of birth among US resident. Second, Figure 1 highlights that comparing across countries yields wider variation in homeownership when comparing to states. Third, the only dataset we are aware of that maps more granular place-of-birth, at the county level, to adult locations is the National Longitudinal Survey of Youth 1979 (NLS79) or 1997 (NLS97). Neither of these samples has much variation in expected homeownership over the credit boom period we study in the paper (2000-2005), due to the participants having already largely transitioned to homeownership (NLS79), or being too young to have started (NLS97), as highlighted in Figure BIV. In sum, using the immigrant sample ACS as our core dataset yields high variation in HOCO, broad coverage of lifecycle demand for housing, and large sample sizes with detailed geographic variation within the US for current locations.

¹¹There is an “ancestry” variable, but it does not allow us to distinguish between the children of recent immigrants and those of families with a long history in the country.

¹²See Appendix Table AII for the full results for the children of immigrants with different combinations of control variables corresponding to the columns in Table II.

¹³The prediction uses the following household characteristics: county fixed effects; Real HH Income (linear and ventile indicators); Quadratic function of Age of HH head and 5-year age group indicators, indicators for educational achievement of HH head, Indicators for # of children and # of relatives in the HH, Indicators for discretized years since immigration of head, marital status of HH head, and the interaction between marital status and years-since-immigration categories.

¹⁴See Appendix Table AIII for the coefficient estimates corresponding to the linear models shown in the graphs in Figure 5.

¹⁵While we do estimate separate coefficients for the Mexican-born and no-HOCO data foreign-born groups, we do not present them as we do not have specific hypotheses for their outcomes, which has the additional benefit of making the table more concise.

¹⁶The AHS only initiated questions on missed rent or mortgage payments in 2013, and asks them every 4 years, such that researchers can only observe these in the 2013, 2017, and 2021 waves.

¹⁷Computed as $0.12 * (-0.366) = -4.4\%$.

¹⁸The type of unit includes whether a home is in one of the following structure types: a mobile/manufactured home, single family detached, singly family attached, 2 unit, 3-4 units, 5-9 units, 10-19 units, 20-49 units, 50+ units.

¹⁹This figure corresponds to the same sample and control variables as used in column (2) of Table VII. For the same analysis but without including any control variables, i.e. a raw comparison of group averages, see Appendix Figure BVII, which shows very similar patterns.

²⁰From “A Proclamation on National Homeownership Month, 2024”. Source: <https://www.whitehouse.gov/briefing-room/presidential-actions/2024/05/31/a-proclamation-on-national-homeownership-month-2024/>

Tables

Table I: Homeownership rates for selected countries of origin computed from international census data

Rank	Origin	%Pop ⁰⁰ _{USA}	% Married	% <i>HOCO</i> ^{All}	% <i>HOCO</i> ^{Married}
<i>Census 2000 Birthplace</i>					
1	Mexico	3.24	67	66	69
2	Philippines	0.51	64	77	77
3	India	0.47	75	79	79
4	China	0.38	71	74	86
5	Canada	0.38	55	68	81
⋮	⋮	⋮	⋮	⋮	⋮
	United States	86.5	54	67	83

Notes: This table summarizes our homeownership rates, by country of origin (*HOCO*) for the largest country-of-origin groups among foreign-born immigrants residing in the U.S. These homeownership rates are computed from harmonized country census microdata (obtained through IPUMS) by marriage status for household heads. Only countries for which this data was available are shown. The table is ranked by the population share of each immigrant group among U.S. household heads who are 18-69 years old in 2000 (according to Decennial Census data), and also shows the share of household heads from that origin who are married. It shows the average homeownership rate in the country of origin, as well as the rate among married household heads. We also provide the US native-born population data for the year 2000. Data for homeownership by marriage status in China is from the 2016 China Family Panel Studies Survey. Due to a lack of population weights by marriage status in the origin country, we approximate China “All” population marriage rates by taking a weighted average of married and single marriage rates, using prevalence of these statuses among Chinese-born US residents. Indian data is obtained from the 2015-2016 wave of the Demographic and Health Survey. Full table of 70 countries with homeownership data available upon request. Population shares of the 5 largest country-of-origin groups for which we could not compute homeownership data by marriage status: Germany, 0.47%; Cuba, 0.36%; South Korea, 0.29%, Japan, 0.21%, Taiwan, 0.15%.

Table II: Baseline effects of Homeownership in Country of Origin on Tenure Choice

<i>Dependent var.:</i>	Household Homeownership _{it} (in %)			
	(1)	(2)	(3)	(4)
HOCO (%)	0.897*** (10.33)	0.208*** (3.21)	0.168*** (3.05)	0.153*** (3.15)
Observations	3,448,546	3,448,436	3,436,300	3,436,252
R-Squared	0.06	0.27	0.27	0.31
	<i>Addl. controls & fixed effects</i>			
Household Characteristics		X	X	X
Imputed HH Wealth			X	X
Commuting Zone × Year FE				X

Notes: This table shows the results of estimating Equation 1 in a pooled household level sample of households with foreign-born heads who are 18-69 years old, for the years 2000 and 2005-2019. The dependent variable is an indicator of homeownership in percent (so 100 indicates that the household owns their home). The independent variable is the homeownership rate in the country of origin (*HOCO*) of the household head for people with the same marriage status (single, married, separated/divorced, widowed). Household characteristic control variables in column 2 consist of: HH Income, (linear and ventile indicators), Quadratic function of Age of HH head and 5-year age group indicators, indicators for educational achievement of HH head, Indicators for # of children and # of relatives in the HH, Indicators for discretized years since immigration of head, marital status of HH head, and the interaction between marital status and years-since-immigration categories. Column 3 adds predicted non-housing household net worth as an additional control variable, and column 4 includes CZ-by-year fixed effects. T-statistics based on heteroskedasticity-robust standard errors clustered at the origin country level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table III: Effects of Homeownership in Country of Origin on Tenure Choice: Robustness Checks and Addl. Specifications

<i>Dependent var.:</i>	Household Homeownership _{it} (in %)				
	(1) Orig. Ctry Controls	(2) Mig. Char. Controls	(3) Excl. 5 largest Orig. Ctries	(4) Gov. source HOCO Data	(5) AHS Data
HOCO: by marriage status	0.167*** (5.35)	0.144*** (3.05)	0.151*** (2.98)		0.134** (2.53)
HOCO: gov. sources				0.148* (1.94)	
Observations	2,060,006	3,436,252	1,627,032	3,485,350	16,320
R-Squared	0.30	0.32	0.33	0.31	0.31
	<i>Addl. controls & fixed effects</i>				
Household Characteristics	X	X	X	X	X
Imputed HH Wealth	X	X	X	X	X
Commuting Zone × Year FE	X	X	X	X	
CBSA × Year FE					X
Origin Ctry Characteristics	X				

Notes: This table shows the estimation results from additional specifications that vary particular aspects of the baseline regression in column 4 of Table II. Column 1 includes origin country characteristics and indicators for whether the householder belongs to different major religions (Buddhism, Christianity, Islam, Hinduism, Other). Column 2 controls for average migrant characteristics by origin country. Column 3 excludes the five largest immigrant groups in the U.S. as of the year 2000 (Mexico, India, China, Vietnam, Philippines) from the sample. Column 4 uses hand collected data from governmental reports for *HOCO*, rather than estimating it from IPUMS microdata. Column 5 uses the AHS sample instead of the baseline ACS sample to estimate HOCO effects. T-statistics based on heteroskedasticity-robust standard errors clustered at the origin country level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table IV: Mechanisms for the Effect of Homeownership in Country of Origin on Tenure Choice

<i>Dependent var.:</i>	Household Homeownership _{it} (in %)					
	(1) Immigrated as Child	(2) By Gender	(3) By Educ.	(4) By Years since Immig.	(5) US-born Spouse	(6) Same Origin Spouse
HOCO	0.122*** (2.96)	0.171*** (3.57)	0.178*** (3.77)	0.004 (0.05)	0.240** (2.31)	0.162** (2.25)
HOCO × Male		-0.037*** (-4.59)				
HOCO × College Educ.			-0.075** (-2.29)			
HOCO × LT Resident (> 10 Yrs)				0.196** (2.46)		
HOCO × U.S.-born Spouse					-0.229*** (-3.26)	
HOCO × Same Origin Spouse						-0.058 (-0.86)
Spouse HOCO						0.157*** (3.12)
Observations	1,220,817	3,436,252	3,436,252	3,436,252	2,050,435	1,595,887
R-Squared	0.31	0.31	0.31	0.31	0.29	0.30
	<i>Addl. controls & fixed effects</i>					
Household Characteristics	X	X	X	X	X	X
Imputed HH Wealth	X	X	X	X	X	X
Commuting Zone × Year FE	X	X	X	X	X	X
Same Origin Spouse Indicator						X

Notes: This table shows the estimation results from additional specifications that vary particular aspects of the baseline regression in column 4 of Table II. Column 1 shows the results in a sample that is restricted to householders *who immigrated to the US when they were 18 years old or younger* as inferred from their age and immigration year. Columns 2-4 check for differential impacts by characteristics of the household head: whether the head is male, is college-educated, and whether the head has been in the U.S. longer than 10 years. Columns 5-6 estimate differential impacts by characteristics of the householder's spouse: Column 5 includes all married householders and interacts the HOCO variable with whether the spouse is U.S.-born. Column 6 includes households where the data also have information on the spouse's HOCO, and shows the independent effect of the spouse's affinity for homeownership, as well as an interaction for whether both spouses are from the same country.

T-statistics based on heteroskedasticity-robust standard errors clustered at the origin country level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table V: Credit Supply Shocks and Homeownership by HOCO

<i>Dependent var.:</i>	$\Delta \text{ Log Loans}_{c,g,t}$ (in pp)		$\Delta \text{ Homeownership}_{c,g,t}$ (in pp)			
	(1)	(2)	(3)	(4)	(5)	(6)
Nonlocal Lending Shock $_{c,t}$	0.451*** (12.23)	0.404*** (9.98)	1.691*** (6.60)	0.918** (2.42)	1.481 (0.74)	0.468 (0.24)
Native \times Nonlocal Lending Shock $_{c,t}$					0.193 (0.09)	0.458 (0.23)
High HOCO \times Nonlocal Lending Shock $_{c,t}$					3.914** (2.10)	3.860** (2.06)
Observations	7,908	7,617	34,044	34,043	34,044	34,043
R-Squared	0.88	0.96	0.24	0.31	0.24	0.31
	Fixed Effects & Controls					
Period FE	X	X	X	X	X	X
County FE	X	X				
County \times Group FE			X	X	X	X
CZ \times Period FE		X		X		X

Notes: Columns 1-2 show county-by-period panel estimation results and columns 3-6 show estimation results using a sample of county-period-group cells, where the periods are 2000-2005, 2005-2010, and 2010-2015. Groups include: US natives, foreign-born w/out HOCO, above-median HOCO, below-median HOCO (reference group), and Mexican-born. Coefficients for Mexican and no-HOCO foreign born groups omitted for legibility. The dependent variable is the county (columns 1-2) or group-level (columns 3-6) change in lending or homeownership rates in percentage points. *Nonlocal lending shock* is a credit supply shock for the local number of mortgage loans originated. T-statistics based on heteroskedasticity-robust standard errors clustered at the commuting zone level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table VI: Delinquencies and HOCO, by Purchase Period

	(1)	(2)
	≥ 1 Missed Pmt.	≥ 1 Missed Pmt.
HOCO	-0.080** (-2.02)	-0.049 (-1.27)
Bought 2006-2007		0.317*** (2.70)
Bought 2006-2007 X HOCO		-0.366** (-2.43)
R^2	0.032	0.039
Observations	4043	4043
	FEs & Controls	
CBSA X Year	X	X
Demographics	X	X
Years in US	X	X
CTI & DTI	X	X

Notes: Data from the 2013 and 2017 waves of the AHS. Sample covers households with foreign-born heads. Left-hand side variable for columns (1) and (2) is whether household has missed at least 1 mortgage payment in the past three months. Missed payments are rare events among these populations, with a missed payment rate of 6.13%. Controls include household head education, race, ethnicity, marital status, number of family members and children in the home, and quadratics in household income, years in the US of household head, and age of household head, as in prior specifications and add debt-to-income and cost-to-income ratios. We include year-by-CBSA fixed effects. T-statistics based on heteroskedasticity-robust standard errors clustered at the commuting zone level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table VII: Effect of HOCO on Real Estate Portfolio Holdings Among Foreign-Born at Age 60

Dep. Var.:	Homeownership		Home Value % of Wealth	
	(1)	(2)	(3)	(4)
HOCO	0.774** (2.036)	0.747* (1.909)	0.931** (2.694)	0.913** (2.692)
Mean of dep. var.	65.2	64.1	60.6	46.1
R-squared	0.11	0.17	0.17	0.19
Observations	466	223	372	181
Demographic controls	X	X	X	X
Non-housing wealth control	X	X	X	X
School outside US	X	X	X	X
Excl. Mexican		X		X

Notes: This table shows the results of estimating Equation 5 in a sample of households with foreign-born financial respondents, who are (as close as possible to) 60 years old and did not live in the US while they were in school, in the Health and Retirement Study. The independent variable is the household’s financial respondents homeownership in their country of origin (HOCO) matched to the respondent’s marital status. The dependent variable in columns 1 and 2 indicates whether the respondent owns a primary residence worth at least \$5K. The dependent variable in column 3 and 4 represents the percent of the respondent’s total wealth that the value of the primary residence represents. Demographic controls include: categorical variable dummies for respondent marriage status, age, total household income, a dummy for gender, a dummy for whether they are at least a high school graduate, and a dummy for whether the respondent was born in Mexico. All regressions also include a dummy indicating the survey wave that the respondent data is from. Where indicated, specifications also control for total non-housing wealth of the household, exclude Mexican-born respondents, and exclude respondents who lived in the U.S. during their school time. Note: Standard errors are clustered at the country-of-origin level. T-statistics shown in parentheses.

Table VIII: Effect of HOCO on Wealth and Risky Asset Portfolio Holdings Among Foreign-Born at Age 60

<i>Dep. Var.</i>	All							
	Only Homeowners							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Net Value of Primary Res.	Total Wealth excl. Primary Res.	$\mathbb{1}[\text{OwnsStocks}]$	Stock Share of Wealth (excl. RE)	Net Value of Primary Res.	Total Wealth excl. Primary Res.	$\mathbb{1}[\text{OwnsStocks}]$	Stock Share of Wealth (excl. RE)
HOCO	3475.2*** (2.909)	4585.4** (2.310)	0.185 (1.588)	0.095 (1.005)	3280.8** (2.191)	6044.3** (2.295)	0.300 (1.551)	0.084 (0.630)
Mean of dep. var.	110,452	128,912	6.4	3.9	169,255	183,287	9.2	4.9
R-squared	0.22	0.29	0.35	0.18	0.23	0.29	0.35	0.18
Observations	466	466	466	296	304	304	304	215
Demographic controls	X	X	X	X	X	X	X	X
Non-housing wealth control			X	X			X	X
School outside US	X	X	X	X	X	X	X	X

Notes: This table shows the results of estimating Equation 5 in a sample of households with foreign-born financial respondents, who are (as close as possible to) 60 years old and did not live in the US while they were in school, in the Health and Retirement Study. The independent variable is the household's financial respondents homeownership in their country of origin (HOCO) matched to the respondent's marital status. Columns 5-8 further restrict the sample to only include homeowners. The dependent variables are as follows: in columns 1 and 5 they are the net value in dollars of the respondent's primary residence; in columns 2 and 6 they are the total respondent wealth excluding the primary residence; in columns 3 and 7 they are an indicator that is 100 if the respondent owns stocks in their portfolio, and 0 otherwise; in columns 4 and 8 they are the share of the respondent's wealth (excl. primary and secondary residence, and other real estate investments) that is invested in stocks. Demographic controls include: categorical variable dummies for respondent marriage status, age, total household income, a dummy for gender, a dummy for whether they are at least a high school graduate, and a dummy for whether the respondent was born in Mexico. All regressions also include a dummy indicating the survey wave that the respondent data is from. Where indicated, specifications also control for total non-housing wealth of the household, and exclude respondents who lived in the U.S. during their school time. Note: Standard errors are clustered at the country-of-origin level. T-statistics shown in parentheses.

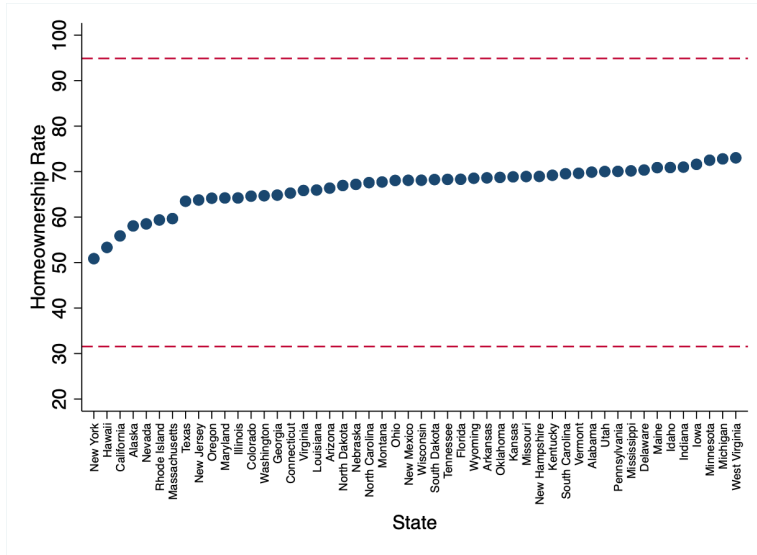
Table IX: Effect of HOCO on Retirement Wealth By Regional House Price Growth

Dep. Var:	Net Value of		Total Wealth	
	Primary Res. (1)	Net Value of Real Estate (2)	excl. All RE (3)	$\mathbb{1}$ [Owns Stocks] (4)
HOCO \times 1[1st Terc. 20-Yr House Price Growth]	2099.460 (1.234)	778.507 (0.241)	-176.744 (-0.090)	0.160 (0.630)
HOCO \times 1[2nd Terc. 20-Yr House Price Growth]	237.241 (0.105)	2876.333 (0.714)	2624.532 (0.884)	-0.183 (-0.483)
HOCO \times 1[3rd Terc. 20-Yr House Price Growth]	6800.043** (2.025)	16575.391** (2.529)	10147.291 (1.575)	1.034* (2.012)
Mean of dep. var.	110,740	168,059	72,883	6.3
R-squared	0.27	0.28	0.27	0.27
Observations	457	457	457	457
Demographic controls	X	X	X	X
School outside US	X	X	X	X
State HP Growth Tercile	X	X	X	X

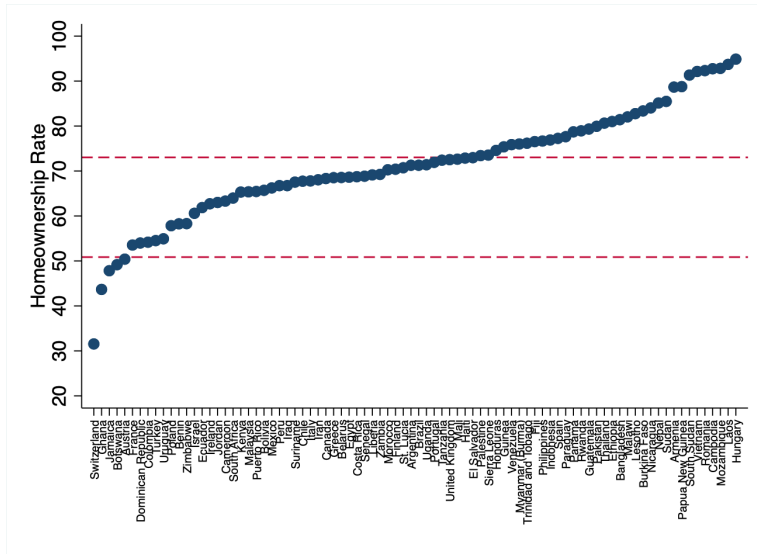
Notes: This table shows the results of estimating Equation 5 in a sample of households with foreign-born financial respondents, who are (as close as possible to) 60 years old and did not live in the US while they were in school, in the Health and Retirement Study. The independent variable is the household's financial respondents homeownership in their country of origin (HOCO) matched to the respondent's marital status. The independent variable is interacted with terciles of *20-Year State House Price Growth*, which is the growth rate in percent of the FHFA HPI measure of house prices between the year of the survey wave and 20 years earlier. It is computed at the state level and assigned to each respondent based on the survey year and their state of residence at the time. The terciles reflect the distribution among the full sample of respondents at age 60. The dependent variables are as follows: in column 1 it is the net value in dollars of the respondent's primary residence; in column 2 it is the total value of the respondent's primary residence, secondary residence, and other real estate investments; in column 3 it is the total respondent wealth excluding the value of the primary residence, secondary residence, and other real estate investments; in column 4 it is an indicator for whether the respondent owns any stocks, scaled by 100. Demographic controls include: categorical variable dummies for respondent marriage status, age, total household income, a dummy for gender, a dummy for whether they are at least a high school graduate, and a dummy for whether the respondent was born in Mexico. All regressions also include a dummy indicating the survey wave that the respondent data is from. Specifications also exclude respondents who lived in the U.S. during their school time. Note: Standard errors are clustered at the country-of-origin level. T-statistics shown in parentheses.

Figures

Figure 1: Variation in Homeownership Across U.S. States vs. Countries



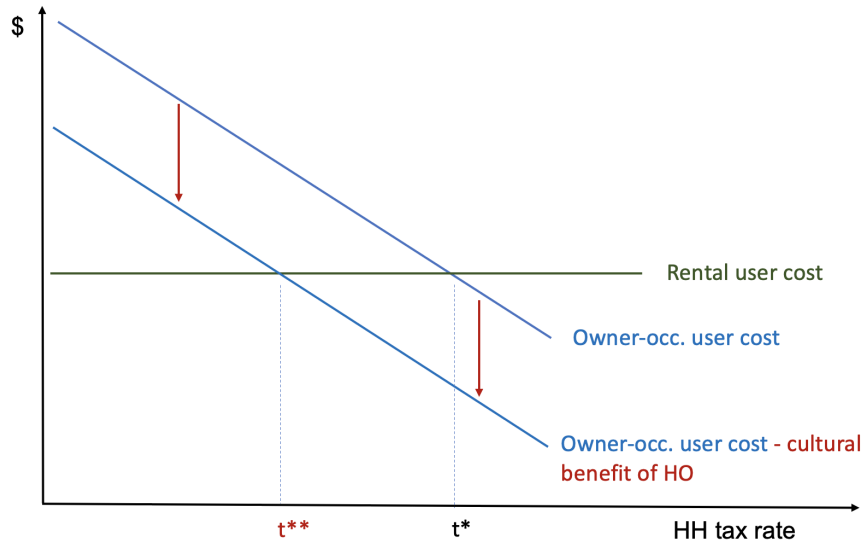
(A) U.S. States



(B) Origin Countries

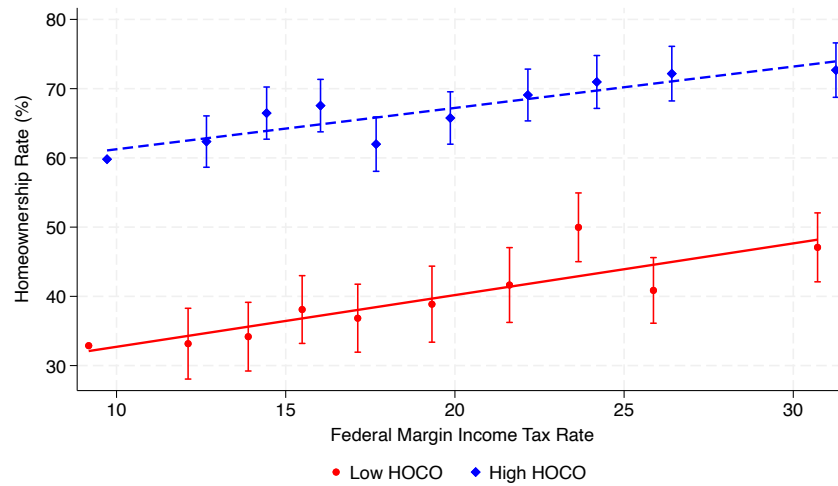
Notes: These graphs show the variation in homeownership rates across U.S. states, among U.S. residents, and 84 origin countries, among populations in those countries. The top graph shows the state-level homeownership rate using the average homeownership rate within a state across the 1970, 1980, 1990, 2000, and 2010 censuses. We also plot the country-level minimum and maximum homeownership rates, for ease of comparison to the bottom graph. The bottom graph draws on the IPUMS harmonized censuses among households with heads aged 18-69, for the available year closest to 2015. We plot the state-level minimum and maximum homeownership rates, for ease of comparison to the top graph. In both graphs, homeownership rates are ordered from smallest (L) to largest (R).

Figure 2: Static tenure choice with cultural affinity for homeownership



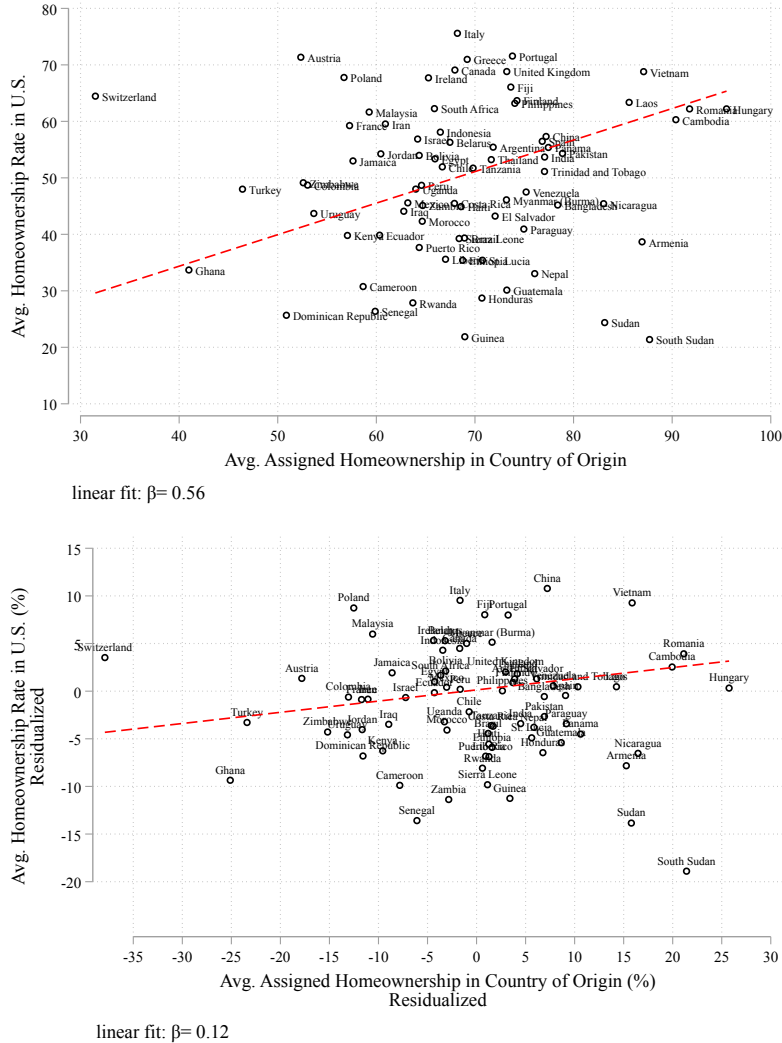
Notes: This figure shows the user-cost vs. tax rate profile for households that rent or own, as outlined in Section II. Here “cultural benefit of homeownership” is analogous to ϕ .

Figure 3: Homeownership Rate by Federal Marginal Income Tax Rate



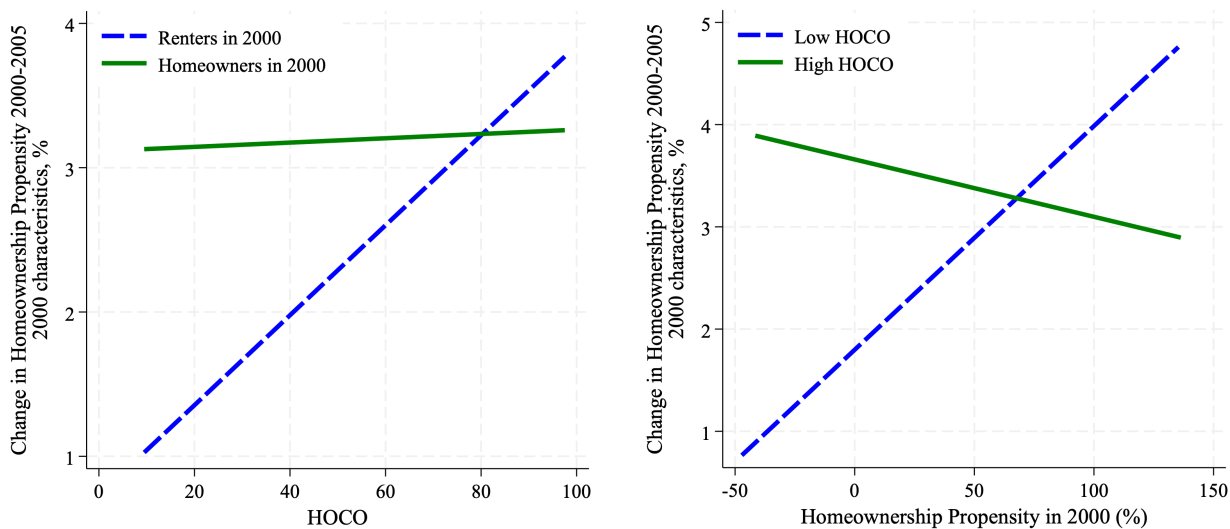
Notes: This table shows a binned scatter of homeownership rates by federal marginal income tax rate, split by immigrants with above- or below-median HOCO. Socioeconomic characteristics incorporated into the TAXSIM model used to produce marginal federal income tax rates. We control for CBSA-by-year fixed effects, as well as housing unit characteristics: number of bedrooms, bathrooms, lotsize in acres, unit size in ft², decade built, structure type (i.e. single family or multifamily), and presence of a garage or carport. Data from the AHS 2001-2019.

Figure 4: Relationship between Homeownership in the U.S. and *HOCO* among Immigrants at the Origin Country Level



Notes: These graphs shows origin country averages of homeownership among U.S. households with foreign-born heads on the vertical axis, and the average of these households' associated homeownership in the country of origin by marital status on the horizontal axis. The left graph shows these averages for the raw ACS data from the years 2000 and 2005-2019, while the right graph residualizes both homeownership and *HOCO* at the household level before aggregating. The control variables for the residualization consist of: HH Income, (linear and ventile indicators), Quadratic function of Age of HH head and 5-year age group indicators, indicators for educational achievement of HH head, Indicators for # of children and # of relatives in the HH, Indicators for discretized years since immigration of head, marital status of HH head, and the interaction between marital status and years-since-immigration categories. The linear fits and slopes in each graph weight each country by the number of foreign-born residents with that origin in the sample. Note that the aggregation means that the regression slopes at the origin country level are merely illustrative here; they are not directly comparable to the household level effect estimates, which should be relied on for the quantitative magnitude of the effect.

Figure 5: Change in Homeownership Propensity during the 2000-2005 Credit Supply Expansion by HOCO

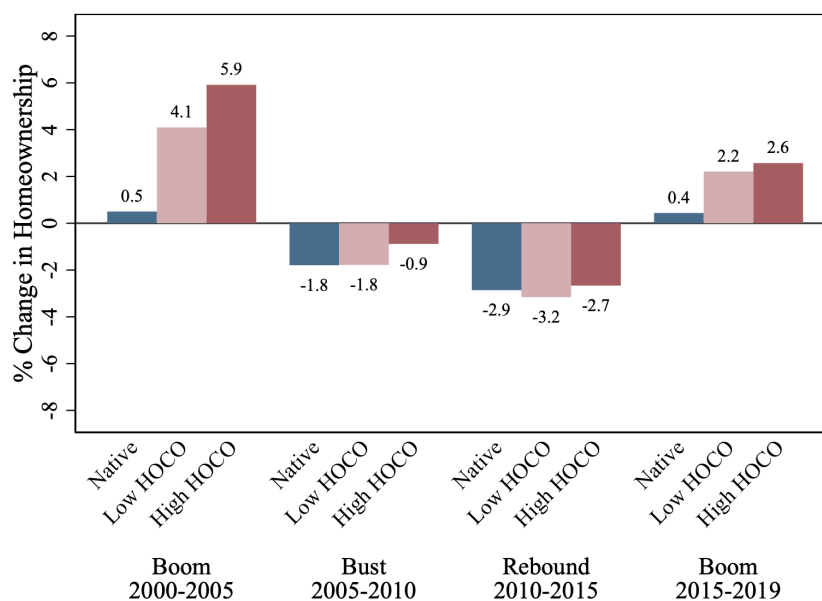


(A) HO advances by HOCO and Tenure

(B) HO advances by HO propensity in 2000

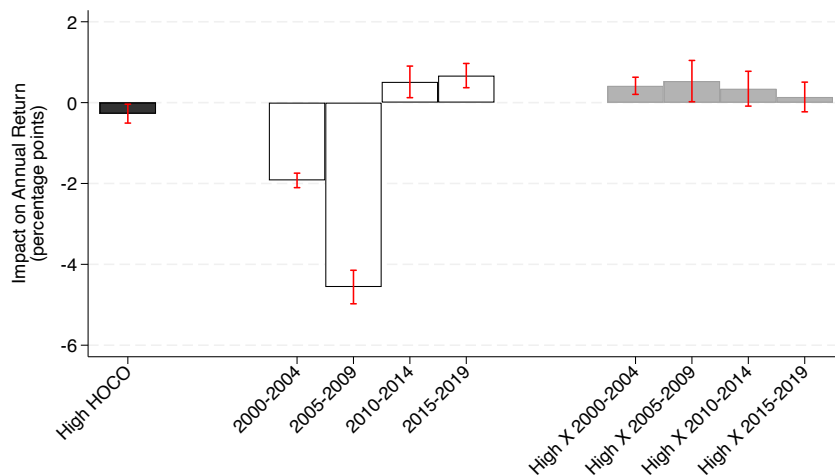
Notes: These graphs show how the homeownership propensity associated with particular demographics changed during the 2000-2005 boom as a function of different household characteristics. Each panel shows estimated linear fits between a household characteristic and the change in the homeownership propensity for different subgroups of foreign-born household heads. The propensity to own a home is the linear predicted homeownership rate for a household using the following demographic variables for the prediction: county fixed effects; Real HH Income (linear and ventile indicators); Quadratic function of Age of HH head and 5-year age group indicators, indicators for educational achievement of HH head, Indicators for # of children and # of relatives in the HH, Indicators for discretized years since immigration of head, marital status of HH head, and the interaction between marital status and years-since-immigration categories. The change in the homeownership propensity 2000-2005 in both graphs uses ACS 2000 data, and uses coefficients for predicting homeownership propensity estimated separately using cross-sectional year 2000 data and year 2005 data for foreign-born headed households to predict how likely a household with those characteristics would be to own a home. The change in this propensity holds household characteristics constant at their 2000 level and estimates the predicted change in homeownership due to changes in the coefficients between the two time periods. Low and High HOCO groups in the right panel are defined as households with foreign-born heads with HOCO below or above the median among all foreign-born heads. Renters and homeowners in the left panel are defined based on 2000 characteristics. Both panels exclude Mexican-origin household heads. The linear prediction for homeownership propensity does not constrain homeownership propensity to range from 0 to 100%.

Figure 6: Homeownership Change by HOCO and Period



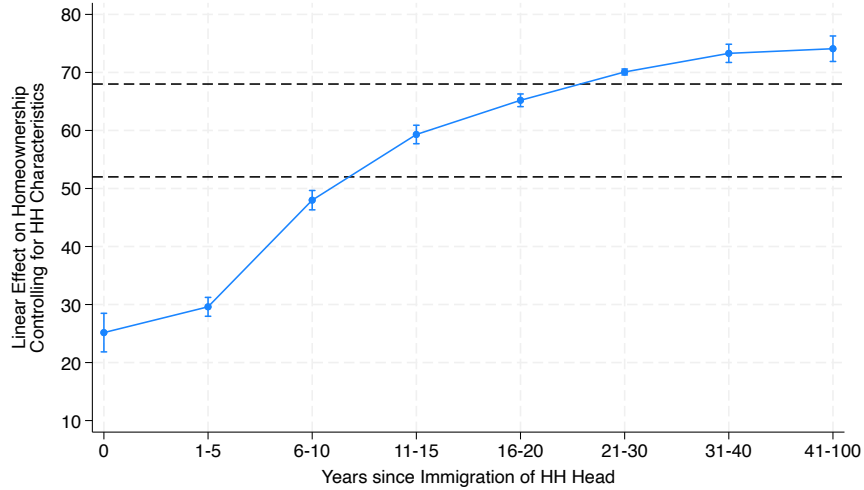
Notes: This figure shows the average change by period in the homeownership in percentage points among U.S.-born American residents, and foreign-born residents with a homeownership in their country of origin (*HOCO*) that is below-median or above-median. The graph and the assignment into HOCO groups exclude Mexican origin foreign-born and residents where we were unable to assign a *HOCO* score.

Figure 7: Differential Annual Returns, by Year of Home Purchase

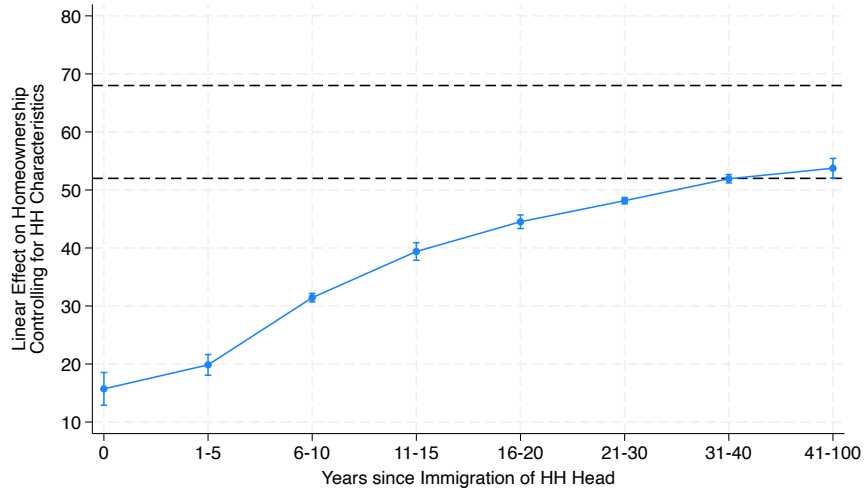


Notes: This figure shows the additional annual returns to homeownership for above-median HOCO households, relative to below-median HOCO households, by era of home purchase. The annual returns are imputed by first computing the annual appreciation in the real value of the home by converting the ratio between the current reported real (CPI-adjusted) market value of a respondent’s home in the American Housing Survey in a particular year and the original real purchase price of the home into an annual growth rate in the real house price. Then, we run regressions that predict this expected annual appreciation on HOCO group-by-purchase period indicators and other variables, including the standard socioeconomic and demographic controls used in other specifications, as well as CBSA by year fixed effects. To control for differential returns to housing by quality, we additionally control for bedrooms, bathrooms, lot size, square footage, the year built, presence of a garage, and the unit size (i.e. single family, duplex, small multifamily). The differential returns shown in the figure correspond to the estimated coefficient estimates for the HOCO group-by-purchase period indicators in this prediction of annual returns due to real price appreciation.

Figure 8: Foreign Born Homeownership in the US, by Years in US (foreign-born only))



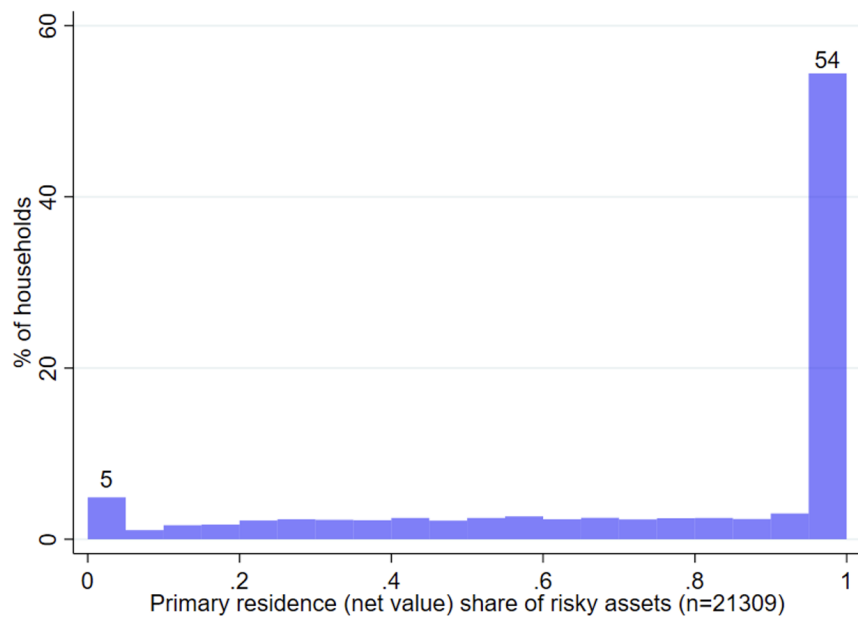
(A) Above-Median HOCO



(B) Below-Median HOCO

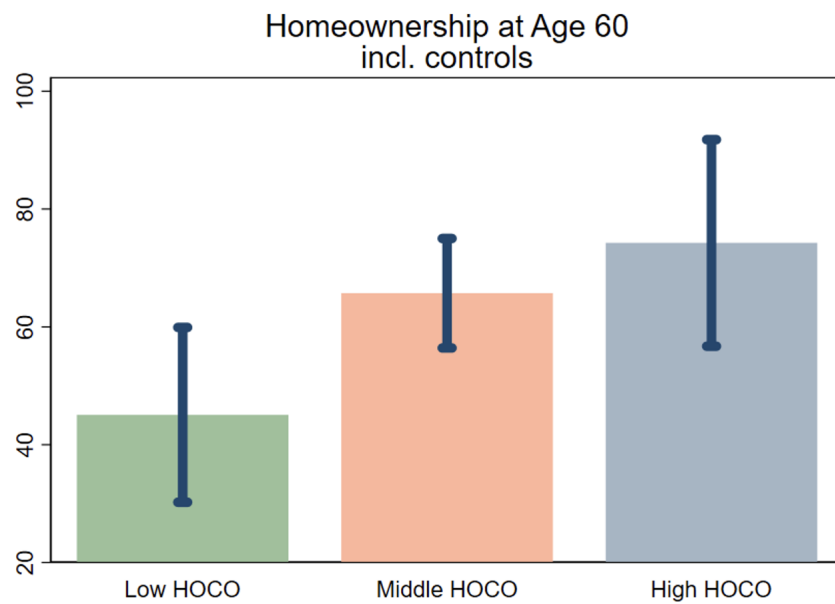
Notes: These graphs show the estimated homeownership rate among foreign-born households by years elapsed since arriving in the US, split by having above- median (top panel) or below-median (bottom panel) HOCO. We control for commuting zone-by-year fixed effects as well as household income, household head’s age and age squared, educational attainment, and marital status, and cluster standard errors by commuting zone and year. Data from the ACS microdata 2005-2019. Horizontal dashed lines represent the mean homeownership rates among households with foreign-born heads (lower) and homeownership rates of households with US-born heads (upper).

Figure 9: Primary residence share of discretionary risky assets at retirement



Notes: This figure shows the distribution of respondents at (or near) age 60 in the Health and Retirement Study by the share of their "discretionary" risky assets (stocks, bonds, and real estate assets held outside of pensions or retirement accounts) that consists of the net value of the primary residence.

Figure 10: Homeownership by HOCO Tercile at Retirement Age



Notes: This figure shows the average homeownership in percentage points among foreign-born respondents in the Health and Retirement Study with a homeownership in their country of origin that is in the lowest, middle, or highest tercile of homeownership in the country of origin (*HOCO*). The graph and the assignment into terciles exclude Mexican origin foreign-born and residents where we were unable to assign a *HOCO* score. The homeownership rates are adjusted for the effect of the same control variables as included in column 2 of Table VII.

ONLINE APPENDIX

Appendix A. Wealth Imputation

In this section, we discuss the imputation of household non-housing networth using the PSID and ACS data. The PSID structure follows an initial set of households recruited in 1968, adding new households as the children of the original households form their own households, or as households otherwise split. This structure does not accommodate population representativeness as immigration evolves over time. As such, we limit the PSID sample to the 1997 and 2017 waves. These refresh the sample of households being surveyed by adding new households that had migrated to the United States between 1968 and 1997 and between 1997 and 2017, respectively. We end up with 1,046 immigrant households whose heads are between the ages of 18 and 69, in keeping with our other sample restrictions.

We construct non-housing networth following Pfeffer, Schoeni, Kennickell, and Andreski (2016) and predict networth using a suite of socioeconomic and demographic variables. We add up checking and savings accounts, the value of business assets, stocks, IRAs, private annuities, vehicle networth (value less remaining auto loans) and any other real estate that is not the primary residence. We then compute the household's liabilities, adding credit card debt, loans from relatives, medical debt, legal debt, student debt and other non-specified debt, excluding primary and secondary mortgages on the primary residence. Non-housing networth is just non-housing assets less non-housing liabilities.

The predictive variables include fourth order polynomial in household income, race and hispanic status of household head and spouse (if applicable), household head's level of education, age and gender, household head's age interacted with education to control for differential income profiles, whether there are children present in the home, whether the householder is married, as well as the interaction between marriage and children. Finally, we include census division fixed effects to control for differential non-housing wealth accumulation patterns across U.S. geographies, as well as continent fixed effects for immigrants' country of origin reflecting the fact that those immigrating from Europe may have differential wealth profiles to those from South America, for example. We are limited in the amount of geographic variation we can control for as our sample of $\sim 1,000$ immigrant households hail from only 38 states. Additionally, only 22 states have more than 10 households, limiting the residual variation should we control for state directly. For similar reasons, we cannot control for country of origin; the PSID only provides 16 broad regions which we map to continents as many of these regions have very few households. Observations are weighted by the cross-sectional family weight provided by the PSID.

Figure BII plots the predictive ability of our model. The regression's R^2 is 0.22, suggesting that many wealth determinants have not been directly accounted for. Nonetheless, the model fits the data quite well up to networths of \$800,000, suggesting that the remaining wealth determinants are most predictive of the right tail (we have only 17/1,049 households with non-housing networths

above \$800,000), not the main sample.

Next, we store the point estimates associated with each of our socioeconomic and demographic predictors of wealth and apply them to our sample of migrants in the ACS data. This allows us to impute non-housing networth as the sum of the model inputs times their respective parameters. Figure BIII compares the actual distribution of non-housing networths observed in the PSID to the imputed non-housing networths we build using the ACS data and PSID-derived coefficients. The ACS immigrant sample, composed of nearly 4 million households, yields a distribution with more networth variation than that in the PSID, which is based on only about 1,000 households. Nonetheless, visual inspection of the two distributions shows many similar characteristics: mean networth near 0, with a significant right tail after about \$500,000, and small mass of households with negative networth between \$0 and -\$500,000.

Appendix B. Health and Retirement Study Data

The overall goal of our analysis using Health and Retirement Study data is to estimate the effect of the homeownership in foreign-born respondents’ countries of origin on their portfolio holdings at the end of their working life (around age 60). The two data sources used in this analysis are:

1. RAND HRS Longitudinal File (provided by HRS)
2. Homeownership by marriage status and country of origin (“HOCO”), collected from public data sources.

We transform and restrict the RAND HRS data in the following way to obtain our analysis sub-sample:

- Merge RAND HRS and HOCO data based on country of birth information (only available in the restricted access files) and also by marriage status
- Identify which survey waves occur closest to age 60 for each respondent and keep only respondents that are 55-65 years old at the time of at least one wave
- Retrieve all variables of interest from the survey wave when the respondent is (closest to) 60 years old (see variable definitions below)
- All final analyses also subset the sample to only include the foreign-born (for whom we have HOCO data), to only include one respondent per household (in couples, the financial data respondent) and only include respondents that are married or partnered at the time of the wave (in order to make households comparable)
- Some analyses additionally subset the sample to exclude respondents born in Mexico to show that the results are robust to excluding this largest country of origin; some analyses also subset only to respondents who did not live in the US when they were going to school.
- Analyses that concern wealth shares additionally condition on the numerator and denominator of portfolio share measures being non-negative and positive, respectively.

The analyses shown in the corresponding results tables consist of cross-sectional regressions that regress the portfolio holdings of the respondent’s household at age 60 on their HOCO and additional household-level control variables – and the tables only report the coefficient on the HOCO variable. The sample sizes of the regression samples are indicated in the row “Observations” in the tables. The smallest regression sample represents 41 different countries of origin. The independent variable is continuous, so it represents an average effect for the full group indicated by that sample size.

Country HOCO terciles. For some analyses, we tabulate average homeownership rates by tercile of HOCO (low, medium, high), with or without controlling for the effect of demographic control variables. The HOCO tercile groups (Low, Middle, High) shown in the homeownership figures consist of the following numbers of respondents from the HRS:

The first group represents respondents from 9 different countries, the second and third group represent 16 different countries each. For reasons of data privacy, we are not allowed to provide further details on the countries included in each group.

Hoco Tercile	Respondents	Percent
1	49	21.97
2	97	43.50
3	77	34.53
Total	223	100.00

Definitions of key outcome and independent variables in HRS analyses

We use or construct the following variables for the HRS analysis, with variable names in the HRS RAND Longitudinal data noted in parentheses where applicable:

HOCO: Homeownership in country of origin, manually collected from government statistics agencies or computed from public IPUMS data and imported into the VDI at the country-by-marriage-status level. Continuous variable

Homeownership at Age 60: Binary indicator for whether the gross value of the respondent's primary residence (*HwAHOUS*) at the time of the survey wave closest to their 60th birthday is greater than \$5K.

Home Value Share of Total Wealth: Percentage measure that indicates the ratio of the net value of the respondent's primary residence (*HwATOTH*) to total wealth (*HwTOTB*) at the time of the survey wave closest to their 60th birthday. Set to missing if the numerator is negative, or the denominator is non-positive.

Total Wealth: Total wealth (*HwTOTB*) in dollars at the time of the survey wave closest to their 60th birthday.

Net Value of Primary Residence: Net value of the respondent's primary residence (*HwATOTH*) in dollars at the time of the survey wave closest to their 60th birthday.

Total Wealth excl. Primary Residence: Total wealth (*HwTOTB*) in dollars minus net value of the respondent's primary residence (*HwATOTH*) in dollars at the time of the survey wave closest to their 60th birthday.

1[Owns Stock]: Binary indicator for whether the gross value of the respondent's assets in stocks (*HwASTCK*) at the time of the survey wave closest to their 60th birthday is greater than \$5K.

Stock share of Wealth excl. Real Estate: Percentage measure that indicates the ratio of respondent's assets in stocks (*HwASTCK*) to total non-housing wealth (*HwTOTN*) at the time of the survey wave closest to their 60th birthday. Set to missing if the numerator is negative, or the denominator is non-positive.

Definitions of additional control variables used in HRS analyses:

Demographic controls: These control variables (included where indicated in the tables and figures) are all for the respondent characteristics at age 60 and consist of categorical variable dummies for respondent marriage status (*RwMSTAT*), a continuous variable for age (*RwAGEY*), a dummy for gender (*RwAGENDER*) and dummies for the survey wave that the Age 60 data for

the respondent is from. It also includes a dummy for whether the respondent was born in Mexico. Additionally, it controls for continuous total household income ($HwITOT$) and an indicator for respondent education, which is a binary for whether they are a high school graduate ($RwAEDUC$ is 3, 4, 5).

Non-housing wealth controls: These control variables (included where indicated in the tables and figures) are all for the respondent characteristics at age 60 and consist of the total non-housing wealth ($HwTOTN$) in dollars.

School outside US: This is a binary indicator for whether the foreign-born respondent lived in a US state when they went to school. It is constructed from the restricted access geographic variable WHRLIV10 and is a binary indicator. Where indicated, the analyses exclude anyone who lived in the U.S. when they went to school from the sample.

Appendix C. Additional Tables

Table AI: Effects of HOCO on Tenure Choice: Country Characteristic Controls

<i>Dependent var.:</i>	Household Homeownership _{it} (in %)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
HOCO (%)	0.092*** (0.03)	0.165*** (0.05)	0.157*** (0.05)	0.163*** (0.05)	0.108** (0.05)	0.179*** (0.04)	0.168*** (0.04)	0.167*** (0.03)	
GDP per capita (2021 USD)		0.082 (0.05)						-0.098 (0.13)	
Urban population (% of total)			0.008 (0.04)					0.118** (0.05)	
Consumer price inflation (annual %)				-0.147 (0.12)				0.012 (0.24)	
Domest. credit by fin. sector (% of GDP)					0.030*** (0.01)			0.022 (0.02)	
Property rights index						0.032 (0.05)		0.106 (0.10)	
Investment freedom index							-0.001 (0.06)	-0.074 (0.07)	
Observations	3,436,252	3,436,252	3,385,624	3,223,662	2,100,486	3,175,701	3,175,701	2,060,006	
R-Squared	0.32	0.31	0.31	0.31	0.30	0.31	0.31	0.30	
				<i>Addl. controls & fixed effects</i>					
Household Characteristics	X	X	X	X	X	X	X	X	
Imputed HH Wealth	X	X	X	X	X	X	X	X	
Commuting Zone × Year FE	X	X	X	X	X	X	X	X	
Major religion FEs	X							X	

Notes: This table shows the results of estimating the same specification as column 4 of Table II, and additionally controlling for country-of-origin characteristics. Each column adds in one characteristic as a control variable. On the one hand, we add Major religion fixed effects (column 1) and 2000-2020 averages of economic characteristics from the World Bank's World Development Indicators: GDP per capita (column 2), urbanization rate (column 3), annual inflation rate (column 4), and domestic credit to private borrowers as a share of GDP (column 5). On the other hand, we include indicators of property rights (column 6) and investment freedom (column 7) from the Heritage Foundation, which are averaged over 2000-2008, and for which higher numbers indicate more rights. Column 8 jointly includes all of these characteristics and fixed effects. Heteroskedasticity-robust standard errors clustered at the origin country level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table AII: Effects of HOCO on Tenure Choice: Immigrated as Children

<i>Dependent var.:</i>	Household Homeownership _{it} (in %)			
	(1)	(2)	(3)	(4)
HOCO (%)	0.995*** (0.09)	0.147*** (0.05)	0.138*** (0.04)	0.122*** (0.04)
Observations	1,224,869	1,221,160	1,221,160	1,220,817
R-Squared	0.08	0.26	0.29	0.31
Household Characteristics	No	Yes	Yes	Yes
Commuting Zone FE	No	No	Yes	No
Commuting Zone \times Year FE	No	No	No	Yes

Notes: This table shows the results of estimating Equation 1 in a pooled household level sample of households with foreign-born heads who are 18-69 years old, for the years 2000 and 2005-2019, and *who immigrated to the US when they were 18 years old or younger* as inferred from their age and immigration year. The dependent variable is an indicator of homeownership in percent (so 100 indicates that the household owns their home). The independent variable is the homeownership rate in the country of origin (*HOCO*) of the household head for people with the same marriage status (single, married, separated/divorced, widowed). Household characteristic control variables consist of: HH Income, (linear and ventile indicators), predicted non-housing household net worth, Quadratic function of Age of HH head and 5-year age group indicators, indicators for educational achievement of HH head, Indicators for # of children and # of relatives in the HH, Indicators for discretized years since immigration of head, marital status of HH head, and the interaction between marital status and years-since-immigration categories. Heteroskedasticity-robust standard errors clustered at the origin country level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table AIII: Changes in homeownership propensity 2000-2005 by HOCO

Dep. Var.:	Δ Predicted Homeownership 2000-2005 (%)			
	Renters	Owners	Low Hoco	High HOCO
Sample:	(1)	(2)	(3)	(4)
HOCO	0.031*** (2.822)	0.001 (0.182)		
Homeownership propensity (2000)			0.022*** (4.321)	-0.006 (-1.173)
Constant	0.731 (1.020)	3.114*** (5.516)	1.798*** (4.503)	3.659*** (6.492)
Observations	177,432	201,270	183,135	195,567

Notes: This table shows the estimated coefficients for the lines of best fit shown in Figure 5. This table shows the linear estimate of the association between a household's HOCO (columns 1 and 2) and predicted propensity to own a home in the year 2000 (columns 3 and 4) with the change in predicted homeownership, holding the household's demographics constant, over the 2000-2005 period. Column 1 includes only renter households as of 2000, and column 2 includes only homeowners in the sample. Column 3 only includes households with below median HOCO in the sample, and column 4 only includes above median HOCO households. The predicted homeownership in 2000 and 2005 is computed from a linear prediction of homeownership based on household demographics in the cross-section in each year - see Figure 5 for further details. T-statistics based on standard errors clustered at the country of origin level are shown in parentheses.

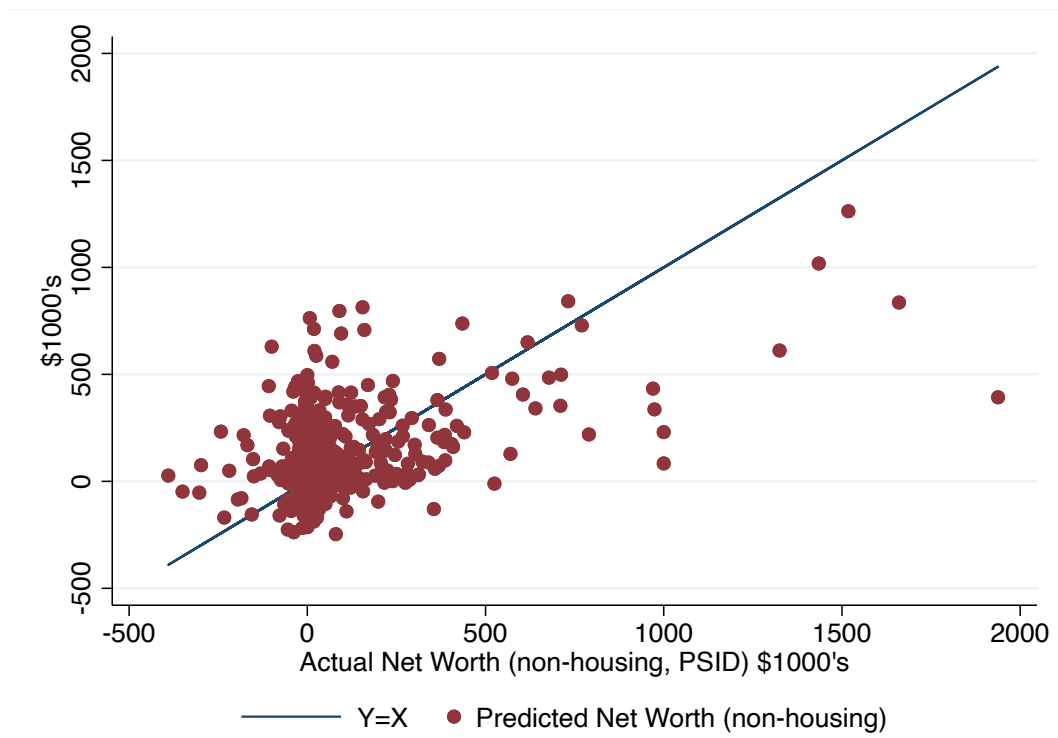
Appendix D. Additional Figures

Figure BI: Correlations between *HOCO* and Country Characteristics

Homeownership rate (% of households)						
-0.419	GDP per capita (2021 US\$), 2010					
-0.475	0.524	Urban population (% of total population)				
0.185	-0.184	-0.096	Consumer price inflation (annual %)			
-0.103	0.436	0.165	-0.354	Domestic credit provided by financial sector (% of GDP)		
-0.368	0.778	0.502	-0.367	0.403	Property rights (Heritage index)	
-0.229	0.486	0.468	-0.429	0.051	0.703	Investment freedom (Heritage index)

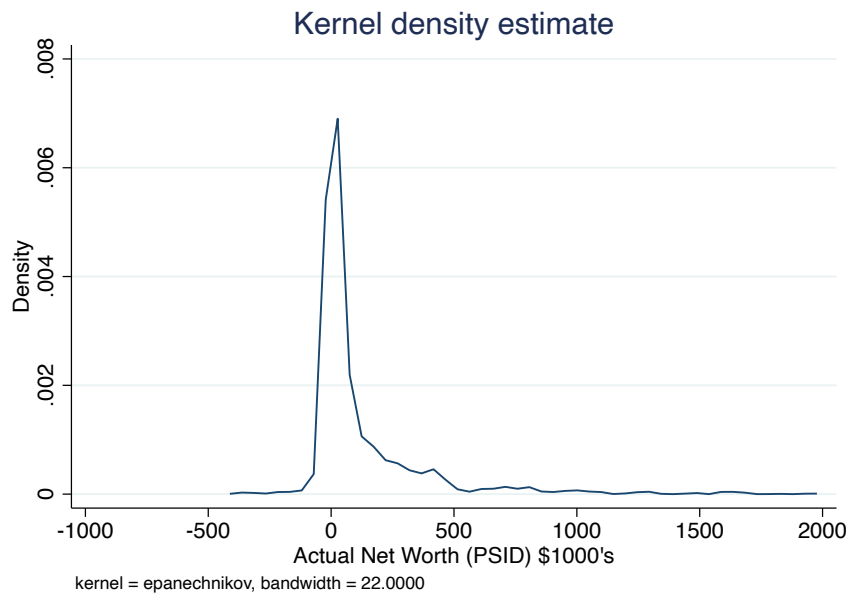
Notes: This matrix shows the correlations between the average (not marital-status-specific) homeownership rate in different countries and other country characteristics. On the one hand, the characteristics consist of economic characteristics like GDP per capita (in constant 2021 USD), urbanization rate, annual inflation rate, and domestic credit to private borrowers as a share of GDP - all of which are computed as averages of all available data for 2000-2020 from the World Bank's World Development Indicators. On the other hand, we include indicators of property rights and investment freedom from the Heritage Foundation, which are averaged over 2000-2008, and for which higher numbers indicate more secure rights and greater freedom.

Figure BII: Actual vs. Predicted Net worth in the PSID

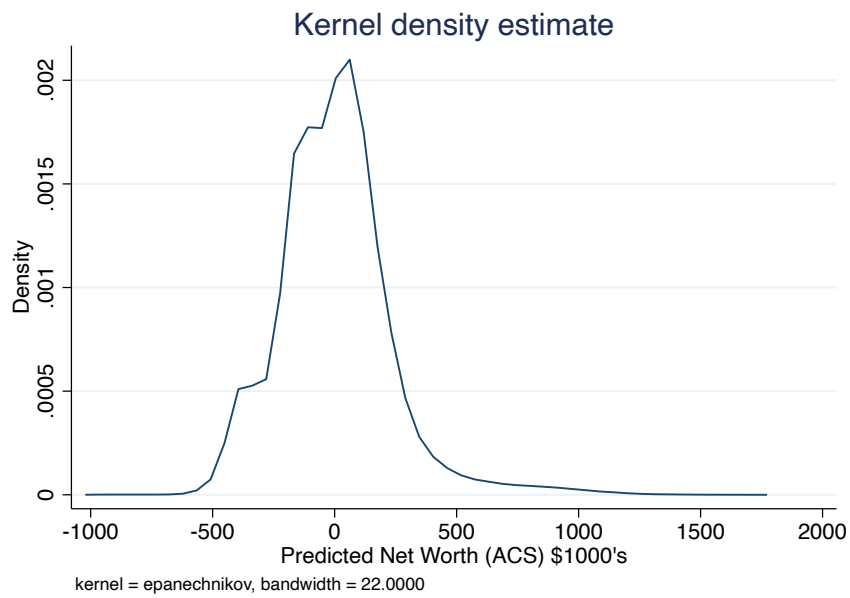


Notes: These figure plots a scatterplot of actual non-housing network vs. predicted housing net worth as in section A. We overlay the scatterplot with the 45 degree line which would reflect perfect prediction. Data comes from the 1997 and 2017 PSID waves, restricted to immigrant households. Net worth constructed according to Pfeffer et al. (2016), excluding assets and liabilities associated with the primary residence. Distribution of actual non-housing net worth trimmed at \$2 million for ease of inspection due to the long right tail in wealth, which drops 8/1,049 households from the figure.

Figure BIII: Distributions of Net worths, by Dataset



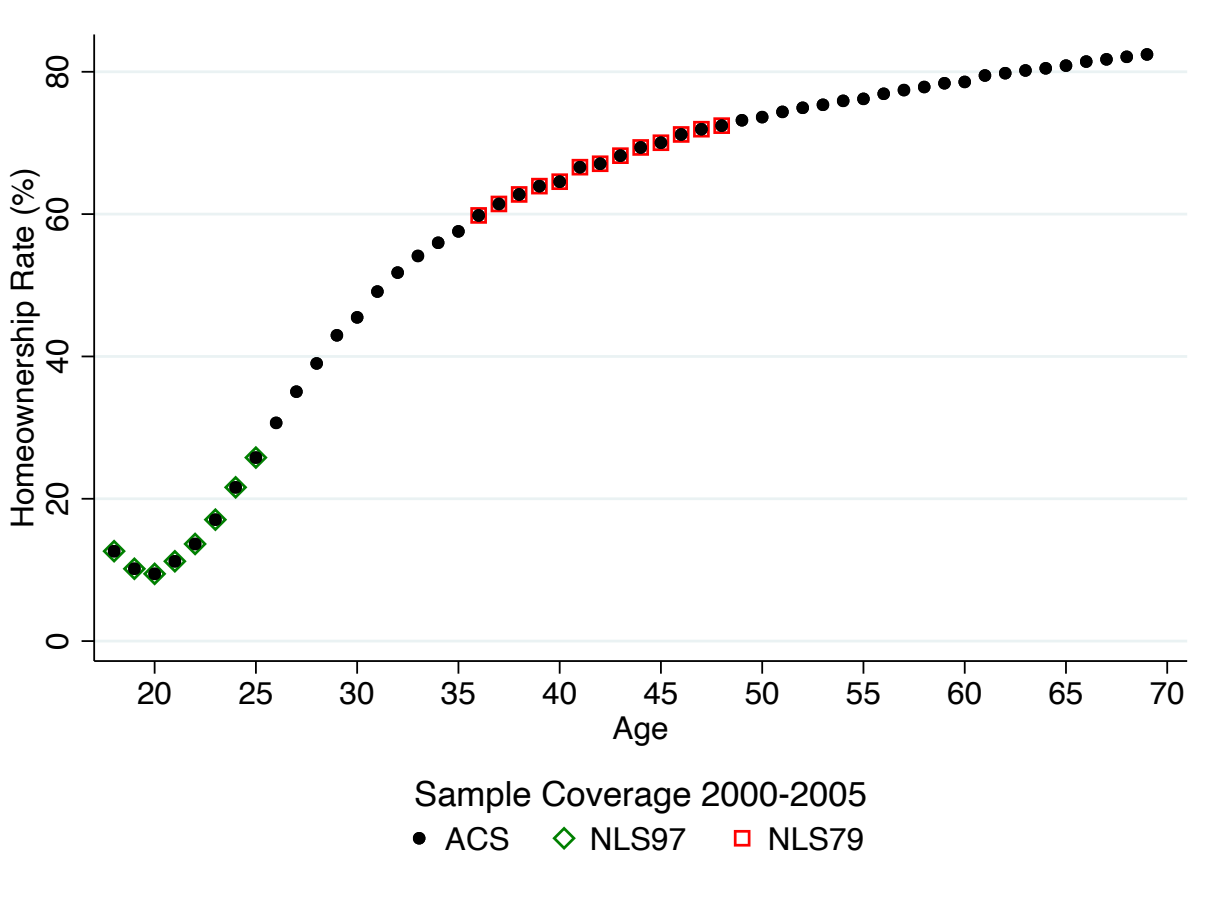
(A) PSID



(B) ACS

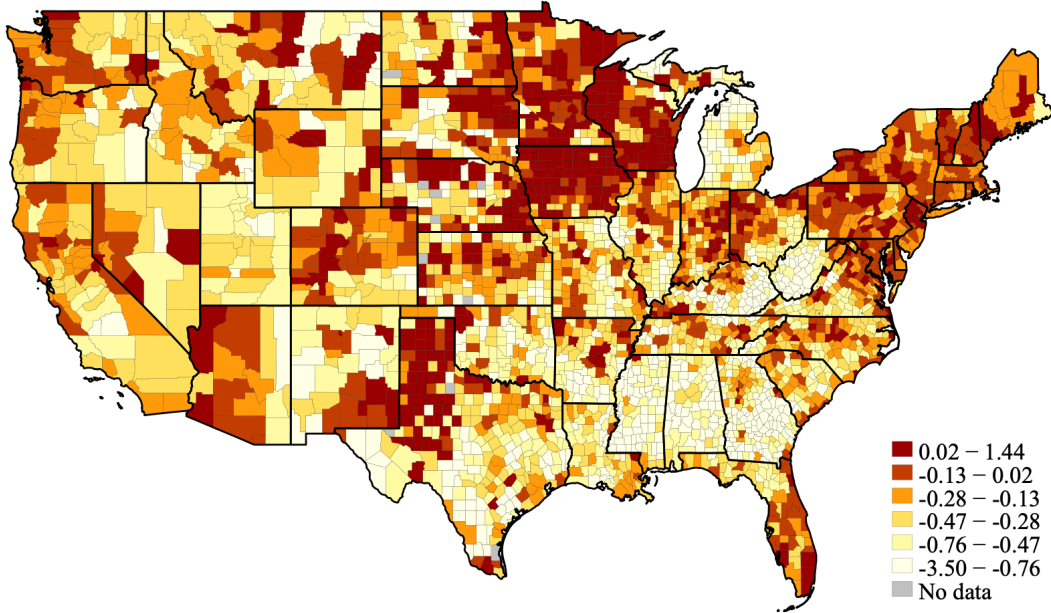
Notes: These figures plot the kernel density distributions of non-housing net worths from the PSID and ACS immigrant samples. We show actual non-housing net worth for the PSID, and imputed net worth according to Appendix A for the ACS sample. Distribution of actual non-housing net worth in the PSID trimmed at \$2 million for ease of inspection due to the long right tail in wealth, which drops 8/1,049 households from the figure.

Figure BIV: Sample Coverage of Candidate Datasets

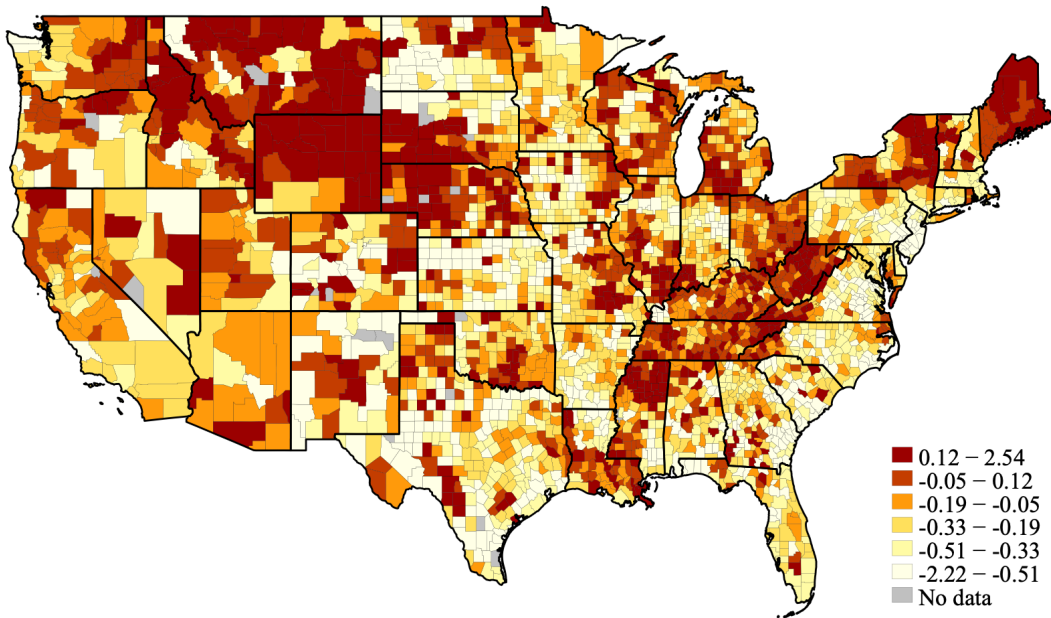


Notes: This plot shows the available sample coverage between 2000 and 2005, our period of study for credit shock analysis as in Section V. The green diamonds highlight the ages available for study in the National Longitudinal Survey of Youth, 1997 (NLS97) between 2000-2005. The red squares note the ages available for the NLS79. The solid dots plot the ages available using the US Census and ACS microdata.

Figure BV: Spatial Distribution of credit supply shocks



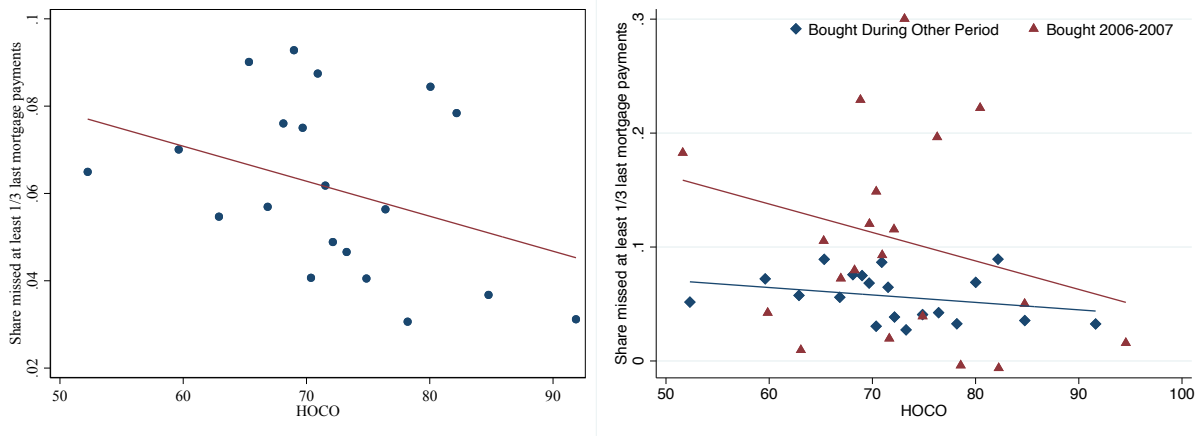
(A) Credit supply shocks: 2000-2005



(B) Credit supply shocks: 2010-2015

Notes: These maps plot the distribution of credit supply shocks at the county level. Darker colors have higher values of implied loan growth, and lighter colors denote lower values.

Figure BVI: Delinquencies by HOCO, and period of Home Purchase

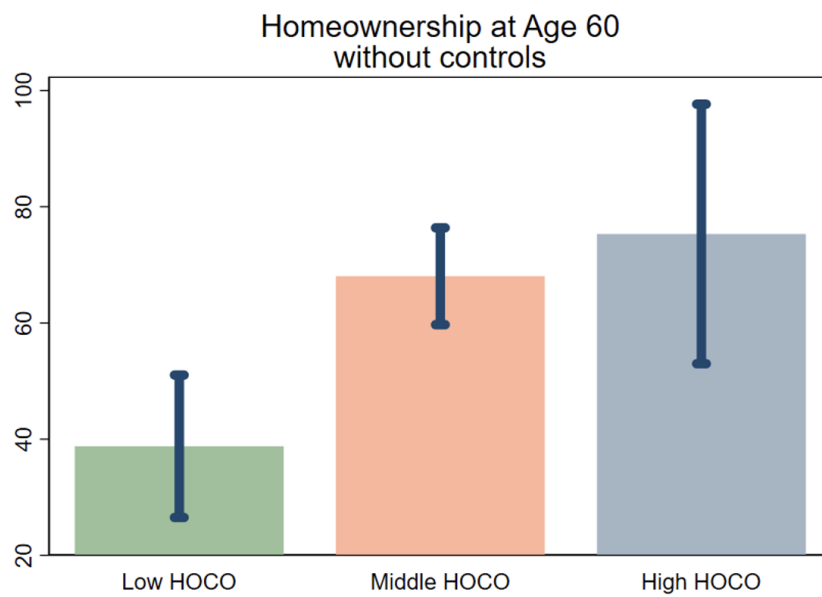


(A) Delinquencies by HOCO

(B) Delinquencies by HOCO, period Bought

Notes: Binscatter using data from the 2013 and 2017 waves of the AHS. Sample covers homeowners with foreign born heads. Controls include household head education, race, ethnicity, marital status, number of family members and children in the home, and quadratics in household income, years in the US of household head, and age of household head, as in prior specifications and add debt-to-income and cost-to-income ratios. We include year-by-CBSA fixed effects. Heteroskedasticity-robust standard errors clustered at the commuting zone level shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Figure BVII: Homeownership by HOCO Tercile at Retirement Age



Notes: This figure shows the average homeownership in percentage points among foreign-born respondents in the Health and Retirement Study with a homeownership in their country of origin that is in the lowest, middle, or highest tercile of homeownership in the country of origin (*HOCO*). The graph and the assignment into terciles exclude Mexican origin foreign-born and residents where we were unable to assign a *HOCO* score. The homeownership rates are adjusted for the effect of the same control variables as included in column 2 of Table VII.