LAND USE REGULATION AND INDIVIDUAL WELFARE

BRIAN BEACH

Department of Economics
College of William & Mary
and NBER
bbbeach@wm.edu

TATE TWINAM

Interdisciplinary Arts and Sciences
University of Washington Bothell
twinam@uw.edu

This draft: November 12, 2017

ABSTRACT. We provide some of the most disaggregated estimates of the welfare effects of land use regulation. To do so, we link individuals between the 1920 and 1940 censuses in order to analyze how households responded to the introduction of Chicago's comprehensive zoning ordinance. Drawing on pre-zoning demographic and land use microdata, we construct a plausibly exogenous household-level measure of zoning mismatch. We find that zoning played a meaningful role in shaping the demographic composition of the neighborhood, with the largest effects appearing in black neighborhoods. Movers (native, black, and foreign) were able to completely offset the changes brought about by zoning. While commercial and manufacturing zoning lowered overall neighborhood quality for blacks, blacks that stayed behind benefited from increased job access.

1. Introduction

It is hard to find a local governmental policy more ubiquitous and controversial than zoning. Although zoning has the potential to increase efficiency (e.g., by reducing negative externalities), it can also undermine efficiency by imposing large costs of developers and by preventing market forces from ordering economic activity in an efficient manner. In this respect, Turner, Haughwout and Van Der Klaauw (2014) analyze the impact of municipal land use regulations on land values and show that land use regulations tend to decrease social welfare.[1] In terms of equity effects, critics also argue that zoning reinforces segregation and that, by reorganizing the distribution of uses across neighborhoods, low-income households have been disproportionately harmed by zoning (e.g., Rothwell and Massey (2009), Rothwell (2011), and Shertzer, Twinam and Walsh (2016a).

While the bulk of existing literature has focused on aggregate city or neighborhood-level outcomes (e.g., new housing starts, home values, etc.), much less is known about how individuals respond to changes in zoning or the extent to which zoning affects individual welfare. Because of sorting, these results do not immediately follow from existing literature. Zoning affects the spatial distribution of housing, employment opportunities, and other neighborhood amenities, but utility-maximizing households face a trade-off between housing costs and neighborhood amenities. Thus, to the extent that zoning's effects on neighborhood development and access to nearby amenities are capitalized into rents and housing values, households that value those gains the least will have the strongest incentive to relocate.[2]

[1]Turner et al. (2014) are, of course, not the first to think about the net effect on land use regulation on welfare, but they offer a substantial methodological improvement over the existing literature. See also, Mayer and Somerville (2000), Ihlanfeldt (2007), Glaeser and Ward (2009), and Zhou, McMillen and McDonald (2008). More recently, Hsieh and Moretti (2017) show that density restrictions also impose costs on the broader economy as workers cannot move to high productivity cities; this misallocation of labor in turn lowers aggregate production. While not about regulation per se, Hornbeck and Keniston (2017) and Siodla (2017) both show that barriers to redevelopment have large aggregate welfare effects.

[2]See Kuminoff, Smith and Timmins (2013) for an overview of these types of models

In light of this, whether zoning has a meaningful effect on individual outcomes is unclear. This is, of course, a difficult question to answer as most cities have zoning restrictions in place today and modern changes to zoning are highly endogenous. Setting identification issues aside, one further complication is that privacy concerns limit our ability to track individuals over time. For this reason, the bulk of the existing literature on zoning has focused on neighborhood or city-level outcomes with land or home values being the primary metric for measuring welfare. In this paper, we overcome these limitations by analyzing the introduction of zoning. By considering a historical setting we are able to link publicly available Census records to track individuals over time. A second advantage of considering the introduction of zoning, which we discuss in greater detail below, is that it allows us to adopt a novel identification strategy to recover the causal effect of zoning. Together, this allows us to examine whether zoning meaningfully affects the type of neighborhood that an individual resides as well as the extent to which zoning affects homeownership and employment.

The setting we consider is Chicago between the years 1920 and 1940. Chicago introduced one of the nation's first comprehensive zoning ordinances in 1923. This ordinance imposed land use and density regulations for all areas of the city, setting in motion a pattern of development that persists until today (Shertzer, Twinam and Walsh (2016b)). By considering the introduction of zoning, we are able to observe the distribution of households and land uses in the absence of zoning. This is important, as it allows us to construct a measure of zoning mismatch, which we define as the difference between actual neighborhood characteristics and the pattern of development envisioned by the zoning commission. With this measure in hand, we exploit zoning mismatch as a shock to neighborhood development and examine the extent to which this affects individual outcomes.

Our empirical approach is as follows: we begin by linking males aged 10 to 40 between the 1920 and 1940 censuses. We then geocode Chicago households in the

1910, 1920, and 1940 censuses, which allows us to compute (at the individual level) detailed neighborhood characteristics by identifying all households that fall within a half-mile radius of each census record. We also draw on a comprehensive land use survey conducted in 1921 and detailed historical crime reports to identify each household's proximity to pre-existing land uses and crime.[3] Important for our analysis is that we observe the pre-existing spatial distribution of residents, land use, and crime before the introduction of zoning. We then attempt to recover the zoning commission's decision-making process by regressing actual zoning on a flexible model that incorporates each of these neighborhood-level variables.

The residuals from these regressions offer a measure of zoning mismatch, which we exploit as a source of identification. These residuals reflect the feasibility constraints facing the zoning commission at the time. In their desire to segregate land uses, the commission constructed use and density districts that would span multiple city blocks, and so some blocks would ultimately be incorporated into a use district that did not match current built environment (e.g., a residential block near manufacturing might be zoned as manufacturing, or vise versa). In a placebo test we show that these residuals are not systematically related to neighborhood characteristics in 1910. We also show that our process for generating these residuals does equally well for neighborhoods where black, foreign born, and native residents lived. In contrast, we show that existing empirical approaches, which tend to simply control for confounding factors, are significantly related to 1910 demographic characteristics.

Results indicate that zoning played an important role in shaping neighborhood development. To highlight one example, a neighborhood that received excess manufacturing zoning in 1920 had a smaller population, fewer white households, more low-income workers, and lower average rents in 1940. This finding matches results in the existing literature and lends support for our underlying assumption that zoning mismatch

[3]The land use survey was recently digitized by Shertzer et al. (2016a).

provides a shock to neighborhood development - or, at a minimum, that individuals responded as if zoning would eventually affect neighborhood development.

Turning to individual outcomes, however, we find that movers completely offset the impact of zoning by choosing new neighborhoods that better matched their preferences. The impact of zoning on neighborhood development appears to be most pronounced for neighborhoods where blacks lived in 1920. Black residents that didn't move saw the largest changes in demographic composition. Relative to apartment & residential zoning, a one standard deviation increase in excess commercial zoning would reduce the white share by 5 percentage points and increase the neighborhood population by 1,766 residents. Blacks that moved, however, were still able to largely offset the effects of zoning on neighborhood change. Turning to individual outcomes, we see that blacks that did not move were rewarded with greater homeownership opportunities and greater labor force participation.

These results are consistent with the fact that, while our measure of excess zoning represents an exogenous shock to zoning, how firms respond to zoning remains endogenous. If commercial and manufacturing firms were disproportionately drawn to black neighborhoods, perhaps because the cost of redevelopment were lower, then this would explain why the impact of zoning is more pronounced for blacks that did not move.

While the primary goal of this paper is to understand how zoning affects individual welfare, the results in this paper also relate to the larger literature on environmental justice. As in the zoning literature, the environmental justice literature has been forced to examine the impact of changes in environmental amenities on neighborhood-level outcomes because of lack of individual-level data. That literature tends to find a large sorting response to changes in amenities, as measured by changes in neighborhood demographic composition and house prices.[4] The literature tends to conclude that

[4]See, for instance, Banzhaf and Walsh (2008), who examine how changes in air quality affect neighborhood composition in California. Depro, Timmins and O'Neil (2015) build a structural model to argue that the resorting associated with an increase in air pollution is driven both by white flight as well as in-migration of minority groups who are less willing to trade consumption to avoid exposure to

environmental justice policies are likely to be ineffective because of sorting. Our results highlight the fact that not everyone moves, and that those who are left behind end up bearing the burdens (and gains) from changes in neighborhood amenities. This suggests that policymakers may want to pay particular attention to an individual's ability to move when crafting policy to address environmental justice issues.

2. Land use regulation and neighborhood change

In 1916, New York City became the first US city to pass a comprehensive zoning ordinance, but by 1930 nearly 500 American cities had zoning ordinances. This surge in popularity was largely driven by a desire to combat many of the negative externalities associated with urban living. As McMillen and McDonald (1999) show in their analysis of Chicago, a frequent and unfortunate result of unregulated land use was the placement of manufacturing and commercial uses on or near residential blocks. Noxious uses were of course only part of the problem. Failure to regulate density often resulted in overcrowding. Given the infectious disease environment of the time, overcrowding was a particularly salient problem in major cities, as without universal water and sewerage systems, increases in population density were often accompanied by an increased incidence of disease (e.g., Troesken (2004), Ferrie and Troesken (2008), and Alsan and Goldin (2015)).

Chicago's experience with comprehensive zoning began in earnest in 1921 with the creation of the Chicago Zoning Commission.[5] Calls for regulation, however, occurred much earlier. As stated in the Zoning Commission's 1922 pamphlet on the zoning

pollution. The literature aimed at assessing the incidence of changes in neighborhood amenities typically focuses on housing values because they don't have individual-level data, an issue we overcome in this paper. Bento, Freedman and Lang (2015) show that the incentives introduced by the Clean Air Act Amendments disproportionately benefited homeowners in low income areas and Grainger (2012) shows that the pass-through of these benefits to renters was on the order of 50%. More recently, Sullivan (2016) offers a number of improvements in measuring exposure to air pollution and concludes that the benefits associated with improvements in air quality likely favor higher-income households.

[5]The City of Chicago had made earlier attempts to regulate undesirable uses, but those efforts were ultimately insufficient at controlling development, Shertzer et al. (2016b).

ordinance: "*Zoning [was] the culmination of a ten-year campaign on the part of real estate boards, homeowners, civic organizations, improvement societies and clubs.*" The primary goal of these interest groups was, of course, to preserve their own land and home values, and the zoning commission was clearly established with this goal in mind. Evidence in support of this appears on the title page of the widely circulated zoning pamphlet, which states:

> *"This ordinance will protect private residence blocks against apartments and stores; apartment house blocks against stores and public garages; shopping streets against warehouses, public garages and laundries; commercial and light manufacturing streets against offensive industries; and manufacturing against fear of molestation."*

The zoning commission conducted its task in a highly technical and transparent fashion. Immediately following its creation, the commission spent 18 months surveying the existing land use for every block in Chicago – data we ultimately use for our analysis. Surveyors collected information on current uses, building heights, setbacks, and age of buildings. Additionally, the surveyors divided the city into 0.25 mile non-overlapping blocks and then calculated the share of each block being used for residences, apartment, commercial, or manufacturing, as well as the share of each block that was vacant. The commission also tabulated the number of families residing in each 0.25 mile block. Particular attention was paid to the "objectionableness" of each use. Pre-existing land uses played an important role in determining zoning, but the commission also held public meetings and solicited feedback from civic organizations. The creation and circulation of the commission's 1922 pamphlet speaks to the goal of transparency. That pamphlet outlined: the goal of the commission, what zoning means, the legality of zoning (including the fact that it is not retroactive), how districts will be changed, how zoning has worked in other cities, and the process with which Chicago is conducting its own zoning ordinance.

Chicago's 1923 zoning ordinance was designed to regulate building density as well as the location of different types of uses. Four use districts were created: Residential (single-family housing), apartment, commercial, and manufacturing. The use districts were hierarchical, with areas zoned to allow manufacturing allowing any other use (commercial, apartment, and residential). Commercially zoned areas permitted any type of residential use as well, while the most restrictive zoning (residential) permitted only single-family homes. The density regulations restricted maximum lot coverage, aggregate volume, and height. The five volume districts created are best described by their restrictions on building height. District 1 capped the height of a building at 3 stories, while districts 2-5 allowed up to 8, 11, 16, and 22 stories, respectively. The allowed height within each district was tied to volume and lot coverage, so in some cases height would be sacrificed to allow greater building footprints.

For the 1923 zoning law to induce a sorting response it would have to have impacted either real land use leading up to 1940 or expectations about the development of neighborhoods in the near future. While we currently lack data on land use in 1940 or city residents' expectations of land use change, there has been a considerable amount of research on the effects of zoning laws. Most of this work has focused on zoning's impact on land values and its malleability in the face of market forces. Zoning seems to exert an effect on land values, suggesting the potential for real effects on land use (Ihlanfeldt 2007, Koster, van Ommeren and Rietveld 2012, McMillen and McDonald 2002, Turner et al. 2014). However, there is also evidence that zoning changes endogenously, responding to market forces or political power (McMillen and McDonald 1991, Munneke 2005, Murray and Frijters 2016, Wallace 1988).

Shertzer et al. (2016b) show that Chicago's 1923 ordinance strongly influenced the evolution of land use up to the present day, even in areas of the city that were fully developed at the time. However, it is not clear how quickly those changes materialized, or whether residents expected large changes in the short run. In this respect, some insight can be gained from Twinam's (2017) study of the response to Seattle's 1923 zoning code, which was similar in design to Chicago's. That study found a sizable response to zoning, with roughly half of the (large) long-run effect of the law accruing in the three decades after implementation. While Seattle was less developed than Chicago, this nonetheless suggests that the zoning could have had a substantial real and anticipated impact. Later, we assess whether zoning affected the demographic composition of neighborhoods – a proxy for assessing the impact of zoning – and find that this was indeed the case. Thus, even if the effect of zoning on land use had not been felt by 1940, Chicago residents responded as if zoning would meaningfully affect neighborhood development.

3. Data

The central goal of this paper is to understand the extent to which Chicago's zoning ordinance affected the well being of its residents. More specifically, did the zoning ordinance affect the types of neighborhoods individuals resided in? And to what extent did the ordinance affect homeownership and income? To answer these questions, we combine two broad sets of data: full count census data from 1910, 1920, and 1940 and digitized land use and zoning maps. The next two subsections describe how we use these sources of data to construct our final sample.

3.1. **Linking individuals between the 1920 and 1940 censuses.** Historical US Census records are only considered confidential for the first 72 years after enumeration. Between this rule and the digitization efforts of Ancestry.com, it is now possible to identify the name and address for every census record between 1910 and 1940. We focus our attention on the 1920 and 1940 censuses. The 1920 Census is closest to 1923, the year in which Chicago's introduces its zoning plan. Thus, by drawing on the 1920 full count census data we can identify the spatial distribution of Chicago residents prior to the introduction of zoning. To identify how zoning affected this spatial distribution, we simply need to analyze one of the future censuses (1930 or 1940). We focus on the 1940 census because it is the first census that asks about income and education, which will prove useful for trying to identify the quality of the neighborhood. Further, to the extent that individuals may not fully comprehend the impact of these new ordinances until they observe the construction of new businesses and housing projects that are forced to comply with those restrictions, extending our time horizon from 1930 to 1940 alleviates concerns that there may simply not have been enough time for individuals to fully respond to these ordinances.

We use the full count census data in two ways. First, we use the data to identify the spatial distribution of Chicago's residents. Second, we rely on the information contained in both datasets to track individuals over time. One issue that hinders our ability to track individuals is that the census does not contain a unique identifier – individuals have never been asked to report their social security number, and even if they were, social security numbers did not exist in 1920. To overcome this barrier, we link individuals between the 1920 and 1940 censuses by exploiting the fact that many individuals are uniquely identified by their first name, last name, place of birth, and year of birth.

Our linking procedure closely follows that of Feigenbaum (2016). We begin by identifying every male residing in Chicago that was between the ages of 10 and 40 in 1920. As is standard in the literature, we focus on males because women tend to change their last name once they are married, which decreases the likelihood of finding a valid match. We limit our analysis to those between the ages of 10 and 40 because we want to identify individuals that will be old enough to have control over their location decision but not so old that they are likely to die before enumeration in 1940.[6] For each of the 751,858 records meeting these three criteria, we begin by standardizing

[6]While individuals under the age of 20 likely had little say in location decisions in 1920, they should have considerable say by 1930. Further, to the extent that zoning may have affected family outcomes (e.g., income, employment opportunities, or homeownership), it is useful to link younger individuals to assess whether those changes carried through to the next generation.

first names. For instance, Bill, Will, and Billy are all recoded as William. Further, because information was spoken to the enumerator we standardize spelling variations (e.g., Eliot and Elliot are both recoded as Elliott).

Next we attempt to identify the relevant record in 1940. To do so, we first restrict our sample to individuals with a unique standardized first name, last name, birthplace, and age. 726,273 of our records are unique across these four dimensions.[7] We first compare each these records against every 1940 Census record in Illinois to identify a set of plausible matches.[8] A plausible match in 1940 must meet each of the following criteria. First, the reported birth state or country must be the same. Second, the birth year inferred from the 1940 census must fall within three years (plus or minus) of the birth year inferred from the 1920 census.[9] Finally, the first and last name in 1940 must be reasonably close to the first and last name in 1920. Because enumerators recorded information that was spoken to them, we do not want to rule out a potential match because of a slight mispelling (e.g., Anderson instead of Andersen). As our measure of closeness, we follow Feigenbaum (2016) and require the Jaro-Winkler string distances between both the standardized first names and the last names to be less than 0.20.[10]

[7]For those without unique criteria, there are 9,369 combinations that appear exactly twice, 1,346 that appear three times, and 400 combinations that appear four times. The most frequent combination (Arthur Johnson, aged 26, and born in Illinois) appears 13 times.

[8]We consider records outside of the city of Chicago because individuals may have moved outside of the city between 1920 and 1940. Individuals can, of course, move out of state as well but extending our analysis to consider every US state is computationally intensive. Further, increasing the number of states that we consider decreases the likelihood that any name, birthplace, age pairing is unique, which decreases our successful match rate. While, only restricting to Illinois does introduce the possibility of a false positive (identifying a successful link when in fact the true record is located outside of Illinois) we feel that the increased likelihood of a false positive is small. Furthermore, because we are ultimately interested in the differential effect of zoning, these false positives are only a concern to the extent that they are systematically related to our treatment variable.

[9]The census asks about an individual's age but not their birth year, which means that birth year has to be inferred based on reported age. Because the incentive to obtain a birth certificate was much smaller in the early 20th century, individuals may misremember their true age. This issue is potentially compounded by the fact that the information is not always reported by the individual (e.g., the information might be reported by the individual's mother in 1920 and his wife in 1940). Because of these issues, we allow for a potential match to have a slightly different age.

[10]Technically we use 1-Jaro-Winkler distance, as two identical strings will have a Jaro-Winkler distance of zero. One common alternative to identifying similar strings is to use the Soundex score, which is a phonetic algorithm that classifies strings based on how they are spoken in English. Because names

Under these criteria, 536,650 of our 726,273 records have at least one plausible match in 1940. While the median record has five potential matches, 119,275 (or 22%) of our records have only one potential match. At the other end of the distribution, 25% of the sample has more 16 matches with one record having a maximum of 418 potential matches. These numbers can be large because we relax restrictions on both age and the spelling of the name to increase the likelihood that the true 1940 record is contained within our set of plausible matches. Of course, now the challenging task is to accurately link the 1920 and 1940 records. To do so, we again follow the suggestions of Feigenbaum (2016) and adopt a machine learning approach. Specifically, we begin by taking a random 0.5% sample of the 1920 records and hand match them against the set of potential matches in 1940. For this sample, we were able to successfully link 1438 of the 2562 records from 1920 to a record in 1940. Unsuccessful links occur because either 1) there was not a 1940 record with a first and last name that was close enough to be considered a convincing match or 2) there were too many records with a first and last name that could be considered as a convincing match and thus it was not possible to confidently identify the true link.

After hand matching these records we run a probit regression to generate predicted links. Here we try to model our probit regression to encompass the criteria that allowed us to determine whether a record was a valid link or not. First we include a series of indicators for whether the absolute value of the difference between the inferred birth years is 0, 1, 2, or 3 years. We then include an indicator for whether the first and last names are an exact match and another indicator for whether the first and last name plus the birth year is an exact match. For those born outside of the state of Illinois, birthplace provides a bit more context for the match, and so we fully interact an indicator for migrant status (i.e. born outside of Illinois) with each of our birth

were spoken to the enumerator, Soundex scores an appealing alternative. However, those names were ultimately written (often in cursive) and then transcribed years later. Small transcription errors are not accounted for with Soundex, and so for that reason we focus on Jaro-Winkler scores.

year difference indicators as well as the exact name match indicators. Next we include the Jaro-Winkler string distance between both the first names and the last names as well as an indicator for whether the last names sound the same, as measured by Soundex.[11] Because migrants, particularly those born outside of the United States, likely pronounce their name with an accent we interact this variable with our migrant indicator. Sometimes an individual reports their middle initial or middle name, and so when both records have a middle initial we include an indicator that equals one if those initials match. To help overcome transcription and spelling errors, we include a series of indicators for whether the first and last initials of both the first and last names match. Finally, to accommodate the fact that a successful link is less likely to occur when there are multiple relevant matches we include the number of potential 1940 records that: are an exact name match, are an exact name and birth year match, have the exact same last name Soundex score, and the square of the number of potential records with the same last name Soundex score.

Using the parameters estimated from running this probit on our hand-matched sample, we generate predicted probit scores for every potential link. This allows us to rank potential matches. Of course, we don't want to simply take the record with the highest predicted link probability. This is because there can be multiple matches that are relatively close (e.g., Elliot, Eliot, and Elliott all with the exact same last name and birth year). Thus, we only identify a record as an acceptable link if it meets two criteria. First, the link must be a "good" link, which we measure as having a predicted link probability that is sufficiently high. Second, the link must be distinct, which we measure as having a predicted link probability that is sufficiently far away from the

[11]Soundex is a phonetic algorithm that indexes names based on how they are spoken in English. Ashcraft and Ashcroft, for instance, both receive the same Soundex index of A261.

predicted probability of the second best match. We define this as the simple difference between the first and second best links.[12]

We perform a grid search on our training sample to choose the correct values for these two criteria. Recall that our training sample, which was hand matched, has both a predicted match and a distance to the next best match. Predicted matches are constrained between zero and one, as is the difference between the first and second best match.[13] Our grid search loops over all possible threshold combinations and then computes two statistics. The first statistic is the share of true positives, identified as the "efficiency" statistic in the terminology of Feigenbaum (2016). The second statistic is the "accuracy" statistic, which is the share of positive links identified that are true positives. In an ideal world, there would be no false positives and this statistic would equal one, as the algorithm would only identify true positives. Obviously there is a trade-off between efficiency and accuracy; however, because we are starting with such a large sample we decide to weight the accuracy statistic twice as much as efficiency in order to minimize our false positive rate. After performing this grid search, we are left with a final sample of 202,601 unique links between 1920 and 1940.[14]

Before turning to summary statistics, it is worth attempting to quantify the prevalence of false positives in our final dataset. We do this by examining the extent to which our 1920 and 1940 records match on variables that should not change over time but were also not used for linking. One such variable is race. While we never imposed that the races of the two records match, it turns out that 98.2% of our records have a

[12]Feigenbaum (2016) recommends using the ratio of the first and second best predicted probabilities. We have used both, but have found that that the simple difference generates fewer false positives. Our procedure for identifying false positives is discussed below.

[13]For records with only one plausible match, we define the difference as one.

[14]It is theoretically possible for our algorithm to match multiple 1920 records to the same 1940 record. Consider, for instance, two John Smiths born one year apart. If the older John Smith died before enumeration in 1940, this would be unobservable to us, and so our algorithm might assign both the older and younger John Smith to the same 1940 record. This, by construction, adds false positives to the dataset and so, in these instances, we omit both observations. This removes 19,101 potential 1920 links, which leaves us with our final sample of 202,601 records.

matching race. Of course, given that we are analyzing a majority white city, perhaps this is not too surprising. One other set of variables that can be analyzed are mother's birthplace and father's birthplace. While the 1920 census asked everyone where their mother and father were born, this question was only asked for 5% of the population in 1940. 17,422 of our linked individuals (or about 8%) were asked this question. We therefore define a link as a false positive if their mother's and father's birthplaces are inconsistently reported. It is, of course, possible that both mother and the father's birthplace match and the link is still a false positive. It is also possible for a link to be valid even if the parental birthplaces do not match. This issue reflects that some misreporting is to be expected, perhaps because the information comes from two different sources (e.g., mother or father in 1920 but spouse in 1940). Nevertheless, for our sample 18% of records have parental birthplaces that are inconsistently reported.

Upon closer inspection, we identified two broad patterns with these mismatches. The first pattern has to do with aggregation. Example mismatches include birthplaces of "Illinois" and "Indiana" or "Germany" and "Poland". In these types of situations, it seems likely that our identified link is still valid. One way to account for these issues is to recode birthplaces at a higher level of aggregation. For instance, we could recode each birth state as its census division and we could assign immigrants the broader birthplace codes that are available on IPUMS.org.[15] After aggregating birthplaces, our set of potential false positives falls to 12%.

This brings us to the second broad pattern, which is for immigrant parents (as reported in 1920) to be reported as being born in Illinois in 1940. 880 of our remaining 2,116 "false positives" reflect this pattern. To put this number in perspective, only 287 of our potential false positives report being born in Illinois in 1920 and report being born abroad in 1940. One can tell a number of stories in which these mismatches are

[15]These categories are: US Territories, Other North America, Central America and Caribbean, Northern Europe, UK and Ireland, Western Europe, Southern Europe, Central/Eastern Europe, Russian Empire, East Asia, Southeast Asia, India/Southwest Asia, Middle East/Asia Minor, Africa, and Oceania.

justifiable: immigrant parents may have felt an incentive to report to census officials (i.e., representatives of the federal government) that they were indeed a citizen of the United States. Alternatively, a child (or the child's spouse) of an immigrant who migrated when they were very young may be particularly likely to misremember whether that parent was born in Illinois or simply spent most of their childhood there.

If we ignore these distinct immigrant-related issues – so treat those mismatches as missing rather than as an incorrect or correct match – our potential false positive rate falls to just under 6%. To put this in perspective, Bailey, Cole, Henderson and Massey (2017) estimate that most automated linking procedures have a false positive rate as high as 30%. Using hand-matched samples, Bailey et al. (2017) estimates that their own false positive rate is likely on the order of 1%. Of course, hand matching is not feasible in most settings as it is both slow and costly. Our ability to link approximately 37% of our initial sample with such a low false positive rate is likely attributable to the fact that, unlike a majority of the linking literature, we focus on adolescent to working age males (as opposed to young children) over a relatively short time horizon (20 years as opposed to 40 or more). Because false positives should attenuate results by introducing measurement error, in a robustness check we will adopt stricter thresholds in order to further limit the presence of false positives.[16]

3.2. **Land use, zoning, and social indicators data.** With our linked sample in hand, we draw on the 1910, 1920 and 1940 full count data to identify, for each linked individual, their neighborhood composition in 1920 and 1940 as well as demographic trends between 1910 and 1920.[17] The census data provide us with a host of demographic variables, which we then supplement with land use data from 1922 and data

[16]To see this more clearly, note that false positives will add 1940 observations that are, by definition, independent of our zoning treatment. If our sample was entirely made up of false positives then this would lead us to believe that zoning had no effect on 1940 location decisions and outcomes.

[17]Chicago renamed and renumbered its streets in 1909, which limits our ability to geocode data from the 1900 census.

on the presence of homicides and gang activity to account for the pre-existing built environment and overall community health.

As mentioned earlier, the zoning commission conducted a citywide land use survey prior to the drafting of its zoning ordinance. From that survey, Shertzer et al. (2016a) geocoded all commercial establishments, warehouses, and manufacturing uses (in five classes) in the city, as well as building heights for all buildings with four or more stories. Appendix Figure 1 illustrates a portion of the resulting survey map, while Appendix Figure 2 shows the geocoded survey for the area around Chicago's central business district, the Loop. There were ultimately 33,622 commercial uses, 9,022 manufacturing uses, and 5,715 buildings over 3 stories tall. The density of commercial and manufacturing uses are depicted in the top panels of Figure 1. Both commercial and manufacturing uses appear to have been concentrated in and around the central business district as well as the Chicago river. Commercial uses are less concentrated and ubiquitous throughout the portion of the city that was substantially developed at that time.

The zoning ordinance used the dual-map system that was standard at the time; Appendix Figure 3 shows samples of the use and density maps. The ordinance established four use districts (Residential (single-family housing), apartment, commercial, and manufacturing) and five volume districts. The digitized districts appear in the bottom panels of Figure 1. There we see that the high volume districts (districts 4 and 5) were reserved for the central business district while manufacturing uses were pushed away from the central business district.

Figure 1. 1922 Land Use Density and 1923 Zoning Extent

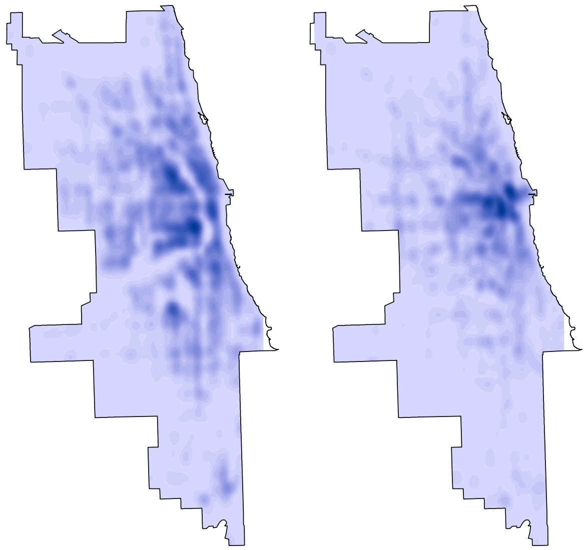

(A) Commercial uses (B) Manufacturing uses

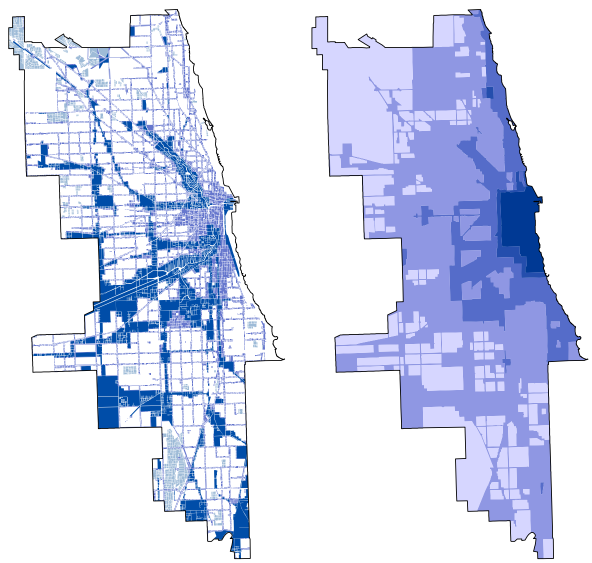

(C) Use zoning (D) Volume zoning

Panel (A) shows the density of commercial uses in 1922. Panel (B) shows the density of manufacturing uses in 1922. In Panels (A) and (B), darkness is proportional to use intensity. Panel (C) maps 1923 use zoning districts: The lightest blue areas were zoned for residences, white areas were zoned for apartments, medium blue areas were zoned commercially, and the darkest blue areas were zoned for manufacturing. Panel (D) maps 1923 volume zoning districts: The lightest areas were zoned for the lowest density (district 1), while the darkest area was zoned for the highest density (district 5).

We draw on two final source of data to better capture neighborhood quality. First, we use data from Chicago's Historical Homicide Project, which digitized the Chicago Police Department's records for its nearly 11,000 homicide cases occurring between 1870 and 1930. Because many of these records included the address of the crime, we were able to successfully geocode 5,270 of these homicide cases. Of the geocoded cases, 4,290 are dated between 1910 and 1930. To provide a more complete picture of the crime environment, we supplement these data with data from Frederic Thrasher's study of Chicago street gangs over the 1923-1926 period Thrasher (1927). Thrasher recorded the locations of 1,313 gangs on a map, which we then geo-referenced to obtain accurate latitude and longitude information. While not all of these gangs engaged in criminal activity, many did. In support of this, the distribution of gangs closely matches that of other crime proxies as well as data from the historical record on the location of criminal activity in this era (Shaw, Cottrell, McKay and Zorbaugh 1929).

3.3. **Neighborhood characteristics and summary statistics.** The full count census data contain two other variables that are essential for our analysis: house number and street address. This allows us to geocode households, and in turn, identify preexisting neighborhood characteristics. Geocoding also allows us to identify the extent to which these ordinances influenced individual location decisions, affecting broader neighborhood change.

There were 653,422 unique households in Chicago's 1920 census, 95% of which report a street address and 62% which report both a house number and a street address. We begin by geocoding each of these 401,253 addresses. Three issues affect our ability to geocode. First, some streets from 1920 no longer exist today. A 14 mile stretch of Grand Boulevard, for instance, was renamed Martin Luther King Jr Drive in 1968. Second, the transcribed street names match exactly what the enumerator wrote, and so slight misspellings or transcription errors will result in an unmatched address. Finally, many transcribers wrote the street address but did not specify the direction of the address

(North, South, East, or West). We deal with these issues by analyzing and correcting all rejected street names that affect more than 150 households. Steve Morse's street name change database reports all street names and, if applicable, how those street names have evolved over time. This allows us to identify street name changes. The street name change database also allows us to correct for misspellings. Because Chicago's street system is on a grid and Chicago's enumeration districts rarely span the main division streets, we are able to assign all streets a direction based on their enumeration district.

The above procedure allows us to geocode 334,982 of the 401,253 addresses with complete street and house number information. Next, we rely on a feature that is well known to economic historians, which is that households that appear next to each other on the census manuscripts are highly likely to be neighbors. This is because enumerators collected information by walking from house to house, which would be done most efficiently if an enumerator did not skip households while walking down a given street (see, for instance, Logan and Parman (2017) and Grigoryeva and Ruef (2015)). We exploit this feature by organizing households by reel, page, and line number, and then interpolating between accurately geocoded addresses, so long as both geocoded addresses fall on the same page. This procedure allows us to recover an additional 144,262 households, bringing our geocoded match rate to 73.3%. Repeating this procedure for the 1910 and 1940 census allows us to geocode 64% of the 467,265 household heads in 1910 and 71% of the 963,898 household heads in 1940.

Next, we attach a number of neighborhood characteristics. We define a neighborhood as the half-mile radius surrounding each household head. The full count census data allows us to gather demographic information. Specifically we calculate the total number of: households, family households, number of children in the household, number of foreign born households (tabulated by northern-, southern-, eastern-, and western Europe, as well as Russia, and other), white and black households (where southern and northern blacks are separately identified), number of households that owned their own home, households that are in the labor force, households that are in professional, white collar, crafts, service, or labor occupations, and the average income (as measured by occupational income). Turning to our land use and crime data, we calculate exposure to homicides, gangs, commercial and manufacturing uses, and buildings over four stories. Finally we use our zoning data to identify the share of each half-mile buffer that is zoned for: residential, apartment, commercial, or manufacturing uses as well as density districts 1 through 5. When calculating these shares we remove the proportion of each band that falls into Lake Michigan. We also calculate each household's distance to the nearest railroad, distance the the central business district, distance to the Chicago river, distance to Lake Michigan, and distance to the city limits.

Because our geocoded data are incomplete, we are, unfortunately, not able to use all of our linked observations. 50,221 of our initial 202,601 linked records do not have a geocoded address for 1920 and so we are forced to drop them because we are unable to calculate their exposure to zoning. We lose an additional 21,283 records because they resided in an area not subject to the zoning ordinance.[18] Table 1 reports summary statistics for our final regression sample, unlinked males, and all 1920 Chicago households that were successfully geocoded. Columns 1-4 present individual characteristics. There we see that our linked and unlinked samples appear to be remarkably similar in terms of age, race, and homeownership status. Columns 5-10 present statistics on the demographic composition of the neighborhood in 1920 while columns 11-20 present statistics on neighborhood land use. Our linked and unlinked samples appear to have lived in neighborhoods that were representative of all households in Chicago during this time.

[18]While the ordinance was comprehensive, the census and the city of Chicago may have had different definitions of the proper city, which is why it is possible for an individual to have a geocoded address but not be subject to the zoning ordinance

Table 1. Summary statistics

	Individual characteristics				Neighborhood demographics					
	Birth year (1)	Foreign born (2)	Black (3)	Family owns home (4)	Pop. (5)	White share (6)	Foreign share (7)	Share homeowners (8)	Share in lab force (9)	Avg. Occ. income (10)
Linked Sample	1895.23	0.209	0.034	0.333	15,978.8	0.966	0.516	0.308	0.882	33.651
	[8.772]	[0.406]	[0.181]	[0.471]	[6537.274]	[0.103]	[0.178]	[0.152]	[0.025]	[2.551]
Unlinked Males	1895.28	0.221	0.024	0.329	16,214.09	0.968	0.541	0.299	0.883	33.389
	[8.750]	[0.415]	[0.152]	[0.469]	[6552.416]	[0.097]	[0.187]	[0148]	[0.025]	[2.608]
All Chicago Households					16,217.65	0.962	0.520	0.293	0.881	33.576
					[6485.36]	[0.107]	[0.188]	[0150]	[0.025]	[2.719]

	Neighborhood land use									
	Com. uses (11)	Warehouses (12)	Large industrial (13)	Nuisance manuf. (14)	Non-nuisance manuf. (15)	4-8 story buildings (16)	9-11 story buildings (17)	12-16 story buildings (18)	17-22 story buildings (19)	buildings over 22 stories (20)
Linked Sample	283.217	6.79	0.491	2.476	48.345	32.811	0.333	0.267	0.087	0.001
	[189.384]	[12.502]	[1.470]	[4.893]	[52.271]	[66.204]	[1.877]	[2.267]	[1.099]	[0.038]
Unlinked Males	306.137	8.058	0.603	2.969	53.693	38.414	0.428	0.362	0.123	0.003
	[205.482]	[14.549]	[1.712]	[5.402]	[57.348]	[73.807]	[2.309]	[2.621]	[1.296]	[0.528]
All Chicago Households	298.452	7.711	0.553	2.747	51.896	38.006	0.408	0.327	0.106	0.002
	[199.819]	[13.672]	[1.606]	[5.207]	[55.377]	[71.236]	[2.062]	[2.491]	[1.164]	[0.044]

Standard deviations reported in brackets. Linked sample contains 129,130 individuals. The sample of males between the ages of 10-40 that could not be successfully linked to 1940 contains 193,946 observations. There are 485,810 households for the full Chicago sample.

4. Results

4.1. **Empirical approach.** Because zoning was not randomly assigned, it is important to adopt an empirical approach that helps deal with concerns about endogeneity. Our approach draws on the vast amounts of pre-zoning data that we collected and then discussed in Section 3. Specifically, we begin by taking as our left-hand side variable each household's exposure to use and density zoning. We measure exposure as the share of the half-mile radius falling around the household's location in 1920 that was zoned for each use and the share zoned for each density. Recall that there were four types of use zoning (single-family residential, apartments, commercial, and manufacturing) and five density/volume districts, which are best defined by their height restrictions (3 stories, 8 stories, 11 stories, 16 stories, and 25 stories). In part to ease interpretation for our later regressions and also because some use and volume definitions were very similar (with one being used more prominently than the other) we combine like definitions. Specifically, we combine 1) residential and apartment use zoning, 2) volume districts 2 and 3 (the districts with 8 and 11 story height restrictions) and 3) volume districts 4 and 5 (height restrictions of 16 and 25 stories, respectively).

For each of our use and density definitions, we estimate the following equation,

$$Share_i = \alpha + \beta_1 LU_i + \beta_2 Geo_i + \beta_3 Neigh_i + \epsilon_i \,, \tag{4.1}$$

Where $Share_i$ is the share of a household's neighborhood that was zoned for either apartment & residential, commercial, and manufacturing uses, as well as density volume 1, 2 & 3, or 4 & 5. LU_i is a vector of land use controls that includes (for each household's half-mile neighborhood) the following variables as well as their square: number of commercial uses, number of warehouses, number of manufacturing uses (further separated by type a or b (general manufacturing that does not cause a nuisance), large scale industrial (type s), storage (type d), or nuisance manufacturing, which emits noise, smoke, or odor (type c)), and the number of tall buildings, which were categorized to match each density restriction (3-8 stories, 9-11 stories, 12-16 stores, 17-22 stories, and over 23 stories). We also fully interact each of these variables with neighborhood population. The vector Geo_i includes the share of the neighborhood that falls in Lake Michigan, as well as each of the following distance variables and their square: distance to nearest railroad, distance to the central business district, distance to the Chicago river, distance to a major street, and distance to Lake Michigan.

Finally, the vector $Neigh_i$ includes each of the following neighborhood variables (as measured in 1920): total residents as well as the square of the number of residents, average dwelling size, child share of neighborhood population, share of nearby households that were families, southern black share, northern black share, Northern-, Southern-, Western-, and Eastern European share, Russian share, share of households that owned their home, share of households in the labor force, share of household heads employed in professional, white collar, crafts, services, or laborer positions, average income (as measured by occupational income scores) and the standard deviation of neighborhood income (again, measured with occupational income scores). We also include the change (between 1910 and 1920) in population, family share, white share, foreign share, homeowner share, share in the labor force, and average occupational income. These variables account for expectations about neighborhood development. Each of our 1920 race and ethnicity share variables are also fully interacted with neighborhood population.

It is, of course, not surprising that households located close to manufacturing uses are more likely to be zoned to allow manufacturing uses. Indeed, the zoning commission took pre-existing land use into account when they constructed the zoning ordinances, and so residents located near manufacturing uses were likely not surprised to be zoned accordingly. To the extent that zoning was determined by the demographic composition of the neighborhood (Shertzer et al. (2016a)), it may also not be surprising that blacks or immigrants received more density and manufacturing zoning than native whites. By including microdata on land use and neighborhood demographics, however, we can effectively tease out the primary determinants of each zoning type, leaving us with a plausibly exogenous measure of zoning.

The reason we begin by regressing true zoning on all of these predictors is that we want to account for the endogeneity of both neighborhood choice and zoning. By modeling zoning separately we can use data from the entire city of Chicago to better recover the goals of the zoning commission. We then take the residuals from these regressions as a measure of zoning mismatch. Given the extensive controls we are able to include, these residuals reflect the fact that when drawing boundaries for each use and density district, the commission likely faced a trade-off between 1) perfectly matching the pre-existing spatial distribution of residents and land uses and 2) having a zoning ordinance that could be easily communicated because each use and density district corresponded to a continuous collection of blocks. The bottom panels of Figure 1 illustrate that the the zoning districts tended to take the shapes of standard polygons.

After obtaining these residuals, we turn to our linked records. These regressions take individual-level outcomes as our dependent variable and include all of the same right-hand side variables, however, instead of zoning we use the residuals, which again represent the idiosyncratic component of zoning. By running the regressions separately, we are able to cut down on measurement error because each of the right hand side variables is simply trying to capture aspects related to household characteristics and preferences without trying to also capture the endogeneity of zoning.

Because each zoning share is bounded between 0 and 1, we use a fractional logit specification when estimating equation (4.1). Figure 2 plots the distribution of residuals for each use and density zoning regression. Aside from volume 4 & 5 density, which was largely reserved for the central business district, there is reasonable support over the -0.5 to 0.5 range for each type of use and density zoning.

FIGURE 2. Distribution of zoning mismatch

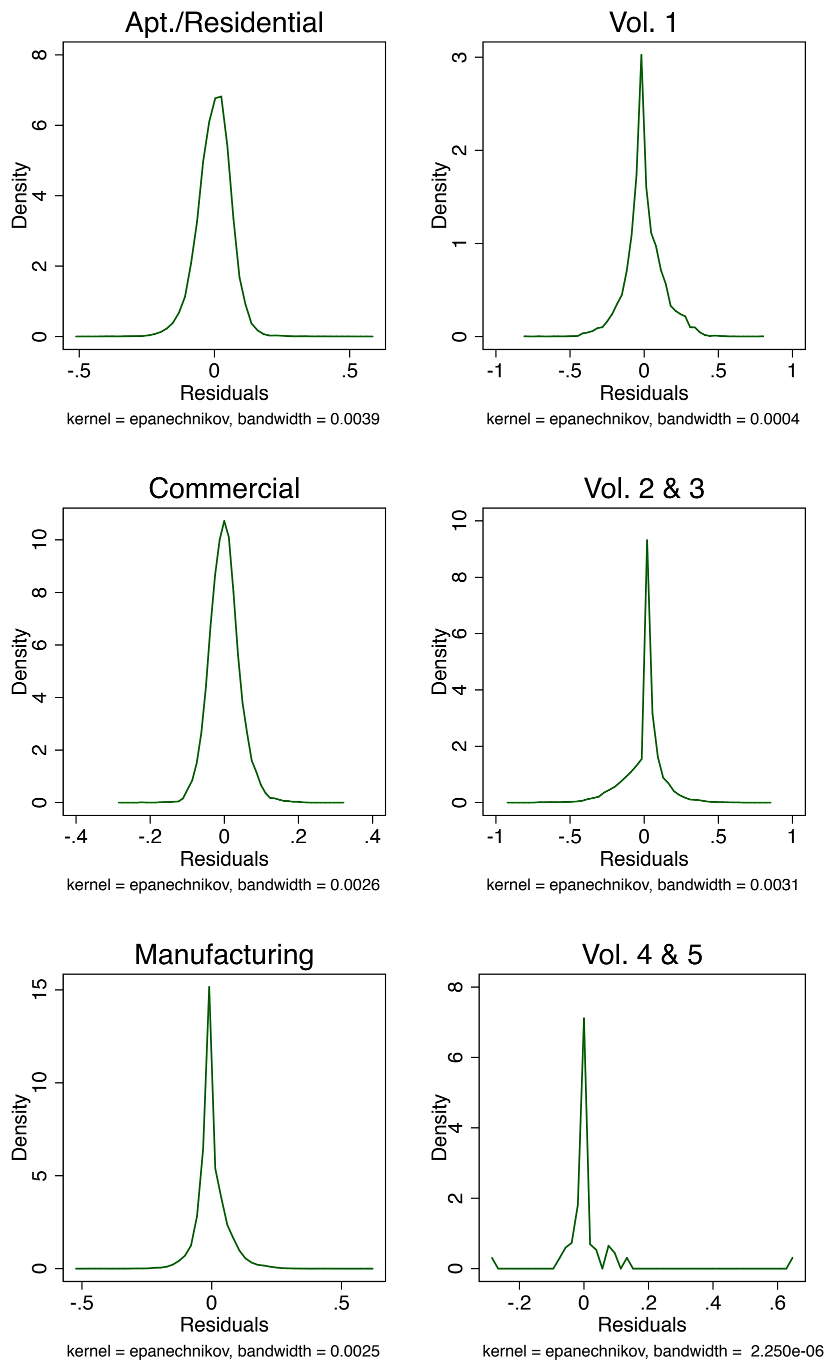

The excess zoning measures are calculated following equation (4.1).

One may be concerned that our predicted zoning model systematically fails for certain sub-groups (e.g., immigrants, and blacks). We assess this in Figure 3 by fitting a local polynomial regression line between predicted and actual zoning (for each use and density zoning) for black, native white, and foreign born households. The top panels of Figure 3 presents results for use zoning while the bottom panels present results for density zoning. With respect to use zoning, our model does a reasonable job at predicting actual zoning for about 99% of observations. In each panel we include a vertical line for either the top or bottom 1%. For low-levels of Apartment/Residential zoning and for high levels of commercial zoning and high levels of manufacturing zoning our model starts to break down. Thus, we truncate our sample by dropping observations that fall in the bottom 1% of predicted Apartment/Residential zoning or the top 1% of either commercial of manufacturing zoning. After truncating, we can see that our model predicts actual zoning equally well for native whites, blacks, and the foreign born, which suggests that our residuals are not systematically related to an individual's ethnicity, which could obviously play an important role in determining neighborhood choice and long-run outcomes.

Another notable feature of Figure 3 is the extent to which predicted zoning and actual zoning tend to align on a 45-degree line. This is important because it suggests that our model is indeed capturing the primary determinants of zoning, and so the residuals are leaving us with plausibly exogenous variation in zoning. We empirically test this assumption below and show that this appears to be the case.

FIGURE 3. Actual vs. predicted zoning by race and ethnicity

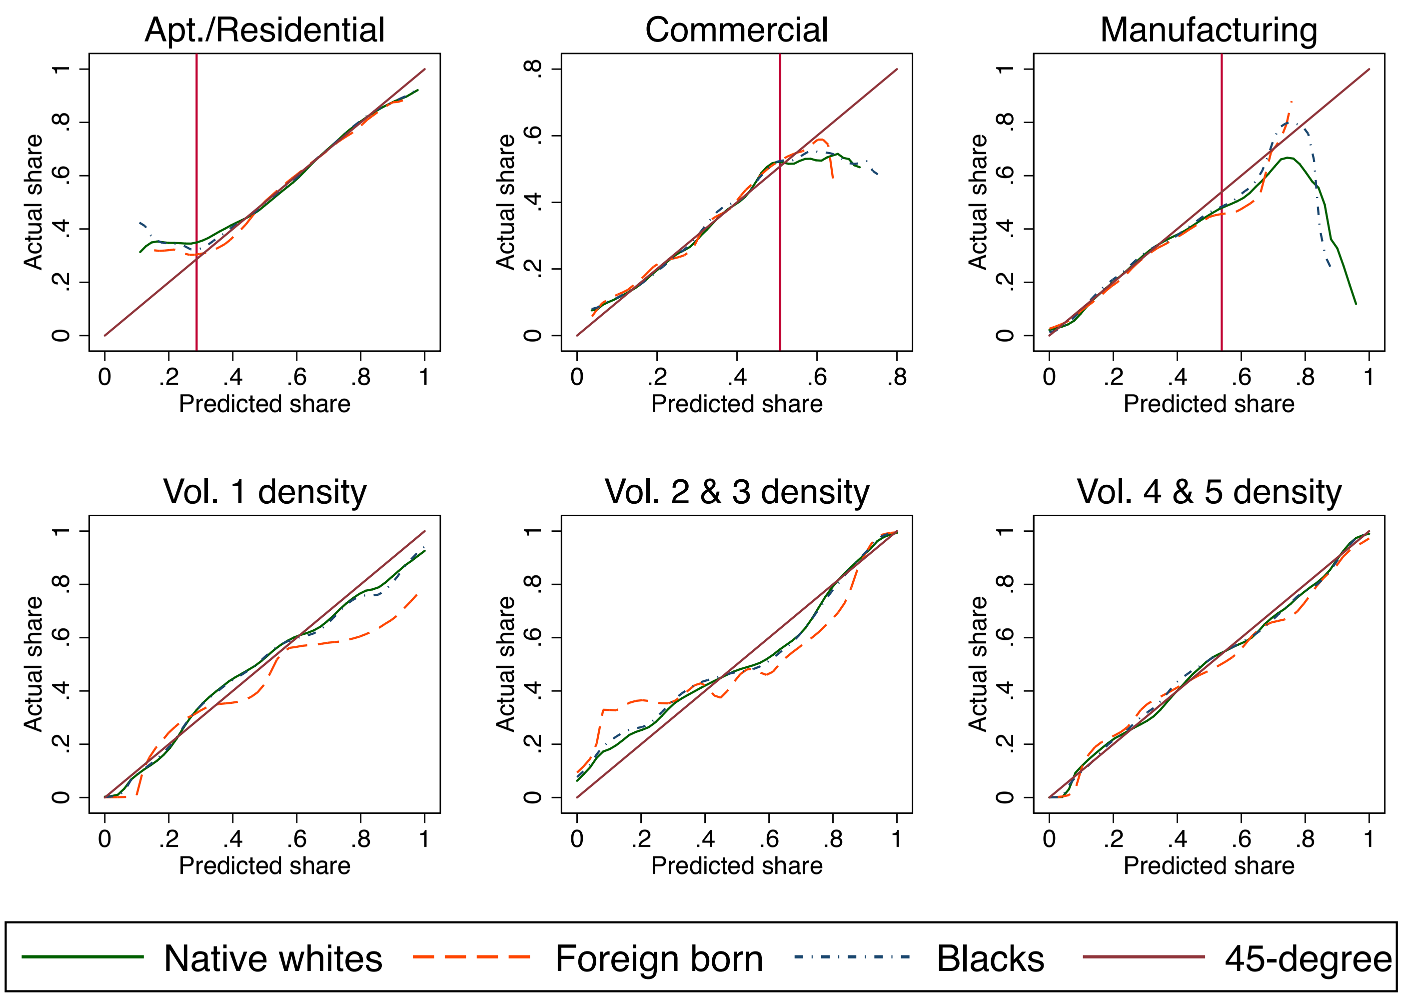

The predicted zoning measures are calculated following equation (4.1). The vertical red lines correspond to either the top or bottom 1% of the data.

4.2. **Did zoning affect neighborhood development?** Before turning to our main results, it is worth assessing the extent to which zoning and zoning mismatch affected neighborhood development. Since we lack land use data from 1940 we are unable to observe whether zoning shaped the built environment between 1920 and 1940. Instead, we draw on our geocoded census data to see whether our measure of zoning mismatch brought about changes in neighborhood composition. If we see any changes, then this provides support for our key assumption that changes in zoning represented – at the very least – a change in expectations about how the neighborhood would develop. We should point out that the existing literature on Tiebout sorting is also limited to assessing how changes in amenities affect neighborhood composition. That literature has consistently found that changes in neighborhood amenities affect the demographic composition of the neighborhood and so it is important that we replicate those findings before turning to individual-level outcomes.

To assess the impact of zoning, we being by taking the XY-coordinates for each of geocoded household in 1920. With those XY-coordinates in hand, we then turn to the 1940 geocoded census data and calculate the following variables based on households that fall within a half-mile radius: total population (as measured by number of members in the household), share of households that were family households, and then based on the identify of household heads we calculate: white population share, foreign born share, share of households that never attended high school, share of household that own their home, share of households in the labor force, average neighborhood income (as measured by occupational income scores), and the average rent in the neighborhood. We then regress each of these outcomes on our zoning shares along with the 1920 land use characteristics, 1920 geographic characteristics, and 1920 neighborhood demographic characteristics that were also used to obtain our measure of zoning mismatch.

The results of this exercise are presented in Table 2. Panel A presents results using the actual zoning shares as our primary explanatory variables of interest, where our omitted use zoning is Apartment & Residential zoning and our omitted density zoning is Volume district 2 & 3 (the 8-11 story height restriction category). Panel B presents results using our indices of zoning mismatch. To ease interpretation, we standardize the zoning and zoning mismatch variables. Many of the results are similar between the two measures, however, we prefer the excess zoning approach as it offers cleaner identification. Notable differences between the two measures relate to the extent to which use and density zoning affected a neighborhood's population, white share, and average rent. Nevertheless, it appears that across almost every outcome, use and density zoning had a meaningful impact on the demographic composition of the neighborhood (including both the total population and the population shares) as well as the local rental market. This indicates that Chicago residents did respond to changes in zoning as if those changes represented a meaningful change in neighborhood amenities.

Of course, because zoning was not randomly assigned, one might worry that places zoned for commercial, manufacturing, or Apartment & Residential uses vary in systematic ways that could be driving the results in Table 2. As mentioned earlier, our approach to overcome this challenge is to use zoning mismatch as a source of identification. To illustrate that zoning mismatch provides us with cleaner identification, we run a placebo test to see how zoning and zoning mismatch affected neighborhood composition in 1910. If our zoning mismatch variable is indeed capturing the exogenous components of zoning then it should not be systematically related to 1910 characteristics. This is exactly what we find in Table 3.

Table 2. Impact of zoning on neighborhood composition in 1940

	Neigh. pop. (1)	Fam. share (2)	White share (3)	Foreign share (4)	No HS share (5)	Owner share (6)	Lab. Force share (7)	Avg. Occ. Inc. (8)	Avg. Rent (9)
Panel a: Zoning and long-run neighborhood development									
Std. Com. share	-491.751***	0.004***	0.014***	-0.000	-0.004***	0.005***	0.002***	0.120***	-5.965***
	(61.270)	(0.001)	(0.004)	(0.002)	(0.001)	(0.001)	(0.001)	(0.043)	(1.512)
Std. Manuf. share	-1133.555***	-0.002**	-0.002	-0.011***	0.005***	0.014***	-0.007***	-0.273***	-5.147*
	(78.757)	(0.001)	(0.004)	(0.003)	(0.001)	(0.001)	(0.001)	(0.052)	(2.800)
Std. Vol. 1 share	-589.584***	-0.000	-0.006**	-0.012***	0.002	0.018***	0.000	-0.156***	-16.274***
	(76.157)	(0.001)	(0.003)	(0.002)	(0.001)	(0.002)	(0.001)	(0.046)	(3.069)
Std. Vol. 4 and 5 share	-174.223**	-0.005***	0.004	0.007***	0.006***	0.003**	-0.002**	-0.066*	2.250
	(76.864)	(0.001)	(0.003)	(0.002)	(0.001)	(0.001)	(0.001)	(0.036)	(1.654)
Panel b: Excess zoning and long-run neighborhood development									
Std. Excess com. share	-119.136***	0.001***	0.008***	0.001	-0.001***	0.001	0.001***	0.047**	-2.310***
	(29.200)	(0.000)	(0.002)	(0.001)	(0.000)	(0.001)	(0.000)	(0.019)	(0.731)
Std. Excess manuf. share	-289.962***	-0.000	-0.006***	-0.004***	0.003***	0.004***	-0.002***	-0.167***	-3.977***
	(35.181)	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)	(0.000)	(0.020)	(1.261)
Std. Excess vol. 1 share	-154.493***	0.001*	0.000	-0.003***	0.001	0.006***	0.000	-0.037**	-7.490***
	(33.317)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.016)	(1.320)
Std. Excess vol. 4 and 5 share	-34.900	-0.000	-0.001	0.000	-0.000	0.000	-0.000	-0.011	0.264
	(23.417)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.008)	(0.283)

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures used in Panel b are calculated following equation (4.1). The unit of observation in each of these regressions is the half-mile radius surrounding each Chicago household. All regressions include a broad set of geographic, land use, and demographic controls. The geographic controls include: the share of the neighborhood that falls into Lake Michigan, the household's distance to the central business district, Chicago river, Lake Michigan, nearest railroad and nearest major street. We also include the square of each of these distances. The land use controls are: the number of buildings in the neighborhood that are between 4-8 stories, 9-11 stories, 12-16 stories, 17-22 stories, and over 23 stories, as well as the number of commercial uses, manufacturing uses, and warehouses. We also include the square of each of these categories as well as the interaction between each category and neighborhood population. The demographic controls include: population (and its square), average dwelling size, child population share, share of households with children, northern and southern black share, share of population born in Eastern, Western, Northern, and Southern Europe, share born in Russia, professional occupation share, white collar share, crafts share, service share, average neighborhood income (as measured by adjusted occupational income scores), standard deviation of neighborhood income, and share of households that own their home. For each of our race/nativity shares we fully interact the share with the neighborhood population. We also include the change in population, white share, foreign share, income, homeowner share, and share in the labor force between 1910 and 1920 as well as the number of homicides and number of gangs.

Panel a of Table 3 presents results assessing the impact of 1920 zoning on 1910 neighborhood demographics. Remarkably, even after conditioning on geography, 1920 land use, 1920 neighborhood demographics, crime, and presence of gang activity, there is still a strong and statistically significant relationship between zoning in 1923 and neighborhood composition in 1910. The reason that this is noteworthy is that this is the exact type of identification used in many existing zoning papers (e.g., McMillen and McDonald (2002) and Shertzer et al. (2016b)). In contrast, Panel B, indicates that our excess zoning measure is largely independent of 1910 neighborhood characteristics. Together, Tables 2 and 3 indicate that excess zoning is capturing the causal effect of zoning on neighborhood development.

Table 3. Impact of zoning on neighborhood composition in 1910

	Neigh. pop. (1)	Fam. share (2)	White share (3)	Foreign share (4)	Owner share (5)	Lab. Force share (6)	Avg. Occ. Inc. (7)
Panel a: Zoning and long-run neighborhood development							
Std. Com. share	-59.313	0.001***	-0.003***	0.001	0.007***	0.007***	-0.129***
	(91.986)	(0.000)	(0.001)	(0.002)	(0.002)	(0.001)	(0.034)
Std. Manuf. share	-39.918	0.003***	0.001	0.009***	0.006***	0.010***	-0.061
	(89.654)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.040)
Std. Vol. 1 share	320.134***	0.002***	-0.006***	-0.005*	0.000	0.001	0.139***
	(77.034)	(0.001)	(0.001)	(0.003)	(0.003)	(0.002)	(0.051)
Std. Vol. 4 and 5 share	-220.963**	-0.002***	-0.005***	0.006***	0.000	0.003***	-0.085***
	(111.365)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.033)
Panel b: Excess zoning and long-run neighborhood development							
Std. Excess com. share	0.899	-0.000*	-0.000	-0.000	0.000	0.000	0.004
	(41.053)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.016)
Std. Excess manuf. share	-22.000	0.000	0.000	0.000	0.001	0.000	0.001
	(36.046)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.019)
Std. Excess vol. 1 share	-22.766	-0.000	-0.000	-0.000	0.001	0.000	-0.003
	(35.608)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.023)
Std. Excess vol. 4 and 5 share	21.383	-0.000	0.000	0.000	-0.000	0.000	-0.004
	(17.994)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.008)

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures used in Panel b are calculated following equation (4.1). The unit of observation in each of these regressions is the half-mile radius surrounding each Chicago household. All regressions include a broad set of geographic, land use, and demographic controls. The geographic controls include: the share of the neighborhood that falls into Lake Michigan, the household's distance to the central business district, Chicago river, Lake Michigan, nearest railroad and nearest major street. We also include the square of each of these distances. The land use controls are: the number of buildings in the neighborhood that are between 4-8 stories, 9-11 stories, 12-16 stories, 17-22 stories, and over 23 stories, as well as the number of commercial uses, manufacturing uses, and warehouses. We also include the square of each of these categories as well as the interaction between each category and neighborhood population. The demographic controls include: population (and its square), average dwelling size, child population share, share of households with children, northern and southern black share, share of population born in Eastern, Western, Northern, and Southern Europe, share born in Russia, professional occupation share, white collar share, crafts share, service share, average neighborhood income (as measured by adjusted occupational income scores), standard deviation of neighborhood income, and share of households that own their home. For each of our race/nativity shares we fully interact the share with the neighborhood population. We also include the number of homicides and number of gangs.

4.3. **Zoning and individual neighborhood choice.** The results in Table 2 confirm that zoning had a meaningful effect on neighborhood development. These results are quite similar to existing results in the environmental justice literature, which suggest that improvements in neighborhood amenities may do little to improve the well-being of current residents because of sorting responses. However, due to a lack of individual data, the existing literature has been unable to say much about the extent to which changes in amenities affect individual well-being. This is the advantage of considering an historical setting, which allows us to track individuals over time. The remainder of our analysis considers these individual outcomes to assess the impact of zoning on individual outcomes.

We begin by estimating the impact of zoning mismatch on the likelihood that an individual moves. An individual is classified as having moved if they are living outside of Chicago in 1940 or if the individual's 1940 geocoded residence is more than 0.1 miles away from their 1920 geocoded residence. As in the previous analysis, our primary explanatory variables of interest are: excess commercial, manufacturing, volume district 1, and volume district 4 & 5 zoning. Our omitted use zoning is Apartment & Residential while our omitted density zoning is Volume district 2 & 3 (where buildings are limited to a height of 8-11 stories). As in our estimation of zoning mismatch in equation (4.1), all regressions include our same three broad categories of controls for geography, pre-existing land use, and neighborhood demographics in 1920. However, we also include the following individual-level controls: birth year fixed effects, race fixed effects, a marital status indicator, indicators for being born in Russia, Northern-, Western-, Southern-, or Eastern Europe, and indicators for whether the individual's father was born in Russia, Northern-, Western-, Southern-, or Eastern Europe. We also control for the family size, the number of children in the household, and the number of children in the household that are under the age of five.

The first column of Table 4 presents our baseline results while columns 2 and 3 restrict the sample to foreign born and black residents, respectively. Relative to Apartment & Residential use zoning, we find little evidence that commercial or manufacturing zoning affected the likelihood that an individual moved between 1920 and 1940. We do see a relationship between low density zoning, but the effect is only meaningful for blacks. For blacks, an increase in low density zoning increases the likelihood of moving by nearly 4 percentage points. In columns 4-6 we restrict the sample to those that remain in Chicago and find very similar results.

In Table 5 we examine the extent to which zoning affected the demographic composition of an individual's neighborhood in 1940. Each column considers a different outcome: neighborhood population, white share, foreign share, average occupational income, homeowner share, and average rent. Because individuals can move in response to zoning, we also fully interact each of our excess zoning measures with an indicator for whether the individual moved between the 1920 and 1940 censuses. Across each outcome and for each type of zoning, it appears that movers ended up in neighborhoods that largely offset the demographic changes that zoning mismatch brought about. For stayers, a standard deviation increase in excess manufacturing zoning reduced the neighborhood population by about 575 residents, decreased average incomes by about $21 (roughly equivalent to one-half of the average Chicago rent in 1940), and lowered rent by a somewhat imprecisely measured $4. For movers, however, the net impact is effectively zero. This suggests that while zoning affected neighborhood development, households that moved were able to successfully mediate the effects of excess zoning.

In Tables 6 and 7 we examine the impact of zoning for foreign born and black residents. The logic for considering these groups separately is twofold. First, while our measure of zoning mismatch represents an exogenous shock to zoning, the firm and household response to that zoning is not exogenous. While manufacturing and

Table 4. Impact of excess zoning on moving propensity

				Remains in Chicago		
	Full Sample (1)	Foreign born (2)	Black (3)	Full Sample (4)	Foreign born (5)	Black (6)
Std. Excess com. share	-0.001	-0.001	0.000	0.000	-0.001	0.001
	(0.001)	(0.003)	(0.005)	(0.002)	(0.003)	(0.007)
Std. Excess manuf. share	-0.002	0.001	0.001	-0.002	0.002	0.007
	(0.001)	(0.002)	(0.007)	(0.002)	(0.003)	(0.010)
Std. Excess vol. 1 share	-0.003*	-0.005*	0.038***	-0.002	-0.004	0.054***
	(0.001)	(0.003)	(0.013)	(0.002)	(0.004)	(0.020)
Std. Excess vol. 4 & 5 share	0.000	0.001	-0.002	0.000	0.001	-0.002
	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)
Observations	98,625	20,090	3,206	71,094	16,685	2,361
R-squared	0.036	0.056	0.087	0.046	0.069	0.115

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures used in Panel b are calculated following equation (4.1). The unit of observation in each of these regressions is the half-mile radius surrounding each Chicago household. All regressions include a broad set of geographic, land use, and demographic controls. The geographic controls include: the share of the neighborhood that falls into Lake Michigan, the household's distance to the central business district, Chicago river, Lake Michigan, nearest railroad and nearest major street. We also include the square of each of these distances. The land use controls are: the number of buildings in the neighborhood that are between 4-8 stories, 9-11 stories, 12-16 stories, 17-22 stories, and over 23 stories, as well as the number of commercial uses, manufacturing uses, and warehouses. We also include the square of each of these categories as well as the interaction between each category and neighborhood population. The demographic controls include: population (and its square), average dwelling size, child population share, share of households with children, northern and southern black share, share of population born in Eastern, Western, Northern, and Southern Europe, share born in Russia, professional occupation share, white collar share, crafts share, service share, average neighborhood income (as measured by adjusted occupational income scores), standard deviation of neighborhood income, and share of households that own their home. For each of our race/nativity shares we fully interact the share with the neighborhood population. We also include the change in population, white share, foreign share, income, homeowner share, and share in the labor force between 1910 and 1920 as well as the number of homicides and number of gangs.

commercial firms are able to choose amongst a set of appropriately zoned areas, they may be disproportionately drawn to areas with more black or foreign residents, perhaps because of lower development costs (e.g., lower land values, more structures that are closer to the end of their usable life, or because of less neighborhood opposition) or because those residents are also their most likely employees. The second reason for considering these groups separately is that these groups may have faced barriers to moving because of formal and informal segregation mechanisms.

Table 6 presents results for foreign born residents. Results for foreign born residents are remarkably similar to the population as a whole. As in Table 5 we continue to see that manufacturing mattered the most and that movers compensated for zoning mismatch by moving to neighborhoods that better matched their preferences. Turning to blacks, we see effects that are much larger in magnitude. Relative to apartment & residential zoning, a standard deviation increase in manufacturing would reduce neighborhood white share by nearly 6 percentage points for blacks that did not move. For comparison, in Tables 5 and 6 the effects of zoning on white share were effectively null. Blacks that moved, however, were able to largely offset these demographic changes.

Table 7 reveals that the largest impact of zoning on blacks came through commercial and low density zoning. A one standard deviation increase in commercial zoning increased population by 1,656 residents, decreased the white share by 6 percentage points and foreign share by 2 percentage points. It also decreased the average neighborhood income by $46. For a standard deviation increase in low-density zoning, average incomes would fall by $130 dollars while the white and foreign shares would decrease by 18 and 7 percentage points, respectively. The large decrease in white and foreign shares for low-density zoning is consistent with the fact that blacks were systematically excluded from most residential areas during this time. Thus, an increase in low-density zoning where a black household was already living in 1920 may have represented an increase in the set of single-family homes that blacks were allowed to live in. As blacks tried to take advantage of this, it is perhaps not surprising that native and foreign born whites, who faced less constraints on single-family residential neighborhood choice, would move in response, leaving the neighborhood to experience a large decline in its white and foreign shares.

Table 5. Impact of excess zoning on 1940 neighborhood demographics

	Pop. (1)	White share (2)	Foreign share (3)	Avg. Occ. income (4)	Share w/o HS edu. (5)	Homeowner share (6)	Avg. rent (7)
Std. Excess com. share	-151.630	-0.000	-0.005**	-0.009	-0.001	0.004*	-2.919*
	(120.230)	(0.002)	(0.002)	(0.041)	(0.002)	(0.002)	(1.501)
Moved X Std. Excess com. share	47.242	-0.001	0.003	0.016	0.000	-0.001	0.817
	(129.971)	(0.002)	(0.002)	(0.045)	(0.002)	(0.002)	(1.453)
Std. Excess manuf. share	-573.901***	-0.001	0.000	-0.206***	0.007***	0.006***	-3.784
	(96.784)	(0.001)	(0.001)	(0.030)	(0.001)	(0.002)	(2.465)
Moved X Std. Excess manuf. share	556.762***	-0.001	-0.001	0.184***	-0.007***	-0.004*	3.837
	(107.511)	(0.001)	(0.002)	(0.033)	(0.001)	(0.002)	(2.338)
Std. Excess vol. 1 share	-212.338**	-0.001	-0.002	-0.041	0.001	0.003	-8.811***
	(101.515)	(0.001)	(0.001)	(0.033)	(0.001)	(0.003)	(2.105)
Moved X Std Excess vol. 1 share	170.139	0.001	0.001	0.055	-0.001	-0.000	6.403***
	(118.765)	(0.001)	(0.001)	(0.037)	(0.002)	(0.003)	(1.944)
Std. Excess vol. 4 & 5 share	-49.869	-0.004**	0.004***	-0.120***	0.006***	-0.003*	-1.374
	(121.421)	(0.002)	(0.002)	(0.038)	(0.002)	(0.002)	(0.981)
Moved X Std. Excess vol. 4 & 5 share	22.374	0.004**	-0.004**	0.111***	-0.005***	0.003*	1.163
	(125.154)	(0.002)	(0.002)	(0.039)	(0.002)	(0.002)	(1.004)
Observations	71,094	71,094	71,094	71,094	71,094	71,094	71,094
R-squared	0.195	0.685	0.395	0.381	0.321	0.260	0.105

*** $p<0.01$, ** $p<0.05$, * $p<0.1$ Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls. The individual controls include: birth year fixed effects, race fixed effects, indicators for whether the individual or their father was born in Russia, Eastern, Western, Northern, and Southern Europe. The household controls include homeownership, family size, number of children in the household, number of children under five in the household, marital status indicator, and an indicator for living in a large dwelling or a multi-family dwelling. The geographic controls include: the share of the neighborhood that falls into Lake Michigan, the household's distance to the central business district, Chicago river, Lake Michigan, nearest railroad and nearest major street. We also include the square of each of these distances. The land use controls are: the number of buildings in the neighborhood that are between 4-8 stories, 9-11 stories, 12-16 stories, 17-22 stories, and over 23 stories, as well as the number of commercial uses, manufacturing uses, and warehouses. We also include the square of each of these categories as well as the interaction between each category and neighborhood population. The demographic controls include: population (and its square), average dwelling size, child population share, share of households with children, northern and southern black share, share of population born in Eastern, Western, Northern, and Southern Europe, share born in Russia, professional occupation share, white collar share, crafts share, service share, average neighborhood income (as measured by adjusted occupational income scores), standard deviation of neighborhood income, and share of households that own their home. For each of our race/nativity shares we fully interact the share with the neighborhood population. We also include the change in population, white share, foreign share, income, homeowner share, and share in the labor force between 1910 and 1920 as well as the number of homicides, and number of gangs.

Table 6. Impact of excess zoning on 1940 neighborhood demographics Foreign born sample only

	Pop. (1)	White share (2)	Foreign share (3)	Avg. Occ. income (4)	Share w/o HS edu. (5)	Homeowner share (6)	Avg. rent (7)
Std. Excess com. share	-92.433	0.005*	-0.004*	0.058	-0.003	0.009***	-1.937
	(171.863)	(0.002)	(0.003)	(0.061)	(0.002)	(0.003)	(1.831)
Moved X Std. Excess com. share	-145.762	-0.004	0.002	-0.044	0.003	-0.003	-0.464
	(180.649)	(0.003)	(0.003)	(0.066)	(0.003)	(0.003)	(1.826)
Std. Excess manuf. share	-492.696***	-0.000	0.002	-0.239***	0.008***	0.005*	-4.212**
	(132.662)	(0.001)	(0.002)	(0.046)	(0.002)	(0.003)	(2.040)
Moved X Std. Excess manuf. share	454.061***	-0.000	-0.002	0.206***	-0.007***	-0.002	4.544**
	(150.521)	(0.001)	(0.002)	(0.051)	(0.002)	(0.003)	(2.046)
Std. Excess vol. 1 share	-255.222*	0.001	-0.001	-0.054	0.001	0.005	-8.259***
	(140.502)	(0.001)	(0.002)	(0.050)	(0.002)	(0.003)	(2.303)
Moved X Std Excess vol. 1 share	248.500	-0.001	0.001	0.047	0.000	-0.003	4.502**
	(162.867)	(0.001)	(0.002)	(0.057)	(0.002)	(0.004)	(2.154)
Std. Excess vol. 4 & 5 share	237.741*	-0.004	0.003	-0.090**	0.004*	-0.005*	-1.757*
	(142.283)	(0.003)	(0.002)	(0.044)	(0.002)	(0.002)	(0.951)
Moved X Std. Excess vol. 4 & 5 share	-232.909	0.003	-0.002	0.076	-0.003	0.004	1.759*
	(148.533)	(0.003)	(0.002)	(0.047)	(0.002)	(0.003)	(0.921)
Observations	16,685	16,685	16,685	16,685	16,685	16,685	16,685
R-squared	0.195	0.101	0.294	0.290	0.296	0.267	0.111

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls, as described in the notes of Table 5.

Table 7. Impact of excess zoning on 1940 neighborhood demographics Black sample only

	Pop. (1)	White share (2)	Foreign share (3)	Avg. Occ. income (4)	Share w/o HS edu. (5)	Homeowner share (6)	Avg. rent (7)
Std. Excess com. share	1656.491***	-0.059***	-0.022***	-0.459**	-0.014**	-0.015*	-2.972
	(558.482)	(0.020)	(0.007)	(0.205)	(0.006)	(0.008)	(3.019)
Moved X Std. Excess com. share	-1356.863**	0.042*	0.018**	0.380*	0.011*	0.013	1.506
	(551.120)	(0.022)	(0.008)	(0.221)	(0.006)	(0.008)	(2.953)
Std. Excess manuf. share	877.214	-0.062**	-0.012	-0.486	-0.005	-0.004	-4.661
	(746.433)	(0.027)	(0.011)	(0.297)	(0.008)	(0.010)	(3.561)
Moved X Std. Excess manuf. share	-1212.144*	0.032	0.006	0.202	0.007	0.002	1.563
	(652.287)	(0.026)	(0.010)	(0.299)	(0.008)	(0.009)	(3.494)
Std. Excess vol. 1 share	-1352.543	-0.178***	-0.074***	-1.307**	-0.016	-0.009	-3.110
	(911.165)	(0.067)	(0.022)	(0.649)	(0.014)	(0.031)	(3.832)
Moved X Std Excess vol. 1 share	2742.623***	0.149**	0.071***	1.019	0.016	-0.014	3.078
	(876.776)	(0.067)	(0.021)	(0.665)	(0.016)	(0.034)	(3.515)
Std. Excess vol. 4 & 5 share	-951.663**	-0.011	-0.000	-0.122	0.012**	0.005	-4.873***
	(367.055)	(0.012)	(0.006)	(0.103)	(0.005)	(0.007)	(1.143)
Moved X Std. Excess vol. 4 & 5 share	845.374**	0.013	0.002	0.158	-0.013***	-0.004	4.592***
	(391.453)	(0.011)	(0.006)	(0.102)	(0.005)	(0.007)	(1.141)
Observations	2361	2361	2361	2361	2361	2361	2361
R-squared	0.102	0.161	0.125	0.120	0.113	0.159	0.144

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls, as described in the notes of Table 5.

4.4. **Zoning and individual outcomes.** Thus far we have seen that zoning affects neighborhood development. We have also consistently found that movers offset the impact of zoning. These neighborhood changes likely reflect the constraints that zoning places on development. Because of this, one natural question is whether zoning also affects economic outcomes like homeownership and labor market attachment. In Tables 8 - 10 we consider the impact of zoning on homeownership, being in the labor force, income, occupational income scores, and weeks worked. Results for the entire sample appear in Table 8, while results for foreign born residents appear in Table 9 and results for blacks appear in Table 10.

In Table 8 we see that, for those that did not move, a standard deviation increase in commercial zoning increased the likelihood that an individual owned a home by about 2.5 percentage points. For those that moved, the net effect on the likelihood of owning a home is effectively zero. These results are consistent with results in the previous section, which illustrated that commercial and manufacturing zoning made the neighborhood less desirable and that movers moved to neighborhoods that did not experience a decrease in desirability. To the extent that this manifests in property values, it is perhaps not surprising to see that stayers were rewarded with an increase in homeownership. Across all other outcomes and types of zoning, we find little evidence that zoning affected individual outcomes.

Table 9 presents results for foreign born residents. There we see results that are nearly identical to Table 8: commercial zoning increased the likelihood of owning a home for stayers but movers saw little gains in homeownership.

As in the tables on neighborhood development, when we focus our attention on blacks, we see effects that are much larger in magnitude. These results appear in Table 10. There we see that, for stayers, an increase in commercial zoning increases the likelihood of owning a home by 7.7 percentage points. Part of this increase in homeownership may be attributable to an increase in labor force attachment. In Column 2 we see that blacks that didn't move were 4.7 percentage points more likely to be in the labor force, given a standard deviation increase in commercial zoning. In Column 3 we see that those blacks saw their incomes increase by about 16 percent and that they worked about 3.8 more weeks in 1939. As in the previous tables, blacks that moved completely offset these effects.

Table 8. Impact of excess zoning on individual outcomes

	Owns home (1)	In Lab. force (2)	ln(Income) (3)	Occ. Inc. score (4)	Weeks worked (5)
Std. Excess commercial share	0.024***	-0.001	-0.000	0.144	-0.023
	(0.006)	(0.003)	(0.012)	(0.143)	(0.235)
Moved X Std. Excess com. share	-0.022***	0.001	0.007	-0.208	0.005
	(0.007)	(0.003)	(0.012)	(0.153)	(0.247)
Std. Excess manufacturing share	0.001	-0.001	-0.002	0.098	-0.092
	(0.005)	(0.002)	(0.008)	(0.110)	(0.191)
Moved X Std. Excess manuf. share	0.003	0.001	0.004	-0.117	0.115
	(0.005)	(0.002)	(0.009)	(0.116)	(0.200)
Std. Excess vol. 1 share	-0.002	-0.002	-0.010	-0.066	-0.245
	(0.004)	(0.002)	(0.009)	(0.102)	(0.176)
Moved X Std. Excess vol. 1 share	0.003	0.003	0.016*	-0.010	0.308
	(0.005)	(0.002)	(0.009)	(0.113)	(0.194)
Std. Excess vol. 4 & 5 share	-0.004	0.003	-0.005	-0.051	-0.123
	(0.008)	(0.003)	(0.013)	(0.157)	(0.228)
Moved X Std. Excess vol. 4 & 5 share	0.006	-0.003	0.005	0.072	0.067
	(0.008)	(0.003)	(0.014)	(0.165)	(0.237)
Observations	71,094	71,094	54,917	60,266	71,094
R-squared	0.112	0.027	0.059	0.073	0.029

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls, as described in the notes of Table 5.

Table 9. Impact of excess zoning on individual outcomes
Foreign born sample only

	Owns home (1)	In Lab. force (2)	ln(Income) (3)	Occ. Inc. score (4)	Weeks worked (5)
Std. Excess commercial share	0.036***	-0.006	0.008	0.020	0.772
	(0.012)	(0.007)	(0.024)	(0.308)	(0.479)
Moved X Std. Excess com. share	-0.024*	0.005	0.002	-0.192	-0.831*
	(0.013)	(0.008)	(0.025)	(0.324)	(0.499)
Std. Excess manufacturing share	-0.012	-0.004	-0.025	-0.079	-0.346
	(0.009)	(0.005)	(0.018)	(0.203)	(0.355)
Moved X Std. Excess manuf. share	0.019*	0.003	0.017	0.006	0.291
	(0.010)	(0.006)	(0.019)	(0.221)	(0.370)
Std. Excess vol. 1 share	-0.003	-0.000	-0.019	-0.143	-0.160
	(0.007)	(0.005)	(0.017)	(0.219)	(0.337)
Moved X Std. Excess vol. 1 share	0.007	0.002	0.017	0.220	0.338
	(0.008)	(0.006)	(0.018)	(0.241)	(0.364)
Std. Excess vol. 4 & 5 share	-0.026**	0.011*	0.006	0.060	0.246
	(0.011)	(0.007)	(0.013)	(0.311)	(0.383)
Moved X Std. Excess vol. 4 & 5 share	0.029**	-0.013*	-0.005	0.112	-0.432
	(0.012)	(0.007)	(0.013)	(0.340)	(0.393)
Observations	16,685	16,685	11,755	14,122	16,685
R-squared	0.142	0.036	0.066	0.069	0.045

*** p<0.01, ** p<0.05, * p<0.1 Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls, as described in the notes of Table 5.

Table 10. Impact of excess zoning on individual outcomes Black sample only

	Owns home (1)	In Lab. force (2)	ln(Income) (3)	Occ. Inc. score (4)	Weeks worked (5)
Std. Excess commercial share	0.072*	0.051**	0.147**	-0.922	3.935**
	(0.040)	(0.024)	(0.072)	(1.256)	(1.676)
Moved X Std. Excess com. share	-0.064	-0.059**	-0.142**	0.573	-3.946**
	(0.041)	(0.024)	(0.070)	(1.242)	(1.645)
Std. Excess manufacturing share	-0.055	0.066**	0.081	-0.981	3.498*
	(0.049)	(0.026)	(0.090)	(1.786)	(1.988)
Moved X Std. Excess manuf. share	0.058	-0.080***	-0.100	0.656	-3.993**
	(0.044)	(0.026)	(0.088)	(1.781)	(1.999)
Std. Excess vol. 1 share	0.048	-0.024	-0.575***	-2.102	-4.050
	(0.082)	(0.041)	(0.145)	(1.990)	(4.210)
Moved X Std. Excess vol. 1 share	-0.073	-0.036	0.617***	0.636	3.801
	(0.080)	(0.038)	(0.138)	(1.809)	(3.704)
Std. Excess vol. 4 & 5 share	0.030*	-0.016	-0.047	1.729*	0.040
	(0.017)	(0.033)	(0.057)	(0.936)	(1.968)
Moved X Std. Excess vol. 4 & 5 share	-0.023	0.018	0.054	-1.538	0.095
	(0.018)	(0.033)	(0.053)	(0.942)	(1.937)
Observations	2,361	2,361	1,822	1,269	2,361
R-squared	0.144	0.099	0.086	0.169	0.080

*** $p<0.01$, ** $p<0.05$, * $p<0.1$ Robust standard errors, clustered at the 1920 enumeration district level, reported in parenthesis. The omitted use zoning is Apartment & Residential and the omitted density zoning is volume 2 & 3. The excess zoning measures are calculated following equation (4.1). All regressions include a broad set of individual, household, geographic, land use, and demographic controls, as described in the notes of Table 5.

5. Conclusion

This paper provides one of the most disaggregated estimates of the welfare effects of zoning. Specifically, we study how the introduction of zoning in Chicago affected neighborhood choice and individual outcomes. By studying the introduction of zoning, we are able to construct a novel and plausibly exogenous measure of zoning, which exploits the fact that in laying out their plan, the zoning commission tended to use polygons which of course did not perfectly reflect the built environment. Further, by studying a historical setting we are able to overcome privacy concerns by tracking individuals with publicly available census data.

Our results indicate that zoning played a meaningful role in shaping neighborhood development. However, individuals were able to largely offset the changes brought about by zoning by moving to neighborhoods that better matched their preferences. Thus, the costs and benefits of zoning appear to be borne by those that stayed. We find that black neighborhoods in particular saw the largest changes – perhaps because developers responding to the zoning ordinance selected into black neighborhoods because of lower development costs. While blacks that stayed behind saw a large decline in neighborhood quality, they appear to have benefited from greater job access.

References

Alsan, Marcella and Claudia Goldin (2015). Watersheds in infant mortality: The role of effective water and sewerage infrastructure, 1880 to 1915.Technical report, National Bureau of Economic Research.

Bailey, Martha, Connor Cole, Morgan Henderson, and Catherine Massey (2017). How well do automated linking methods perform in historical samples? Evidence from new ground truth. *Unpublished manuscript*.

Banzhaf, H. Spencer and Randall P Walsh (2008). Do people vote with their feet? An empirical test of Tiebout's mechanism. *The American Economic Review*, **98**(3), pp. 843–863.

Bento, Antonio, Matthew Freedman, and Corey Lang (2015). Who benefits from environmental regulation? evidence from the clean air act amendments. *Review of Economics and Statistics*, **97**(3), pp. 610–622.

Depro, Brooks, Christopher Timmins, and Maggie O'Neil (2015). White flight and coming to the nuisance: can residential mobility explain environmental injustice? *Journal of the Association of Environmental and resource Economists*, **2**(3), pp. 439–468.

Feigenbaum, James (2016). Automated Census Record Linking: A Machine Learning Approach. *Unpublished manuscript*.

Ferrie, Joseph and Werner Troesken (2008). Water and Chicago's mortality transition, 1850–1925. *Explorations in Economic History*, **45**(1), pp. 1–16.

Glaeser, Edward L. and Bryce A. Ward (2009). The causes and consequences of land use regulation: Evidence from Greater Boston. *Journal of Urban Economics*, **65**(3), pp. 265–278.

Grainger, Corbett A (2012). The distributional effects of pollution regulations: Do renters fully pay for cleaner air? *Journal of Public Economics*, **96**(9), pp. 840–852.

Grigoryeva, Angelina and Martin Ruef (2015). The historical demography of racial segregation. *American Sociological Review*, **80**(4), pp. 814–842.

Hornbeck, Richard and Daniel Keniston (2017). Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872. *American Economic Review*, **107**(6), pp. 1365–1398.

Hsieh, Chang-Tai and Enrico Moretti (2017). Housing Constraints and Spatial Misallocation.

Ihlanfeldt, Keith R. (2007). The effect of land use regulation on housing and land prices. *Journal of Urban Economics*, **61**(3), pp. 420–435.

Koster, Hans R.A., Jos van Ommeren, and Piet Rietveld (2012). Bombs, boundaries and buildings: A regression–discontinuity approach to measure costs of housing supply restrictions. *Regional Science and Urban Economics*, **42**(4), pp. 631–641.

Kuminoff, Nicolai, V. Kerry Smith, and Christopher Timmins (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, **51**(4), pp. 1007–1062.

Logan, Trevon and John Parman (2017). The national rise in residential segregation. *The Journal of Economic History*, **77**(1), pp. 127–170.

Mayer, Christopher J and C Tsuriel Somerville (2000). Land use regulation and new construction. *Regional Science and Urban Economics*, **30**(6), pp. 639–662.
McMillen, Daniel P. and John F. McDonald (1991). A Markov chain model of zoning change. *Journal of Urban Economics*, **30**(2), pp. 257–270.
McMillen, Daniel and John McDonald (1999). Land use before zoning: The case of 1920's Chicago. *Regional Science and Urban Economics*, **29**(4), pp. 473–489.
McMillen, Daniel P. and John F. McDonald (2002). Land values in a newly zoned city. *Review of Economics and Statistics*, **84**(1), pp. 62–72.
Munneke, Henry J. (2005). Dynamics of the urban zoning structure: An empirical investigation of zoning change. *Journal of Urban Economics*, **58**(3), pp. 455–473.
Murray, Cameron K. and Paul Frijters (2016). Clean money, dirty system: Connected landowners capture beneficial land rezoning. *Journal of Urban Economics*, **93**, pp. 99–114.
Rothwell, Jonathan T (2011). Racial enclaves and density zoning: the institutionalized segregation of racial minorities in the United States. *American Law and Economics Review*, **13**(1), pp. 290–358.
Rothwell, Jonathan and Douglas S Massey (2009). The effect of density zoning on racial segregation in US urban areas. *Urban Affairs Review*, **44**(6), pp. 779–806.
Shaw, Clifford, Leonard Cottrell, Henry McKay, and Frederick Zorbaugh (1929). *Delinquency areas: A study of the geographical distribution of school truants, juvenile delinquents, and adult offenders in Chicago*. University of Chicago Press.
Shertzer, Allison, Tate Twinam, and Randall Walsh (2016a). Race, ethnicity, and discriminatory zoning. *American Economic Journal: Applied Economics*, **8**(3), pp. 1–31.
——— (2016b). Zoning and the economic geography of cities. *Unpublished manuscript*.
Siodla, James (2017). Clean slate: Land-use changes in San Francisco after the 1906 disaster. *Explorations in Economic History*.
Sullivan, Daniel M (2016). Residential Sorting and the Incidence of Local Public Goods: Theory and Evidence from Air Pollution. *Resources for the Future Working Paper*.
Thrasher, Frederic M. (1927). *The gang: A study of 1,313 gangs in Chicago*. University of Chicago Press.
Troesken, Werner (2004). *Water, race, and disease*. MIT Press.
Turner, Matthew A., Andrew Haughwout, and Wilbert Van Der Klaauw (2014). Land use regulation and welfare. *Econometrica*, **82**(4), pp. 1341–1403.
Twinam, Tate (2017). The long-run impact of zoning: Institutional hysteresis and durable capital in Seattle, 1920-2015. *Unpublished manuscript*.
Wallace, Nancy E. (1988). The market effects of zoning undeveloped land: Does zoning follow the market? *Journal of Urban Economics*, **23**(3), pp. 307–326.
Zhou, Jian, Daniel P McMillen, and John F McDonald (2008). Land values and the 1957 comprehensive amendment to the Chicago zoning ordinance. *Urban Studies*, **45**(8), pp. 1647–1661.

6. Appendix

Figure 1. 1922 Land Use Survey Sample

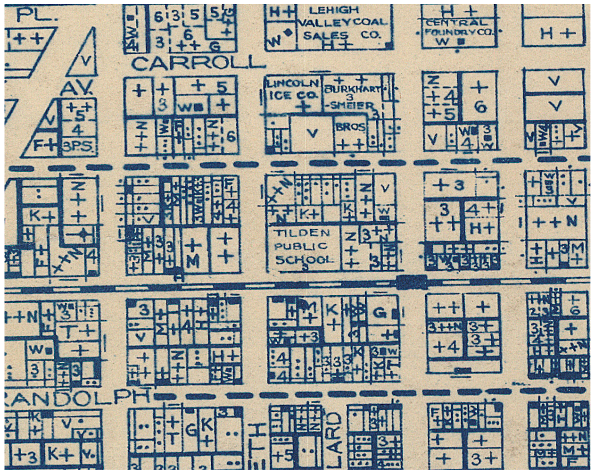

A portion of the 1922 land use survey map created by the Chicago Zoning Commission. These blocks are located just across the Chicago River to the west of the downtown. Numbers indicate building heights in stories. Black squares within parcels indicate commercial uses; letters sometime accompany these to indicate a specific commercial activity. V indicates a vacant lot/building. Letters followed or preceded by a single + indicate light industrial uses. Letters preceded by ++ indicate heavier industrial uses; in particular, ++N indicates uses which “by reason of excessive noise, odor, fumes, gases, etc., affect the adjacent territory.”

Figure 2. 1922 Land Use Survey Digitized

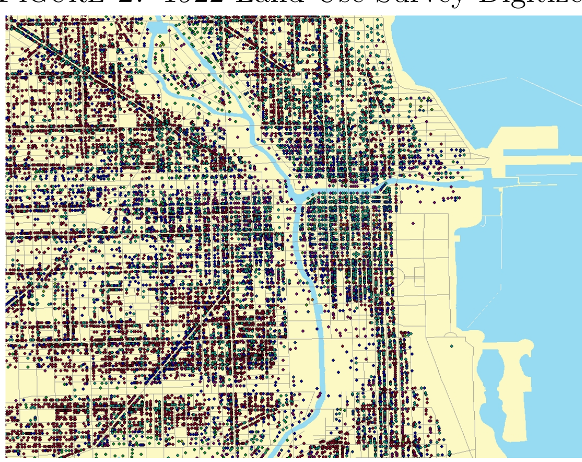

Figure depicts land use in the Loop, Little Italy, Near West Side, and Near North Side.

FIGURE 3. 1923 Zoning Map Samples

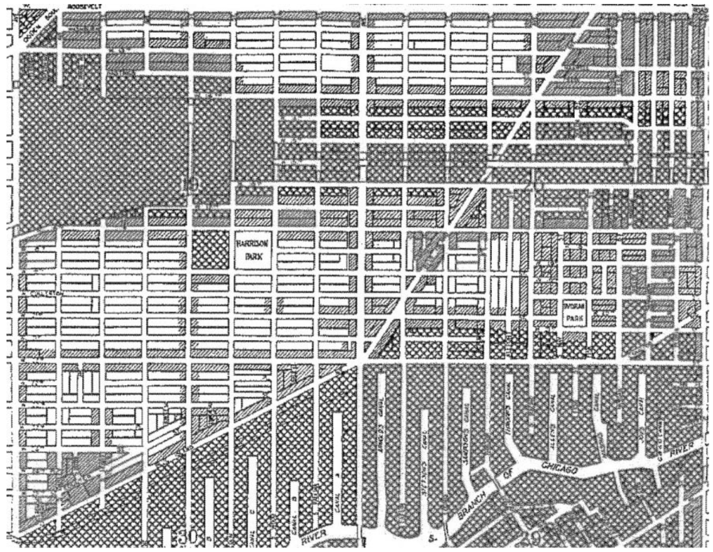

(A) Portion of 1923 use zoning map. Unhatched areas are zoned for apartments, hatched areas are zoned for commercial uses, and cross–hatched areas are zoned for manufacturing.

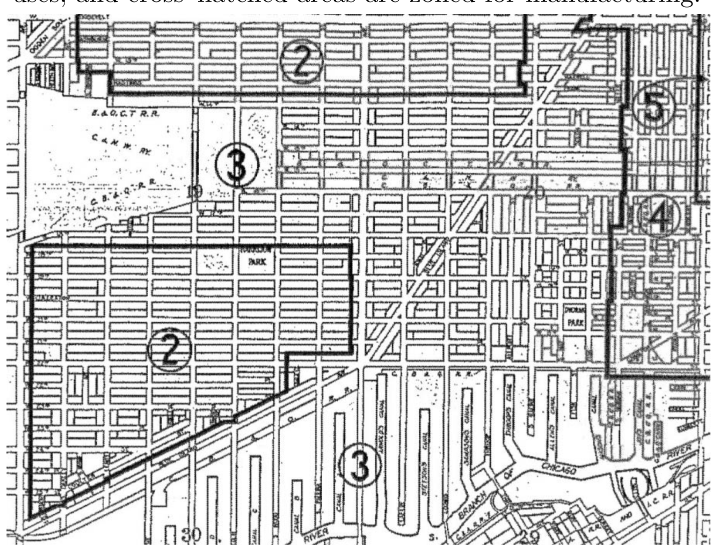

(B) Portion of 1923 volume zoning map. Zone 2 is the lowest density area depicted here, accommodating low–rise apartment buildings. Zone 5 is the highest density area, allowing for skyscrapers over 20 stories.