Consumption Segregation*

Corina Boar† Elisa Giannone‡

July 2023

Abstract

This paper introduces consumption segregation, a new margin of residential segregation, and examines its patterns, causes, and discusses its aggregate consequences. We use new longitudinal and highly granular data to measure consumption segregation in the United States and document that it is high but relatively stable over the past 15 years, with substantial regional variation. We find that income segregation plays a more prominent role than other forms of segregation in driving consumption segregation, mainly due to the inability to smooth shocks to income. We illustrate a new mechanism through which, in the presence of social comparisons, consumption segregation can exacerbate wealth inequality.

*We thank Federico Kochen for superb research assistance, our discussants Rebecca Diamond, Eduardo Morales, and David Cutler for insightful feedback, as well as seminar participants at the ASSA Meetings, SED Meeting, the Federal Reserve Bank of Richmond, Dartmouth Mini Conference, CREI, OSUS, Barcelona Summer Forum, NBER SI Income distribution and NBER SI Urban for their helpful comments. AElisa Giannone acknowledges nancial support from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in RD (Barcelona School of Economics CEX2019-000915-S)"

†New York University and NBER, corina.boar@nyu.edu.

‡Center de Recerca en Economia Internacional and Barcelona School of Economics, egiannone@crei.cat.

1 Introduction

Over the past few decades, economic inequality in the United States has been on the rise, accompanied by an increase in spatial segregation by income. These trends have created pressing socio-economic issues, which have led to extensive research aimed at understanding their causes and effects.[1]

While there is ample work on income and other forms of segregation, less is known about the segregation of consumption. This is important because consumption is arguably more directly related to well-being than income, and can provide a better understanding of the segregation of welfare. Additionally, as we show in this paper, the residential segregation of consumption can exacerbate wealth inequality through a new channel. Due to market imperfections or permanent income differences, income segregation leads to consumption segregation. In the presence of social comparisons, in less segregated areas, this creates an externality through which the desire to keep up with the consumption of richer neighbors reduces wealth inequality.

Despite its importance, analyzing differences in consumption over time and across space remains challenging because of the considerable obstacle of data availability. In this paper, we use new longitudinal and highly granular consumer-level data that allows us to progress in this direction: we characterize the patterns of residential consumption segregation, study its determinants, and highlight its potential effects on aggregate outcomes. We proceed in three steps. First, we provide the first account of consumption segregation in the United States after validating the new data. Second, we study the drivers of consumption segregation and document a strong relationship with income segregation. Third, through the lens of a stylized model of consumption and geography, we argue that consumption segregation may lead to higher wealth inequality in the presence of social comparisons.

We begin our analysis by providing a thorough characterization of consumption segregation. We utilize two sources of data. First, we leverage a newly built, large-scale dataset from Infutor, a private company. This dataset covers nearly 250 million individuals, 160 million vehicles, and over 150 million properties. We use the extensive information on car and home ownership in this dataset, connected to demographic information, to analyze the extent of durable consumption segregation. Second, we complement it with Nielsen Homescan data to

[1] A large body of work has studied the relationship between the two, emphasizing the role of human capital formation as a driver (Fernández and Rogerson, 1998, Fogli and Guerrieri, 2019, Zheng and Graham, 2020, Eckert and Kleineberg, 2021).

measure the segregation of non-service retail consumption, which we refer to as non-durable consumption for brevity throughout the paper. One key limitation of the data we use is that it does not include services. Therefore, our focus is on durable and non-service retail consumption, which together account for 51% of total household spending.

One of the primary goals of this work is to provide researchers with measures of consumption differences across the US by item, over time, and across demographic groups. Therefore, before conducting our measurement of consumption segregation, we conduct an extensive validation exercise of the Infutor data which, to the best of our knowledge, has not been used to measure durable consumption. We find that it captures well the demographic composition of widely used datasets such as ACS. It also matches a host of statistics regarding house values, homeownership shares and the vehicle fleet in the United States.

The Infutor dataset has multiple advantages. It has detailed geographic information which, combined with the large sample size, allows for a precise analysis of consumption differences across space. It has a relatively extended time coverage which allows us to explore trends in consumption differences. Additionally, it includes both values and characteristics of properties, as well as types of vehicles, thus allowing us to make a distinction between spending and consumption and bypass the lack of fine regional price data over time.

We measure consumption segregation using the entropy and rank-order index, which capture how representative the distribution of consumption is in a geographic unit (such as a PUMA) compared to that in a broader region (such as a CBSA, a state, or the country). The index ranges from 0 to 1, with 0 indicating perfect integration (i.e., each geographic unit is representative of the broader region), and 1 indicating perfect segregation (i.e., there is no diversity within a geographic unit). This index is widely used in segregation literature and performs better than other indices in terms of its conceptual and mathematical properties.[2]

We find that, unlike the segregation of income, the segregation of consumption has been relatively steady during the last fifteen years, concealing a decline in the segregation of non-durable consumption, and an increase in the segregation of vehicle consumption. The category with the most significant segregation is housing consumption. Furthermore, we uncover significant regional heterogeneity: consumption is most segregated in New York, where the average PUMA is 30% less diverse than the state. In contrast, Wyoming has the lowest consumption segregation, with levels approximately ten times lower than in New York. We also find that consumption segregation is higher in regions that are richer, larger,

[2]See Reardon and Firebaugh (2002) for an in-depth comparison of segregation indices.

and younger, and where a significant proportion of the population has a college degree and a smaller share of the population is white. Lastly, we find that consumption segregation is a shared characteristic of socio-economic groups and only to a much smaller extent a reflection of the differences between these groups.

Having characterized the patterns of consumption segregation, we then turn to examine its causes and implications. Specifically, we explore the extent to which consumption segregation is related to other dimensions of segregation that have been previously studied. By doing so, we establish a connection between these dimensions of segregation and the segregation of welfare. We find an important role for income segregation as a determinant of consumption segregation, and a more muted role for racial segregation. Our result is consistent with standard consumption-saving theory, which predicts that income and consumption are positively related due to market incompleteness or permanent income differences that prevent consumption insurance. We apply the insights of Krueger and Perri (2006) to establish a more predominant role for incomplete markets rather than persistent income differences in explaining the patterns in the data.

We argue that consumption segregation might exacerbate wealth disparities among individuals in the presence of social comparisons, where people have the desire to keep up with their neighbors' spending patterns (i.e., a "keeping up with the Joneses" consumption motive). We conceptualize this hypothesis in a stylized framework with regions and wealth accumulation. Intuitively, in areas with low levels of segregation, poor households have the incentive to save to afford purchasing visible (and often durable) consumption goods that allow them to keep up with their richer neighbors, thus reducing the wealth gap between the rich and the poor. We test the predictions of the model regarding the relationship between consumption segregation and wealth inequality and find empirical support for the theory.

Related Work. By studying the segregation of consumption, we contribute to the vast literature on residential segregation that studies trends in segregation by race (e.g., Cutler and Glaeser (1997), Cutler et al. (1999), Cutler et al. (2008)) or by income (e.g., Reardon and Firebaugh (2002)). We refer the reader to Trounstine (2018) for a comprehensive description of this work. We contribute to this literature by moving beyond income and characterizing the trends and spatial patterns of consumption segregation.

More recently, a new set of studies have focused on consumption differences across space. Agarwal et al. (2017) studies how consumers' willingness to travel shapes the characteristics of industries that deliver final consumption. Davis et al. (2017) uses Yelp reviews to study the role of spatial and social frictions in influencing restaurant visits in New York City. Allcott et al. (2019a) studies the extent to which the presence of supermarkets affects nutritional inequality. Lastly, Diamond and Moretti (2022) provides novel evidence on expenditure differences across individuals and commuting zones using a cross-section of credit card transaction data, and connects income and consumption inequality through sorting across space. We complement these studies in two main dimensions. First, we focus on the segregation of the main durable consumption categories, vehicles and housing, as well as that of non-service retail consumption for the entire country. Second, we explore the time dimension of spatial differences in consumption.

Our paper is also related to the work that studies the relationship between income segregation and economic inequality. This work has emphasized the role of human capital formation in shaping this relationship, via neighborhood choice and schooling (Fernández and Rogerson, 1998, Zheng and Graham, 2020, Eckert and Kleineberg, 2021), or more general spillovers at a local level (Fogli and Guerrieri, 2019). We complement this line of work by identifying a new channel through which income segregation can lead to economic inequality that operates through the residential segregation of consumption.

By emphasizing the role of income segregation as a driver of consumption segregation, our paper is also related to papers that study the relationship between income and consumption inequality over time: Krueger and Perri (2006), Blundell et al. (2008), Fisher et al. (2013), Aguiar and Bils (2015). Differently from these papers, we focus on segregation as the geographical element of inequality. We also highlight a new channel through which differences in consumption can affect inequality in the presence of social comparisons.

The remainder of the paper proceeds as follows. Section 2 describes the data sources used in the analysis. Section 3 documents the patterns of consumption segregation. Section 4 analyzes the drivers of consumption segregation. Section 5 investigates a new channel through which consumption segregation intermediates the link between income segregation and wealth inequality. Finally, section 6 concludes.

2 Data and Validity

This section introduces the data used in our analysis. A primary objective of our paper is to construct a comprehensive dataset on consumption across space and time that other researchers can download and utilize. As a result, a substantial portion of this section focuses on explaining and validating the data and the construction our primary variables of interest.

Our empirical assessment of the segregation of consumption and its relationship with other dimensions of residential segregation is based on three datasets that enable us to observe the income and consumption choices of individuals, their demographic characteristics, and track their locations precisely. Specifically, we analyze the consumption segregation of durables such as cars and houses using the Infutor data, non-service retail consumption segregation using the Nielsen Homescan data, and income segregation using data from the American Community Survey (ACS). We provide a detailed discussion of these datasets below, with particular emphasis on the Infutor data, which is relatively less used, especially in the study of durable consumption. A noteworthy feature of this dataset is that it allows us to examine the quality of goods, which is closer to effective consumption than spending itself. For validation, we compare key summary statistics with counterparts from the ACS and the Bureau of Transportation Statistics (BTS).[3]

2.1 Infutor Data

This dataset is compiled by Infutor, a private data vendor based in Chicago that aggregates publicly available information on individuals and makes it available for purchase to both researchers and for business.[4]

Data Sources Infutor uses multiple data sources: phone books, voter files, credit header files, public government records, property deeds, county property records, vehicle warranties, and data from vehicle repair and maintenance providers. Most of these sources are publicly available. The key innovation is that Infutor developed an algorithm to identify individuals across these different datasets and link them to each other. The resulting data contains information on more than 250 million individuals, which are linked to 160 million vehicles and more than 150 million properties spanning from 2000 to 2018. This dataset includes demographic information for consumers (gender, age, marital status), vehicle information (make, model, year), household income, property information (value, construction quality), length of residence, type of dwelling, geographic delineations (time zone, county), and address information (street address, state, city, zip). Identifiable information within this dataset includes Social Security number, Vehicle Identification Number, date of birth, and name.

These characteristics and the size of the sample make for an unprecedented dataset. Additionally, the dataset also contains person identifiers that could facilitate merging with

[3]In Appendix A, we provide further details about these three datasets.

[4]Penn State University purchased the data in 2019 from Infutor.

other datasets, thus increasing its value. While other studies have used the component on the history of the address of Infutor, to the best of our knowledge, this is the first study that employs the vehicle, properties, and demographic components together. As a result, we devote a significant part of our analysis to testing the validity and representativeness of the data. In the following sections, we describe the structure of the dataset and discuss its representativeness, comparing it with data from the Census, American Community Survey (ACS), and Bureau of Transportation Statistics (BTS) data, as well as its advantages and shortcomings. A contains a more detailed description of our treatment of the data.

2.1.1 Structure of the Infutor Data

The Infutor data is organized into four linkable files: Auto Profiles, Property Profiles, Demographic Profiles, and History of Addresses. In this project, we primarily use the first three files but describe the latter for completeness. We have access to different snapshots of this data. Starting in 2012 through 2018, we observed the data at an annual frequency. After March 2019, we observed the data at a monthly frequency. For housing, the initially available snapshot was in 2015. Although the snapshots we observed start in 2012 or 2015, ownership records and the history of addresses go much further back in time. For example, we observe ownership records of vehicles starting in the 1980s, while property deed records go back to the 1950s. The longitudinal aspect of the data enables us to observe the history of car ownership and residences, and also study trends in consumption segregation in addition to the regional dimension.

We next describe the content of the Infutor data files that are most relevant to our paper.[5]

Auto Profiles. This file provides information on individual car ownership, including characteristics such as the car's manufacturer (make), model, manufacturing year, and the first 10 digits of the Vehicle Identification Number (VIN). It also includes the owner's address, allowing us to identify the vehicle's location and select demographic and economic characteristics. This data file contains approximately 160 million observations during any given period (i.e. November 2012-2018 and every month of 2019), which we use to measure the segregation of car consumption.

In Section 3.1, we compute two measures of segregation: along categorical dimensions, such as car model, and along continuous dimensions, such as car value. To measure car values

[5]Figure 5 in the Appendix contains a visual summary of this section.

in the Infutor data, we merge it with transaction prices for all dealer sales of new and used cars in Texas from 2012 to 2019, based on the VIN. When the VIN is not available, we infer prices using the average transaction price of a make-model-year combination.

While we observe Auto Profile snapshots starting in 2012, information on vehicle ownership records dates back to the 1980s. With additional assumptions, this allows us to study a more extended time dimension of segregation. Specifically, for years prior to the first time we observe a vehicle, we set the location of the vehicle to be the zipcode of the first snapshot in which we observe it. We are confident that this does not introduce substantial noise in our measurement, as between 2012 and 2018 only 3.3% of the vehicles in the sample changed PUMAs, while only 1.1% changed states. For years prior to 2012, we impute the value of a vehicle using a VIN-level annual depreciation rate. Appendix A contains a detailed discussion of these procedures.

Property Profiles. This file contains information on residential properties, including the home's address, value, and characteristics such as the year it was built, square footage, and building quality. The dataset includes approximately 150 million observations in any given period, and we focus only on properties used for residential purposes and not for business. We use this information to measure the segregation of housing consumption.

The dimension of housing segregation we study is the value of the home. To measure this, we use property deeds data that provide information on the date and price at which a property was acquired by the current owner. While only snapshots of Property Profiles from 2015 are observed, property deeds records are available from as early as the 1950s, allowing for a generous time dimension. However, the market value of a home is only available during the period in which a home was transacted and a deed exists. To capture the most relevant aspects of segregation, we use zipcode-level price indices based on homes that are transacted to impute the value of a home in periods between two consecutive deeds. Appendix A.2 shows that this imputation procedure generates levels and dynamics of house prices that align with those published by the Federal Housing Finance Agency.

Demographic Profiles. This file contains information on the demographic and economic characteristics of approximately 250 million individuals in any given period, along with their addresses. The file includes data on individuals' age, gender, ethnicity, marital status, educational attainment, number of children, and estimated income, wealth, and home value. However, it should be noted that the Infutor data does not contain information on race. In order to study the relationship between race, income, and consumption, we probabilistically assign race to individuals in the data using an algorithm proposed by the Census, which assigns race based on surnames according to the 2010 Census Surname Table.[6] We verify the accuracy of this imputation by comparing the race distribution in the Infutor data with that in the American Community Survey.

History of Addresses. This dataset contains information about individuals' past and current addresses in the United States, including the time period in which they lived there between 1990 and 2019. Although we do not utilize this data for our project, we mention it for completeness. Diamond et al. (2019) use the detailed information in this dataset to examine the effects of rent control on relocation in San Francisco, while Qian and Tan (2020) investigate how the entry of new firms into a location impacts individuals' choices.

2.1.2 Representativeness and Validity of Infutor Data

Naturally, there is a concern about the representativeness of the Infutor dataset for the US population and specific demographic groups. Although the large number of observations may be indicative of the reliability of the data, researchers such as Bernstein et al. (2018) and Phillips (2019) have systematically compared it with Census data and found that it covers 78% of the US Census estimated adult population and matches the cross-sectional distribution of the population across counties. In this section, we aim to answer two questions: first, how representative is the Infutor data of the overall US population and the different demographic groups we are interested in studying? Second, is the Infutor data suitable for studying questions related to durable consumption segregation across space and over time?

Demographics. To answer the first question, Table 1 presents a comparison of summary statistics for various demographic groups between the Infutor data and the ACS. Column 1 reports the distributions of gender, race, educational attainment, age, and household income in the ACS, while column 2 reports the corresponding distributions in the Infutor data. To assess the spatial representativeness of Infutor data for the different demographic groups of

[6] We used file B from https://www.census.gov/topics/population/genealogy/data/2010_surnames.html. Census reports last names that occurred more than 100 times in its collected 2010 data, and the proportion of people with each last name that belongs to each race category (White, Black, Asian/Pacific Islander, American Indian/Alaskan Native, bi/multi-racial, and Hispanic). We assigned the likelihood of belonging to each race group based on individuals' last names using the Census race data. We were able to assign race probabilistically to 90.25% of the sample.

interest, in columns 3 and 4 we report the cross-state and cross-PUMA correlations of the shares of each demographic group in each state and PUMA between the two datasets.

In the top panel of Table 1, we present the results for gender. The share of women in the Infutor data is 52%, very close to the 51% in the ACS. The cross-state and cross-PUMA correlations of gender shares between the two datasets are 0.8 and 0.59, respectively, indicating a high level of correspondence.

In the second panel, we report the results for race. In the ACS data, 75% of individuals identify as white, 12% as Black, and the remaining 13% as belonging to other racial groups. In the Infutor data, the corresponding shares are 62%, 12%, and 26%, respectively. Although Infutor uses an imputation procedure to assign race probabilistically, the high correlations between the race shares in the two datasets for U.S. states and PUMAs suggest that Infutor accurately captures the distribution of racial groups across space. For example, the state-level correlation of the share of the white population is 0.71, while those of Blacks and other racial groups are 0.8 and 0.86, respectively. The accurate representation of the racial distribution across space is crucial for our analysis of the role of racial segregation in driving consumption segregation.

In the third panel, we report the shares of the population with and without a college degree. The share of the college-educated population in the Infutor data is higher at 51%, compared to 29% in the ACS. However, the high correlations in Columns 3 and 4 suggest that Infutor data can effectively capture which U.S. states and PUMAs have the highest share of college-educated individuals.

The fourth panel of the table shows that the age distribution in the Infutor data closely resembles that in the ACS, except for the under-representation of the age group 20-29 in the Infutor data. This factor may explain the higher college share in the Infutor data since the dataset may not cover very young individuals who are still in college.

Lastly, the fifth panel of Table 1 reports the income distribution in the two datasets. Although income is imputed in the Infutor data, the income distribution in both datasets is similar, with approximately 10% of individuals earning less than $20,000 per year, and the most common income range being between $50,000 and $75,000 per year. However, we note that the Infutor data under-represents individuals with an annual income between $20,000 and $50,000, as well as those with an annual income over $125,000. The high correlations in Columns 3 and 4 indicate that the Infutor data can capture the distribution of income not only nationally but also at the level of U.S. states and PUMAs.

Table 1: Summary Statistics of Demographics and Representativeness of the Infutor Data

		Population shares		Correlation	
		ACS	Infutor	State	PUMA
Gender					
	Female	0.51	0.52	0.80	0.59
	Male	0.49	0.48	0.80	0.59
Race					
	White	0.75	0.62	0.71	0.59
	Black	0.12	0.12	0.80	0.79
	Other	0.13	0.26	0.86	0.74
Education					
	Non-college	0.71	0.49	0.70	0.87
	College degree	0.29	0.51	0.70	0.87
Age					
	20-29	0.19	0.07	-0.13	0.38
	30-39	0.19	0.17	0.25	0.65
	40-49	0.18	0.20	0.83	0.67
	50-59	0.19	0.23	0.31	0.63
	60-69	0.16	0.20	0.37	0.78
	70-80	0.10	0.13	0.72	0.90
Income (USD)					
	<20,000	0.11	0.10	0.79	0.85
	20,000-29,999	0.08	0.06	0.75	0.82
	30,000-39,999	0.08	0.12	0.59	0.71
	40,000-49,999	0.08	0.12	0.83	0.68
	50,000-74,999	0.18	0.27	0.85	0.71
	75,000-99,999	0.14	0.16	0.52	0.55
	100,000-124,999	0.10	0.09	0.82	0.64
	≥125,000	0.22	0.08	0.95	0.89
Number of observations		249,773,525	239,077,422	0.99	0.74

The table reports population shares by demographic group in the Infutor data and in the ACS, as well as the spatial correlation of these population shares across the two datasets, both at the state and at PUMA level.

Auto. To answer the second question, Table 2 presents summary statistics on car ownership and the age of cars in the Infutor data, contrasting them with statistics from The Bureau of Transportation Statistics (BTS). The table reports the number of vehicles and the average age of a vehicle, with the BTS reporting the average age of a car as 11.5 years.

The Infutor dataset's cars are, on average, seven months older. According to the BTS, there are 254,582,694 vehicles, with Infutor covering 144,863,872 of them, accounting for approximately 57% of the total number of cars in the United States. The state-level correlation between the number of vehicles reported by the BTS and those in the Infutor data is 0.54, suggesting that our data accurately capture the spatial distribution of vehicles, making it suitable for studying the segregation of this component of durable consumption. To provide further validation of the Infutor data, Table 11 of Appendix B reports the share of vehicles by make and model for 2009 and 2017, comparing Infutor and the National Household Travel Survey (NHTS). Overall, we find that the two datasets have similar shares of cars by make and model at both the aggregate level and the state level.

Table 2: Summary Statistics of Vehicles and Representativeness of the Infutor Data

	National level		Correlation
	BTS	Infutor	State-level
Vehicle age	11.5	12.2	
N	254,582,694	144,863,872	0.54

Note: The table reports the average age and number of vehicles in the BTS and in Infutor, as well as the state-level correlation between the number of vehicles in the two datasets.

Properties. To further address the second question, Table 3 presents summary statistics on homeownership rates, house value, house size, and year of construction in the Infutor dataset. We compare the homeownership rates and house values in the ACS with those in the Infutor data in columns 1 and 2, respectively. In the ACS, the homeownership rate is 67%, while in Infutor, it is 62%. We also observe a high correlation between the homeownership rates at the state and PUMA level in the two datasets, indicating that the Infutor data captures not only the average homeownership rate but also regional patterns.

Furthermore, we find that the Infutor data accurately captures the value of houses. The average house value is $319,000 in the ACS and $360,000 in the Infutor data, while the median home value is $218,000 in the ACS and $222,000 in the Infutor data. The regional correlation of house values is also high, as shown in columns 3 and 4. Finally, the table reports that the average house size is 1,900 square feet, and the average year of construction is 1974, although this information is not available in the ACS.

Table 3: Summary Statistics of Properties and Representativeness of the Infutor Data

	National level		Correlation	
	ACS	Infutor	State	PUMA
Homeownership share	0.67	0.62	0.88	0.88
House value (USD)				
Mean	319,081	361,336	0.58	0.60
Median	218,000	221,667	0.96	0.98
Average house size (SqFt)		1900		
Year of construction		1974		
Number of observations	137,389,926	109,691,219	0.98	

Note: The table reports homeownership shares and house values in the ACS and in the Infutor data, as well as the spatial correlation of these statistics across the two datasets.

2.1.3 Advantages and Disadvantages of the Infutor Data to Study Consumption

A considerable body of literature employs the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) as the primary sources of consumption data (Krueger and Perri, 2006, Aguiar and Hurst, 2007, Aguiar and Bils, 2015, Andreski et al., 2014, Boar, 2021, Aguiar et al., 2020). The advantage of these datasets is that they cover a significant share of spending: PSID covers roughly 70% of NIPA consumption, and CEX covers an even more substantial proportion. However, the sample sizes of both PSID and CEX are around 5,000 households, too small to undertake our exercise at fine levels of geography (finer than a U.S. state).

Recent studies of consumption patterns rely on credit and debit card transactions or personal finance management applications (Agarwal et al., 2017, Ganong and Noel, 2019, Dolfen et al., 2019, Olafsson and Pagel, 2018, Diamond and Moretti, 2022). Compared to these studies, the Infutor data has three distinct advantages. First, it enables us to accurately capture durable consumption such as cars and housing, which are likely not purchased using credit cards. Second, it has a relatively extended time coverage, enabling us to infer trends for two decades. Third, besides observing the value of the goods, the Infutor data allows us to observe the types of goods, thereby allowing us to distinguish between consumption and expenditure, which is particularly useful in light of the challenges of observing prices at such a granular level of geography.

However, one drawback of using the Infutor data to study consumption segregation is that it only covers durable consumption. To overcome this limitation, we supplement our analysis with Nielsen Homescan data, which provides rich information about household purchasing patterns at various retailers, such as Walmart or Whole Foods, enabling us to study non-durable consumption segregation. Although using the Nielsen Homescan data expands our coverage of consumption categories, we cannot measure services such as restaurant meals, haircuts, taxi rides, or other leisure activities. This dimension of consumption segregation has been partly studied by Davis et al. (2017) in the context of restaurant meals in New York City using Yelp data and theoretically by Couture et al. (2019), so we view our work as complementary. Overall, the consumption categories we capture in the Infutor and Nielsen datasets correspond to 51% of total household spending.

2.2 Nielsen Homescan Data

Our data source for non-durable consumption is the Nielsen Homescan Data from the Kilts Marketing Data Center at the University of Chicago Booth School of Business. This dataset contains longitudinal panel data from approximately 40,000-60,000 U.S. households between 2004 and 2007. The data provide information about the products purchased by households, as well as when and where they were purchased. The Nielsen Homescan data covers the entire nation and includes a wide range of household locations and demographics. Overall, the dataset comprises purchases of almost 250 million unique items. One advantage of this dataset is that it records barcodes at a very fine level and the expenditure on each item, which allows us to construct multiple measures of non-durable consumption segregation.

2.3 American Community Survey

The ACS is a yearly survey conducted by the Census Bureau, which randomly samples individuals in each state, the District of Columbia, and Puerto Rico. We use information on income, education, age, race, and geography to validate the representativeness of the Infutor data and to measure residential segregation by race, income, age, and education, which are all important variables in our analysis but are imputed in the Infutor data. We limit the ACS sample to individuals between the ages of 22 and 80 and exclude observations with non-positive income. The income variable we focus on is total household income. Although the car and home characteristics from the Infutor data are at the individual level, we consider these consumption categories likely to be used by all household members.

3 Patterns of Consumption Segregation

In this section, we describe the patterns of consumption segregation across space, time, consumption categories, and socio-economic groups. We highlight four main findings. First, consumption segregation has remained relatively flat over the past 15 years. Second, there is a great degree of regional variation in consumption segregation. Third, consumption is more segregated in richer, younger, larger, more educated, and less white regions. Fourth, consumption segregation reflects to a large extent segregation within rather than across demographic groups.

3.1 Measuring Segregation

Throughout the paper, we consider two measures of segregation, depending on whether the variable we consider is categorical or continuous: the *entropy index* for categorical variables and the *rank-order index* for continuous variables. Both indices capture how representative is the distribution of a given economic outcome in a geographic sub-unit relative to a broader geographic unit. They are commonly employed in the literature on segregation and have a large number of desirable properties. For example, Reardon and Firebaugh (2002) evaluate multiple indices of segregation, including the dissimilarity index and the Gini index which are also commonly used, and conclude that the entropy index is the most conceptually and mathematically satisfactory along a wide range of considerations.[7] We next describe these two measures in more detail.

Entropy Index. We use the entropy index to measure segregation along categorical dimensions, such as car model, characteristics of houses, race, and categories of groceries.

To construct the entropy index for a broad geographic unit, we begin by describing the construction of the entropy score, a measure of diversity. The entropy score of a geographic unit i (e.g., i can be a street, Census tract, zip code, city, etc.) is

$$h_i = -\sum_{j=1}^{J} p_{ij} \ln\left(p_{ij}\right), \tag{1}$$

where J is the number of mutually exclusive groups (car or house characteristics in the case of the Infutor data) and p_{ij} is the share of the population of geographic unit i that belongs

[7] Reardon and Firebaugh (2002) evaluate segregation indices along seven criteria: organization equivalence, size invariance, the principle of transfers, the principle of exchanges, composition invariance, additive organizational decomposability, and additive group decomposability. We make use of some of these criteria in our analysis and refer the interested reader to Reardon and Firebaugh (2002) for a more detailed discussion.

Figure 1: Measuring Segregation: An Example

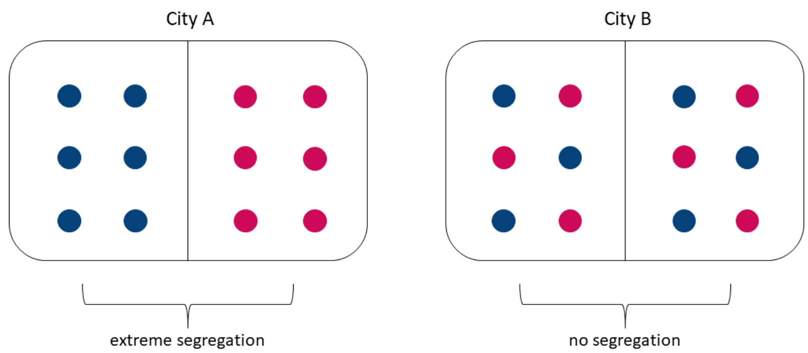

Note: Figure 1 illustrates with an example our measure of segregation. Cities A and B are divided into two neighborhoods: left and right. Each neighborhood contains individuals that are either blue or pink. In city A there is extreme segregation and in city B there is no segregation.

to group j. Specifically, $p_{ij} = \frac{n_{ij}}{n_i}$, where n_i is the population of the geographic unit i and n_{ij} is the population of the geographic unit i that belongs to group j. For example, if i is a zip code, then in our case n_i is the total number of cars in geographic unit i and n_{ij} is the total number of cars of model j in geographic unit i. The maximum value for h_i is $\ln(J)$. A high value of h_i indicates more diversity. The extreme case $h_i = 0$ indicates that the geographic unit i contains only one group.

Based on the measure above, one can calculate segregation for a broader geographic unit, such as a state. Specifically, the entropy index of a broader geographic unit is defined as

$$H = \frac{\hat{H} - \bar{H}}{\hat{H}}, \tag{2}$$

where $\hat{H}$ is the entropy index calculated at the level of a broader geographic unit using equation (1) and $\bar{H}$ is the population share weighted average of h_i, for all geographic units i belonging to the broader unit. The index H is bounded between 0 and 1. A value of $H = 1$ indicates extreme segregation, with each geographic unit i containing only one group (i.e., $\bar{H} = 0$), as in the left panel of Figure 1. A value of $H = 0$ indicates no segregation, with each geographic unit i being representative of the broader geographic unit (i.e., $\hat{H} = \bar{H}$), as in the right panel of Figure 1.

Rank-order Index. We use the rank-order information theory index to measure segregation along continuous dimensions such as income or expenditure (Reardon et al., 2006;

Reardon and Bischoff, 2011). The rank-order index of a continuous variable y is given by

$$H^R(y) = 2\int_0^1 \left\{ \sum_{i=1}^{I} \omega_i \left[\bar{F}_{i,p} \ln\left(\frac{\bar{F}_{i,p}}{p}\right) + (1-\bar{F}_{i,p}) \ln\left(\frac{1-\bar{F}_{i,p}}{1-p}\right) \right] \right\} \mathrm{d}p \tag{3}$$

where ω_i is the population share of geographic unit i, $\bar{F}_{i,p} = F_i(y(p))$ is the cumulative function in the geographic unit i evaluated at $y(p)$, where note that both y and the percentiles p correspond to the income distribution of a geographic unit that is broader than i. The rank-order index H^R is also bounded between 0 and 1, with $H^R = 1$ indicating extreme segregation and $H^R = 0$ indicating no segregation. We note that we obtain similar results if, instead, we measure the segregation of continuous variables by using the entropy index applied to bins of the continuous variable, such as income deciles.

Measuring Segregation in the Data. When constructing the moments of segregation reported in the paper, the geographic unit i is a PUMA (Public Use Microdata Area), which is the smallest unit of geography available across all the datasets we use. The broader unit of geography is either a CBSA (Core Based Statistical Area) or a US state.[8] In reporting trends, the broader unit of geography is the country.[9]

We exclude CBSAs that cover less than two PUMAs, as the entropy index calculation requires that the broader geographic unit contains at least two sub-units of geography i. Out of the 933 CBSAs initially considered, we drop 295 CBSAs that do not include any PUMA (very small CBSAs) and 422 CBSAs that map one-to-one into a PUMA. The remaining 216 CBSAs include, on average, 9 PUMAs and cover approximately 80% of the total population.

Since our consumption data contains information on both the dollar value of the goods consumed (e.g., dollars spent on groceries, the value of the car, the value of the home) and on the actual goods consumed (type of groceries purchased, type of car owned), we report statistics on both the segregation of spending (using the rank-order index) and the segregation

[8]Because CBSAs and states have different numbers of PUMAs, a natural concern that arises is that a higher number of PUMAs mechanically leads to a higher segregation index at the level of the CBSA or the state. We performed Monte Carlo simulations to show that this does not drive our results. In each simulation, we randomly partitioned states into geographic sub-units and computed the resulting segregation index. We did this using the ACS for categorical variables (race using the entropy index) and continuous variables (income using the rank-order index). We then regressed the entropy index on the number of sub-units and compared the distribution of slope coefficients across 500 simulations with the slope coefficient observed in the actual data. While, in general, we do find a positive relationship between the number of geographic sub-units and the segregation index, the empirical slope coefficient is much larger than the slope coefficients from the simulated data, suggesting that the regional differences in segregation that we report below are not driven by differences in the number of PUMAs across states.

[9]Table 14 in Appendix B shows that our results are robust to considering zip codes or counties as the geographic unit i.

of consumption (using the entropy index). We also report statistics on the segregation of durable and non-durable consumption, which are known to respond differently to economic shocks and policy (Berger and Vavra, 2014, 2015).

We aggregate the segregation indices for the different consumption categories using expenditure shares as weights. Using data from the CEX, we calculate that the spending share on rent and rent equivalent for homeowners is 32%, the spending share on food at home, alcohol at home, tobacco, personal care products, and other personal goods is 15%, and the spending share on new and used motor vehicles is 4%. These are the three largest components of the household consumption basket according to (Theloudis, 2020).[10]

3.2 Trends in Consumption Segregation

We begin by examining the trends in consumption segregation in the United States. Figure 2 depicts these trends since 2000. Panel (a) shows the evolution of three broad measures of consumption: total consumption, durable consumption, and non-durable consumption, measured in terms of expenditure. Two key observations can be made from Panel (a). First, durable consumption is more segregated than non-durable consumption. The average durable consumption entropy index is 25%, approximately 2.5 times larger than non-durable consumption segregation. This means that the PUMA in the United States is 25% less diverse than the country as a whole. We will return to this point shortly and show that the large segregation of durable consumption is mainly due to the segregation of housing consumption. Given the large expenditure share on durables, the segregation of total consumption is close to that of durables: the average total consumption entropy index is 20%. Second, the segregation of non-durable consumption has declined by approximately 25% since 2005, the beginning of our Nielsen Homescan sample, whereas the segregation of durable and total consumption exhibits no trend growth. Fluctuations in durable consumption segregation are cyclical, reflecting the boom-bust cycle in the housing market.

In Panels (b)-(d) of Figure 2, we zoom in on the segregation of consumption categories. Panel (b) depicts the evolution of non-durable consumption segregation measured both using data on non-service retail expenditure, as in Panel (a), as well as the specific product categories purchased. The average non-durable consumption expenditure segregation is higher than that measured using product categories that households purchase. Both measures of non-durable segregation are declining over time and are lower than the average consumption

[10] Our entropy indices are conditional on having positive expenditure on groceries, own at least a vehicle or a house.

segregation. The segregation of non-durable consumption at the level of product categories is non-zero, suggesting a role for supply-side channels stemming from product availability. That the level of segregation is low and declining is therefore consistent with Allcott et al. (2019b), who find that only 10% of food inequality between rich and poor household derives from supply-side channels such as food deserts. That the segregation of non-durable spending is higher suggests there are forces beyond product availability that shape the segregation of non-durable consumption. We investigate which are these forces below.

Panels (c) and (d) depict the evolution of housing and car consumption segregation between 2000 and 2018. The first striking feature that emerges is that housing consumption segregation is much higher than car consumption segregation. This result is true irrespective of whether we use house values or house characteristics to measure the entropy index over time. This result is robust also to using the value of homes directly reported values in the ACS and Census.[11]. Car consumption is, in fact, the least segregated dimension of consumption, whether we measure it using the values of the cars or car models. A second feature that emerges is that the segregation of housing consumption is quite stable over time. When it is measured in house values it mirrors the dynamics of house prices more generally, consistent with differential house price appreciation across US regions (Guren et al., 2020). In contrast, the segregation of car consumption displays a mild but steady upward trend.

3.3 The Geography of Consumption Segregation

In this section, we examine consumption segregation across U.S. regions. Figure 3 provides a visual representation of the geography of consumption segregation. Firstly, we compute entropy indices for each consumption category, where we assume the narrow geographic unit is a PUMA and the broader geographic unit is a U.S. state. Subsequently, we report in the figure, for each state, the average of the entropy index over the period 2016-2018.

Panel (a) of Figure 3 reveals substantial heterogeneity in how segregated total consumption is across U.S. states, with the most segregated state being 11 times more segregated than the least segregated state. The entropy index ranges from 28% in New York to 2.5% in Wyoming, indicating that the average PUMA in New York is almost 30% less diverse than the state, while in Wyoming, the average PUMA is almost representative of the state. Overall, total consumption is most segregated in the West Coast, Northeast, and Texas. Conversely, the Rocky Mountain Region shows the lowest levels of consumption segregation.

[11] See Figure 7 in Appendix B

Figure 2: Consumption Segregation Over Time

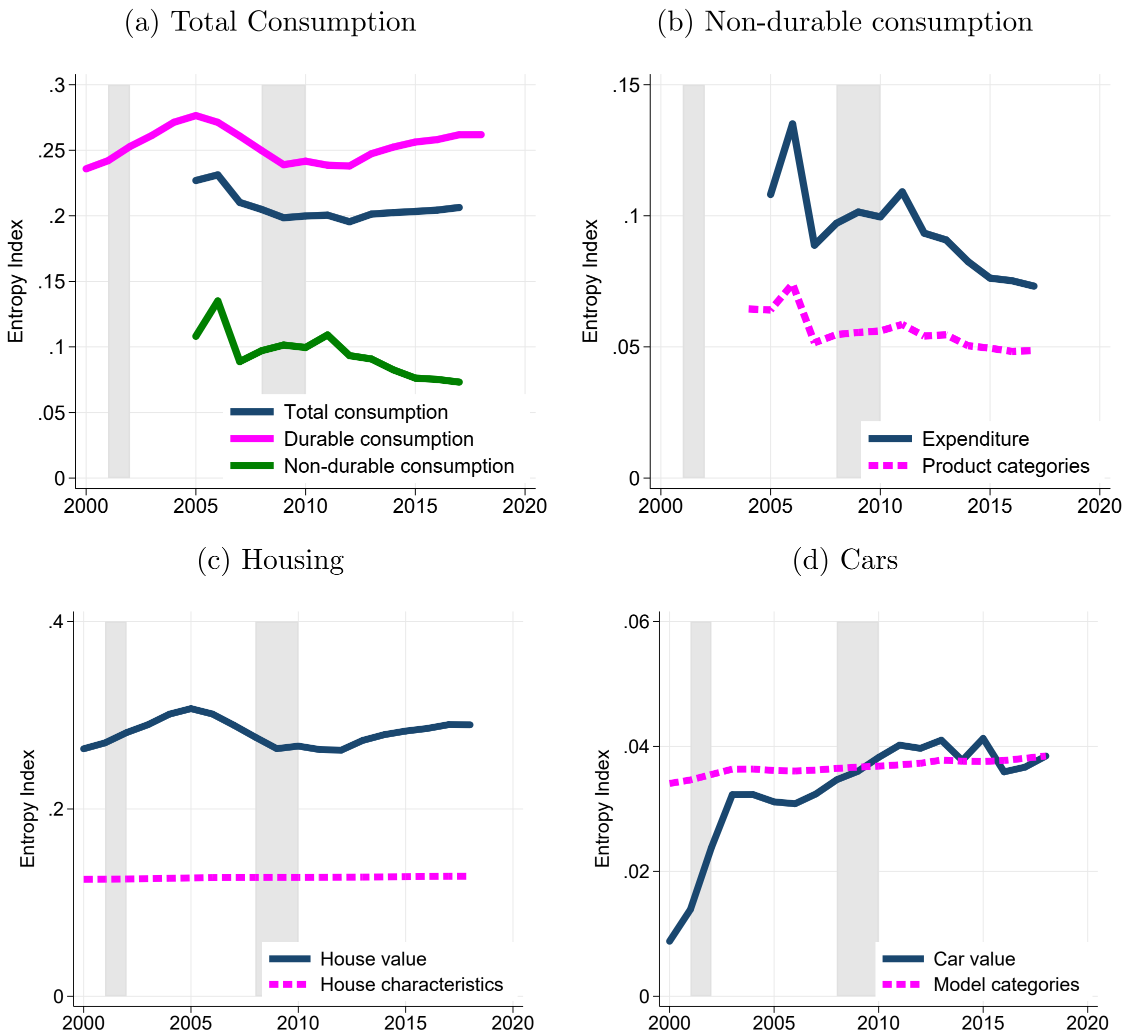

Panel (b) of Figure 3 depicts the geography of durable consumption segregation. The segregation of durable consumption displays similar patterns as that of total consumption, reflecting the large share of durable spending in total spending. Specifically, durable consumption is most segregated in New York, while Wyoming, South Dakota, and Alaska are the least segregated. Figure 4 zooms in on the sources of durable consumption segregation. Panel (a) of the figure shows that the segregation of durable consumption reflects, to a large extent, segregation in housing consumption. The segregation of vehicle consumption, shown in Panel (b) of the figure, exhibits slightly different geographic patterns, with higher segregation in the South and less segregation in the Northeast. Still, the level of segregation in this consumption category is generally much smaller and close to perfect diversity.

Figure 3: State-level Consumption Segregation

(a) Total consumption

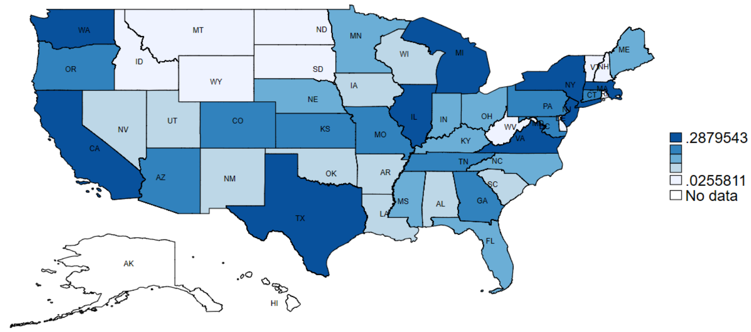

(b) Durable consumption

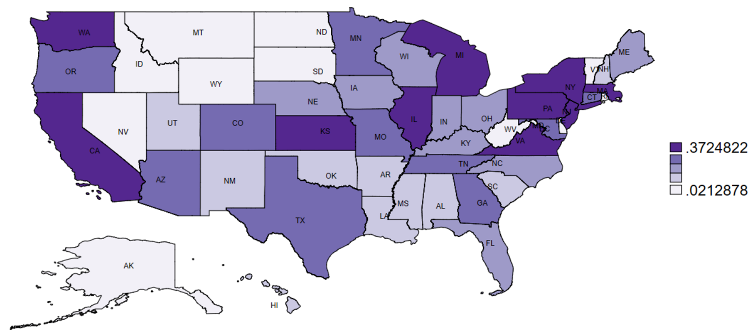

(c) Non-durable consumption

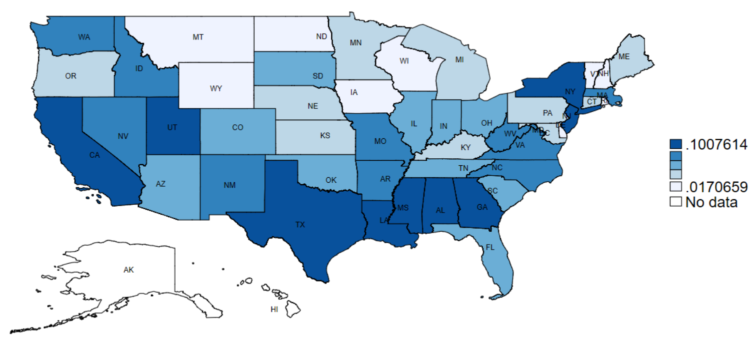

Notes: The figure plots state-level consumption entropy indices averaged over the period 2016-2018.

Turning to non-durable consumption, Panel (c) of Figure 3 shows that the geography of the segregation of non-durable consumption is different than that of durable consumption. Specifically, while New York and California continue to exhibit among the highest levels of segregation, the figure reveals very high levels of non-durable consumption segregation in the South. Texas has the highest level of non-durable consumption segregation, where the

Figure 4: State-level Durable Consumption Segregation

(a) Housing (b) Cars

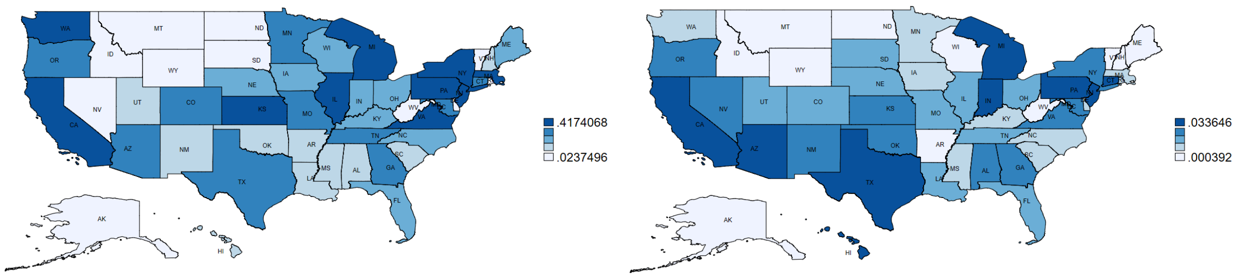

Notes: The figure plots state-level consumption entropy indices averaged over the period 2016-2018.

average PUMA is 10% less diverse than the state. At the other extreme, in Vermont, the average PUMA is nearly representative of the state.

We next investigate the locus of consumption segregation. To that end, we leverage the additive organizational decomposability of the entropy index. Specifically, as shown in Reardon and Firebaugh (2002), assuming that J geographic units are clustered in K clusters ($K < J$), the entropy index H can be rewritten as the sum of $K+1$ additive components

$$H = H_K + \sum_k \omega_k \frac{h_k}{h} H_k, \tag{4}$$

where ω_k is the population share of the cluster k, h_k is the entropy score of cluster k, h is the national-level entropy score, and H_k is the information theory index of cluster k. The term H_K represents the *between* cluster component, while the remaining K terms represent the portion of total segregation due to segregation *within* cluster. We derive a similar decomposition for the rank-order index H^R. Specifically,

$$H^R = H^R_K + \sum_k \omega_k H^R_k, \tag{5}$$

where, relative to equation (4), the term corresponding to $\frac{h_k}{h}$ is equal to 1 in the case of the rank-order index. The share of total segregation due to the segregation within cluster k represents the amount by which total segregation would fall if segregation within cluster k were eliminated by rearranging individuals among its geographic units while leaving all other geographic units unchanged.

We report the resulting decomposition in Table 4. The top panel of the table shows that two-thirds of consumption segregation occurs within states. That is, total consumption segregation would decrease by approximately 70% if segregation within each state was eliminated by appropriately rearranging individuals within the PUMAs of each state. The second and third columns of the table decompose the segregation of durable and non-durable consumption into within and between state components. While it is the case that most consumption segregation is accounted for by segregation within states, even for the two broad consumption categories, we note that in the case of non-durable consumption, nearly all segregation is explained by the within-state component. Finally, the last two columns of the table provide further detail into the segregation of durable consumption and show that the fact that most durable consumption segregation is attributable to the within-state component is a reflection of housing consumption, which is more segregated within than across states, while segregation of vehicle consumption reflects a very large extent segregation between states.

The bottom panel of Table 4 performs this decomposition assuming that clusters are CBSAs. Given that CBSAs are geographic units smaller than US states and that the within/between state decomposition of segregation reveals a large role for segregation within states, the decomposition of segregation within and between CBSAs offers further detail on the locus of the within-state segregation. We find nearly equal contributions of the within and between CBSA components in explaining total consumption segregation, as well as segregation of durable and non-durable consumption.

Table 4: Consumption Segregation Within and Between Regions

	Total consumption	Durable consumption	Non-durable consumption	Housing	Cars
			State-level segregation		
Between states	0.285	0.392	0.036	0.350	0.727
Within state	0.715	0.608	0.964	0.650	0.273
			CBSA-level segregation		
Between CBSAs	0.444	0.569	0.557	0.670	0.166
Within CBSA	0.556	0.431	0.443	0.330	0.834

Notes: The table reports averages of the between and within region components of the consumption entropy indices over the period 2016-2018.

Table 5 investigates the relationship between consumption segregation and socioeconomic characteristics of CBSAs, building on the observed heterogeneity in segregation across the US. We calculate the average entropy index for total, durable, and non-durable consumption for CBSAs in the top 25% or bottom 25% of the national distribution of income, education, race, age, and population size. The results show that consumption segregation is higher in CBSAs that are richer, younger, larger, and have a higher share of college-educated individuals and a lower share of whites. For instance, total consumption is 2.5 times more segregated in CBSAs in the top 25% of the income distribution than those in the bottom 25%. The relative segregation between rich and poor CBSAs persists across durable and non-durable consumption categories, with durable consumption being 2.7 times more segregated and non-durable consumption 1.9 times more segregated in rich CBSAs. The findings for education closely mirror those for income. The results for race show that total consumption is twice as segregated in CBSAs in the bottom 25% of the white population share. The heterogeneity is more pronounced for durable consumption than for non-durable consumption, while differences by age are minimal. Finally, consumption segregation is 4.3 times higher in CBSAs in the top 25% of the population distribution. These results suggest that consumption segregation varies considerably across core demographic characteristics.

We now delve into a more detailed analysis of how consumption segregation varies with individuals' demographic characteristics. In Table 6, we present consumption segregation for different demographic groups. The top panel of the table divides individuals in our sample into three racial groups and calculates the entropy index H^R defined in equation (3) for each group. We observe little variation in total consumption segregation across racial groups, mainly due to a similar level of durable consumption segregation. However, we note more prominent differences for non-durable consumption. Specifically, white individuals exhibit the highest segregation of non-durable consumption, similar to that of Black individuals, but 37% higher than that of other racial groups. The middle panel of the table displays the average consumption segregation for college-educated and non-college-educated individuals. Total consumption is 13% more segregated among the college-educated, with an 11% higher durable consumption segregation and a 19% higher non-durable consumption segregation among this group.

These findings suggest a minor role for differences across demographic groups regarding consumption segregation. To investigate this further, we decompose consumption segregation into within and between demographic group components. We examine both the broad demographic groups defined in Table 6 and narrow demographic groups defined as (race, education, age) tuples.[12] We employ a decomposition that is analogous to that described in equation (5) and utilize the additive group decomposability property of the entropy and

[12]We consider all combinations of three racial groups (white, Black, other), two education groups (with a college degree, without a college degree) and two age groups (younger than 35, older than 35).

Table 5: Consumption Segregation and Regional Characteristics

	Total consumption	Durable consumption	Non-durable consumption
	By income		
Bottom 25%	0.058	0.063	0.040
Top 25%	0.144	0.169	0.075
	By education (college share)		
Bottom 25%	0.063	0.068	0.044
Top 25%	0.134	0.154	0.075
	By race (share white)		
Bottom 25%	0.136	0.152	0.087
Top 25%	0.060	0.069	0.036
	By age (share $\leq$ 35)		
Bottom 25%	0.076	0.088	0.046
Top 25%	0.103	0.110	0.072
	By population		
Bottom 25%	0.028	0.032	0.018
Top 25%	0.121	0.136	0.076

Notes: The table reports population-weighted averages of CBSA-level consumption entropy indices over the period 2016-2018.

rank-rank order indices (Reardon and Firebaugh, 2002).[13] Table 7 summarizes the results of this decomposition and indicates that the majority of the segregation in consumption is due to the within demographic group component. Consequently, there is minimal influence of differences in segregation across demographic groups in understanding overall consumption segregation patterns. These results imply that geographic variations, particularly with respect to the income and education of regions, hold greater explanatory power in explaining consumption segregation.

[13]For the purpose of this decomposition, the geographic unit i in equation (1) is a combination between a PUMA and a demographic group.

Table 6: Consumption Segregation by Demographic Group

	Total consumption	Durable consumption	Non-durable consumption
		By race	
White	0.102	0.111	0.071
Black	0.108	0.116	0.068
Other	0.115	0.124	0.052
		By education	
No college degree	0.102	0.107	0.079
College degree	0.115	0.119	0.094
		By age	
Younger than 35	0.104	0.118	0.058
Older than 35	0.104	0.117	0.064

Notes: The table reports population-weighted averages of CBSA-level consumption entropy indices by demographic groups over the period 2016-2018.

4 What Drives Consumption Segregation?

In this section, we examine potential drivers of consumption segregation. We document that differences in consumption segregation across space and time are driven mostly by differences in income segregation, reflecting households' inability to arbitrage away income differences through financial markets.

4.1 Drivers of Consumption Segregation

A long line of work studies segregation in the United States over the past century along various socio-economic dimensions (race, education, income, ethnicity), focusing on racial and income segregation.[14] We next analyze to what extent residential segregation along these previously documented margins is a driver of consumption segregation, thus providing a bridge between previously studied dimensions of residential segregation and the residential segregation of welfare. To that end, we project consumption segregation on other dimension of residential segregation and estimate the following regression specification

$$\ln H_{C,rt} = \alpha_0 + \alpha_1 \ln H_{Y,rt} + \alpha_2 \ln H_{R,rt} + \alpha_3 \ln H_{E,rt} + \alpha_4 \ln H_{A,rt} + \boldsymbol{\gamma}\mathbf{X}_{rt} + \delta_t + \varepsilon_{rt}, \quad (6)$$

[14] See Trounstine (2018) for a comprehensive summary of this work.

Table 7: Consumption Segregation Within and Between Demographic Groups

	Total consumption	Durable consumption	Non-durable consumption
		Race	
Between group	0.033	0.009	0.091
Within group	0.967	0.991	0.909
		Education	
Between group	0.028	0.034	0.012
Within group	0.972	0.966	0.988
		Age	
Between group	0.014	0.007	0.029
Within group	0.986	0.993	0.971
		Race × Education × Age	
Between group	0.033	0.042	0.007
Within group	0.967	0.958	0.993

Notes: The table reports averages of the between and within demographic group components of the CBSA-level consumption entropy indices by demographic groups over the period 2016-2018.

where $H_{C,rt}$, $H_{Y,rt}$, $H_{R,rt}$, $H_{E,rt}$, and $H_{A,rt}$ denote the entropy indices of consumption, income, race, education, and age in CBSA r and year t, respectively, $\mathbf{X}_{rt}$ is a vector of time-varying controls for CBSA r, and δ_t are year fixed effects. The measures of consumption and income segregation are calculated as described in equation (3), while the measures of racial, education and age segregation are calculated as described in equation (2), considering the same racial, education and age groups as in Section 3.3. Motivated by the analysis in Section 3.3, the vector of controls includes average income in CBSA r in year t, as well as the population share of whites, Blacks, college educated and younger than 35.

Table 8 reports the results of the estimation, sequentially augmenting the set of controls for segregation.[15] Specifically, the specification in column 1 only controls for segregation by income, in addition to the set of controls $\mathbf{X}_{rt}$ and the time fixed effects. Column 2 adds controls for racial segregation and column 3 adds controls for segregation by education and age. Focusing on the specification in column 1 for total consumption we find an elasticity of total

[15] See Table 17 in Appendix B for the corresponding results exploiting variation across U.S. states.

consumption segregation to income segregation of 0.67. That is, 67% of income segregation translates into consumption segregation. The R^2 we obtain from estimating equation (6) without any residential segregation controls is 0.41. Comparing this with the R^2 reported in column 1 suggests that differences in income segregation across CBSAs explain 30% of the differences in consumption segregation. In column 2, we include in the specification a control for racial segregation and find that while the estimated elasticity of consumption segregation with respect to racial segregation is also positive, 0.14, it is four times smaller than that with respect to income segregation. Additionally, the R^2 only increases by two percentage points. In column 3, we add to the specification controls for segregation by education and by age. We continue to find that income segregation is the main driver of consumption segregation. We also find a sizable role for segregation by education, but note that segregation by income and by education are strongly correlated. Specifically, the correlation coefficient between the two is 0.77. We interpret the two as broadly representing economic status. Segregation by education and by age explains an additional 4% of the difference in consumption segregation across CBSAs, further strengthening the role of differences in income segregation as the main driver of differences in consumption segregation.

That income segregation emerges as the main driver of consumption segregation is in line with the findings of the existing literature on residential segregation and the standard consumption theory. Specifically, Massey et al. (2009) document that during the last third of the twentieth century, the United States moved from a regime of segregation based on race and ethnicity to a regime of segregation based on economic status, where economic status is given by income and education. Standard consumption theory predicts that differences in income across individuals are expected to translate into differences in consumption if income differences are persistent and/or if credit market frictions or another form of market incompleteness prevent individuals to insure against income fluctuations. Extended to a spatial dimension, as in Giannone et al. (2019), this argument suggests that, under the same conditions, residential segregation by income will translate into residential segregation by consumption.

Turning to the drivers of durable and non-durable consumption segregation, we draw a similar conclusion. Specifically, we find that in the most conservative specification 43% of income segregation translates into durable consumption segregation. This elasticity is approximately half for non-durable consumption: only 24% of income segregation translates into non-durable consumption segregation. Relative to the R^2 from a regression that ab-

Table 8: Drivers of Consumption Segregation

	Total consumption			Durable consumption			Non-durable consumption		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
$\ln H_Y$	0.67*** (0.03)	0.58*** (0.04)	0.39*** (0.06)	0.87*** (0.05)	0.73*** (0.05)	0.43*** (0.07)	0.44*** (0.03)	0.35*** (0.04)	0.24*** (0.05)
$\ln H_R$		0.14*** (0.05)	0.10** (0.05)		0.21*** (0.07)	0.16*** (0.06)		0.14*** (0.04)	0.12*** (0.03)
$\ln H_E$			0.21*** (0.05)			0.33*** (0.06)			0.11*** (0.03)
$\ln H_A$			0.02 (0.02)			0.01 (0.03)			0.04 (0.03)
R^2	0.72	0.74	0.78	0.65	0.66	0.73	0.57	0.59	0.61

Notes: The table reports the estimates of α_1, α_2, α_3 and α_4 in equation (6). The estimate of α_0, the time fixed effects and the estimate of the vector $\boldsymbol{\gamma}$ are omitted. Robust standard errors are reported in parentheses. The regression is estimated with population weights based on the 2010 Census. ***, **, and *, represent statistical significance at 1%, 5% and 10%, respectively.

stracts from any controls from segregation, we find that differences in income segregation across CBSAs account for 33% of the differences in durable consumption segregation, while segregation by race, education and by age account for an additional 8%. Differences in non-durable consumption segregation across CBSAs are accounted for in proportion of 22% by differences in income segregation and by an additional 4% by differences in segregation by race, by education and by age.[16]

4.2 Inspecting the Relationship Between Income and Consumption Segregation

Motivated by the results above, in this section we inspect what is the mechanism driving the strong positive relationship between income and consumption segregation. To that end, we scrutinize the two potential explanations in standard consumption theory: market incompleteness and persistent income differences between individuals.[17] Market incompleteness or the lack of access to some financial products (either because of these products being absent or because of underlying informational frictions) and persistent income shocks both have the

[16] The R^2 from regressing durable consumption segregation on the controls in equation (6), excluding the segregation controls, is 0.32. The corresponding R^2 for non-durable consumption segregation is 0.35.

[17] We conjecture that our results are amenable to persistent preference differences between individuals.

potential to generate such a positive relationship since they prevent consumption insurance among individuals with different income.

In Appendix B.2, we outline a simple model with borrowing constraints and heterogeneous agents that guides the empirical analysis. Based on that model, if income differences are transitory, the only scenario under which they translate into consumption differences is if markets are incomplete. In the extreme case of autarky, they fully translate into consumption differences. If instead, income differences are persistent, then they are not insurable and translate to a large extent into consumption differences. In the extreme case of permanent differences, they fully translate into consumption differences.

We propose an empirical test to assess whether the correlation we observe in the data between income and consumption segregation reflects permanent income differences between socio-economic groups or incomplete asset markets that prevent consumption insurance against potentially insurable income shocks. To that end, we decompose the mapping between income and consumption segregation into two components meant to capture how income segregation translates into within- and between-group consumption segregation. The rationale behind this decomposition follows Krueger and Perri (2006) and is based on the assumption that income differences between socio-economic groups can be attributed to fixed characteristics of households. Since it is likely very hard to insure against these differences, an increase in income segregation should translate into an increase in between-group consumption segregation. Within a socio-economic group, the assumption is that income differences arise as a consequence of idiosyncratic income shocks. If there exist financial markets that can facilitate partial insurance against these shocks, then income segregation should translate into within-group consumption segregation only to a limited extent.[18]

Recall that in Section 3.3 we decomposed consumption segregation into a within demographic group and a between demographic group component, where a demographic group is defined as a $(race, education, age)$ tuple. Letting $H_{C,rt}$ denote the level of consumption segregation in region r in year t, the decomposition allows us to write

$$H_{C,rt} = H^B_{C,rt} + H^W_{C,rt},$$

where $H^B_{C,rt}$ denotes the between-group component of consumption segregation and $H^W_{C,rt}$ denotes the within-group component. We then investigate the relative extent to which in-

[18]An extreme example is that of hand-to-mouth consumers, who cannot insure at all against transitory income shocks. For such consumers, income segregation translates one-to-one into consumption segregation. At the other extreme, when markets are complete, income segregation does not translate at all into consumption segregation.

come segregation affects these two components of consumption segregation by estimating the following regression specification

$$H^i_{C,rt} = \alpha^i_0 + \alpha^i_1 \; H_{Y,rt} + \boldsymbol{\gamma^i}\mathbf{X}_{rt} + \delta^i_t + \varepsilon^i_{rt},$$

where $i \in \{B, W\}$ and $\mathbf{X}_{rt}$ is the same vector of controls as in equation (6). The coefficients α^B_1 and α^W_1 can be shown to sum up to the coefficient α_1 from estimating

$$H_{C,rt} = \alpha_0 + \alpha_1 \; H_{Y,rt} + \boldsymbol{\gamma}\mathbf{X}_{rt} + \delta_t + \varepsilon_{rt}.$$

Therefore, the share $\frac{\alpha^B_1}{\alpha_1}$ is informative of the relative importance of permanent income differences in explaining consumption segregation, while $\frac{\alpha^W_1}{\alpha_1}$ is indicative of the importance of market incompleteness.

Table 9 presents estimates of α_1, α^B_1 and α^W_1 for total consumption and consumption sub-categories. The between-group component accounts for 14% of the relationship between income and total consumption segregation at the CBSA-level, while the within-group component accounts for the majority of this relationship at 86%. This pattern holds true for broad consumption categories, durable and non-durable consumption, as well as for specific durable consumption categories. The within-group component is particularly important in explaining the relationship between income segregation and the segregation of durable consumption, especially housing, which is consistent with market frictions that limit households' ability to smooth consumption fluctuations (Boar et al., 2020). Specifically, 90% of the mapping between income segregation and housing consumption segregation is accounted for by the within-group component, while in the case of non-durable consumption, this component explains 80% of the relationship.

Overall, we conclude that while there is a role for permanent income differences in explaining why consumption is not equalized across space, market incompleteness plays a major role in explaining why income segregation translates into consumption segregation.

5 A Theory of Segregation and Inequality

We propose a simple mechanism through which income segregation that translates into consumption segregation can lead to higher wealth inequality in the presence of social externalities that result in households valuing their social status relative to their neighbors. When segregation is low, poor households are motivated to save to purchase status-enhancing consumption goods, which enables them to *keep up* with their wealthier neighbors. As a result,

Table 9: Consumption and Income Segregation

	Total	Between	Within
Total consumption	1.52	0.21	1.32
α_1^i/α_1		13.8%	86.2%
Non-durable consumption	0.64	0.13	0.51
α_1^i/α_1		20.3%	79.7%
Durable Consumption	1.94	0.16	1.76
α_1^i/α_1		9.7%	90.3%
Cars	0.17	0.03	0.14
α_1^i/α_1		17.6%	83.4%
Housing	2.16	0.21	1.94
α_1^i/α_1		9.7%	90.3%

Notes: The table reports the estimates of α_1, α_1^B and α_1^W. All estimates are significant at 1% significance. Robust standard errors, time fixed effects and the estimates of $\boldsymbol{\gamma^i}$ and $\boldsymbol{\gamma}$ are omitted but are included in the estimation. The regression is estimated with population weights based on the 2010 Census, using data for the 2016-2018 period.

the wealth gap between the poor and rich is reduced. To outline this mechanism, we propose a theoretical framework based on Drechsel-Grau and Greimel (2018). We have kept the framework simple for clarity, but it can be generalized for quantitative exploration.[19]

5.1 Environment

We consider the utility maximization problem of a household of type $i \in \{P, R\}$, who can be either income poor (P) or income rich (R) and who lives in a location j. Each period the household earns constant income y^i, with $y^R > y^P$, and chooses how to optimally divide it between spending on non-durable consumption, spending on a durable and visible consumption good (e.g. housing, car, jewelry, etc.) and saving.[20] We purposefully assume income to be the same across locations not to confound the mechanism we explore with differences in access to opportunity that arise as a consequence of earnings differentials.

Each period, the household derives utility from non-durable consumption and social sta-

[19]See Drechsel-Grau and Greimel (2018) for a generalization without regions.

[20]Our assumption of permanent income differences is for analytical tractability only. At the sacrifice of analytical tractability, one could obtain a similar mechanism as described here arising from partially uninsurable transitory income differences, as suggested by our findings in Section 4.

tus, which depends on the household's stock of visible and durable goods and a reference consumption measure assumed to be the average stock of richer agents' visible and durable goods in location j. We note two things. First, that social status is related to consuming visible and durable goods aligns with the Charles et al. (2009) observation that conspicuous consumption items include jewelry, clothes, and cars, which are all durables. Second, for the purpose of analytical tractability, we assume that richer households determine the reference group instead of the average household. Given our constant income assumption across locations, the household's choices of non-durable consumption, visible (and durable) consumption and saving will depend on the location j only through the effect that the population composition of the location has on these choices, thus isolating the mechanism we want to highlight. We let c^i_{jt}, h^i_{jt} and a^i_{jt+1} denote these choices, respectively and we let $s\left(h^i_{jt}, \bar{h}^i_{jt}\right)$ denote social status.

We assume that social status is given by

$$s\left(h^i_{jt}, \bar{h}^i_{jt}\right) = h^i_{jt} - \bar{h}^i_{jt},$$

and is therefore increasing in the household's stock of the visible good and decreasing in the reference measure $\bar{h}^i_{jt}$ which, in our simple example with two types of agents, is the stock of the visible consumption good of the richer agents in the location j. This implies that poor households compare themselves with rich households and rich households do not compare themselves with anyone.[21]

The objective of the household is to maximize lifetime utility

$$\sum_{t=0}^{\infty} \beta^t \frac{\left((1-\xi)\left(c^i_{jt}\right)^{\varepsilon} + \xi s\left(h^i_{jt}, \bar{h}^i_{jt}\right)^{\varepsilon}\right)^{\frac{1-\sigma}{\varepsilon}}}{1-\sigma},$$

where β is the discount factor, ξ is the relative weight on social status in the utility function, $\frac{1}{1-\varepsilon}$ is the elasticity of substitution between non-durable consumption and social status, and σ is the coefficient of relative risk aversion. We assume that the visible and durable consumption good depreciates at rate δ and that maintaining it involves expenditure x^i_{jt}.[22] The stock of durable goods, therefore, evolves according to

$$h^i_{jt+1} = (1-\delta)\, h^i_{jt} + x^i_{jt},$$

[21] This could be generalized by assuming that the reference measure is $\bar{h}^P_{jt} = g_{PR} h^R_{jt}$ for poor households and $\bar{h}^R_{jt} = g_{RP} h^P_{jt}$ for the rich, where g_{ik} is the weight that the household of type i places on the household of type k. Here, to keep the notation as simple as possible, we simply assume that $g_{PR} > 0$ and $g_{RP} = 0$.

[22] If h^i_{jt} is thought of as housing, then x^i_{jt} represents home improvements and maintenance expenses.

and the budget constraint of the household is

$$c_{jt}^i + x_{jt}^i + a_{jt+1}^i = y^i + (1+r)\, a_{jt}^i,$$

where r is the real interest rate on savings. Lastly, we assume the household is subject to the borrowing constraints $a_{jt+1}^i \geq \underline{a}$ and $h_{jt}^i > 0$.

5.2 Optimal Choices

We next discuss how social externalities affect the consumption and saving decisions of households. The patterns uncovered here are key for establishing the link between income, consumption segregation and wealth inequality.

We show in Appendix C that the choices of the durable good h_j^R and h_j^P in location j are constant over time and given by

$$h_j^R = \frac{1}{r+\delta+\kappa} Y^R, \qquad \text{and} \qquad h_j^P = \frac{1}{r+\delta+\kappa} Y^P + \underbrace{\frac{\kappa}{r+\delta+\kappa} h_j^R \mu_j^R}_{>0\text{, keeping up effect}},\ \forall t,$$

where $\kappa = \left(\frac{(r+\delta)(1-\xi)}{\xi}\right)^{\frac{1}{1-\varepsilon}} > 0$ is a constant that depends on the parameters of the model, μ_j is the measure of rich households in location j, and Y^i is the permanent income of a household of type i and is equal to the sum of the initial asset position and the present discounted value of the earnings stream. Both types of households' stock of the visible consumption good is increasing in the household's own permanent income. Importantly, for poor households it is also increasing in the stock of the visible consumption good of the rich as captured by the positive second term, which we call the *keeping up effect*. This reflects the poor households' desire to keep up with their richer neighbors to increase their social status.

In order to keep up with the visible consumption of their richer neighbors, poor households save more. In particular, the asset position ra_j^i of the two types of households living in location j is[23]

$$ra_j^R = \frac{\kappa+\delta}{r+\delta+\kappa} Y^R - y^R \tag{7}$$

$$ra_j^P = \frac{\kappa+\delta}{r+\delta+\kappa} Y^P - \frac{\kappa}{r+\delta+\kappa}\left(\frac{\kappa+\delta}{r+\delta+\kappa} - 1\right) Y^R \mu_j^R - y^P \tag{8}$$

Equation (8) reveals that in the presence of social comparisons, having more rich neighbors induces the poor to save more than otherwise. To see why, note that $\frac{\kappa+\delta}{r+\delta+\kappa} - 1 < 0$, so the

[23]See Appendix C for the derivation.

second term of equation (8) is positive. For illustration, in the extreme case where there are no rich households in location j, the asset position of a poor household in that location is $\frac{\kappa+\delta}{r+\delta+\kappa}Y^P - y^p$ and is strictly smaller than that in equation (8). As we show below, this intuition is key for establishing the link between segregation and wealth inequality.

5.3 Segregation and Wealth Inequality

We now turn to discussing the effect of segregation on wealth inequality in the presence of social comparisons. We illustrate this relationship in the context of an abstract broad geographical unit J with two locations j_1 and j_2, populated by rich and poor households. For example, in line with the level of aggregation in our data, J could be thought of as a state or a CBSA and j_1 and j_2 as PUMAs. We assume, without loss of generality, that the mass of rich and poor households is $\mu_J^R = \mu_J^P = 1$ and define wealth inequality as the difference in the asset position between rich and poor households

$$\Delta_W = \underbrace{ra_{j_1}^R \mu_{j_1}^R + ra_{j_2}^R \mu_{j_2}^R}_{\equiv \bar{W}^R} - \underbrace{\left(ra_{j_1}^P \mu_{j_1}^P + ra_{j_2}^P \mu_{j_2}^P\right)}_{\equiv \bar{W}^P},$$

where $\mu_{j_1}^i + \mu_{j_2}^i = \mu_J^i = 1$, $i \in \{P, R\}$.

Extreme segregation. Consider first the case of *extreme segregation*, which is a situation in which all poor households live in location j_1 and all rich households live in location j_2. Then, within location $j \in \{j_1, j_2\}$ nobody makes social comparisons and the average wealth position of poor and rich households in J is

$$\bar{W}^P = \frac{\kappa+\delta}{r+\delta+\kappa}Y^P - y^P \qquad \text{and} \qquad \bar{W}^R = \frac{\kappa+\delta}{r+\delta+\kappa}Y^R - y^R,$$

respectively. Wealth inequality is therefore given by

$$\Delta_W^{extreme\ segregation} = \frac{\kappa+\delta}{r+\delta+\kappa}\left(Y^R - Y^P\right) - \left(y^R - y^P\right) > 0.$$

No segregation. Consider next the case of *no segregation*, that is a scenario in which rich and poor households are equally distributed across locations j_1 and j_2, i.e. $\mu_{j_1}^R = \mu_{j_1}^P = \mu_{j_2}^R = \mu_{j_2}^P = 0.5$. In this case, poor households in each location compare their consumption of the visible good with that of rich households which, as discussed above, leads them to save more relative to a counterfactual poor household that does not make such social comparisons.

The average wealth of the poor and the rich is then equal to

$$\bar{W}^P = \frac{\kappa+\delta}{r+\delta+\kappa} Y^P - y^P - \frac{\kappa}{r+\delta+\kappa}\left(\frac{\kappa+\delta}{r+\delta+\kappa} - 1\right) Y^R \qquad \text{and} \qquad \bar{W}^R = \frac{\kappa+\delta}{r+\delta+\kappa} Y^R - y^R,$$

respectively. Consequently, wealth inequality in this case can be shown to equal

$$\Delta_W^{no\ segregation} = \Delta_W^{extreme\ segregation} + \frac{\kappa}{r+\delta+\kappa}\left(\frac{\kappa+\delta}{r+\delta+\kappa} - 1\right) Y^R.$$

Because $\frac{\kappa+\delta}{r+\delta+\kappa} - 1 < 0$, the second terms is negative and

$$\Delta_W^{no\ segregation} < \Delta_W^{extreme\ segregation}.$$

That is, wealth inequality is increasing in income segregation. If income differences are permanent, or if markets are incomplete, income segregation translates into consumption segregation.[24] In the presence of externalities that create a motive for *keeping up with the Joneses*, when the poor have rich neighbors, as it is the case in less segregated locations, they increase their stock of visible consumption in order to keep up with the visible consumption of the richer neighbors. This requires poor households to save in order to afford to purchase the visible consumption goods, thereby reducing the wealth gap relative to the rich.

5.4 Empirical Evidence

The simple framework above predicts a positive relationship between income segregation, consumption segregation and wealth inequality. We next attempt to provide some empirical support for this prediction by exploiting regional variation income and consumption segregation and wealth inequality.

Measuring wealth inequality in the United States is challenging even at the aggregate level (Saez and Zucman, 2016, Smith et al., 2021), so regional measures of wealth inequality are scarce. We use two imperfect measures and show they have similar implications. The first is a state-level measure of the share of total wealth held by those with net worth higher than \$30 million. This statistic is calculated by the Institute on Taxation and Economic Policy (ITEP) using data from the ITEP Microsimulation Tax Model, survey data from the Survey of Consumer Finances (SCF) and information on billionaires in the United States published by Forbes.[25] The second is a state-level measure of the wealth differential between rich and

[24] In the previous section, we found evidence for both channels, with the latter being more important quantitatively. The framework we presented in this section focuses on the former for analytical tractability.

[25] See https://itep.org/the-geographic-distribution-of-extreme-wealth-in-the-u-s/ for more details.

poor households, similar to the definition of the wealth gap Δ_W above. We calculate this statistic using data from the Panel Study of Income Dynamics (PSID) and follow Aguiar et al. (2020) in our sample selection and definition of wealth.[26] We define the wealth gap as the difference between the average wealth of households with real disposable income above the median and the average wealth of those with real disposable income below the median. We calculate the median disposable income in each year of the sample (1999-2019, bi-annual) and average (real) wealth across all years before computing the wealth gap.

Table 10 reports pairwise correlation coefficients between state-level indices of income and consumption segregation, H_Y and H_C, and the two measures of wealth inequality. Income and consumption segregation are positively correlated, in line with the results in Section 4. Importantly, income and consumption segregation are positively correlated with the two measures of wealth inequality, in line with the predictions of the simple model above.

Table 10: Correlation between Segregation and Wealth Inequality

	H_Y	Top W share	Δ_W
H_C	0.86	0.72	0.34
H_Y		0.46	0.32
Top W share			0.31

Notes:The table reports correlation coefficients between state-level entropy indices for income and consumption and two measures of wealth inequality: the share of wealth held by those with wealth above $30 million (Top W share) and the wealth gap between households with income above and below the median (Δ_W).

6 Conclusions

We leverage a new data source in conjunction with existing data to examine consumption segregation in the United States and to investigate its underlying factors and consequences. We analyze how consumption segregation varies across time, regions, consumption categories, and demographic groups. Our findings reveal that, unlike income segregation, consumption segregation has remained relatively stable over the past two decades. However, there is significant spatial heterogeneity, with New York, the most segregated state, exhibiting 11

[26] Specifically, wealth is the sum of net liquid and illiquid assets. Net liquid assets are the sum of checking and savings balances, money market funds, certificates of deposit, treasury bills, and stocks outside of pension funds, net of debts in the forms of credit and store cards, student loans, medical or legal debt, and debt owed to the family. Net illiquid assets are the value of the home and other real estate equity, IRA/pension holdings, non-government bonds, insurance equity, and the net value of any business, farm, or vehicle.

times more consumption segregation than Wyoming, the least segregated state. Additionally, consumption segregation is more pronounced in richer, more educated and larger regions, as well as in those where a smaller percentage of the population is white. Interestingly, consumption segregation is not primarily driven by various demographic groups consuming differently in distinct locations, but instead is a shared characteristic across groups.

We examine the drivers of consumption segregation and find that income segregation is the main factor, even when accounting for previously documented dimensions of segregation such as race, education, and age. We also find a role, albeit more muted, for racial segregation. The high correlation between income and consumption segregation is due to incomplete markets that prevent consumption insurance. In light of this, we propose a simple mechanism that suggests consumption segregation could worsen wealth inequality by decreasing the motivation for low-income households to save and narrow the wealth gap.

Our work aims to provide researchers with measures of consumption differences across the US by item, over time and demographic groups. We hope this encourages research on the causes and consequences of regional differences in consumption and welfare, more broadly. Most studies that analyze socio-economic disparities consider income as the main factor driving location decisions. Our results suggest that individuals might also choose where to live based on the consumption of others. The interplay between income, consumption, migration and geographic disparities is therefore a natural avenue for future research.

References

Agarwal, Sumit, J. Bradford Jensen, and Ferdinando Monte, "Consumer Mobility and the Local Structure of Consumption Industries," Working Paper 23616, National Bureau of Economic Research July 2017.

Aguiar, Mark A, Mark Bils, and Corina Boar, "Who Are the Hand-to-Mouth?," Technical Report, National Bureau of Economic Research 2020.

Aguiar, Mark and Erik Hurst, "Measuring trends in leisure: The allocation of time over five decades," *The Quarterly Journal of Economics*, 2007, *122* (3), 969–1006.

_ **and Mark Bils**, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, September 2015, *105* (9), 2725–56.

Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell, "Food deserts and the causes of nutritional inequality," *The Quarterly Journal of Economics*, 2019, *134* (4), 1793–1844.

_ , _ , Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell, "Food Deserts and the Causes of Nutritional Inequality," *The Quarterly Journal of Economics*, 05 2019, *134* (4), 1793–1844.

Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni, "Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CE," *American Economic Review*, 2014, *104* (5), 132–35.

Berger, David and Joseph Vavra, "Measuring How Fiscal Shocks Affect Durable Spending in Recessions and Expansions," *The American Economic Review*, 2014, *104* (5), 112–115.

_ and _ , "Consumption Dynamics During Recessions," *Econometrica*, 2015, *83* (1), 101–154.

Bernstein, Shai, Rebecca Diamond, Timothy J. McQuade, and Beatriz Pousada, "The Contribution of High-Skilled Immigrants to Innovation in the United States," Working Paper 2018.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5).

Boar, Corina, "Dynastic Precautionary Savings," *The Review of Economic Studies*, 03 2021, *88* (6), 2735–2765.

_ , Denis Gorea, and Virgiliu Midrigan, "Liquidity Constraints in the U.S. Housing Market," Working Paper 23345, National Bureau of Economic Research April 2020.

Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov, "Conspicuous consumption and race," *The Quarterly Journal of Economics*, 2009, *124* (2), 425–467.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst, "Income Growth and the Distributional Effects of Urban Spatial Sorting," Working Paper 26142, National Bureau of Economic Research August 2019.

Cutler, David M and Edward L Glaeser, "Are ghettos good or bad?," *The Quarterly Journal of Economics*, 1997, *112* (3), 827–872.

_ , _ , and Jacob L Vigdor, "The rise and decline of the American ghetto," *Journal of political economy*, 1999, *107* (3), 455–506.

_ , _ , and _ , "When are ghettos bad? Lessons from immigrant segregation in the United States," *Journal of Urban Economics*, 2008, *63* (3), 759–774.

Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales, "How Segregated is Urban Consumption?," Technical Report, National Bureau of Economic Research 2017.

Diamond, Rebecca and Enrico Moretti, "Where is Standard of Living the Highest? Local Prices and the Geography of Consumption," Technical Report 2022.

_ , Tim McQuade, and Franklin Qian, "The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–94.

Dolfen, Paul, Liran Einav, Peter J Klenow, Benjamin Klopack, Jonathan D Levin, Laurence Levin, and Wayne Best, "Assessing the Gains from E-Commerce," Working Paper 25610, National Bureau of Economic Research February 2019.

Drechsel-Grau, Moritz and Fabian Greimel, "Falling Behind: Has Rising Inequality Fueled the American Debt Boom?," 2018 Meeting Papers 1032, Society for Economic Dynamics 2018.

Eckert, Fabian and Tatjana Kleineberg, "Saving the American Dream? Education Policies in Spatial General Equilibrium," *Education Policies in Spatial General Equilibrium (March 9, 2021)*, 2021.

Fernández, Raquel and Richard Rogerson, "Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform," *The American Economic Review*, 1998, *88* (4), 813–833.

Fisher, Jonathan D., David S. Johnson, and Timothy M. Smeeding, "Measuring the Trends in Inequality of Individuals and Families: Income and Consumption," *American Economic Review*, May 2013, *103* (3), 184–88.

Fogli, Alessandra and Veronica Guerrieri, "The End of the American Dream? Inequality and Segregation in US Cities," Working Paper 26143, National Bureau of Economic Research August 2019.

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *American economic review*, 2019, *109* (7), 2383–2424.

Giannone, Elisa, Qi Li, Nuno Paixao, and Xinle Pang, "Unpacking Moving," Working Paper 2019.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson, "Housing Wealth Effects: The Long View," *The Review of Economic Studies*, 04 2020.

Krueger, Dirk and Fabrizio Perri, "Does income inequality lead to consumption inequality? Evidence and theory," *The Review of Economic Studies*, 2006, *73* (1), 163–193.

Massey, Douglas S., Jonathan Rothwell, and Thurston Domina, "The Changing Bases of Segregation in the United States," *The ANNALS of the American Academy of Political and Social Science*, 2009, *626* (1), 74–90.

Olafsson, Arna and Michaela Pagel, "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *The Review of Financial Studies*, 04 2018, *31* (11), 4398–4446.

Phillips, David, "Measuring Housing Stability with Consumer Reference Data," 2019.

Qian, Franklin and Rose Tan, "The Effects of High-skilled Firm Entry on Incumbent Residents," 2020.

Reardon, Sean F and Glenn Firebaugh, "Measures of multigroup segregation," *Sociological methodology*, 2002, *32* (1), 33–67.

_ **and Kendra Bischoff**, "Income inequality and income segregation," *American Journal of Sociology*, 2011, *116* (4), 1092–1153.

_ **, Glenn Firebaugh, David O'Sullivan, and Stephen Matthews**, "A new approach to measuring socio-spatial economic segregation," in "29th general conference of the International Association for Research in Income and Wealth, Joensuu, Finland" 2006.

Saez, Emmanuel and Gabriel Zucman, " Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data *," *The Quarterly Journal of Economics*, 02 2016, *131* (2), 519–578.

Smith, Matthew, Owen M Zidar, and Eric Zwick, "Top Wealth in America: New Estimates and Implications for Taxing the Rich," Working Paper 29374, National Bureau of Economic Research October 2021.

Theloudis, Alexandros, "Consumption inequality across heterogeneous families," Working Paper 2020.

Trounstine, Jessica, *Segregation by Design: Local Politics and Inequality in American Cities*, Cambridge University Press, 2018.

Zheng, Angela and James Graham, "Public Education Inequality and Intergenerational Mobility," *Available at SSRN 3092714*, 2020.

Appendices

A Data Appendix

A.1 Data Sources

In this section, we describe in further detail the rest of the data sources we use. We first give a broad overview of the ACS and the Nielsen Homescan datasets. We also provide more detail on the structure of the Infutor data and the variables we use to measure the housing and car consumption segregation. Finally, we explain our imputation procedure for car and house prices and provide additional detail on the geographic crosswalk.

A.1.1 American Community Survey

The ACS is an annual survey conducted by the Census Bureau since 2003 where individuals are randomly sampled in each state, the District of Columbia and Puerto Rico. We use the information on income, education, age, race and geography (PUMA and state) from the ACS for two reasons. First, we use all the information above to validate the representativeness of the Infutor data. Second, we use the information on race and income to measure residential segregation by race, income, age and education, all of which are important variables in our analysis but are imputed in the Infutor data. We restrict the ACS sample to individuals between 22 and 80 and exclude observations with non-positive income. The income variable we focus on is total household income. Although the car and home characteristics data from Infutor is at the level of a car-individual owner or home-individual owner level, we view these two consumption categories as likely to be used jointly by all members of a household.

Although the ACS is available since the year 2000 we exclude the years before 2005 as the publicly available data does not report individuals' locations besides the state of residence. In all our analysis we use the finest unit of geography available in the ACS, which is the Public Use Microdata Area (PUMA). There are 2,351 different PUMAs covering all the U.S. territory and they are state-specific, which means that PUMAs do not cross state borders.

The main ACS variables we use in the analysis are total household income (`HHINCOME`) and house value (`VALUEH`). Our results for income segregation are robust to considering alternative income variables such as total personal earned income (`INCEARN`) or total personal income (`INCTOT`).

A.1.2 Census

To complement the ACS we use the 5% microdata sample from the 1990 and 2000 Census available in IPUMS. This sample contains information on approximately 10 million individuals and 5 million households. As with the ACS, from the Census we use in the analysis information on total household income and house value.

A.1.3 Infutor

The Infutor data is organized as described in Figure 5. Here we supplement the description of the dataset from the main text with additional details.

Figure 5: Structure of Infutor Data

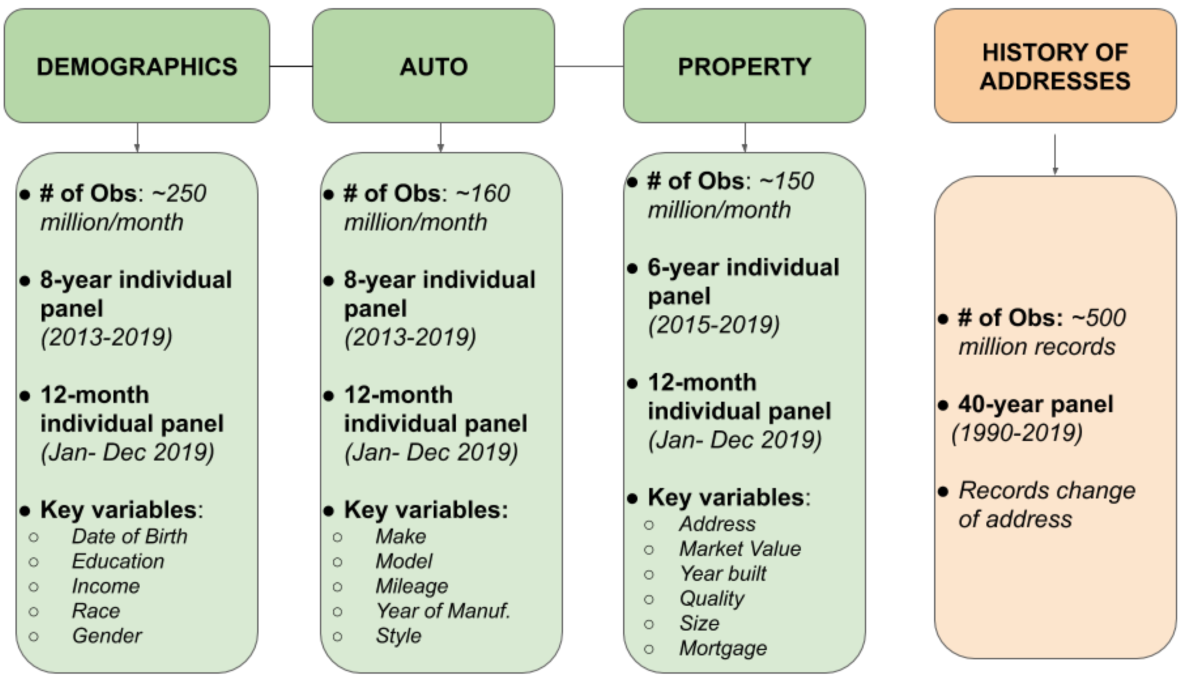

Notes: Figure 5 describes the structure of the Infutor data divided in four main files. The first three files (Demographic Profiles, Auto Profiles and Property Profiles) have been linked through individual identifiers.

For both housing and vehicles we observe in the dataset the ZIP code where the house or the vehicle is located. ZIP codes are more than 10 times more disaggregated than PUMAs. However, given that our analysis for income is done at the PUMA level we aggregate ZIP codes to PUMAs, which we use as the finest geographic aggregation level for all the analysis, i.e. the geographic unit i in equation (1). Our segregation results for housing and vehicles are qualitatively robust to considering alternative definitions of the geographic unit i, such as ZIP codes or counties, instead of PUMAs. In Section A.4 we describe the ZIP-PUMA crosswalk we use to assign the ZIP codes in the Infutor data to PUMAs.

Housing. For each property we observe a unique identifier (given by the property address), an owner unique identifier, its location, its characteristics, such as the year of construction (`PROP_YRBLD`) or the number of rooms, information about the property deed, and its history of mortgages. Infutor also reports whether the property is used as a business or as a residence. For all our analysis we restrict to properties that are used for residential purposes only.

For houses we mostly rely on the deed data from which we observe information of the date (`PROP_SALEDATE`) and the price (`PROP_SALEAMT`) at which the property was acquired by the current owner in a given snapshot. Furthermore, for approximately 68% of the deed records we also observe these two variables for the previous deed of the property (`PROP_SALEDATE_PRIOR` and `PROP_SALEAMT_PRIOR`). Combined with the multiple snapshots of the data, this implies that for several properties we observe more than one transaction. On average we observe 1.86 transactions per property with a maximum of 7 transactions for some of the properties.

Sample selection. Using the variables described above, we include a house in the year t sample if:

$$\min\left\{\texttt{PROP\_YRBLD}, \min_{\text{snapshot}}\{\texttt{PROP\_SALEDATE\_PRIOR}, \texttt{PROP\_SALEDATE}\}\right\} \leq t, \tag{9}$$

where all the date variables are in years and the minimum is computed allowing for possible missing values. This definition implies that we include in our sample all the properties which we can verify that have been built or sold at any previous, or current, date. In Section A.2 we explain how we interpolate selling prices to compute the value of all in-sample houses across time.

Vehicles. For vehicles we observe unique owner-vehicle identifier together with the owner's address which allows us to identify vehicles' location. Unlike houses, we cannot track the same vehicle across multiple owners. We also observe relevant characteristics of the vehicle such as the manufacturer (`MAKE`), the model (`MODEL`), the year (`YEAR`), and the first 10-digits of the Vehicle Identification Number (`VIN`). For example, an entry in our data is a BMW-5 Series-2015 with VIN code WBA5M6C5xF.

As with housing, we combine all the information available at the different snapshots. This allows us to increase the number of vehicles across time, particularly as the coverage significantly increased in the most recent snapshots. Regarding the vehicle location, for the years between 2012 to 2018 we compute the segregation index using the current snapshot ZIP code, when available. For the rest of the years prior to the first time we observe a vehicle we impute the ZIP code of the first snapshot in which we observe it. This assumption is likely to only have a limited impact as between 2012 and 2018 on average only 3.3% of vehicles changed PUMAs and only 1.1% changed state.

Sample selection. A vehicle is included in the year t sample if

$$\texttt{YEAR} \leq t \leq \max\{\text{snapshot}\}, \tag{10}$$

where the max operator is computed across, potentially different snapshot years for which we observe the same vehicle.

Two comments about the vehicle sample selection are in order. First, the backward-looking nature of equation (10), going back to the vehicle's manufacturing year, allows us to run our segregation analysis even for years before Infutor's first snapshot, which corresponds to the year 2012 for the case of vehicles. Thus, for example, if we observe a Toyota-Corolla-2007 in the 2014 snapshot we include this vehicle in our segregation analysis starting from the year 2007 even though we observe this vehicle for the first time in 2014. In other words, what definitions (9) and (10) are doing is to include houses and vehicles since the first year we verify that they exist. Second, definition (10) allows for vehicle exit given by the last snapshot in which the vehicle was observed. We allow for that as for the cases in which Infutor confirmed that the vehicle was sold the owner-vehicle entry will stop appearing in the data in all subsequent snapshots. In Section A.3 we describe the detail of how we obtain the value of the vehicles in our sample and across time.

To test the validity of our vehicles' dataset, we exploit information on age and manufacturers. Specifically, the top panel of Table 11 reports the average age of vehicles in Infutor and BTS between 2012 and 2018. The average age in Infutor is slightly higher but both can be approximated to be at 12 years. The table also reports the correlation at state level of the average age. We find this correlation to be 0.54. The central (bottom) panel reports the age of vehicles in Infutor and NHTS in 2009 (2017). We find that the average age is 2 years higher (3 years lower) in NHTS than in Infutor and the state-level correlation is 0.37 (0.33). We also analyze the distribution of make in both datasets. We find that overall the distributions are very similar both at national- and state-level. Table 12 reports the average age of vehicles at national level for each year and the BTS datasets and in both versions of the Infutor datasets that we built. Overall we find that comparing the snapshot Infutor dataset to the BTS, the latter has on average slightly younger vehicles. Instead, when we compare the baseline dataset to the BTS, we find that the latter has slightly older cars for the years before 2016 and it reverts afterward. Overall, the differences are not very large.

A.1.4 Nielsen Homescan

For non-durable consumption we use the Nielsen Homescan Data from the Kilts Marketing Data Center at the University of Chicago Booth School of Business. These data consist on a longitudinal panel of approximately 40,000-60,000 U.S. households and contain information about the products they buy, as well as when and where they were purchased between 2004 and 2017. It tracks consumers' grocery purchases by asking them to scan the bar codes for each product they purchase after each shopping trip. The Nielsen Homescan data has national coverage and provides wide variation in household location and demographics. Overall, the data include purchases of almost 250 million different items. One of the advantages of this dataset is that it records the bar codes at a very fine level, as well as the expenditure on each of them. Moreover, it covers longer period of time than Infutor. A disadvantage of the Nielsen Homescan dataset relative to the Infutor data is that the sample is orders of magnitude smaller, which means that we would not be able to use these data at very fine geographical detail.

The Nielsen-Kilts data reports detailed location of households. In particular it reports households' ZIP code of residence. We use the same crosswalk used for the Infutor data to assign households to the different PUMAs.

Household purchases are reported at a very granular level. Specifically, at the Universal Product Code (UPC) level. This data reports both the quantity and the total expenses made at the UPC level. Besides UPC there are other two additional aggregation levels which are product modules, such as soap or beer, and ten different department codes such as dairy or health and beauty consumption.

A.2 House Prices

This section explains how we compute prices for in-sample houses in the Infutor data, defined in equation (9), at any time t. We use observed selling prices to compute ZIP code level price indices with which we interpolate property prices at different time periods.

In our analysis, we restrict attention to properties for which we observe the date and the price of at least one transaction. As explained above, for each property we keep all the selling

Table 11: Representativeness in Infutor and BTS datasets

		National-level			State-level	
		BTS	NHTS	Infutor	FHA	NHTS
2012-2018						
Vehicle Age						
	Mean	11.5		12.2		0.54
2009						
Vehicle Age						
	Mean	10.3	9.4	7.1		0.37
Vehicle Make*						
	Chrysler		0.12	0.15		0.74
	Ford		0.19	0.19		0.59
	GM		0.22	0.22		0.62
	Honda		0.09	0.08		0.40
	Toyota		0.12	0.10		0.22
	Other		0.26	0.26		0.36
N		254,582,694	211,501,318	98,996,193	0.32	0.44
2017						
Vehicle Age						
	Mean	11.7	10.4	13.3		0.33
Vehicle Make*						
	Chrysler		0.11	0.13		0.71
	Ford		0.15	0.18		0.65
	GM		0.19	0.22		0.73
	Honda		0.11	0.08		0.57
	Toyota		0.15	0.11		0.40
	Other		0.30	0.28		0.55
N		254,582,694	222,578,947	158,898,357	0.56	0.54

Notes: The table reports the average age of a car across datasets, as well as the distribution of car makes.

Table 12: Age distribution of cars in Infutor and BTS

	National-level		
Vehicle Age	BTS	Snapshot	Baseline
2000	8.9		5.7
2001	8.9		5.9
2002	9.6		6.0
2003	9.7		6.1
2004	9.8		6.3
2005	9.8		6.6
2006	9.9		6.9
2007	10		7.2
2008	10.1		7.4
2009	10.3		8.0
2010	10.6		8.4
2011	10.9		8.8
2012	11.2	9.7	9.1
2013	11.4	12.0	9.5
2014	11.4	11.6	10.2
2015	11.5	12.1	10.9
2016	11.6	12.4	11.7
2017	11.7	13.3	12.6
2018	11.7	14.4	13.5
Corr with BTS			
2012-2018		0.950	0.952
2000-2018			0.935

Notes: The table reports the age distribution of cars in BTS and Infutor. For Infutor, we report both the "Snapshot" and the "Baseline" version of the data. The bottom of the table reports correlation coefficients of the average age by year between datasets.

dates and prices available in the data. The total number of transactions we observe varies by year. For example, for the years before 1990, we observe around 600,000 annual transactions. For the time period between 1990 and 2018, we observe more than 3 million transactions per year.

Using the selling prices we compute a time series for the median transacted house price at different geographies: at the national, state, PUMA and ZIP code level. To assess the quality of our house price data, Figure 6 displays the level and the annual log changes of the national-level median price against the House Price Index (HPI) from the Federal Housing Finance Agency. This figure shows that the national-level house prices obtained from Infutor are consistent with aggregate-level house price dynamics and properly capture the different boom and bust episodes observed in the last four decades.

Figure 6: House Prices

(a) Total consumption (b) Durable consumption

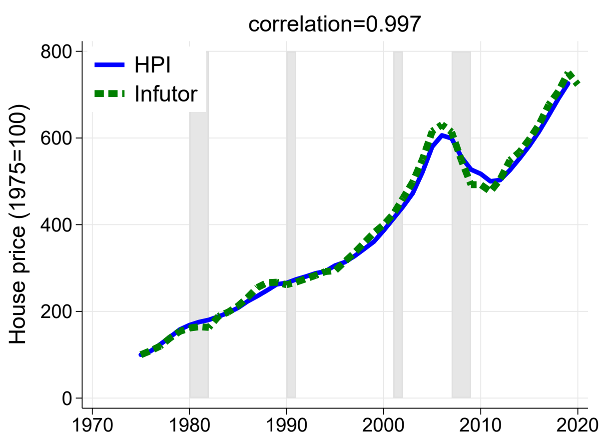

Notes: The figure compares the evolution over time of house prices and house price growth in Infutor and HPI. HPI denotes the All-Transactions House Price Index from the Federal Housing Finance Agency. Infutor denotes the national-level median selling price.

ZIP Code House Price Indices. We use the data to construct ZIP code-level price indices from 1950 to 2018. To that end, we first compute annual price changes across all four geographies for which we observe at least 50 transactions. If this constraint is not met for a given ZIP-year pair we try to use the price change of the immediate broader unit of geography level, which is the PUMA. If the corresponding PUMA year does not have the price change available we continue to the state or even the national level. These imputations are more common as we go back in time, particularly, before 1990.

Price Interpolation. With the corresponding ZIP code level price index in which the property is located we interpolate prices to other years different from the property observed transactions. This strategy takes into account different boom and bust house price cycles observed at different geographies. To exemplify our interpolation strategy, suppose that we observe that house h, located in ZIP code z, was sold at price $p_{h,z,\tau}$ in period τ. We compute the price of house h at any other year t by assuming that its price between period t and τ followed the zip code level price dynamics. This assumption can be written as

$$\Delta_{t-\tau} \log p_{h,z,\tau} = \Delta_{t-\tau} \log p_{z,\tau},$$

which implies that house h price at year t is equal to

$$p_{h,z,t} = \exp\left(\log p_{h,z,\tau} + \Delta_{t-\tau} \log\log p_{z,\tau}\right),$$

where $p_{z,\tau}$ is ZIP code z price index in period τ.

Note that in the expression above we could have $t < \tau$ as long as equation (9) is satisfied. For the properties for which we observe more than one transaction we interpolate the price at period t using the closet (in terms of minimum time distance) selling price.

A.3 Vehicle Prices

One limitation of the Infutor vehicle data is that we do not observe prices or any other estimate for the value of vehicles. To circumvent this limitation we merged the Infutor data with transaction prices for all the dealer sales of new and used cars in Texas from 2012 to 2019. Dealer sales refer to all the sales done by licensed dealers.

We perform this merge in two steps, trying to use the most disaggregated data available. To that end, we first compute time series for 10-digit VIN-level mean selling prices for all the VINs for which we observe at least 50 transactions. Additionally, we compute mean selling prices, also restricting to the minimum 50 observations constraint, for all the make-model-year combinations. VIN codes are considerably narrower specifications compared to make-model-year tuples as they also distinguish other technical characteristics such as the vehicle engine or the type of fuel it uses. To put this in perspective, in Infutor, we observe more than 130,000 different 10-digit VINs and close to 15,000 make-model-year combinations. In our data around 50% of the make-model-year tuples have at most 4 different 10-digit VINs.

Using both sets of time series we compute VIN-level prices for 2012 to 2018. As we did for houses, we impute missing VIN-level prices using the make-model-year price when the VIN-level price is not available. For these years we drop observations below the 10^{th} and above the 90^{th} percentile annual price change distribution. This restriction rules out, for example, price increases in used cars, which probably reflect changes in the composition of sold vehicles over time.

Using prices from 2012 to 2018, we then interpolate prices for missing observations and for the years before 2012, for the older vehicles in our sample. For that, we use a VIN-level annual depreciation rate. Specifically, for each VIN v we compute the average depreciation rate between 2012 and 2018 as

$$\delta_v = -\frac{1}{\#\mathcal{T}_v} \sum_{t \in \mathcal{T}_v} \Delta \log p_{v,t},$$

where $\mathcal{T}_v \subset \{2012, \dots, 2018\}$ is the set of years for which we observe VIN v's prices. The average depreciation rate (or average annual price decrease) in our VIN-level prices is 0.108.

With this VIN-level depreciation rate we interpolate missing prices for our in-sample vehicles in year t that satisfy equation (10). For example, for a 2005 VIN ($\texttt{YEAR}= 2005$) for which we observe selling prices starting in 2012 we compute its time $2005 \leq t < 2012$ price as

$$p_{v,t} = \exp\left[\, \log p_{v,2012} + (2012 - t)\delta_v \,\right],$$

where $p_{v,2012}$ is the VIN-level selling price observed in 2012.

A.4 Geographic Crosswalks

In addition to the country and U.S. states, the economic geography and segregation literature usually consider five other levels of geography: (1) ZIP-code, (2) county, (3) Public Use Microdata Area (PUMA), (4) Core Based Statistical Area (CBSA), and (5) Commuting Zone (CZ). These categories are presented in Table 13, ordered from left to right in terms of their level of disaggregation.

We make some remarks about these units of geography. First, ZIP-codes do not cover the entire territory, approximately 0.002% of population is not assigned to a ZIP-code. Additionally, a ZIP-code can cover more than one state. Second, counties cover the entire territory and are state-specific. Third, PUMAs also cover the entire territory and are state-specific. As mentioned above, this is the finest aggregation level available in the ACS and Census public use microdata. Fourth, CBSAs do not cover the entire territory, approximately 6% of population is not assigned to a CBSA. CBSAs can cover more than one state. For example, the CBSA New York-Newark-Jersey City, spans three states: New York, New Jersey and

Table 13: Number of Locations by Geographic Aggregation Level

	ZIP	County-10	PUMA	CBSA-15	CZ
MCDC	32,845	3,143	2,351	933	741
ACS [2005-2018]		376-430	2,351		
Census 1990		385	1,726		
Census 2000		377	2,351		
Infutor [2012-2019]	39,131				

Notes: The table reports the number of geographic units from smaller to larger. MCDC denotes the Missouri Census Data Center. County-10 and CBSA-15 denote the list of counties and CBSAs in 2010 and 2015 respectively. The Missouri Census Data Center Geographic Correspondence Engine is a comprehensive source of definitions and geographical cross-walks for the United States. We compare the zipcodes available in Infutor with the ones in the MCDC data.

Pennsylvania. Fifth, MSAs (not reported in the table), are, broadly, a subgroup of CBSAs that only restricts to metropolitan areas. There are 384 MSAs (vs. 933 CBSAs). Lastly, CZs are clusters of U.S. counties that are characterized by strong within-cluster and weak between-cluster commuting ties.

For our analysis we constructed a crosswalk that uniquely assigns ZIPs to PUMAs (an injection correspondence), and consequently, to states. We started from the ZIP-PUMA crosswalk available in the MCDC. This crosswalk, besides assigning ZIPs to PUMAS, reports the ZIP-code population share that is assigned to the potentially different PUMAs. Indeed, for approximately 28% of the ZIP-codes there is more than one PUMA assigned. In these cases, we assigned each ZIP to a single PUMA using the one with the largest population share. For example, if a ZIP-code is assigned to two PUMAs with 70-30% population shares we assigned this ZIP-code to the former PUMA with a 70% population share. With this procedure we constructed an injective crosswalk between ZIP codes to PUMAs. This crosswalk is available upon request.

B Robustness

This section provides several robustness tests for the entropy measure. First, we change the subunit of geography from PUMA to zipcodes and report the correlation coefficients between entropy indices at the national or state level calculated using zipcodes, counties and PUMAs as the underlying geographic unit. Table 14 suggests that the correlations between the entropy indexes for car and housing consumption when the smallest unit of geography is the zipcode, county or PUMA are very highly correlated both at the national level and at the state level. The correlations range between 0.639 and 1, suggesting that the analysis at PUMA level is robust of the analysis at the zipcode level. Second, Figure 7 compares the results for entropy measure in house values both in Infutor and ACS. Third, in section B.1 we change the larger unit of geography from states to CBSAs and report the main results for the entropy index at CBSA level rather than at the state level.

Table 14: Correlations between Entropy Indexes at Different Levels of Geography

	State-level		National-level	
	Housing - PUMA	Car - PUMA	Housing - PUMA	Car - PUMA
Housing - County	0.717***		0.971***	
Housing - Zipcode	0.971***		0.976***	
Car - County		0.812***		1.000***
Car - Zipcode		0.639***		0.992***
Observations	2499	1224	49	24

Notes: The table reports the correlations between H-indexes for housing and car consumption measured using the Infutor data for different levels of underlying geographies. On the left panel, we report correlations between H-indexes aggregated at the national level. On the right, we report correlations for H-indexes aggregated at state-level.

Figure 7: Comparing Trends in Entropy Index for Infutor and ACS Housing Data

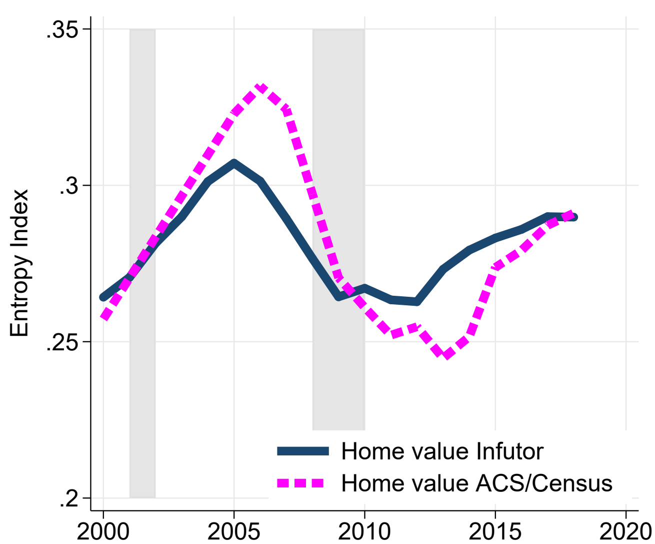

Note: The figure plots consumption entropy indices for house values in Infutor and ACS.

B.1 Entropy Index for CBSAs

Figure 8: CBSA-level Consumption Segregation

(a) Total consumption

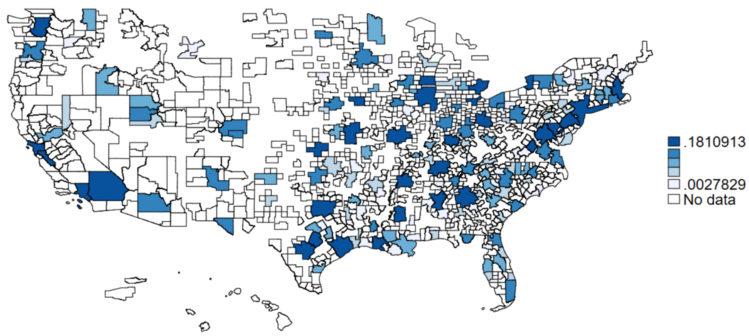

(b) Durable consumption

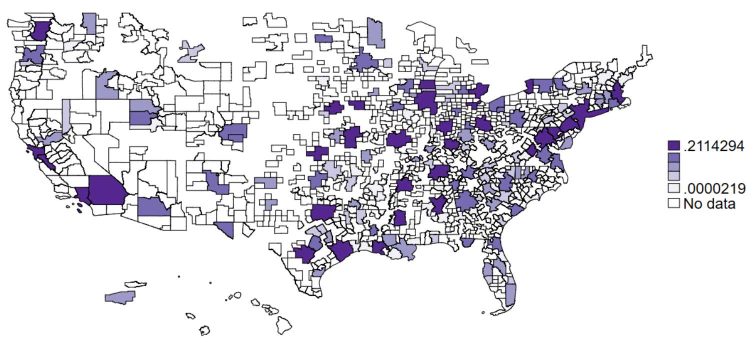

(c) Non-durable consumption

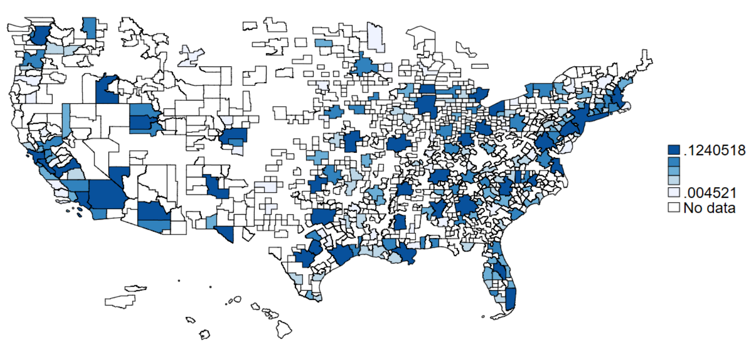

Notes: The figure plots CBSA-level consumption entropy indices for durables, housing and cars, averaged over the period 2016-2018.

Figure 9: CBSA-level Durable Consumption Segregation

(a) Housing (b) Cars

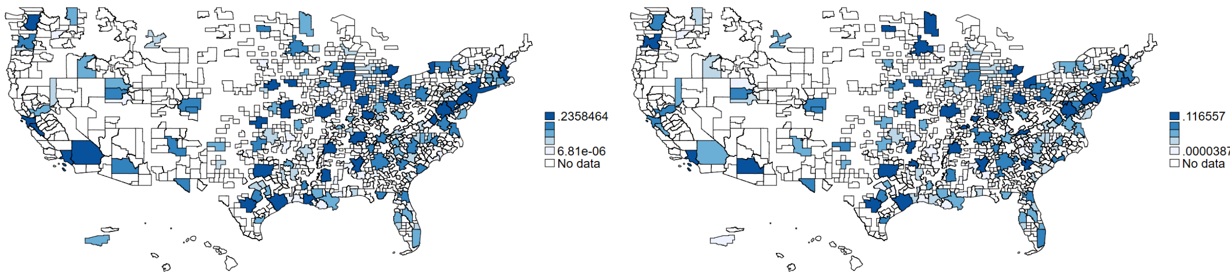

Notes: The figure plots CBSA-level consumption entropy indices averaged over the period 2016-2018.

Table 15: Consumption Segregation and Regional Characteristics

	Total consumption	Durable consumption	Non-durable consumption
	By income		
Bottom 25%	0.096	0.107	0.067
Top 25%	0.200	0.251	0.085
	By education (college share)		
Bottom 25%	0.101	0.113	0.068
Top 25%	0.188	0.236	0.072
	By race (share white)		
Bottom 25%	0.200	0.246	0.085
Top 25%	0.098	0.115	0.056
	By age (share $\leq$ 35)		
Bottom 25%	0.118	0.141	0.060
Top 25%	0.172	0.207	0.089
	By population		
Bottom 25%	0.061	0.068	0.044
Top 25%	0.169	0.206	0.078

Notes: The table reports population weighted averages of state-level consumption entropy indices over the period 2016-2018.

Table 16: Consumption Segregation by Demographic Group

	Total consumption	Durable consumption	Non-durable consumption
		By race	
White	0.143	0.166	0.083
Black	0.156	0.170	0.118
Other	0.152	0.183	0.070
		By education	
No college degree	0.148	0.168	0.095
College degree	0.166	0.178	0.135
		By age	
Younger than 35	0.167	0.184	0.123
Older than 35	0.148	0.176	0.076

Notes: The table reports population weighted averages of state-level consumption entropy indices over the period 2016-2018.

Table 17: Drivers of Consumption Segregation at State Level

	Total consumption			Durable consumption			Non-durable consumption		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
$\ln H_Y$	0.60***	0.49***	0.47***	0.77***	0.65***	0.49***	0.20***	0.25***	0.21**
	(0.05)	(0.06)	(0.08)	(0.05)	(0.06)	(0.09)	(0.06)	(0.07)	(0.09)
$\ln H_R$		0.14***	0.14*		0.16***	0.14**		-0.06	-0.07
		(0.07)	(0.08)		(0.06)	(0.07)		(0.05)	(0.06)
$\ln H_E$			0.04			0.22**			0.07
			(0.12)			(0.11)			(0.09)
$\ln H_A$			-0.05			-0.08			-0.02
			(0.05)			(0.05)			(0.07)
R^2	0.82	0.84	0.84	0.79	0.81	0.82	0.53	0.53	0.53

Note: The table reports the estimates of α_1, α_2, α_3 and α_4 in equation (6). The estimate of α_0, the time fixed effects and the estimate of the vector $\boldsymbol{\gamma}$ are omitted. Robust standard errors are reported in parentheses. The regression is estimated with population weights based on the 2010 Census. ***, **, and *, represent statistical significance at 1%, 5% and 10%, respectively.

Table 18: Drivers of Consumption Segregation, Heterogeneity

	Total consumption	Durable consumption	Non-durable consumption
	By income		
$\ln H_Y$	0.60***	0.83***	0.34***
	(0.05)	(0.08)	(0.05)
$\ln H_Y \times \mathbf{1}_{\text{Top 25\%}}$	0.23	0.09	0.15*
	(0.16)	(0.18)	(0.08)
	By education (college share)		
$\ln H_Y$	0.61***	0.82***	0.41***
	(0.05)	(0.08)	(0.05)
$\ln H_Y \times \mathbf{1}_{\text{Top 25\%}}$	0.10	0.11	0.008
	(0.08)	(0.13)	(0.07)
	By race (share white)		
$\ln H_Y$	0.69***	0.87***	0.48***
	(0.06)	(0.05)	(0.07)
$\ln H_Y \times \mathbf{1}_{\text{Top 25\%}}$	-0.06	-0.12	-0.11
	(0.09)	(0.09)	(0.08)
	By age (share $\leq$ 35)		
$\ln H_Y$	0.67***	0.77***	0.52***
	(0.05)	(0.06)	(0.04)
$\ln H_Y \times \mathbf{1}_{\text{Top 25\%}}$	-0.03	0.06	-0.02
	(0.10)	(0.12)	(0.07)
	By population		
$\ln H_Y$	0.14	0.23**	-0.02
	(0.09)	(0.09)	(0.07)
$\ln H_Y \times \mathbf{1}_{\text{Top 25\%}}$	0.60***	0.71***	0.26**
	(0.15)	(0.16)	(0.11)

Notes: The table reports slope coefficients from regressing consumption segregation on income segregation and income segregation interacted with a dummy variable that equals 1 if a CBSA is in the top 25% of the national distribution of income, the share of the population that is college educated, white, younger than 35, and population, respectively. All regressions include separate controls for the dummy variable previously described and, as well as the set of controls in equation (6). All other coefficients are omitted in the interest of space. With the exception of the last panel, all regressions are estimated with population weights based on the 2010 Census. ***, **, and *, represent statistical significance at 1%, 5% and 10%, respectively.

B.2 Inspecting the Relationship Between Income and Consumption Segregation

Motivated by the results above, in this section we inspect what is the mechanism driving the strong positive relationship between income and consumption segregation. To that end, we scrutinize two potential explanations: market incompleteness and persistent income differences between individuals. Market incompleteness, or the lack of access to some financial products (either because of these products being absent or because of underlying informational frictions) and persistent income shocks both have the potential to generate such a positive relationship since they prevent consumption insurance among individuals with different income.

B.2.1 An Analytical Framework

We begin by outlining a simple framework to guide our empirical investigation and then test the predictions of this theory. We consider the utility maximization problem of a household i who lives in region r in year t. Each period the household earns stochastic income y^i_{rt} and chooses the optimal path of consumption c^i_{rt} and savings a^i_{rt+1} that maximizes its expected life-time utility

$$\mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c^i_{rt}\right),$$

subject to the budget constraint

$$c^i_{rt} + \frac{a^i_{rt+1}}{r} = y^i_{rt} + a^i_{rt}$$

and a borrowing constraint $a^i_{rt+1} \geq \underline{a}$. Here, β is the discount factor of the household, r is the real interest rate on savings and $\underline{a}$ is the borrowing limit. The optimal consumption-saving choice is characterized by the Euler equation

$$u'(c^i_{rt}) = \beta(1+r)\mathbf{E}_t u'(c^i_{rt+1}) + \mu_t,$$

where μ_t is the multiplier on the borrowing constraint. That is, when deciding how much to save, the household compares the cost of saving an additional unit of income, expressed in terms of the marginal utility of lost consumption, with the discounted benefit of having an additional unit of resources available for consumption in the following period.

In order to derive an analytical characterization of optimal consumption choices, we specialize this model to what is commonly called the strict permanent income hypothesis. Specifically, we assume that *(i)* households have quadratic utility $u(c) = b_1 c + \frac{1}{2} b_2 c^2$, where the parameters b_1 and b_2 are such that the utility function is strictly increasing and strictly concave, *(ii)* the interest on savings equals the inverse of the discount rate $\beta(1+r) = 1$, and *(iii)* there is no borrowing constraint. Under these assumptions, it can be easily shown that consumption is a martingale

$$c^i_{rt} = \mathbf{E}_t c^i_{rt+1} \quad \text{and, more generally,} \quad c^i_{rt} = \mathbb{E}_t c^i_{rt+j}, \forall j \geq 0.$$

Iterating forward on the budget constraint, taking the limit as $t \to \infty$ and using a No-Ponzi scheme condition gives

$$c^i_{rt} = \frac{r}{1+r}\left[a^i_{rt} + \sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \mathbb{E}_t y^i_{rt+j}\right],$$

that is, consumption is the annuity value of human and financial wealth. It can then be shown that the change in consumption $\Delta c^i_{rt} \equiv c^i_{rt} - c^i_{rt-1}$ is equal to

$$\Delta c^i_{rt} = \frac{r}{1+r}\sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \left(\mathbb{E}_t - \mathbb{E}_{t-1}\right) y^i_{rt+j}$$

and is proportional to the revision in expected earnings due to the new information accruing between $t-1$ and t.

To make further analytical progress, we specialize the process for income by assuming that

$$y^i_{rt} = \rho y^i_{rt-1} + \varepsilon^i_{rt},$$

where ε is a mean-zero shock with variance σ^2_ε that is iid across households, regions and over time, and ρ is the persistence parameter. Under this specialized income process, consumption dynamics are given by

$$\Delta c^i_{rt} = \frac{r}{1+r-\rho}\varepsilon^i_{rt}.$$

This expression nests two special cases. First, when $\rho = 0$ shocks to income are transitory and consumption evolves according to $\Delta c^i_{rt} = \frac{r}{1+r}\varepsilon^i_{rt}$, implying that the households consume only the annuity value of the shock and that most of it is saved. However, if there are no asset markets, then, $c^i_{rt} = y^i_{rt}$ and, consequently, $\Delta c^i_{rt} = \varepsilon^i_{rt}$, so income shocks translate entirely into consumption. Second, when $\rho = 1$ shocks to income are permanent, and consumption evolves according to $\Delta c^i_{rt} = \varepsilon^i_{rt}$, implying that the shocks translates entirely into consumption.

Zooming out and focusing on the distributions of income and consumption in a given region, the discussion above makes it clear that if income shocks are permanent, then the two distributions mirror each other. If we think about the regions r as being part of a broader geographic unit, such as a state or a CBSA, then the segregation of income within the broader geographic unit should be reflected entirely in the segregation of consumption. In general, how closely the distribution of consumption tracks that of income depends on how persistent income shocks are: the more persistent they are, the closer are the two distributions.

If, instead, income shocks are transitory, the distribution of consumption is less reflective of the distribution of income (i.e. the term $\frac{r}{1+r} \approx 0.02$ for values of r typically used in the literature), as agents can use saving and borrowing to smooth consumption in response to this type of shocks. Once more, thinking of regions r as being part of a broader geographic unit, this implies that income segregation translates into consumption segregation only to a limited extent. However, if asset markets that allow for saving and borrowing are absent (this is what we refer to as market incompleteness), then the two distributions become more intimately linked, as in the case of permanent/persistent income shocks.

Table 19: Consumption and Income Segregation, State Level

	Total	Between	Within
Total consumption	2.05	0.04	2.01
α_1^i/α_1		2%	98%
Non-durable consumption	0.71	0.25	0.47
α_1^i/α_1		35.2%	64.8%
Durable Consumption	2.99	0.15	2.84
α_1^i/α_1		5%	95%
Cars	0.14	0.02	0.12
α_1^i/α_1		14.3%	85.7%
Housing	3.34	0.13	3.21
α_1^i/α_1		3.9%	96.1%

Notes: The table reports the estimates of α_1, α_1^B and α_1^W. All estimates are significant at 1% significance. Robust standard errors, time fixed effects and the estimates of $\boldsymbol{\gamma^i}$ and $\boldsymbol{\gamma}$ are omitted but are included in the estimation. The regression is estimated with population weights based on the 2010 Census, using data for the 2016-2018 period.

C Model Appendix

Let $w^i_{jt} = a^i_{jt} + h^i_{jt}$ denote the household's total stock of wealth. We assume that $a_{j0} = a_0$ and $h_{j0} = 0\ \forall t$. That is, initial wealth is the same across locations (no ex-ante heterogeneity) and households are endowed with no units of the visible good. Combining the budget constraint and the law of motion for the stock of durables gives

$$c^i_{jt} + w^i_{it+1} = y^i + (1+r)\,w^i_{jt} - (r+\delta)\,h^i_{jt}. \tag{11}$$

The first order conditions with respect to c^i_{jt}, h^i_{jt} and w^i_{jt+1} are, respectively:

$$\begin{aligned}
\left[(1-\xi)\left(c^i_{jt}\right)^{\varepsilon} + \xi\left(h^i_{jt} - \bar{h}^i_{jt}\right)^{\varepsilon}\right]^{\frac{1-\sigma}{\varepsilon}-1}(1-\xi)\,\varepsilon\left(c^i_{jt}\right)^{\varepsilon-1} &= \lambda_{jt} && (12)\\
\left[(1-\xi)\left(c^i_{jt}\right)^{\varepsilon} + \xi\left(h^i_{jt} - \bar{h}^i_{jt}\right)^{\varepsilon}\right]^{\frac{1-\sigma}{\varepsilon}-1}\xi\varepsilon\left(h^i_{jt} - \bar{h}^i_{jt}\right)^{\varepsilon-1} &= \lambda_{jt}\,(r+\delta) && (13)\\
\lambda_{jt} &= (1+r)\,\lambda_{jt+1}, && (14)
\end{aligned}$$

where λ_{jt} is the Lagrange multiplier of the budget constraint (11) and where we assume that the household is at an interior solution. Combining the first order conditions (12) and (13) gives

$$\frac{1-\xi}{\xi}\left(\frac{c^i_{jt}}{h^i_{jt} - \bar{h}^i_{jt}}\right)^{\varepsilon-1} = \frac{1}{r+\delta},$$

or, equivalently,

$$c^i_{jt} = \underbrace{\left(\frac{(r+\delta)(1-\xi)}{\xi}\right)^{\frac{1}{1-\varepsilon}}}_{\equiv\kappa} h^i_{jt} - \left(\frac{(r+\delta)(1-\xi)}{\xi}\right)^{\frac{1}{1-\varepsilon}} \bar{h}^i_{jt} \tag{15}$$

From (14) we get that $\lambda_{jt} = \lambda_j,\ \forall t$. Plugging the above in (13) yields a left-hand side that is a function of h^i_{jt} and a right-hand side that is constant. Therefore, $h^i_{jt} = h^i_j,\ \forall t$, $c^i_{jt} = c^i_j,\ \forall t$ and $a^i_{jt} = a^i_0,\ \forall t$, where a^i_0 denotes the initial wealth of a household of type i.

Iterating forward on the budget constraint (11) and applying the transversality condition allows gives

$$c^i_j + (r+\delta)\,h^i_j = \underbrace{rw^i_0 + y^i}_{Y^i \equiv \text{permanent income}},\ \forall t. \tag{16}$$

Combining equations (15) and (16) and solving for h^i_j we obtain

$$\begin{aligned}
h^R_j &= \frac{1}{r+\delta+\kappa}Y^R && (17)\\
h^P_j &= \frac{1}{r+\delta+\kappa}Y^P + \frac{\kappa}{r+\delta+\kappa}h^R_j = \frac{1}{r+\delta+\kappa}\left(Y^P + \frac{\kappa}{r+\delta+\kappa}Y^R\right) && (18)
\end{aligned}$$

Using the budget constraint, we can express the household's asset position as

$$ra_j^i = c_j^i + \delta h_j^i - y^i,$$

or, equivalently,

$$\begin{aligned}
ra_j^R &= c_j^R + \delta h_j^R - y^R = (\kappa+\delta)\, h_j^R - y^R = \frac{\kappa+\delta}{r+\delta+\kappa} Y^R - y^R && (19)\\
ra_j^P &= c_j^P + \delta h_j^P - y^P = (\kappa+\delta)\, h_j^P - h_j^R - y^P \\
&= \frac{\kappa+\delta}{r+\delta+\kappa}\left(Y^P + \frac{\kappa}{r+\delta+\kappa} Y^R\right) - \frac{1}{r+\delta+\kappa} Y^R - y^P \\
&= \frac{\kappa+\delta}{r+\delta+\kappa} Y^P - \frac{\kappa}{r+\delta+\kappa}\left(\frac{\kappa+\delta}{r+\delta+\kappa} - 1\right) Y^R - y^P && (20)
\end{aligned}$$

Lastly, using equation (16) we can solve for non-durable consumption, which equals

$$\begin{aligned}
c_j^R &= Y^R - (r+\delta)\, h_j^R = Y^R \frac{\kappa}{r+\delta+\kappa} && (21)\\
c_j^P &= Y^P - (r+\delta)\, h_j^P = Y^P - (r+\delta)\frac{1}{r+\delta+\kappa}\left(Y^P + \frac{\kappa}{r+\delta+\kappa} Y^R\right) \\
&= Y^P \frac{\kappa}{r+\delta+\kappa} - \frac{r+\delta}{r+\delta+\kappa}\frac{\kappa}{r+\delta+\kappa} Y^R && (22)
\end{aligned}$$

We can also characterize the relative spending on the two types of consumption goods. Specifically,

$$\begin{aligned}
\frac{c_j^R}{h_j^R} &= \frac{Y^R - (r+\delta)\, h_j^R}{h_j^R} = \frac{Y^R}{h_j^R} - (r+\delta) = \frac{Y^R}{\frac{1}{r+\delta+\kappa} Y^R} - (r+\delta) = \kappa && (23)\\
\frac{c_j^P}{h_j^P} &= \frac{Y^P}{h_j^P} - (r+\delta) = \frac{Y^P}{\frac{Y^P}{r+\delta+\kappa}\left(1 + \frac{\kappa}{r+\delta+\kappa}\frac{Y^R}{Y^P}\right)} - (r+\delta) \\
&= \frac{r+\delta+\kappa}{1 + \frac{\kappa}{r+\delta+\kappa}\frac{Y^R}{Y^P}} - (r+\delta) < r+\delta+\kappa - (r+\delta) = \kappa, && (24)
\end{aligned}$$

where the inequality follows from the fact that $1 + \frac{\kappa}{r+\delta+\kappa}\frac{Y^R}{Y^P} > 1$.

In a case with no segregation the poor do not compare their consumption with that of the rich and we have that $\frac{c_j^P}{h_j^P} = \kappa$. So the higher segregation is, the smaller is the relative spending on non-durables by the poor, as they substitute towards the status enhancing visible consumption good.