One size fits all? The value of standardized retail chains*

Ben Klopack

Texas A&M University

May 31, 2021

Abstract

Multi-outlet firms, or chains, make up a large and growing part of the US retail sector and are the subject of important local policy; over 30 US cities ban or restrict the entry of chain firms. This paper quantifies the welfare and profit effects of standardized chains: chains face higher demand than independent firms, in part because of economies of scale in branding and advertising, but at the same time chains are less flexible in customizing product selection or prices across locations. I quantitatively assess the effects of this tradeoff on firm revenue and consumer welfare in the restaurant industry using a large credit card dataset that covers 20% of US consumption. I find that on average chains could earn 20% higher variable profits if they could customize their product optimally to local tastes, but they would lose 30% of their variable profits if they were to lose their demand advantages. Policies that ban chain restaurants would result in a loss of consumer welfare equivalent to 1.5%-6.4% of restaurant spending and would disproportionately impact lower income consumers.

Keywords: retail chains, standardization, market structure, demand heterogeneity, preference externalities

JEL Classification: L11, L25, R32

*I would like to thank Liran Einav, Tim Bresnahan, Rebecca Diamond, and Matt Genztkow for their invaluable advice and support. I am also grateful to Panle Jia Barwick, Pete Klenow, Brad Larsen, Fernando Luco, Luca Maini, Guido Martirena, Nicola Pierri, Steve Puller, Katja Seim, Tom Wollmann, and the participants in the IO^2 seminar and the Stanford Industrial Organization workshop for comments that significantly improved this paper. I gratefully acknowledge financial support from the Patricia Liu McKenna and Kenneth McKenna Graduate Fellowship through a grant to the Stanford Institute for Economic Policy Research. Finally, I thank the data science team at the payment cards company, as well as Haris Bacic at Pricelisto, for generously providing the data used in this paper. All errors are my own.

1 Introduction

Large, national retail chains that operate in many markets may have important demand advantages relative to smaller firms. Chains may have economies of scale both in building an experience that consumers value, through product development and branding, and in communicating the reputation of that brand through marketing and advertising. The chain structure also enables a firm to share its reputation across stores, allowing it to more effectively compete for mobile consumers.[1] However, chains also face a potentially important diseconomy of scope. In order to realize these reputational advantages, chains often standardize their product offering across stores. If consumer tastes differ across markets, standardization may impose a costly strategic constraint for chain firms.

Chains are also the subject of important local policy. In recent years, over 30 cities in the US have enacted regulations that ban or restrict the entry of chain firms, including restaurants, and many more have debated such policies.[2] Chain bans are largely motivated by a desire to protect local businesses and maintain local character. The effect of these policies on consumers, and the distribution of their welfare impacts across different groups, depend on how consumers value chains relative to the types of firms that would enter if chains were banned.

This article quantifies the profit and welfare impacts of standardized chains in the restaurant industry. The costs and benefits of standardization have clear implications for retailer strategy, competition, and consumer welfare. All else equal, when demand is very heterogeneous across markets, standardization becomes more costly for a firm, and thus fewer large standardized firms will exist in equilibrium. I find that restaurant chains that could be flexible in choosing their product characteristics across markets could substantially increase their variable profits. However, if chains were to lose their sizable demand advantages over

[1] Chains may also have important cost advantages, including economies of scale in procurement and distribution, logistics, and bargaining power with upstream suppliers.

[2] I include a list of known cities with chain bans in Table A.1 in the Appendix. Horgan (2017) and City of Los Angeles (2019) detail discussions of similar proposals in New York, Toronto and Los Angeles. The largest US city with a strict chain ban policy is San Francisco, which bans any retail firm (including restaurants) with more than 11 nationwide locations from operating in three central neighborhoods (Hayes Valley, Chinatown, and North Beach) and requires Conditional Use Authorization in many other areas of the city (San Francisco Planning, 2020).

independent firms to gain that flexibility, they would be significantly worse off.

By the same token, standardization can affect consumer surplus; national chains may provide important benefits for consumers, but their standardized nature may cause a mismatch between local tastes and equilibrium product offerings. Chain bans might remedy this by better aligning firm characteristics with local preferences, but remove a set of firms that many consumers like. If restaurant preferences are heterogeneous, then these policies may also have important distributional impacts. I find that chain bans reduce overall consumer welfare and disproportionately impact the lowest income consumers, who have the highest preference for chains. To my knowledge, this work is the first to study the effects of chain bans on consumers.

This paper leverages a large and novel transaction-level dataset from a payment cards company that includes approximately 20% of total US consumption. I construct a panel of consumer restaurant purchases during dinner hours in seven midsize US cities. I merge these data with supplementary samples from Yelp and Pricelisto that contain detailed restaurant-level information and menu prices. The primary analysis uses this combined sample to fit a discrete choice model of consumer demand for restaurants. The estimation strategy exploits the unique breadth and richness of these data to quantify heterogeneity in consumer preferences over restaurant characteristics. In the demand system, restaurants are differentiated in their cuisine type, quality level, price, physical location, and chain affiliation. Consumer preferences over those restaurant attributes vary along both observable and unobservable dimensions.

I first use the demand estimates to quantify the value of chain affiliation for firms and consumers. I exploit variation in the spatial distribution of cardholders that changes how far a given consumer must travel to visit a particular restaurant. The estimates imply that large chains are highly valued by many consumers. Consumers with annual household income less than $50,000 and between $50,000 and $100,000 would be willing to travel an additional 1.2 and 0.7 miles respectively to eat at a large chain relative to an independent restaurant with similar characteristics, relative to median travel distances of 3.8 and 5.1 miles. However, preferences for chains are sharply declining in income; cardholders in the two highest income groups (income between $150,000-$200,000 and above $200,000) actually prefer independent restaurants to chains.[3] This heterogeneity in preferences for chains suggest that policies that restrict their entry may have important distributional consequences.

I then investigate the possible mechanisms for this demand premium. I show evidence that information and branding appear to be important contributing factors. I analyze behavior from a subset of credit cards in the data that move from one sample city to another. I show that consumers who were exposed to a restaurant brand in their past state are 13% more likely to visit that chain than a consumer who moved from a state where the chain had no presence. The effect I measure is large enough to explain nearly half of the chain demand advantage. These results suggest that consumer familiarity and information—developed through branding and advertising—account for an important part of the chain premium.

The costs of standardization, both in terms of firm profits and consumer welfare, depend upon the degree of heterogeneity in preferences across markets. I identify this heterogeneity using variation in the relative popularity of chains that are available in multiple cities, which account for 40% of all restaurants in my sample. I find evidence of significant taste differences across markets along both horizontal (restaurant cuisine type) and vertical (willingness to pay for quality) dimensions. For example, the median consumer in Madison, WI prefers European restaurants with an average entree price of about $14, while a typical consumer in Cleveland, OH consumer would rather visit an American restaurant with an average price of $16.50. These across-city taste differences suggest that the standardization constraint may be costly for chains that operate in multiple markets.

I first quantify the tradeoff between the benefits of chain affiliation and the costs of standardization in terms of firm variable profits. I specify a model of restaurant supply in which firms choose product characteristics and prices. To quantify the costs of standardization, I ask how a given chain's demand would change if it could optimally choose its cuisine type and quality level separately in each city, but keep the demand advantage afforded by its chain affiliation. I find that the average large chain could earn 20% higher variable profits

[3] Because my data include only consumers with credit cards, I am not able to observe preferences for very low-income consumers. For example, the Federal Reserve reports that 22% of US adults were unbanked or underbanked in 2018, nearly all of which have annual household income of $40,000 or below (Federal Reserve Board of Governors, 2019). However, given that I find preferences for chains that are monotonically decreasing in income, my estimates are likely a lower bound for the chain preferences for consumers without access to credit markets. I discuss representativeness of my sample further in Section 2.

if it could flexibly choose its cuisine type and quality level in each market. However, chains also have important demand advantages over independents. If chains had to give up these advantages to ease the standardization constraint, they would lose an average of 30% of their variable profits. These estimates imply that, on net, the advantages of chain affiliation for the firm outweigh the costs of standardization for the chains in my sample.

Next I turn to quantifying the welfare impacts and distributional consequences of chain ban policies. In this framework, chain bans have two opposing effects. On one side, these policies remove a set of valuable chain firms, which decreases welfare for most consumers. The losses tend to be largest for consumers in the lowest income groups, who have the highest preferences for chains. On the other hand, the policy may result in a change in restaurant characteristics that improves the match between restaurants and local tastes. Moreover, because preferences are heterogeneous within a market, some consumers will benefit from this compositional shift while others will be harmed.

I simulate the long run effects of a chain ban by computing a new equilibrium in which chains are replaced by independents that endogenously choose their product characteristics. I compare consumer welfare under this counterfactual to a baseline in which each chain chooses a single set of product characteristics that is fixed across markets. I find that the policy decreases consumer surplus by about $12M, equivalent to 6.4% of sample spending, with 87% of the losses accruing to consumers in the two lowest income groups.[4] In particular, consumers with income less than $50,000 per year suffer a loss in welfare equivalent to more than 10% of their restaurant spending.

While this analysis focuses on restaurants, the same basic forces are present in standardized chains in other categories to varying degrees. Standardization is likely to be an important constraint in categories where differentiation happens at the retailer level as opposed to the product level. In clothing, for example, many retailers sell private label goods and cater to a particular segment of consumers.[5] Restaurants are an attractive setting to

[4]Concerns about the potential for chain bans to have outsize impacts on lower income groups have also surfaced in local policy debates. Dee Dee Workman, vice president of policy for the San Francisco Chamber of Commerce argues of the San Francisco ban on chain stores: "the people who have traditionally lived in that neighbourhood, who are lower-income and are hanging on by their fingernails, there's really nothing there for them...the formula retail ban is counter-productive." (Horgan, 2017)

[5]For instance, while Old Navy and Banana Republic share the same parent company, Old Navy sells less expensive clothing targeted towards younger shoppers, while Banana Republic sells higher priced products

study consumer preferences for chains for several reasons. First, product differentiation in restaurants is straightforward to measure using these data—Yelp records a measure of cuisine type, and I impute a restaurant's average price by combining menu data and information from the average dollar amount of its credit card transactions. Second, both chains and independent restaurants are available in the choice sets of most consumers; firms with only one location accounted for about a third of total restaurant sales, significantly more than other large consumer categories in the credit card data. This variation in chain affiliation is important, since identification of preferences for chain firms relies on variation in consumer choice patterns between chains and independents. Finally, restaurants make up a significant amount of consumer spending (approximately 14% of all transactions in the payment card data).

There is a broad empirical literature that studies the supply-side advantages of chains.[6] A primary contribution of this work is to propose and quantify a complementary demand-side channel to explain the importance of large firms in retail industries. My work is most closely related to Hollenbeck (2017), who studies the advantages of chains in the hotel industry. He finds that chain-affiliated hotels earn 20% higher revenues than similar independents, but that chains do not appear to have meaningful cost advantages. I estimate a similar demand premium for chains, but I study a setting in which horizontal differentiation is quantitatively more important, and show that the demand advantages of large firms are naturally limited by the costs of uniformity.

A second contribution of this paper is to quantify the welfare impacts and distributional consequences of local policy on firm entry. These policies are pervasive, and existing empirical literature on their effects is limited.[7] This work shows that one such policy, a ban on large

marketed to young professionals. While many clothing chains do customize the set of products across stores (for example, by offering more swimsuits and sandals in warmer climates as in Quan and Williams (2018)), they tend to sell clothes of a similar quality level across outlets.

[6] Papers that examine the cost advantages of chains include Doms et al. (2004), Holmes (2001), Foster et al. (2006), Jia (2008), Holmes (2011), Ellickson et al. (2013), and Houde et al. (2017). In the theoretical literature, Loertscher and Schneider (2011) and Cai and Obara (2009) examine the reputational benefits of chains in attracting consumers. Additional relevant empirical papers include Mazzeo (2004) and Hollenbeck and Giroldo (2020).

[7] Sadun (2015) studies the impact of chain ban policies in the UK on firm competition. Other papers that consider the impacts of local entry restrictions include Bertrand and Kramarz (2002), Viviano (2008), and Griffith and Harmgart (2012). Papers that study entry subsidies include Allcott et al. (2019), Mast (2020), and Slattery (2018).

chains, has substantial negative welfare effects and undesirable distributional consequences.

My work is also connected to a set of papers that estimate the degree of demand heterogeneity across places and the implications for retailers. DellaVigna and Gentzkow (2019) show that retail chains in the grocery, drug store, and mass merchandise categories set uniform prices across stores, despite significant heterogeneity in consumer demographics and competition across markets. Quan and Williams (2018) focus on the implications of across-market taste differences on the size of consumer benefits from online retail. In their context, retailers respond to this heterogeneity through extensive customization across outlets. In contrast, the degree of restaurant chain customization across places is relatively limited.[8]

This paper also connects to the work on preference externalities, the idea that nearby consumers with similar tastes benefit each other by increasing the market size for potential sellers (Waldfogel, 2003, 2010). The existing literature has focused on preference externalities as a local phenomenon; my work demonstrates that because of standardized national chains, consumers can affect each other across markets as well.

Finally, this works builds on several recent papers that study the geography of consumption. Davis et al. (2019) use data on Yelp reviews to analyze the role that spatial and social frictions play in consumer choice of restaurants in New York City. Eizenberg et al. (2021) use aggregated credit card data to study the role of spatial frictions in explaining differences in grocery prices across neighborhoods in Tel Aviv. I bring a new rich dataset to look at these questions. I use a similar empirical strategy to Davis et al. (2019), but focus on how consumers value chains and independent restaurants, and the effect of chain ban policies on consumer welfare.

The rest of the paper is organized as follows. In Section 2, I describe the data. Section 3 presents statistics about the aggregate importance of chains and facts about market structure in the restaurant industry. Section 4 introduces a model of consumer demand for restaurants and shows results from estimation. Section 5 shows the results of two sets of counterfactuals. Section 6 concludes.

[8] Adams and Williams (2019) also examines the effects of uniform pricing policies by chains. Other work that estimates the degree of demand heterogeneity across markets includes Bronnenberg et al. (2009), Bronnenberg et al. (2012), Choi and Bell (2011), and Waldfogel (2003).

2 Data

2.1 Sources

The analysis of retail chains leverages a novel dataset provided by a large payment cards company that includes the universe of credit and debit transactions on the network in 2016. Each observation in the underlying data is a transaction between a cardholder and a merchant. On the merchant side, the data contain the merchant identity and store, which are mapped to a business category and location. On the card side, each transaction is linked to a unique card identifier. The data contain no information on the specific goods or services that were purchased, nor the prices of those items. The sample is completely anonymized, and does not contain the name, address, or any other personally identifiable information about the cardholder. For about 50% of active credit cards in 2016, the company has access to a measure of estimated household income and the 9-digit billing zipcode of the card, which I use in some of the analysis. I am unable to link multiple accounts to the same individual, and thus treat each credit account as a separate consumer.

I supplement these data with additional data from Yelp and Pricelisto. Yelp is a major consumer review platform that allows users to review businesses and collects information about business category, hours, locations, and other characteristics. Yelp makes a sample publicly available for academic use that contains reviews and business attributes from a 2017 snapshot of all businesses in seven US cities: Las Vegas, NV; Phoenix, AZ; Charlotte, NC; Pittsburgh, PA; Cleveland, OH; Madison, WI; and Urbana-Champaign, IL. The data provide additional merchant information, including a more granular business category, price level, and average user rating. I merge the Yelp sample with the credit card data using the merchant name, zipcode and address for these seven cities to get additional restaurant attributes, including the cuisine type and hours. I report additional details of the merge in Appendix A.

Finally, the analysis uses a sample of data from Pricelisto, a website that aggregates restaurant menus and prices. The Pricelisto data contain prices collected in 2019 and 2020 at the menu item level for a subset of outlets belonging to twenty restaurant chains. I discuss these data in more detail in Appendix A.

2.2 Representativeness

The payment card dataset includes a significant share of total US consumption.[9] However, consumers that pay via cards may differ systematically in their restaurant choices from those that pay with cash. A 2018 survey from the payment processor TSYS suggests that consumers use cash in 15% of dine-in restaurant transactions and 32% of fast food transactions.[10] Further, while the broader data include both debit and credit transactions, the analysis sample is limited to credit transactions only (about half of sales in the payment card data), that contain the consumer's billing zipcode and household income, key covariates in the empirical analysis. If the restaurants visited by consumers who use credit cards are different from those that pay cash, then the estimates of restaurant preferences may be biased. For example, if independent restaurants tend to discourage credit card usage relative to chains, then the consumer preferences I recover for chains may be larger than the underlying preference in the population. On the other hand, if low income consumers are more likely to pay cash, and tend to visit chains more than independents, these preference parameters may be too small.

To assess the representativeness of the credit card data, I compare it to a sample of cell phone visits provided by Safegraph in the Houston area in 2018. Safegraph uses GPS data from a panel of smartphones to count the number of customers that visited local businesses. The cell phone data is unlikely to be subject to the same selection issues as credit card purchases; low income consumers are more likely to have smartphones than credit cards and consumers carry their smartphones even when they pay with cash.[11] The data use agreement prohibits merging these two datasets at the business level. Instead, I merge each of the sources to the Yelp dataset. This exercise shows that the market shares of chains, as well as the share of visits to restaurants of different Yelp price levels, look comparable across the two datasets, with the credit card data containing slightly more purchases at chains and

[9] Dolfen et al. (2019) report that spending on debit and credit cards issued by the provider made up about 20% of total US consumption in 2016.

[10] See https://www.creditcards.com/credit-card-news/payment-method-statistics-1276.php (all links in this paper were last referenced March 27, 2020).

[11] According to the 2018 Diary of Consumer Payments administered by the FRB of Atlanta, 75% of consumers with annual household income below $25,000 and 84% of those with income between $25,000 and $50,000 had smartphones, relative to 50% and 75% who had credit cards, respectively.

at higher priced restaurants.[12]

2.3 Variable construction

I lean on several data constructs throughout the paper to characterize merchants and cards. First, I define a consumer's location as the centroid of its 9-digit home billing zipcode. The primary demand model uses the distance between a consumer's home and a restaurant's location as an explanatory variable.[13]

To characterize each restaurant brand, or merchant, I construct three measures: the number of store locations within its brand, its restaurant cuisine type, and its average price. To construct the number of store locations l_m for each merchant m, I count the number of distinct store IDs belonging to a given merchant ID. I assign each restaurant merchant a cuisine type using the Yelp data. Yelp associates each restaurant with one or more tags that describe the type of food that it sells. I manually map this set of tags into eight categories and assign each restaurant to one of these. See Table A.3 for examples of popular tags in each category.

To construct a measure of restaurant price, I combine information from the payment card data and the Pricelisto menu price data. For each restaurant in the payment card data, I compute the average transaction size at the merchant level, defined as the sum of dollars divided by the sum of transactions. If the average quantity and composition of items purchased does not vary across restaurants, average transaction size will be perfectly correlated with price. Otherwise, average transaction size will capture variation in both prices and average quantities.

I compare average transaction size to an alternative measure of price from the Yelp data. Yelp assigns each restaurant a price rating between one and four dollar signs. If variation in average quantities is relatively large, the ranking between average transaction

[12]The share of large chains and the share of two dollar sign Yelp restaurants are about three percentage points higher in the credit card data relative to the smartphone data. I report additional details on this exercise in Appendix A.2.

[13]Consumers may not always travel from home when visiting a restaurant. To take this into account, I construct an alternate measure of card location based on the zipcodes in which the card transacts (using non-restaurant purchases). The two location measures are highly correlated, and my results are unaffected by which location is used in estimation.

size and the Yelp price measure may be reversed for some restaurants. Figure A.4 shows the the distribution of log ticket size by Yelp dollar sign rating. The four distributions are monotonically ordered by transaction size and largely non-overlapping. This suggests that average transaction size (and thus the imputed price) tends to preserve the relative price rankings of different restaurants.

However, average transaction price may distort the relative prices of restaurants. In particular, the ratio of average transaction size to price appears to be systematically higher for more expensive restaurants than for cheaper restaurants.[14] To correct for this, I calculate the ratio between average transaction size and the average price of a dinner entree for each chain in the menu data. This ratio varies between just under two for fast food restaurants to about 5 for fine dining chains. I then deflate average transaction size for each restaurant in the sample by this factor, which I allow to vary by average transaction size. I discuss this in more detail in Appendix A.

2.4 Sample construction

The primary analysis sample contains transactions made by cards active in 2016 with matched income and home billing zipcode data. Section 3 shows some basic facts about chains using all transactions that occurred in a set of augmented retail categories—in addition to traditional retail (corresponding to 2-digit NAICS codes 44 and 45), the sample includes restaurant and hotel transactions (NAICS 72). The largest categories within this group by total spending were restaurants (25%), grocery stores (16%), and general merchandise stores (13%).

In later sections, the analysis focuses on restaurants. In Section 4, I construct a sample of restaurant transactions made by consumers living in or near the seven US cities included in the Yelp data that transacted at a sample restaurant. I restrict my analysis to restaurants within city limits that are matched to the Yelp data. I eliminate cards with a billing zipcode further than 25 miles away from the nearest sample restaurant and keep only evening

[14] Lower priced restaurants may be more likely to have counter service rather than full service, may make it easier for groups of customers to pay individually, may attract parties of smaller sizes, and may sell fewer auxiliary food items per customer, such as drinks or dessert.

transactions (those made between 5pm and 11pm local time).[15] In estimating demand, I define the outside good as restaurant transactions by this sample of cards at restaurants that are outside of city limits, but within 25 miles of the consumer's home.[16] An observation in this dataset is a single transaction that occurred between a card and a restaurant. I show summary statistics for this sample by city in Table 1a. In total, the sample included over a million accounts who made about 21M transactions for $768M dollars, spread over about 4000 restaurants (plus those aggregated in the outside good) across the seven cities.

Table 1b shows card-level activity measures by bins of estimated household income. Consumers in the highest income group (>$200,000 per year) make about 40% more restaurant transactions than those in the lowest group (<$50,000). Higher income groups also tend to go to higher-priced restaurants. The average restaurant visited by the top income group had an average price of about $16, compared to about $13 for the lowest income group.

One group of interest in this paper will be consumers that move from one state to another. I identify movers by constructing a panel of all active credit cards that change their billing zipcode between 2017 and 2020.[17] Each mover is associated with an origin state and a destination state. I define a consumer's move date as the first month in which the account reports a billing address in the destination state. For this subsample of movers, I keep all evening restaurant transactions between 2017 and 2020. I use the behavior of these cards to investigate possible mechanisms for the chain demand advantage.

3 Chains vs. independent stores: empirical patterns

3.1 Importance of chains

To motivate the focus on retail chains, I first document the importance of chains across merchant categories. Figure 1a shows the share of aggregate spending by firm size. Less

[15] To eliminate restaurants that are not truly available to consumers at dinner, I require that a restaurant have at least 100 total transactions in a year and be marked as open by Yelp during dinner hours for at least four days of the week. I provide additional details in Appendix A.

[16] I calculate the number of outside good restaurants available to each consumer within a 2, 5, 10, and 25 mile radius of each consumer's home and use this vector as an explanatory variable in estimation.

[17] Consumer billing zipcodes are reported yearly between 2017, 2018, and 2019, and quarterly in 2020. I verify that these consumers live in their reported states by eliminating cards that primarily transact in some other state. See Appendix A.7 for additional details on construction of this sample.

than 20% of spending in retail categories is made up by single-establishment firms, with the rest spent at chains. In particular, nearly 40% of spending occurred at large, national chains with more than 1000 locations.

Figure 1b shows the breakdown of spending by firm size separately for four large merchant categories in the payment card data. Among the four, restaurants is the most fragmented, with about a third of spending at single-establishment firms, while in the general merchandise sector, chains accounted for over 98% of all spending.[18]

3.2 Characteristics of chain restaurants

I illustrate several stylized facts about chain and independent restaurants to motivate the model of demand in Section 4.2. First, large chain restaurants have lower average prices than independent restaurants. Figure 2a shows the mean and quantiles of average price separately for independents and restaurant chains of different sizes, with each merchant weighted by its number of transactions in 2016. The mean and median price are monotonically decreasing in the number of locations. A dinner entree at an average independent restaurant has a price of about \$15, while chains with more than 1000 locations have an average price of about \$7. The figure also shows that the distribution of prices for these large chains is relatively compressed—the difference between the 25th and 75th percentile is only \$2, compared to about \$9 for independents.

Second, independent restaurants and small chains tend to have a larger share among higher income consumers. Figure 2b plots the share of sample transactions that went to restaurant chains of different sizes by cardholder annual income. Consumers in the highest income group (those with annual income above \$200k), conducted 25% of their restaurant transactions at chains in the two largest size groups (100-1000 locations or more than 1000 locations), relative to 33% of transactions for the lowest income group.

Together, these facts indicate that while chain restaurants are generally popular, there is a large fringe of independents that continue to operate. Chains also tend to target

[18] Other work in the literature has documented the growth of chains during the 20th century. Analysis of longitudinal Census data shows that chains accounted for less than 30% of retail sales in 1948. That figure rose to 40% by 1976, and over 60% by 1997, with most of the growth at firms with more than 100 locations (Jarmin et al., 2009).

different types of consumers than independents. Chains are most popular among lower-income consumers who may be more price sensitive and more receptive to advertising (Dutta-Bergman, 2006). This heterogeneity suggests that consumer preferences for chains may be heterogeneous across demographic groups. It also suggests that some chains may have room to increase profits in some markets if they could customize—for example, by offering higher quality food in markets in which higher income consumers are more prevalent.

4 Demand and supply for restaurants

Now I turn to quantifying the revenue and profit implications of chains. I write down a model of demand and supply for restaurants to quantify the fundamental demand-side trade-off faced by a firm considering standardization: standardized chains have important demand advantages, but the product attributes they choose may not be optimal in a particular market. In the model, each restaurant plays a two-stage game in which they first choose product attributes (quality and cuisine type) which are fixed across markets for each firm, and then set prices to maximize variable profits.

I estimate the parameters of demand using a sample of restaurant transactions from consumers in seven mid-size US cities described in Section 2 to recover preferences for chains and for these two product characteristics separately in each city. Differences in consumer tastes across cities measure the degree of heterogeneity in demand. I show suggestive evidence that consumer familiarity with chains is driving part of the demand advantage. I then use the recovered parameters to conduct counterfactual analysis: I estimate the degree to which a chain could increase its variable profit by choosing its cuisine type and quality level flexibly in each city, and assess the impacts and distributional consequences of a ban on large chain restaurants.

4.1 Descriptive evidence

To illustrate some of the data patterns that drive my demand results, this section shows model-free evidence of several of the most important findings. First, the distribution of consumer preferences is heterogeneous across markets. The key variation that identifies the cost of standardization is the extent to which consumers in different cities prefer different types of restaurants. If demand is homogeneous across markets, the optimal choice of product characteristics (quality and cuisine type) is the same in each market, and the standardization constraint is not binding. In my sample, I find evidence of important heterogeneity across cities in tastes for both quality and cuisine type.

I first show that different types of restaurants are popular in different cities. To control for unobserved restaurant heterogeneity, I focus on the relative popularity of chains with locations in more than one sample city conditional on the chain's average popularity. This is the same variation that drives the estimates of taste differences across markets in the demand model. For each restaurant, I calculate its average number of transactions and divide it by the city average. I then regress this on a set of chain fixed effects and save the residuals for each chain restaurant.[19] Figure 3a shows the residuals from this regression averaged across restaurants of different price levels, normalized to one in each city. The magnitude of this measure can be interpreted as the performance for restaurants in a given price category compared to an average restaurant in that city, relative to the popularity of restaurants in the same chain in other cities. The figure shows that restaurants with average prices less than $10 are much more successful in Madison, which is a college town with many students. Madison restaurants in the lowest price category received about 17% more transactions than the average restaurant, while restaurants in the highest price category received about 30% less than the average restaurant. In Pittsburgh and Cleveland, higher end restaurants with prices above $15 tend to be most successful (about 5% more successful than the average, while restaurants in the cheapest category received 7-10% fewer transactions than the average).

The data show a similar degree of heterogeneity across restaurants of different cuisine types. Figure 3b shows the average (residualized) transactions for restaurants of different cuisine types across sample cities. In Las Vegas, Burger chains were very popular (+27% relative to the average), compared to Asian or Latin-American chains (-23% and -20% respectively). In Phoenix and Charlotte, Asian and Latin-American restaurants performed better than Burger restaurants. This variability in the success of restaurants of different

[19] 1,420 sample restaurants belonged to chains with at least two locations in the sample, out of a total of 3,944 restaurants.

types suggests that a chain that could be fully flexible in its choice of product characteristics across markets might choose to sell different types of food at different prices in each city.

Finally, consumers in my sample are quite sensitive to distance. In total, over 70% of sample transactions were at restaurants that were less than five miles away from the consumer. In Figure A.7 in the Appendix, I show the probability that a restaurant is chosen as a function of the distance between a consumer's home billing zipcode and the restaurant, aggregated across consumers and cities. The figure shows that a given restaurant is about half as likely to be chosen at two miles away versus one mile away, and an additional 40% less at three miles versus two miles. This sensitivity to distance combined with variation in the spatial distribution of consumers and restaurants provides useful identifying variation in estimating demand.

4.2 Econometric specification

4.2.1 Consumers

In the main demand specification, restaurants are differentiated in their cuisine type, quality level, price and distance from consumers' homes. Each restaurant can take one of eight cuisine types.[20] Preferences over cuisine types vary by consumer income group and city. I impute each restaurant's price as described in Section 2; the demand system allows for unobserved heterogeneity in consumer preferences over price within a city-income group. Restaurant quality is unobserved, but is assumed to be constant across restaurants within a brand. Each consumer's choice also depends on her distance in physical space from each restaurant. A consumer chooses between visiting a restaurant within the city or some other restaurant within 25 miles of her home; transactions at these other restaurants are aggregated into an outside option.

The utility that consumer i with income y in city c receives from visiting restaurant j belonging to brand m in trip t is given by:

$$u_{ijt} = x_j\beta_{cy} - \gamma_{cy} dist_{ij} + \alpha_i^l \log l_m + \alpha^q \log q_m - \alpha_i^p p_m + \xi_m + \Delta\xi_{jyc} + \varepsilon_{ijt} \tag{1}$$

[20] Table A.7 shows summary statistics on the number of restaurants in each cuisine type.

x_j is a vector of restaurant characteristics that includes a set of cuisine type and chain size dummy variables.[21] $dist_{ij}$ is the distance in miles between restaurant j and consumer i's 9-digit billing zipcode. $\log l_m$ is the natural log of the number of locations belonging to brand m, which captures variation within a city-income group in preference for chains. q_m and p_m are the quality and average price of restaurant brand m. Quality q_m is a measure of vertical differentiation of the bundle of inputs used by m. All consumers prefer high quality to low quality, with diminishing marginal returns, but differ in their willingness to pay. Finally, ξ_m is the unobservable product characteristic for brand m, $\Delta\xi_{jyc}$ is the restaurant-income group-city deviation from brand level unobservable ξ_m, and ε_{ijt} is a random preference shock that is distributed iid extreme value type I.

The outside good for each consumer i is defined as a visit to a restaurant outside of city limits, but within 25 miles of the consumer's home. The utility of choosing the outside option is given by:

$$u_{i0t} = v_{i0} + \log R_i \pi^0_{cy} + \varepsilon_{i0t} \tag{2}$$

v_{i0} is an intercept term that captures i's unobserved preference for the outside option. $\log R_i$ is a vector that contains the natural logarithm of the number of outside good restaurants within four distance bands around consumer i's home.[22]

I assume that the random parameter vector for each consumer $v_i \equiv (\alpha^l_i, \alpha^p_i, v_{i0})$ is distributed multivariate normal with a mean and covariance matrix that varies flexibly across cities c and income groups y:[23]

$$\begin{pmatrix} \alpha^l_i \\ \alpha^p_i \\ v_{i0} \end{pmatrix} \sim \mathcal{N}\left[\begin{pmatrix} 0 \\ \bar{\alpha}^p_{cy} \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^{\alpha^l}_{cy} & \rho_{cy} & 0 \\ \rho_{cy} & \sigma^{\alpha^p}_{cy} & 0 \\ 0 & 0 & \sigma^v_{cy} \end{pmatrix}\right]$$

[21] I separate restaurant brands into four bins based on the number of US locations belonging to that brand: 1 location, 2-100 locations, 101-1000 locations, and more than 1000 locations.

[22] R_i includes the log of the number of open restaurants within 2, 5, 10, and 25 miles.

[23] I normalize α^l_i to have a mean of zero and include fixed effects for different chain sizes in x_j to allow additional flexibility in preferences over chain sizes.

The goal of the demand estimation is to quantify the degree of heterogeneity in consumer preferences across different cities. As such, the specification given in equations (1) and (2) is chosen to allow for as much flexibility as is feasible across groups. I leverage the richness of the credit card transaction data by allowing for heterogeneity in preference parameters across cities and income groups over price, chain affiliation, travel time, and restaurant cuisine. Within a city-income group, equation (1) allows for unobserved heterogeneity in the form of a random coefficient on chain preferences, price sensitivity, and preference for the outside option.

Preferences for other attributes, such as cuisine type, may also be heterogeneous within a city-income group. However, Monte Carlo experiments with simulated data suggest that these are not empirically identified.[24] Estimation of additional random coefficients also significantly increases the computational burden.

Relative to studies that use aggregate data, the transaction-level data used here offer two advantages. First, they contain panel-like information about which sequences of choices were made by which individuals. This allows the estimation of a richer heterogeneity structure than could be identified from only aggregate market shares, which depends on variation in choice sets across markets and over time.[25] Second, my data identify where individual consumers live, which generates useful idiosyncratic variation in travel cost within a market.

A natural identification concern is that restaurants in different cities may be of different unobserved quality levels that are correlated with observable characteristics. For example, if the data show that burger restaurants are popular in Pittsburgh, it could be the case that Pittsburgh consumers love burgers or that the burger places in Pittsburgh are especially good and would be equally successful if they were accessible to consumers in other cities. The estimation strategy will control for this by including brand-specific fixed effects, under the assumption that this set of unobservables is fixed within a chain. Thus, the identification of preference heterogeneity comes only from differences in the relative popularity of chains

[24]This is consistent with the finding of Sebastien (2008) that random coefficients on categorical variables are not identified without variation in choice sets over time.

[25]For example, I allow the variance and covariance parameters to vary freely for every city-income group combination. Studies that use aggregate data are typically only able to allow parameters to vary in a restricted way across markets (for example, variance parameters may vary linearly with aggregate demographics as in Nevo (2001)). See additional discussion in Berry et al. (2004) and Hess and Train (2011).

that operate in multiple cities. Chains with restaurants in at least two sample cities account for about 40% of total sample restaurants.

A related concern in many settings is that price may be correlated with unobserved product quality, even conditional on a brand fixed effect. If a branch of a chain is especially popular in a market (i.e. it has a high $\Delta\xi_{jyc}$), it may charge higher prices in equilibrium. Because I use a measure of price that is imputed from average transaction size and likely contains measurement error, I use average prices computed at the brand level, rather than at the restaurant level. As such, price cannot be correlated with $\Delta\xi_{jyc}$ in this setting, conditional on brand fixed effects.

However, since neither cuisine type nor price vary within a chain, the mean level of preferences across all cities is not separately identified from the brand fixed effects. I proceed as in Nevo (2001) by first recovering estimates of preference heterogeneity across cities and then regressing the estimated brand fixed effects on the vector of product characteristics. I describe this procedure in more detail in Section 4.3.

Quality q_m is an unobserved product attribute that describes the bundle of a restaurant's raw ingredients and service. I back out quality from a brand's estimated marginal cost. A fine dining restaurant that serves filet mignon and employs a sommelier will have a higher level of q_m than a fast food burger joint that cooks frozen hamburger patties and offers counter service. Because q_m is fixed within a brand m, it does not bias the estimates of preference heterogeneity across cities (conditional on brand fixed effects). While the parameter α^q is fixed across individuals, this is just a normalization, as the price coefficient varies flexibly across and within groups. I assume that valuation for quality is concave while distaste for price is linear; this ensures every consumer prefers a finite level of quality given the cost.

4.2.2 Restaurants

I assume that restaurant brands play a two stage game. In the first stage, restaurant firms make choices over two product attributes, cuisine type and quality. In the second stage, restaurants play a Nash-Bertrand pricing game and compete for consumers. The profits for brand m are given by the sum of profits coming from restaurants in the brand J_m:

$$\pi_m = \sum_{j \in J_m} N_c \, s_j(p_m, q_m) \Big(p_m - mc_m(q_m)\Big) \tag{3}$$

where N_c is the number of transactions in j's city c, $mc(q_m)$ is the marginal cost for brand m of producing quality q_m, and s_j is the market share of j in c. A restaurant's quality level q_m is a composite measure of vertical differentiation that captures differences in inputs across firms. Marginal cost mc_m is increasing in quality; I assume that cost and quality are linearly related:

$$mc_m(q_m) = \zeta_m q_m$$

Quality is unobserved in the data, and so ζ_m is not separately identified from the other parameters of consumer demand. I proceed by normalizing $\zeta_m = 1$ for all m. If part of the appeal of chains is that they are able to more efficiently produce quality (i.e. chains have lower ζ_m than do independents), then this will appear as part of the estimated demand premium for chains, contained in β.

I assume that product attribute decisions are made at the brand level, rather than by each restaurant. While many restaurant chains have franchised outlets, contractual terms typically restrict franchisees from changing restaurant menus.[26]

4.3 Estimation

I proceed with estimation in three stages. I first estimate the nonlinear parameters of demand using maximum simulated likelihood allowing for correlation of choices within an individual (Revelt and Train, 1998). For each guess of the parameters, I solve for the mean utility for each restaurant that matches its observed market share, as in Berry et al. (1995). I treat each city-income group as a market, which allows me to heavily parallelize estimation. I rewrite equation (1) as

[26] For example, a recent franchise agreement for McDonald's specifies that a franchisee "may sell only products authorized by McDonald's," using "packaging, paper goods, ingredients, and handling and preparation methods that meet the McDonald's System specifications and quality standards." Source: https://www.bluemaumau.org/sites/default/files/MCD%202013%20FDD.pdf

$$u_{ijt} = \delta_{jcy}(x_j, p_m, q_m, \xi_m, \Delta\xi_{jyc}; \theta_1) + \mu_{ijt}(l_m, p_m, dist_{ij}; \theta_2) + \varepsilon_{ijt}$$

where δ_{jcy} is the mean utility of restaurant j in city c and income group y and μ_{ijt} is consumer i's deviation from that mean utility:

$$\delta_{jcy} = x_j\beta_{cy} + \alpha^q \log q_m - \bar{\alpha}^p_{cy} p_m + \xi_m + \Delta\xi_{jyc}$$
$$\mu_{ijt} = -\gamma_{cy} dist_{ij} + \alpha^l_i \log l_m - \tilde{\alpha}^p_i p_m$$
$$\tilde{\alpha}^p_i = \alpha^p_i - \bar{\alpha}^p_{cy}$$

$\theta_1 \equiv (\beta_{cy}, \alpha^q, \bar{\alpha}^p_{cy})$ collects the parameters that are absorbed by δ_{jcy} and $\theta_2 \equiv (\gamma_{cy}, \sigma^{\alpha^l}_{cy}, \sigma^{\alpha^p}_{cy}, \rho_{cy}, \sigma^v_{cy}, \pi^0_{cy})$ contains the parameters that enter the likelihood in a nonlinear way.

From equation (1), let $V_{ijt} \equiv u_{ijt} - \varepsilon_{ijt}$. Conditional on the vector of random coefficients v_i, the probability that consumer i chooses restaurant j in trip t is given by:

$$P_{ijt}(v_i) = P(y_{it} = j|v_i) = \frac{\exp V_{ijt}}{\sum_{j'} \exp V_{ij't}}$$

The conditional probability of observing a sequence of choices in different time periods $\boldsymbol{m} = m_1, \dots, m_T$ by consumer i is:

$$P_{i\boldsymbol{m}}(v_i) = \prod_t P_{im_t t}(v_i)$$

and the unconditional probability of the sequence $\boldsymbol{m}$ is:

$$L_{i\boldsymbol{m}} = \int P_{i\boldsymbol{m}}(v_i) f(v_i) dv_i \tag{4}$$

I approximate the integral in equation (4) with 50 scrambled Halton draws, indexed by r, that are fixed within a consumer. The log likelihood is given by:

$$\log L(\theta_2, \delta) = \sum_i \log(\frac{1}{50} \sum_r P_{i\boldsymbol{m}r}) \tag{5}$$

where $P_{imr}(v_i^r) = \prod_t \frac{\exp V_{ijt}(v_i^r)}{\sum_{j'} \exp V_{ij't}(v_i^r)}$.

In the second stage, I regress the vector of mean utilities δ_{jcy} on x_j and p_j interacted with city and income group dummies, a full set of brand fixed effects δ_m, and city-income group fixed effects. This recovers the market-level taste for restaurant characteristics in each city and income group, relative to the mean utility across all markets, which is absorbed by δ_m.

The parameters that describe mean utility across markets are not identified separately from brand fixed effects with a single cross section of data (Berry et al., 2004). The most challenging coefficient to identify is price, which is likely to be correlated with unobserved product characteristics q_m and ξ_m. Estimation via instrumental variables requires an instrument that is correlated with price but not with unobserved quality. Traditional cost-shifting instruments, such as wholesale food prices or wages, do not vary within a market. Variables that might affect markups, such as the number of competing restaurants nearby, have little predictive power on prices relative to the unobservables (and may themselves be endogenous if restaurants cluster in high demand areas).

Another approach would be to use variation in restaurant prices over time due to changes in input costs. Unfortunately, price changes in the data are very small from year-to-year—between 2016 and 2018, average restaurant transaction sizes increased by about 3% annually. Further, there is very little variation in relative prices during this time period. This issue is compounded by the scope for measurement error in the imputed prices used in estimation.

Instead, I take an approach similar to that suggested in Berry et al. (2004) and match moments from external data to calibrate this parameter. I use the firms' first order pricing condition to back out marginal costs[27] and then match them to data I collect for a set of 12 publicly traded restaurant chains. On average, these firms report that the sum of labor and food costs account for 61% of their total revenues at company-operated restaurants. I choose the price coefficient to match this cost share for these 12 firms. The mean level of the price parameter does not materially affect the results of any of the counterfactuals. Intuitively, the returns to operating a standardized chain relative to a set of customized restaurants

[27]The vector of restaurant-level marginal costs for city c is given by $\boldsymbol{mc}_c = \boldsymbol{p}_c - \Delta_c(\boldsymbol{p}_c)^{-1}\boldsymbol{s}_c(\boldsymbol{p}_c)$ where $\Delta_c(\boldsymbol{p}_c) = -\boldsymbol{H}_c \odot \frac{d\boldsymbol{s}_c}{d\boldsymbol{p}_c}(\boldsymbol{p}_c)$, the element-wise product of the matrix of demand derivatives and the $J_c \times J_c$ ownership matrix $\boldsymbol{H}_c$, which each (j, k) entry equals one if j and k belong to the same chain. I compute marginal costs at the restaurant level and take the average for each brand m.

depend on the dispersion of tastes, rather than the absolute level, and thus this parameter is of secondary importance.

After calibrating the mean price sensitivity to match empirical data on marginal costs, I recover the remaining parameters of mean utility in the third stage by regressing the brand fixed effects δ_m, less the part related to price $\alpha^p p_m$, on the vector of product characteristics x_m and quality $\log q_m$. I compute standard errors using a block bootstrap procedure with 50 replications. I discuss additional details of the estimation procedure in Appendix B.

4.4 Results

I present estimates of a subset of the parameters by city and income group in Table 2 (see Tables B.10-B.12 for the full set of 543 parameters with bootstrapped standard errors). Estimates of γ, the utility from traveling one mile, and $\bar{\alpha}^p$, the utility from a \$1 increase in average entree price, are negative in every city and income group, as expected. Lower income groups are significantly more price sensitive and slightly more sensitive to travel distance than are higher income groups. α^q, the parameter on the natural log of quality which is assumed to be constant across cities and income groups, is estimated to be 4.31, suggesting that the average consumer is willing to pay about \$1.30 for a 10% increase in quality.

The key parameters for quantifying the costs and benefits of chains are the extent to which consumers value chains relative to independents with similar product characteristics and the degree of preference heterogeneity across cities. The estimates of β^{1001+} in panel 2a suggest that large chains face significantly higher demand than independent firms or small chains in five out of seven markets. Chains enjoy the largest demand advantages in Madison, Phoenix and Champaign, and are actually less popular than independents in Cleveland and Pittsburgh.

Table 2b reveals that preferences for chains are highly heterogeneous across income groups. Low income groups have large, positive tastes for chains, while consumers in the highest income group prefer independents. To quantify the magnitude of these coefficients, I convert the vector of chain preferences for each income group to "mile equivalents" by dividing each element of β by the corresponding γ_{cy} for that city-income group.[28] I plot the chain premium by income group (averaged across cities) in Figure 4b. The chart shows that consumers with household income below \$50,000 and between \$50,000 and \$100,000 are indifferent between eating at an independent restaurant and traveling an additional 1.2 and 0.7 miles to eat at a large chain, respectively. Consumers in the two highest income bins, however, have much weaker preferences for chains; those with income levels above \$200,000 prefer independents by about 1.2 miles. This chain premium is positive and statistically different from zero at the 5% confidence level in 24 out of 35 city-income group markets (see Table B.13). The size of the chain premium, and the gradient across income levels, have important implications for the counterfactual exercises later in the paper. Policies that ban chains are likely to have negative effects on lower income groups, but may benefit high income consumers.

In addition to the differences in preference over chains along observable margins described above, there is also substantial unobserved heterogeneity within city-income groups. The standard deviation of α_i^l, the random coefficient on the log of the number of firm locations, is about 0.194 averaged across all markets, which is nearly three times as large as the mean preference in even the lowest income group.[29]

Consumers may value chains differently from independents for a variety of reasons; chains may offer better or more predictable food, faster service, larger menus, or a more consistent experience, among other things. The data do not allow me to estimate consumer valuations separately for each of these factors. In Section 4.5, I show that at least part of the demand for chains is related to past consumer exposure to a chain's branding.

By comparison, the degree of unobserved heterogeneity in price sensitivity is more modest, but still quantitatively important. The standard deviation of α_i^p is between 0.04 and 0.06 for each income group, which is about 15% of the mean price sensitivity parameter, implying that the difference in price sensitivity between consumers in the 90th percentile relative to

[28] Alternatively, the coefficients could be expressed relative to the parameter on price. However, because the prices used in estimation are for a single entree, and a typical transaction will contain more than one unit, this does not map directly to money metric utility, and must be scaled up to account for average quantities.

[29] A standard deviation of $\alpha_i^l = 0.194$ implies an increase in utility from chains relative to independents of $0.194 \times (\log(1000) - \log(1)) = 1.34$. The mean preference for chains in the lowest income group is 0.454. $1.34/0.454 = 2.95$.

those in the 10th percentile within a city-income group is about 40% of the mean.[30]

Table 2 also shows that ρ is negative in every market and for every income group, suggesting that the segment of consumers who like chains also tend to be the most price sensitive. This is consistent with the empirical pattern that chains tend to set lower prices on average relative to independents (Figure 2a). This also has important implications for the counterfactuals, because it suggests that under a chain ban, the set of independents that would enter might choose higher levels of price and quality than those chosen by chains.

The demand estimates also imply significant heterogeneity in consumer tastes over price and cuisine. The average price sensitivity parameter in the most sensitive market (Madison) is about 25% larger than in the least sensitive (Cleveland). An average consumer in Madison would be in the 97th percentile of the price sensitivity distribution in in Cleveland.[31] This suggests that a restaurant chain that optimizes for the population in Madison will not be well suited to the preferences of most Cleveland consumers.

To better illustrate how this heterogeneity impacts the optimal choices of firms, I calculate the utility that an average consumer in each market receives from restaurants of different price and quality levels. Figure 4a shows that the average consumer in Cleveland prefers a restaurant with an average meal price of about \$19, while an average Madison consumer prefers a meal price of around \$15.

There is also important taste heterogeneity in preferences over restaurant categories across markets. In Figure 5, I show a heat map containing the rankings of preferences of different cuisine types by city (averaged across income groups within a city). In Champaign, Charlotte, Las Vegas and Madison, the most popular cuisine category was European, but this was one of the least popular categories in Phoenix and Cleveland. In Phoenix and Cleveland, the most popular category was American.

[30] The difference between the price sensitivity at the 10th and 90th percentile is $\bar{\alpha}^p + \phi^{-1}(0.9) \times \sqrt{\sigma^{\alpha^p}} - (\bar{\alpha}^p + \phi^{-1}(0.1) \times \sqrt{\sigma^{\alpha^p}}) = 2 \times 1.282 \times \sigma^{\alpha^p} = 0.128$ evaluated at $\sigma^{\alpha^p} = 0.049$, the mean standard deviation across all markets. As a share of the mean price coefficient -0.330, this difference is $0.126/0.330 = 0.381$.

[31] The difference between $\bar{\alpha}^p$ in Cleveland and Madison is $0.357 - 0.280 = 0.077$. The standard deviation of the random price coefficient in Cleveland is $\sqrt{0.0016} = 0.04$. $0.077/0.04 = 1.925$, and $\Phi(-1.925) = 0.027$.

4.5 Mechanisms for chain demand advantage

My demand estimates imply that chain restaurants have a large and important demand advantage over smaller firms. However, the mechanism for this advantage is unclear. I consider two classes of channels. First, chains may have informational advantages through branding and advertising; risk-averse consumers may get higher expected utility from chains because they have more information about their product (Erdem and Keane, 1996). Second, upon visiting the restaurant, consumers may enjoy the food or experience more at a chain relative to a similar independent (for example, because of selection into which restaurants grow into chains or because chains provide higher quality food at a given price). In this section, I provide suggestive evidence that information and branding are an important part of the effect that I measure.[32]

4.5.1 Evidence from movers

I first show evidence from a sample of consumers that move across state lines. I identify about 295,000 movers between 2017 and 2020 that made about 12.7M restaurant transactions (see Appendix A for additional details on this sample). Movers are a particularly useful population to study in this context; moving to a new state induces a sharp change in a consumer's choice set, while her information and accumulated advertising exposure adjust more slowly.

Figure 6a shows the share of restaurant transactions that occur at chains and independents around the date that the consumer moves. Immediately after moving, a consumer is likely to have less information about the set of nearby restaurants, and thus the chain reputation may be particularly valuable. The figure shows that in the first year after moving, the share of transactions at large chains increases by about seven percentage points on a base of 30%. Most of this increase comes at the expense of independent restaurants, whose share of transactions falls by about four percentage points on a base of 31%.

[32] Existing evidence suggests that incomplete information may be an important reason why some restaurants are more successful than others. For example, Luca (2011) finds evidence that favorable restaurant reviews on Yelp lead to significantly higher revenue for independent restaurants, but do not affect chains, about which consumers conceivably have more familiarity, and thus are less likely to update their priors based on internet reviews.

These observed changes in spending patterns after moving may be driven in part by differences in the choice set between the origin and destination states, or by compositional changes in which cards tend to visit restaurants before and after moving. To rule out these alternative explanations, I compute the share of transactions at chains and independents at the card-month level and regress each share on card and current-state fixed effects, plus a set of dummies for months relative to move date. Figure A.8 plots the coefficients on the months since move dummies, which confirms the pattern and magnitudes in Figure 6a. This exercise suggests that chains are more valuable in an environment where consumers have less information. However, because moving may also result in changes to a consumer's job, income, or other habits, it is difficult to fully rule out the importance of other factors.

To further isolate and measure the effect of information, I focus on the propensity of movers to visit regional chains after moving. Because consumers move from different states, they have differential past exposure to regional chains depending on whether the chain has a presence in their origin state.[33] For this exercise, I use the set of post-move transactions that occur in one of the seven sample cities used in the main estimation. I provide summary statistics on this subsample in Table A.8. For each consumer, I classify each restaurant in the city by whether it belonged to a chain that had at least one location in the consumer's origin state. I estimate a simple multinomial logit regression with restaurant-income group-city fixed effects, pooling data across cities:

$$u_{ijt} = \delta_{jyc} + \beta Avail_{ij} + \varepsilon_{ijt} \tag{6}$$

where cards are indexed by i, restaurant outlets by j, y is the income group of consumer i, and $Avail_{ij} = 1$ if j's brand m had locations in i's origin state. Identification of β in equation 6 comes from variation in how consumers in the same income group that move to the same city from different states visit a restaurant as a function of whether the chain had a presence in their origin state. If the sole driver of the demand advantage is an unobserved quality dimension common to large firms, then consumers should be as likely to visit familiar

[33] Many chains have locations only in a subset of regions, particularly in the middle size buckets (2-100 locations and 101-1000 locations). I show statistics on the geographic footprint of chains of different sizes in Table A.9.

chains as unfamiliar ones. If the main driver of the chain premium is related to information communicated through the brand, restaurants that were unavailable in the consumer's prior city should have a smaller demand premium after moving.

Estimation of equation 6 via maximum likelihood yields an estimate of β of 0.123, with a bootstrapped standard error of 0.029 (see additional details of estimation in Appendix B). This implies that a restaurant that had a location in the mover's origin state is about 13% more likely to be chosen than a restaurant that was unavailable.[34] A back-of-the-envelope calculation using the estimates from the full model implies that the familiarity effect is large enough to account for 48% of the demand premium for large chains relative to independents.[35]

The impact of past familiarity on consumption suggests that at least part of the chain demand premium is related to information and branding. If instead the preference for chains was related solely to cost or quality differences that operate through the chain's scale, we should expect demand to be similar for both familiar and unfamiliar chains.[36]

4.5.2 Evidence from new restaurant entries

As a final exercise to investigate the mechanisms of the chain demand advantage, I show evidence from new restaurant entries. Using the credit card data, I construct a sample of new restaurants that enter between 2015 and 2018 in the seven sample cities.[37] I compute monthly sales for each new restaurant and normalize it to the average monthly sales for that restaurant in in its first year after opening. I plot the evolution of entrant sales against time

[34]The ratio of choice probabilities for a restaurant j that had a presence in the origin state of consumer i ($Avail_{ij} = 1$) but not in the origin state of i' ($Avail_{i'j} = 0$) is $\frac{P(y_{it}=j)}{P(y_{i't}=j)} \approx \exp(0.123) = 1.13$.

[35]Averaged across income groups and cities, the demand premium for chains with more than 1000 locations relative to independents is 0.259. 0.123/0.259 = 47.5%

[36]Two other facts suggest that the advantage is related to advertising and branding. First, the largest restaurant chains advertise extensively. In 2018, 8 of the 25 most-advertised brands on television were restaurant chains, per reports by IdenTV, with McDonalds spending approximately $1.5B on US advertising in 2017, or about 25% of its US revenue, according to Ad Age. Second, this interpretation is broadly consistent with the findings of the literature on branding. In particular, see Hollenbeck (2017) and Tsai et al. (2015) who analyze the hotel industry. Hollenbeck (2017) finds that a hotel that switches from independent to chain gets a 21% increase in their revenue, while Tsai et al. (2015) find that hotels that rebrand from one chain to another increase their revenues about 4%.

[37]I consider the first date that a restaurant begins to report transactions its opening date. To verify that these are real entries, I merge the sample with the Yelp data and eliminate entries where a restaurant's first Yelp review does not occur within three months of the first credit card transaction.

since open by chain size in Figure 6b.[38]

The figure shows that new chain restaurants (>1000 locations and 101-1000 locations) nearly reach their steady state level of sales in the first year and grow relatively little between year 1 and year 3—they end the third year about 20% higher than their first year average. Independents and small chains with 2-100 locations, in contrast, grow steadily over the first 3 years; by the end of year 3, they are generating 80% and 60% more revenue than their first year averages, respectively. Independent restaurants also tend to exit at higher rates than do chains. These patterns point to the importance of consumer learning for independent restaurants, who build their customer base over multiple years, while new chain establishments benefit from the existing chain reputation and are able to quickly acquire sales.

5 Counterfactuals

I use the demand estimates described above to quantify the revenue and welfare impacts of chains. This section contains two sets of counterfactuals. The first set analyzes the decision of a firm considering whether to operate a standardized chain or a network of customized restaurants. Restaurants affiliated with chains have higher demand, but they must pick a single quality level and cuisine type that is fixed across markets. In these counterfactuals, I consider the effect of a unilateral change in product characteristics by one firm at a time and hold the behavior of competing restaurants fixed. My results suggest that the benefits of chain affiliation are substantially larger than the costs of standardization in my sample markets.

The second set of counterfactuals quantifies the welfare impact of a chain ban on consumers. Policymakers face a similar tradeoff to that of the firm—a ban on chain restaurants removes a set of firms that many consumers like, but may improve the match between restaurant characteristics and local preferences. In the primary counterfactual, I assume that large standardized chains are replaced by independents that choose their quality

[38] The entries in Figure 6b include both firms that survive and those that don't. Independents and small chains tend to exit at higher rates than do chains.

and cuisine type to maximize their variable profits. I find that the equilibrium that results from the ban policy reduces total consumer surplus by $12M, equivalent to 6.4% of total sample spending on inside good restaurants, with nearly 90% of the welfare losses accruing to consumers in the two lowest income groups.

5.1 Chain customization

I first quantify the costs and benefits of operating a standardized chain from the perspective of the firm. I consider the choice of quality and cuisine type of large chains with more than 1000 locations nationwide. For each chain, I compute variable profits under three scenarios. I take full standardization to be the baseline case—the chain chooses one cuisine type and quality level that is fixed across all of its restaurants and gets the chain demand premium β^{1001+}. In the first counterfactual scenario—"flexible chain"—the chain chooses its cuisine type and quality flexibly in each market while keeping its demand premium β^{1001+}. In the second counterfactual scenario—"chain becomes independent"—the chain chooses its cuisine type and quality flexibly in each restaurant but loses its demand premium β^{1001+}. I conduct this exercise for each of the 23 large chains operating in at least four of seven cities and report averages across brands. Throughout the counterfactuals, I assume that firms play a two stage game in which they choose cuisine type and quality in the first stage and then play a Nash Bertrand pricing game in the second stage.

The flexible chain scenario—in which a chain like McDonalds keeps its brand advantage but can sell Mexican food in one market and Chinese in another—is not intended as an evaluation of a feasible strategy for restaurant chains. Changing a key attribute of a chain's product would conflict with its existing brand image, hurting its value in attracting consumers. McDonalds' brand advantage is likely tied to awareness among consumers that it sells a specific set of menu items at a given price point. Rather, evaluating the performance of a chain if it could be flexible is a way to quantify the impact of preference heterogeneity on the returns to forming a chain. The comparison between the baseline scenario, where the chain standardizes across markets, and the flexible scenario, in which it can fully customize, holds fixed the chain's demand advantage and isolates the effect of taste dispersion on firm variable profits.

The demand estimates imply that consumers receive higher utility from visiting a chain than an independent with the same product characteristics. However, the interpretation of the counterfactuals depends on why consumers like chains. One set of reasons is related to some investment made by the chain that is justified by its size (for example, in advertising or product development). If these are the primary channels, then a chain that decides to operate its restaurants independently would likely forfeit these advantages. Another set of reasons is related to selection into which firms become chains. If the chain advantage is related to some unobservable quality or management skill, some of this advantage may be retained if the firm were independent. I interpret the demand advantage as primarily being the former, consistent with the analysis of movers and new entries in Section 4.5.

This set of counterfactuals considers unilateral changes by one chain at a time, holding the characteristics of all other firms fixed.[39] For every chain restaurant, I find the quality level and cuisine type (and resulting price) that maximizes variable profits at the firm level when chains are standardized, and at the restaurant level when they can choose flexibly. In computation of the counterfactuals, I am calculating firm profits for a set of product characteristics and prices that were not chosen by the chain in the data, and thus I set the unobservable characteristics ξ_m and $\Delta\xi_{ijl}$ to zero.[40]

Table 3a shows the quality, cuisine type, and price that the chains in the data would choose under the three scenarios described above if they were optimizing based only on consumers in the estimation sample. When chains standardize across markets, they choose a quality level of about \$12 and sell American food (columns 2-4). The resulting average per-entree price is \$14.61.[41] However, these choices trade off demand in cities like Cleveland,

[39] A previous version of the paper computed an equilibrium that allowed a subset of nearby firms to respond to the changes by the chain with minimal effects on the main results.

[40] I assume that the unobservable traits of the firm are unlikely to be preserved if the restaurant were to make major changes to its characteristics and prices. The main results do not meaningfully change if I instead assume that each restaurant keeps its ξ_m and $\Delta\xi_{ijl}$ in every counterfactual.

[41] The optimal prices implied by my demand model are somewhat higher than the observed values in the data for major fast food chains (though similar to the average prices for many national casual dining brands with between 100 and 1000 locations). There are several possible reasons for this divergence: for example, chains that also operate during non-dinner hours (where average prices are lower) may want to limit intertemporal differences in prices. Chains also consider the tastes of consumers in places not in my sample, including rural markets in which consumers may be more price sensitive. Finally, the optimal set of product attributes for the chain reflects both demand and supply conditions; a chain might increase its profits by choosing a higher quality level either because most consumers prefer higher quality meals, or because few restaurants cater to that segment of consumers, relative to the level of demand for that restaurant type.

where consumers prefer a higher level of quality, and Madison and Las Vegas, where consumers tend to be more price sensitive. If the chain could be flexible in optimizing for the tastes of consumers in Cleveland, it would choose a quality level of \$13.49, with an average price of \$16.56 (columns 5-7). In Madison or Las Vegas, the chain would choose a level of price and quality that is slightly lower than the standardized choice. There is similar heterogeneity in tastes for cuisine types—European cuisine was most popular in four of the seven cities, but Cleveland and Phoenix consumers tend to prefer American cuisine, while Burgers were most popular in Pittsburgh. These differences in tastes across cities result in frictions for standardized firms.

In the last three columns of Table 3a, I show the product characteristics that chains would choose if they were to allow each outlet to become a flexible independent. In this scenario, firms would choose uniformly higher levels of quality and price. While independents face lower demand than chain firms on average, they tend to be relatively more attractive to high income consumers, who are less price sensitive. Within an income group, the negative estimate of ρ implies that consumers who have the lowest taste for chains also tend to be least price sensitive. These two factors result in restaurants setting prices between \$2 and \$8 higher when they operate as independents compared to their choices as standardized chains.[42]

In Table 3b, I show the change in variable profits for a chain that could choose its product characteristics flexibly across markets. I compute profits under the three scenarios described above for every large chain that operates in at least four of the seven sample cities. I take full standardization as the baseline scenario and show the change in profits for each chain if it could become flexible with and without the chain demand β^{1001+}. Table 3b shows that the returns to customization are substantial. The average chain in my data could increase its variable profits by about 20% if it could be fully flexible while keeping β^{1001+}. For about half of the chains, this flexibility is worth between 10% and 20% of its profits, but one chain would increase its variable profits by 37% if it could sell different products in different markets.

I then consider the "chain becomes independent" counterfactual—each chain can choose

[42] The optimal cuisine type does not change when the firm becomes independent. While I allow preferences for cuisine type to vary across income groups within a market, in practice the variation across markets tends to be larger than the variation within a market across income groups.

its product characteristics in every market in which it operates without constraint, but faces the demand of an independent firm (column 3). On average, large chains would lose 30% of their variable profits without the brand premium β^{1001+}, with some chains losing more than 50%. This highlights the important demand advantages that large chains enjoy relative to independent firms.

Together, the results of the counterfactuals highlight the importance of heterogeneity across markets, consistent with prior work (Quan and Williams, 2018; DellaVigna and Gentzkow, 2019). Although the average chain in the data enjoys a substantial demand premium over an independent firm with similar characteristics, consumers in different cities prefer different types of restaurants. This heterogeneity limits the degree to which chain firms can dominate the market, creating space for independent firms to survive in equilibrium alongside large, popular chains. Standardized chains will tend to choose product characteristics that are popular in many markets, but will be relatively poorly matched for consumers in the tails of the preference distribution. Further, because chain preference and price sensitivity are correlated, large chains will tend to focus on consumers who have low willingness to pay for quality. This helps to rationalize the observation that there are many national chains that serve low-priced hamburgers, but few or no chains that focus on Ethiopian cuisine or high-end French food. Nevertheless, the firms I observe earn much higher profits as standardized chains than they would as independents.

5.2 Chain bans

Finally, I consider the welfare effects of a ban on chain restaurants with more than 1000 locations. Chain bans or entry restrictions have been implemented in a number of small and large cities across the US and the majority of these policies apply explicitly to restaurants.[43] I will consider a "hard ban" most similar to the policy enacted in a number of neighborhoods in San Francisco.[44] I assume that the outcome of the policy is to close existing chain restaurants

[43]See, for example, the policies in Jersey City (City of Jersey City, 2020); San Francisco (San Francisco Planning, 2020); Fredericksburg, TX; Coronado, CA; Port Townsend, WA; Arcata, CA; and McCall, ID (Institute for Local Self Reliance, 2020).

[44]The San Francisco policy is more restrictive than my simulated policy, as it bans chains that have more than 11 locations from operating. I focus on the effects of closing very large chains, as my demand estimates imply these are the firms that are most valuable for consumers. The hard ban operates only in certain

and prevent new chain entries, leading to long run entry of independents.

Ex ante, the welfare impacts of chain bans are unclear. A chain ban has two opposing effects: it removes a set of restaurants that consumers value, but may result in a better match between restaurants and local tastes. Because preferences over both chains and restaurant characteristics are heterogeneous, chain bans may also have important distributional consequences.

The welfare impacts of chain bans depend on whether chains are replaced, and the mix of product characteristics chosen by the restaurants that enter in the aftermath of the policy. In my primary counterfactual scenario, I assume that chains are replaced one-for-one by independents.[45] Each independent chooses its quality and cuisine type to maximize its profits, given the estimated parameters of demand. I compute the new market structure using an iterated best response algorithm, as in Fan and Yang (2016)—I allow each firm to sequentially choose its quality and cuisine type in the first stage, and its resulting price in the second stage, given what other firms have already chosen. I loop through each (former) chain restaurant and iterate until no firm would like to change its product characteristics. I hold fixed the total number of restaurants in the market, the geographic locations of all restaurants, and the characteristics of all non-chain restaurants. As in the firm-level counterfactuals described above, I set the unobserved product characteristics ξ_m and $\Delta\xi_{ijl}$ to zero.

As I discuss above, a chain may consider a number of factors outside of my demand estimation exercise when choosing its product characteristics, including consumer preferences at breakfast and lunch times and in markets outside of my sample. To isolate the effect of the demand heterogeneity that I can measure, I show the welfare effects of chain bans relative to a baseline scenario in which a chain chooses a standardized quality and cuisine type to maximize profits based only on the preferences of consumers in my sample.

A well known challenge in the literature on firm entry and endogenous product characteristic choice is the potential for multiple equilibria (Seim, 2006; Fan, 2013; Wollmann,

areas of the city but other restrictions on chain retail are in place in most San Francisco neighborhoods (San Francisco Planning, 2020).

[45] If independent restaurateurs are more credit constrained than chain-affiliated outlets, or if fewer independents enter because they face lower demand, chains may not be fully replaced, in which case my estimates are a lower bound for the welfare impacts of the policy.

2018). I test for this by rerunning the chain ban counterfactuals with firms choosing in the reverse order and find very small differences in equilibrium price and quality levels. The welfare estimates of the impact of chain bans are essentially unchanged.

A chain ban results in important differences in the equilibrium distribution of quality levels and cuisine types. In the counterfactual, chains are replaced by independents who choose higher quality levels, and subsequently set higher prices. This is driven by two parameters in the demand model: first, β^{1001+} is systematically larger for low income groups, and thus independent firms are relatively more appealing to high income consumers, who are less price sensitive. Second, I estimate ρ to be negative in every market, implying that consumers who prefer independents to chains also have a higher willingness to pay for quality. This shift to higher quality and price levels in the counterfactual tends to benefit consumers who prefer high quality restaurants, but harms those who prefer lower prices. Independent firms also choose different cuisine types than standardized firms. This adjustment in cuisine type is welfare increasing for nearly all consumers.[46]

The second way in which a chain ban impacts welfare is through loss of the utility that comes from visiting a large chain firm, β^{1001+}. This parameter is large and positive on average for low income groups, but smaller or negative for consumers in the highest income groups. Preferences for chains are also heterogeneous within a city-income group. This heterogeneity in taste for both chain status and restaurant price and quality implies that a chain ban can have differential impacts across consumers along both observable and unobservable dimensions.

Using the new market structure, I compute expected consumer utility for every consumer in the data. From equation 1, consumer surplus in mile equivalents for each consumer i in trip t is given by:

$$CS_{it} = \frac{1}{\gamma} \int \left(\ln \sum_j \exp(V_j(\theta_i)) \right) f(\theta_i) d\theta_i \tag{7}$$

I approximate the integral in equation (7) using 50 Halton draws per consumer, as in the

[46] The average product characteristics chosen in the chain ban equilibrium in each city are similar to those shown in Table 3a.

main estimation. I then monetize this measure by multiplying by \$3.44 per one-way mile.[47] An alternative way of monetizing utility is to divide by the price coefficient, which gives the disutility from a \$1 increase in entree price, scaled up by the average number of entrees in a transaction. This approach would deliver a similar conversion rate.[48]

The counterfactual results imply that a chain ban would have stark distributional effects. In Table 4, I show the impact of a chain ban on consumer welfare by income (column (4)). Consumers in the lowest income groups suffer a large negative welfare impact as a result of the ban. Cardholders with income less than \$50,000 and between \$50,000 and \$100,000 per year are worse off by an amount equivalent to 10.4% and 7.6% of their restaurant spending, respectively. Consumers with income above \$200,000, on the other hand, are essentially unaffected. Within an income group, the welfare effects are even larger for the most price sensitive and chain-loving consumers. Summed across all income groups, a ban would decrease aggregate consumer surplus by \$12M, an amount equivalent to 6.4% of sample spending.

The welfare impacts of the ban come from two sources: a change in equilibrium product characteristics and the loss of the chain utility β^{1001+}. To decompose the contribution of these two factors, I compute the change in consumer welfare under two additional scenarios. First, I hold fixed restaurant characteristics at the values chosen by standardized firms but set β^{1001+} to zero (column (5) in Table 4). All groups are worse off without the change in product characteristics that would result from the ban, which partially reflects the fact that a chain ban results in a better match between local tastes and equilibrium product

[47] I convert mile equivalent utility into dollar terms by multiplying the number of (one-way) miles between a consumer and a restaurant by two to get the roundtrip distance and assuming that each mile costs \$0.91 in time costs and \$0.81 in direct costs, for a total of \$3.44 for each one-way mile between the consumer and the store. To obtain the monetary cost of a mile, I follow Dolfen et al. (2019) in using estimates from Einav et al. (2016), who report summary statistics for a large number of short-distance trips of breast cancer patients. They report that an average trip takes 10.9 minutes to travel 5.3 straight-line miles, with an actual driving distance of 7.9 miles. The BLS reports that the average after-tax hourly wage in 2016 was \$26 per hour. As an estimate for the driving cost, I use the IRS 2016 reimbursement rate of \$0.54 per mile, which considers the cost of fuel and depreciation of the car. Thus, the time cost of driving one mile is given by $\$26/60 \cdot 10.9/5.3 = \$0.91$ and the driving cost of of one mile is $\$0.54 \cdot 7.9/5.3 = \$0.81$.

[48] The ratio of average transaction size to entree price ranges from roughly two to five (see Appendix A.4 for additional details). The average price coefficient across all income groups is 0.330 and the average distance coefficient is 0.302. To transform utility to dollars using miles, I divide utility by the distance coefficient and multiply by \$3.44 per one way mile, which gives a conversion ratio of 11.4. To transform utility into dollars using the price coefficient, I divide by 0.302 and multiply by about 3 (approximately the ratio of transaction size to entree price for the mean restaurant), which gives a conversion ratio of 9.93.

characteristics. However, the difference between columns (4) and (5) are largest for higher income groups, as the mix of product characteristics (in particular, the higher quality levels and prices) chosen by independents are especially attractive to richer consumers.

To isolate the effect of the loss of the chain utility, I set product characteristics to the values chosen by firms in the chain ban equilibrium but allow the new independent entrants to keep β^{1001+} (column (6)). Lower income groups recover much of the welfare losses that they would suffer under the full ban; the surplus of the lowest group drops by only 2.6% of spending compared to 10.5%. Higher income groups are also better off than under the full ban, but the differences are significantly smaller than for the lowest income consumers. Consumers with income above $200,000 are better off by 2.6% of spending, while their welfare is largely unchanged in the full ban.[49]

The chain ban results rely on the estimates of consumer preference heterogeneity across cities to predict the characteristics that would be chosen in the counterfactual equilibrium. To assess the robustness of these findings, I consider several alternative assumptions about the types of restaurants that would replace chains. First, I assume that large chains are replaced by independents with the same quality and cuisine type. Second, I replace each large chain with a randomly drawn independent (with replacement) from the set of independents in the city. I compare both scenarios to the set of characteristics chosen in the data. I report welfare effects under these alternative assumptions in Table C.14. While the welfare effects are smaller in magnitude than the estimates reported in Table 4, the gradient with respect to income is robust to these alternative scenarios.[50]

Arguments for chain bans tend to focus on three themes: aesthetic concerns, positive local externalities from independent businesses, and a desire to help independent business owners. The first two justifications—preserving the character of downtown areas and positive labor market and other externalities from independents—are difficult to quantify and outside the scope of my model. However, my estimates suggest that chain bans are a costly tool for

[49]While high income consumers have a negative mean preference for chains, there is significant within-group taste dispersion, which results in a positive overall welfare effect for this top income group in column (6).

[50]When chains are replaced by identical independents or a randomly selected independent, aggregate consumer surplus falls by 4% and 1.5%, respectively, with the welfare losses decreasing monotonically in income. When discussing the net welfare impact of a chain ban, I consider a range of welfare impacts between 1.5% and 6.4%.

redistributing surplus to local business owners. Assuming that the overall level of restaurant spending remains fixed under the policy, a chain ban would redirect the 7% of sample spending going to large chains to independents and small chains. Industry sources estimate average accounting profit margins at about 6% of revenues, implying that these policies would increase independent profits by roughly 0.4% of total spending. I estimate that the welfare loss to consumers is between 1.5% and 6.4% of restaurant spending, or between 3 and 16 times larger than the transfer, implying that these policies are a costly way to help independent owners.

The results of the counterfactuals also highlight important distributional consequences of chain ban policies. Chain bans tend to hurt consumers with lower incomes, and within an income group, those who are most price sensitive. This suggests a political economy explanation for chain bans. The average consumer in a high income city, or in a market with idiosyncratic tastes that are not well served by standardized chains, may be better off with the ban. Indeed, about three quarters of the 30 cities that have enacted these policies have median household income above the national median. Nevertheless, my results highlight the outsize impact that these policies can have on low income residents within these places.

6 Conclusion

In this paper, I study the role of large chains in the restaurant market. As a complement to existing work on the supply-side advantages of large firms, I show that chains also have important demand side advantages. However, these demand advantages are counterbalanced by an important cost for standardized firms. Because consumers are heterogeneous, a chain that standardizes sacrifices a significant amount of potential profits in some markets. I quantify both sides of this demand side trade-off using a rich, transaction-level dataset that includes about half of credit and debit transactions in the United States. I show that if chains could choose their product characteristics flexibly in each market but keep their current demand, they could increase their variable profits by 20% on average. However, if they were to give up that demand advantage in exchange for flexibility, they would lose 30% of their profits.

Despite the advantages of large national chains, small chains and independent firms still account for a significant portion of restaurant sales. My work suggests that differences in consumer tastes across markets are an important reason for this. Absent heterogeneity across markets, we might expect to see differentiated large chains capturing all transactions, even when consumer tastes are dispersed within a market. Restaurants are naturally limited in their ability to cater to different tastes within the same outlet; restaurants that offer food from many cuisine types at many price levels are uncommon. Retail categories in which chains account for nearly all of sales, like in the general merchandise category, tend to offer a wide array of different products, and thus the natural constraint that standardization imposes may not be as important.

This work also quantifies the effect of a ban on chain restaurants on consumer welfare. Chain bans remove a set of high quality firms that many consumers prefer, but also result in better matches between firms and consumer preferences in local markets. Overall, I find that bans would decrease consumer welfare by between 1.5% and 6.4% of restaurant spending, with nearly 90% of the welfare losses falling on consumers with incomes below $100,000 per year. The magnitudes of my estimates suggest that the loss in consumer surplus that would result from a chain ban is between 3 and 16 times as large as the additional profits that would flow to independent businesses, and thus these policies are justified only if independent firms bring large positive externalities to local downtown areas.

Figure 1: Share of chains

(a) Aggregate share of spending by number of firm locations in augmented retail categories

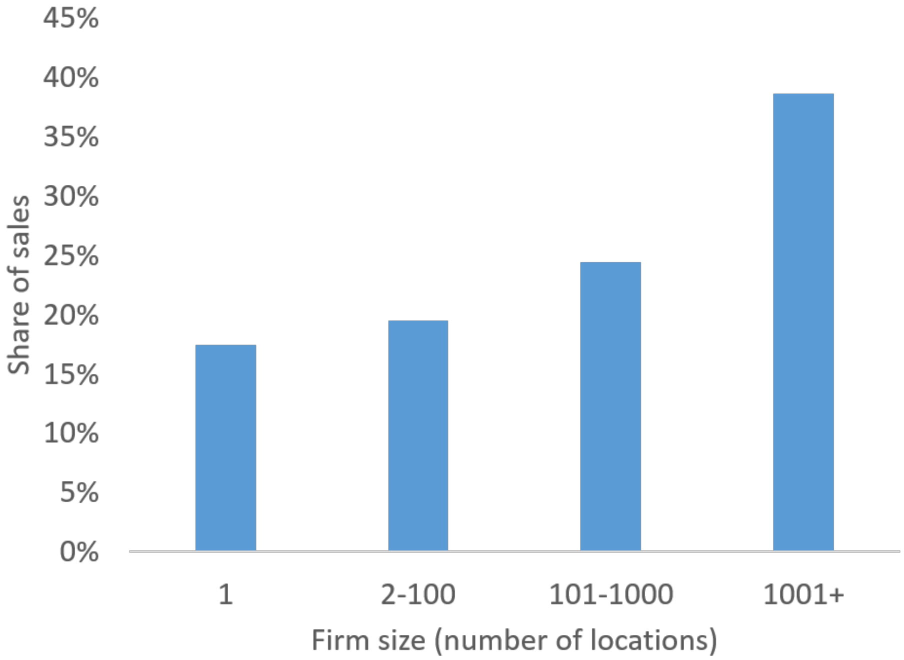

(b) Share of spending by number of firm locations by merchant category

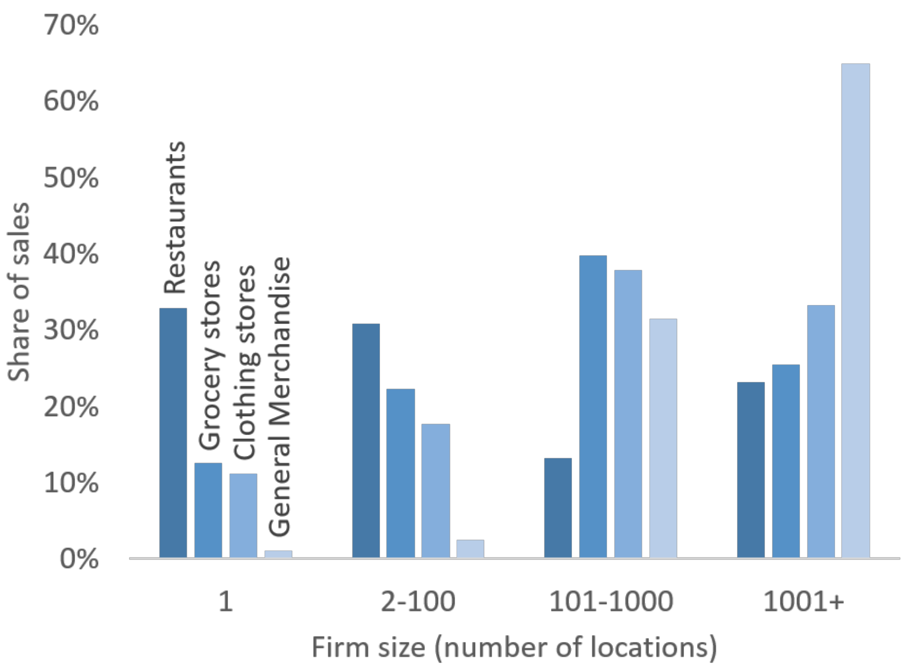

The figure shows the share of US payment card spending in 2016 that went to merchants with the number of locations given on the x-axis. Panel (a) shows the aggregate shares by firm size in the retail, restaurants, and hotel categories (corresponding to NAICS codes beginning with 44, 45, and 72) . The largest categories wihin this set by share of spending were restaurants (25%), grocery stores (16%), and general merchandise stores (13%). Panel (b) shows the share of spending by firm size for four large retail categories. Each group of bars corresponds to a firm size bin, while each bar within the group gives the share of spending within a firm category that went to firms with that number of locations.

Figure 2: Characteristics of chain restaurants

(a) Quantiles of average restaurant price by chain size

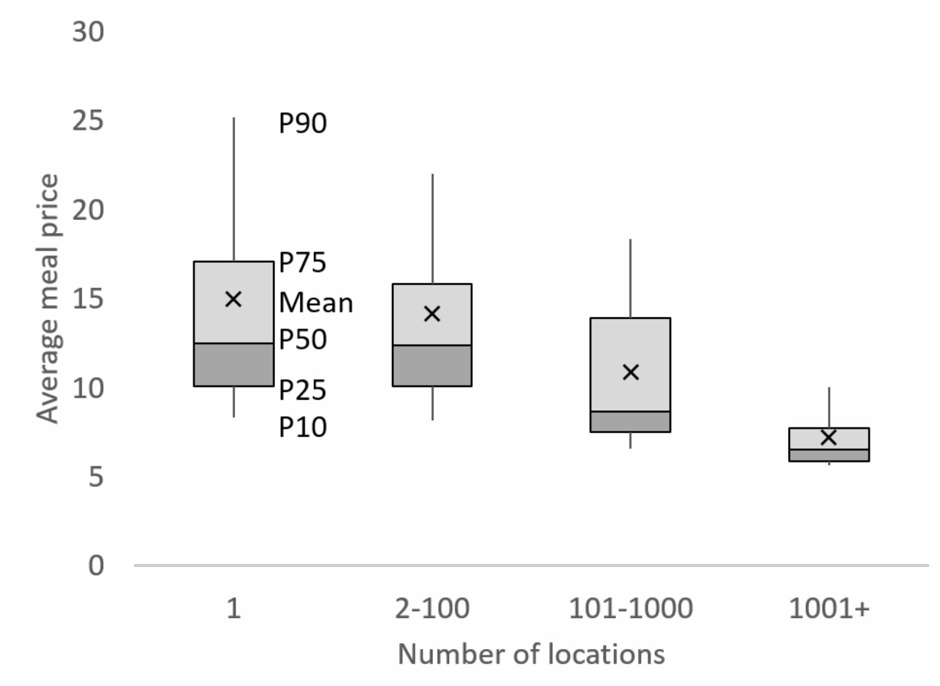

(b) Share of chain restaurants by annual cardholder income

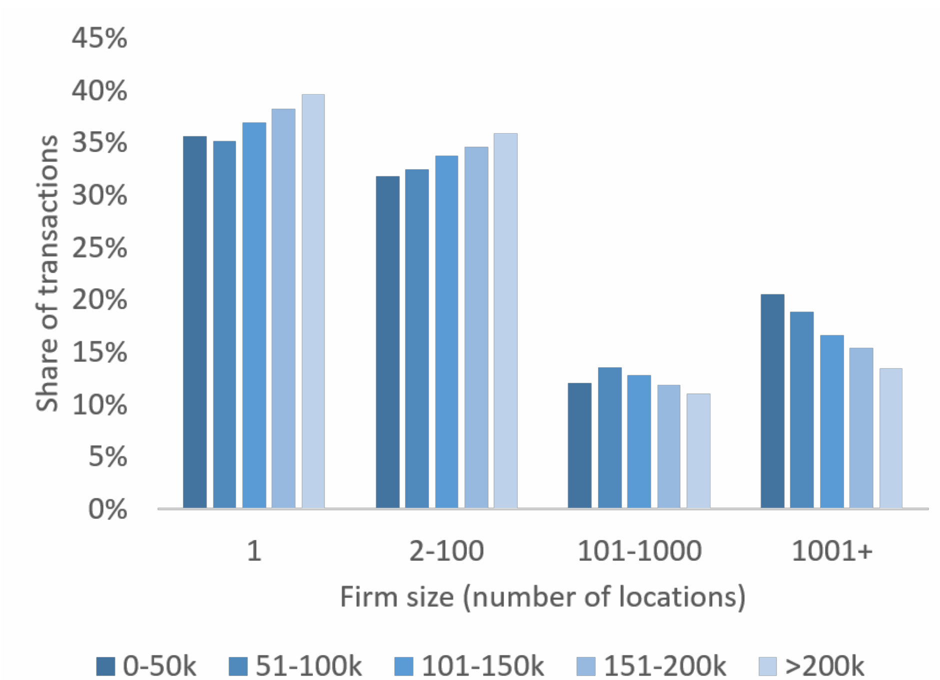

The figure shows characteristics of chain restaurants and their customers by chain size. Panel (a) shows the mean and quantiles of the distribution of average restaurant entree price for restaurants in the seven sample cities used in estimation. An observation in the underlying dataset is a restaurant in 2016. I impute average meal price by combining a sample of detailed price data for 20 large restaurant chains with the average transaction size at each sample restaurant. I describe this calculation in detail in Section 2 and Appendix A.4. Panel (b) shows the share of sample spending in 2016 that went to merchants with the number of US locations given on the x-axis by estimated cardholder income.

Figure 3: Popularity of restaurants by average price and cuisine type across cities

(a) Per-restaurant transactions by average transaction size and city

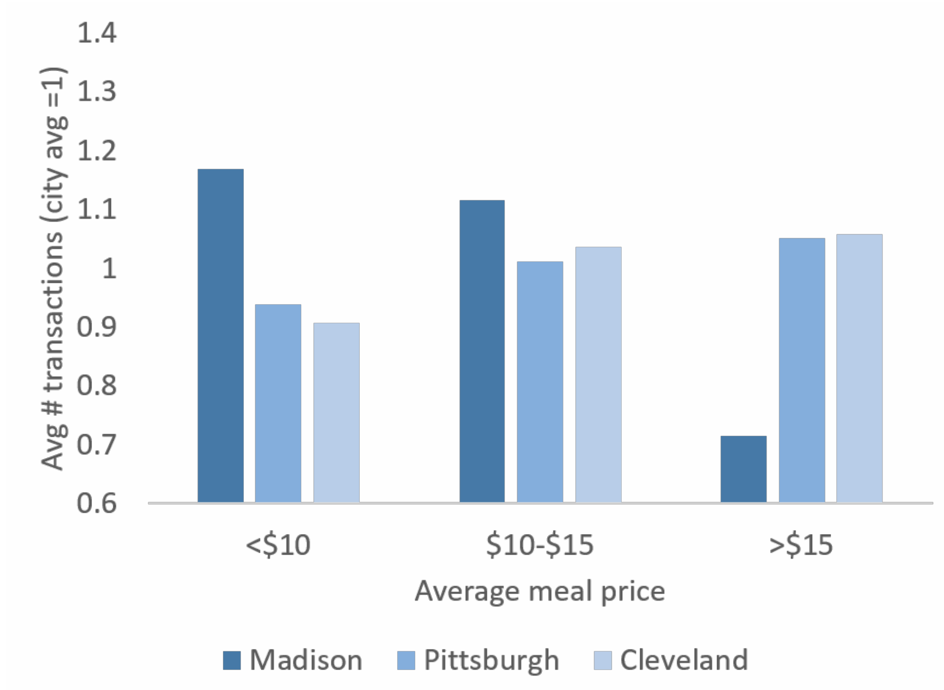

(b) Per-restaurant transactions by cuisine type and city

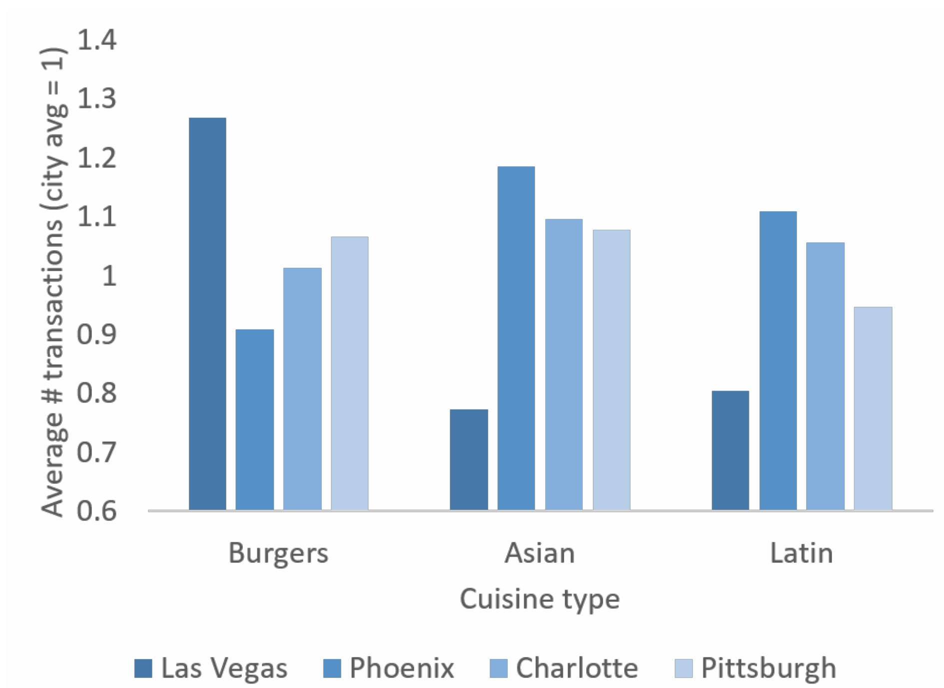

The figure shows the relative success of chain restaurants of different price levels and cuisine types in four cities included in the main estimation sample described in Section 2. For each sample restaurant, I divide its total number of transactions by the average per-restaurant transactions in its city and regress this variable on a set of merchant fixed effects. Using the set of chains that have multiple sample locations (1,420 out of 3,944 total restaurants in the sample), I average the residuals in each city for restaurants of different price levels (panel (a)) and cuisine types (panel (b)). In each plot, I normalize the residuals to average to 1 for each city.

Figure 4: Demand estimates

(a) Utility for average consumer by restaurant price

(b) Chain premium in mile equivalents

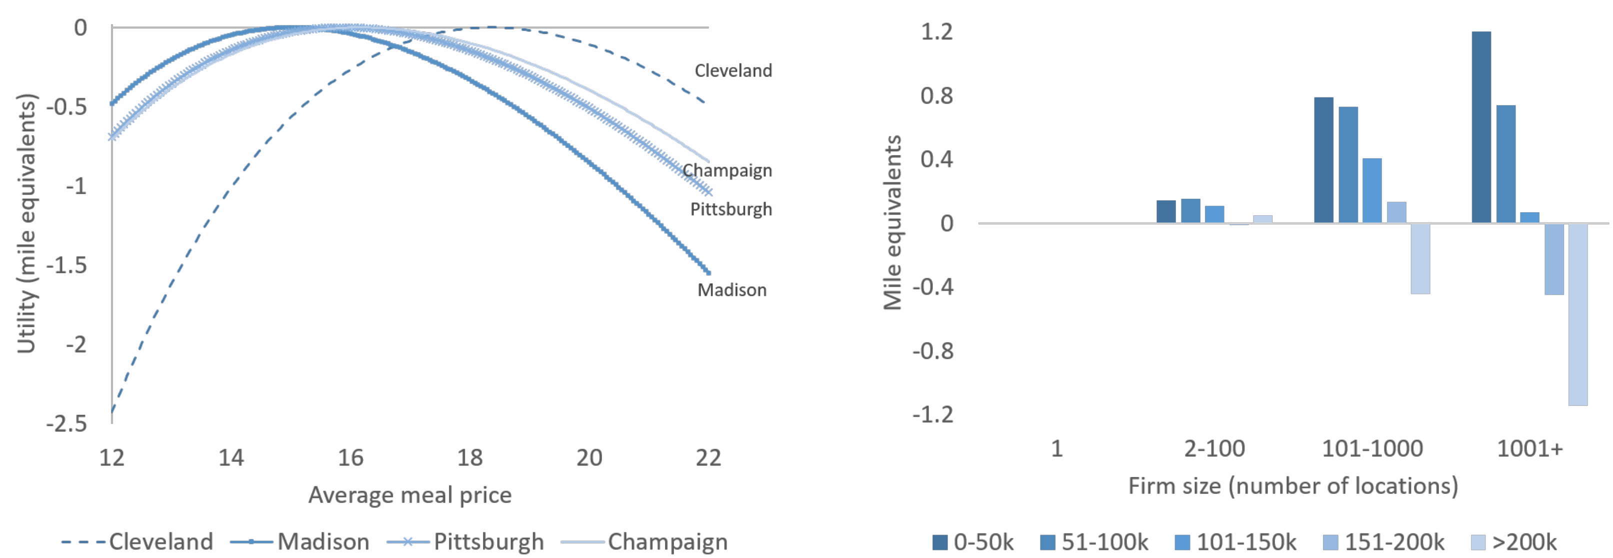

Panel (a) shows the utility that the average consumer receives from visiting restaurants of different price levels (with their associated level of quality) in four sample cities. Utility is measured in mile equivalents, which is defined as the change in utility when a restaurant is moved one mile closer to a consumer's home. To produce the figure, I back out restaurant quality for each price level using the average markup (I assume quality is linear in a restaurant's marginal cost). I then calculate the utility for the average city c consumer from visiting restaurant j with price p_j and quality $q_j(p_j)$ as $u_{ij} = \alpha^q \log q_j - \alpha^p_c p_j$, where α^p_c is the average price sensitivity in city c. I transform this utility to mile equivalents by dividing it by the utility cost of traveling one mile γ_c and normalize it to zero at the maximum for each city. Panel (b) shows the additional utility in mile-equivalent units that a consumer receives from visiting a chain restaurant relative to an independent restaurant as a function of her income. To calculate this, I compute the chain premium in each city-income group as the sum of the chain size fixed effects in Tables B.11a, B.11b, and B.12 and divide by the cardholder cost of traveling one mile γ_{cy}. I average this preference across cities by income group, weighting by the number of transactions. The fixed effect for firms with one location is normalized to zero. Each group of bars corresponds to a firm size bin, while each bar within the group gives the relative chain premium corresponding to cardholders with a given estimated household income.

Figure 5: Cuisine type heatmap with ranks

	American	Asian	Burgers	European	Latin	Other	Pizza	Sandwiches
Champaign	3	7	6	1	4	5	8	2
Charlotte	2	4	3	1	6	8	5	7
Cleveland	1	7	4	8	3	5	6	2
Las Vegas	2	7	4	1	8	3	6	5
Madison	2	5	3	1	6	4	8	7
Phoenix	1	5	2	7	6	3	8	4
Pittsburgh	5	7	1	2	6	3	8	4

The figure shows estimates of consumer tastes for different cuisine types by city. Each cell in the matrix gives the ranking of a given cuisine type in that city, averaged across income groups. A rank of 1 indicates that the category was the top-ranked category in the city and a rank of 8 indicates that it was the lowest ranked category in the city. The heat map is color coded with dark blue boxes corresponding to higher rankings (higher demand) while white boxes correspond to lower rankings (lower demand).

Figure 6: Channels of chain demand advantage

(a) Restaurant transactions of movers by months since move

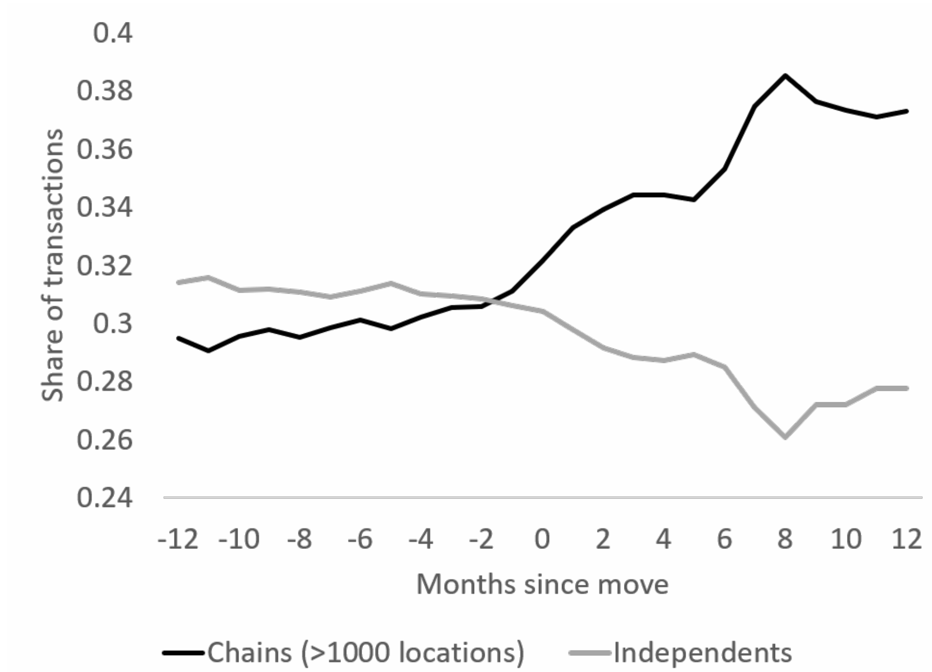

The figure shows the share of restaurant transactions for consumers that move across state lines at large chains (>1000 locations) and independents (1 location) in the 12 months before and after their date of move. Each consumer's move date is defined as the first month that they are reported with a billing zipcode in a new state. See Appendix A for additional details on sample construction.

(b) Monthly sales for new restaurant entries by months since open

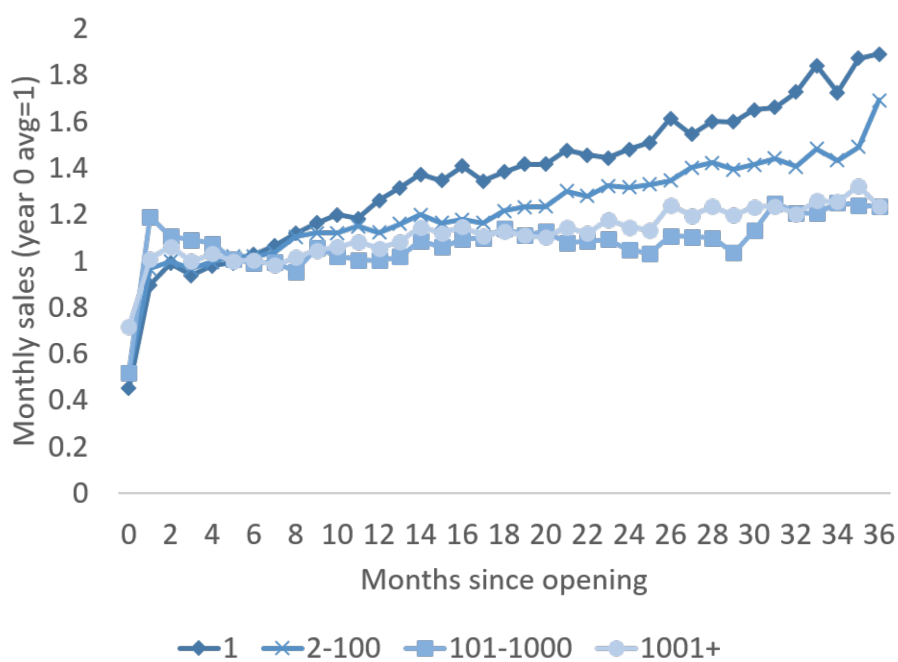

The figure shows monthly sales for new restaurant entries in the seven sample cities broken out by chain size. The y-axis gives the share of sales for restaurant j in a given month divided by the average monthly sales for j over the first 12 months, averaged across restaurants of a given chain size.

Table 1: Summary statistics

(a) By city

City	Accounts (K)	Transactions (M)	Dollars (M)	Restaurants
Champaign	26.8	0.5	15.4	147
Charlotte	186.2	3.3	115.7	817
Cleveland	114.3	3.1	125.3	378
Las Vegas	174.2	2.9	109.0	754
Madison	90.0	2.0	65.4	358
Phoenix	240.2	5.8	198.3	724
Pittsburgh	176.3	3.3	139.2	766
Total	1,008.0	21.0	768.4	3,944

(b) By household income

Household Income	Accounts (K)	Avg. Transactions	Avg. Dollars	Avg. Price
<$50k	349.7	19.1	601.5	13.42
$51-100k	356.6	19.8	704.4	14.29
$101-150k	168.8	23.2	916.1	14.92
$151-200k	63.7	24.3	986.5	15.28
>$200k	69.2	26.5	1,292.7	16.14

Panel (a) shows the number of accounts, transactions, dollars and restaurants included in the main analysis sample used in Section 4 by city. Panel (b) shows summary statistics on the activity of cardholders included in the main analysis sample used in Section 4 by bin of household income. Both panels include purchases at outside good restaurants.

Table 2: Parameter estimates from main estimation by city and income

(a) By city

Parameter	Description	Champaign	Charlotte	Cleveland	Las Vegas	Madison	Phoenix	Pittsburgh
β^1	Firm size - 1 loc.	0.000	0.000	0.000	0.000	0.000	0.000	0.000
β^{2-100}	Firm size - 2-100 loc.	0.038	0.035	0.035	0.038	0.036	0.036	0.035
$\beta^{101-1000}$	Firm size - 101-1000 loc.	0.322	0.400	-0.650	0.118	0.824	0.557	-0.015
β^{1001+}	Firm size - 1001+ loc.	0.685	0.473	-1.280	0.407	1.144	0.863	-0.270
γ	Physical distance (miles)	-0.385	-0.296	-0.211	-0.310	-0.338	-0.302	-0.354
$\bar{\alpha}^p$	Average price sensitivity ($)	-0.326	-0.354	-0.280	-0.347	-0.357	-0.326	-0.333
α^q	Log(quality) ($)	4.316	4.316	4.316	4.316	4.316	4.316	4.316
σ^{α^p}	Variance(price sensitivity)	0.006	0.003	0.002	0.002	0.003	0.002	0.002
σ^{α^l}	Variance(chain preference)	0.034	0.033	0.048	0.031	0.037	0.041	0.035
ρ	Cov(price sens.,chain pref.)	-0.008	-0.008	-0.007	-0.007	-0.008	-0.009	-0.007

(b) By income

Parameter	Description	<50k	50-100k	100-150k	150-200k	>200k
β^1	Firm size - 1 loc.	0.000	0.000	0.000	0.000	0.000
β^{2-100}	Firm size - 2-100 loc.	0.046	0.044	0.030	-0.003	0.014
$\beta^{101-1000}$	Firm size - 101-1000 loc.	0.287	0.256	0.184	0.123	-0.014
β^{1001+}	Firm size - 1001+ loc.	0.454	0.291	0.143	0.028	-0.128
γ	Physical distance (miles)	-0.325	-0.297	-0.285	-0.284	-0.293
$\bar{\alpha}^p$	Average price sensitivity ($)	-0.359	-0.335	-0.311	-0.299	-0.279
α^q	Log(quality) ($)	4.316	4.316	4.316	4.316	4.316
σ^{α^p}	Variance(price sensitivity)	0.003	0.003	0.002	0.002	0.002
σ^{α^l}	Variance(chain preference)	0.037	0.039	0.038	0.036	0.040
ρ	Cov(price sens.,chain pref.)	-0.008	-0.008	-0.007	-0.006	-0.006

The table shows the aggregated point estimates of a subset of the parameters in equation 1. Panel (a) shows the weighted average of each parameter by city, where each income group is weighted by its number of transactions. Panel (b) shows the average parameters for each income group, where each city is weighted by its number of transactions. β^1 is normalized to zero for each group. See Tables B.10-B.12 for the full set of parameter estimates with bootstrapped standard errors.

Table 3: Chain reoptimization counterfactuals

(a) Optimizing ticket size and cuisine type choices by city

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		Std. chain			Flex. chain			Ind.	
City	Quality	Cuisine	Price	Quality	Cuisine	Price	Quality	Cuisine	Price
Champaign	11.85	American	14.61	13.75	European	16.96	19.10	European	22.60
Charlotte	11.85	American	14.61	11.40	European	14.06	13.66	European	16.34
Cleveland	11.85	American	14.61	13.49	American	16.56	16.89	American	19.98
Las Vegas	11.85	American	14.61	11.34	European	13.98	13.50	European	16.14
Madison	11.85	American	14.61	11.38	European	14.02	13.66	European	16.33
Phoenix	11.85	American	14.61	11.80	American	14.53	14.70	American	17.45
Pittsburgh	11.85	American	14.61	11.96	Burgers	14.76	14.24	Burgers	17.05

(b) Counterfactual change in variable profits

	Flex. chain	Chain becomes ind.
Avg.	20%	-30%
Std.	8%	14%
Min.	7%	-58%
P25	15%	-41%
P50	19%	-27%
P75	24%	-21%
Max.	37%	-9%
N	23	23

The table shows results of the chain reoptimization counterfactuals described in Section 4. Panel (a) shows the average quality, cuisine type, and price chosen by large chains in each market under three different scenarios: standardization, "flexible chain", and "chain becomes independent". In each scenario, restaurants play a two stage game in which they first choose quality and cuisine type, and then set prices a la Nash Bertrand. Under standardization, each chain chooses a quality level and cuisine type that is fixed across markets to maximize its variable profits (columns 2-4). In "flex chain", each individual restaurant chooses its cuisine type and quality separately and keeps chain demand β^{1001+} (columns 5-7). In "chain becomes independent", every restaurant can choose flexibly, but loses chain demand β^{1001+} (columns 8-10). In panel (b), I show the change in variable profits when the chain moves from standardization to flexible chain and full independent, respectively. I compute the change in profits for each of the 23 large chain firms that operate in at least four out of seven sample cities. Panel (b) presents statistics for the change in profits at the firm level.

Table 4: Welfare effects of a a chain ban by consumer income

	(1)	(2)	(3)	(4)	(5)	(6)
				% of spending		
Income group	Accounts (K)	Spending ($M)	Baseline CS ($M)	Chain ban	Stand. chain without β^{1001}	Chain ban with β^{1001}
<50k	349.7	55.5	31.8	-10.4%	-11.3%	-2.4%
50-100k	356.6	61.4	32.1	-7.6%	-9.6%	-0.7%
100-150k	168.8	35.1	16.0	-3.7%	-6.9%	1.2%
150-200k	63.7	14.2	6.1	-1.4%	-5.3%	2.2%
>200k	69.2	20.1	7.2	-0.0%	-4.0%	2.7%
Total	1,008.0	186.3	93.1	-6.4%	-8.7%	-0.3%

The table shows the effect of a ban of large chains with more than 1000 locations on consumer welfare. Column (2) gives the total amount of spending on inside good restaurants. Column (3) shows the level of consumer surplus in millions of dollars when large chains choose a standardized level of quality and cuisine type that is fixed across markets (the baseline scenario). Columns (4)-(6) show the change in consumer surplus relative to the baseline that would result from each counterfactual scenario, as a share of total spending. In column (4), I assume that each large chain is replaced by an independent restaurant (without the chain demand advantage β^{1001+} that chooses its cuisine type, quality, and price to maximize its profits. In column (5), I assume that the replacement keep the quality and cuisine type would have chosen as standardized firms, but lose β^{1001+}. In column (6), I assume that the replacement restaurants choose the same product attributes as in column (4), but keep β^{1001+}. Consumer surplus is computed by converting utility to mile equivalent units by dividing by γ and then monetized by multiplying by $3.44 per one-way mile (see Section 5 for more details).

References

Adams, B. and K. R. Williams (2019). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics 11*(1), 124–56.

Allcott, H., R. Diamond, J.-P. Dubé, J. Handbury, I. Rahkovsky, and M. Schnell (2019). Food deserts and the causes of nutritional inequality. *The Quarterly Journal of Economics 134*(4), 1793–1844.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.

Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy 112*(1), 68–105.

Bertrand, M. and F. Kramarz (2002). Does entry regulation hinder job creation? Evidence from the French retail industry. *The Quarterly Journal of Economics 117*(4), 1369–1413.

Bronnenberg, B. J., S. K. Dhar, and J.-P. H. Dubé (2009). Brand history, geography, and the persistence of brand shares. *Journal of Political Economy 117*(1), 87–115.

Bronnenberg, B. J., J.-P. H. Dubé, and M. Gentzkow (2012). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review 102*(6), 2472–2508.

Bureau of Labor Statistics ((accessed April 28, 2021)). *High-income households spent half of their food budget on food away from home in 2015.*

Cai, H. and I. Obara (2009). Firm reputation and horizontal integration. *The RAND Journal of Economics 40*(2), 340–363.

Chen, M. K. and D. G. Pope (2020). Geographic mobility in america: Evidence from cell phone data. Working paper, National Bureau of Economic Research.

Choi, J. and D. R. Bell (2011). Preference minorities and the internet. *Journal of Marketing Research 48*(4), 670–682.

City of Jersey City ((accessed July 8, 2020)). City ordinance 15.053. `https://library.municode.com/nj/jersey_city/ordinances/code_of_ordinances?nodeId=741092`.

City of Los Angeles (2019). Regulating Formula Retail in Los Angeles. `https://www.venicenc.org/productphotos/Formula%20Retail%20Presentation_8-26-19.pdf`.

Conlon, C. and J. Gortmaker (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics 51*(4), 1108–1161.

Davis, D. R., J. I. Dingel, J. Monras, and E. Morales (2019). How segregated is urban consumption? *Journal of Political Economy 127*(4), 1684–1738.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in US retail chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Dolfen, P., L. Einav, P. J. Klenow, B. Klopack, J. D. Levin, L. Levin, and W. Best (2019). Assessing the gains from e-commerce. Working paper, National Bureau of Economic Research.

Doms, M. E., R. S. Jarmin, and S. D. Klimek (2004). Information technology investment and firm performance in US retail trade. *Economics of Innovation and New Technology 13*(7), 595–613.

Dutta-Bergman, M. J. (2006). The demographic and psychographic antecedents of attitude toward advertising. *Journal of Advertising Research 46*(1), 102–112.

Einav, L., A. Finkelstein, and H. Williams (2016). Paying on the margin for medical care: Evidence from breast cancer treatments. *American Economic Journal: Economic Policy 8*(1), 52–79.

Eizenberg, A., S. Lach, and M. Oren-Yiftach (2021). Retail prices in a city. *American Economic Journal: Economic Policy 13*(2), 175–206.

Ellickson, P. B., S. Houghton, and C. Timmins (2013). Estimating network economies in retail chains: A revealed preference approach. *The RAND Journal of Economics 44*(2), 169–193.

Erdem, T. and M. P. Keane (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science 15*(1), 1–20.

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the US daily newspaper market. *American Economic Review 103*(5), 1598–1628.

Fan, Y. and C. Yang (2016). Competition, product proliferation and welfare: A study of the US smartphone market. *Available at SSRN 2506423*.

Federal Reserve Board of Governors (2019). Report on the Economic Well-Being of US Households in 2018, May 2019.

Foster, L., J. Haltiwanger, and C. J. Krizan (2006). Market selection, reallocation, and restructuring in the US retail trade sector in the 1990s. *The Review of Economics and Statistics 88*(4), 748–758.

Griffith, R. and H. Harmgart (2012). Supermarkets competition in England and planning regulation. *The International Review of Retail, Distribution and Consumer Research 22*(1), 1–25.

Hess, S. and K. E. Train (2011). Recovery of inter-and intra-personal heterogeneity using mixed logit models. *Transportation Research Part B: Methodological 45*(7), 973–990.

Hollenbeck, B. (2017). The economic advantages of chain organization. *The RAND Journal of Economics 48*(4), 1103–1135.

Hollenbeck, B. and R. Giroldo (2020). Winning big: Scale and success in retail. *Available at SSRN*.

Holmes, T. J. (2001). Bar codes lead to frequent deliveries and superstores. *RAND Journal of Economics*, 708–725.

Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica 79*(1), 253–302.

Horgan, C. (2017). Retail revolution: should cities ban chain stores? *The Guardian*.

Houde, J.-F., P. Newberry, and K. Seim (2017). Economies of density in e-commerce: A study of Amazon's fulfillment center network. Technical report, National Bureau of Economic Research.

Institute for Local Self Reliance ((accessed July 8, 2020)). Formula business restrictions. `https://ilsr.org/rule/formula-business-restrictions`.

Jarmin, R. S., S. D. Klimek, and J. Miranda (2009). The role of retail chains: National, regional and industry results. In *Producer dynamics: New evidence from micro data*, pp. 237–262. University of Chicago Press.

Jia, P. (2008). What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica 76*(6), 1263–1316.

Loertscher, S. and Y. Schneider (2011). Chain stores, consumer mobility, and market structure. *Journal of Institutional and Theoretical Economics JITE 167*(2), 236–246.

Luca, M. (2011). Reviews, reputation, and revenue: The case of yelp. com. Working paper, Harvard Business School.

Mast, E. (2020). Race to the bottom? local tax break competition and business location. *American Economic Journal: Applied Economics 12*(1), 288–317.

Mazzeo, M. J. (2004). Retail contracting and organizational form: Alternatives to chain affiliation in the motel industry. *Journal of Economics & Management Strategy 13*(4), 599–615.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Quan, T. W. and K. R. Williams (2018). Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics 49*(4), 877–913.

Revelt, D. and K. Train (1998). Mixed logit with repeated choices: households' choices of appliance efficiency level. *Review of Economics and Statistics 80*(4), 647–657.

Sadun, R. (2015). Does planning regulation protect independent retailers? *Review of Economics and Statistics 97*(5), 983–1001.

San Francisco Planning ((accessed July 8, 2020)). Chain stores/business. `https://sfplanning.org/permit/chain-stores-businesses`.

Sebastien, L. (2008). Variations in choice sets and empirical identification of mixed logit models: Monte carlo evidence. Working paper, ALISS-08, 1-23.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics 37*(3), 619–640.

Slattery, C. (2018). Bidding for firms: Subsidy competition in the US. *Available at SSRN 3250356*.

Tsai, Y.-L., C. S. Dev, and P. Chintagunta (2015). What's in a brand name? Assessing the impact of rebranding in the hospitality industry. *Journal of Marketing Research 52*(6), 865–878.

Varadhan, R. and C. Roland (2008). Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scandinavian Journal of Statistics 35*(2), 335–353.

Viviano, E. (2008). Entry regulations and labour market outcomes: Evidence from the Italian retail trade sector. *Labour Economics 15*(6), 1200–1222.

Waldfogel, J. (2003). Preference externalities: An empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics 34*(3), 557–568.

Waldfogel, J. (2010). Who benefits whom in the neighborhood? Demographics and retail product geography. In E. Glaeser (Ed.), *Agglomeration Economics*. Chicago: University of Chicago Press.

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review 108*(6), 1364–1406.

Online appendix: One size fits all? The value of standardized chains

Appendix A Data

A.1 Payment card data overview

My primary source of data is the universe of 2016 transactions on a major payments card network. The payments card provider is among the largest in the US. Total transaction volume on the network in 2016 was approximately 20% of all US consumption.

An observation in the underlying data is a transaction between a card and a merchant. The key variables used from this dataset are a unique card identifier, date and time of transaction, merchant identifier (defined at the brand level), store identifier (defined at the outlet level), latitude and longitude of the store location, and the dollar amount of the transaction. I exclude transactions that are not sales drafts or that occur at a non-US merchant. For 55% of active 2016 credit cards issued by the payment cards network, the company has access to a measure of estimated household income and the cardholder's billing zipcode. Household income is estimated by a third party from information available in a credit report. I use all credit card transactions for which the payment card company observes estimated household income and billing zip+4. This excludes all spending by prepaid cards and debit cards, as I do not observe cardholder income for these transactions. To describe the aggregate importance of chains and the variation across store category (Section 3), I use the sample described above. In the demand estimation and entry model analysis, I further restrict this sample to specific geographies, described below.

A.2 Representativeness of credit card data

As I report in the main text, credit and debit card spending make up between 70% and 85% of all restaurant spending.[51] However, the purchases of consumers with cards may differ systematically from those without. This can induce bias into my estimates of consumer

[51] See https://www.creditcards.com/credit-card-news/payment-method-statistics-1276.php

preferences across places if the propensity to hold or use credit cards differs across cities. There are two primary concerns. First, low-income consumers are significantly less likely to have credit cards than high-income consumers. In Table A.2, I show the share of consumers by income group that have credit cards, calculated from the Atlanta FRB's 2018 Survey of Consumer Payment Choice. Over 95% of consumers with annual household income above $100,000 per year had at least one credit card in 2018, compared to only 50% of consumers with income below $25,000. This may lead me to understate demand for the types of restaurants preferred by low income consumers (per my demand estimates, they tend to be cheaper restaurants and large national chains).[52]

Second, consumers may be more likely to use credit cards at some restaurants than others. Independent restaurants in particular may not accept credit cards or discourage the use of cards by imposing additional fees or minimum purchase amounts in order to avoid processing fees.[53] Consumers may also be less likely to use cards for small purchases, which could lead me to understate demand for low priced restaurants.

To assess the extent of these potential biases, I compare the patterns of consumer behavior observed in the credit card data to that recorded in a dataset that tracks users based on their smartphone locations from Safegraph. A consumer in the Safegraph data is recorded as making a "visit" to a business when that consumer's GPS location is recorded within the coordinates of the business for at least 15 minutes (Chen and Pope, 2020). While smartphones are still more prevalent among higher income consumers, they have a higher penetration among lower income groups than do credit cards. In addition, smartphone visits are unlikely to be biased towards venues that are more willing to accept credit cards. Of course, visits recorded from smartphones may be subject to their own set of biases and sources of mismeasurement. Businesses that are located in dense venues such as shopping

[52]Low income consumers also spend much less on restaurants than consumers with higher incomes. Data from the 2015 Consumer Expenditure Survey show that the bottom 40% of households by income accounted for less than 20% of all spending on food away from home, with the bottom 20% accounting for only 8% of spending (Bureau of Labor Statistics, 2021). This suggests that selection into which consumers use credit cards is likely to have a limited impact on the demand estimates.

[53]Using the Yelp data, I calculate that about 5% of restaurants in the seven US cities included in my sample do not accept credit cards. Restaurants with no credit card transactions will not appear in the choice set of the consumers used in estimation, and thus are less likely to bias results of estimation than are restaurants that merely discourage the use of cards.

malls, stadiums, theme parks, or hospitals may receive more cell phone visits than actual purchases if consumers spend time in proximity to those businesses without purchasing. Conversely, consumers that make purchases quickly, such as in a drive-through or for a to-go order, may be undercounted by cell phone data relative to actual purchases.

There are several important caveats to this exercise. First, the Safegraph sample does not match up exactly with the sample I use in the main estimation. It contains data only from the Houston market in 2017, rather than 2016 consumption in the seven US cities used in my estimation exercise. It also includes transactions across all meal times, rather than only transactions that occur during dinner hours. I create a corresponding sample for Houston using the credit card data for comparison purposes. Second, I am unable to link the credit card and Safegraph data at the merchant level per the terms of the data use agreement. In order to compare them, I instead merge each separately to the Yelp dataset and report aggregate statistics.

Safegraph reports a name for every restaurant. If the restaurant belongs to a larger chain, it also lists the parent brand. In the Houston-area sample I use for this exercise, there are 294 distinct brands. I proceed by first matching the brand name to the Yelp dataset to get the number of locations in the chain and the Yelp price classification. Because the Yelp data does not contain information for Houston, I am only able to match restaurants that belong to chains that have at least one location in one of the Yelp cities. I do the same exercise for the corresponding sample of the credit card data. I then separately compute the share of transactions that went to restaurants of different chain sizes and price levels (using the Yelp $ rating for prices) in the two datasets.

In spite of these limitations, the two samples look quite comparable. I show the share of credit card transactions and Safegraph visits by chain size in Figure A.1. Because I am only able to match names for chains, I assume that all unmatched restaurants belong to independents or chains with fewer than 100 locations, and I aggregate these two categories together in the figure. I show the share of transactions and visits by price category in Figure A.2. In the second figure, I am only able to report shares for chain restaurants, where I can observe their price category from the Yelp data. The credit card data show slightly more visits to chains (about three percentage points higher for large chains and four percentage points higher for medium chains relative to the Safegraph data) and to $$ restaurants (about 2.5 percentage points higher relative to the Safegraph data), but the observed differences are small given the significant scope for measurement error in this matching exercise.

A.3 Choice set construction

Yelp publishes a sample of data for academic purposes.[54] The data contain reviews and business characteristics for 11 cities, 7 of which are in the US. I construct the choice set in each city by matching restaurants in the payment card data to this sample. I include all matched restaurants where the listed city in the Yelp dataset is Champaign, IL; Charlotte, NC; Pittsburgh, PA; Cleveland, OH; and Madison, WI; this excludes suburbs and cities adjacent to these places. Because complexity of the estimation procedure grows quickly with the size of the choice set, I further restrict the geographic area in the two largest cities, Las Vegas (6,937 Yelp-listed restaurants) and Phoenix (4,220 Yelp-listed restaurants).

n Las Vegas, I restrict the sample to restaurants in the Downtown, Spring Valley, and Southeast neighborhoods, according to a map of Las Vegas neighborhoods, which I reproduce in Figure A.6. These neighborhoods correspond to the following zipcodes: 89101, 89102, 89103, 89104, 89106, 89107, 89109, 89119, 89123, 89146, 89147, and 89169. In Phoenix, I limit restaurants to those available within a 10 mile radius of the downtown area, as defined by the Google Maps pin.

I match restaurants based on restaurant name, zipcode, and address. For restaurants that do not have an exact name match, I perform fuzzy matching using a Jaro-Winkler algorithm. If a restaurant does not have an exact address match but matches on name, I match two entities if their latitude-longitude coordinates in the credit card data and the Yelp data are within 0.25 miles of each other.

I show summary statistics on the result of the merging process in Table A.4 (after limiting Yelp restaurants to the geography described above and dropping observations without a zipcode, address, or valid name). Overall, I am able to match 42% of Yelp restaurants across the seven cities to an entity in the payment card data, containing 44% of the total Yelp reviews. There are several potential reasons that businesses may not be matched. In the

[54]The current version can be freely downloaded here: https://www.yelp.com/dataset/challenge

payment card data, about 25% of entities do not have a valid merchant name, and thus cannot be matched using the procedure described above. Other common reasons for failure to match restaurants include discrepancies in the recorded business name between the two data sources (in the payment card data, businesses self-populate the fields), unpopulated or incorrect address information (e.g. food trucks), the presence of Yelp businesses that do not accept credit cards, restaurants in which the name does not correspond to its management group (e.g. restaurants inside hotels or office parks tend to be identified by the ownership entity in the payment card data and the restaurant name in the Yelp data), and businesses that closed in early 2016, but still had some Yelp reviews. To reduce the presence of false positives that could bias my demand results, I require a close match on the business name and location for inclusion in the analysis sample, which is likely to lower the absolute number of businesses I am able to match. I am also unable to link transactions to specific restaurants when they occur through payment processors like Square.[55]

My main analysis focuses on transactions during dinner hours, which I define as 5pm to 11pm. To construct an accurate choice set for each consumer, I eliminate restaurants that that are not open during these hours. Specifically, I require that each sample restaurant receive at least 100 total transactions during 2016 in the payment cards data and be listed as open during at least four days of the week in the Yelp data. I define a restaurant as open during dinner on a day of the week if their listed hours extend until at least 8pm on a given day. I exclude a small number of restaurants do not have hours listed on Yelp.

A.4 Restaurant price

The payments card data do not contain information on prices of the items purchased. I construct a measure of average entree price during dinner service by combining average transaction size, which I can compute from the credit card data for every restaurant, and a sample of more detailed menu price data, which is available only for 20 large chains. Using the credit card data, I first compute average transaction size, or ticket size, for each restaurant brand as the sum of all dollars spent at brand m divided by the number of swipes

[55] In the payment card data, all businesses that use Square appear as part of a single chain. Because I am not able to reliably parse these transactions, I exclude them from the analysis.

at brand m, summed across all outlets belonging to m.

Average transaction size is highly correlated with Yelp's measure of price, defined on a dollar sign scale that goes from one to four dollar signs. The Yelp measure is based on a survey of Yelp users that review a restaurant. The survey asks about the approximate price for a meal for one person including drink, tax and tip. The translation of the dollar sign measure to prices is as follows: one \$ implies a cost under \$10, two \$\$ between \$11 and \$30, three \$\$\$ between \$31 and \$60, and four \$\$\$\$ above \$60. I show the distribution of log(ticket size) by Yelp dollar sign rating in Figure A.4. The four distributions are monotonic and largely non-overlapping, implying that almost no one dollar sign restaurant has a higher average ticket size than a two dollar sign restaurant, and the same for two and three dollar sign restaurants. There are few four dollar sign restaurants in the data, and thus the distribution of their average ticket size tends to be noisier. This suggests that the average transaction size measure preserves the price rankings of different restaurants.

However, even if the ticket size measure preserves the rankings, the scale of relative prices may be skewed. A steak at a fine dining restaurant may cost five times as much as a fast food hamburger, but its average transaction size may be ten times as high. This may induce bias in my estimates of consumer price sensitivity.

I account for this by adjusting average transaction size to match average entree price computed from a sample of restaurant menu price data for 20 chains from Pricelisto. Pricelisto is a company that collects local pricing information across many business categories. They obtain prices from a variety of sources, including user submission, business submission, and through third parties. They provided me with a sample of restaurant menu data for ten limited service chains and ten full service chains. These data are provided at the chain-location-menu item level. In total, the data contain 4.7M prices that were collected in 2019 and 2020.

For each chain, I clean this data using the following procedure. I start by removing items that were sold in fewer than half of the chain's restaurants. I then remove items that are not dinner entrees to create a standardized price measure for each chain. I remove the following items: appetizers and sides (e.g. french fries, side vegetables, dips, sauces, side salads, and other items that are marked as sides or appetizers); desserts (pies, cakes, ice cream, milk shakes); beverages (soda, coffee, tea, smoothies, and alcoholic drinks); lunch only items (lunch combos and items marked explicitly as only available during lunch); and multi-person meals (family combos, bundles, and catering platters). Using the remaining entree items, I compute the average price for each chain. I show the average prices for limited service and full service restaurants in Tables A.5 and A.6.

I then use these average entree prices, which I observe for only 20 chains, to predict average entree prices from average ticket size for the remaining restaurants in the sample. I do this by computing the ratio between average ticket size and entree price for each of the 20 chains in the Pricelisto data, which I call the deflation factor. This ratio is systematically higher for higher priced restaurants than for lower priced restaurants. It ranges from under two for fast food chains to over six for fine dining chains.[56] I regress this ratio on an intercept, average ticket size, and the square of average ticket size, and use this predicted ratio to convert average restaurant ticket size into average entree price for every restaurant in my sample. I show this ratio as a function of ticket size in Figure A.5.[57]

A.5 Locations

I define the number of locations for each merchant using data from the payments card company. The company records a unique merchant and store identifier associated with each transaction. For each merchant, I define the number of nationwide locations as the number of distinct store identifiers that had at least 100 transactions in 2016.

A.6 Restaurant categories

I use the category "tags" from the Yelp data to assign each restaurant to a one of eight cuisine type categories: Latin-American, European, Pizza, Sandwiches, Asian, Burgers, American, and Other. Restaurants in Yelp are assigned detailed category tags that describe the type of

[56] The data use agreement restricts me from showing any credit card data, including average ticket size, for any individual merchant.

[57] For restaurants with very high transaction sizes, the predicted ratio between ticket size and price begins to decline because of the quadratic term in the regression. I keep this ratio fixed at the maximum predicted ratio for these very high priced restaurants (the maximum ratio is about 5 for restaurants with average ticket size over $175). This is shown in Figure A.5 as the dotted section of the line.

food they sell. Tags can be assigned by restaurant owners or Yelp users, and are sometimes populated algorithmically from review content by Yelp. Restaurants are frequently assigned multiple tags. I show a screenshot from Yelp highlighting these tags for one restaurant in Figure A.3.

I manually map each tag into one of these eight categories. Within a merchant, I assign all outlets to the modal category if different outlets are tagged with different categories. In practice, these are few cases. In Table A.3, I briefly describe the types of restaurants in each category and show some of the most common Yelp tags for each category.

A.7 Movers sample

In Section 4.5.1, I study the behavior of consumers that move across state lines. In addition to data on where a consumer transacts, the payment card company also has information on a consumer's 9-digit billing zipcode over time. The billing zipcode is reported once per year in 2017, 2018, and 2019, and quarterly in 2020. Each report is associated with an exact month (for example, a card may be reported to live in zipcode X in 2017, with the billing zipcode active as of 9/2017). To construct the movers sample, I start with the set of credit cards that have a billing zipcode reported in at least three different years. I define a mover as an account that reported billing zipcodes in exactly two different states in different years. I call the state corresponding to the earliest reported billing zipcode for the account its origin state and the subsequent state its destination state.

Because billing zipcodes are reported annually for much of the sample, I am not able to identify the exact month in which a consumer moved. I proceed by defining each card's move date as the first month it was listed as residing in the destination state. For example, if a credit card was first reported to live in California in 9/2017, California in 9/2018, and then Texas in 9/2019, its origin state would be California, its destination state would be Texas, and its move date would be 9/2019.

I also impose a set of restrictions to ensure that the consumer is primarily transacting in the origin state prior to the move and the destination state after the move. I eliminate cards that conducted less than 80% of their pre-move offline transactions in the origin state or less than 80% of their post-move transactions in the destination state (where pre-move transactions are those that occurred before the last date that the card was reported in the origin state and post-move transactions are those that occurred after the first month that the card was reported in the destination state). I also eliminate cards that conducted more than 5% of their pre-move transactions in the destination state, and cards that were reported returning to their origin state after one or more reports in a new state. These sample restrictions are designed to remove cards that travel frequently, have incorrect billing addresses, or do not update their billing address after moving.

For this sample of movers, I keep transactions that occurred in the card's current state of residence - i.e. its origin state in the pre-move period, its destination state in the post-move period, and in both the origin and destination state during the year of the move. In total, there are 295,777 credit cards that meet these sample restrictions that made 12,785,648 evening time restaurant transactions between 2017 and 2020. I use all transactions by these mover cards in the descriptive analysis. In the estimation exercise described in 4.5.1, I further limit the sample to only the set of cards that transacted at a restaurant in one of the seven sample cities. I present summary statistics on the sample used in estimation in Table A.8.

Figure A.1: Share of credit card transactions and smartphone visits by chain size

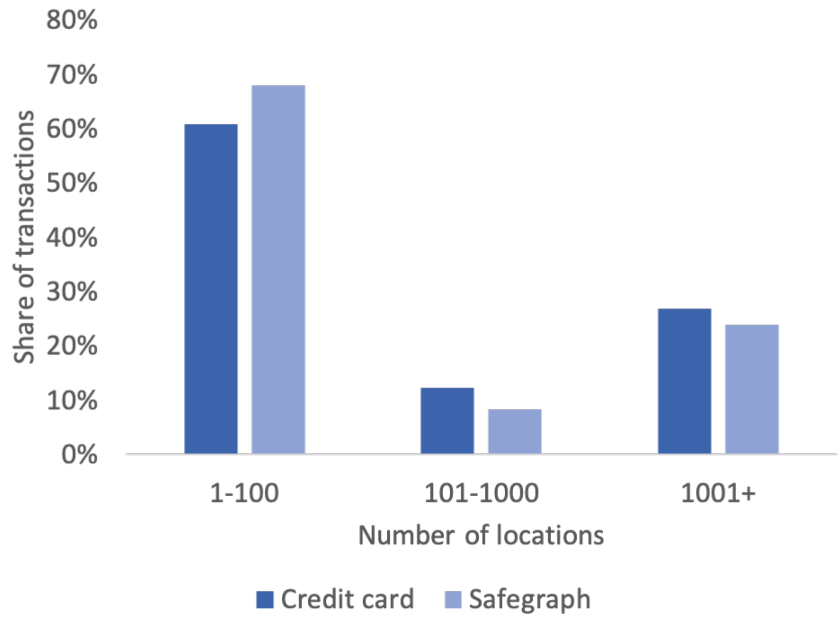

The figure shows the share of transactions recorded in the credit card dataset and smartphone visits recorded in data provided by Safegraph for Houston area restaurants across all mealtimes. I match the Safegraph data to the Yelp data to calculate the number of locations. I am unable to match small chains and independents, as the Yelp data does not cover the Houston area directly, and thus I assume all unmatched restaurants are in the bottom chain size bin (between 1 and 100 locations).

Figure A.2: Share of credit card transactions and smartphone visits by price category

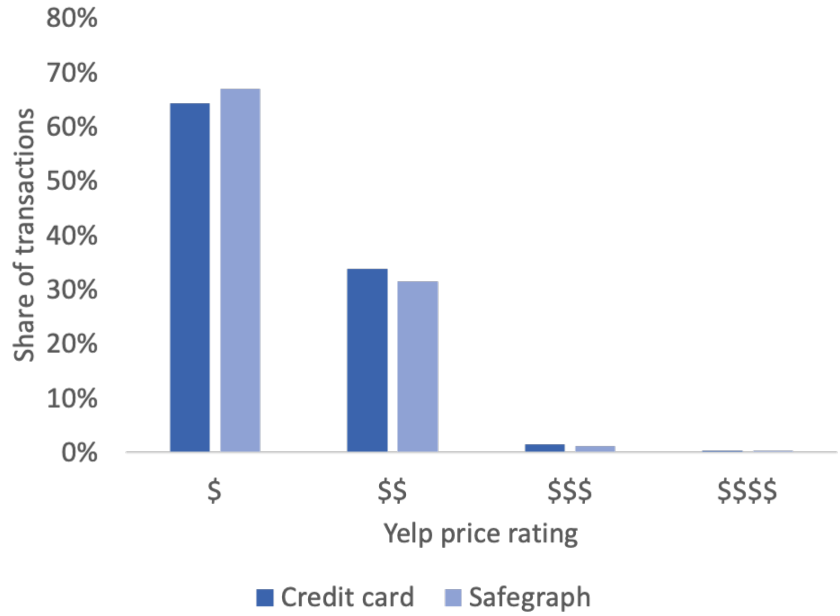

The figure shows the share of transactions recorded in the credit card dataset and smartphone visits recorded in data provided by Safegraph for Houston area restaurants across all mealtimes. I match the Safegraph data to the Yelp data to get the modal price rating for each chain. I am unable to match small chains and independents, as the Yelp data does not cover the Houston area directly, and thus they are excluded from the figure.

Figure A.3: Example of Yelp review with category tags

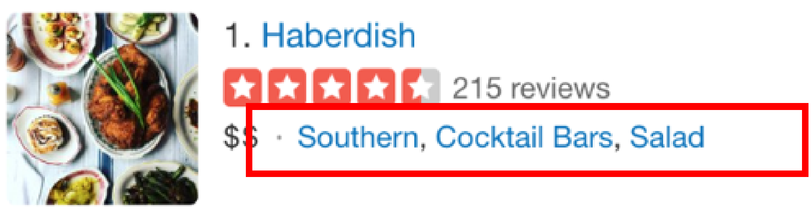

The figure shows a screenshot for the Yelp listing of a restaurant in Charlotte, NC. The red box shows the three category tags for this restaurant. I map these category tags to eight restaurant categories, described above.

Figure A.4: Distribution of log(ticket size) for restaurants by Yelp dollar sign rating

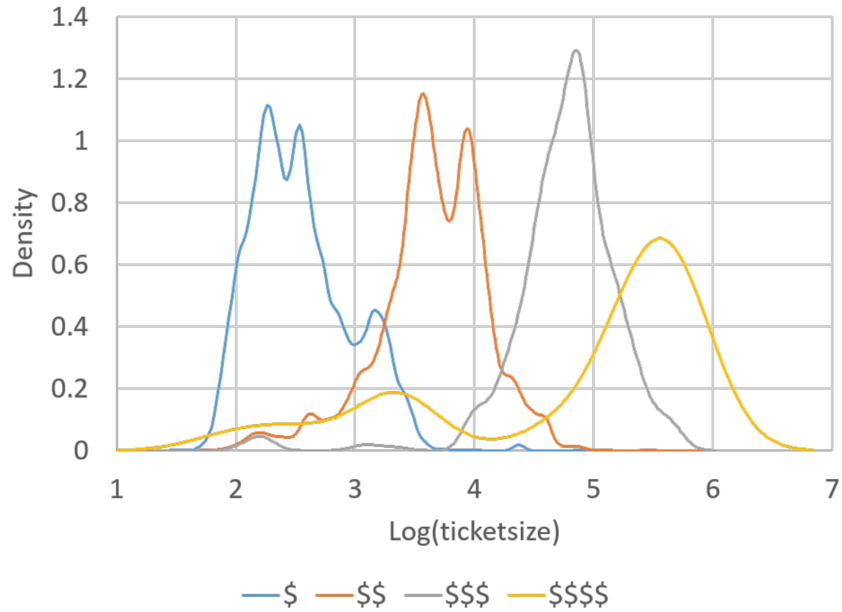

The figure shows the distribution of log(ticket size) for restaurants included in my urban consumer sample by their dollar sign rating on Yelp. I calculate ticket size at the merchant level (so that all outlets within a chain have the same ticket size).

Figure A.5: Ratio of average ticket size to price

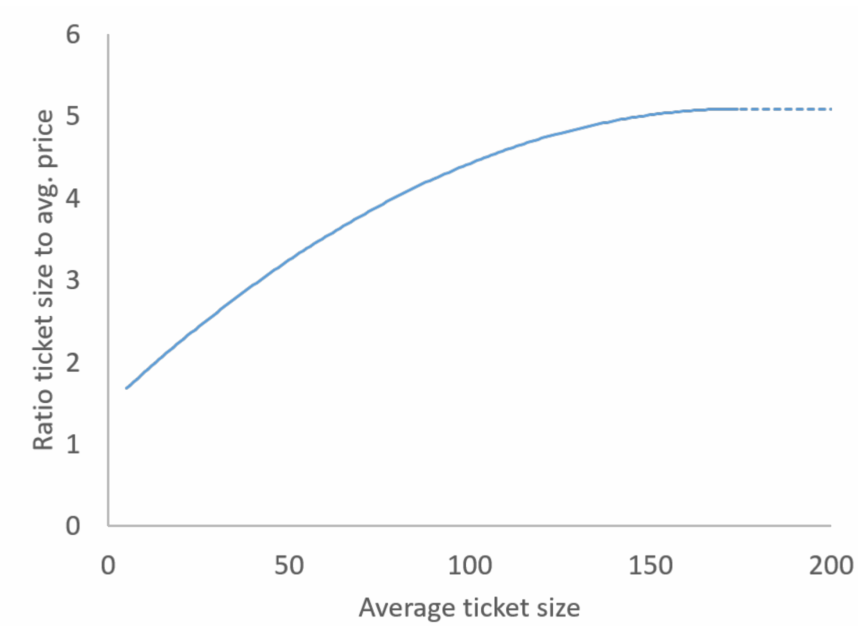

The figure shows the predicted relationship between the ratio of average restaurant ticket size (computed in the credit card data) to average entree price (computed from the Pricelisto sample). The relationship is predicted from a regression of the ratio of ticket size to average price on an intercept, average ticket size, and the square of average ticket size. For average ticket sizes above $175, I hold fixed the ratio at the maximum (the dotted section of the line in the plot).

Figure A.6: Las Vegas Neighborhood Zipcode Map

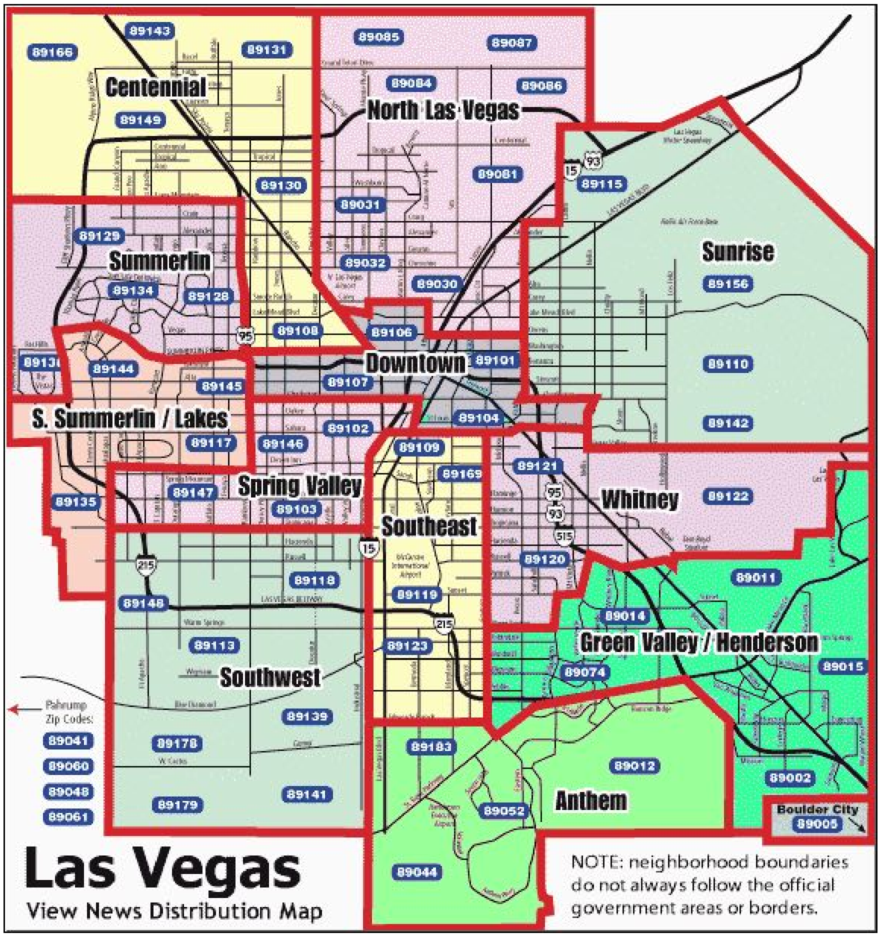

The figure shows a map of Las Vegas by neighborhoods, with their accompanying zipcodes. In my empirical analysis of Las Vegas, I restrict attention to restaurants located in the downtown area, including the Downtown, Spring Valley, and Southeastern neighborhoods on the above map. This includes the official Downtown area, as well as the Strip (primarily in Southeast) and adjacent neighborhoods in Spring Valley. My sample includes consumers that live within 25 miles of this area.

Table A.1: Cities with chain bans or entry restrictions

State	City
CA	Arcata
CA	Benicia
CA	Calistoga
CA	Carmel by the Sea
CA	Coronado
CA	Ojai
CA	Pacific Grove
CA	San Francisco
CA	San Juan Bautista
CA	Sausalito
CA	Solvang
CT	Fairfield
FL	Sanibel
ID	McCall
MA	Nantucket
MD	Chesapeake City
ME	Ogunquit
ME	York
NJ	Jersey City
NY	Port Jefferson
RI	Bristol
TX	Fredericksburg
WA	Bainbridge Island
WA	Port Townsend

Sources: `https://ilsr.org/rule/formula-business-restrictions/`
`https://www.malibucity.org/DocumentCenter/View/4882/PC130729_Item-6D_Correspondence_DWaite2`

Table A.2: Smartphone and credit card penetration by household income

HH income (thousands)	Has credit card	Has smartphone
<25	0.50	0.75
26-50	0.74	0.84
51-100	0.88	0.90
101-150	0.95	0.95
150+	0.97	0.97

The table reports the share of participants in the 2018 Atlanta FRB Survey of Consumer Payment Choice that had at least one credit card or smartphone. Additional details on the survey can be accessed at `https://www.frbatlanta.org/banking-and-payments/consumer-payments/survey-of-consumer-payment-choice`.

Table A.3: Description and Popular Tags of Cuisine Types

Category	Description	Popular Tags
Latin American	Restaurants specializing in cuisines from South and Central America and the Caribbean.	Mexican, Tex-Mex, Latin American, Seafood, Caribbean, Cuban
American	Restaurants serving American cuisine, but excluding restaurants specializing in Burgers and Sandwiches, and excluding restaurants that were also tagged as another type.	American (Traditional), American (New), Breakfast & Brunch, Chicken Wings, Diners
Asian	Restaurants specializing in cuisines from South Asian, East Asian, and Southeast Asian countries, as well as Pacific Islands.	Chinese, Japanese, Sushi Bars, Asian Fusion, Thai, Indian, Hawaiian
Burgers	Restaurants with tag "Burgers".	Burgers, Hot Dogs, Sports Bars, Steakhouses
European	Restaurants specializing in Italian, French, or other European cuisines, except for restaurants also tagged "Pizza".	Italian, French, Irish, Wine Bars, Noodles, Mediterranean
Pizza	Restaurants with tag "Pizza".	Pizza, Italian, Salad
Sandwiches	Restaurants with tag "Sandwiches", "Deli", or "Cheesesteaks".	Sandwiches, Deli, Cheesesteaks
Other	Restaurants not tagged as any of the above categories.	Cafe, Bar, Breakfast & Brunch

Table A.4: Share of Yelp restaurants matched to credit card data

	Restaurants			Reviews		
City	All	Matched	% Matched	All	Matched	% Matched
Champaign	284	152	54%	19,491	12,009	62%
Charlotte	1,874	873	47%	207,271	119,348	58%
Cleveland	1,088	425	39%	81,069	43,621	54%
Las Vegas	2,509	851	34%	906,623	327,926	36%
Madison	783	366	47%	71,976	39,589	55%
Phoenix	1,862	790	42%	311,775	152,235	49%
Pittsburgh	1,738	799	46%	160,359	86,211	54%
Total	10,138	4,256	42%	1,758,564	780,939	44%

The table shows summary statistics for the share of Yelp businesses and reviews by city that I am able to match to an entity in the credit card data in 2016. In the table above, I report all entries in the Yelp data that were categorized as restaurants within the geographic areas that I study. I also filter out Yelp restaurants that did not have a valid zipcode, address, or name, were not open during dinner time, or that had their last Yelp review before 1/1/2016 or after 12/31/2016 (to remove restaurants that were not open during the sample period). The number of restaurants used in estimation is slightly smaller than that reported above, as I further require each restaurant to have at least 100 credit transactions during the year.

Table A.5: Average entree prices for limited service chains

Chain	Average Entree Price
Burger King	5.85
Chipotle	9.63
Domino's	10.84
Five Guys	6.94
KFC	7.39
McDonald's	5.70
Panera Bread	6.42
Pizza Hut	12.95
Subway	7.22
Taco Bell	4.64

The table shows the average entree prices computed from the Pricelisto sample for the ten limited service chains included in my sample.

Table A.6: Average entree prices for full service chains

Chain	Average Entree Price
Applebee's	12.54
Buffalo Wild Wings	10.56
Chili's	12.06
Fleming's Steakhouse	46.98
Morton's Steakhouse	30.83
Olive Garden	14.27
Outback Steakhouse	18.19
P.F. Chang's	14.61
Red Lobster	17.42
Ruth's Chris Steakhouse	36.86

The table shows the average entree prices computed from the Pricelisto sample for the ten full service chains included in my sample.

Table A.7: Summary statistics by cuisine type

Category	# Restaurants	Avg. Transactions	Avg. Dollars	Avg. Accounts
American	325	1,369	87,851	967
Asian	797	943	40,480	616
Burgers	402	1,495	33,600	950
European	353	1,206	76,387	849
Latin	489	1,289	39,559	818
Other	506	1,641	74,677	1,099
Pizza	576	907	31,733	584
Sandwiches	496	765	19,288	479

The table shows summary statistics on restaurants included in the sample used in Section 4 by restaurant cuisine type. Each observation used to create the table is a restaurant. Dollars, accounts, and transactions are computed from sample cards, as described in Section 2.

Table A.8: Summary statistics for movers sample

City	Accounts	Transactions	Dollars	# Unique Origin States
Champaign	196	1,672	45,693	35
Charlotte	2,379	16,555	639,680	46
Cleveland	611	2,952	121,581	42
Lasvegas	1,847	8,455	348,625	47
Madison	785	6,198	203,760	47
Phoenix	2,463	12,164	440,798	47
Pittsburgh	1,178	8,432	333,516	46
Total	9,459	56,428	2,133,653	-

The table shows summary statistics for a sample of cards that move between 2017 and 2020 that transact in one of the seven sample cities. A mover is defined as a card that reports billing zipcodes in two different states (see Appendix A for additional details on sample construction). The sample does not include cards with their origin state in Hawaii or Alaska.

Table A.9: State presence by chain size

Chain size (# locations)	1	2-100	101-1000	1001+
mean	1.0	4.6	27.3	43.4
p25	1	1	17	43
p50	1	2	31	47
p75	1	2	31	47
count	1,517	1,201	110	34

The table shows statistics on the number of states where chains of different sizes had at least one location. The table does not include locations in Hawaii and Alaska.

Figure A.7: Probability that a restaurant is chosen as a function of distance

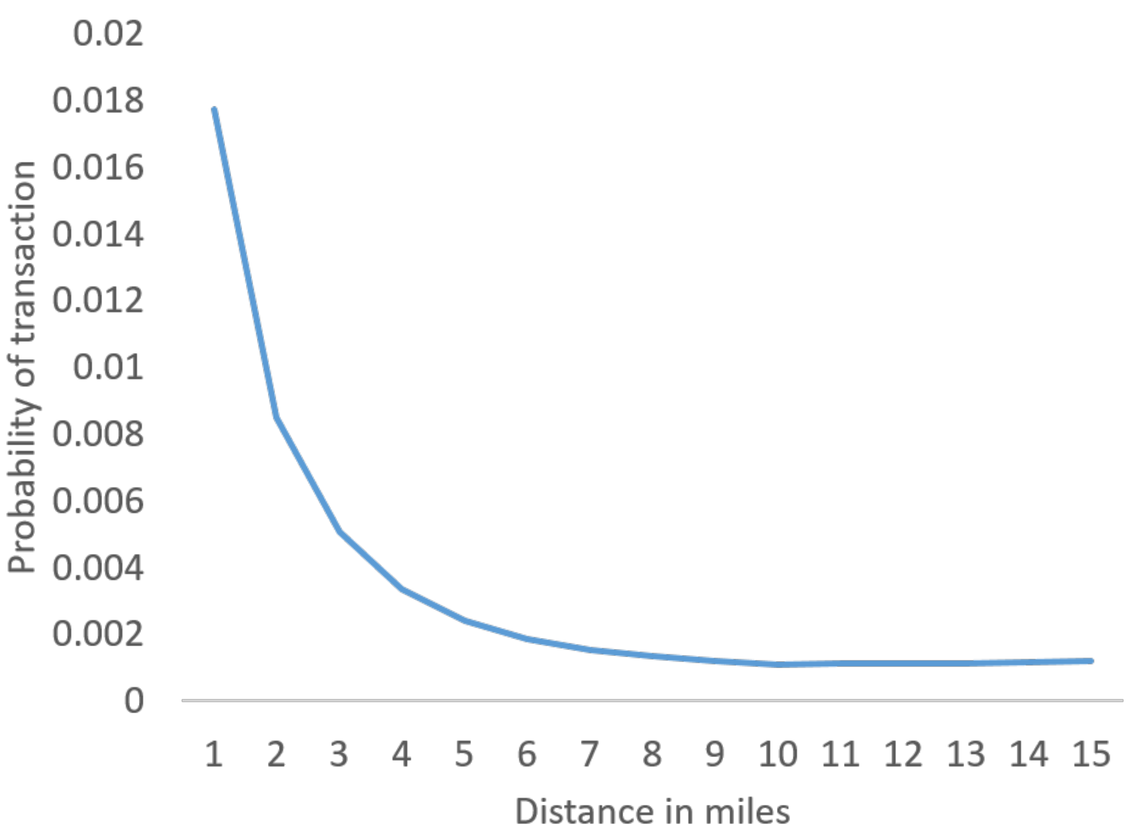

The figure shows the probability that a given restaurant is chosen as a function of the distance between the restaurant and the consumer's home billing address. Each observation in the underlying data is a restaurant-consumer-trip combination for a given cardholder. I average across all possible restaurant choices and consumers in the sample for each one mile bin.

Figure A.8: Share of transactions by movers at chains and independents before and after moving

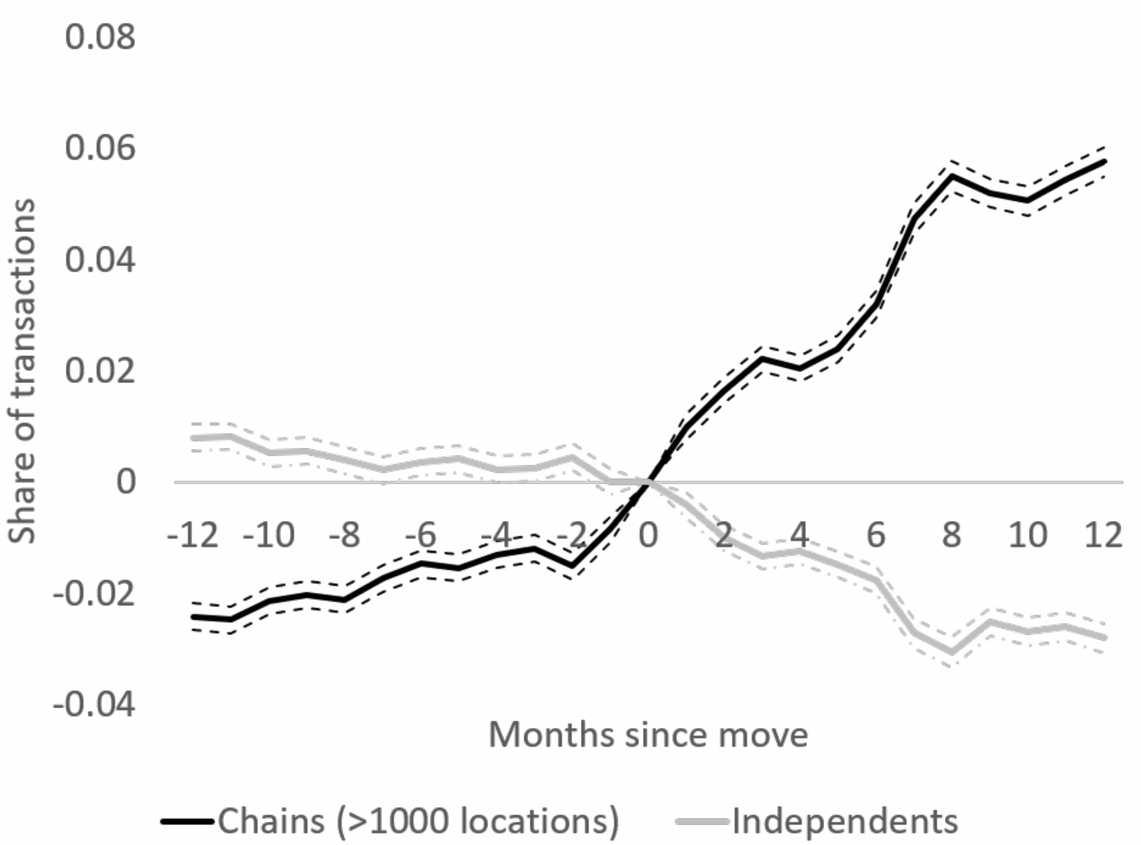

To produce the figure, I calculate the share of transactions by each card in the mover sample at large chains (>1000 locations) and independents (1 location) and regress each share measure on a set of card and state-of-residence fixed effects and set of dummies that correspond to number of months relative to the consumer's move date. The figure plots the dummies on months since move date. I define a mover as a card that reports billing zipcodes in exactly two states in different years, subject to additional restrictions described in Appendix A.

Appendix B Estimation

B.1 Main estimation

In this section, I provide additional details of the estimation of the parameters of demand in equation 1 described in Section 4.2 in the main text. Estimation proceeds in three stages.

In the first stage, I perform maximum simulated likelihood for each city-income group combination with a set of restaurant x income group fixed effects δ_{jcy}. This step is computationally challenging because of the dimension of the choice set for each consumer, which varies between 147 (Champaign) and 817 (Charlotte). Each computation of the choice probabilities requires $J+1 \times M \times R$ calculations, where $J+1$ is the number of restaurants plus the outside good, M is the number of choices, and R is the number of random draws used to approximate the integral. I implement a number of computational tips suggested in Conlon and Gortmaker (2020) (hereafter CG).

I begin by simulating $R = 50$ draws that are fixed for each consumer using a scrambled Halton sequence.[58] Then for each guess of the parameters θ_2 that enter the likelihood in a nonlinear way, I solve for the vector of δ_{jcy} that equate the predicted and observed market shares. This entails finding a fixed point in a system of J equations in J unknowns. The original contraction mapping suggested by Berry et al. (1995) (hereafter BLP) to find the vector of δ_{jcy} converges very slowly in my context, and is a significant bottleneck in estimation. I instead implement the SQUAREM algorithm developed by Varadhan and Roland (2008) (and suggested by CG), which significantly accelerates convergence to the fixed point. I find the values of θ_2 and corresponding vector of mean utilities δ_{jcy} that maximize the log likelihood given in equation 5 using the gradient-based L-BFGS-B algorithm with box constraints on the parameter values, supplying an analytical gradient.[59] This step can be run for each of the 35 city-income group combinations in parallel. Estimation time ranges from 30 minutes in the smallest markets to about 24 hours in the largest using a server instance with 16 CPU cores and a Tesla V100 GPU.

[58] Experiments with larger numbers of Halton draws yielded similar results and were significantly slower, in particular in computing counterfactuals.

[59] My experience is consistent with the findings of CG that the use of an analytical gradient in the BLP-type estimator significantly improves the performance and reliability of estimation.

In the second stage, I regress the vector of mean utilities δ_{jcy} on a vector of brand fixed effects δ_m and city x income, cuisine x city, cuisine x income group, chain size x city, and chain size x income group dummies, plus price times city x income group dummies (collected in x_j):

$$\delta_{jcy} = \delta_m + x_j \beta_{cy} - \bar{\alpha}^p_{cy} p_m + \Delta\xi_{jcy}$$

This stage yields parameter estimates that measure the degree of taste heterogeneity over product characteristics across markets, as well as the differential preference for chains across cities and income groups. Identification of the parameters that vary across cities comes from differences in the popularity of restaurant brands that are available in multiple places. I find that the high preference for large chains among low income consumers is robust to the inclusion of restaurant, rather than brand, fixed effects. The restaurant-level unobservable characteristic $\Delta\xi_{jcy}$ is the residual from the linear regression described above.

In the third and final stage, I recover the mean tastes across all markets for cuisine types, chain size, price, and quality, which are absorbed in the brand dummy variables δ_m. δ_m contains two unobserved product characteristics: q_m and ξ_m. q_m measures the (unobserved) quality of the restaurant's inputs, including its raw ingredients and labor, and is assumed to be linearly increasing in brand m's marginal cost. ξ_m measures any additional utility that consumers receive from purchasing brand m, which could come from the appeal of its menu or concept, skill of its chef, or its accumulated brand capital. Both q_m and ξ_m are likely to be correlated with price p_m.

I do not have an attractive source of exogenous variation to identify the mean coefficient on price $\bar{\alpha}^p$. I instead choose $\bar{\alpha}^p$ to match the marginal cost shares for a set of 12 publicly traded restaurant chains. I take marginal costs from their annual reports as the sum of food costs and labor costs at company owned restaurants. This cost share averages to 61%. The vector of restaurant level marginal costs as a function of the mean price coefficient $\bar{\alpha}^p$ are:

$$\boldsymbol{mc}_c = \boldsymbol{p}_c - \Delta_c(\boldsymbol{p}_c; \bar{\alpha}^p)^{-1} \boldsymbol{s}_c(\boldsymbol{p}_c)$$

where $\Delta_c(\boldsymbol{p}_c; \bar{\alpha}^p) = -\boldsymbol{H}_c \odot \frac{d\boldsymbol{s}_c}{d\boldsymbol{p}_c}(\boldsymbol{p}_c; \bar{\alpha}^p)$ is the element-wise product of the matrix of demand derivatives and the $J_c \times J_c$ ownership matrix $\boldsymbol{H}_c$, which each (j,k) entry equals one if j and k belong to the same chain. I compute marginal costs at the restaurant level and take the average for each brand m. I match the cost shares for the set of 12 chains described above, which gives a mean price parameter of 0.3. I experiment with lower and higher values and find that the counterfactual results do not depend heavily on this parameter.

I then regress $\tilde{\delta}_m = \delta_m + \bar{\alpha}^p p_m$ on the log of quality q_m and a set of chain size and cuisine type dummy variables. The brand level unobservable ξ_m is the residual from this regression. I show the full set of 543 parameters ($9 \times 35 = 315$ from stage 1, 214 from stage 2, and 14 from stage 3) in Tables B.10-B.12 below.

In Monte Carlo exercises, I experimented with versions of the main estimation that allow for additional random coefficients on cuisine types and chain size dummies. The results suggest that there is insufficient variation in categorical regressors to identify a random coefficient. I instead include a random coefficient on $\log l_m$, the logged number of locations for brand m that captures unobservable heterogeneity in the preference for chains. I assume that the random coefficient on $\log l_m$ has a mean zero and allow the intercept terms to vary flexibly by chain size bin.

B.1.1 Bootstrap

I compute standard errors using a nonparametric block bootstrap. For each city-income group, I randomly sample cards with replacement and rerun the full estimation procedure described above. I perform 50 replications for each market. I report the standard deviation over this bootstrapped sample in parentheses in Tables B.10-B.12.

B.2 Movers

In this section, I provide additional details on estimation of the parameter β in equation 6 using the movers sample. Using the set of cards that move described in Appendix A.7, I separate them into markets by destination city-income group, where destination city is one of the seven midsize cities used in estimation. Each card is associated with an origin state, which varies across cards. The estimation exercise measures whether a consumer is more likely to visit a particular chain after moving as a function of whether the chain was operating in the consumer's origin state, conditional on a set of restaurant-city-income group fixed effects. From equation 6, the probability that restaurant j is chosen by card i in income group y and city c is:

$$P_{ijt} = P(y_{it} = j) = \frac{\exp(\delta_{jcy} + \beta Avail_{ij})}{\sum_{j'} \exp(\delta_{j'cy} + \beta Avail_{ij'})}$$

Because the movers sample is smaller than the set of cards used in the main estimation, some restaurants included in the main estimation are never chosen by a mover, and so I eliminate these from the choice set. I only model consumer choice conditional on choosing an inside good restaurant. I pool estimation across all markets to estimate β. For each guess of β, I solve for the vector of δ_{jcy} that equate observed and predicted shares in that city-income group as described above and compute the log likelihood summing observations across all markets. I compute a standard error for β using a bootstrap routine with 50 replications, which I report in the main text.

Table B.10: Estimates of nonlinear parameters θ_2

(a) Estimates of γ (distance sensitivity) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	-0.466	-0.310	-0.333	-0.391	-0.362
	(0.019)	(0.014)	(0.022)	(0.037)	(0.041)
Charlotte	-0.299	-0.292	-0.292	-0.298	-0.303
	(0.003)	(0.003)	(0.004)	(0.006)	(0.006)
Cleveland	-0.250	-0.215	-0.185	-0.172	-0.168
	(0.004)	(0.004)	(0.005)	(0.007)	(0.007)
Las Vegas	-0.323	-0.309	-0.293	-0.294	-0.300
	(0.003)	(0.003)	(0.006)	(0.009)	(0.010)
Madison	-0.367	-0.320	-0.321	-0.326	-0.345
	(0.006)	(0.004)	(0.008)	(0.010)	(0.010)
Phoenix	-0.321	-0.296	-0.285	-0.279	-0.305
	(0.003)	(0.003)	(0.004)	(0.005)	(0.006)
Pittsburgh	-0.372	-0.347	-0.338	-0.346	-0.359
	(0.004)	(0.003)	(0.004)	(0.007)	(0.006)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (a) shows estimates of γ, the parameter that describes consumer sensitivity to travel distance. Bootstrapped standard errors are shown in parentheses.

(b) Estimates of $\sigma_{cy}^{\alpha^l}$ (variance of taste for chains) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	0.035	0.033	0.032	0.039	0.043
	(0.003)	(0.004)	(0.004)	(0.008)	(0.009)
Charlotte	0.033	0.034	0.033	0.031	0.033
	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
Cleveland	0.048	0.047	0.047	0.047	0.056
	(0.003)	(0.003)	(0.004)	(0.007)	(0.010)
Las Vegas	0.032	0.031	0.032	0.031	0.030
	(0.002)	(0.002)	(0.004)	(0.005)	(0.006)
Madison	0.033	0.043	0.034	0.036	0.036
	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)
Phoenix	0.039	0.041	0.044	0.039	0.043
	(0.002)	(0.002)	(0.003)	(0.004)	(0.005)
Pittsburgh	0.035	0.037	0.033	0.031	0.034
	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (b) shows estimates of $\sigma_{cy}^{\alpha^l}$, the variance of the random coefficient on log(locations) α_i^l. Bootstrapped standard errors are shown in parentheses.

(c) Estimates of $\sigma_{cy}^{\alpha^p}$ (variance of price sensitivity) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	0.00577	0.00608	0.00545	0.00422	0.00355
	(0.00136)	(0.00148)	(0.00150)	(0.00175)	(0.00257)
Charlotte	0.00336	0.00333	0.00283	0.00245	0.00248
	(0.00030)	(0.00026)	(0.00029)	(0.00046)	(0.00032)
Cleveland	0.00227	0.00172	0.00129	0.00105	0.00077
	(0.00043)	(0.00028)	(0.00032)	(0.00045)	(0.00032)
Las Vegas	0.00227	0.00205	0.00177	0.00152	0.00126
	(0.00021)	(0.00019)	(0.00023)	(0.00046)	(0.00038)
Madison	0.00350	0.00334	0.00269	0.00286	0.00252
	(0.00044)	(0.00040)	(0.00036)	(0.00057)	(0.00063)
Phoenix	0.00306	0.00238	0.00180	0.00166	0.00153
	(0.00023)	(0.00021)	(0.00031)	(0.00030)	(0.00027)
Pittsburgh	0.00268	0.00246	0.00213	0.00211	0.00197
	(0.00029)	(0.00025)	(0.00024)	(0.00042)	(0.00030)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (c) shows estimates of $\sigma_{cy}^{\alpha^p}$, the variance of the random coefficient on price α_i^p. Bootstrapped standard errors are shown in parentheses.

(d) Estimates of ρ_{cy} (covariance of price sensitivity and chain preference) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	-0.008	-0.010	-0.008	-0.009	-0.005
	(0.002)	(0.002)	(0.003)	(0.004)	(0.005)
Charlotte	-0.008	-0.008	-0.007	-0.007	-0.006
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Cleveland	-0.009	-0.007	-0.007	-0.005	-0.005
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Las Vegas	-0.008	-0.007	-0.007	-0.006	-0.006
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Madison	-0.007	-0.009	-0.007	-0.008	-0.007
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Phoenix	-0.010	-0.009	-0.008	-0.007	-0.006
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Pittsburgh	-0.007	-0.007	-0.007	-0.006	-0.006
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (d) shows estimates of ρ_{cy}, the covariance of the random coefficients α_i^p and α_i^l. Bootstrapped standard errors are shown in parentheses.

(e) Estimates of σ^v_{cy} (variance of outside option) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	0.745	0.545	0.518	0.620	0.489
	(0.366)	(0.319)	(0.605)	(0.736)	(0.767)
Charlotte	0.596	0.597	0.608	0.609	0.694
	(0.094)	(0.093)	(0.124)	(0.219)	(0.167)
Cleveland	0.666	0.642	0.698	0.737	0.773
	(0.144)	(0.134)	(0.151)	(0.331)	(0.231)
Las Vegas	0.427	0.418	0.465	0.438	0.546
	(0.114)	(0.069)	(0.154)	(0.208)	(0.250)
Madison	0.691	0.690	0.668	0.575	0.537
	(0.123)	(0.159)	(0.187)	(0.284)	(0.303)
Phoenix	0.650	0.701	0.689	0.675	0.710
	(0.090)	(0.090)	(0.106)	(0.147)	(0.160)
Pittsburgh	0.651	0.661	0.659	0.576	0.705
	(0.099)	(0.106)	(0.121)	(0.180)	(0.214)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (e) shows estimates of σ^v_{cy}, the variance of the random intercept term from visiting the outside option v_{i0}. Bootstrapped standard errors are shown in parentheses.

(f) Estimates of $\pi_{cy}^{0,2}$ (utility of outside good $\times$ # restaurants within 2 miles) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	0.235	0.100	0.177	0.136	0.066
	(0.063)	(0.065)	(0.074)	(0.116)	(0.111)
Charlotte	0.121	0.076	0.065	0.054	0.095
	(0.023)	(0.027)	(0.032)	(0.048)	(0.047)
Cleveland	0.221	0.119	0.104	0.075	-0.044
	(0.031)	(0.026)	(0.051)	(0.058)	(0.058)
Las Vegas	0.075	0.056	0.029	0.113	0.113
	(0.023)	(0.025)	(0.030)	(0.054)	(0.062)
Madison	0.098	0.083	0.079	0.062	0.057
	(0.042)	(0.033)	(0.036)	(0.057)	(0.046)
Phoenix	0.188	0.143	0.125	0.079	0.053
	(0.022)	(0.025)	(0.034)	(0.046)	(0.045)
Pittsburgh	0.021	0.028	0.009	0.047	0.008
	(0.028)	(0.027)	(0.029)	(0.049)	(0.047)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (f) shows estimates of $\pi_{cy}^{0,2}$, the coefficient on the log of the number of outside option restaurants within 2 miles of a consumer's home (included in equation 2). Bootstrapped standard errors are shown in parentheses.

(g) Estimates of $\pi_{cy}^{0,5}$ (utility of outside good × # restaurants within 5 miles) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	0.379	0.222	0.141	0.064	0.398
	(0.114)	(0.071)	(0.135)	(0.197)	(0.211)
Charlotte	0.026	0.096	0.046	0.145	-0.021
	(0.048)	(0.050)	(0.052)	(0.102)	(0.116)
Cleveland	-0.248	-0.195	-0.093	-0.037	-0.061
	(0.087)	(0.065)	(0.077)	(0.153)	(0.142)
Las Vegas	0.081	0.086	0.104	0.136	0.191
	(0.062)	(0.047)	(0.072)	(0.119)	(0.117)
Madison	0.166	0.053	0.017	0.093	0.066
	(0.079)	(0.052)	(0.055)	(0.090)	(0.111)
Phoenix	0.124	0.120	0.103	0.082	0.166
	(0.041)	(0.047)	(0.060)	(0.094)	(0.084)
Pittsburgh	0.046	-0.002	0.006	0.051	0.205
	(0.070)	(0.050)	(0.065)	(0.106)	(0.115)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (g) shows estimates of $\pi_{cy}^{0,5}$, the coefficient on the log of the number of outside option restaurants within 5 miles of a consumer's home (included in equation 2). Bootstrapped standard errors are shown in parentheses.

(h) Estimates of $\pi_{cy}^{0,10}$ (utility of outside good × # restaurants within 10 miles) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	1.782	1.218	1.343	1.840	1.084
	(0.174)	(0.115)	(0.190)	(0.352)	(0.359)
Charlotte	0.764	0.570	0.605	0.535	1.086
	(0.064)	(0.065)	(0.091)	(0.161)	(0.175)
Cleveland	1.913	1.655	1.222	1.169	1.435
	(0.141)	(0.102)	(0.114)	(0.209)	(0.251)
Las Vegas	0.976	0.904	0.829	0.622	0.593
	(0.098)	(0.084)	(0.122)	(0.235)	(0.189)
Madison	1.550	1.361	1.360	1.369	1.517
	(0.090)	(0.065)	(0.088)	(0.152)	(0.211)
Phoenix	0.210	0.125	0.115	0.164	0.011
	(0.060)	(0.067)	(0.067)	(0.136)	(0.115)
Pittsburgh	1.160	1.010	1.001	0.749	0.779
	(0.081)	(0.074)	(0.104)	(0.175)	(0.160)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (h) shows estimates of $\pi_{cy}^{0,10}$, the coefficient on the log of the number of outside option restaurants within 10 miles of a consumer's home (included in equation 2). Bootstrapped standard errors are shown in parentheses.

(i) Estimates of $\pi_{cy}^{0,25}$ (utility of outside good × # restaurants within 25 miles) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	1.871	1.090	1.114	0.941	1.631
	(0.519)	(0.358)	(0.535)	(0.855)	(1.159)
Charlotte	2.582	2.577	2.868	2.541	1.887
	(0.163)	(0.148)	(0.251)	(0.376)	(0.548)
Cleveland	0.669	0.505	0.646	0.494	0.153
	(0.141)	(0.109)	(0.118)	(0.207)	(0.294)
Las Vegas	0.250	0.563	0.716	2.413	1.556
	(0.739)	(0.687)	(1.113)	(2.167)	(1.433)
Madison	2.173	1.915	1.755	1.725	1.098
	(0.296)	(0.274)	(0.343)	(0.639)	(0.788)
Phoenix	3.608	3.430	3.482	3.377	4.372
	(0.128)	(0.091)	(0.140)	(0.290)	(0.226)
Pittsburgh	4.363	4.238	4.071	4.826	5.076
	(0.143)	(0.135)	(0.202)	(0.352)	(0.604)

The table shows estimates of the parameters collected in the vector θ_2, described in Section 4.3 and Appendix B. Panel (i) shows estimates of $\pi_{cy}^{0,25}$, the coefficient on the log of the number of outside option restaurants within 25 miles of a consumer's home (included in equation 2). Bootstrapped standard errors are shown in parentheses.

Table B.11: Estimates of stage 2 linear parameters $(\beta, \bar{\alpha^p})$

(a) Estimates of β_c^l (mean utility of chains) by city

	Chain size (locations)			
City	1	2-100	101-1000	1001+
Champaign	0.000	0.000	0.319	0.915
	(0.000)	(0.000)	(0.084)	(0.090)
Charlotte	0.000	0.000	0.416	0.745
	(0.000)	(0.000)	(0.064)	(0.072)
Cleveland	0.000	0.000	-0.634	-1.008
	(0.000)	(0.000)	(0.083)	(0.103)
Las Vegas	0.000	0.000	0.118	0.649
	(0.000)	(0.000)	(0.054)	(0.076)
Madison	0.000	0.000	0.831	1.400
	(0.000)	(0.000)	(0.086)	(0.097)
Phoenix	0.000	0.000	0.567	1.121
	(0.000)	(0.000)	(0.083)	(0.079)
Pittsburgh	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)

The table shows estimates of the parameters collected in the vector $(\beta, \bar{\alpha^p})$ and estimated by regressing δ_{jcy} on a set of merchant fixed effects and restaurant characteristics interacted with city and income variables, described in Section 4.3 and Appendix B. Panel (a) shows estimates of β_c^l, the categorical variables on chain size bins interacted with city dummies. Demand for restaurants with 1 or 2-100 locations are absorbed by the brand fixed effects and city-income dummies, and are not estimated in this stage. Bootstrapped standard errors are shown in parentheses.

(b) Estimates of β_y^l (mean utility of chains) by income group

	Chain size (locations)			
Income	1	2-100	101-1000	1001+
<\$50k	0.000	0.032	0.255	0.474
	(0.000)	(0.016)	(0.035)	(0.052)
\$50-100k	0.000	0.030	0.232	0.326
	(0.000)	(0.018)	(0.038)	(0.052)
\$100-150k	0.000	0.016	0.166	0.203
	(0.000)	(0.017)	(0.038)	(0.053)
\$150-200k	0.000	-0.017	0.107	0.097
	(0.000)	(0.022)	(0.047)	(0.073)
>\$200k	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)

The table shows estimates of the parameters collected in the vector $(\beta, \bar{\alpha^p})$ and estimated by regressing δ_{jcy} on a set of merchant fixed effects and restaurant characteristics interacted with city and income variables, described in Section 4.3 and Appendix B. Panel (b) shows estimates of β_y^l, the categorical variables on chain size bins interacted with income group dummies. Bootstrapped standard errors are shown in parentheses.

(c) Estimates of $\beta_c^{cuisine}$ (mean utility of cuisine types) by city

City	Cuisine							
	American	Asian	Burgers	European	Latin	Other	Pizza	Sandwiches
Champaign	-0.045	0.227	-0.739	0.063	-0.098	-0.215	0.319	0.000
	(0.089)	(0.080)	(0.051)	(0.086)	(0.059)	(0.082)	(0.076)	(0.000)
Charlotte	0.337	0.707	-0.284	0.396	0.057	-0.147	0.853	0.000
	(0.072)	(0.063)	(0.031)	(0.052)	(0.037)	(0.056)	(0.060)	(0.000)
Cleveland	0.181	-1.447	-0.878	-2.541	-0.078	-0.729	0.012	0.000
	(0.078)	(0.113)	(0.056)	(0.218)	(0.053)	(0.080)	(0.085)	(0.000)
Las Vegas	0.546	0.196	-0.312	0.508	-0.392	0.189	0.716	0.000
	(0.070)	(0.057)	(0.042)	(0.067)	(0.046)	(0.058)	(0.067)	(0.000)
Madison	0.550	0.767	-0.069	0.972	0.081	0.138	0.749	0.000
	(0.063)	(0.140)	(0.041)	(0.054)	(0.040)	(0.059)	(0.066)	(0.000)
Phoenix	0.345	0.502	-0.326	-0.541	-0.297	0.139	0.255	0.000
	(0.062)	(0.074)	(0.040)	(0.085)	(0.044)	(0.053)	(0.062)	(0.000)
Pittsburgh	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The table shows estimates of the parameters collected in the vector $(\beta, \bar{\alpha^p})$ and estimated by regressing δ_{jcy} on a set of merchant fixed effects and restaurant characteristics interacted with city and income variables, described in Section 4.3 and Appendix B. Panel (c) shows estimates of $\beta_c^{cuisine}$, the categorical variables on restaurant cuisine types interacted with city dummies. Demand for restaurants with 1 location for each income group are absorbed by the brand fixed effects. Dummies for Pittsburgh and the Sandwich category are absorbed by market and brand fixed effects. Bootstrapped standard errors are shown in parentheses.

(d) Estimates of $\beta_y^{cuisine}$ (mean utility of cuisine types) by income group

	Cuisine							
Income	American	Asian	Burgers	European	Latin	Other	Pizza	Sandwiches
<\$50k	0.319	0.287	0.120	0.141	0.236	0.239	0.262	0.000
	(0.034)	(0.034)	(0.031)	(0.031)	(0.033)	(0.031)	(0.031)	(0.000)
\$50-100k	0.253	0.114	0.119	0.086	0.161	0.193	0.246	0.000
	(0.034)	(0.035)	(0.033)	(0.032)	(0.036)	(0.032)	(0.032)	(0.000)
\$100-150k	0.083	0.059	0.035	-0.004	0.043	0.061	0.161	0.000
	(0.038)	(0.036)	(0.034)	(0.029)	(0.036)	(0.035)	(0.030)	(0.000)
\$150-200k	0.001	0.033	0.015	-0.005	0.001	0.042	0.060	0.000
	(0.042)	(0.045)	(0.044)	(0.043)	(0.041)	(0.042)	(0.045)	(0.000)
>\$200k	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The table shows estimates of the parameters collected in the vector $(\beta, \bar{\alpha^p})$ and estimated by regressing δ_{jcy} on a set of merchant fixed effects and restaurant characteristics interacted with city and income variables, described in Section 4.3 and Appendix B. Panel (d) shows estimates of $\beta_y^{cuisine}$, the categorical variables on restaurant cuisine types interacted with income group dummies. Dummies for $> 200k$ and the Sandwich category are absorbed by market and brand fixed effects. Bootstrapped standard errors are shown in parentheses.

(e) Estimates of $\bar{\alpha}^p_{cy}$ (mean price coefficient) by city and income group

	Income				
City	<\$50k	\$50-100k	\$100-150k	\$150-200k	>\$200k
Champaign	-0.056	-0.012	-0.014	0.016	0.032
	(0.013)	(0.012)	(0.013)	(0.014)	(0.020)
Charlotte	-0.074	-0.063	-0.042	-0.030	-0.012
	(0.006)	(0.007)	(0.007)	(0.008)	(0.007)
Cleveland	-0.019	0.011	0.043	0.060	0.089
	(0.010)	(0.009)	(0.010)	(0.011)	(0.015)
Las Vegas	-0.061	-0.050	-0.034	-0.024	-0.008
	(0.007)	(0.007)	(0.007)	(0.010)	(0.009)
Madison	-0.089	-0.051	-0.036	-0.031	-0.022
	(0.007)	(0.007)	(0.006)	(0.009)	(0.008)
Phoenix	-0.062	-0.030	0.004	0.017	0.033
	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)
Pittsburgh	-0.055	-0.033	-0.023	-0.017	0.000
	(0.005)	(0.005)	(0.005)	(0.007)	(0.000)

The table shows estimates of the parameters collected in the vector $(\beta, \bar{\alpha^p})$ and estimated by regressing δ_{jcy} on a set of merchant fixed effects and restaurant characteristics interacted with city and income variables, described in Section 4.3 and Appendix B. Panel (e) shows estimates of $\bar{\alpha}^p_{cy}$, the mean price sensitivity in each city-income group relative to the excluded category (Pittsburgh consumers with income $> \$200k$. Bootstrapped standard errors are shown in parentheses.

Table B.12: Estimates of stage 3 linear parameters θ_1

Parameter	Estimate	Std. Error
α^p	0.300	0.000
α^q	4.316	0.103
β^1	-0.441	0.091
β^{2-100}	-0.427	0.087
$\beta^{101-1000}$	-0.649	0.084
β^{1001+}	-1.007	0.104
$\beta^{American}$	-0.200	0.054
β^{Asian}	-0.749	0.075
$\beta^{Burgers}$	0.412	0.054
$\beta^{European}$	0.040	0.085
β^{Latin}	-0.178	0.043
β^{Other}	-0.127	0.064
β^{Pizza}	-0.994	0.058
$\beta^{Sandwiches}$	0.000	0.000

The table shows estimates of the parameters collected in the vector θ_1, which are absorbed in the merchant fixed effects. These coefficients are estimated by regressing the merchant fixed effects, less the the mean utility related to price which is calibrated from data on marginal costs, on $\log(q_m)$ and a set of cuisine type and chain size dummies. Estimation is described in more detail in Section 4.3 and Appendix B. Bootstrapped standard errors are shown in parentheses.

Table B.13: Estimates of chain premium by city and income group

City	Income Group	1	2-100	101-1000	1001+
Champaign	<$50k	0.000	0.046	0.366	0.824
		(0.000)	(0.017)	(0.080)	(0.094)
Champaign	$50-100k	0.000	0.044	0.343	0.676
		(0.000)	(0.017)	(0.085)	(0.095)
Champaign	$100-150k	0.000	0.030	0.277	0.552
		(0.000)	(0.019)	(0.083)	(0.095)
Champaign	$150-200k	0.000	-0.003	0.218	0.446
		(0.000)	(0.020)	(0.085)	(0.091)
Champaign	>$200k	0.000	0.014	0.111	0.349
		(0.000)	(0.020)	(0.088)	(0.109)
Charlotte	<$50k	0.000	0.046	0.463	0.654
		(0.000)	(0.017)	(0.053)	(0.065)
Charlotte	$50-100k	0.000	0.044	0.441	0.506
		(0.000)	(0.017)	(0.055)	(0.064)
Charlotte	$100-150k	0.000	0.030	0.374	0.382
		(0.000)	(0.019)	(0.060)	(0.071)
Charlotte	$150-200k	0.000	-0.003	0.315	0.276
		(0.000)	(0.020)	(0.053)	(0.073)
Charlotte	>$200k	0.000	0.014	0.208	0.180
		(0.000)	(0.020)	(0.062)	(0.077)
Cleveland	<$50k	0.000	0.046	-0.587	-1.100
		(0.000)	(0.017)	(0.079)	(0.107)
Cleveland	$50-100k	0.000	0.044	-0.610	-1.248
		(0.000)	(0.017)	(0.078)	(0.107)

Continued on next page

Table B.13: Estimates of chain premium by city and income group

City	Income Group	1	2-100	101-1000	1001+
Cleveland	\$100-150k	0.000	0.030	-0.676	-1.371
		(0.000)	(0.019)	(0.081)	(0.112)
Cleveland	\$150-200k	0.000	-0.003	-0.735	-1.477
		(0.000)	(0.020)	(0.081)	(0.119)
Cleveland	>\$200k	0.000	0.014	-0.842	-1.574
		(0.000)	(0.020)	(0.094)	(0.122)
Lasvegas	<\$50k	0.000	0.046	0.164	0.558
		(0.000)	(0.017)	(0.044)	(0.074)
Lasvegas	\$50-100k	0.000	0.044	0.142	0.409
		(0.000)	(0.017)	(0.046)	(0.076)
Lasvegas	\$100-150k	0.000	0.030	0.075	0.286
		(0.000)	(0.019)	(0.049)	(0.080)
Lasvegas	\$150-200k	0.000	-0.003	0.016	0.180
		(0.000)	(0.020)	(0.052)	(0.081)
Lasvegas	>\$200k	0.000	0.014	-0.091	0.083
		(0.000)	(0.020)	(0.049)	(0.090)
Madison	<\$50k	0.000	0.046	0.878	1.308
		(0.000)	(0.017)	(0.085)	(0.095)
Madison	\$50-100k	0.000	0.044	0.855	1.160
		(0.000)	(0.017)	(0.090)	(0.102)
Madison	\$100-150k	0.000	0.030	0.789	1.037
		(0.000)	(0.019)	(0.086)	(0.096)
Madison	\$150-200k	0.000	-0.003	0.730	0.931
		(0.000)	(0.020)	(0.087)	(0.099)
Madison	>\$200k	0.000	0.014	0.623	0.834
		(0.000)	(0.020)	(0.090)	(0.110)

Continued on next page

Table B.13: Estimates of chain premium by city and income group

City	Income Group	1	2-100	101-1000	1001+
Phoenix	<\$50k	0.000	0.046	0.614	1.029
		(0.000)	(0.017)	(0.069)	(0.078)
Phoenix	\$50-100k	0.000	0.044	0.591	0.881
		(0.000)	(0.017)	(0.065)	(0.072)
Phoenix	\$100-150k	0.000	0.030	0.525	0.758
		(0.000)	(0.019)	(0.068)	(0.080)
Phoenix	\$150-200k	0.000	-0.003	0.466	0.652
		(0.000)	(0.020)	(0.071)	(0.082)
Phoenix	>\$200k	0.000	0.014	0.359	0.555
		(0.000)	(0.020)	(0.079)	(0.095)
Pittsburgh	<\$50k	0.000	0.046	0.047	-0.091
		(0.000)	(0.017)	(0.040)	(0.058)
Pittsburgh	\$50-100k	0.000	0.044	0.024	-0.239
		(0.000)	(0.017)	(0.041)	(0.054)
Pittsburgh	\$100-150k	0.000	0.030	-0.042	-0.363
		(0.000)	(0.019)	(0.045)	(0.056)
Pittsburgh	\$150-200k	0.000	-0.003	-0.101	-0.469
		(0.000)	(0.020)	(0.048)	(0.068)
Pittsburgh	>\$200k	0.000	0.014	-0.208	-0.566
		(0.000)	(0.020)	(0.046)	(0.071)

The table shows the chain premium by city and income group with bootstrapped standard errors in parentheses. The chain premium for size bin l is defined as the sum of the fixed effects on chain size for l (estimated by income and city in stage 2 of the estimation and across the entire sample in stage 3) minus the sum of the fixed effects for chains with 1 location. See Tables B.10 and B.11 for the parameter estimates from each stage and Section 4.3 and Appendix B for estimation details.

Appendix C Counterfactuals

In this section, I provide additional details for the two sets of counterfactuals described in Section 5.

C.1 Chain reoptimization counterfactuals

The first set of counterfactuals quantifies the variable profits that a chain would earn if it could be flexible in setting its product characteristics in each city. I compute chain profits under three different assumptions about its optimization behavior: (1) full standardization: a chain chooses one quality level and cuisine type that is fixed across markets; (2) flexible chain: each chain outlet chooses its quality and cuisine type to maximize its own profits, keeping chain demand; (3) chain becomes independent: each chain outlet chooses quality and cuisine type flexibly but faces the demand of an independent restaurant. In each of these, I assume that the chain first chooses its quality and cuisine type and then sets prices a la Nash Bertrand. I do this exercise separately for each chain with more than 1000 locations that operates in at least 4 cities.[60] The product characteristics and prices of all other firms remain fixed at their values in the data.

These counterfactuals do not allow for competing firms to respond by changing their own product characteristics. A previous version of the paper estimated a version of these counterfactuals where 20 nearby independent firms were allowed to change their own product characteristics. The impact on chain profits of allowing for this competitive response was minimal and it imposes significant additional computational burden, as it requires the computation of a new market equilibrium in each iteration of the chain's optimization problem.

C.2 Chain ban

To simulate the effects of a chain ban, I assume that each restaurant that belongs to a chain with over 1000 locations is replaced by an independent restaurant that operates in the same

[60]There is no trade-off in the model for chains that operate in only one city. 94% of large chain outlets in the sample belong to firms that have restaurants in at least 4 of 7 cities.

location. I allow for each of these outlets to choose their quality and cuisine type and then their price to maximize their variable profits as described in the main text. I then compute consumer welfare under the new chain ban equilibrium. I compare this consumer welfare to a baseline scenario in which each chain chooses the quality and cuisine type that it would choose if it were to be fully standardized across markets.

The set of chain ban counterfactuals discussed in the main text assumes that restaurants optimize per the demand parameters of the model. Restaurants may face other constraints in choosing their product characteristics—for example, they may consider the preferences of consumers at breakfast and lunch times, or the preferences of consumers in markets that are not included in my sample.

As a robustness check, I perform an alternate version of these counterfactuals that infers product characteristics from those chosen by firms in the data. First, I assume that the independent firm that replaces the chain chooses identical product characteristics to those chosen by the chain in the data (column (4) in Table C.14). This scenario isolates the effect of losing the utility benefits that come from chains, but does not allow for potentially welfare-improving customization. Second, I assume that the replacement firm takes the characteristics of some other independent firm in the data (column (5)). Independents should be more locally customized relative to chains because they do not face the standardization constraint in choosing their product characteristics, and thus this scenario allows for some limited improvement in the match between product characteristics and local tastes. I compare the consumer welfare in these alternate counterfactuals to the surplus that consumers receive when all firms have the set of characteristics observed in the data.

The welfare impacts are somewhat smaller in these two alternative scenarios, and I consider these in the range of magnitudes when discussing the trade-offs facing local policymakers. The distributional effect of a chain ban across income groups is quite robust, and is driven primarily by the disproportionate chain preferences of lower income groups relative to higher income groups.

Table C.14: Welfare effects of a chain ban under alternative assumptions

	(1)	(2)	(3)	(4)	(5)
				% of spending	
Income group	Accounts (K)	Spending ($M)	Baseline CS ($M)	Identical replacement	Random replacement
<50k	349.7	55.5	22.8	-5.3%	-2.7%
50-100k	356.6	61.4	22.7	-4.2%	-1.6%
100-150k	168.8	35.1	11.8	-3.2%	-0.8%
150-200k	63.7	14.2	4.7	-2.9%	-0.4%
>200k	69.2	20.1	5.6	-2.1%	-0.1%
Total	1,008.0	186.3	67.5	-4.0%	-1.5%

The table shows the change in consumer surplus from a chain ban as a share of restaurant spending under two alternative assumptions about the restaurants that would replace chains under a full ban. Column (4) assumes that large chains are replaced by independents with identical product characteristics. Column (5) assumes that large chains are replaced by independents with product characteristics that are randomly drawn with replacement from the set of characteristics chosen by independents in the city.