Welcome! We Have a New Menu: Measuring Product Responses to Competition
Preliminary, Do Not Cite

Jacob Cosman* Nathan Schiff†

December 2017

Abstract

This paper measures the response to new competition using a novel longitudinal dataset of restaurant menus in New York City. We use a technique from computer science to obtain a scalar measure of the pairwise distance between restaurants in product characteristic space based on differences in menu text. This particularly detailed measure of characteristic space allows for precise measurement of price and product changes in the menu of incumbent restaurants responding to entry. We address the endogeneity of location choice by matching "treated" incumbent restaurants facing competition from a new entrant with a "control" group of incumbent restaurants that have similar menus and location characteristics but no new competition. While we observe significant menu changes over our sample period, we do not find any evidence that restaurants facing competition respond by changing prices or product characteristics. We seek to provide some of the first evidence on the response to competition in markets with substantial product differentiation, many firms, and rapid turnover.

Notes for this draft (to delete)

1 Introduction

What is the response of an incumbent firm to competition from a nearby entrant? A vast and distinguished theoretical literature on spatial competition suggests that the response may depend upon transportation costs, the degree of product differentiation between the firms, consumers' geographical distribution as well as their preferences, future expectations, capacity constraints, and a host of other considerations. While the theoretical literature suggests a broad range of potentially salient factors, empirical measurement of this response can illuminate which of these are most important in a given market. However, most of the empirical literature has considered this question in the context of markets with a few firms selling similar products. In this paper we use a novel longitudinal dataset of restaurant menus in New York City, gathered from a large online food delivery service, to study the response to competition in a market with substantial product differentiation, many firms, and rapid turnover.

*Carey Business School, Johns Hopkins University

†School of Economics, Shanghai University of Finance and Economics

This is a worthwhile market structure to better understand. Not only are restaurants increasingly important as an amenity in urban centres, but also responses to new competition in this setting may be quite varied. When a new restaurant enters the market, a nearby incumbent restaurant may lower prices to retain customers. Alternatively, or additionally, they may change their menu — for example, by incorporating popular items from the entrant's menu or by modifying the menu to further differentiate their product from the entrant's. An incumbent might also make no changes whatsoever. A priori, the response is not clear and thus a natural topic for empirical work.

Studying the response to entry is difficult because firm location choice is endogenous. An entering restaurant may choose a specific site because of attractive location characteristics or because they have anticipated the response from incumbent restaurants. If the unobserved determinants of the location choice are correlated with factors affecting the measured outcomes of the incumbents then this introduces selection bias. For example, if new entrants tend to move into areas with rapidly increasing incomes and commercial rents, then incumbent restaurants may be raising prices independent of entry, thus biasing upwards estimates of the response to competition. A related issue is that different types of restaurants may change their menus with different frequencies, or respond differently to changes in city wide input prices; a labor shortage of sushi-chefs should not have the same effect on Japanese and Italian restaurants. While this may not cause bias unless these factors are correlated with entry, it does suggest that one must carefully control for the general heterogeneity of menu changes in order to more precisely measure the effects of competition. Lastly, we expect that the incumbent response to entry is a function of both characteristics of the incumbent and the entrant; the same Italian restaurant could respond differently to the entry of a new sushi restaurant versus a new Italian restaurant.

To address these issues we use a matching technique that exploits the unusual degree of product information in our dataset. We match "treated" incumbent restaurants facing competition from a new entrant with a "control" group of incumbent restaurants that have very similar menus and location characteristics but face no changes to the competitive environment. A central challenge in implementing this matching technique is how to determine the product similarity of two restaurants from the text of their menus. We employ a text processing technique from computer science to calculate a scalar measure of the similarity of two restaurant menus, "cosine similarity," and use this as a metric for distance in product space. We compare this measure with a set of known restaurant characteristics and find that it is a strong predictor of when two restaurants have the same product features. With this measure and additional location characteristics we compile a set of treatment and control observations and examine incumbent responses to entry in a number of channels and settings. Our measure also allows us to explore the determinants of incumbent responses by looking at how these responses vary as a function of the product similarity between entrant and incumbent.

Our preliminary results suggest that restaurants facing a new entrant do *not* change their prices, products, or service in any way different from the changes of restaurants without new competition. The remainder of the paper is organized as follows. First, we discuss the literature of spatial competition and product differentiation. Then, we discuss our data, definition of new competition, and show baseline results without matching. We then discuss the potential endogeneity in our estimation, how we construct a measure of product distance from our menu data, and the implementation of our matching strategy. Finally, we present preliminary results on restaurants' response to competition and conclude with the next steps for our project.

2 Literature Review

What does economic theory suggest should be the response of an incumbent restaurant to competition from a new entrant? Spatial competition models of oligopolistic competition with product differentiation argue that firms generally seek to differentiate, within the constraints imposed by demand, in order to mitigate direct price competition. In *The Theory of Industrial Organization*, Tirole (1988) refers to this idea as "the principle of differentiation[1]". When firms can differentiate in multiple dimensions—such as restaurants choosing geographic location and cuisine—they may choose to differentiate only in the dimension that is most important to the consumer. In an *n*-dimensional Hotelling setup with two firms, Irmen and Thisse (1998) show that firms will maximize differentiation in the most salient dimension while minimizing differentiation in all other dimensions.

From these results one might expect that new restaurants will be sufficiently differentiated from nearby incumbents so that there is little resulting price competition or that after entry incumbents will change menus to increase differentiation. However, the models discussed above describe oligopolistic competition and may not apply to a setting with many firms. Monopolistic competition models yield different predictions about price competition depending upon whether competition is spatial, as in circular city models (Salop 1979), or aspatial, as in Dixit-Stiglitz CES models (Dixit and Stiglitz 1977). Anderson and de Palma (Anderson and de Palma 2000) refer to this distinction as "local" versus "global" competition: are restaurants competing directly with their neighbors in geographic or product space, or do they simply compete indirectly for a share of a consumer's expenditure with all other restaurants in the market? Complicating the issue further is that many of these models assume a uniform distribution of demand and perfectly informed consumers; if demand is lumpy or consumers search then firms selling similar products may still have an incentive to geographically cluster. Ethnic neighborhoods are an example of lumpy demand, where concentration of a particular group leads to a concentration of restaurants of the corresponding cuisine. Konishi (2005) shows that consumer uncertainty can also lead to a concentration of competing firms[2]. In his model consumers have taste uncertainty and do not know the price of a store's product before visiting. When similar stores spatially concentrate, the probability a consumer finds her ideal product increases. Thus, consumers are willing to travel further to shop at the cluster, which increases overall demand. However, the firms in the cluster must now compete more strongly in prices than if they were located further from each other; this is in fact a further benefit to the consumer, which Konishi calls "the market size effect due to the lower price expectation." Under these conditions some new restaurants may locate close to similar incumbent

[1] See Chapter 7 for an overview of relevant models.

[2] Other notable examples of these models include Stahl (1982), Wolinsky (1983), and Dudey (1990).

restaurants, causing incumbents to lower prices. An example of this might be Manhattan's Little Italy, an area that now has few Italian residents, but where many competing Italian restaurants attract tourists in search of Italian food. Therefore the theoretical literature suggests that in different contexts, entry of a new restaurant may lead to increased incumbent differentiation, increased price competition, or no significant change if the entrant is already sufficiently differentiated from the incumbents.

Several empirical studies have investigated the response to competition in differentiated industries, including those where firms may compete in multiple dimensions. Netz and Taylor (2002) examine patterns of location for gasoline stations in Los Angeles and conclude that increased competition leads to increase spatial differentiation, defined as the geographic distance between stations. They also look at the relationship between spatial differentiation and characteristics differentiation, which they measure using attributes of the stations, such as gasoline brand or repair services available. They find a positive relationship between these two types of differentiation, contradicting the max-min hypothesis. Pinkse and Slade (2004) estimate cross-price elasticities of competing British beers and then use their estimates in a structural model to simulate the effects of mergers among brewers. They find that brands of the same beer type (lager, ale, or stout) have the strongest cross-price effects, with significant but weaker cross-price effects for brands with similar alcohol content (one of their measures of distance in product space). In our context, we might expect to find that incumbent menu responses are larger to entrants of the same cuisine. Their data consists of beer sales at the region level and so they do not also examine competition between firms in geographic space. Chisholm, McMillan and Norman (2010) investigate competition between thirteen first-run movie theaters in Boston. They find that theaters closer in geographic space are more distant in product space, as measured by film-programming choices over a one year period. Sweeting (2010) studies mergers between radio stations in the same listening format and geographic market to study the effect of common ownership on product differentiation. He finds that after two stations come under common ownership, the new owner increases differences between the music playlists of the two stations and repositions at least one of the stations closer to other competing stations. He also looks at whether the merger increases (implicit) listener prices, measured as commercials per hour, but finds no statistically significant result.

The markets we study and the data we use share some features with earlier studies but differ in several important ways. First, with the exception of Netz and Taylor (2002), the previous literature has investigated oligopoly markets with tens of firms while the market we study has thousands of firms. Second, most earlier work examines equilibrium outcomes with cross-sectional data or product changes in markets with little entry or exit. In contrast, our work is focused on dynamic responses to new competition in markets with substantial entry and exit, which helps us to more easily control for firm heterogeneity[3]. Third, restaurant menus not only provide extensive detail on product differentiation, as do radio playlists or movie screening schedules, but they also give itemized prices, allowing for a richer study of price competition across firm attributes.

[3]Sweeting (2010) also uses a panel to look at dynamics. However, both his focus on mergers, rather than entry, and the substantial differences between the radio industry and the restaurant industry (geography, number of firms, business model) make it difficult to extrapolate his results to our context.

3 Data and Descriptive Results

We collected data on New York City restaurants from a large online delivery service, which lists the menus of participating restaurants in a standardized text format. An important feature for our study is that customers order and pay for food from a restaurant directly through the website, which implies that the prices and items listed on the menu are current. We collected the data weekly from November 2016 through November 2017 (52 periods), allowing us to observe changes in a restaurant's menu. We observe restaurants joining the website and leaving the website, and therefore we have an unbalanced panel of menus from roughly 9400 unique restaurants. According to the NYC Department of Health website (accessed February 2017), there are about 24,000 active restaurants in NYC, implying about a third of the total restaurants are in our dataset. There is likely some selection on the characteristics of the restaurants in our data; for example, extremely expensive restaurants may not offer delivery. Nonetheless, we believe the size of this dataset is sufficient to allow us to make general statements about restaurant competition.

3.1 Descriptive Statistics

In our dataset there is a fair amount of noise. Given our empirical strategy, which relies on menu changes rather than levels, this noise seems unlikely to result in bias but certainly could obscure any competitive response. Therefore we use several criteria to determine outliers and drop unusual observations likely to result from characteristics of our data source. One feature of the restaurant delivery website is that the menus shown to the user can change depending upon whether the restaurant is open or not when the page is viewed. When a restaurant is closed, users have the option to "pre-order" but the items shown are only those that the restaurant always serves ("core items"). When the restaurant is open the menu may be longer and include daily specials and other items not part of the core set. Since we collect data at different times of day throughout our panel, we sometimes observe just the core menu while other times the full menu for that day. This can generate what looks like large period to period changes in the menu but instead simply reflects the time of day viewed. Roughly 25% of restaurants have a closed menu that differs from their open menu. To ensure that our measurement of the competitive response is not obscured by this time-of-day noise we drop all restaurants that change their menu dramatically according to time-of-day. We also observe some very high item prices (ex: $2000 items), which upon inspection reflect idiosyncratic situations, such as when a restaurant offers a catering menu for 100 people. These cases we also classify as outliers and exclude from the analysis. All of the tables and figures in the remainder of our paper exclude outliers, which total 28% of the sample.

In addition to these outliers, an issue in our data collection process resulted in missing prices for three consecutive weeks: 4/23, 4/30, and 5/7. Again, this should not bias our findings in any way but it does prevent us from using outcome variables in these three weeks to estimate our coefficients.

In Table 1 we show characteristics of the restaurants, averaged across restaurant-periods. The price statistics represent the statistic applied to a restaurant's menu and then averaged across all restaurant-periods.

For example, the variable "median item price" represents the median price across all items on a restaurant's menu in a single period; Table 1 then reports this statistic averaged across all the menus in our dataset. The variable "dollar signs" is a coarse price indicator the website shows to the user while the last five variables come from user reviews of different characteristics of the restaurant. In Table 2 we show the same set of statistics but now look at single period changes; for example, "median price change" reports the change in median item price over two consecutive periods. Table 2 then shows the average across all consecutive restaurant period changes. One can immediately see that most restaurants make very few period to period changes. The mean item count change is 0.05 items while the mean median item price change is just $0.01. Since these statistics are the average of both negative and positive changes we also show mean *absolute* changes for median item price and item count. These measures are significantly higher but still not very large. However, the standard deviation of these statistics is fairly large, indicating that some restaurants are making large changes. In Figure 1 we take a closer look at changes in item count and median item price. In panel a) the blue line shows that between 600 and 1100 restaurants change their item count every period (left axis) while the average *absolute* item count change for those restaurants varies between 4 and 10 items. In panel b) we show an analogous graph for median item price with between 250 and 500 restaurants changing their median item price every period with an absolute price change that varies between $0.4 and $0.7. In both of these figures one can see the periods of missing data as large dips in the series around May 1st. Lastly, in Table 3 we look at changes over time within a restaurant by running regressions of the form:

$$Y_{it} = \beta * weeks_{it} + \eta_i + \eta_t + \varepsilon_{it} \tag{1}$$

The η_i term is a restaurant fixed effect, the "weeks" variable measures the number of weeks since we first observed the restaurant, while η_t captures any factors affecting all restaurants in a given period; since restaurants enter our sample in different periods these last two terms are not collinear. We cluster standard errors by restaurant. From columns 1-4 we can see that restaurants slowly increase their prices at roughly $0.007 per week, with larger changes for the most expensive menu item and no detectable change for the cheapest item. Menus increase in length by 0.06 items per week but this estimate is not significant once we cluster the standard errors. On average, each week restaurants receive 6.7 new reviews. Therefore while generally restaurant menus are quite stable, there is still a fair amount of change, both across restaurants and within restaurants, with which we might measure competitive responses.

3.2 Determining Entry

Unfortunately, the appearance of a new restaurant menu on the delivery website does not imply that the restaurant has just entered the market. In order to determine entry we combine data from two additional sources: restaurant inspections from the City of New York and restaurant reviews from Yelp.com. According to the New York City government website, all restaurants in New York City must have a "Food Establishment

Table 1: Restaurant Characteristics, Levels

	mean	p50	sd	min	p1	p99	max
item count	127.62	104.00	89.7	6.0	13.0	400.0	500.0
median item price	8.46	7.99	3.2	2.5	3.0	18.0	25.0
mean item price	9.14	8.66	3.6	2.3	3.8	21.1	45.4
min itm prc	1.50	1.25	1.3	0.0	0.0	7.0	16.0
max itm prc	29.20	22.00	25.7	3.0	7.5	145.0	250.0
dollar signs	2.12	2.00	0.9	0.0	1.0	4.0	4.0
stars	3.65	4.00	1.3	1.0	1.0	5.0	5.0
food rating	85.63	88.00	9.2	8.0	50.0	99.0	100.0
delivery rating	86.45	90.00	10.9	0.0	47.0	100.0	100.0
order rating	90.01	93.00	8.6	0.0	58.0	100.0	100.0
Observations	298439						

Averaged across all restaurant-periods with no missing value.
Excludes missing item name or missing price periods.

Table 2: Restaurant Characteristics, Changes

	mean	p50	sd	min	p1	p99	max
item count change	-0.11	0.00	17.5	-480.0	-57.0	56.0	480.0
abs item count change	5.45	0.00	16.6	0.0	0.0	74.0	480.0
median price change	0.00	0.00	0.9	-19.5	-3.0	3.0	15.0
abs median price change	0.23	0.00	0.9	0.0	0.0	4.5	19.5
mean price change	0.01	0.00	0.8	-24.9	-2.3	2.4	25.6
min price change	0.00	0.00	0.3	-12.0	-0.3	0.3	15.2
max price change	0.04	0.00	5.1	-197.0	-8.0	9.0	231.0
dollars change	0.00	0.00	0.6	-4.0	-2.0	2.0	4.0
stars change	-0.02	0.00	0.8	-4.0	-3.5	3.0	4.0
food rating change	-0.01	0.00	0.9	-40.0	-3.0	2.0	39.0
delivery rating change	-0.01	0.00	0.9	-27.0	-3.0	2.0	36.0
order rating change	-0.01	0.00	0.8	-25.0	-2.0	2.0	29.0
Observations	271324						

Each variable is measured using consecutive period changes
Excludes periods with missing item names or missing prices.

Table 3: Within Restaurant Changes

	(1)	(2)	(3)	(4)	(5)	(6)
	mean itm prc	p50 itm prc	min itm prc	max itm prc	itm ct	revws ct
weeks observed	0.0068***	0.0053***	0.0001	0.0429***	0.0067	6.4790***
	(0.0008)	(0.0003)	(0.0002)	(0.0038)	(0.0070)	(0.1037)
Observations	333713	333713	333713	331773	333713	339681
R^2	0.000	0.014	0.001	0.007	0.014	0.338

Dependent variables: mean item price, median item price, minimum item price, maximum item price, item count, and reviews count.
All specifications include restaurant FE and period FE.
Standard errors clustered by restaurant,* p<0.1 ** p<0.05 *** p<0.01

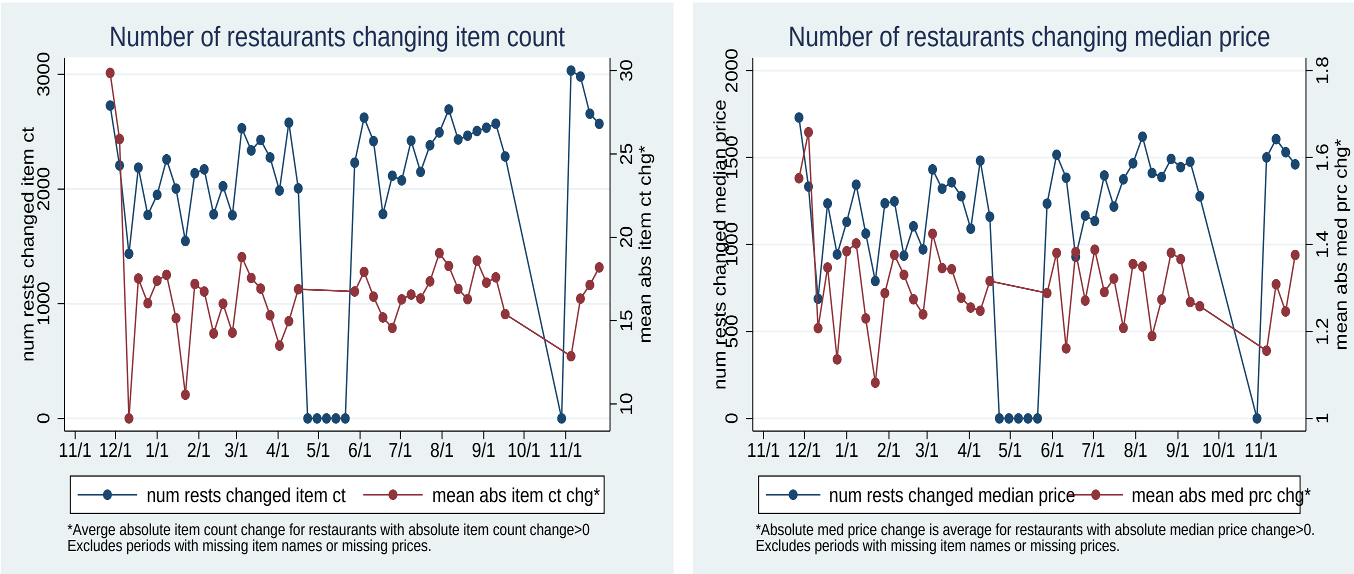

(a) Change in Item Count

(b) Change in Median Price

Figure 1: Menu Changes in NYC Restaurants

Permit" and a pre-permit inspection is required before the restaurant can open[4]. This suggests that pre-permit inspection dates would be a good measure of entry into the market. However, we found that while the inspection data goes back as far as August 2011, there are many restaurants whose first inspection date is in 2014 or later but do not have a pre-permit inspection, implying that there may be entrants without pre-permit inspections[5]. Further, for some restaurants whose first inspection occurs during our sample period, there are restaurant reviews on Yelp.com that far precede this initial inspection date. To ensure we have accurate dates for entry we use the following procedure. For each restaurant which first appears in the inspection data during our sample period, we find the date of the first Yelp review for the restaurant. If the first Yelp review is less than one hundred days before the first inspection or less than twenty days after the first inspection, we assume that this is a newly opened restaurant. These durations were chosen after looking at the histogram of review/inspection date differences. We define the entry date as the earlier of the first inspection date and the first Yelp review date. Figure 2 shows a histogram of entry over our sample period; note that entry necessarily declines towards the end of the sample due to our procedure for assigning entry dates, as well as lags in the inspection data.

3.3 Competition from Entrants and Treatment Assignment

In order to measure the response to competition we first need to clearly define which restaurants have been "treated" by a change in competition, which we measure as a new restaurant opening nearby. However, we do not know a priori the spatial range over which restaurants compete, nor the timescale with which they may

[4]See http://www1.nyc.gov/nyc-resources/service/2578/restaurant-permit.

[5]A call to the New York City Department of Health and Mental Hygiene, which oversees inspections, confirmed that while all restaurants should request an inspection before opening, this does not always happen.

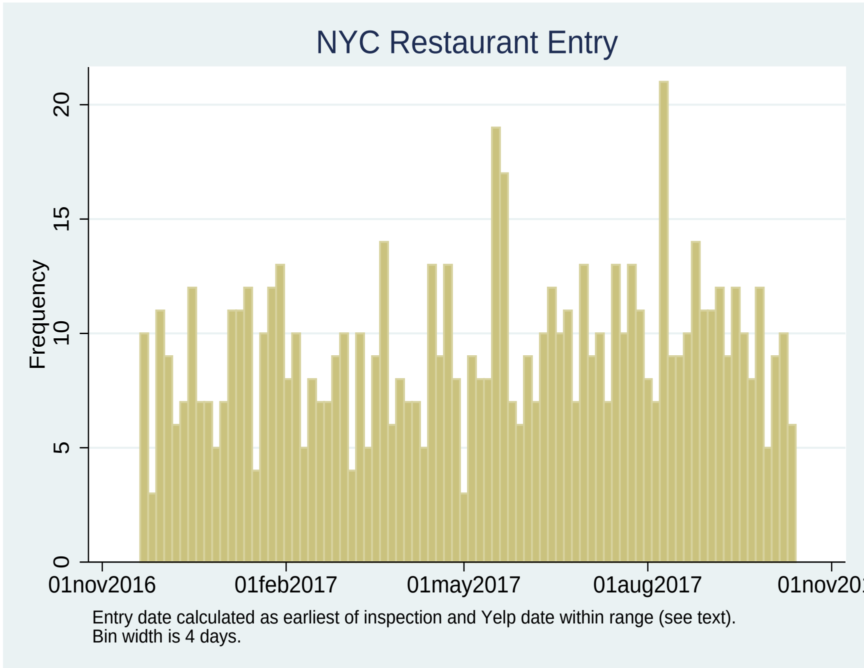

Figure 2: Entry from Inspections and Yelp Data

change their menus in response to the entrant. Further, for incumbent restaurants facing multiple entrants, it could be quite difficult to determine which entrant the incumbent is responding to. Therefore we choose to focus on cases where an entrant is most likely to represent a change in competition and where the response to a specific entrant can be isolated. Specifically, we specify a tuple (d, ρ_T, ρ_C) where d is a duration (measured in weeks), ρ_T is an inner radius, and ρ_C is an outer radius (i.e. $\rho_C > \rho_T$). For a given restaurant in time period t, a restaurant is deemed treated if and only if exactly one entrant within radius ρ_T first operates in period t and no other restaurants open from period $t-2d$ through period $t+2d$ anywhere within the larger radius ρ_C. A restaurant is deemed to be a control if and only if no restaurants open anywhere within radius ρ_C from period $t-2d$ through period $t+2d$. Our regression analysis uses a subset of this window, analyzing changes in a restaurant's menu from period $t-d$ to period $t+d$. Figure 3 provides a schematic of the timing of treatment and control definitions.

These definitions yield conservative samples of treatment and control restaurants. The separate radii ρ_T and ρ_C enforce a "buffer" between situations where the change in competitive environment from the nearby entrant is salient and situations where any new entry is too far away to have a substantial effect. Only including restaurants with exactly one entrant over $2d$ periods ensures that we are including restaurants which have experienced a comparable change in local competitive intensity. The long $t \pm 2d$ window serves a similar function to the distance buffer by helping us to exclude lagged effects and thus isolate effects only due to the observed new entrant. An important aspect of this definition is that treatment is determined by geography *and* timing. Over our entire sample period two incumbent restaurants may receive the same number of entrants within distance ρ_t but for a given duration $t \pm 2d$ one is treated and one is a control.

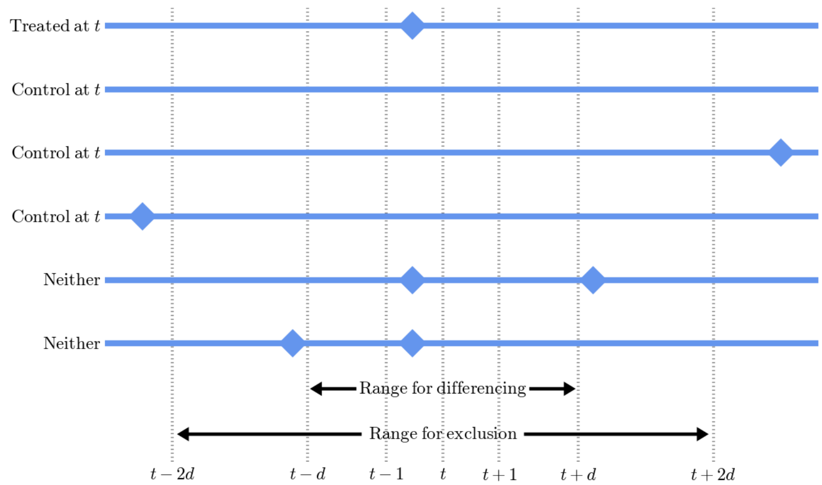

Figure 3: Schematic of the timing for treatment and control assignment.

We discuss this further when describing our identification strategy in section 4. The definitions do mean that some restaurants are assigned neither treatment nor control status and are therefore excluded from the rest of the analysis. Figure 4 provides a visual representation of the spatial aspects of treatment and control definitions.

Figure 5 shows an example of treatment and control for December 4, 2016 with an inner radius of 400 m, an outer radius of 500 m, and a duration of one week. As shown, the parsimonious specification of treatment and control excludes many restaurants for being near to several simultaneous openings. In practice it's infeasible to rigorously implement our matching estimator for a range of radii and so we will focus on an inner radius of $\rho_t = 500$m and $\rho_c = 600$m but test a range of durations d. These radii capture the spatial scale regarded as a reasonable walking distance in the urban planning literature. In the 1995 Nationwide Personal Transportation Survey the median length of a daily walking trip is a quarter mile (Boer, Zheng, Overton, Ridgeway and Cohen 2007). (Krizek 2003) describe this as "a scale sensitive to walking behavior". Our scale corresponds to approximately two long "avenue" blocks or six short "street" blocks in Manhattan (Pollak 2006).

3.4 Pre-matching Results

In this section we show OLS regression results using our definition of treatment but without any matching. We run the following regression for a number of different outcome variables:

$$\triangle_d Y_{it} = \beta * D_{it} + \eta_t + \varepsilon_{it} \tag{2}$$

The D_{it} term is the treatment indicator and the η_t term is a period fixed effect that captures anything which might have affected all restaurants equally in that week. In Table 4 we report results for $d = 4$. Across all outcome variables the treatment is insignificant. On the other hand, the constant is significant for most

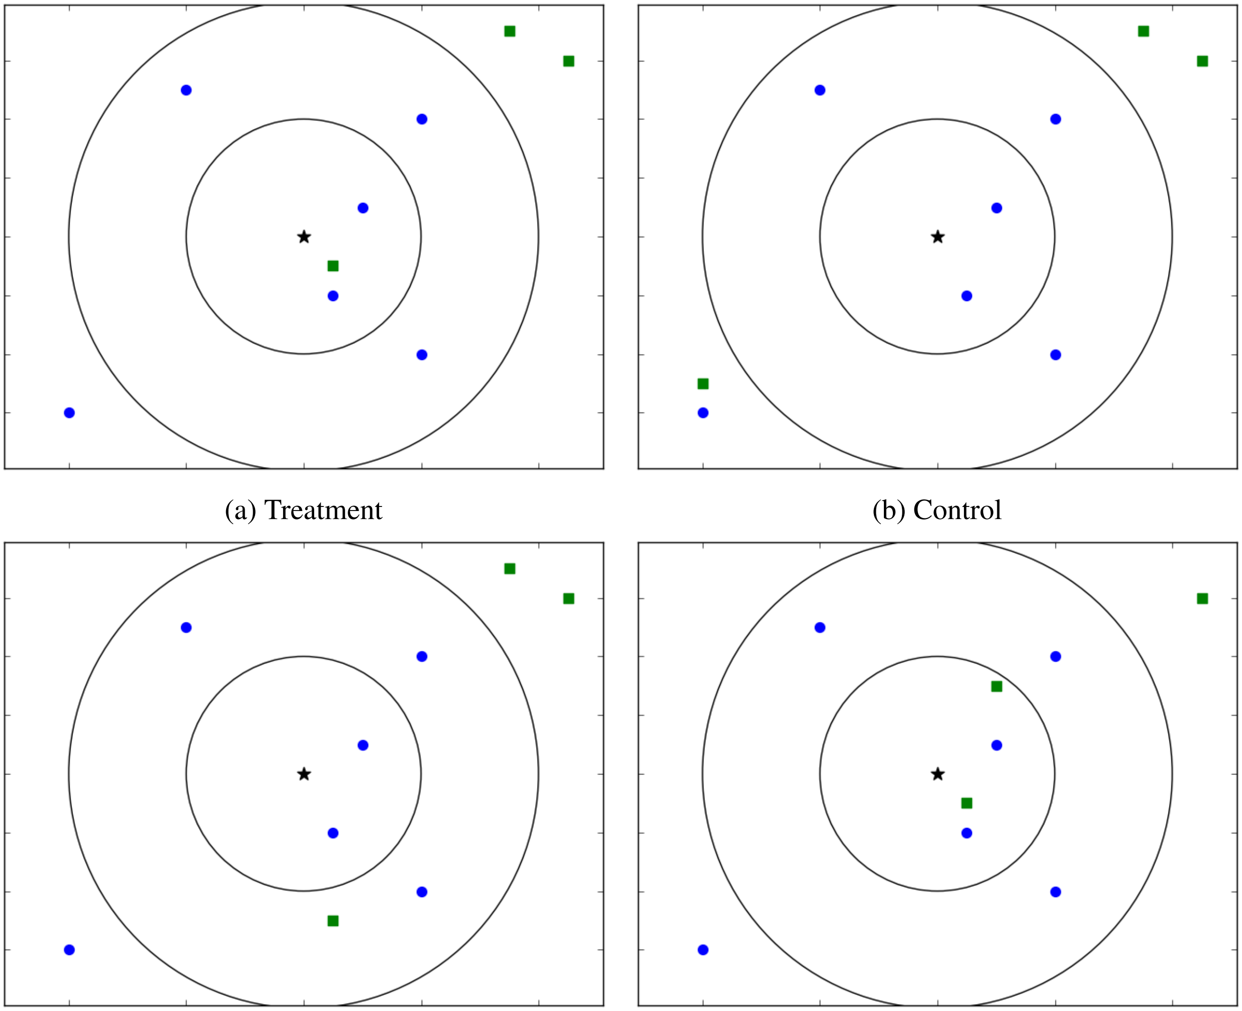

(a) Treatment

(b) Control

(c) Neither treatment nor control

(d) Neither treatment nor control

Figure 4: Examples of treatment and control assignment. The caption for each example indicates the assignment for the restaurant at the centre of the diagram (indicated by a star). Blue circles represent incumbent restaurants and green squares represent entrants. The two concentric circles represent the radii ρ_T and ρ_C.

outcome variables, implying that restaurants are changing over this eight week period, just not differentially between treatment and control. To investigate whether our choice of duration is important we also try running these same specifications for different durations. In panel a) of Figure 6 we show the coefficients and confidence intervals for regressions of differenced absolute median item on the treatment variable, for different durations. Again, nearly all of these coefficients are insignificant and there is no discernible pattern.

To better understand the time series pattern of the data we also run regressions comparing the level of each outcome variable on the treatment indicator for different periods. This allows us to see how treated and control restaurants compare, both before and after treatment.

$$Y_{i,t\pm d} = \beta * D_{it} + \eta_t + \varepsilon_{it} \tag{3}$$

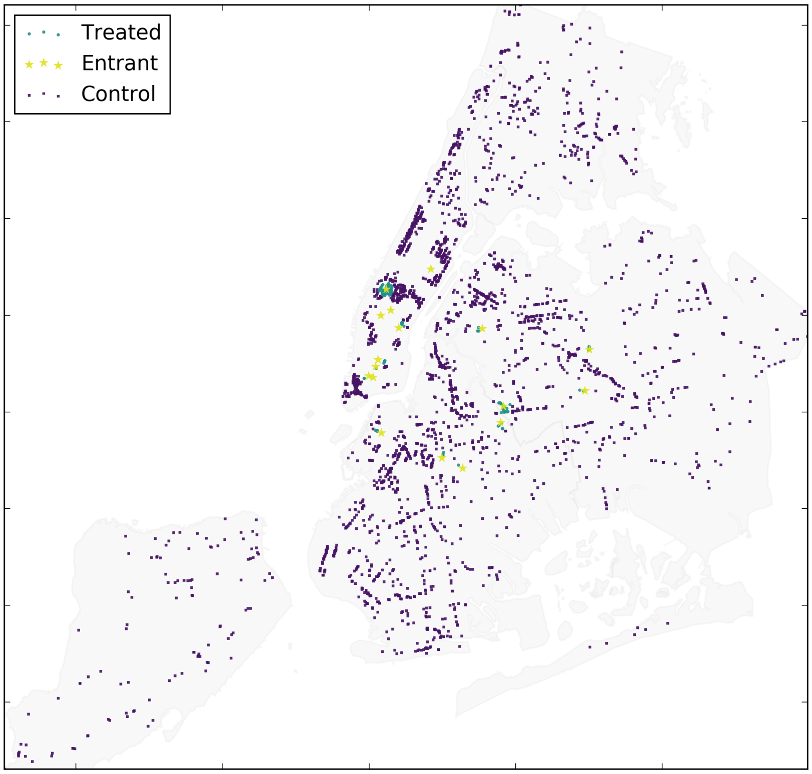

Figure 5: Treatment and control assignment for December 4, 2016 with an inner radius of 400 m, an outer radius of 500 m, and a duration of two weeks

We show the results for median item price in panel b) of Figure 6. Again, there is no discernible time series pattern and no large change after treatment. Further, there does not seem to be any difference between treated and control for median item price throughout the time series, but this turns out to depend on the variable we measure. In fact, it turns out that when measured in levels, treated and control restaurants are different in statistically significant ways for a number of variables. In Table 5 the left panel compares restaurant and menu characteristics for treated and control restaurants while the right panel compares demographic characteristics of the restaurants' locations. Generally treated restaurants have about 8 fewer items, 43 more reviews, and possibly slightly higher ratings, although the ratings differences are imprecise. Treated and control restaurants are also in quite different areas, with treated restaurants locating in younger, wealthier, and more highly educated neighborhoods while control restaurants are found in neighborhoods with a larger black percentage, greater percentage of households married and in families, and a larger percentage of the housing stock as single unit detached. Many of these demographic differences stem from the fact that treated restaurants are in areas with lots of entry and lots of restaurants: a large percentage are located in lower Manhattan. The variable "comp_500" counts the number of restaurants within 500m (for the first available period) and shows that a treated restaurant has about 17 more nearby restaurants than a control restaurant. Given these differences, and additional differences in the location characteristics of treated and control (next section), we now turn to describing our identification strategy.

Table 4: Pre-matching Results

	(1) itemct	(2) abs_itemct	(3) medprc	(4) abs_medprc	(5) prc95	(6) abs_prc95
trt_500_600_4	-0.2220	0.4004	-0.0005	-0.0158	-0.0707	-0.0600
	(0.5758)	(0.6148)	(0.0118)	(0.0112)	(0.0521)	(0.0525)
Constant	-0.2092	5.6164***	0.0493***	0.1355***	0.1841***	0.3388***
	(0.5530)	(0.5311)	(0.0140)	(0.0135)	(0.0506)	(0.0499)
Observations	40338	40338	40338	40338	40338	40338
R^2	0.060	0.052	0.001	0.007	0.002	0.002

Dependent var is (t+4) - (t-4) difference.
Treatment: inner radius=500m, outer=600m, duration=4.
All specifications include period FE.
Standard errors clustered by restaurant,* p<0.1 ** p<0.05 *** p<0.01

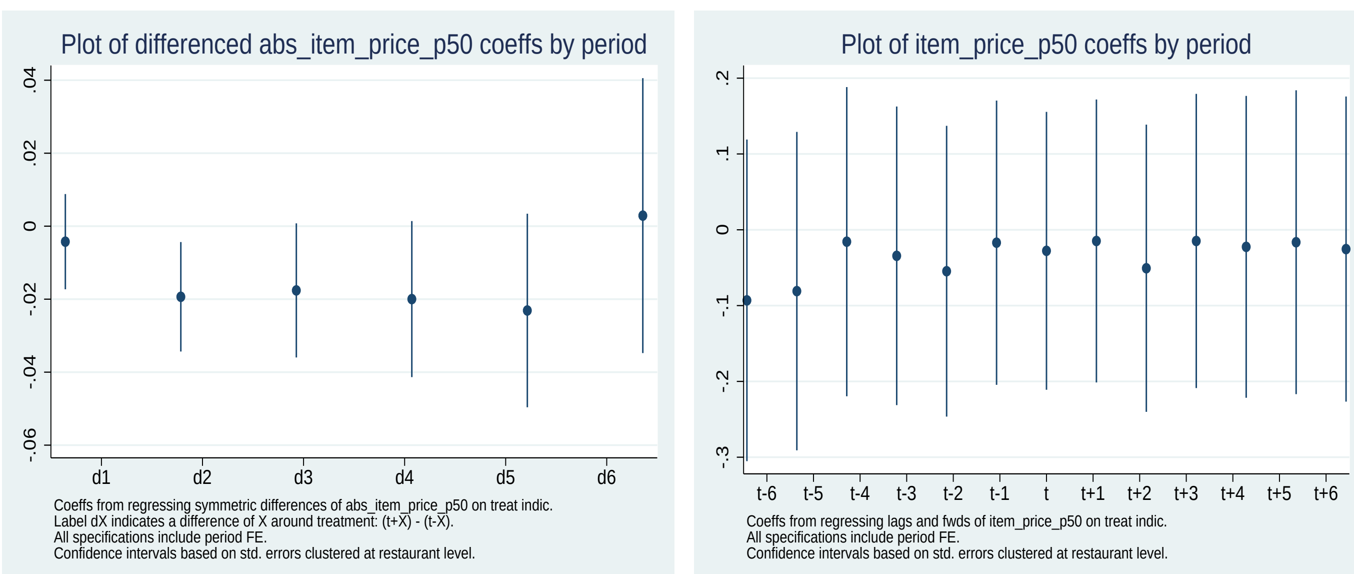

(a) Different Durations

(b) Period-by-Period

Figure 6: Median Price Changes Over Time

	Menu stats t-tests
item_count	-7.93**
item_price_mean	-0.08
item_price_p50	-0.01
item_price_p25	0.05
item_price_p95	-0.58*
stars	0.05
review_count	43.58***
order_rating	0.47
food_rating	0.73*
delivery_rating	0.58
Observations	62736

Tests difference between treated and control.
Calculated using values 4 periods before treatment.
Outliers are removed.

	Demographics t-tests
age2529	0.020***
age3039	0.027***
age7079	-0.003**
racewhite	0.062***
raceblack	-0.031***
hhfamily	-0.073***
hhmarried	-0.032***
educdegree	0.109***
poverty	-0.026***
income100150	0.008***
income150200	0.009***
unitdetached	-0.050***
comp_500	17.438***
N	71956

Tests difference between treated and control.
All demographics calculated as percent of area.
Comp variable is number of competitors within radius.
Outliers are removed.

Table 5: T-tests for Restaurant and Demographic Characteristics

4 Empirical Approach

In the previous section we showed that treated restaurants and control restaurants differ significantly in both location characteristics and menu characteristics, raising concerns that simple OLS comparisons may be biased. To address these concerns, we use a two-stage matching process that seeks to match treated restaurants with control restaurants which have both similar location characteristics and menu characteristics. In the first stage, we match each treated restaurant with a group of control restaurants that have a predicted entrant count (similar to propensity score) within a narrow band of the predicted entrant count of the treated restaurant ("calliper matching"). Then in the second stage we select a restaurant from this control group that has the closest menu to the treated restaurant. We base this two-stage approach on Rubin and Thomas (2000), who (in a different context) use a large set of covariates to get an initial propensity score and then match on a few highly-important covariates within narrow propensity score callipers.

To implement this approach we must choose a bandwidth (callipers) for the first stage but this choice involves a trade-off between balancing location characteristics and menu characteristics. Narrow bandwidths improve the similarity of location characteristics between the treated restaurant and its control group but also decrease the size of this control group. With a smaller control group it may be more difficult to find a control restaurant with a similar menu to the treated, thus leading to more imbalance in menu characteristics. Therefore we perform the two stage matching analysis for a range of bandwidths and then choose the bandwidth that leads to acceptable balance for both location characteristics and menu characteristics. It's important to emphasize we do not estimate any treatment effects during this process and that our choice of bandwidth is based entirely on balancing covariates and uninfluenced by outcome variables.

In this section we first discuss the general endogeneity issues and our matching approach with a reduced form model. We then describe our predicted entrant model and show its effect on location characteristics balance. Because we plan to run specifications with different treatment durations it's infeasible to show that our predicted entrant model improves balance for all treatment specifications. Instead, we show that conditioning on predicted entrants *generally* improves covariate balance for any locations where the observed entrant count differs. Next we explain our process for converting the text of menus into a scalar measure of the distance in product space between two menus. We then show how matching using this measure in the second stage of our process further improves menu characteristics balance.

4.1 Endogeneity and Identification

To illustrate potential endogeneity concerns and our identification strategy, we use the following reduced form selection model of restaurant outcomes and entry. Let Y_{it} be a restaurant level outcome (ex: median price, menu length, indicator for specific menu items) for incumbent restaurant i in neighborhood $j(i)$ at time t. Let D_{it} indicate whether at time t a new competitor (entrant) has entered within radius ρ_t of restaurant i. As discussed in section 3.3, we focus on incumbent restaurants that face at most one new competitor within a long time window and therefore we define $t = k$ as the single period when entry occurs within the window $k \pm 2d$.

$$Y_{it} = \beta_i \mathbb{I}\{t \geq k\} D_{it} + u_i + u_{j(i)} + \xi_{it} + \xi_{j(i)t} + \varepsilon_{it} \tag{4}$$

Our objective is to estimate β in equation 4 but we allow for a variety of restaurant and neighborhood level effects, both time-varying and invariant, to affect restaurant i's outcome. The time-invariant restaurant effect u_i could control for a restaurant's tendency to generally have high prices or a long menu in every period while the location effect $u_{j(i)}$ might capture the average income level or house price for a neighborhood over time. The ξ_{it} variable captures restaurant-specific time-varying shocks, such as the hiring of a new chef or a price increase in some ingredient specific to that restaurant. We assume these shocks are a function of restaurant characteristics but not location. Location specific shocks are captured by $\xi_{j(i)t}$ and might include gentrification in a neighborhood or new road construction that deters customers. Lastly, ε_{it} is an i.i.d. restaurant-time error term.

Following the potential outcomes framework, let Y_{it}^1 be the outcome of a restaurant at time t when there is entry (treatment) and Y_{it}^0 represent the outcome when there is not entry (control). These terms and the switching equation are then:

$$\begin{aligned} Y_{it}^0 &= \mu_i + \mu_{j(i)} + \xi_{it} + \xi_{j(i)t} + \varepsilon_{it} \\ Y_{it}^1 &= \beta_i \mathbb{I}\{t \geq k\} + Y_{it}^0 \\ Y_{it} &= D_{it} * Y_{it}^1 + (1 - D_{it}) * Y_{it}^0 \end{aligned} \tag{5}$$

We want to estimate the effect of new competition on incumbent restaurants, the average treatment effect on the treated (ATT), β:

$$ATT = E[Y_{it}^1 - Y_{it}^0 | D_{it} = 1] = E[\beta_i | D_{it} = 1] = \beta \tag{6}$$

As always, we don't observe what restaurants that faced new competition would have counterfactually done in the absence of this competition ($Y_{it}^0 | D_{it} = 1$). Further, it is highly likely that factors determining restaurant outcomes also affect entry. To model entry we assume that a new competitor enters near restaurant i at time $t = k$ if expected profit, modeled as a latent variable, is positive[6].

$$D_{it} = \mathbb{I}\left\{\theta_i + \theta_{j(i)} + \psi_{it} + \psi_{j(i)t} \geq 0\right\} \mathbb{I}\{t \geq k\} \tag{7}$$

As shown in equation 7, this entry process may also be a function of characteristics of incumbent restaurant i and location j, both time-varying and invariant. If any of these entry variables are also correlated with the restaurant outcome variables in equation 4 then the coefficient β estimated from a simple regression of Y_{it} on the treatment indicator D_{it} will be biased due to selection. For example, u_i and θ_i could be correlated if some type of restaurant, such as coffee shops, tend to always have low prices and attract additional entry. A greater concern is if unobserved changes to a neighborhood, such as gentrification or a neighborhood becoming "trendy", affect both existing restaurants ($\xi_{j(i)t}$) and entry probabilities ($\psi_{j(i)t}$). Relatedly, unobservable restaurant-level shocks could also change outcomes and spur entry. If incumbent restaurant i is struggling because their type of cuisine has suddenly become less popular then the restaurant may try to lower prices to attract consumers while, at the same time, a new entrant may locate nearby because they expect little competition from unpopular restaurant i.

We address these concerns with a difference-in-difference matching strategy (see (Heckman, Ichimura, Smith and Todd 1998) and (Smith and Todd 2005)). Essentially, we first difference the outcomes to remove the time-invariant effects and then use matching to try and control for the time-varying components that may cause selection bias. Given potential entry in period $t = k$, let the difference in an outcome d periods before entry and d periods after be defined as: $\triangle Y_{it} = Y_{i,k+d} - Y_{i,k-d}$. Then we can estimate β from this difference:

$$ATT = E[\triangle Y_{it}^1 - \triangle Y_{it}^0 | \triangle D_{it} = 1] = E[\beta_i | \triangle D_{it} = 1] = \beta \tag{8}$$

This differencing removes any correlation between the time-invariant terms in the outcome equation and the

[6]In equation 7 we are currently treating entry as a process independent of the entrant but we plan to later allow the characteristics of potential entrant l to affect this decision.

selection equation[7]. Entry and outcomes could still both be influenced by the time-varying terms, ξ and ψ, and therefore we now use matching to mitigate this form of selection bias.

We match treated restaurants with control restaurants using both characteristics of the incumbent restaurant's location j and the restaurant's menu text. We use a two-stage matching process where we first calculate the number of predicted entrants for each location using demographic variables, which gives us a subset of control restaurants with similar likelihoods of facing entry, and then choose the restaurant with the closest menu within this subset. We use the predicted entrants in essentially the same way as a propensity score, but for reasons described in the next section, it turns out that this count variable is better suited to our context than a simple binary propensity score. Let $P(X_{j(i)})$ denote the count of predicted entrants using the location characteristics and M_i stand for the menu-text. Then, our key identifying assumption is conditional mean independence (see Smith and Todd (2005)):

$$E[\triangle Y_{ik}^0 | P(X_{j(i)}), M_i, \triangle D_i = 1] = E[\triangle Y_{ik}^0 | P(X_{j(i)}), M_i, \triangle D_i = 0] \quad (9)$$

In our context, equation 9 implies that conditional on the predicted entrants and menu text, competition within this time period is essentially randomly assigned. This allows us to use the observed outcomes of restaurants that do *not* have new competition over a specific duration as a replacement for the counterfactual outcomes of restaurants that did. The idea behind this approach is that the matched control restaurant will be located in a similar neighborhood and sell similar food, and therefore will be subject to location and restaurant-level shocks that are similar to those of the treated restaurant. For example, city-wide trends in taste, such as a fad for cupcakes or kale, may have a similar effect on the demand for restaurants selling these foods, which is captured in their menu text. On the supply-side, increases in the cost of an input specific to certain types of restaurants–sushi grade tuna or the wage of sushi chefs–will impact restaurants with that cuisine on the menu. We can make an analogous argument for location: if neighborhood trends are correlated with underlying demographics then by matching on these demographics we choose control observations that experience the same trends. An example might be if neighborhoods with relatively low rent but well educated residents become hip neighborhoods with lots of new restaurants and changes in incumbent restaurants.

Lastly, when we select a control restaurant using menu-text we are essentially using an outcome variable in the pre-treatment period to improve the match. A recent paper by Chabe-Ferret (2014) argues that matching with pre-treatment outcomes when selection is due to both a fixed effect and transitory shocks can lead to improperly matched observations or misalignment. The author suggests instead matching on covariates that do not vary over time. For this reason we use the earliest period menu for each restaurant, which we believe will capture the general cuisine of the restaurant but is far enough (often months) from the new competitor entry date that the menu is unlikely to include pre-treatment trends.

[7] In equation 8 note that $\triangle Y_{i\tau}^1 = Y_{i,k+\tau}^1 - Y_{i,k-\tau}^1 = \beta_i + Y_{i,k+\tau}^0 - Y_{i,k-\tau}^0 = \beta_i + \triangle Y_{it}^0$

4.2 Matching with Predicted Entrants

As noted earlier, treatment assignment depends on timing and thus a given restaurant may be treated, control, or neither, for different time periods. For this reason time-invariant characteristics of a location cannot accurately predict treatment assignment and thus we do not use a propensity score for matching. However, as we show in this section, some locations have much more entry than others over our sample period and the total number of entrants is correlated with time-invariant location characteristics. Although treatment timing cannot be predicted by location characteristics, we still want to make sure that we are comparing treated restaurants to control restaurants in similar areas. Therefore, we model the total number of entrants over our entire sample period in each location using a poisson model and then use the predicted number of entrants to balance the location covariates. Our use of predicted entrants for matching is very similar to a simple propensity score and we are not estimating multi-valued treatment effects (ex: Hirano and Imbens (2004) or Flores, Flores-Lagunes, Gonzalez and Neumann (2012)). However, we do borrow from this literature when assessing covariate balance.

For each incumbent restaurant ever observed in our sample, we count the number of total entrants n_i over the sample period within $\rho_t = 500$m. Note that this count of entrants is really a characteristic of the location and does not depend on how many periods we observed restaurant i or when it entered our sample. We then model the count of entrants as following a poisson process where the expected count depends on the characteristics of the area j around restaurant i, $X_{j(i)}$.

$$ln(E[n_i|X_{j(i)}]) = X'_{j(i)}\theta \tag{10}$$

As candidates for $X_{j(i)}$, we assembled a large number of census tract variables from the 2009-2014 American Community Survey, "fair market rent" at the zipcode level from the department of Housing and Urban Development (HUD), and the distance to the nearest subway station. We also included the count of competitor restaurants within several different radii, calculated with the first period of data to ensure this measure wasn't correlated with our dependent variables. We then use a penalized poisson model (LASSO) to select the variables and estimate the coefficients, using cross-validation to choose the penalty parameter. We describe the details of this process and show the coefficients estimates in appendix section 7.1. Generally we find that higher incomes, higher rents, lower marriage rates, smaller families, closer subways, and more nearby restaurants are all associated with higher entry counts.

For each restaurant i we can now calculate the number of predicted entrants $P(X_{j(i)})$ using our model. To form a control group for each treated restaurants we will choose a subset of all control restaurants that have a predicted entrant count within a narrow bandwidth ("callipers") of the treated restaurant. After performing an optimization routine, discussed in appendix section 7.2, we find the optimal bandwidth to be 0.2 standard deviations of the log of predicted entrants.

As noted in the introduction to this section, we cannot show balance tables for all of the different treatment durations we plan to examine. Instead, we follow the general procedure of Hirano and Imbens (2004) by grouping restaurants into quintiles of observed entrants and then comparing the covariates for observations in a specific quintile to observations not in that quintile. For example, the first quintile of observed entrant count consists of locations that never had any entrants. We wish to compare the average value of each location covariate for locations with zero entrants (quintile 1) to locations with more than zero entrants (quintiles 2-5). As recommended in Imbens (2015), we compare covariates using the standardized normal mean, which for quintile q and covariate x is $(mean(x)_q - mean(x)_{\neq q})/sd(x)_{all}$. To calculate the post-adjustment balance we need a way to show the standardized normal mean for locations in a quintile compared to the *set* of matched locations within the callipers (0.2 standard deviations of log predicted entrants). To do this we use a montecarlo method where in each simulation run we randomly select one observation within the callipers for each treated and then calculate the standardized mean by observed entrant quintiles. We take the absolute value of this measure and then average it over all simulation runs (100) to get a measure of the average (absolute) difference between observations in the quintile and outside of the quintile. In Table 6 we show the standardized normal difference after adjustment for each covariate, for five quintiles of observed entrants. For nearly all covariates the standardized normal difference is below 0.2, suggested as a reasonable threshold in Imbens (2015). Especially important are the counts of nearby competitors, the first four rows of Table 6, which in Table 5 differed quite strongly across treated and control. As Table 6 shows, using predicted entrants allows us to match restaurants in areas with similar restaurant density. That said, there are still some covariate that are not very well matched, with some age brackets (18-24, 65-69) and housing characteristics (built year) still differing across the quintiles.

4.3 Measuring Menu Distance

The second stage of our matching process requires matching restaurants with similar menus. A challenge for doing this is that our menu data is literally the text of a restaurant menu, with no additional structure, classification, or standardization. Each restaurant usually divides their menu into item sections (e.g. "Vegetables" or "Noodles") and then lists each item in the section with a price. Restaurants may also include an item description, like "Thin noodles. Spicy." For the purposes of economic research, it would be ideal if restaurants classified every one of their dishes into standardized item codes so that menus could be easily compared. Without such a system restaurant comparison becomes quite challenging, but any attempts to create our own item standardization would require a myriad of arbitrary decisions, such as whether a meatball hero sandwich is the same as a meatball submarine sandwich. Instead, we follow the text processing literature in computer science to calculate a measure of the similarity between the overall text of two restaurant menus. Specifically, we use the method in Damashek (1995), which breaks the text of a document into a set of strings of consecutive characters, called "ngrams," and then compares two documents based on the counts of their component ngrams.

An ngram of size n is a text string of n consecutive characters. The phrase "with fries" has seven 4-grams including the space between words: "with", "ith_", "th_f", "h_fr", "_fri", "frie", and "ries". We can decompose the text of any restaurant menu into ngrams of a given size and then count the number of

Table 6: Covariate Balance Table, Post-Adjustment

Variable	Q1	Q2	Q3	Q4	Q5
within_0.025_km	0.07	0.02	0.03	0.06	0.16
within_0.25_km	0.01	0.01	0.00	0.12	0.07
within_0.5_km	0.01	0.05	0.00	0.02	0.01
within_1.0_km	0.03	0.11	0.07	0.04	0.02
rent_4_room_0_year_lag	0.13	0.11	0.04	0.06	0.15
rent_5_room_0_year_lag	0.1	0.11	0.04	0.09	0.12
age.lt.10	0.01	0.03	0.08	0.04	0.3
age.10.17	0.09	0.16	0.29	0.1	0.25
age.18.24	0.22	0.01	0.27	0.32	0.01
age.25.29	0.06	0.15	0.08	0.16	0.08
age.30.39	0.07	0.17	0.11	0.23	0.02
age.40.49	0.23	0.14	0.02	0.13	0.09
age.60.64	0.15	0.09	0.06	0.19	0.02
age.65.69	0.04	0.29	0.3	0.43	0.07
commute.county	0.07	0.05	0.05	0.05	0.15
hh.family	0.02	0.27	0.01	0.03	0.31
school.college	0.12	0.04	0.3	0.06	0.06
school.none	0.03	0.03	0.41	0.1	0.2
ap.lim	0.08	0.05	0.02	0.06	0.05
ap.eng	0.18	0.14	0.06	0.24	0.03
poverty	0.1	0.26	0.06	0.01	0.06
owner.occ	0.02	0.1	0.06	0.14	0.20
unit.detached	0.09	0.11	0.11	0.15	0.03
unit.3.9	0.12	0.07	0.23	0.2	0.15
unit.10.49	0.02	0.17	0.00	0.24	0.03
unit.gt.50	0.03	0.19	0.04	0.22	0.26
built.2000.2009	0.12	0.36	0.15	0.20	0.03
built.1990.1999	0.03	0.18	0.13	0.30	0.08
rent.1250.1500	0.04	0.21	0.07	0.20	0.16
rent.2000.plus	0.01	0.09	0.12	0.23	0.24
rent.income.40.50	0.08	0.07	0.09	0.15	0.07
value.500.750	0.08	0.12	0.06	0.03	0.13
value.750.1m	0.01	0.05	0.06	0.09	0.26
subway_dist	0.18	0.06	0.09	0.25	0.12

occurrences of every specific ngram. For example, if we looked at the 3-gram decomposition of a barbecue restaurant menu there might be a large number of "bar" or "bbq" 3-grams. Dividing the count of any specific ngram by the total count of ngrams in the menu gives us the proportion of the menu represented by that particular ngram. Then a menu with j unique ngrams can be represented as a j-dimensional vector of the ngrams' weights, with each weight representing the relative frequency of the ngram. Once two restaurant menus have been converted into vectors in ngram space, we can then measure the difference between their menus as the angle between their ngram vectors.

Following Damashek's notation, we can compare menu a to menu b by comparing their ngram weights on the set of J ngrams, where J is the superset of ngrams from both menus for some pre-chosen ngram size (ex: 3). If a menu has count m_i occurrences of ngram i then the weight x_i of this ngram is:

$$x_i = \frac{m_i}{\sum_{j=1}^{J} m_j} \tag{11}$$

When a menu has no occurences of ngram i that ngram receives zero weight. Damashek defines the "cosine similarity" S_{ab} of two documents (menus) a and b as the cosine of the angle between their ngram vectors:

$$S_{ab} = \frac{\sum_{j=1}^{J} x_{aj} x_{bj}}{\left(\sum_{j=1}^{J} x_{aj}^2 \sum_{j=1}^{J} x_{bj}^2 \right)^{1/2}} = \cos(\theta) \tag{12}$$

In Damashek (1995) the author uses his method to assign documents to languages (ex: "French") and topic areas for news articles in a given language (ex: "mining"). He finds that equation 12 does a good job for language assignment but has worse performance for topic assignment. He suggests that this is because the ngram vectors of two articles written in the same language will have a great deal of similarity simply due to common and uninformative ngrams in the language or general group to which the documents belong. For example, in English the 3-gram "the" is common but uninformative about topic. To deal with this issue he suggests centering all ngram vectors by subtracting a common vector that captures the ngram distribution of some specific language or subject group. Letting μ represent this common vector of weights the "centered cosine similarity" is:

$$S_{ab}^{c} = \frac{\sum_{j=1}^{J} (x_{aj} - \mu_j)(x_{bj} - \mu_j)}{\left(\sum_{j=1}^{J} (x_{aj} - \mu_j)^2 \sum_{j=1}^{J} (x_{bj} - \mu_j)^2 \right)^{1/2}} = \cos(\theta) \tag{13}$$

This is essentially a correlation coefficient, with interpretation depending on the definition of μ, and thus will vary from -1 to 1. In our context, we wish to subtract out the common distribution of restaurant menu ngrams (for example, ngrams of the word "chicken" appear to be especially frequent) and so we define the vector μ as simply the vector of ngram centroids across all menus M. If we weight each menu equally then the centroid for ngram j is:

$$\mu_j = \frac{1}{M}\sum_{m=1}^{M} x_{mj} \tag{14}$$

Lastly, in order to make our product space metric consistent with geographic distance we convert cosine similarity to cosine distance, with the new measure ranging from 0, no distance between products, to 2, maximum distance between products:

$$cosDist_{ab} = 1 - \frac{\sum_{j=1}^{J}(x_{aj}-\mu_j)(x_{bj}-\mu_j)}{\left(\sum_{j=1}^{J}(x_{aj}-\mu_j)^2 \sum_{j=1}^{J}(x_{bj}-\mu_j)^2\right)^{1/2}} = 1-\cos(\theta) \tag{15}$$

Several previous papers have used uncentered cosine distance as a measure of differentiation in product space. Jaffe (1986) defines the technological position of a firm as a vector of the distribution of its patents over 49 classes and then uses the angle between two of these vectors to measure changes in technological position. Similarly, Sweeting (2010) measures differentiation between radio stations as the angle between vectors of airplay for music artists and Chisholm et al. (2010) measure differentation between first-run theaters as the angle between vectors of movie screenings. If restaurants were in the business of selling ngrams then our measure would be analogous to these other papers. However, it is not clear that ngrams are a good representation of products, nor that the angle beween two restaurants' centered vectors of potentially thousands of ngrams provides any information about the similarity of their products. Therefore in the next section we discuss implementation and testing of this new method.

4.3.1 Text distance implementation and testing

We calculate the cosine distance metric using the names and descriptions of every item on a restaurant's menu. As described later in our section on identification, we want to capture a pre-treatment measure of the menu distance between two restaurants. Therefore we use the first observed menu for every restaurant; for the majority of restaurants this is the first period of our data but varies for later entrants. In order to center each menu's vector we need to choose a "corpus" or distribution of ngrams used for calculating the μ vector. We choose the set of menus from the first period of data as the corpus. This is a very large set of ngrams and choosing different periods or combining periods is unlikely to have any qualitative effect on our measure[8]. With this corpus and the initial menu of every restaurant we calculate one matrix of pairwise distances between all restaurants in our sample.

[8]There are a few ngrams that show up in the menus of later entrants that are not found in this initial corpus. We assign these ngrams a μ weight of zero.

It is worth emphasizing that this cosine distance metric appears to capture salient aspects of menu differentiation. The online delivery service data includes one or more cuisines for each restaurant in the sample. As shown in Figure 7, the distribution of pairwise cosine distances between restaurants with identical cuisine sets first-order stochastic dominates the distribution of restaurants that share at least one, but not all, cuisines. Moreover, the distribution of pairwise cosine distance between restaurants that share at least one cuisine first-order stochastic dominates the distribution of pairwise cosine distance between restaurants that share no cuisines. Pairs of restaurants with a small cosine distance are particularly likely to share all cuisine categories.

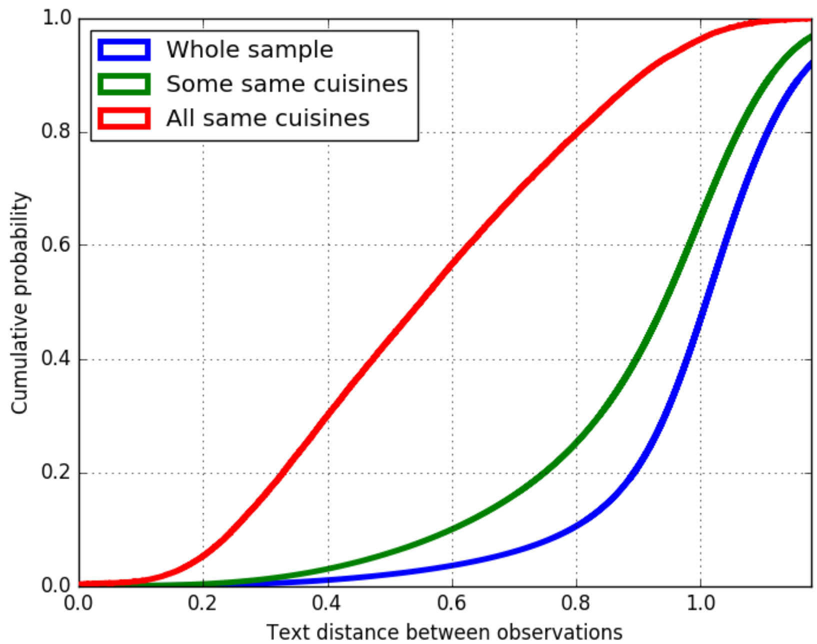

Figure 7: Cumulative distribution function of cosine distance between pairs of restaurants that share all cuisines (in red), at least one cuisine (in green), and all cuisines (in blue).

As well, the cosine metric can provide additional information beyond the cuisine categories provided by the online delivery service. Many of the cuisine categories are very broad. For example, Bella Pizza (which serves items including “10 Piece Chicken Buffalo Chicken Wings”), Genuine (which serves items including “Fries with Turkey Chili and Queso Fresco”), and WINE 34 (which serves items including “Acorn Squash Ravioli”) are all listed with “American” as their sole cuisine category. In fact, as “American” is such a broad cuisine category, two restaurants with “American” as their sole listed cuisine may not have particularly similar menus. Figure 8a shows the distribution of cosine distances between pairs of restaurants with “American” as their sole listed cuisine. As shown, many of these pairs of restaurants have high cosine distances between their menus.

However, for restaurants with more narrowly-defined cuisines, the cosine metric provides a closer match. Figure 8b the distribution of cosine distances between pairs of restaurants that have exactly the two cuisine listings “Asian” and “Chinese”. Most of these cosine distances are much smaller than the cosine distances

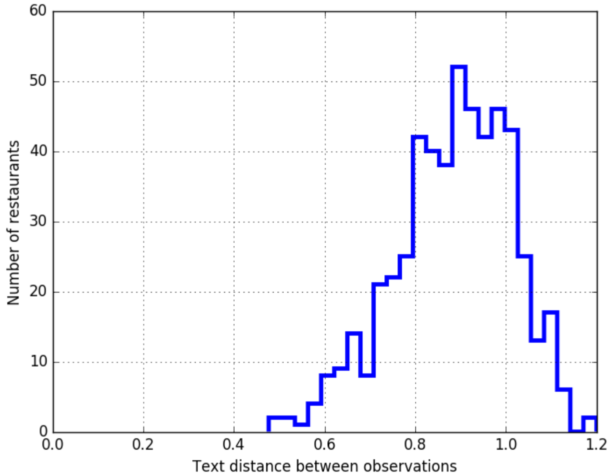

(a) Pairwise cosine distance between restaurants with the cuisine "American".

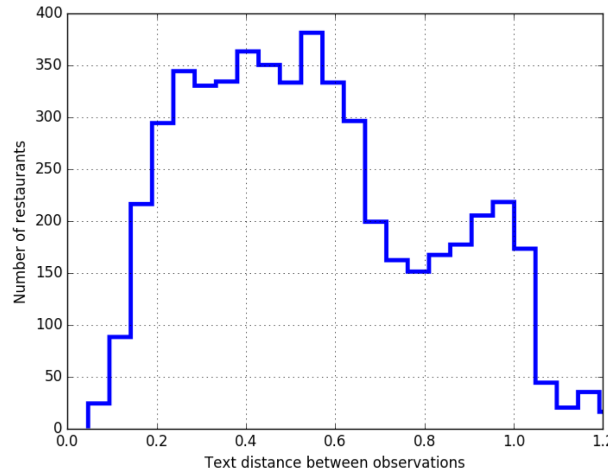

(b) Pairwise cosine distance between restaurants with the cuisines "Asian" and "Chinese".

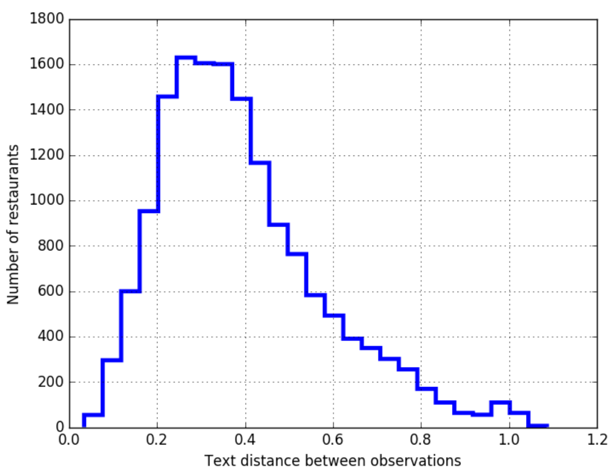

(c) Pairwise cosine distance between restaurants with the cuisines "Asian", "Chinese", and "Lunch Specials".

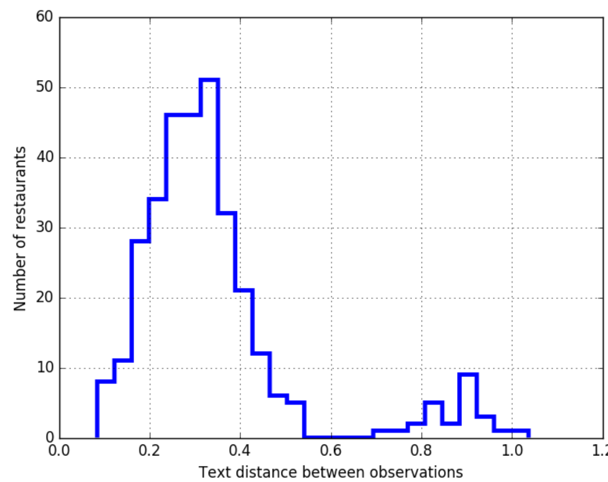

(d) Pairwise cosine distance between restaurants with the cuisines "Asian", "Cantonese", "Chinese", and "Lunch Specials".

Figure 8: Histograms of cosine distances between pairs of restaurants with the same set of cuisines.

between restaurants listed as "American". As shown in Figure 8c, narrowing the set further to pairs of restaurants that have exactly the three cuisine listings "Asian", "Chinese", and "Lunch Specials" gives cosine distances that are smaller yet; restaurants that focus on lunch specials are on average more similar than restaurants that do not. Considering the set of restaurants that have exactly the four cuisine listings "Asian", "Cantonese", "Chinese", and "Lunch Specials" gives an even spread of pairwise cosine distances as shown in Figure 8d.

The text cosine distance was developed to identify the language and topic of documents. In the context of our sample it appears to consistently identify menu differentiation. Not only are restaurants with the same cuisine categories closer in cosine distance than menus with different cuisine categories but also it provides a continuous measure of differentiation that appears to capture combinations of cuisines and within-cuisine differentiation as well.

4.4 Matching with Menu Distance

We now use the menu distance measure just described to select control restaurants with the closest menu to the treated. Simply, for each treated restaurant we find the control restaurant within the predicted entrant callipers that has the smallest menu distance to the treated. To show this effect on the balance of restaurant and menu characteristics we show the standard normalized differences after adjustment. In Table 5 we saw that there were large differences in item counts between treated and control restaurants. Table 7 shows that across different observed quintiles the differences in item counts are fairly small, with the exception of quintile 4. One can also see that the balance is less even for quintile 5, locations with the highest number of observed entrants, compared to the other quintiles. Generally we will not be using many observations in quintile 5, which have up to 46 observed entrants, because the high frequency of entry mostly disqualifies these locations from satisfying our treated or control duration requirement.

Table 7: Menu and Restaurant Characteristics Balance, Post-Adjustment

Variable	Q1	Q2	Q3	Q4	Q5
p_median	0.18	0.12	0.32	0.22	0.09
p_95	0.03	0.02	0.22	0.23	0.26
item_count	0.01	0.13	0.3	0.38	0.13
name_char	0.15	0.21	0.33	0.05	0.24
desc_char	0.02	0.1	0.03	0.17	0.17
total_chars	0.01	0.15	0.1	0.21	0.13
p_median_abs	0.17	0.07	0.04	0.14	0.25
item_count_abs	0.07	0.1	0.11	0.09	0.09
total_chars_abs	0.04	0.11	0.16	0.17	0.22
p_int	0.33	0.15	0.03	0.15	0.14
p_ext	0.01	0.13	0.12	0.05	0.18
food_good	0.04	0.02	0.12	0.00	0.28
on_time	0.2	0.08	0.02	0.05	0.27
order_accurate	0.1	0.24	0.05	0.1	0.26

5 Results and discussion

We now present a series of results using our matching method for the specification with $\rho_T = 500$ metres, $\rho_C = 600$ metres, and $d = 4$ weeks. Our regressions use the same specification from the "pre-matching results" section, regressing the difference in an outcome variable on a treatment indicator. Across all outcome variables the treatment variable is statistically insignificant and very close to zero. To emphasize the bounds on these estimates we show 95% confidence intervals, using clustered standard errors. Interestingly, the constant, which represents the change for all restaurants over this duration, is consistently large and significant. In fact, by comparing the estimated treatment effect to the constant we can see that in many of these specifications the treatment effect is so close to zero that the upper bounds of the confidence interval on these estimates is less than the constant. In other words, the maximum difference between treated and control is still smaller than the simple expected change over this time period. We interpret this magnitude as a tightly estimated zero; all restaurants are changing prices but there is absolutely no difference between treated and control in how they change prices.

Table 8 shows the differential response in menu item prices for treated and non-treated restaurants. As shown, all changes are both statistically and economically insignificant. The coefficient on the treatment indicator for the change in median item price is a particularly tightly estimated zero. Note that this is *not* driven by an inability of restaurants to change their prices on the platform; the positive and significant constant terms indicate that on average restaurants are adjusting their prices upwards by several cents per month. Rather, the change in prices appears to be unaffected by a nearby entrant.

Table 8: Regression results for item prices. Coefficients shown with 95% confidence intervals.

	Median price	95th price	Int marg price	Ext marg price
Treated	0.000	−0.011	0.072	4.691
	(−0.062, 0.063)	(−0.187, 0.165)	(−0.036, 0.179)	(−4.327, 13.710)
Constant	0.142***	0.154***	0.000	0.043*
	(0.067, 0.216)	(0.042, 0.265)	(−0.000, 0.000)	(−0.001, 0.087)
Observations	4,032	4,032	1,913	1,913
Adjusted R^2	−0.000	−0.000	0.001	0.001

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 9 shows the changes to restaurant menu text. Again, changes to the menus in response to treatment are small and statistically insignificant while the overall average change (i.e. the constant terms) are larger and more economically meaningful. Neither the number of items on the menu nor the length of the item names or descriptions appears to be affected by a nearby entrant.

Table 10 shows results for the three quality indicators from restaurant reviews (the quality, punctuality, and accuracy of food orders) as well as the total hours per week the restaurant is open. The first three variables are normalized to $[0, 100]$. As shown, with the exception of the punctuality measure, all coefficients on the treatment indicator are small and not statistically significant. The sole treatment coefficient with statistical significance is the punctuality indicator. However, insofar as the variation here is very small on the scale of the range of observed values and insofar as it is the only statistically significant coefficient among a large number of outcome variables we are hesitant to impute meaning to this coefficient.

Table 11 shows changes in the absolute magnitude of the median price, the item count, and the total number of characters in the menu. As shown, these changes are also generally small in magnitude relative

Table 9: Regression results for menu text characteristics. Coefficients shown with 95% confidence intervals.

	Item count	Name length	Desc length	Total chars
Treated	0.974	−0.023	0.030	40.469
	(−0.753, 2.701)	(−0.223, 0.176)	(−0.361, 0.421)	(−76.010, 156.949)
Constant	−5.917***	−1.728**	1.530***	−266.261**
	(−9.807, −2.026)	(−3.045, −0.411)	(0.808, 2.252)	(−530.147, −2.375)
Observations	4,032	4,032	4,032	4,032
Adjusted R^2	−0.000	−0.000	−0.000	−0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 10: Regression results for restaurant services. Coefficients shown with 95% confidence intervals.

	Quality	Punctuality	Accuracy	Hours
Treated	0.104	0.143**	0.015	−0.621
	(−0.030, 0.238)	(0.027, 0.260)	(−0.097, 0.128)	(−1.767, 0.526)
Constant	−0.152***	−0.115**	−0.072**	0.522
	(−0.246, −0.058)	(−0.205, −0.025)	(−0.143, −0.001)	(−0.411, 1.454)
Observations	3,933	3,933	3,933	1,800
Adjusted R^2	0.000	0.001	−0.000	0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

to the constant coefficients as well.

Table 11: Regression results for absolute menu changes. Coefficients shown with 95% confidence intervals.

	Abs median price	Abs item count	Abs total chars
Treated	0.075**	1.256*	40.581
	(0.001, 0.149)	(−0.206, 2.718)	(−54.844, 136.007)
Constant	0.424***	14.560***	1,085.516***
	(0.355, 0.493)	(11.162, 17.957)	(864.415, 1,306.617)
Observations	4,032	4,032	4,032
Adjusted R^2	0.001	0.000	−0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

It is possible that chain restaurants react differently when the local competitive environment changes than non-chain restaurants. Tables 13 through 16 in Appendix 7.4 contain results with chain restaurants removed[9]. As shown, the results are qualitatively very similar to the results with the full sample.

To account for the possibility that the zero coefficient estimates arise as outlier restaurants attenuate meaningful variation, Tables 17 through 20 in Appendix 7.4 contain results with both chain restaurants and outliers removed. At this point the sample size is appreciably reduced. However, again, the results are qualitatively very similar to the results with the full sample.

To summarize these preliminary results, we are finding no evidence of restaurants responding to new competition by changing their menu or changing service quality. Our finding of a zero coefficient is *not* due to a general lack of change or sticky price effect; we find significant changes for all restaurants over this duration but treated and control restaurants appear to be changing in the same way. Further, our results are also not likely to stem from noise in our data or imprecision in our estimates. The coefficients we estimate have confidence intervals quite close to zero and far smaller than the expected amount of change for any restaurant over this period.

6 Conclusion

Measuring the impact of competition on firm behaviour in the restaurant industry entails quantifying both price and non-price competition in a market with many monopolistically competitive firms. We address this complex environment by applying a text differentiation algorithm from the machine learning literature

[9]We define a chain restaurant very loosely by a binary indicating whether we observe another restaurant in New York with the same name. This measure would pick up large franchises, small local chains, and possibly unrelated restaurants that share the same name.

and calliper matching with a rich set of covariates to a novel data set. Our matching technique yields a conservative and plausibly identified causal effect of additional competition on restaurant behaviour. We demonstrate that the cosine distance metric provides a reasonable measurement of qualitative differences between restaurant menus. Our early results suggest that restaurants are not responding to competition from new entrants by changing their prices or product. This finding contradicts the predictions of spatial competition models but is consistent with with CES models of monopolistic competition. In future research we hope to better understand this lack of change by looking more closely at the location choice of entrants, as well as examining specific situations where competition is most likely to occur, such as between restaurants of the same cuisine.

References

Anderson, Simon P and André de Palma, "From local to global competition," *European Economic Review*, 2000, *44* (3), 423–448.

Boer, Rob, Yuhui Zheng, Adrian Overton, Gregory K Ridgeway, and Deborah A Cohen, "Neighborhood design and walking trips in ten US metropolitan areas," *American journal of preventive medicine*, 2007, *32* (4), 298–304.

Chabé-Ferret, Sylvain, "Bias of Causal Effect Estimators Using Pre-Policy Outcomes," *Working P*, 2014.

Chisholm, Darlene C, Margaret S McMillan, and George Norman, "Product differentiation and film-programming choice: do first-run movie theatres show the same films?," *Journal of Cultural Economics*, 2010, *34* (2), 131–145.

Damashek, Marc et al., "Gauging similarity with n-grams: Language-independent categorization of text," *Science*, 1995, *267* (5199), 843–848.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 1977.

Dudey, Marc, "Competition by choice: The effect of consumer search on firm location decisions," *The American Economic Review*, 1990, pp. 1092–1104.

Flores, Carlos A, Alfonso Flores-Lagunes, Arturo Gonzalez, and Todd C Neumann, "Estimating the effects of length of exposure to instruction in a training program: the case of job corps," *Review of Economics and Statistics*, 2012, *94* (1), 153–171.

Heckman, James, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd, "Characterizing Selection Bias Using Experimental Data," *Econometrica*, 1998, *66* (5), 1017–1098.

Hirano, Keisuke and Guido W Imbens, "The propensity score with continuous treatments," *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 2004, *226164*, 73–84.

Imbens, Guido W., "Matching Methods in Practice: Three Examples," *Journal of Human Resources*, 2015, *50* (2), 373–419.

Irmen, A. and J.F. Thisse, "Competition in multi-characteristics spaces: Hotelling was almost right," *Journal of Economic Theory*, 1998, *78* (1), 76–102.

JAFFE, ADAM B, "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value," *The American Economic Review*, 1986, *76* (5), 984–1001.

Konishi, H., "Concentration of competing retail stores," *Journal of Urban economics*, 2005, *58* (3), 488–512.

Krizek, Kevin J, "Operationalizing neighborhood accessibility for land use-travel behavior research and regional modeling," *Journal of Planning Education and Research*, 2003, *22* (3), 270–287.

Netz, J.S. and B.A. Taylor, "Maximum or minimum differentiation? Location patterns of retail outlets," *Review of Economics and Statistics*, 2002, *84* (1), 162–175.

Pinkse, Joris and Margaret E Slade, "Mergers, brand competition, and the price of a pint," *European Economic Review*, 2004, *48* (3), 617–643.

Pollak, Michael, "Knowing the Distance," *New York Times*, September 17 2006.

Rubin, Donald B and Neal Thomas, "Combining propensity score matching with additional adjustments for prognostic covariates," *Journal of the American Statistical Association*, 2000, *95* (450), 573–585.

Salop, Steven C., "Monopolistic competition with outside goods," *Bell Journal of Economics*, 1979, *10.*

Smith, Jeffrey A and Petra E Todd, "Does matching overcome LaLonde's critique of nonexperimental estimators?," *Journal of econometrics*, 2005, *125* (1), 305–353.

Stahl, Konrad, "Differentiated products, consumer search, and locational oligopoly," *The Journal of Industrial Economics*, 1982, pp. 97–113.

Sweeting, Andrew, "The effects of mergers on product positioning: evidence from the music radio industry," *The RAND Journal of Economics*, 2010, *41* (2), 372–397.

Tirole, Jean, *The Theory of Industrial Organization*, MIT Press, 1988.

Wolinsky, A., "Retail trade concentration due to consumers' imperfect information," *The Bell Journal of Economics*, 1983, *14* (1), 275–282.

7 Appendix

7.1 Predicted Entrant Model

Table of Coefficients from LASSO Poisson Regression HERE

7.2 Choice of Predicted Entrant Bandwidth

The two-stage calliper matching process described in the text requires us to choose a bandwidth for the callipers. This bandwidth determines the range of predicted entrant counts in which we search for the closest control observation match by cosine distance. Choosing a bandwidth implies a tradeoff across match quality on the two stages; a small bandwidth ensures a closer match on predicted entrant count in the first stage but also limits the pool of potential closest cosine matches.

We obtain a suitable bandwidth through a process that ensures a close balance of propensity covariates. The selection procedure proceeds as follows:

1. We divide observations into quintiles of predicted entrant count $q \in \{1,2,3,4,5\}$.

2. For each observation i in quantile q we find the observation in quantile $-q \neq q$ with the smallest cosine distance to observation i. Then, we take the average across each quintile q. We denote the maximum of this average across all quantiles as the "maximum average cosine distance".

3. For each observation i in quantile q we select a random observation j from a quintile $-q \neq q$. For each covariate in the Poisson regressions we take the average of the standardized distance between the covariate value for observations i and j. We denote the maximum of this average across all quantiles as the "average maximum Poisson covariate distance".

Figure 9 shows the resulting cosine distances and propensity covariate distances for a bandwidth of α standard deviations for $\alpha \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$. Based on these results, we select a bandwidth of 0.25 standard deviations of predicted entrant count for the two-stage calliper matching procedure.

7.3 Decomposing Price Changes

Restaurants have two margins along which they can adjust their prices: the *intensive* margin (changing the price of items already on the menu) and the *extensive* margin (adding new items that shift the price distribution). To understand which of these effects dominates, we decompose the average change in item prices into the intensive and extensive components. Let I_k, I_a, and I_r be the set of items present at both $t-d$ and $t+d$, present at $t+d$ but not at $t-d$, and present at $t-d$ but not at $t+d$. Let N_k, N_a, and N_r be the cardinalities of these sets. Furthermore, let $\bar{p}^k_\tau$, $\bar{p}^a_\tau$, and $\bar{p}^r_\tau$ be the average prices of each of these sets of items

Table 12: Demographic Variables

	mean	sd	min	max
age.lt.10	0.09	0.1	0.0	0.4
age.18.24	0.10	0.1	0.0	0.7
age.25.29	0.13	0.1	0.0	0.5
age.30.39	0.20	0.1	0.0	0.6
age.40.49	0.14	0.1	0.0	0.5
age.50.59	0.12	0.1	0.0	0.4
age.60.64	0.05	0.0	0.0	0.3
age.65.69	0.04	0.0	0.0	0.2
age.70.79	0.05	0.0	0.0	0.4
race.white	0.64	0.3	0.0	1.0
race.black	0.12	0.2	0.0	1.0
race.asian	0.16	0.2	0.0	1.0
race.latino	0.22	0.2	0.0	1.0
commute.county	0.34	0.3	0.0	1.0
commute.state	0.04	0.0	0.0	0.5
hh.family	0.48	0.2	0.0	1.0
hh.married	0.32	0.2	0.0	1.0
hh.roommates	0.12	0.1	0.0	0.6
school.college	0.09	0.1	0.0	0.8
school.none	0.79	0.1	0.2	1.0
educ.degree	0.52	0.3	0.0	1.0
span.lim	0.06	0.1	0.0	0.8
span.eng	0.14	0.1	0.0	0.8
ie.lim	0.03	0.1	0.0	0.8
ie.eng	0.11	0.1	0.0	0.8
ap.lim	0.03	0.1	0.0	0.7
ap.eng	0.06	0.1	0.0	1.0
poverty	0.11	0.1	0.0	1.0
income.lt.10	0.09	0.1	0.0	0.8
income.10.20	0.09	0.1	0.0	0.6
income.20.30	0.08	0.1	0.0	0.5
income.30.40	0.07	0.1	0.0	0.5
income.50.60	0.07	0.1	0.0	0.6
income.60.75	0.09	0.1	0.0	0.6
income.75.100	0.11	0.1	0.0	0.5
income.100.150	0.14	0.1	0.0	1.0
income.150.200	0.07	0.1	0.0	0.4
income.gt.200	0.12	0.1	0.0	0.7
owner.occ	0.27	0.2	0.0	1.0
unit.detached	0.05	0.1	0.0	1.0
unit.3.9	0.20	0.2	0.0	1.0
unit.10.49	0.28	0.2	0.0	1.0
built.post.2010	0.01	0.0	0.0	0.3
built.1990.1999	0.03	0.1	0.0	0.7
Observations	187906			

There are roughly 6,270 Census block groups in NYC.
Of these, 2510 have a restaurant in our data.
Statistics in the table are weighted by restaurant-period.

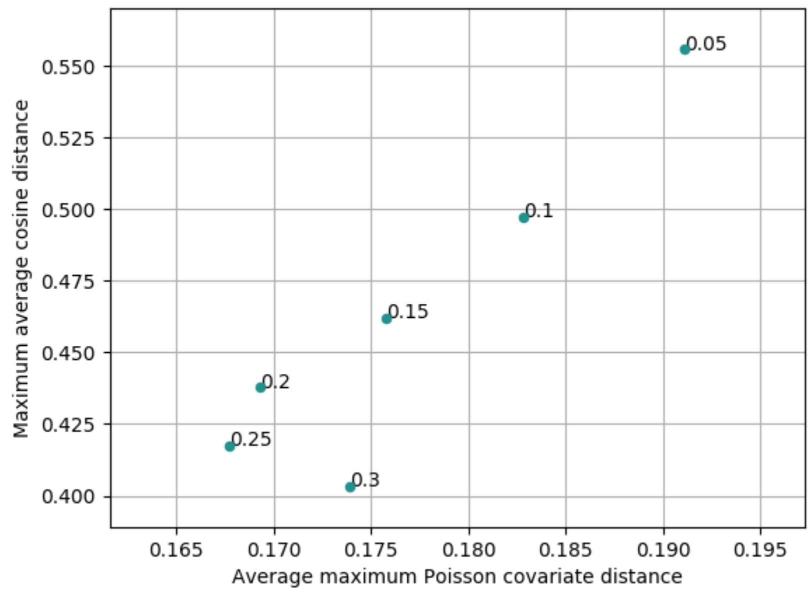

Figure 9: Comparison of cosine distance between treated and control pairs with standardized distance between Poisson regression variables for varying calliper sizes.

at period τ. Then, the intensive and extensive components of the change in price are defined as follows:

$$\Delta^{int}\bar{p} = \frac{N_k^2(\bar{p}^k_{t+d} - \bar{p}^k_{t-d}) + N_k(N_r\bar{p}^k_{t+d} - N_a\bar{p}^k_{t-d})}{(N_k + N_r)(N_k + N_a)} \tag{16}$$

$$\Delta^{ext}\bar{p} = \frac{N_aN_r(\bar{p}^a_{t+d} - \bar{p}^r_{t-d}) + N_k(N_a\bar{p}^a_{t+d} - N_r\bar{p}^r_{t-d})}{(N_k + N_r)(N_k + N_a)} \tag{17}$$

7.4 Regression robustness

Table 13: Regression results for item prices with chain restaurants removed. Coefficients shown with 95% confidence intervals.

	Median price	95th price	Int marg price	Ext marg price
Treated	0.020	−0.081	0.158***	7.058
	(−0.063, 0.103)	(−0.307, 0.146)	(0.040, 0.276)	(−6.629, 20.745)
Constant	0.130***	0.159***	0.000	0.038
	(0.047, 0.213)	(0.039, 0.279)	(−0.000, 0.000)	(−0.013, 0.088)
Observations	2,746	2,746	1,283	1,283
Adjusted R^2	−0.000	−0.000	0.009	0.001

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 14: Regression results for menu text characteristics with chain restaurants removed. Coefficients shown with 95% confidence intervals.

	Item count	Name length	Desc length	Total chars
Treated	1.245	−0.012	−0.017	42.123
	(−0.549, 3.039)	(−0.209, 0.186)	(−0.423, 0.390)	(−90.918, 175.164)
Constant	−6.079***	−1.728**	1.538***	−284.241**
	(−10.212, −1.945)	(−3.099, −0.357)	(0.789, 2.286)	(−562.859, −5.623)
Observations	2,746	2,746	2,746	2,746
Adjusted R^2	−0.000	−0.000	−0.000	−0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 15: Regression results for restaurant services with chain restaurants removed. Coefficients shown with 95% confidence intervals.

	Food good	On time	Order accurate	Hours
Treated	0.092	0.123	0.033	−0.598
	(−0.086, 0.269)	(−0.026, 0.272)	(−0.101, 0.168)	(−1.900, 0.705)
Constant	−0.134**	−0.117**	−0.098**	0.596
	(−0.248, −0.020)	(−0.217, −0.017)	(−0.187, −0.009)	(−0.478, 1.670)
Observations	2,674	2,674	2,674	1,240
Adjusted R^2	0.000	0.001	−0.000	−0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 16: Regression results for absolute menu changes with chain restaurants removed. Coefficients shown with 95% confidence intervals.

	Abs median price	Abs item count	Abs total chars
Treated	0.085*	0.916	30.812
	(−0.005, 0.176)	(−0.733, 2.566)	(−89.586, 151.211)
Constant	0.440***	15.275***	1,138.614***
	(0.364, 0.515)	(11.673, 18.876)	(906.421, 1,370.807)
Observations	2,746	2,746	2,746
Adjusted R^2	0.001	−0.000	−0.000

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 17: Regression results for item prices with chain restaurants and outliers removed. Coefficients shown with 95% confidence intervals.

	Median price	95th price	Int marg price	Ext marg price
Treated	0.024	−0.075	0.101	−0.114
	(−0.083, 0.130)	(−0.281, 0.131)	(−0.062, 0.263)	(−0.286, 0.058)
Constant	0.089*	0.174**	0.000	0.048**
	(−0.007, 0.184)	(0.039, 0.309)	(−0.000, 0.000)	(0.006, 0.090)
Observations	1,282	1,282	550	550
Adjusted R^2	−0.001	−0.000	0.003	0.003
Notes:	***Significant at the 1 percent level.			
	**Significant at the 5 percent level.			
	*Significant at the 10 percent level.			

Table 18: Regression results for menu text characteristics with chain restaurants and outliers removed. Coefficients shown with 95% confidence intervals.

	Item count	Name length	Desc length	Total chars
Treated	0.842	−0.028	−0.244	−6.674
	(−1.260, 2.944)	(−0.306, 0.250)	(−0.715, 0.227)	(−184.671, 171.323)
Constant	−4.273	−0.960	1.489***	−100.452
	(−9.639, 1.093)	(−2.846, 0.926)	(0.449, 2.530)	(−469.269, 268.364)
Observations	1,282	1,282	1,282	1,282
Adjusted R^2	−0.001	−0.001	−0.000	−0.001
Notes:	***Significant at the 1 percent level.			
	**Significant at the 5 percent level.			
	*Significant at the 10 percent level.			

Table 19: Regression results for restaurant services with chain restaurants and outliers removed. Coefficients shown with 95% confidence intervals.

	Food good	On time	Order accurate	hours
Treated	0.049	0.075	0.014	−2.240*
	(−0.200, 0.298)	(−0.153, 0.302)	(−0.190, 0.218)	(−4.716, 0.237)
Constant	−0.240***	−0.202**	−0.203***	1.818
	(−0.389, −0.091)	(−0.355, −0.048)	(−0.344, −0.062)	(−0.571, 4.207)
Observations	1,239	1,239	1,239	496
Adjusted R^2	−0.001	−0.000	−0.001	0.005

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 20: Regression results for absolute menu changes with chain restaurants and outliers removed. Coefficients shown with 95% confidence intervals.

	Abs median price	Abs item count	Abs total chars
Treated	0.097*	−0.047	−15.388
	(−0.016, 0.210)	(−1.969, 1.876)	(−183.076, 152.299)
Constant	0.322***	15.200***	1,133.192***
	(0.239, 0.404)	(10.751, 19.648)	(835.138, 1,431.245)
Observations	1,282	1,282	1,282
Adjusted R^2	0.001	−0.001	−0.001

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.