

**Columbia University**

Department of Statistics

**New York Pizza**

How to Find the Best

BY

Jared P. Lander

December 8, 2008

# Table of Contents

Abstract.....	1
Introduction.....	1
Model Setup and Inference.....	2
Model Diagnostic.....	5
Conclusion.....	8
End Notes.....	9

## Abstract

While pizza was initially invented in Italy, it got propelled to its current status in the streets of New York. As such, New York is often seen as the pizza capital of the world. Here we try to discern what makes a pizzeria more or less favorable to the non-expert pizza eating community based on user reviews at [menupages.com](http://menupages.com). We have information regarding the location, relative price, type of fuel used in the oven, number of reviews and average rating for over 600 pizza serving establishments in Manhattan and parts of Brooklyn. Our results suggest that the average consumer views all pizza to be of the same general quality with the possible exception of Midtown pizza being slightly less desirable. Using the number of user reviews as a proxy for the popularity of a pizzeria, we see that coal fired ovens draw larger crowds than either wood or gas powered ovens.

## 1. Introduction

While pizza, in one form or another, has been eaten since the times of Ancient Greece and Persia, the root of modern pizza was seeded in 1889 when a chef from Naples prepared a pie with mozzarella, basil and tomatoes (the colors of the Italian flag) for King Umberto I and Queen Margherita of Italy. From that moment pizza became a sensation, cooked by pizzaioli (pizza chefs) all over the country in wood burning ovens. Over the next decade, millions of Italians migrated to America bringing their pizza with them. Coal was the most plentiful fuel in the new world, so it was used in pizzeria ovens, giving American pizza a distinctive bend, charred crust and unique flavor.

Gennaro Lombardi opened America's first licensed pizzeria in 1905 at 53 ½ Spring Street in New York City's SoHo neighborhood selling nickel pizza pies to the area's working class Italian denizens. Lombardi's kitchen was a classroom for future pizza innovators who all learned their craft from the master himself. Many of his apprentices went on to open now-famous pizzerias. Anthony "Totonno" Pero left the restaurant and opened Totonno's in Coney Island in 1924. John Sasso followed his former coworker when he established John's of Bleecker in 1929. Patsy Lancieri not only ventured away from Lombardi in 1933 to open his eponymous restaurant in the then Italian dominated Spanish Harlem, but was the first proprietor to sell pizza by the slice. Further, Lancieri's nephew, Patsy Grimaldi, owns and operates Grimaldi's in Brooklyn Heights in the shadow of the Brooklyn Bridge. This family tree of pizzerias grows from there, but these restaurants have three common traits: They do not accept credit cards; do not sell by the slice (except Patsy's), they all use coal ovens.

After World War II, pizza's popularity ballooned as the soldiers returning home from Europe brought back their love for this new delicacy. Pizzerias popped up all over the country and continued to spread like weeds through New York. Gas was the fuel of choice due to its ease of use, low cost and new regulations regarding clean air among other factors. Today there are a plethora of pizza joints with myriad different offerings. A hungry New Yorker can get a quick one dollar slice or wait 40 minutes for a small \$25 pie with no alterations allowed.

With so many factors going into the ingredients, preparation and culture of pizza it is fascinating to question what makes one particular pizza better than others. A big

differentiator is the type of fuel used to fire the oven—typically coal, wood or gas. As mentioned earlier, the price of pizza can fluctuate wildly so it is natural to question if higher prices necessarily indicate better quality. Even geographic location can play a role, whether it is through cultural differences or the quality of the water.

The dataset (pulled in mid October 2008) comes primarily from [menupages.com](http://menupages.com)<sup>1</sup>, a website hosting restaurant menus and allowing users to comment and rate participating merchants. Covering Manhattan and parts of Brooklyn, the site contains information on 699 restaurants tagged as “Pizza” (not a full enumeration) which includes pure pizzerias, Italian restaurants, delis and other types of establishments, but none of the national chains such as Pizza Hut, Domino’s or Sbarro. Each restaurant entry had a number of user reviews (ranging between 1 and 146), address, neighborhood, price level (between 1 and 5) and an average user rating (integer and half values from 1 to 5). From this list 55 restaurants either had no rating or were closed and thusly removed.

The data were augmented with two more explanatory variables. The first is an indicator variable describing whether some version of the word pizza is in the name of the restaurant. The last—and according to many critics, the most important—variable is the type of fuel used. The vast majority of pizzerias in this study utilize gas ovens to cook their pizza while only 52 use wood and 18 utilize coal (information supplied by “Slice,” a pizza blog<sup>2</sup>). Table 1.1 shows a sampling of the data.

**Table 1.1** Sample Data

Rating	Name	Neighborhood	Area	Price	Price.Level	Reviews	Number.Reviews	Fuel	PizzaName
3	Village Restaurant	East Village/LES	Downtown	Cheap	1	Low	3	Gas	0
3.5	Cascina Ristorante	Midtown West	Midtown	Expensive	4	Low	20	Wood	0
3.5	Tuscany	Midtown West	Midtown	Cheap	2	Low	8	Gas	0
4	Vinci's	Upper East Side	Uptown	Cheap	2	Low	21	Gas	0
3.5	Il Corallo Trattoria	SoHo	Downtown	Cheap	2	High	81	Gas	0
3	Fat Sal's	Upper East Side	Uptown	Cheap	2	Low	12	Gas	0
4	Nolita House Restaurant & Bar	SoHo	Downtown	Expensive	3	Low	29	Wood	0
3.5	L'Allegria	Midtown West	Midtown	Expensive	4	Low	29	Gas	0
3.5	Celeste	Upper West Side	Uptown	Expensive	3	Medium	57	Wood	0
3.5	Bravo Pizza	Chelsea	Downtown	Cheap	2	Low	5	Gas	0

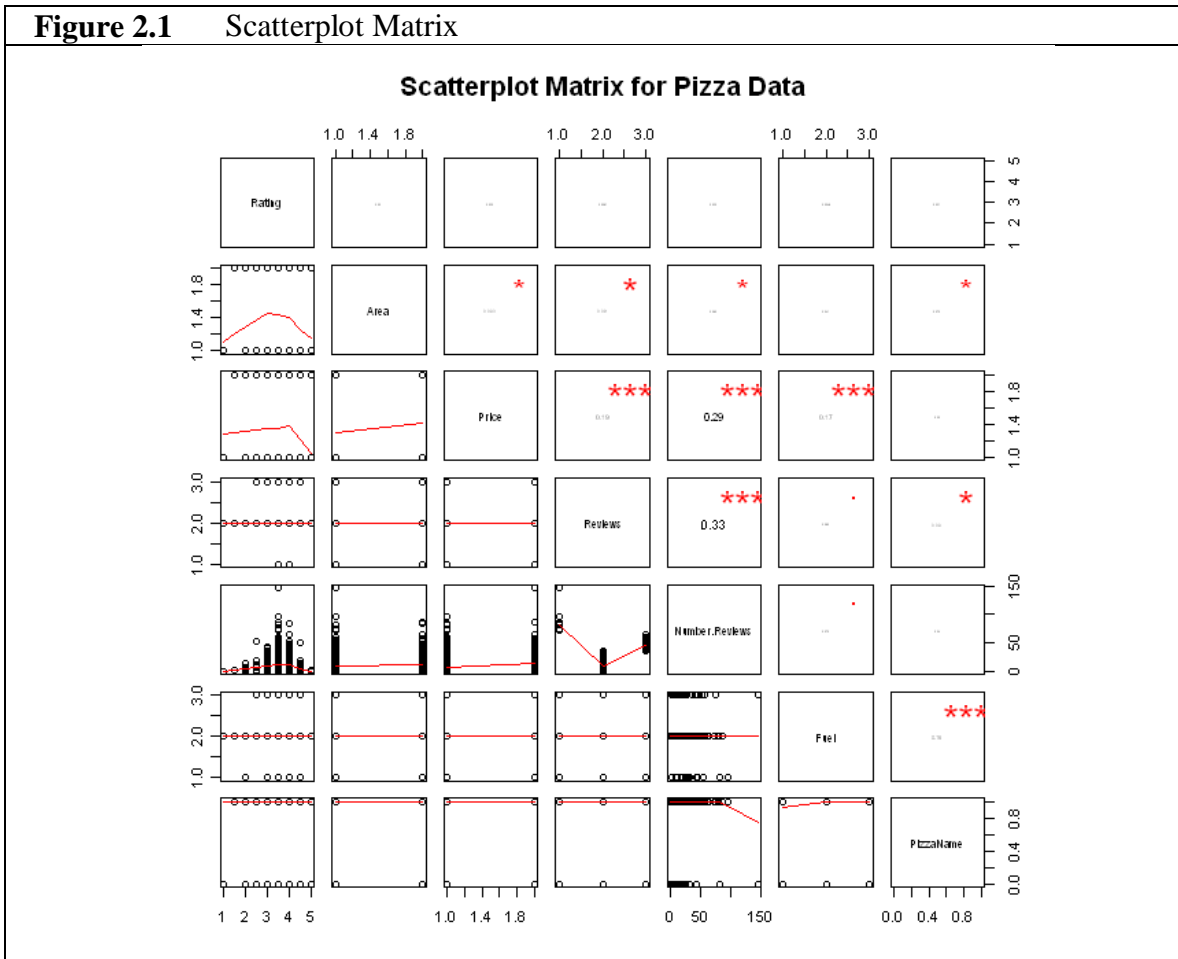
Due to the geographic restrictions and selective nature of the ratings a number of famous coal pizzerias (such as the original Totonno’s in Coney Island and Carbone in Manhattan) were not included. Another unfortunate artifact of the geographic criteria is that Di Fara Pizza in Midwood, Brooklyn—often ranked the single best pizzeria in all of New York City, and possibly the country—is left out of this study.

There are two possible response variables of interest. The rating is suggestive of what the non-expert pizza eating community thinks of a restaurant. The number of reviews of a restaurant can serve as a proxy for popularity though makes no claims as to quality. The latter is logical as a predictor of the former but not the reverse. This study will build models to determine what elements go into the observed quality and popularity of pizzerias in New York.

## 2. Model Setup and Inference

In order to build the two models the data were arranged in a way conducive to analysis. The neighborhoods were grouped together in a variable named Area and labeled as either “Uptown” (north of Manhattan’s Chelsea) or “Downtown” (Chelsea and

farther south plus Brooklyn). An alternative grouping also breaks out Midtown Manhattan and Brooklyn eateries. Price was binned into “Expensive” (levels 3, 4 and 5) and “Cheap” (levels 1 and 2). The number of reviews was also binned into “Low” (1 to 35), “Medium” (36 to 70) and “High” (71 to 148). A scatterplot matrix is shown in Figure 2.1.

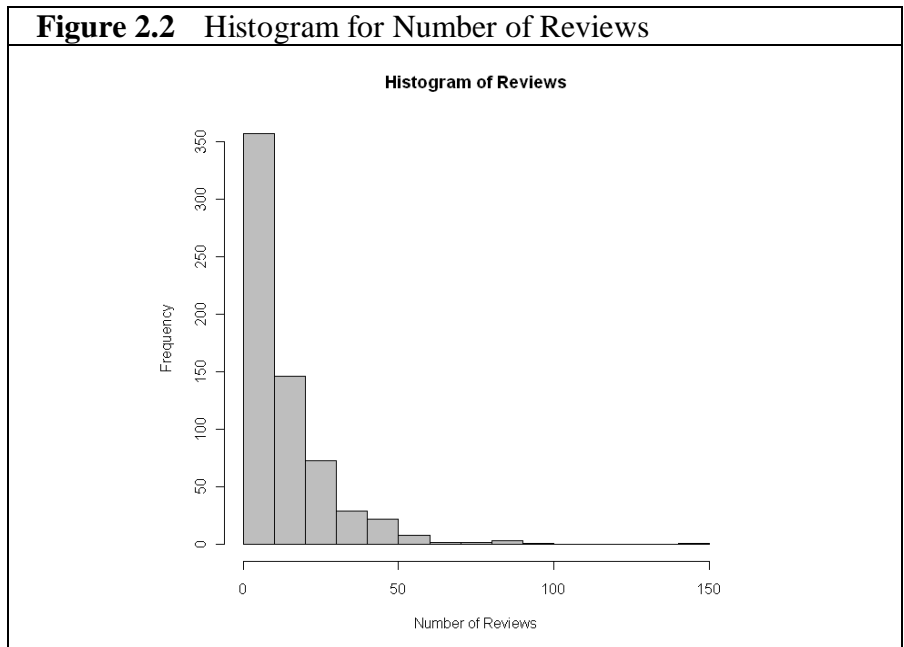


To assess a pizzeria’s perceived quality we regressed  $\text{logit}(\text{Rating}/5.5)$  on Area, Fuel, Price, PizzaName and Reviews as well as various cross terms. None of the coefficients were significant when using the compressed Area groupings, as seen in Table 2.1 (a). When Brooklyn and Midtown were broken out individually, a Midtown location has a slight detrimental effect with a p-value of .007 (Table 2.1 (b)). An ANOVA Test and Wald Test (discussed in the Diagnostic section) both suggest that the model  $\text{logit}(\text{Rating}/5.5) \sim \text{Area}$ , Table 2.1 (c) is a superior fit, indicating that most of our predictors have little effect on the perceived level of quality.

The number of ratings a pizzeria receives (the variable is Number.Reviews) is a count and is approximately distributed as a Poisson(13.75) and thus well suited for Poisson regression. The histogram is in Figure 2.2. To judge the count we used a generalized linear model, regressing Number.Reviews on Area, Fuel, Price and PizzaName. Only the shortened version of Area (delineated as “Uptown” and

“Downtown”) is used to reduce the number of cells with zeros. The final model decided upon, on display in Table 2.1 (d), relies on Area, Fuel and Price with Fuel\*Price as the only cross term which a Wald Test and Coefficient Test confirm as a strong model. A potential outlier is discussed in the Diagnostics section.

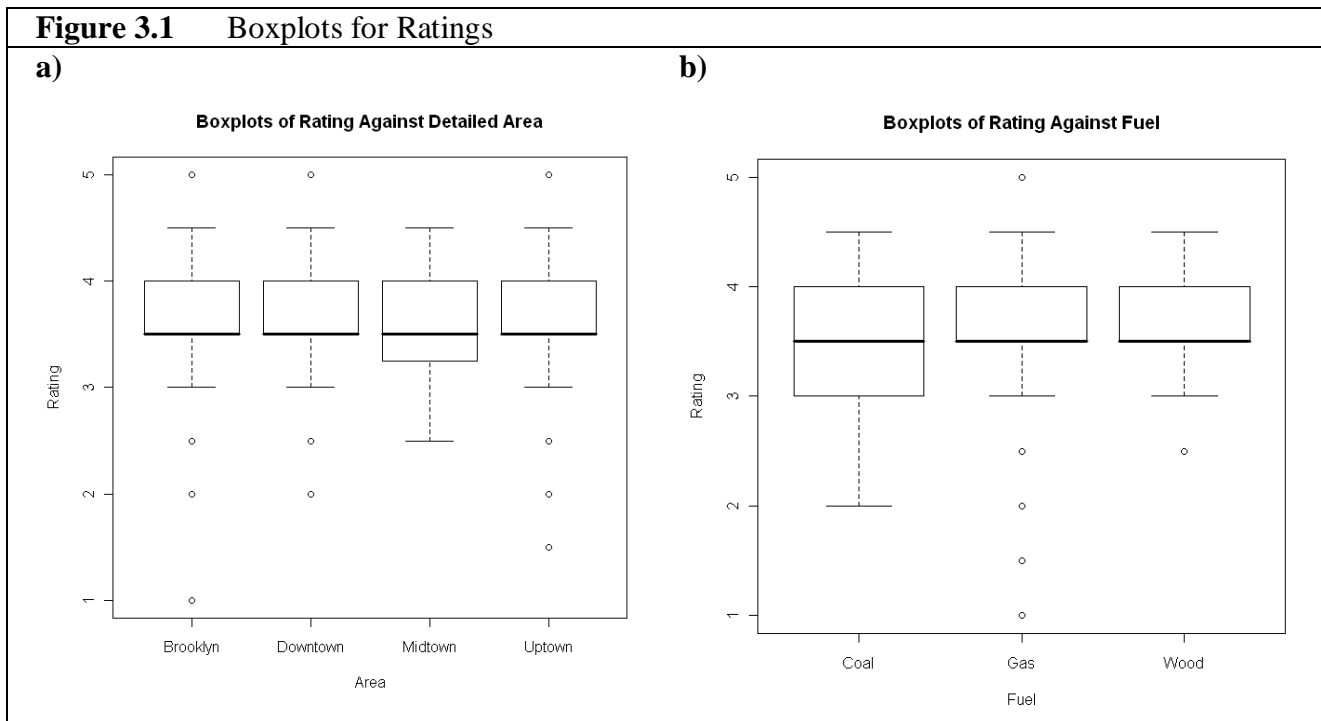
<b>Table 2.1</b> Regression Summaries					
<b>a) Model 1: Ratings Full Model (Collapsed Area)</b>			<b>b) Model 1: Ratings Full Model (Detailed Area)</b>		
	<u>Est</u>	<u>SE</u>	<u>t-stat</u>	<u>P-val</u>	
(Intercept)	0.551	0.196	2.815	0.005	**
Uptown	-0.055	0.037	-1.465	0.143	
Gas	0.136	0.116	1.168	0.243	
Wood	0.083	0.130	0.635	0.526	
Expensive	-0.032	0.040	-0.786	0.432	
ReviewL	0.052	0.180	0.288	0.773	
ReviewM	0.047	0.191	0.244	0.807	
Name	0.016	0.062	0.264	0.792	
AIC: <b>845.8049</b>	BIC: <b>886.0142</b>				
<b>c) Model 1: Ratings Final Model</b>			<b>d) Model 2: Reviews Final Model</b>		
	<u>Est</u>	<u>SE</u>	<u>t-stat</u>	<u>P-val</u>	
(Intercept)	0.776	0.047	16.497	≈0	***
Downtown	-0.058	0.055	-1.057	0.291	
Midtown	-0.178	0.065	-2.733	0.006	**
Uptown	-0.049	0.059	-0.827	0.409	
AIC: <b>835.4253</b>	BIC: <b>857.7638</b>				
	<u>Est</u>	<u>SE</u>	<u>t-stat</u>	<u>P-val</u>	
(Intercept)	3.827	0.059	64.876	≈0	***
Uptown	0.157	0.022	7.297	≈0	***
Gas	-1.642	0.060	-27.273	≈0	***
Wood	-0.735	0.089	-8.271	≈0	***
Expensive	-0.546	0.079	-6.890	≈0	***
Gas:Exp	1.119	0.083	13.501	≈0	***
Wood:Exp	0.670	0.108	6.195	≈0	***
AIC: <b>8799.019</b>	BIC: <b>8830.292</b>				



### 3. Model Diagnostic

Model 1:  $\text{logit}(\text{Rating}/5.5) \sim 0.77571 - 0.1782 * \text{Midtown}$

As seen in Figure 3.1 (a) and (b), there is little visible difference between the areas and the type of fuel used. Stepwise regression indicated regressing only on the variable Area as seen in the Regression Summary in Table 2.1 (c). A Wald Test, in Table 3.1 (a), shows that this model is superior to the null model and an ANOVA Test, Table 3.1 (b), shows that the variable Area is significant. The summary of the regression, Table 2.1 (c), shows that the categorical variable Area is only significant for the Midtown level. The final model chosen had the lowest AIC of all models tried at 835.4253.



**Table 3.1** Regression Summaries

<b>a) Wald Test</b>				<b>b) ANOVA</b>							
Model 1: $\text{logit}(\text{Rating}/5.5) \sim \text{Area}$					<u>Sum</u>	<u>Mean</u>					
Model 2: $\text{logit}(\text{Rating}/5.5) \sim 1$					<u>Sq</u>	<u>Sq</u>	<u>F-stat</u>	<u>P-val</u>			
	<u>Res.Df</u>	<u>Df</u>	<u>F-stat</u>	<u>P-val</u>							
	1	640									
	2	643	-3	2.820	0.038	*					
					Area	3	1.795	0.598	2.820	0.038	*
					Resid	640	135.848	0.212			

To further investigate this model we compared the mean difference in Rating broken up by the two most significant predictors: Area and Fuel. Two-sided Welch Two-Sample t-tests indicated possible significant differences only between Midtown and Uptown, Downtown and Brooklyn (respective p-values: .013, .011 and .040). One-sided t-tests suggested that Midtown pizza is of lower quality than Uptown, Downtown and Brooklyn (respective p-values: .007, .006 and .020). A summary for the one sided tests

is in Table 3.2. T-tests showed no significant difference between coal, wood or gas ovens.

The overall mean rating for the pizzerias in our dataset is 3.646. This average may be more appropriate for assigning quality than the result of Model 1. The lone coefficient in the model does not account for much difference and including it may just lead to unnecessary complication.

<u>Area 1</u>	<u>Area 2</u>	<u>Diff</u>	<u>T-stat</u>	<u>Df</u>	<u>P-val</u>
Midtown	Uptown	-0.141	-2.484	248.096	0.007
Midtown	Downtown	-0.131	-2.558	224.596	0.006
Midtown	Brooklyn	-0.169	-2.067	153.205	0.020

Model 2:  $\text{Number.Reviews} \sim 3.82693 + 0.15733*\text{Uptown} - 1.64163*\text{Gas} - 0.73476*\text{Wood} - 0.54561*\text{Expensive} + 0.11908*\text{Gas:Expensive} + 0.67017*\text{Wood:Expensive}$  (Poisson Regression)

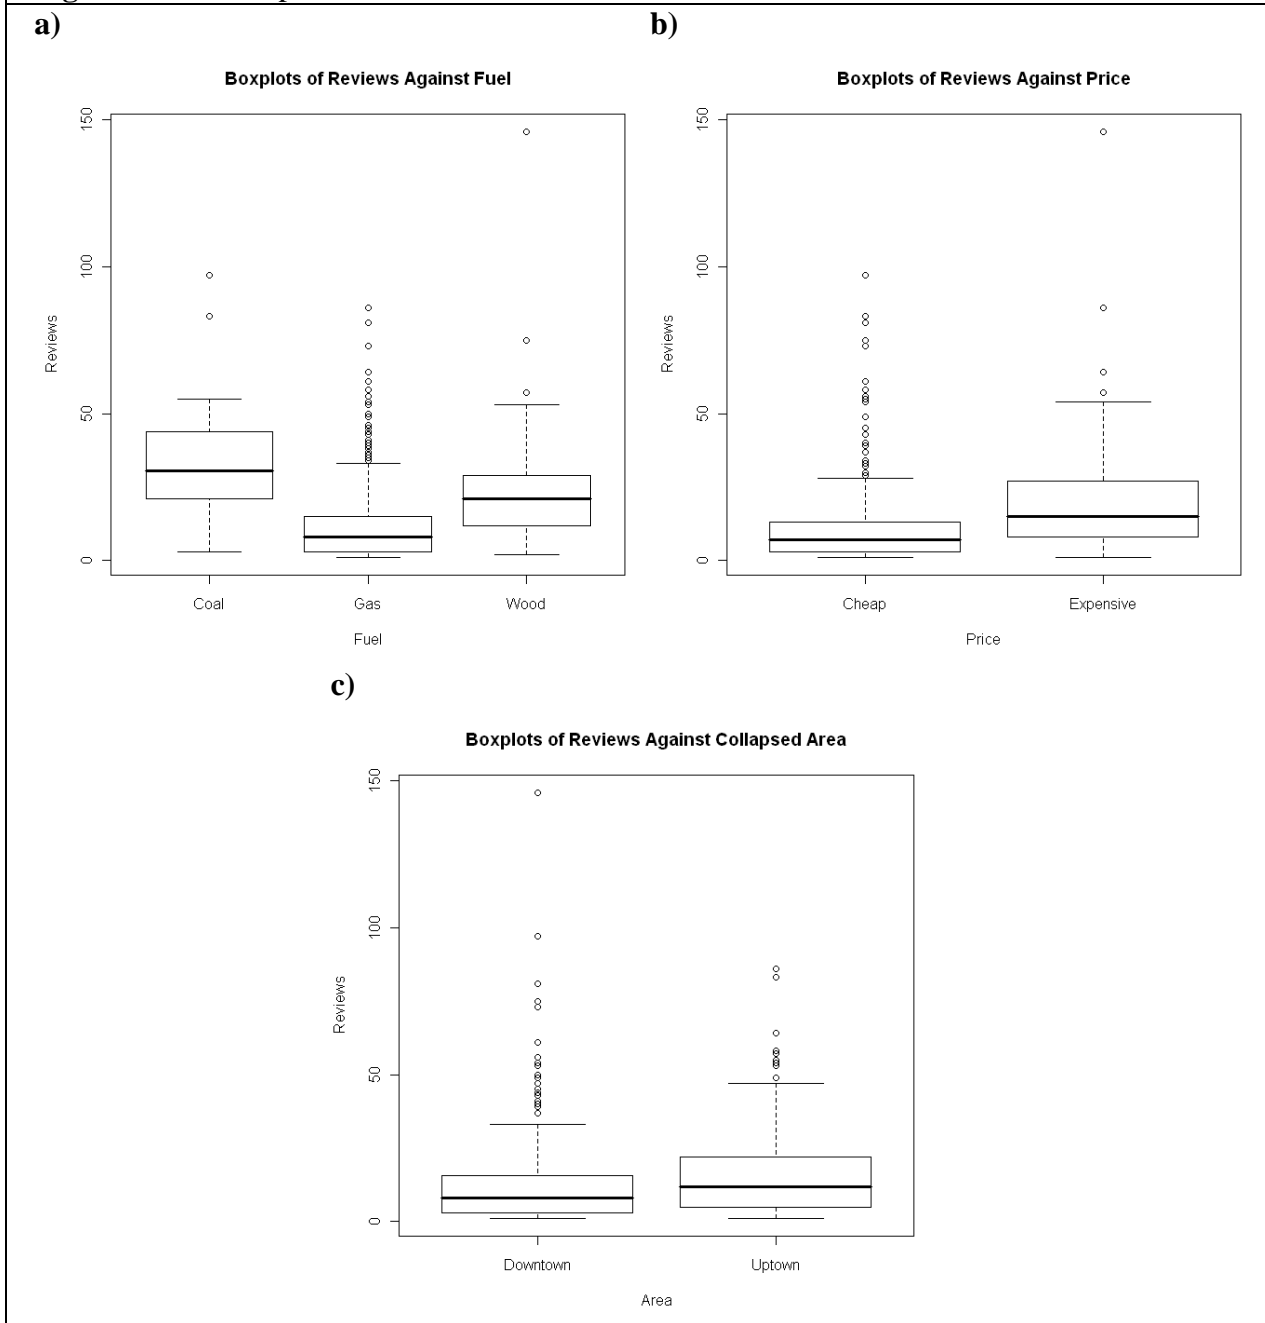
The boxplots in Figure 3.2 (a) clearly reveal that coal pizzerias receive the most reviews followed by wood and then gas. Figure 3.2 (b) shows expensive pizzerias garnering more reviews. The boxplots in Figure 3.2 (c) are less distinctive but suggest uptown customers review their pizza places at a greater rate than their downtown counterparts.

The graphical views are supported by the Regression Summary table in Table 2.1 (d) which exhibits the coefficient estimates and their p-values. The Coefficient Test in Table 3.3 (a) further corroborates that each coefficient at their current level is significant and the Wald Test in Table 3.3 (b) confirms that the model is strong. This model was reached by starting with a fuller model featuring more variables and cross terms that were narrowed down through a combination of stepwise regression, Wald Tests and Coefficient Tests. The variable PizzaName was notably left out which is all the better as many restaurants even have their own internal inconsistencies regarding their names. A cross term, Fuel\*Price, was added because historic coal and fancy wood restaurants often have higher prices than the corner pizza shop. Further, a Drop-in Deviance Test, Table 3.3 (c), showed it to be significant with a p-value near 0. The model decided upon had a BIC of 8830.292.

One-sided t-tests strongly suggest that both coal and wood outdraw gas with respective estimated margins of 23.64 and 13.87 reviews. The t-test does not suggest a necessarily significant difference between coal and wood. However, Otto Enoteca & Pizzeria is a potential outlier drawing 146 reviews, which is far outside the norm. This is most likely due to owner Mario Batali's many other successful and trendy locales. Removing this observation leads to an estimated mean difference between coal and wood of 12.12. The difference between Wood and Gas is lessened with the outlier removed, though it is still significant. Removing the outlier does not significantly affect the difference between different Areas nor does it affect the selected models.



**Figure 3.2** Boxplots for Number of Reviews



The outlier's influence in the first model was 0.173 and 0.047 in the second model which further suggests that the observation can safely be omitted from the data without great impact. The outlier is legitimate because it did not originate out of error. However, it does possibly misrepresent the popularity of wood burning ovens in general.

a) Coefficient Test						b) Wald Test					
	Est	SE	z-stat	P-val		Model 1: Number.Reviews ~ Area+Fuel+Price+Price*Fuel					
(Intercept)	3.827	0.059	64.876	≈0	***	Model 2: Number.Reviews ~ 1					
Uptown	0.157	0.022	7.297	≈0	***		<u>Res.Df</u>	<u>Df</u>	<u>F-stat</u>	<u>P-val</u>	
Gas	-1.642	0.060	-27.273	≈0	***	1	637				
Wood	-0.735	0.089	-8.271	≈0	***	2	643	-6	309.94	≈0 ***	
Expensive	-0.546	0.079	-6.890	≈0	***						
Gas:Exp	1.119	0.083	13.501	≈0	***						
Wood:Exp	0.670	0.108	6.195	≈0	***						
c) Drop-in Deviance Test						d) ANOVA					
Model 1: Number.Reviews ~ Area+Fuel+Price+Fuel*Price							<u>Df</u>	<u>Dev</u>	<u>Resid. Df</u>	<u>Resid. Dev</u>	
Model 2: Number.Reviews ~ Area+Fuel+Price						NULL			643	7923.4	
	<u>Resid. Df</u>	<u>Resid. Dev</u>	<u>Df</u>	<u>Dev</u>	<u>P-val</u>	Area	1	80.9	642	7842.5	
1	637	6225.1				Fuel	2	1026.8	640	6815.7	
2	639	6418.2	-2	-193.1	≈0 ***	Price	1	397.5	639	6418.2	
						Fuel:Price	2	193	637	6225.1	

		With Otto Enoteca & Pizzeria				Without Otto Enoteca & Pizzeria			
Fuel 1	Fuel 2	Diff	T-stat	Df	P-val	Diff	T-stat	Df	P-val
Coal	Gas	23.642	4.181	17.289	≈0	-	-	-	-
Coal	Wood	9.765	1.509	28.661	0.071	12.121	2.006	22.358	0.029
Wood	Gas	13.878	4.297	53.725	≈0	11.522	5.119	55.733	≈0

## 4. Conclusion

Typical Top 10 lists for New York City pizza—which are subjectively compiled by individuals—all vary but for the most part contain at least some of the pizzerias randomly arranged in Table 4.1. Of these typical entries, only Patsy’s in East Harlem is among the top 10 pizzerias, by Rating, in our dataset (for records with five or more reviews). However, Lombardi’s, Pizza 33 and John’s of Bleecker’s Times Square location are part of our data’s 10 most reviewed. This by itself shows a big discrepancy between critical opinions and those of the general public.

Lombardi's	Patsy's (East Harlem)
Di Fara	Pizza 33
Joe's	Nicks' Pizza (Queens)
John's of Bleecker	Artichoke
Totonno's (Coney Island)	Maffei's
Vinny Vincenz	New York Pizza Suprema
Denino's Pizzeria & Tavern (Staten Island)	Una Pizza Napoletana
No. 28	Franny's
Grimaldi's	Joe & Pat's

The fitted values for Model 2, shown in Table 4.2, agree more with the typical Top 10 lists with Totonno's (both Manhattan locations), Patsy's (East Harlem location) and Grimaldi's in the Top 10. This suggests that our model does a good job of capturing the truth about the relative popularity of pizzerias. It must be noted, that it is very likely the professional Top 10 lists drive people to the more famous pizzerias, thus increasing their number of reviews, so the two are interlinked.

**Table 4.2** Fitted Values for Model 2

Fitted	Number.Reviews	Rating	Name	Neighborhood	Area	Price	Price.Level	Fuel
44.35367	44	3	Totonno's	Upper East Side	Uptown	Expensive	3	Coal
44.35367	25	3.5	Patsy's Pizzeria	Midtown East	Uptown	Expensive	3	Coal
44.35367	32	3.5	Patsy's Pizzeria	Upper West Side	Uptown	Expensive	3	Coal
44.35367	13	3.5	Patsy's Pizzeria	Upper East Side	Uptown	Expensive	3	Coal
44.35367	10	4.5	Patsy's	Harlem	Uptown	Expensive	3	Coal
37.61977	30	3	Totonno's	Murray Hill	Downtown	Expensive	3	Coal
37.61977	20	3.5	Patsy's Pizzeria	Chelsea	Downtown	Expensive	3	Coal
37.61977	21	3.5	Patsy's Pizzeria	Village	Downtown	Expensive	3	Coal
37.61977	25	3	Patsy's Pizzeria	Murray Hill	Downtown	Expensive	3	Coal
37.61977	47	4	Luzzo's	East Village/LES	Downtown	Expensive	3	Coal
37.61977	44	3	Grimaldi's	Brooklyn	Downtown	Expensive	3	Coal
37.61977	31	4	Arturo's	Village	Downtown	Expensive	3	Coal
30.77105	19	3	Serafina Fabulous Pizza	Upper East Side	Uptown	Expensive	4	Wood
30.77105	20	3.5	Serafina Fabulous Grill	Upper East Side	Uptown	Expensive	4	Wood
30.77105	44	3.5	Serafina Broadway	Midtown West	Uptown	Expensive	4	Wood
30.77105	23	4	Petaluma	Upper East Side	Uptown	Expensive	4	Wood
30.77105	47	3.5	Nick's	Upper East Side	Uptown	Expensive	4	Wood
30.77105	19	3.5	Naples 45	Midtown East	Uptown	Expensive	4	Wood
30.77105	14	3.5	Mezzaluna	Upper East Side	Uptown	Expensive	4	Wood
30.77105	6	3.5	Leone Pizza & Pasta	Midtown West	Uptown	Expensive	3	Wood

According to Model 2, an Uptown location adds to a pizzeria's observed popularity (based on the number of reviews) while high prices are not necessarily a detriment. Even more so than those two variables, a coal oven is a big draw. This could be due to their rarity, historic nature or the general affinity the pizza cognoscenti have for charred pies.

The numerous variables and factors have very little affect on the average ratings attributed to pizzerias. For the most part, all the pizzerias were rated on a fairly level playing field, hence using the mean as a simple model. This could indicate, as is the case with wine, that people in general do not have a sophisticated enough palate to fully appreciate the many facets of pizza.

Our findings were able to discern the factors that go into a pizzeria's popularity but did not discover much differentiation in quality. Popularity and quality are not always equivalent. It is likely that we may have just proved the old adage about pizza: "Even when it's bad, it's still good."

<sup>1</sup> <http://www.menupage.com/>

<sup>2</sup> <http://slice.seriousseats.com/>