Does a Daily Deal Promotion Signal a Distressed Business? An Empirical Investigation of Small Business Survival

Ayman Farahat
Adobe Inc.
farahat@adobe.com

Nesreen K. Ahmed
Purdue University
nkahmed@purdue.edu

Utpal Dholakia
Rice University
dholakia@rice.edu

ABSTRACT

In the last four years, daily deals have emerged from nowhere to become a multi-billion dollar industry world-wide. Daily deal sites such as Groupon and Livingsocial offer products and services at deep discounts to consumers via email and social networks. As the industry matures, there are many questions regarding the impact of daily deals on the marketplace. Important questions in this regard concern the reasons why businesses decide to offer daily deals and their longer-term impact on businesses. In the present paper, we investigate whether the unobserved factors that make marketers run daily deals are correlated with the unobserved factors that influence the business, In particular, we employ the framework of seemingly unrelated regression to model the correlation between the errors in predicting whether a business uses a daily deal and the errors in predicting the business' survival. Our analysis consists of the survival of 985 small businesses that offered daily deals between January and July 2011 in the city of Chicago. Our results indicate that there is a statistically significant correlation between the unobserved factors that influence the business' decision to offer a daily deal and the unobserved factors that impact its survival. Furthermore, our results indicate that the correlation coefficient is significant in certain business categories (e.g. restaurants).

Categories and Subject Descriptors

D.2.8 [**Computational Advertising**]: Economics of Daily Deals

Keywords

Daily deals, Consumer ratings, Seemingly unrelated regression

1. INTRODUCTION

Daily deal sites such as Groupon represent a novel approach to Internet marketing that tap into local markets,

and based on the massive scale and rapid growth of such sites, the business model has gained rapid adoption from a wide range of businesses.

Despite this success, a vocal contrarian view of the daily deals model has emerged. Its chief criticism is probably skepticism about the value of daily deals to the merchants whose goods and services are promoted. To be sure, enthusiastic advocates are easy to find; Groupon has claimed that 97% [11] of businesses using its service want to be featured again. But an independent study estimates repeat intent at only 48.1%. Some anecdotal reports are surprisingly harsh, including a highly publicized blog posting by the owner of a New York bakery cafe, who described her Groupon promotion as "the single worst decision I have ever made as a business owner".

The diverging views about the profitability and long term impact of the multi-billion daily deals industry calls for a thorough study and more details than previous attempts. Any study pertaining to evaluate the impact of daily deals on key business metrics needs to isolate the *causal impact* of daily deals from other factors that might be correlated with daily deal adoption and also impact business metrics.

In this paper, we examine the impact of daily deals on the business from the lens of modern econometrics. It is well known that the gold standard for evaluating a treatment effect such as daily deals is through randomized experiments. In practice however, it is not possible to run the experiments needed to isolate the impact of daily deals from other confounding factors. Therefore, most studies including ours will have to contend with observational data. Working with observational data poses its own set of challenges and it is exactly these challenges that call for tools from modern econometrics.

One of the main challenges of working with observational data is unobserved heterogeneity. While two businesses might look identical along the observed dimensions such as category, location..*etc*, they may differ along dimensions that can have a significant impact on the business. For example, a struggling business with a high staff turnover will be more likely to fail than a similar business. Also, struggling businesses might be more tempted to use a daily deal to help shore sales. A naive application of standard statistical techniques will show daily deal adoption to be a factor in business failure while in reality it is staff turnover that contributed to the failure. To alleviate the problems of unobserved heterogeneity, we treat the daily deal adoption as a *dependent variable*. A business's decision to adopt a daily deal will depend on a number of observed independent vari-

ables such as the business category, how popular are deals in category, etc. The daily deal decision will also be impacted by unobserved factors. Similarly, business failure will depend on observed dependent variables and unobserved factors. For example, consider a business that runs a daily deal despite of having no apparent reason to do so. Does that signal anything about the unobserved reasons that might lead the business to fail down the line? If the unobserved factors that make a business run a daily deal are correlated with the unobserved factors that lead to business failure then indeed the daily deal adoption conveys additional information about the business and should be taken into account. On the other hand, if there is no correlation between the unobserved factors in daily deal and failure, then knowing whether a business offered a daily deal does not convey additional information.

We resort to techniques from modern econometrics to help test whether the unobserved factors in daily deal adoption and failure are correlated. In particular, we use the framework of multiple equation models. One equation in our setting models the business's decision to offer a daily deal and the second equation models the business failure. In order to model both equations, we need to identify the dependent variables (daily deal adoption and failure) and a set of covariates that best describe each model.

In this work, we tried to insure that our results were statistically significant, robust to modeling assumptions and scalable. To help achieve statistical significance, we compiled a large data set that had information on daily deal adoption, business failure and information about the businesses. To test robustness of our models, we used two frameworks that make different modeling assumptions. To insure scalability, we developed a data driven approach that does not require expensive interviews with customers or lab experiments. In particular, we crawled Yelp to get information about the business and whether a business was closed. For daily deal adoption, we used a data set from [14].

In this paper we develop a model for business failure in Chicago where both the failure data and business information are derived from Yelp. We also develop a separate model for daily deal adoption where we join the data set from [14] with Yelp data. We decided to use Chicago because in addition to being the third largest city in the U.S, it is the home town of the largest daily deal provider "Groupon" and had a large number of business adopting daily deals. We then developed joint models of daily deal adoption and business failure. Our results show that a joint model of business failure and daily deal adoption does a better job in explaining the data than the two separate models. We tested the robustness of our results by specifying two different modeling paradigms with different model assumption ; bivariate probit and seemingly unrelated regression.

Our results show that models of business failure and daily deals adoption using Yelp based features provide good performance. Furthermore, we find a positive and statistically significant correlation between the unobserved factors that make a business offer a daily deal and the unobserved factors that contribute to failure. Our results indicate that the correlation is strongest in case of Restaurants **0.281** and smaller in case of Spas **0.24**. These results are consistent with results of Gupta [11] that found that deals are generally good for Spas and bad for Restaurants.

In summary, the main contributions of this paper are the following:

- We conducted a data driven large scale study to test and quantify whether the unobserved factors that make a business decide to make daily deals are correlated to the unobserved factors that impact the business survival.

- We developed and analyzed business survival and daily deal adoptions models based on business features collected from Yelp.

- We conducted statistical tests and two different econometrics frameworks to insure statistical significance of results.

- Our results are consistent with findings from previous research work that used labor intensive surveys.

This paper is organized as follows. In Section 2, we survey related work. In Section 3, we describe our data and the methods employed to collect it. In Section 4, we describe the necessary background from econometrics that is needed to develop our models. In Section 5, we describe our experiments, results and evaluations. Finally, in Section 6 we summarize our findings and give directions for future work.

2. RELATED WORK

Recently, there has been an increasing number of both empirical and theoretical research on daily deals or what is sometimes called voucher discounting. Most of this work focused on studying Groupon and LivingSocial discounted deals. For example, Dholakia studied the question of whether Groupon promotions are profitable for businesses and which businesses fare the best and worst after offering a Groupon promotion [7, 8]. He found mixed results where some business owners reported their Groupon promotion was profitable and others regretted making the promotion based on their experience of lower spending and return rates from Groupon users. In another study, Arabshahi provided a detailed analysis that explains the Groupon business model and its underlying principles [3]. In his paper, he explained that the main challenge facing merchants lies in identifying price-sensitive potential customers and offering them discounts. Therefore, Groupon can help merchants to apply price discrimination through the "highly discounted deals" provided to a massive scale of price-sensitive subscribers. In a similar work, Edelman *et al.* provided a theoretical study of the economics of Groupon deals from the perspective of participating merchants rather than from the perspective of the deal service provider [9]. Their results indicate that voucher discounts are naturally good fits for certain types of merchants, and poor fits for others.

Similarly, Gupta *et al.* [11] investigated when are daily deals profitable for business by interviewing over 2000 business that offered daily deal through "Groupon". They found that the success of a daily deal is far from certain and that the return on investment varies widely. They identified the types of businesses as a reliable predictor of profitability and that daily deals are good for spas but bad for restaurants.

The work presented in [7, 8] and [11] can be viewed as complementary to our work; while we focused on a data driven scalable approach, [7, 8] and [11] focused on a more labor intensive interview process. These two approaches are not exclusive with the findings from one guiding the other.

Business Category	Total No.	Multiple Deals (%)	Closed (%)
RESTAURANTS & BARS	337	7.1	8.3
BEAUTY & SPAS	189	10	4.7
ACTIVE LIFE	151	9.2	1.9
SHOPPING	116	6.4	0.9
FITNESS & INSTRUCTION	87	5.3	0.0
FOOD	84	7.1	11.9
HEALTH & MEDICAL	77	12.1	0.0
NIGHTLIFE	73	5.5	1.4
HAIR SALONS	73	9	6.5
ARTS & ENTERTAINMENT	72	6.8	1.4

Table 1: Statistics of Groupon Businesses in Chicago (Jan-July 2011)

Also, Ye *et al.* studied the group purchasing behavior of daily deals in Groupon and LivingSocial and they proposed a predictive dynamic model for group buying behavior [20]. Their model was able to predict the popularity of group deals as function of time. They also found that the different incentive mechanisms applied in Groupon and Living Social (individual threshold versus collective threshold) lead to different propagation behavior, which finally lead to different predictability.

While studying daily deals is interesting in itself, another trend of research started to study the marriage between daily deal sites and the growing consumer phenomena such as Yelp. Byers *et al.* initiated the study of how daily deal sites affect the reputation of the business and in particular the business Yelp reviews [5]. In their first research paper, the authors studied the interplay between social networks and daily deal sites. They found that daily deal sites benefit from significant word-of-mouth effects during sales events. They also studied the effects of daily deals on the long-term reputation of merchants, based on their Yelp reviews before and after they run a daily deal. They found that the Yelp ratings of Groupon-bearing consumers were on average 10% lower than those of their peers.

In another study, Byers *et al.* rigorously evaluated various hypotheses about underlying consumer and merchant behavior to understand the Groupon effect on businesses [6]. They examined a number of hypotheses to justify the Groupon effect. For example, they illustrated a poor business behavior, and Groupon user experimentation to be possible root causes of the Groupon effect. They also found an evidence that on average Groupon users are no more critical than their peers.

Similarly, Zervas tried to establish basic facts regarding the evolving quality of the deals that Groupon offers [22]. He used Yelp ratings as a proxy for measuring the quality. Using simple regression analyses, Zervas found a statistically significant negative correlation between the time deals that have been offered and the Yelp ratings of the merchants who offered them. Further, he discussed some possibilities that might cause these trends. For example, as Groupon is expanding the number of deals it offers, it has to work with some lower-rated merchants. Also, it is possible that better-rated merchants dropping out of running Groupon deals, and Groupon has to substitute them for merchants with some lower-rated merchants.

Our work builds on [6] and [22] by explicitly modeling the decision to offer a daily deal and leveraging the error in the model to explain part of the unobserved factors in modeling business performance.

Business reviews collected from Yelp have also been studied by Luca [16]. Luca evaluated the impact of Yelp reviews on restaurant's quarterly earnings in Seattle using the framework of regression discontinuity. Luca finds that the observed response to Yelp rating are consistent with bayesian learning. Under the bayesian hypothesis, reactions to signals are stronger when the signal is more precise (i.e., the Yelp average rating contains more information when the number of reviews is high). Moe precisely, a change in a restaurant's average rating has 50% more impact when the restaurant has at least 50 reviews (compared to a restaurant with fewer than 10 reviews).

Luca [16] also tests whether restaurants are gaming the rating system using the McCrary [17] test. The intuition of the test is as follows. Suppose that restaurants were gaming Yelp in a way that would bias the results. Then, one would expect to see a disproportionately large number of restaurants just above the rounding thresholds. Luca [16] finds that this is not the case. The results presented in [16] are related to our work in two ways. First, we use the concept of "Bayesian Learning" to derive statistically significant predictors of both daily deal adoption and business failure. Second, the McCrary test [17] suggests that the reviews on Yelp truly reflect the opinion of the Yelp community and are not being manipulated by the businesses on Yelp.

Pindyck and Rubinfeld [18] model the relation between private school attendance and voting for property tax increases that are used in part to finance public schools. In this application, the variables are whether children attend private school, number of years the family has been at the present residence , log of property tax , log of income and whether the head of the household voted for an increase in property taxes. Pindyck and Rubinfeld [18] wanted to test the hypothesis that parents of children who attended private school will have no incentive from an increase in property taxes that finance public schools and will vote against any such increases. Pindyck and Rubinfeld [18] model the bivariate outcomes of whether children attend private school and whether the head of the household voted for other covariates. They conclude that the two outcomes are independent and that the voting patterns of parents of children attending private schools do not differ from parents of children attending public schools. Our work is related to [18] in that we test for the independence of two binary outcomes; daily deal adoption and business failure.

3. DATA COLLECTION & ANALYSIS

Our dataset collection has two major components. First,

Business Category	Total No.	Closed (%)
RESTAURANTS & BARS	8490	17.7
SHOPPING	4961	10.6
FOOD	3259	14.9
BEAUTY & SPAS	2692	5.9
HEALTH & MEDICAL	2666	2.14
NIGHTLIFE	1851	16.0
ACTIVE LIFE	1301	6.3
ARTS & ENTERTAINMENT	1267	6.0
HAIR SALONS	948	5.3
FITNESS & INSTRUCTION	740	7.6

Table 2: Statistics of Yelp Businesses in Chicago as of July 2012

we collected data from Groupon as one of the top deal sites that offers daily deals in Chicago. Second, we collected data from Yelp for all the businesses in Chicago.

Groupon Data.

We used the Groupon data set compiled by Byers et al. [5] which includes the web links of $16,692$ deals offered by Groupon in 20 U.S. cities between January and July 2011. In this paper we focus only on the subset of Groupon deals offered in Chicago. We selected Chicago not just because it is the third largest city in U.S but also it the home town of Groupon. When the Groupon business is featured on *Yelp*, Groupon occasionally uses that information to promote the deal by including a link to the Yelp site as well as other information (e.g. star rating and selected customer reviews). However, in some cases Groupon does not mention the Yelp link on the deal page even if the business has a Yelp link.

We are interested in Yelp since it provides a wide range of information about the business. For Example, Yelp provides business location, number of reviews, date of review, star ratings, review text and other features such as alcohol license and price range. Moreover, Yelp indicates whether the business is still in operation or whether the business has closed by adding the string "CLOSED" next to the business name. Previous research has shown the potential of Yelp to indicate business parameters and performance [16] as we explained in section 2.

Groupon provides a convenient API [1] to collect information about the deals, however, we decided to develop our web crawler to extract features that are not supported by the API. For example, whether there is a link to Yelp or not. We initially had 1861 Groupon deals from chicago, with approximately 60% of them had their Yelp links listed. For the deals without Yelp links, we used the Yelp search feature to find a match for the Groupon business on Yelp. Specifically, we searched Yelp by the business name and the zip code listed on the deal webpage. Typically, Yelp return search results for relevant matches within the given zip code and other nearby zip codes. However, we report only the query results that exactly matched both the business name and its zip code. By the end of this matching process, we successfully associated 1184 Groupon deals with Yelp links, we call them *"GrouponDealsWithYelp"*. We also observed that some businesses offered multiple deals while others offered only one deal. While these deals are supposed to be all in Chicago, we found few cases where the deal zip code was outside Chicago (for other branches of the business in other states). Since we focus only on Chicago business we decided

---

[1]http://www.groupon.com/pages/api

to filter these cases. Finally, we developed our web crawler to extract the Yelp information for the businesses in the set *"GrouponDealsWithYelp"*. After filtering the businesses that had zero reviews (since they don't provide any information about the business), we observed 985 businesses with Groupon deals and Yelp links. We call this set *GrouponBusiness"*. Table 1 provides statistics of the set *"GrouponBusiness"* for the top 10 business categories.

Yelp Data.

We crawled the Yelp site to collect all the businesses that appear in Chicago (regardless whether they offer a deal or not). Yelp uses a structured format that arranges business names by alphabetical order. We initially had $38,000$ businesses listed in Yelp. After filtering all the cases that have zero reviews, we had $32,424$ with approximately 9% failed businesses (closed) and approximately 4% offered Groupon between January and July 2011. We refer to this set as our *Yelp Population* to represent the real population of businesses in Chicago. However, we expect Yelp data would represent certain business categories (e.g. restaurants) more than others (e.g. Insurance). Therefore, we decided to analyze only the top business categories. In Table 2 we provide statistics of the dataset we collected from Yelp for the top 10 business categories.

In the next sections, We proceeded to build a model for predicting failure using Yelp data. The details of the model are shown in Section 5.2. To build the bivariate model of Groupon adoption and business failure, we had to restrict our analysis to the businesses that did not offer a Groupon but were operating during the same period as the *"GrouponBusiness"*. We use the date of the last review posted for the business as a proxy for the closing date. We refer to the set that includes all of the businesses that did not offer Groupon deals and did not fail before January 2011 as *"nonGrouponBusiness"*.

Analysis.

Table 1 provides statistics of the set *"GrouponBusiness"* for the top 10 business categories. Also, Table 2 provides statistics of our data *Yelp Population* which includes both the two sets *"GrouponBusiness"* and *"nonGrouponBusiness"*. From Table 1 and Table 2, we observe some interesting patterns. First, we observe a difference in the ranks of the business categories between the total population of Yelp data and the *"GrouponBusiness"* except for the first category "Restaurants & Bars". We conjecture that some business categories could be popular on Yelp but they don't have the incentives to offer daily deals. We also observe that the highest percentage of closed businesses comes from the categories "Restaurants & Bars", "Food", and "Nightlife". In addition, we observed that some business in the set *"GrouponBusiness"* had the incentives to offer multiple daily deals during the six month period we analyzed.

As we discussed in section 2, previous research work emphasized the potential of Yelp as a proxy measure for business key performance indicators (e.g. survival, consumer appeal, and revenue) [16, 5, 6, 22]. Luca in [16] found that 69% of restaurants in Seattle are listed on Yelp. Also, Luca showed that changes in Yelp ratings are associated with changes in revenues [16]. These studies indicate the potential of Yelp data as representative of the true population. However, to the best of our knowledge, none of

this research analyzed Yelp as a source of business failure information. Therefore, to test the representativeness of Yelp data as a source of business failure information, we compared the number of closed businesses collected from Yelp to the number of bankrupted businesses as reported by bankruptcy filings of the Northern Illinois (which includes Chicago) open court records collected by the bankruptcy data project at Harvard [2]. Both the data from Yelp and court bankruptcy filings are between January 2006 and July 2012. Figure 1 shows the plot of the two normalized time series data . As shown in the plot, there is a strong correlation between the number of closed businesses computed from Yelp and the number of bankrupted businesses (correlation coeff. = 0.7164). Also, the two datasets have a similar trend as shown in Figure 1.
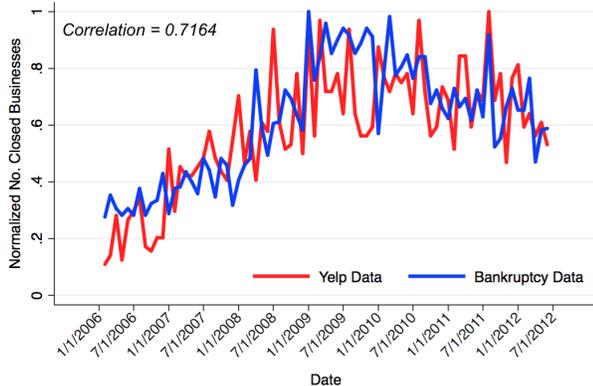

Figure 1: Normalized Yelp Closed Businesses versus Normalized Bankrupted Businesses between Jan 2006 and July 2012

4. ECONOMETRIC FRAMEWORK

In this section, we develop the econometric model. To motivate the need for an econometric model, consider the task of deciding whether running a daily deal increases the risk of business failure. We can model this as a regression problem of the form

$$failure_i = \beta * dailydeal_i + e_i \tag{1}$$

where "$failure_i$" is whether business "i" fails, , "$dailydeal_i$" indicates whether the business ran a daily deal and e_i is the error term and captures all the factors that are not included in the model. These factors are referred to as the "unobserved heterogeneity". The unobserved heterogeneity can include factors such as location specific risk, category risk .. .*etc.* It can also include factors that are correlated with the business decision to run a daily deal such as whether the business is struggling. Consider for example the class of struggling businesses that use daily deals as a last resort to attract more customers. In expectation, a struggling business will have a higher risk of failure. However, if we were to perform a regression analysis using Equation 1, we will over estimate the impact of daily deals (β) because a large number of the struggling businesses that used daily deals failed. While in fact, these businesses did not fail because of daily deals, they failed because of internal problems that happened to be correlated with the decision to offer a daily

---

[2]http://bdp.law.harvard.edu

deal. In this case, we have an *omitted variable bias* [10]. The omitted variable bias will only arise however, if the omitted variable (which in our case is whether the business was struggling), is correlated with one of the regressors (the decision to use daily deal).

The problem can also be viewed from the lens of endogeneity. A perquisite to any regression model is that the regressors are exogenous, [10] *i.e.,* the regressor variable comes from outside of the model and cannot be explained by any of the variables of the model. However, as we have seen in the above example, the regressor *dailydeal* was correlated with the errors. Consequently, the variable *dailydeal* is not exogenous as it it can be explained in part by the errors. The omitted variable bias is a form of endogeneity that results in a biased estimator.

If we wanted to estimate the impact of daily deals on the business survival, we will need to find an *instrument* [10]. An instrument is a variable that is correlated with the endogenous variable (dailydeal) but is not correlated with error (struggling business). For example, we could use the size of the daily deals providers sales force as an instrument since it is likely to be correlated with decision to adopt daily deals but unlikely to be correlated with the error term. A model with an *instrumental variable* is estimated using *2 stage least squares (2sls)*. In future work, we will address the use of instrumental variables approach.

However, in this paper, we don't directly attempt to model the impact of the daily deal on business survival. Instead, we focus on assessing the correlation between the unobserved factors that make a business offer a daily deal and the unobserved factors that influence failure. In that vein, we consider two different econometric models. In section 4.1, we use a bivariate probit model to check whether the errors in predicting if the business decision to run a daily deal are correlated with the errors in predicting the business failure. Since the errors model the unobserved heterogeneity, the bivariate model will help identify whether the unobserved factors that make a business run a daily deal are also correlated with the unobserved factors that contribute to the business failure. In section 4.2, we relax the conditions that the dependent variables (run a daily deal and failure) are binary and use the framework of seemingly unrelated regression " (SURE) to answer the same question addressed by the bivariate profit.

Our decision to use the SURE model is motivated by two factors. First, the bivariate probit model makes the explicit assumption that the error terms are jointly normal. While this might sound like a reasonable assumption, we have no way of checking its validity. On the other other hand, the SURE model does not make any assumptions about the joint distributions of the error, but rather uses the variance of these distributions to derive its estimate. Therefore, the use of two models provides us with a robustness test as in [19]. The second reason we consider the SURE model is that recent developments in econometrics [2] that used actual experimental data do not show any advantages of enforcing the limited range of dependent variables.

An important question is whether a single model best explains the two dependent outcomes (daily deals adoption and survival) or whether we need two separate models. In section 4.3, we describe a specification test based on the likelihood of the data and show how they relate to parametric technique, that test the null hypothesis of no correlation

between the error terms.

4.1 Bivariate Probit

Assume that we have a random sample of N observations where each observation is donated by i such that $i = 1, ..., N$. In ordinary regression models, we typically observe only one dependent variable for each observation $Y = (Y_1, ..., Y_N)$. However, in the general case general case, we can observe multiple dependent variables for each observation. Let Y_{ji} denote the response of the i^{th} observational unit for the j^{th} dependent variable. A typical situation is the case when we observe 2 variables, such that Y_{1i} and Y_{2i} are two binary dependent variables.

Traditional probit models can generally be described as latent variable models, in which we define a latent variable Y^* such that $Y = 1_{(Y^*>0)}$. In this section, we consider the bivariate probit model. The bivariate probit model belongs to the generalized class that is usually used to estimate several correlated binary variables jointly. These often represent two interrelated decisions, for example, to adopt two different, but related, policy initiatives. In the bivariate probit model, we have two separate probit models with correlated error terms. Specifically, we have two binary dependent variables for each i^{th} observational unit : $Y_{ji} = Y_{1i}, Y_{2i}$ such that $j = 1, 2$. Therefore, we have the following model

$$Y_{1i}^* = X_{1i}\beta_1 + \varepsilon_{1i} \qquad (2)$$
$$Y_{2i}^* = X_{2i}\beta_2 + \varepsilon_{2i} \qquad (3)$$

where Y_{1i}^* and Y_{2i}^* are the latent variables and they are related to the dependent variables by the following equation

$$Y_{1i} = 1_{(Y_{1i}^*>0)} \qquad (4)$$
$$Y_{2i} = 1_{(Y_{2i}^*>0)} \qquad (5)$$

The vectors X_{1i} denotes the $[N_1 \times 1]$ vector of exogenous regressor for dependent variable "1" . X_{2i} denotes the $[N_2 \times 1]$ vector of exogenous regressor for dependent variable "2" . In the bivariate probit model, the main assumption is that error terms $\varepsilon_{1i}, \varepsilon_{2i}$ are independent across observations "i" but may have cross-equation correlations. Therefore, we have $E[\varepsilon_{ji}\varepsilon_{jk}|x] = 0 \ \forall i \neq k$. In addition, the error terms are drawn from a bivariate normal distribution [10]

$$\begin{pmatrix} \varepsilon_{1i} \\ \varepsilon_{2i} \end{pmatrix} |X \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right) \qquad (6)$$

Where ρ is a correlation parameter denoting the extent to which the two ε's covary. The conditional expectation for the bivariate normal distribution is given by

$$E(\varepsilon_{2i}|\varepsilon_{1i} > z) = \rho \frac{\phi_1(z)}{\Phi_1(-z)} \qquad (7)$$

where Φ_1, ϕ_1 are the univariate normal cumulative distribution and density functions respectively. Equation 7 is the Inverse Mills ratio and has been used extensively in econometrics , [13] [10]. Equation 7 shows that the in case of bivariate probit, the errors are not independent, for example the error in estimating the probability of offering a Groupon deal corresponds in expectation to a large error in estimating the probability of business failures. The errors typically correspond to unobserved variables and by implication the Inverse Mills ratio indicates that these two variables move in synchronization. Fitting the bivariate probit model involves estimating the values of the parameters $\beta_1, \ \beta_2, and \ \rho$. We

use maximum likelihood estimation to estimate the parameters. The likelihood function L is defined as:

$$\begin{aligned} \mathcal{L} = \ & \prod_{i \in N} P(Y_{1i} = 1, Y_{2i} = 1)^{Y_{1i}Y_{2i}} \\ & P(Y_{1i} = 0, Y_{2i} = 1)^{(1-Y_{1i})Y_{2i}} \\ & P(Y_{1i} = 1, Y_{2i} = 0)^{Y_{1i}(1-Y_{2i})} \\ & P(Y_{1i} = 0, Y_{2i} = 0)^{(1-Y_{1i})(1-Y_{2i})} \qquad (8) \end{aligned}$$

Substituting the latent variables Y_1^* and Y_2^* in the Probability functions and taking the logarithm gives the log likelihood function LL:

$$\begin{aligned} \mathcal{LL} = \ & \sum_{i \in N} Y_{1i}Y_{2i} \ln P(\varepsilon_{1i} > -X_{1i}\beta_1, \varepsilon_{2i} > -X_{2i}\beta_2) \\ & + (1-Y_{1i})Y_{2i} \ln P(\varepsilon_{1i} < -X_{1i}\beta_1, \varepsilon_{2i} > -X_{2i}\beta_2) \\ & + Y_{1i}(1-Y_{2i}) \ln P(\varepsilon_{1i} > -X_{1i}\beta_1, \varepsilon_{2i} < -X_{2i}\beta_2) \\ & + (1-Y_{1i})(1-Y_{2i}) \ln P(\varepsilon_{1i} < -X_{1i}\beta_1, \varepsilon_{2i} < -X_{2i}\beta_2) \end{aligned}$$
$$(9)$$

After rearranging the terms, the log-likelihood function becomes:

$$\begin{aligned} \mathcal{LL} = \ = & \sum_{i \in N} Y_{1i}Y_{2i} \ln \Phi(X_{1i}\beta_1, X_{2i}\beta_2, \rho) \\ & + (1-Y_{1i})Y_2 \ln \Phi(-X_{1i}\beta_1, X_{2i}\beta_2, -\rho) \\ & + Y_{1i}(1-Y_{2i}) \ln \Phi(X_{1i}\beta_1, -X_{2i}\beta_2, -\rho) \\ & + (1-Y_{1i})(1-Y_{2i}) \ln \Phi(-X_{1i}\beta_1, -X_{2i}\beta_2, \rho) \end{aligned}$$
$$(10)$$

Note that Φ is the cumulative distribution function of the multivariate normal distribution|bivariate normal distribution and ϕ is the corresponding density function. Y_{1i} and Y_{2i} in the log-likelihood function are observed variables being equal to one or zero.

4.2 Seemingly Unrelated Regression

In the previous section, we enforced the constraint that the dependent variable have a limited range *limited dependent variables*. The bivariate probit frameworks explicitly uses bivariate normal distribution to model the joint distribution of the error terms. To test whether this assumption is indeed valid, we relax these constraints and use the framework of seemingly unrelated regression (SURE) proposed by Zellner [21]. Furthermore, there is mounting evidence in the econometrics literature [2] that argues in favor of using Ordinary least squares (OLS) even when the dependent variables are binary.

$$Y_{1i} = X_{1i}\beta_1 + \varepsilon_{1i} \qquad (11)$$
$$Y_{2i} = X_{2i}\beta_2 + \varepsilon_{2i} \qquad (12)$$

As before, the assumption of the model is that error terms $\varepsilon_{1i}, \varepsilon_{2i}$ are independent across observations "i" but may have cross-equation correlations. Therefore we have $E[\varepsilon_{ji}\varepsilon_{jk}|x] = 0 \ \forall i \neq k$ and

$$\begin{pmatrix} \varepsilon_{1i} \\ \varepsilon_{2i} \end{pmatrix} |X \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{11}^2 & \rho\sigma_{11}\sigma_{22} \\ \rho\sigma_{11}\sigma_{22} & \sigma_{22}^2 \end{pmatrix} \right) \qquad (13)$$

$$\begin{pmatrix} y_{11} \\ y_{12} \\ \vdots \\ y_{1N} \\ y_{21} \\ y_{22} \\ \vdots \\ y_{2N} \end{pmatrix} = \begin{pmatrix} X_{11} & 0 \\ X_{12} & 0 \\ \vdots & \vdots \\ X_{1N} & 0 \\ 0 & X_{21} \\ 0 & X_{22} \\ \vdots & \vdots \\ 0 & X_{2N} \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} + \begin{pmatrix} \varepsilon_{11} \\ \varepsilon_{12} \\ \vdots \\ \varepsilon_{1N} \\ \varepsilon_{21} \\ \varepsilon_{22} \\ \vdots \\ \varepsilon_{2N} \end{pmatrix} \quad (14)$$

Or in a more compact notation $\mathbf{Y} = \mathbf{X}\beta + \varepsilon$, $\quad \mathrm{E}[\varepsilon|X] = 0$, $\mathrm{Var}[\varepsilon|X] = \Omega$.

The SURE regression differs from the OLS in that the covariance matrix is not spherical *i.e.,* $\mathrm{Var}[\varepsilon|X] \neq \sigma^2 I_n$ where I_N is the identity matrix. In ordinary least squares, the Best Linear unbiased estimator for the parameters β is given by $\widehat{\beta}_{OLS} = (X'X)^{-1}X'y$. Once we have a non-spherical covariance matrix OLS is not efficient. To overcome this restriction, we use *feasible generalized least squares* which is a two stage estimator. In the first stage, we run ordinary least squares estimation assuming that the two equations are independent. The residuals from the OLS are used to estimate the elements of the covariance matrix $\hat{\sigma}_{ij} = \frac{1}{N}\hat{\varepsilon}_\mathbf{i}^T \hat{\varepsilon}_\mathbf{j}$. In the second stage, we run *weighted least squares* using the previously estimated covariance matrix. Feasible generalized least squares method estimates β by minimizing the squared Mahalanobis length of the residual vector $\hat{\beta}_{FGLS} = \arg\min_b (Y - Xb)' \Omega^{-1} (Y - Xb)$ (note the in case of ordinary least squares Ω is diagonal and therefore $\hat{\beta}_{OLS} = \arg\min_b (Y - Xb)'(Y - Xb)$). The explicit form of the estimator is given by

$$\hat{\beta}_{FGLS} = (X'\Omega^{-1}X)^{-1}X'\Omega^{-1}Y. \quad (15)$$

To test whether the two equations are best modeled using a SURE or can be modeled using ordinary least squares, it suffices to test whether the errors $\varepsilon_{1t}, \varepsilon_{2t}$ are correlated. The Breusch-Pagan test [4] which is widely used in detecting heteroskedasticity can be applied.

4.3 Specification Testing

We use the Akaike Information Criteria (AIC) [10] to test whether the data is best described by two separate models or a single joint model. In the case of the bivariate probit model, the joint model will have one additional parameter (the correlation coefficient). We therefore compare the log likelihood of the joint models in Equation 10 to the log likelihood of the separate equations and test whether the difference is greater than "1.0". In the case of SURE model, we replace the log likelihood with the the sum of squared residuals [10]. For the parametric approach, we employ the Cramer-Rao bound to compute parameter's mean and variance from the maximum likelihood estimator [10]. We use a "t-test" to test whether the parameters including the correlation coefficient ρ are different from "0". In the SURE setting, we use the Breusch -Pagan framework to test the correlation coefficient.

5. EXPERIMENTS AND RESULTS

In this section, we present the experimental setup and results of several regression models both in a univariate framework and a multiple equation framework. In the univariate framework, we use *a probit model* to analyze the significance of the different factors that may impact the business survival. Similarly, we use *a probit model* to analyze the significance of the different factors that influence the business decision to make a daily deal in Groupon. In the multiple equations framework, we use bivariate probit and seemingly unrelated regression frameworks to jointly model the business survival and the business daily deal decision. Here, our goal is to specifically investigate whether the unobserved factors that make the business offer a daily deal are correlated with the unobserved factors that impact the business survival. We conduct several statistical tests to test if the data is best described by a joint model and whether the correlation between the unobserved factors is significant. Our results indicate that joint models fit the data better than single univariate models and there is strong significant correlation for the two business categories restaurants and spas.

5.1 Sampling

When modeling business survival and Groupon decision we need to take into account the relative frequency of both events. On one hand, the probability of a business failure in any given year is fairly low (in order of 8%). On the other hand, daily deals is a relatively new phenomena which is still being evaluated, the number of businesses that leverage a daily deal are fairly low compared to the businesses that don't offer daily deals. In that vein, we are attempting to model two rare events : daily deal decision and business failure. Statistical models tend to underestimate the probability of rare events[10]. Since the vast majority of the businesses will be non-groupon and non-closed, the model will assign a large negative constant to the two equations that model the daily deal and failure models. The large constants will make the error terms small and can impact the model's ability to detect the correlation between the error terms in the two equations. To address this problem, we first apply the Bivariate Probit and the SURE to the entire data set. We then restrict the models such that there is no constant terms. We find that in the first case (unrestricted full sample), a positive but not statistically significant correlation between the error terms. In case of the full sample but no-constant , we find that there is a strong positive correlation between the error terms. This confirms our intuition about the inability of the full data model to capture the two rare events. Therefore, we randomly sampled from the population of "nonGrouponBusiness" to account for the sparsity problems and selection biases that can be caused by data collection.

5.2 Univariate Probit Models

A *probit model* is a type of regression used to model a binary dependent variable. Here, we define two binary dependent variables. First, we define the variable *isClosed* to represent whether the business has failed (*isClosed*=1) or the business is still operating (*isClosed*=0). Second, we define the variable *isGroupon* to represent whether the business has made at least one daily deal (*isGroupon*=1) at Groupon or the business did not make any deals (*isGroupon*=0). We develop a business survival model and Groupon decision model to analyze the factors that influence *isClosed* and *isGroupon* respectively.

Business Survival Model.

Symbol	Variable	Description
isClosed	ISCLOSED	A binary outcome indicating whether the business failed (isClosed =1) or operating (isClosed = 0)
isGroupon	ISGROUPON	A binary outcome indicating whether the business made a deal (isGroupon =1) or no deals (isGroupon = 0)
fzrisk	FAIL ZIPRISK	The percentage of failed businesses in the same zip code
fprisk	FAIL PRICERISK	The percentage of failed businesses in the same price category
gzrisk	GROUPON ZIPRISK	The percentage of businesses that made Groupon deals in the same zip code
gprisk	GROUPON PRICERISK	The percentage of businesses that made Groupon deals in the same price category
rate	RATING	Average Yelp Rating
nreview	REVIEWS	Number of Yelp Reviews
price	PRICE CATEGORY	{1,2,3,4} From cheap to most expensive

Table 3: Description of Variables

Table 4: Yelp Population Data

Dependent Variable: isClosed		AUC = 0.674
Variable	**Coefficient**	**(Std. Err.)**
Fail Ziprisk	6.786**	(0.317)
Fail Pricerisk	7.039**	(0.380)
Rating × Reviews	-0.003**	(0.000)
Reviews	0.008**	(0.002)
Rating	-0.061**	(0.010)
Intercept	-2.287**	(0.060)
Significance levels :	† : 10% * : 5%	** : 1%

Table 5: Yelp Population Data

Dependent Variable: isGroupon		AUC = 0.894
Variable	**Coefficient**	**(Std. Err.)**
Groupon Ziprisk	4.521**	(0.136)
Groupon Pricerisk	10.771**	(0.888)
Rating × Reviews	-0.001**	(0.000)
Reviews	0.006**	(0.001)
Rating	0.057**	(0.021)
Intercept	-2.848**	(0.093)
Significance levels :	† : 10% * : 5%	** : 1%

We model the business survival variable *isClosed* as a probit function of a number of business variables we collected from Yelp (refer to Table 3 for notations):

$$isClosed = f(rate \times nreview, rate, nreview, fzrisk, fprisk) \quad (16)$$

We tried a number of other specifications and selected the specification of Equation 16 based on the AIC. Table 4 shows the results of the model trained on Yelp population data we collected from Chicago. Although the model is simple, it fits the data well (AUC=0.674) and all the variables are statistically significant ($p-value < 0.01$).

We can gain further insight by examining the marginal contributions of each of the factors to the probability of survival. Therefore, we make the following observations from the results:

- When the average Yelp rating is higher, the risk of failure gets lower. Higher rated business tend to be more successful.

- When the number of reviews is higher, the risk of failure increases. A business with high number of reviews has been around for a longer time and its risk increases with time.

- When the average Yelp rating weighted by the number of reviews $rate \times nreview$ increases, the risk of failure gets lower. This is consistent with the theory of Bayesian learning presented in [16]. This is because the average Yelp rating weighted by the number of reviews gives more precise information compared to rating or number of reviews only.

- The business location makes a difference, some zip codes are riskier than others. This is consistent with previous work on restaurant failure by zip code.

- The price risk computed for the business price category is also significant similar to the zip code risk.

Groupon Adoption Model.

Similar to the survival model, we model the Groupon decision variable *isGroupon* as a probit function of a number of business variables from Yelp (as in Table 3):

$$isGroupon = f(rate \times nreview, rate, nreview, gzrisk, gprisk) \quad (17)$$

We selected the specification of Equation 17 based on the AIC . Table 5 shows the results of the model trained on Yelp population data from Chicago. The model has a good accuracy of AUC = 0.894, that shows the ability of the model to predict whether a business will decide to make a daily deal based on some business parameters collected from Yelp. Also, the regressors are statistically significant ($p-value < 0.01$). We can gain further insight by examining the marginal contributions of each of the factors to the probability of offering a daily deal. Therefore, we make the following observations from the results:

- When the average Yelp rating is higher, the probability of daily deal increases. We conjecture that this is due in part to how daily deal sales force selects the business to approach.

- When the number of reviews is higher, the probability of daily deal increases. We conjecture that this is also due in part to how daily deal sales force selects the business to approach.

- When the average Yelp rating weighted by the number of reviews $rate \times nreview$ increases, the probability of daily deal gets lower. This is an indication of a more

successful business that would not need to make a deal at Groupon, especially that Groupon takes 50% of the deal revenue [3]. This results is also consistent with the theory of Bayesian learning presented in [16].

- The business location in terms of zip code makes a difference, some zip codes are more likely to offer daily deals. The higher the number of businesses that make a deal in the same zip code, the higher the chance of a business to adopt a deal. This is part due to completive pressure and in part due to how daily deal sales team targets geographical areas.

- The business price category is also significant as we mentioned before.

5.3 Joint Model

We jointly modeled the Groupon and failure models using Equations 17 and 16 using a bivariate probit model. Table 6 shows that the data is best modeled by a bivariate probit since the log likelihood of the bivariate model differs from that of the two separate equations by more than "1.0". In addition, Tables 7 and 8 show that for the two largest daily deal categories, the correlation between the unobserved factors "ρ" is positive and significant (0.281 in Restaurants/Bars, 0.24 in Beauty/Spas). As a robustness test, we used the SURE model to compute the correlation between the unobserved errors and tested its significance using the Breusch-Pagan test (correlation=0.054 in Restaurants/Bars, 0.060 in Beauty/Spas). It should be noted that the bivariate probit model operates in the space of latent variables that can be in range $[-\infty, 0]$ when the dependent variable is "0" and in the range of $[0, \infty]$ when the dependent variable is "1". The residuals are therefore computed in that space and can assume large values. On the other hand, the SURE model operates directly in the space of dependent variable that assume range between "0" and "1". Therefore the residuals are smaller in case of SURE. This explains why the correlation coefficient in case of bivariate probit differs from that in case of SURE model.

Previous work has shown that restaurants with their high marginal cost, low fixed cost and inability to schedule the arrival of daily deal customers [11] are not well suited to daily deals. On the other hand, Spas with their low marginal cost, high fixed cost and ability to schedule daily deal customers are better suited for daily deals. Conversely, a daily deal is a more desperate measure for a restaurant than a Spa. This is validated by our results, if a restaurant offers a daily deal with out having any strong reason to do so, its probability of failing increases more than a Spa that had some motivation for offering a daily deal. Similar to what we did in the univariate case, we analyzed the contributions of factors, and we have the following observations:

- When the ratings weighted by the number of reviews $rate \times nreview$ increases, the risk of failure gets lower. Also, the probability of offering a daily deal gets lower.

- Unlike the univariate case, the ratings $rate$ is not a significant factor. However, the weighted ratings are significant. This is consistent with the theory of Bayesian learning presented in [16].

- As in the univariate case, when the number of reviews is higher, the risk of failure increases.

Restaurants and Bars		
Model Name	**No. Params**	**AIC**
isGroupon+isClosed	10	1976.46
Bivariate Probit	11	**1968.33**
Beauty and Spas		
Model Name	**No. Params**	**AIC**
isGroupon+isClosed	10	1029.97
Bivariate Probit	11	**1028.89**

Table 6: AIC for univariate versus bivarite probit

We have also tried other business categories, for example, Health & Medical, Active Life,..*etc*. However, there was a lack of significance of correlation between the unobserved factors for these categories. We conjecture the lack of significance is because we did not have enough samples in those categories. In future work, we aim to extend our study to daily deal data from Yipit [1] as well as collecting data from other sources such as city department of revenue to gain a better access to businesses that are not represented well by our current data.

6. CONCLUSION AND FUTURE WORK

In this work, we studied whether daily deal adoption signals additional information about the business. For the two largest daily deal categories restaurants and spas, we found that the unobserved factors that contribute to a business decision to offer a daily are positively correlated with unobserved factors that contribute to a business failure. Restaurants had a higher correlation while spas had a lower correlation. These results indicate that daily deal provide a strong signal of business survival for restaurants and to a lesser extent for spas. Our results also show that social media sites such as Yelp provide a rich set of information that can be used to model business. In particular, we found that consistent with Bayesian learning theory, the rating of business weighted by the number of reviews provides a statistically significant predictor of business failure.*Ceteris paribus*, a business with a high number of positive reviews has higher odds of survival.

In future, we plan to extend our work in multiple directions. The first direction is to consider other business categories. In this work, we were limited to the daily deal data from [5]. Second, we plan to consider daily deal providers other than Groupon. In the case of Groupon, we modeled the marketers decision to use daily deal as a binary decision. We also modeled the business failure as a binary decision. This allowed us to leverage the framework of bivariate probit to jointly model daily deal and business failure. In the more general setting, the marketer can choose to offer a daily deal through a number of providers or not to offer a deal. In the general case, the marketer's choice is best modeled as a multinomial. We plan to investigate techniques from Multilevel Multiprocess Models (MLMP) [15]. The data needed to undertake these two directions can be obtained from daily deal aggregators such as Yipit [1]

We also plan to address the dynamics of Groupon adoption and business failure. In particular, we plan to test whether daily deal adoption is stationary, whether business failures are stationary and whether the two time series are co-integrated [12]. This will allow us to investigate whether there is a long term equilibrium between daily deal adoption and business failure.

Table 7: Restaurants and Bars: Bivariate Probit

Bivariate Probit		
Variable	**Coefficient**	**(Std. Err.)**
Equation 1 : isGroupon		
Groupon Pricerisk	13.978**	(2.122)
Groupon Ziprisk	5.715**	(0.535)
Rating × Reviews	-0.002**	(0.001)
Reviews	0.010**	(0.002)
Rating	0.180*	(0.073)
Intercept	-3.385**	(0.293)
Equation 2 : isClosed		
Fail Ziprisk	8.791**	(1.021)
Fail Pricerisk	12.928**	(3.541)
Rating × Reviews	-0.002*	(0.001)
Reviews	$0.007^{\dagger}$	(0.004)
Rating	-0.013	(0.070)
Intercept	-4.202**	(0.495)
Equation 3 : Joint		
athrho	**0.288****	(0.09161)
rho	**0.281**	(.08440)
SURE: Breusch-Pagan test of independence		
Correlation	**0.0543****	
Significance levels : $\dagger$: 10% $*$: 5% $**$: 1%		

Table 8: Beauty and Spas: Bivariate Probit

Bivariate Probit		
Variable	**Coefficient**	**(Std. Err.)**
Equation 1 : isGroupon		
Groupon Pricerisk	17.203**	(3.326)
Groupon Ziprisk	4.460**	(0.373)
Rating × Reviews	-0.004**	(0.001)
Reviews	0.024**	(0.006)
Rating	-0.087	(0.059)
Intercept	-2.613**	(0.288)
Equation 2 : isClosed		
Fail Ziprisk	5.199**	(1.556)
Fail Pricerisk	8.906*	(4.538)
Rating × Reviews	-0.007*	(0.003)
Reviews	0.026**	(0.010)
Rating	0.103	(0.090)
Intercept	-4.201**	(0.719)
Equation 3 : Joint		
athrho	$\mathbf{0.25^{\dagger}}$	(0.1433)
rho	**0.246**	(0.1346)
SURE: Breusch-Pagan test of independence		
Correlation	**0.060****	
Significance levels : $\dagger$: 10% $*$: 5% $**$: 1%		

Last but not least, we plan to investigate the causal impact of Groupon on metrics other than business failure. We plan to investigate the Groupon impact on sales, number of orders and order size. This will help us better understand how the daily deals impact revenue: a) do they impact revenue through a change in the number of orders. or b) do daily deals have a stronger impact on order size or c) do daily deals have an impact on both number and size of orders. We will also look at using instruments [10] to get an estimate of the casual impact of Groupon. We are investigating Groupon sales force as a potential instrument.

7. REFERENCES

[1] http://www.yipit.com/.

[2] J. D. Angrist and J.-S. Pischke. *Mostly Harmless Econometrics: An Empiricists Companion.* Princeton University Press, 2008.

[3] A. Arabshahi. Undressing groupon: An analysis of the groupon business model, 2010.

[4] T. Breusch and A. Pagan. Simple test for heteroscedasticity and random coefficient variation. *Econometrica*, 47(5), 1979.

[5] J. Byers, M. Mitzenmacher, and G. Zervas. Daily deals: Prediction, social diffusion, and reputational ramifications. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 543–552. ACM, 2012.

[6] J. W. Byers, M. Mitzenmacher, and G. Zervas. The groupon effect on yelp ratings: a root cause analysis. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 248–265, New York, NY, USA, 2012. ACM.

[7] U. Dholakia. How effective are groupon promotions for businesses? 2010.

[8] U. Dholakia. How businesses fare with daily deals: A multi-site analysis of groupon, livingsocial, opentable, travelzoo, and buywithme promotions. 2011.

[9] B. Edelman, S. Jaffe, and S. Kominers. To groupon or not to groupon: The profitability of deep discounts. 2010.

[10] W. H. Green. *Econometric Analysis.* Pearson, 6 edition, 2008.

[11] S. Gupta, T. Keinungham, R. Weaver, and L. Williams. Are daily deals good for merchants?

[12] J. D. Hamilton. *Time Series Analysis.* Princeton University Press, 1994.

[13] J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–162, 1979.

[14] G. Z. John Byers, Michael Mitzenmacher. Daily deals: Prediction, social diffusion, and reputational ramifications. In *5th ACM International Conference on Web Search and Data Mining (WSDM)*, 5th ACM International Conference on Web Search and Data Mining (WSDM).

[15] L. Lillard and L. Waite. Til death do us part: marital disruption and mortality. *The American Journal of Sociology*, 100, 1995.

[16] M. Luca. Reviews, reputation, and revenue: The case of yelp.com. Technical Report 12-016, Harvard Business School Working Paper, 2011.

[17] J. McCrary. Manipulation of the running variable in the regression discontinuity design. *Journal of Econometrics*, 2008.

[18] R. Pindyck and D. Rubinfeld. *Microeconomics.* Prentice Hall, 2008.

[19] K. Tucker and A. Goldfarb. Online display advertising: Targeting and intrusiveness. *Marketing Science*, 30(3):389–404, 2011.

[20] M. Ye, T. Sandholm, C. Wang, C. Aperjis, and B. A. Huberman. Collective attention and the dynamics of

group deals. In *Proceedings of the 21st international conference companion on World Wide Web*, WWW '12 Companion, pages 1205–1212, New York, NY, USA, 2012. ACM.

[21] A. Zelner. An efficient method of estimating seemingly unrelated regression equations and tests for aggregation bias. *Journal of the American Statistical Association*, 57:348–368, 1962.

[22] G. Zervas. The evolving quality of groupon deals (draft). 2012.