Groupon Fatigue: Search & Learning in a Daily Deals Site *

Mantian (Mandy) Hu † Chu (Ivy) Dang ‡ Pradeep K. Chintagunta §

August 20, 2017

*The authors gratefully acknowledge the helpful comments from Andrew Ching, Tulin Erdem, Ali Hortaçsu, Jun Kim, Xing Li, Tesary Lin, Nathan Olivia, Stephan Seiler, S. Sriram, Raluca Ursu; attendees of the 2016&2017 Marketing Science Conference, 2016&2017 Marketing Dynamics Conference; seminar participants at Guanghua School of Management, Peking University and CUHK Business School. Chintagunta thanks the Kilts Center for Marketing at the Booth School of Business, The University of Chicago for financial support. This work was also supported by grants from the Research Grants Council of the Hong Kong SAR (Ref. No. CUHK24500214). The usual disclaimer applies.

†Assistant Professor of Marketing, CUHK Business School, The Chinese University of Hong Kong, mandyhu@baf.cuhk.edu.hk

‡Ph.D. Candidate, CUHK Business School, The Chinese University of Hong Kong, dangchu@baf.cuhk.edu.hk

§Joseph T. and Bernice S. Lewis Distinguished Service Professor of Marketing, Booth School of Business, The University of Chicago, pradeep.chintagunta@chicagobooth.edu

Abstract

Daily deal "fatigue" refers to a phenomenon in which consumers get tired of deals from daily deal sites (e.g., Groupon). Such fatigue, if it exists, poses a challenge to the daily deal industry. We study purchase behavior of daily deals with individual clickstream data on browsing sessions of consumers who newly subscribe to Groupon between January and March 2011. The data reveal two patterns. First, consistent with the notion of "fatigue," the probability of a consumer clicking on a merchant in the emailed newsletter declines over time. Second, the probability that the consumer makes a purchase conditional on clicking increases over time. Our objective is to propose a model that rationalizes these patterns and to then provide insights for companies to deal with "fatigue".

When consumers first subscribe to a daily deal site, they are unlikely to be fully informed about the quality of the deals offered on the site. The daily newsletter provides price and some limited information about that day's featured deals. To learn more about quality, consumers need to click on the newsletter, go to the deal's website and invest time and effort to learn about quality. Such a search for information is costly and only provides a noisy signal of quality to consumers. Further, consumers do not know the future deals they may receive. Given the uncertainty about quality and the cost of searching, consumers are more likely to search early on (i.e., click on the newsletter). As they learn about deal quality, they need to search less; resulting in clicks declining over time. As learning accumulates, consumers are better at recognizing deals of higher quality to click on and they are more likely to click in order to purchase rather than to learn about quality. This results in an increase in the conditional probability of purchasing. We formulate a dynamic model of search and learning based on the above characterization of consumer behavior. We show that the model is able to replicate the patterns in the data. Next, we estimate the parameters of the model and provide managerial insights to daily deal websites based on our findings and policy simulations.

Keywords: Learning model; Dirichlet updating; Dynamic search model; Daily deal fatigue

1 Introduction

Daily deal websites, which offer discounts on goods and services for a given time period (typically 24 to 36 hours), have attracted considerable attention in the global marketplace in recent years. In 2015, the revenues of daily deal sites in the US reached $4 billion, up from $2 billion in 2011, with 329 businesses in this line of work and employing 20,757 people (IBISWorlds Daily Deals Sites market research report) [1]. Roughly 4 in 10 American adults use online deals according to a survey by YouGov in 2014. Groupon is currently the dominant player in the daily deal industry. By the third quarter of 2016, it had about 50.8 million active users worldwide. It works with over a million merchants [2] generating about $3.1 billion revenue by the end of 2015 [3]. However, despite the fast growth in the early years, with the passage of time, the industry has started to suffer from what the media calls "Daily Deal Fatigue" or "Groupon Fatigue" (Dholakia and Kimes, 2011) [4]. "Fatigue" is the phenomenon in which consumers get tired of deals from these daily deal sites. Such a concern, if valid, challenges the very sustainability of the business model. The concern is so widespread that some have even predicted the death of daily deals and daily deal websites.[5] The issue of fatigue however, goes beyond the daily deal phenomenon and has been observed in other contexts as well. Indeed, businesses that are built based on a similar model e.g., online dating websites (for example, Tinder and OKCupid) and flash sales websites (for example, Gilt) seem to suffer from the same 'issue.[6] [7]

To fix ideas, we first describe how daily deal websites, including Groupon, operate. Upon registration, new subscribers receive an emailed newsletter everyday from the website (as shown in Figure 1 (a)). One of several featured deals of the day for a particular city would be highlighted in the email, with a captioned subject line and more information in the email body. The deals provided usually have a short expiry window (typically ranging from 1 to 3 days). So subscribers face a trade-off between purchasing the current deal or waiting for another ("better") one. A consumer receiving the newsletter has to decide whether or not (s)he wants to know more about the deal by clicking on the newsletter and getting redirected to the webpage of the deal (as show in Figure 1 (b)). Here

[1] http://www.ibisworld.com/industry/daily-deals-sites.html
[2] https://www.statista.com/statistics/273245/cumulative-active-customers-of-groupon/
[3] https://biz.yahoo.com/e/160212/grpn10-k.html
[4] The terms have been used in articles in several major mass media such New York Times, CNN, Forbes, and etc.
[5] http://business.time.com/2012/11/16/is-the-daily-deal-model-dying-a-slow-death/
[6] https://www.theatlantic.com/health/archive/2016/10/the-unbearable-exhaustion-of-dating-apps/505184/
[7] https://www.cnbc.com/id/47507737

the consumer receives more information regarding the deal itself such as the number of people who have bought the deal already, the time left before it expires, whether the deal had "tipped" yet, etc.[8] The information on the webpage helps the consumer assess the "quality" of the deal. The consumer now has to decide whether to purchase the deal or not. While new subscribers may initially not know what to expect in terms of the qualities of the merchants offering the deal, over time they can use the visits to the deal webpages to better understand the (distribution of) qualities of the various merchants on the site. An important feature of this process is that in order to obtain information about the deal, the consumer needs to expend some effort in order to visit the webpage of the deal. In this context, "fatigue" would refer to a decline in clicks and visitation of the webpages associated with the deals featured in the newsletter.

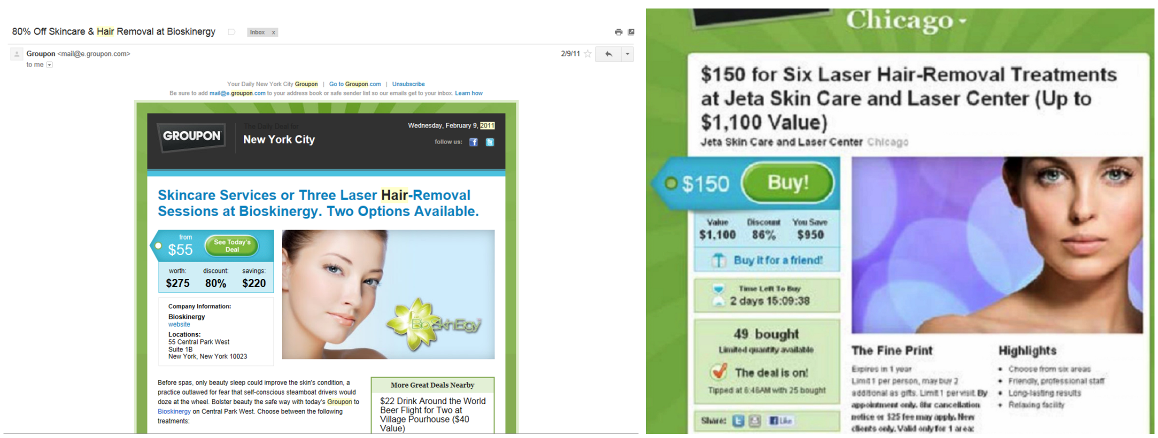

(a) Groupon Newsletter in early 2011 (b) Groupon Webpage in early 2011

Figure 1: Example of Groupon Newsletter and Webpage in Early 2011

If the concern regarding fatigue associated with the daily deals is indeed valid, the industry would need to formulate ways to overcome this issue. At the same time, establishing the validity of the fatigue phenomenon is itself important from both academic and practical perspectives. After all, how could consumers get tired of deals given that marketers consistently use promotions to acquire and retain customers? Specifically, understanding how consumers behave when it comes to such promotions could help to shed light on whether the industry is indeed headed for doom or whether what observers are seeing is the natural evolution of customer behavior in the category. Despite its importance for the very survival of the industry, daily deal fatigue has not been investigated much in the literature so far. Dholakia and Kimes (2011), for example, surveyed consumer perceptions on

[8] "Tipping" means that enough people have purchased the coupon so the merchant will actually run the promotion - this ensures that the merchant generates enough revenues from the entire promotion to justify running it.

daily deals and found no evidence of daily deal fatigue. Thus our main objective in this paper is to provide some empirical insight into how acquired customers behave on daily deal sites and to examine the empirical validity of the concern regarding deal fatigue among those who have subscribed to the site. Based on what we observe in the data, we then propose an empirical model to characterize such behavior. This model will then help us to examine the consequence of eliminating the effort that the consumer needs to expend in gathering information about the featured deal (i.e., visiting the deal webpage to obtain information) and the approach daily deal websites may take to encourage consumers to purchase deals. Deal fatigue, as we look at in this paper is a "within" consumer phenomenon over time associated with subscribers to the site. Alternatively, fatigue can be viewed as an "across" consumer phenomenon, i.e., that consumers acquired more recently are less likely to click (and purchase) on the site. As we describe below, since our data come from a relatively short period in time, we are not able to provide any insights into this alternative characterization of fatigue although we acknowledge that this could be a critical issue for the companies.

We obtained our data from a third party online research firm in the US. This proprietary dataset consists of the complete clickstream within the browsing sessions of people who newly subscribed to the Groupon service between January and March 2011. Thus we observe consumers from the beginning of their associations with Groupon. The data reveal two patterns. First, as noted in the popular press, we find that the probability that a consumer clicks on a merchant in the emailed newsletter indeed declines over time. This is consistent with the notion of fatigue. By itself, this finding does not bode well for Groupon. After all, if there is less and less interest in the service over time then this could adversely impact the revenues the site can expect to generate from its subscriber base. The second pattern revealed by the data however, is a potential source of optimism for Groupon. In particular, we find that the probability that the consumer makes a purchase, conditional on clicking, is actually increasing over time. Together, these patterns suggest that even though the consumer is getting more selective in terms of exploring the offers received (s)he is more likely to yield the site revenue as time passes.

Our proposed structural model of search and learning tries to explain the empirical patterns of the within-consumer, over-time decline in clickthroughs to the webpages of the deals as well as the increased probability of purchasing on the site. To allow for the feature of the daily deal sites that require consumers to expend effort to obtain information on the deal and to purchase it, we postulate that consumers incur a "search cost". Unlike traditional search models (Weitzman,

1979) where consumers actively seek out information on a number of alternatives in order to resolve their uncertainties regarding these alternatives in order to make a choice, in our case the search is "passive" in that deals arrive on a daily basis to the consumer with the latter merely deciding whether or not to expend the cost associated with obtaining the information on that particular deal. While this search cost enables us to explain whether or not the consumer clicks on the newsletter to go to the deal's webpage, it does not explain the decline in clicks over time. To rationalize this behavior, we postulate that new subscribers to a site are uncertain about the nature of deals available to them. By clicking on the newsletter and visiting the webpage of the deal, they are able to (partially) resolve this uncertainty about the "quality" of deals on the site. Over time, they learn about the deals that the site offers; as we describe later we operationalize the learning in terms of the distribution of deal qualities on the website. This resolution of uncertainty over time combined with the cost of obtaining the deal information results in the declining pattern of clicks (or "deal fatigue"). Since they are also now better informed about the deals as time progresses, consumers are better at choosing the deal with better quality to click. So instead of clicking to learn about the quality, they click to purchase thereby contributing to the increased probability of making a purchase on the site conditional on click. The "waiting to learn" motivation makes our model inherently forward-looking in nature.

Based on the model estimates, we are able to replicate the two patterns observed in the data. Consumers who are uncertain about the quality of future deals face the tradeoff between whether to click on the deal and purchase it now or to wait for the next one. Consumers tend to click on more deals in the beginning to learn about the quality distribution of deals. As their knowledge accumulates, the incentive to learn declines so they click fewer deals. Over time, since clicking is less about learning the quality of deals, the motivation for clicking is more likely to be to actually purchase the deal resulting in the data pattern of increasing purchase probabilities. Learning is a key driving force generating the data patterns in our setup; the model without learning fails to replicate these patterns. With the model in place, we then turn to counterfactuals that may be of interest to the firm. Specifically, the daily deal websites can generate more revenue by doing the following two things: reducing search costs, and increasing the variance in deal qualities to prolong the learning process.

The rest of this paper is organized as follows: In Section 2, we review the literature. Section 3 describes the data and empirical data analysis. Section 4 presents the model. The empirical application is in Section 5. Section 6 is the managerial implications. And in Section 7, we conclude and discuss future extensions of our model.

2 Literature Review

Our research mainly relates to three streams of literature - on search, on learning and on daily deals. We summarize them in Table 9. In the literature on search, one area focuses on understanding consumers' search behaviors: whether they engage in sequential search (Stigler, 1961; McCall, 1970; Weitzman, 1979; Kim et al., 2010; Chen and Yao, 2015) or in simultaneous search (Honka, 2014). In addition, research like De los Santos et al. (2012) and Honka and Chintagunta (2015) also propose identification strategies to distinguish between the two types of search. Another area focuses on search costs: how to identify them and how to quantify their impact on consumer demand (Hortaçsu and Syverson, 2004; Hong and Shum, 2006; Koulayev, 2010; Seiler, 2013). While studies such as those by Rust (1987), Erdem et al. (2003), Song and Chintagunta (2003) and Hendel and Nevo (2006) utilize dynamic structural models to account for consumer forward-looking behavior, search is not considered, so consumers only need to make a purchase decision by trading off making a purchase this period and waiting to obtain a higher utility in the next period. Seiler (2013) considers consumers' forward-looking behavior under a two-stage decision process: first, the consumer decides whether to search for price; if yes, the consumer decides wether to purchase. At the search stage, consumers look ahead to the purchase stage; at the purchase-stage, they look ahead to the next period's search stage. Our model takes Seiler's model as a benchmark, since we also explicitly model consumers' joint click (or "passive search") and purchase decisions.

Our research also relates to the Bayesian learning models in the discrete-choice demand literature (c.f. Erdem and Keane (1996), Crawford and Shum (2005)). In those studies, consumers are uncertain about some aspect of their demand functions that they resolve over time via signals they obtain. Such learning has been modeled in a Bayesian fashion using a variety of distributional assumptions. Perhaps the most widely used is the parametric normal Bayesian updating rule (c.f. Erdem and Keane (1996), Crawford and Shum (2005), Zhang (2010)). For example, Zhao et al. (2013) and Ching and Lim (2016) incorporate learning into consumer choice models, but consumers are not assumed to be forward looking. Erdem and Keane (1996), Crawford and Shum (2005) and Zhang (2010) embed consumer learning within dynamic discrete-choice models. Rothschild (1974) theoretically studies the use of a non-parametric Bayesian method to study optimal search under two types of uncertainty: the uncertainty about current prices and the uncertainty about the underlying process that generates prices. Specifically, he utilizes Dirichlet priors to model a searcher's belief about the unknown price distribution, and updates these priors in a non-parametric Bayesian fashion as new price quotes arrive. Koulayev (2013) was the first to bring the search model with Dirichlet priors into the empirical domain. He developed a novel characterization of optimal search that leads to closed-form, easily computable, ex-ante probabilities of purchasing products. De los Santos et al. (2013) took a step further by modeling non-parametric learning using more general Dirichlet process priors. Taking advantage of an individual-level dataset which contains information on consumers' search sequences, the study develops a method to estimate consumer search costs for differentiated products when consumers have only partial information about the distribution that they sample from.

In this study, we embed consumer learning within a two-stage dynamic search model. The two stages of the model involve the (passive) search decision and the purchase decision. Consumers are uncertain about the distribution of qualities of the merchants offering the deals on Groupon and learn about this distribution by seaching the various deals that are offered. We use a non-parametric Dirichlet updating rule by which consumers update their beliefs about the quality distribution. The Dirichlet prior places fewer restrictions than the typically assumed normal distribution; the updating rule is simple to implement and quick to compute; and in our context, consumers learn about the *distribution* of deal quality, rather than the "true" quality of the product since in this case the "product" changes over time as the same deal does not repeat in the data.

Our research also contributes to the literature on daily deals. Dholakia and his colleagues wrote a series of reports on the daily deal industry (c.f. Dholakia (2010), Dholakia (2011), Dholakia and Kimes (2011), and Dholakia (2012)). Dholakia (2010) and Dholakia (2012) surveyed the business side of the market and found no evidence of deterioration for small and medium-sized business in the performance of daily deal promotions over the surveyed year. Dholakia and Kimes (2011) surveyed consumer perceptions on daily deals and found no evidence of daily deal fatigue. Dholakia and Kimes (2011) developed a conceptual framework specifying the determinants of a profitable Groupon promotion and they empirically tested it using their surveyed data. Wu et al. (2015) quantified the economic value of daily deals for merchants by proposing a structural model; they empirically confirmed the long-run profitability of the business model for merchants. Luo et al. (2014) investigated the effect of deal popularity and social-influence related factors on the purchase likelihood and redemption time using a dataset of 30,272 customers of a group-buying website. Hu and Winer (2016) tested the effects of the"tipping point" on Groupon using augmented clickstream data and found tipping point can alter consumer behavior and affect sales. Our research also utilizes augmented clickstream data in which we observe the detailed click history of individual consumers, but we focus only on *newly* registered customers and their behaviors. We are the first to structurally model how acquired customers behave on daily deal sites with the objective to examine the empirical validity of the concern regarding deal fatigue among those who have subscribed to the site.

3 Data and Preliminary Analysis

The data comes from a third party research firm in US that tracks the online browsing behavior of a sample of over 2 million people in US and collects their clickstream data. For new Groupon users who subscribe between January and March 2011, we extract the complete clickstream data within their browsing sessions whenever they logged onto the Groupon website. Thus we observe consumers from the beginning of their association with Groupon. Our data contain, but are not limited to, the individual's entire history on Groupon such as subscription, click and view history, purchase behavior along with the associated time stamps. Furthermore, based on the clickstream information, we also crawled and obtained detailed information on deals from the Groupon website.This information includes the price of the deal, discount rate, and the total number of coupons being sold for the deal. Overall, the dataset contains 26,523 records,14,096 deals and 10,951 new subscribers who visited Groupon at least once during the sample period. The deals cover 46 states and 171 cities in the U.S. and contain a variety of categories including restaurants, arts and entertainment, shopping, beauty and spas. The detailed statistics of the dataset are shown in Table 1.

Table 1: Data Summary

DATA	SUMMARY
Covered time period	Jan 1, 2011 to Mar 31, 2011
No. of records	26,523
Average record size	180 Bytes
New subscribers within 3 months	10,951
Repeated visitors within three months	9,611
Distinct deals	14,096

To proxy for the (time invariant) quality of a deal, we retrieve the final number of coupons sold for a deal from the back end Groupon database. We then normalize this number by the local population to get a comparable metric across locations. We then assign the normalized number to one of 9 bins [9] and use that as the "quality" index in the analysis. Figure 2 graphs out the conditional purchase probability against the quality distribution of current deals following this definition. The horizontal axis represents the 9 quality bins. The higher the number the better the quality. The bars reflect the quality distribution of the deals, and the line represents the conditional purchase probability. The conditional probability is increasing with the bins, indicating that higher the quality the more likely the deal is purchased (see Orhun et al. (2015) for a similar operationalization of quality).

One might be concerned that there is a mechanical relationship between the number of coupons sold and the conditional probability of making a purchase. Note however that purchases are happening throughout the duration in which the deal is active whereas the number of coupons corresponds to when the deal ends. In other words, subscribers are inferring some aspect of quality when they visit a merchant's page and that aspect of quality ends up resulting in a higher number of coupons sold at the end of the deal period. This makes our choice a reasonable proxy for quality.

We explore several alternative measures of quality to proxy for deal quality, such as the Groupon ratings and customer referral behavior to ensure that they provide consistent results with those we obtain from the sales measure. Based on the relative fit we report the results of the best-fitting operationalization which is the one described above. We recognize that using a measure being based on the dependent variable maybe a concern. So, as a robustness check, we provide results from an alternative measure of quality as well.

3.1 Behavior in Daily Deal Websites

The data reveal several behavioral patterns of consumers that are relevant when specifying the model. First, through the sampling period, 90% of subscribers click only one newsletter deal on a given day (Figure 3 (a)). Second, among those who click only one newsletter deal a day, 79.50% only click on the featured deal in the newsletter in the entire sampling period (Figure 3 (b)).In a typical Groupon newsletter in 2011 (Figure 1 (a)) , the main body is taken by the featured deal and on the sidebar, there were several additional deals. For the subscribers who click on newsletter deals (including

[9] The quality bins from 1 to 9 are assigned to the following intervals of the normalized sold-out quantity: $(0,50]$, $(50,100]$, $(100,200]$, $(200,400]$, $(400,600]$, $(600,800]$, $(800,999]$, $(999,4999]$, and $4999+$.

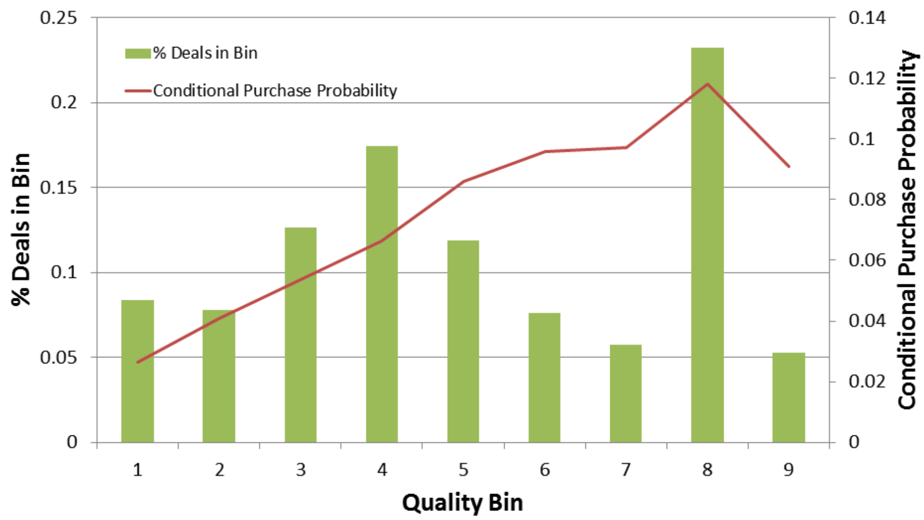

Figure 2: Conditional Purchase Probability vs % Deals in Bin

those who clicked on more than one newsletter deal per day), 82.66% only click on the featured deal. Combining this with the first observation, without loss of generality, we focus on each subscriber's daily decision of whether or not to click on the featured deal.

Third, 99.09% of subscribers in the data purchased at most one deal a day (Figure 3 (c)). And 93.66% of those making a purchase only purchased a featured deal in the sampling period. So combining all this information, we focus on the subscribers' decisions regarding the featured deal from the newsletter on a specific day - whether they click on it or not and whether they purchase it or not.

Fourth, 83.61% of the subscribers only checked a deal once (Figure 3 (d)). 89.86% of the deals were only checked once by a given subscriber. And we do not observe purchases of previously clicked deals. This shows that for subscribers, the purchase of Groupon deals is a sequential decision and there is no recall.

On average, each subscriber clicked and viewed about 2.3 deals over the 3-month period with an average purchase rate, conditional on a click, of 8%. The low purchase rate indicates that subscribers rarely purchase even when they receive deals that could be better in price and quality than previously purchased deals. Our data show that, following a purchase decision of a subscriber, there were, on average, 7.68 deals, better than the previously purchased deal in terms of both price and quality, that were not clicked. Similarly, purchase rates are lower conditional on a purchase (Figure 3 (e)). One plausible explanation is that the deals offered by Groupon can be redeemed over a certain time period (200 days on average in the data). So once a subscriber purchases a deal, (s)he may not want to purchase another deal from the same category again very soon. This can be viewed as deal satiation or a manifestation of the subscriber's budget constraint. Thus making one deal purchase seems to affect subsequent clicks and purchases by the subscriber. In order to account for why subscribers do not purchase deals that prima facie, are better than those previously purchased, we will need to account for previous purchases in our model and analysis.

The data also reveal that from the time consumers first click on the newsletter deals, on average, it takes about 49 minutes (s.d. 140 min) for each individual to complete the purchase if a purchase is made (Figure 3 (f)). So consumers tend to spend some time on the Groupon webpages after they check the newsletter. This provides some support for the notion of search costs associated with purchasing on this site.

An obvious concern when studying the behavior of consumers on Groupon is such behavior may be potentially influenced by activities undertaken by Groupon to encourage clicking and purchase behavior; specifically the targeting of deals by Groupon in early 2011. As we show here, it does not appear that Groupon was targeting at this stage except for sending deals in the category of interest specified by customers when they registered. This means that everyone saw the same deals in a given category. First, in the presence of targeting, we would expect a significant difference between the probability of receiving a deal in the same category before and after purchasing a deal from that category. This is a within person comparison. We find that probability to be 25.59% before purchase and 20.67% after purchase, respectively; a t-test shows no significant difference between the two (p-value = 0.26). Second, we compare two groups of people: those who bought a certain category ("buy") and those who never bought in this category ("no buy"). We calculate the probability of receiving featured deals of the purchased category from the newsletter for both groups in different cities. 92 (53.80%) out of 171 cities have the "buy" group with a higher probability of receiving deals in the same category, while there are 79 cities (46.20%) where the "no buy" groups have a higher probability. We cannot reject the null (p-value = 0.16) that the two groups are the same. A plausible reason why we find no evidence for targeting is probably because Groupon focused on local deals. And in the earlier days, there were usually not that many deals from the local area on a specific day from a specific category. So it would not have been possible to customize deals. Indeed, if targeting was getting better we would expect to see more clicking over time instead of the opposite unless consumer learning about the deals was faster than the firm's learning about the consumer.

We summarize the findings in the empirical data analysis and implications for our modeling effort as follows:

1. Without loss of generality, we can focus on consumers' decisions to purchase the featured deal in the newsletter sent to subscribers each day.

2. We need to account for the sequential nature of decision making and to allow for multiple purchases.

3. (i) The number of previously purchased deals, and (ii) search costs, in addition to deal characteristics, play an important role in consumer deal purchase behavior.

3.2 Groupon Fatigue ?

In the data, we track individual subscribers over time from registration on Groupon. Thus we observe within-subscriber over-time patterns of behavior from inception on Groupon. We see two patterns of interest in the data. First, we find that the probability that a consumer clicks on a merchant in the emailed newsletter declines over time. This is consistent with the notion of fatigue. However, we also find that the probability that the consumer makes a purchase is actually increasing over time conditional on click.

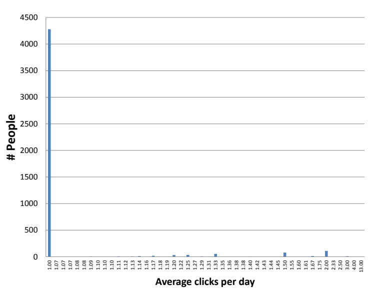

(a) Deal Clicks per Day

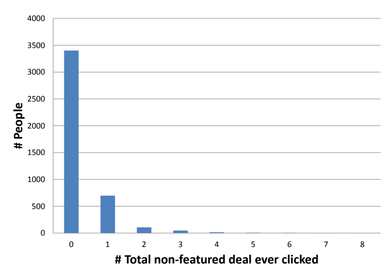

(b) Featured Deal Clicks per Day

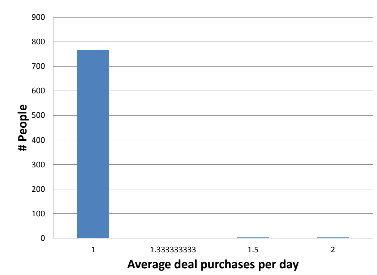

(c) Purchase per Day

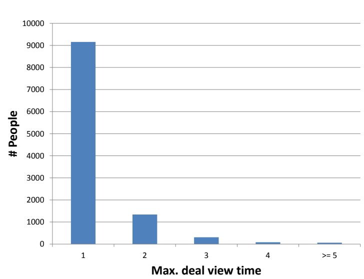

(d) Deal View

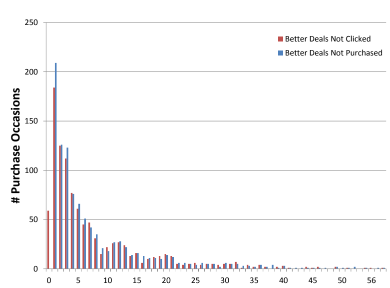

(e) Better Deals Not Clicked or Purchased

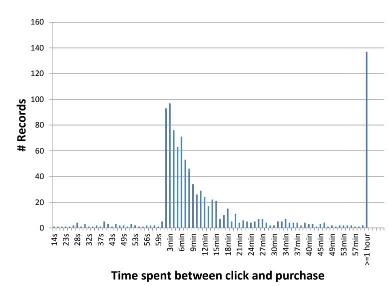

(f) View Purchase Time

Figure 3: Empirical Data Analysis

We graph out the two patterns in Figure 4 (a) and (b). Figure 4 (a) plots the click probability (y-axis) against tenure on the site (x-axis) and Figure 4 (b) shows the conditional and unconditional purchase probabilities (y-axis) against tenure (x-axis). The probabilities are averaged across all individuals with the same tenure [10]. We also include two solid lines which are the linear fitted trend line for each data series. We ran a simple (subscriber and deal) fixed-effects logit model [11] on the click and purchase data and controlled for the price, quality of the deal and the number of previously purchased deals. The estimate of tenure in the click model is -0.0320 (s.e. = 0.0004, average marginal effect = -0.0013) and the estimate of tenure in the conditional on click purchase model is 0.0373 (s.e. = 0.0028, average marginal effect = 0.0012). The results are consistent with declining clickthroughs to the webpages of the deals as well as the increasing conditional probabilities of purchasing on the site over time. Following previous research (Seiler, 2013) that looks at the role of inventory, we also examined whether the number of previous purchases of Groupon deals had an effect on click and purchase rates. These estimates are negative for both the click decision (coef. = -0.15, s.e. = 0.02, average marginal effect = -0.0063) and the conditional purchase decision (coef. = -1.69, s.e. = 0.08, average marginal effect = -0.05), indicating the presence of deal satiation or a budget constraint. This result motivates our subsequent inclusion of previous purchases in our model.

3.3 Search with learning

To reconcile these two data patterns we invoke the notion of consumer learning. Upon receiving the newsletter from daily deal websites, consumers can quickly obtain some information regarding the featured deal, such as the original price ($220 in the illustrated example in Figure 1a), the discount rate (80%), the the price after discount ($55) and some brief discription of the retailer (e.g. the address). If the consumer is interested in the deal, he can click the newsletter and get redirected to the webpage of the deal (i.e. Figure 1b). The webpage contains the complete information of the deal. In addition to the information revealed in the newsletter, the consumer can obtain the information regarding the time left to purchase of the deal (“2 days 15 hours and 9 minutes”), the number of coupons sold (“49 bought”), the status of the tipping point (“The deal is on!”) and the fine print on the deals (such as the redemption period, the maximum number of coupons that can be purchased

[10] The drop in the second period is because the registration with Groupon website is usually accompanied by the intention to view and purchase the deal. So the click and purchase rates in the first period are extremely high.

[11] We control for both consumer and deal fix effects.

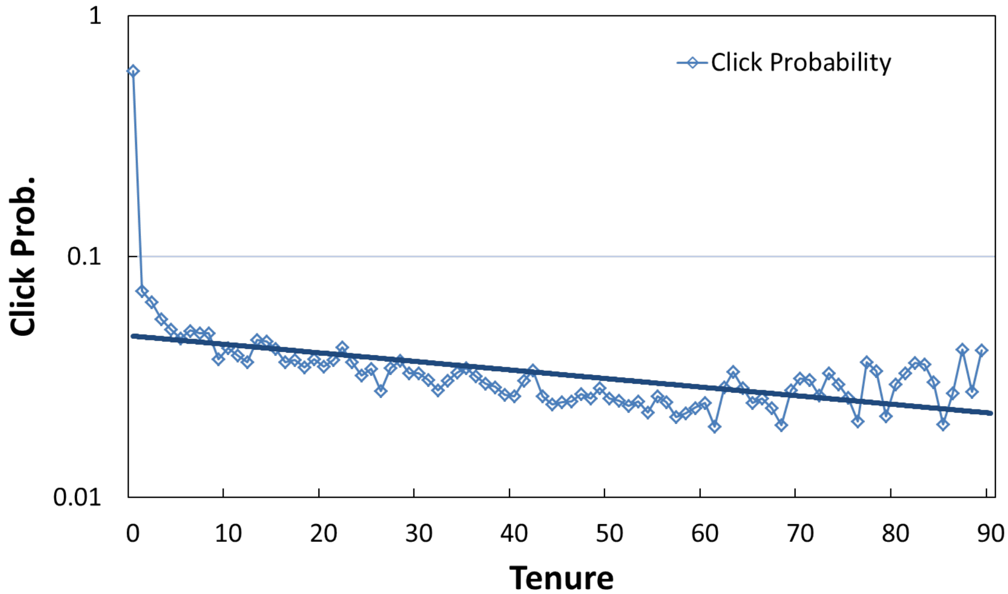

(a) Deal Clicks

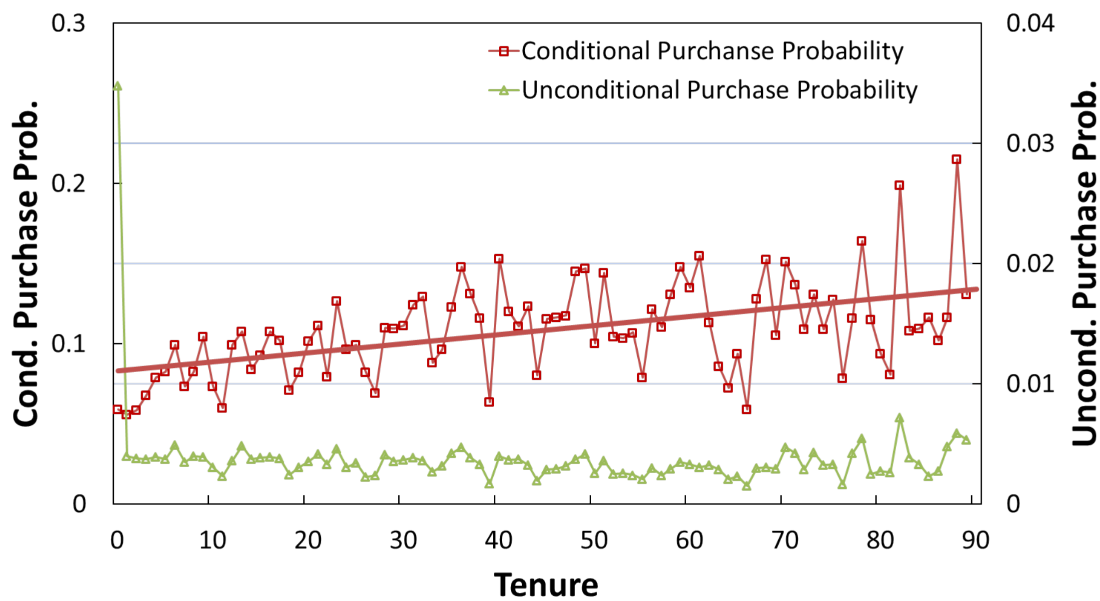

(b) (Un)Conditional Deal Purchases

Figure 4: Data Trends

per person, etc).

When making the purchase decision, the consumer faces the tradeoff between whether to purchase the current deal based on the revealed quality of the deal in the webpage or to wait for another deal in the future. However, as a new subscriber, the consumer is uncertain about the quality of future deals so expectations regarding the deal quality offered by Groupon need to be formed. This, in turn, is based on the prior knowledge of the Groupon deals as well as the deals the consumer clicks and views after subscription. Consumers use the information on the deal from the webpage to update their knowlege regarding the quality distribution of the deals offered by Groupon. This comprises the “learning” stage of our model.

Upon receiving the newsletter from the daily deal website, the consumer decides whether or not to click on the newsletter to obtain more information on the deal. A search cost is incurred if the consumer decides to vist the page to obtain the information. This could be associated with the loading time of the webpage or the effort spent reading through the details of the deal. This is the “search” stage. Consumers tend to click on more deals in the beginning to learn about the quality distribution of deals. As their knowledge accumulates, the incentive to learn declines so they click fewer deals. Over time, since clicking is less about learning the quality of deals, the motivation for clicking is more likely to be to purchase - hence the data pattern of increasing purchase probabilities.

A key question to be answered is: what is being learnt at the webpage? Note that several characteristics are available directly in the newsletter. By visitng the webpage the consumer can learn about some additional features but also about where this deal stacks up relative to other deals (s)he may have encountered on the site. We are not interested in the purchase of a specific deal (since that deal is unlikely to reappear in the data); rather we are more interested in the consumer’s knowledge of the distribution of deals and how that might be influenced by information from this specific deal as this, in turn, would influence subsequent behavior on the site. As noted in the data section, we refer to this as the distribution of deal quality (prices are already observed in the newsletter so this is known prior to investing the search cost); we use the total number of coupons sold for the deal as representing that deal’s quality. Clearly, this information may or may not be available to a speecific subscriber since subscribers arrive at different times. We are merely using the normalized total number sold (at the end of the deal) as a proxy for all the information that the subscriber obtains at the site since, as shown earlier, this measure of quality correlates well with the probability of any given subscriber purchasing on the site. Another feature of the quality of deals being provided by

Groupon is that the distribution remains quite stable over the time period of the data. We verified this by comparing the quality distributions of the deals at different times within our observation window; we found no significant difference in the quality distribution over time. Later, in the model specification we assume subscribers are learning from a *stationary* quality distribution.

3.4 Reduced Form Analysis for Learning

In this section we try to show two patterns in the data - first, that learning seems to be occurring among the subscribers; second, that this learning is influencing click and purchase behavior. First, we take each subscriber's tenure with Groupon and divide into two equal temporal regimes. Then, in Figure 5a, we plot the mean and standard deviation of the purchased deal quality for the earlier and later time periods averaged across all subscribers. Compared to the first half of tenure, subscribers purchase deals with higher quality (5.00 vs 6.36, $p < .000$) and with lower quality variance (1.67 vs 1.45, $p < .05$) in the second half. Figure 5b uses the quality/price ratio to control for price and the effects are even more pronounced.

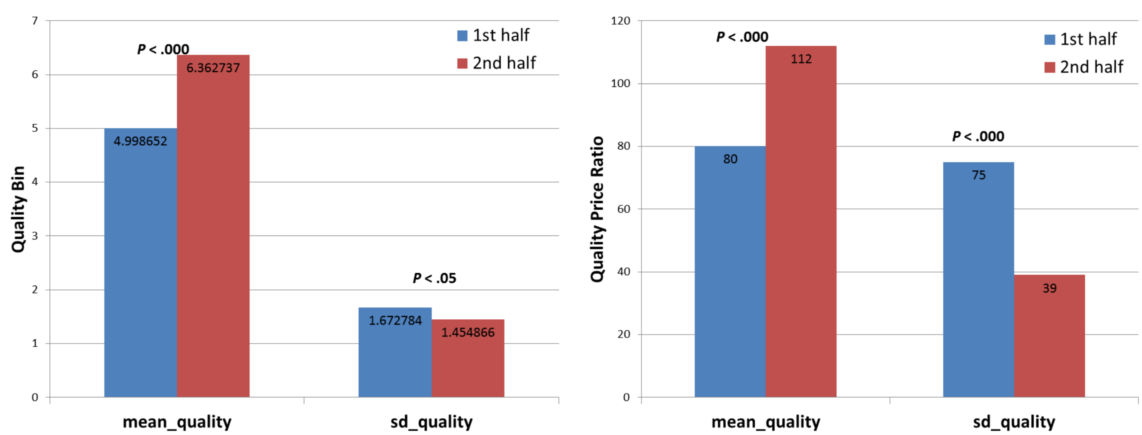

(a) Median Split Analysis on Quality (b) Median Split Analysis on Quality Price Ratio

Figure 5: Empirical Evidence of Consumer Learning

Next, we ran a reduced form analysis for consumers' click and purchase behaviors on the cumulative number of clicks (s)he has had before, while controlling for deal price, quality (using the above operationalization), and number of previously purchased deals. If learning is manifested in the way being described above, the coefficient on lagged cumulative clicks should be negative for click and positive for purchase, so the more clicks the individual has had before, the less likely (s)he will be to click the deal but more likely to purchase a deal once (s)he clicks. Indeed, this is what we find in the data. Table 2 presents the reduced form analysis results. Column I reports the results for a logit model with deal and subscriber fixed effects on the click behavior and controls for price, quality but replaces tenure with the cumulative number of clicks, whose coefficient is negative. Similarly, we report the results of logit model on purchase in Column II. The coefficient of cumulative clicks on purchase (conditional on click) is positive. So the two observed data patterns over time appear to be consistent with the notion of learning. This provides us with motivation to incorporate learning into the two-stage dynamic search model (although we recognize that there might be other plausible hypotheses for these findings as well).

Table 2: Reduced Form Analysis

DV	(I) **Click**		(II) **Conditional Purchase**	
	coeff.	ave.marginal.eff	coeff.	ave.marginal.eff
Price/100	0.0005	0.00002	-1.0579***	-0.0397
	(0.0065)		(0.1718)	
Quality	0.0374***	0.0016	0.2071***	0.0078
	(0.0033)		(0.0188)	
Cumu. Click	-0.0142***	-0.0006	0.0174**	0.0007
	(0.0023)		(0.0078)	
Deal FE	Yes		Yes	
Individual FE	Yes		Yes	
Observations	527,298		25,459	

Notes: Standard errors in parentheses. *** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

4 Methodology

We build a two stage dynamic model (Seiler, 2013) with consumer learning. Figure 6 illustrates the decision flow we want to capture in the model. In the first stage of the model, consumers form expectations about the deal quality distribution and decide whether they want to click the webpage of the deal or not. Once they click and visit the webpage, they decide whether to purchase the deal or wait for another one based on the information they gather about the deal. Visiting the page also allows consumers to to learn about that information and use it to update their posterior perceptions of the quality distribution of deals provided by the daily deal website.

We use a non-parametric learning approach (i.e. Dirichlet learning) to characterize the learning process (Koulayev,2013; De los Santos et al., 2013). One distinction we draw with a majority

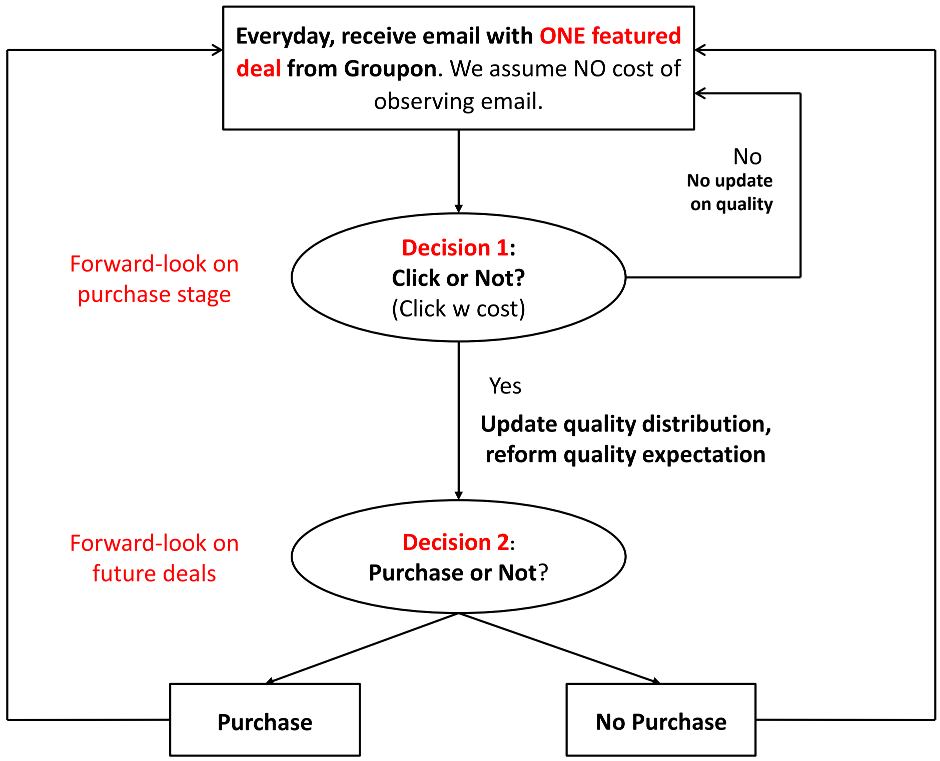

Figure 6: Structural Model Flow

of the existing learning literature (e.g. Erdem and Keane,1996) that assumes consumers know the distribution of quality (e.g.,normal) and learn about the mean (and sometimes the variance) of the distribution, is that we assume that new subscribers to Groupon do not know the distribution of deal qualities and need to learn about the distribution itself. As a result, we incorporate Dirichlet learning into the model (Rothschild, 1974; Koulayev, 2013; De los Santos et al., 2013). A Dirichlet prior is less restrictive in terms of the shape of the distribution compared to e.g., a normal distribution. We assume a consumer has a prior on the quality distribution of Groupon deals before registration. In particular, the consumer knows that quality can take a finite number of values but the consumer does not know the probability of each quality level occurring. Based on the deals clicked on, the consumer learns about the probability of each quality level occurring, so the consumer learns about the *distribution* of deal quality, rather than the "true" quality of the product (as in e.g., Erdem and Keane (1996)). However, the consumer can never predict the exact deal quality of the next featured deal with certainty even after learning about the distribution of deals.

4.1 Two Stage Forward Looking Search Model

Our structural model is incorporating non-parametric Bayesian learning to a two stage dynamic search model (Seiler, 2013). Seiler (2013) built a two-stage dynamic demand model which takes into account search cost. In the first stage, consumer decides whether to search for price; in the second stage, consumer decides whether to make a purchase. However, he did not consider consumer's learning behavior. Our model incorporate consumer learning into the two-stage dynamic model.

In section 1, we have introduced consumer's decision making process. Correspondingly, we will build a model to reflect such a process. Based on our preliminary analysis, we focus on consumers' click and purchase decisions for the featured deal in the emailed newsletter they receive on each day. As we will show below, for consumers, the daily decision can be formulated as a dynamic discrete choice problem. We note that the basic model borrows heavily from Seiler (2013); we then extend the model to account for subscribers' learning behavior.

4.1.1 Flow Utility

We use the subscript "*cs*" to denote the click stage, and "*ps*" to denote the purchase stage; "*i*" to denote consumers, and "*t*" to denote time period. "d_{it}^{C}" is the decision indicator at the click stage. $d_{it}^{C} = 1$ ($d_{it}^{C} = 0$) means the consumer decides to (not) click the featured deal in the newsletter. Similarly, "d_{it}^{P}" is the decision indicator at purchase stage. $d_{it}^{P} = 1$ ($d_{it}^{P} = 0$) means the consumer decides to (not) purchase the featured deal. At each stage, corresponding to each decision, the consumer will have a flow utility [12].

At the click stage, to obtain information on the featured deal, the consumer incurs a search cost associated with clicking and going to the web page of the deal. The click cost is denoted by c_{it}. The flow utility (s)he gets is

$$
\begin{aligned}
u_{cs,it}^{1} &= -c_{it} + \varepsilon_{cs,it}^{1} \\
&= \bar{u}_{cs,it}^{1} + \varepsilon_{cs,it}^{1} \qquad (1)
\end{aligned}
$$

Similarly, if (s)he decides not to click the deal, no click cost is incurred and (s)he will get the

[12] The definition of flow utility follows Seiler (2013), which is the utility at either decision stage in a particular time period.

following flow utility

$$u^0_{cs,it} = \eta \cdot pp_{it} + \varepsilon^0_{cs,it}$$
$$= \bar{u}^0_{cs,it} + \varepsilon^0_{cs,it} \tag{2}$$

where pp_{it} is the number of previous deals purchased by individual i before period t. Daily deals are typically coupons that can be redeemed within a certain period. So the existing deals purchased could influence the subscriber's opportunity costs before they are redeemed which in turn affects whether (s)he clicks on a subsequent deal. We use this term to capture the effects of deal satiation or budget constraint. Adding this term is motivated by our earlier finding in the raw data that Groupon subscribers skipped better deals after making a purchase decision and a negative coefficient on the number of previously purchased deals in our regressions (see Table 2 where η captures the corresponding effect). We re-set $pp_{it} = 0$ after the redemption periods of the purchased deals expire. The average redemption period of the deals in the data is about 195 days (with s.d. of 82 days). We also checked for sensitivity of our results to this reset period. The pp_{it} variable is different from the "inventory" effect for storable goods in Seiler (2013) or the "the level of depreciation" effect for tied goods in Hartmann and Nair (2010), in that we allow the value to increase over time but reset it periodically. With this term, we are able to explain why people do not keep purchasing deals even when they face better ones after a purchase. But this effect will not last when we re-set the variable; this ensures stationarity.

If the consumer decides to click, (s)he enters the second stage where more information about the featured deal is revealed. Upon opening the emailed newsletter, the price of the featured deal P_{it} is revealed to the subscriber; we assume that there is no cost to obtain the information on price. We use Q_{it} to denote the quality level of the specific deal revealed to the consumer after clicking through to the merchant's site. Based on the newly obtained information (i.e., the quality level of the deal), (s)he updates his/her beliefs about the quality distribution (we will introduce the learning mechanism in Section 4.2). We denote the prior quality distribution at period t for consumer i as $f^{pr}_{it}(Q)$ and the corresponding posterior as $f^{po}_{it}(Q)$. We also assume the consumer is risk neutral. Therefore, if (s)he decides to purchase the deal, the flow utility obtained will be determined by the current deal (i.e., its price and quality) and the previous purchase status (i.e., number of previously purchased deals):

$$
\begin{aligned}
u^1_{ps,it} &= \beta_0 + \beta_P \cdot P_{it} + \beta_Q \cdot Q_{it} + \eta \cdot pp_{it} + \varepsilon^1_{ps,it} \\
&= \bar{u}^1_{ps,it} + \varepsilon^1_{ps,it}
\end{aligned} \tag{3}
$$

If, however, (s)he decides not to purchase, (s)he gets utility

$$
\begin{aligned}
u^0_{ps,it} &= \eta \cdot pp_{it} + \varepsilon^0_{ps,it} \\
&= \bar{u}^0_{ps,it} + \varepsilon^0_{ps,it}
\end{aligned} \tag{4}
$$

When deciding to click the deal in the newsletter, the consumer anticipates (s)he will have to make another decision within the same time period and knows that (s)he will receive more *non-discounted* utility as defined by the purchase-stage flow utility. Therefore the total utility in the current time period if the consumer decides to search is simply the sum of the flow utilities from both decision-making stages:

$$
\begin{aligned}
U_{it}(d^C_{it} = 1) &= \max\{u^0_{ps,it}, u^1_{ps,it}\} + u^1_{cs,it} \\
&= \max\{u^0_{ps,it}, u^1_{ps,it}\} + \bar{u}^1_{cs,it} + \varepsilon^1_{cs,it}
\end{aligned} \tag{5}
$$

When deciding not to click, the total utility in the current time period is just equal to the flow utility at click stage (since the consumer does not enter into the second stage, therefore no purchase utility obtained).

$$
\begin{aligned}
U_{it}(d^C_{it} = 0) &= u^0_{cs,it} \\
&= \bar{u}^0_{cs,it} + \varepsilon^0_{cs,it}
\end{aligned} \tag{6}
$$

It is worth mentioning that prior to clicking, the current period flow utility in the purchase stage is unknown. Therefore, the consumer has to form an expectation over the flow utility in the purchase stage, knowing (s)he will make the optimal purchase decision conditional on the information that (s)he will obtain in the purchase stage [13]. This information is not yet available to the consumer in

[13] As in Seiler (2013), the use of the max-operator in the above equation constitutes a slight abuse of notation. To be more accurate, in the dynamic setting, the consumer will make a choice in the purchase stage that maximizes the present discounted value and not the flow utility. The maximization is therefore with respect to the choice-specific value

the click stage, and therefore the consumer cannot perfectly predict his/her decision at the purchase stage.

Let $\tilde{\varepsilon}_{cs,it} = (\varepsilon^0_{cs,it}, \varepsilon^1_{cs,it})$ be the vector of idiosyncratic shocks at click stage in time t and $\tilde{\varepsilon}_{ps,it} = (\varepsilon^0_{ps,it}, \varepsilon^1_{ps,it})$ is the corresponding idiosyncratic shock vector at purchase stage. Note that different sets of error terms enter the model before and after the clicking decision.

4.1.2 The Dynamic Optimization Problem

Formally, a consumer chooses an infinite sequence of decision rules $D = \{d_{it}\}_{t=0}^{\infty}$ in order to maximize his/her expected present-discounted sum of utilities [14]

$$\max_{\{d_{it}\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty} \delta^t \cdot U_{it}(d_{it})|x_{it}, \tilde{\varepsilon}_{it}\right\} \tag{7}$$

where $d_{it} = \{d^C_{it}, d^P_{it}\}$, $\tilde{\varepsilon}_{it}$ includes all the idiosyncratic shocks and δ is the discount factor. x_{it} is a vector of observed state variables for subscriber i in period t.

4.1.3 Value Functions

In the model, the consumer potentially has two consecutive decisions to make in one time period. First, the consumer has to decide whether to click the deal in the newsletter to obtain more information about the deal. If (s)he does not click, (s)he does not to make any other decision in the current period. If (s)he decides to click the deal, she receives information about the quality distribution, and has to then decide whether to make a purchase. Therefore we define two different value functions depending on whether the consumer has clicked the deal: 1) value function at click stage $V_{cs,it}$; 2) value function at purchase stage $V_{ps,it}$. They depend on one another and must be solved simultaneously. As in Rust (1987), we define the expectation of value function at click (purchase) stage

function, which we will make clear subsequently.

[14] Specifically, the *total* flow utility gained in period t with respect to each decision is expressed as follows.

$$U_{it}(d^C_{it} = 0) = u^0_{cs,it}$$
$$U_{it}(d^C_{it} = 1, d^P_{it} = 0) = u^1_{cs,it} + u^0_{ps,it}$$
$$U_{it}(d^C_{it} = 1, d^P_{it} = 1) = u^1_{cs,it} + u^1_{ps,it}$$

$EV_{cs,it}$ ($EV_{ps,it}$), integrated over the realization of $\tilde{\varepsilon}_{cs,it}$ ($\tilde{\varepsilon}_{ps,it}$) as

$$EV_{cs,it} = EV_{cs}(x_{cs,it}) = \int_{\tilde{\varepsilon}_{cs,it}} V_{cs,it}(x_{cs,it}, \tilde{\varepsilon}_{cs,it}) \cdot dF_{\tilde{\varepsilon}_{cs,it}} \tag{8}$$

$$EV_{ps,it} = EV_{ps}(x_{ps,it}) = \int_{\tilde{\varepsilon}_{ps,it}} V_{ps,it}(x_{ps,it}, \tilde{\varepsilon}_{ps,it}) \cdot dF_{\tilde{\varepsilon}_{ps,it}} \tag{9}$$

By assuming the error terms ($\tilde{\varepsilon}_{cs,it}$, $\tilde{\varepsilon}_{ps,it}$) follow i.i.d. extreme value distributions and also making the standard conditional independence assumptions [15] (c.f. Rust (1987)), we obtain closed-form expected value functions with respect to each decision stage [16]:

$$\begin{aligned}
EV_{cs,it} &= EV_{cs}(x_{cs,it}) \\
&= \log\{\exp(\bar{v}^0_{cs,it}) + \exp(\bar{v}^1_{cs,it})\} \\
&= \log\left\{\exp\left(\bar{u}^0_{cs,it} + \delta \cdot E\left[EV_{cs}(x_{cs,it+1})|x_{cs,it}, d^C_{it} = 0\right]\right) + \right. \\
&\qquad \left. \exp\left(\bar{u}^1_{cs,it} + E\left[EV_{ps}(x_{ps,it})|x_{cs,it}, d^C_{it} = 1\right]\right)\right\}
\end{aligned} \tag{10}$$

$$\begin{aligned}
EV_{ps,it} &= EV_{ps}(x_{ps,it}) \\
&= \log\{\exp(\bar{v}^0_{ps,it}) + \exp(\bar{v}^1_{ps,it})\} \\
&= \log\left\{\exp\left(\bar{u}^0_{ps,it} + \delta \cdot E\left[EV_{cs}(x_{cs,it+1})|x_{ps,it}, d^P_{it} = 0\right]\right) + \right. \\
&\qquad \left. \exp\left(\bar{u}^1_{ps,it} + \delta \cdot E\left[EV_{cs}(x_{cs,it+1})|x_{ps,it}, d^P_{it} = 1\right]\right)\right\}
\end{aligned} \tag{11}$$

The expected value functions only depend on observed state variables [17]. The term $EV_{ps,it}$ can be interpreted as the inclusive value of clicking the featured deal in time period t, excluding the click cost. At the click stage, the consumer will make a trade-off between the expected utility of this inclusive value, less the click cost, with the utility of not clicking.

[15] Specifically, these conditional independence assumptions are:

$$\begin{gathered}
p(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}|x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 0) = p(\tilde{\varepsilon}_{cs,it+1}|x_{cs,it+1}) \cdot p(x_{cs,it+1}|x_{cs,it}, d^C_{it} = 0) \\
p(x_{ps,it}, \tilde{\varepsilon}_{ps,it}|x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 1) = p(\tilde{\varepsilon}_{ps,it}|x_{ps,it}) \cdot p(x_{ps,it}|x_{cs,it}, d^C_{it} = 1) \\
p(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}|x_{ps,it}, \tilde{\varepsilon}_{ps,it}, d^P_{it} = j) = p(\tilde{\varepsilon}_{cs,it+1}|x_{cs,it+1}) \cdot p(x_{cs,it+1}|x_{ps,it}, d^P_{it} = j), \quad j \in \{0, 1\}
\end{gathered}$$

[16] The detailed derivation of value functions can be found in Section 7 of the Appendix.

[17] We postphone the detailed introduction to state variables in Section 4.1.4

4.1.4 State Variables and Transition Rules

There are two decision stages in one time period, and the consumer only enters the second stage if and only if (s)he clicks, otherwise, (s)he just make one click decision in the current period and then enter into the next period. There are two sets of state variables: click-stage state variables $x_{cs,it}$ and purchase-stage state variables $x_{ps,it}$. Corresponding to the decision sequences, we have the following possible state transitions:

1. From the *click stage* in the current period to the *click stage* in the next period: $x_{cs,it} \to x_{cs,it+1}$
2. From the *click stage* in the current period to the *purchase stage* in the current period: $x_{cs,it} \to x_{ps,it}$
3. From the *purchase stage* in the current period to the *click stage* in the next period: $x_{ps,it} \to x_{cs,it+1}$

Prior to clicking in period t, the consumer has information on the price of the current featured deal from the newsletter, as well as some knowledge about the quality levels of the deals [18]. We use the vector $\vec{\alpha}_{it}^{pr}$ to denote a consumer's knowledge of the distribution of quality ($f_{it}^{pr}(Q)$) on the Groupon website at click stage in the current period. The expected value function therefore, depends on $\vec{\alpha}_{it}^{pr}$. We provide the derivation in Section 4.2. Also, we assume customers know the number of deals they have previously purchased with certainty (i.e., their inventory). Therefore, the state variables at this stage are the price of the deal, number of previous purchased deals, and the prior knowledge of the quality levels:

$$x_{cs,it} = \{P_{it}, pp_{it}, \vec{\alpha}_{it}^{pr}\} = \{P_{it}, pp_{it}, \{\alpha_{it}^{pr,1},\ \alpha_{it}^{pr,2}, ..., \alpha_{it}^{pr,N}\}\} \tag{12}$$

If the consumer enters the purchase stage, more information is revealed, (s)he receives the quality signal Q_{it} for this deal, and (s)he uses this signal to update her knowledge about the distribution of quality. After updating, the distribution becomes $f_{it}^{po}(Q)$ and the knowledge about this distribution is denoted as $\vec{\alpha}_{it}^{po}$. The purchase-stage state variables are

$$x_{ps,it} = \{P_{it}, Q_{it}, pp_{it}, \vec{\alpha}_{it}^{po}\} = \{P_{it}, Q_{it}, pp_{it}, \{\alpha_{it}^{po,1},\ \alpha_{it}^{po,2}, ..., \alpha_{it}^{po,N}\}\} \tag{13}$$

[18] if (s)he is a fresh new user without any click history, this knowledge corresponds to the subscriber's prior beliefs. That is $f_{it}^{pr}(Q)$ where $t = 0$.

To complete the definition of state variables, we note that P_{it} remains unchanged within period t so there is no transition in P_{it} from the click stage to purchase stage. P_{it+1} is the price revealed to consumer i (and researcher) the next period. We make the assumption that the consumer has rational price beliefs and follows the empirical distribution of prices which we approximate with a normal distribution with the empirical mean and empirical standard deviation in the data. The evolution of pp_{it} is straightforward. pp_{it+1} is equal to $pp_{it} + \Delta pp_{it}$, where Δpp_{it} is the number of deals purchased in period t. In our case, Δpp_{it} is either 0 or 1. We lay out the detailed transition rules in Section 7 of the Appendix.

4.2 Dirichlet Learning Process

The subscriber updates his or her belief about the quality distribution according to a Dirichlet learning process (DLP). The Dirichlet distribution is a widely used in Bayesian statistics, because it is a conjungate prior to the multinomial distribution.

Assume there are N quality levels: $\vec{q} = \{q_1, q_2, ..., q_N\}$ and the sampling probability of each level at peroid t is denoted using the vector $\vec{\rho}_t = \{\rho_{1t}, \rho_{2t}, ..., \rho_{Nt}\}$. $\vec{\rho}_t$ is distributed according to a Dirichlet distribution of order N with density

$$f(\rho_{1t}, \rho_{2t}, ..., \rho_{Nt}) = \frac{\Gamma(\sum_{n=1}^{N} \alpha_{nt})}{\prod_{n=1}^{N} \Gamma(\alpha_{nt})} \prod_{n=1}^{N} \rho_{nt}^{\alpha_{nt}-1} \tag{14}$$

If the prior expected value of each ρ_{nt} is given by $E[\rho_{nt}] = \frac{\alpha_{nt}}{W_t}$ and $W_t = \sum_{n=1}^{N} \alpha_{nt}$, then the corresponding posterior is

$$E[\rho_{nt+1}] = \begin{cases} \frac{\alpha_{nt}}{W_t+1} & if\ q_n\ is\ not\ sampled \\ \frac{\alpha_{nt}+1}{W_t+1} & if\ q_n\ is\ sampled \end{cases} \tag{15}$$

This completes our definitions of state variables and their transition rules. For a more detailed description of the Dirichlet learning process, please refer to Appendix 7. We now explicitly write down the click- and purchase-stage expected value functions (a detailed derivation can be found in Appendix 7.)

$$
\begin{aligned}
EV_{cs,it} &= EV_{cs}(P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}) \\
&= \log\left\{\exp\left(\eta \cdot pp_{it} + \delta \cdot \int\limits_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{pr}, pp_{it+1}) \cdot dF(P_{it+1})\right) + \right. \\
&\qquad \left. \exp\left(-c_{it} + \sum_{n=1}^{N} \left(EV_{ps}(P_{it}, \vec{\alpha}_{it}^{po}(q_n), pp_{it}) \cdot \rho_e^n\right)\right)\right\}
\end{aligned} \tag{16}
$$

$$
\begin{aligned}
EV_{ps,it} &= EV_{ps}(P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}) \\
&= \log\left\{\exp\left(\eta \cdot pp_{it} + \delta \cdot \int\limits_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{po}, pp_{it+1}) \cdot dF(P_{it+1})\right) + \right. \\
&\qquad \left. \exp\left(\beta_0 + \beta_P \cdot P_{it} + \beta_Q \cdot Q_{it} + \eta \cdot pp_{it} + \delta \cdot \int\limits_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{po}, pp_{it+1}) \cdot dF(P_{it+1})\right)\right\}
\end{aligned} \tag{17}
$$

4.3 Terms in the Likelihood Function

Based on our assumptions for the idiosyncratic shocks in both stages, the probability of clicking and the conditional probability of purchasing can be derived as follows using their corresponding choice-specific counterparts:

$$
P_{it}^{c} = \frac{\exp(\bar{v}_{cs,it}^{1})}{\exp(\bar{v}_{cs,it}^{0}) + \exp(\bar{v}_{cs,it}^{1})} \tag{18}
$$

$$
P_{it}^{pur|c} = \frac{\exp(\bar{v}_{ps,it}^{1})}{\exp(\bar{v}_{ps,it}^{0}) + \exp(\bar{v}_{ps,it}^{1})} \tag{19}
$$

- The probability of observing a consumer click and purchase the deal is ($d_{it}^{C} = 1$, d_{it}^{P}=1): $P_{it}^{c} \cdot P_{it}^{pur|c}$
- The probability of observing a consumer click the deal but not purchase is ($d_{it}^{C} = 1$, d_{it}^{P}=0): $P_{it}^{c} \cdot (1 - P_{it}^{pur|c})$

- The probability of observing a consumer not clicking is ($d_{it}^{C}=0$, d_{it}^{P}=0): $1-P_{it}^{c}$

The probabilities derived above can now be used to form the likelihood function using the decision indicator at both stages.

$$L_{it}=\left(P_{it}^{pur|c}\cdot P_{it}^{c}\right)^{d_{it}^{P}}\times\left(\left[\left(1-P_{it}^{pur|c}\right)\cdot P_{it}^{c}\right]^{d_{it}^{C}}\cdot[1-P_{it}^{c}]^{1-d_{it}^{C}}\right)^{1-d_{it}^{P}} \tag{20}$$

The total log-likelihood function is simply $LL=\sum_{i}\sum_{t}\log L_{it}$.

4.4 Estimation and Identification

The two-stage dynamic learning model is estimated using the nested fix-point algorithm (Rust, 1987). As in Rust (1987), given a set of parameter values, the inner loop computes the value function and evaluates the likelihood function; the outer loop searches for parameters which maximize the likelihood function. In particular, We use Berndt-Hall-Hall-Hausman (Berndt et al., 1974) algorithm to conduct the outer loop optimization. When computing the value functions, we use an adapted version of the Keane and Wolpin (1994) approximation method used in Crawford and Shum (2005). [19].

4.4.1 Empirical Identification

Now we briefly consider the structural features of the model and the variation in the data that helps to identify the model parameters. From the purchase decision, the variation in deal prices and quality levels along with the number of deals previously purchased help us identify the price (β_P), quality (β_Q) and previous purchase number (η) coefficients. As in Koulayev (2010), search cost (c_{it}) is identified through the variation in click decisions (i.e., click or not click) together with the information of previously clicked deals. Unlike Seiler (2013), where the search stage is latent, in our case we observe individual click behavior in the first stage so this allows for identification of the search cost parameter. The rate of the declining click trend and increasing conditional purchase trend along with a stationary quality level distribution (shown previously) helps us identify the learning parameters [20], that is the Dirichlet priors ($\vec{\alpha}_{it=0}^{pr}=\{\alpha_{it=0}^{pr,1},\alpha_{it=0}^{pr,2},...,\alpha_{it=0}^{pr,N}\}$). Since the discount factor is difficult to estimate (c.f. Rust (1994); Erdem and Keane (1996)), we fix the value of δ at 0.995.

[19]A detailed explanation of the method can be found in the Appendix of Crawford and Shum (2005).

[20]However, the underlying learning mechanism - e.g., whether consumers learn about the mean of the quality distribution by normal priors or learn about the distribution of quality as we assume - is hard to empirically identify. In our application, we estimate both models: a two-stage dynamic model with normal Bayesian learning (Erdem and Keane, 1996; Zhang, 2010) and with Dirichlet learning (Koulayev, 2013; De los Santos et al., 2013). We find that the Dirichlet

4.4.2 Monte Carlo Simulation

We conduct Monte Carlo simulations to verify our estimation algorithm and to examine our estimation strategy. The main objective is to confirm that our estimation procedure is able to recover the unknown parameters of the model. The simulation is run as follows. We randomly generate a panel data of 500 agents with 20-day observations. Each day each agent receives a newsletter with one featured deal. Then the agent make decision: whether to click the deal; and if a click occurs, whether to purchase. An observation includes data on the deal price, quality level, click decision and purchase decision. The price and quality are generated from a normal distribution. The stochastic term in the utility specification is randomly drawn from a standard Type-I extreme value distribution and each observation has a different draw for each stage-decision combination. The flow utility gained at each stage follows the definition in Section 4.1.1.

Table 3 presents the results of the Monte Carlo simulation. Column (I) provides the search cost and utility parameters used to generate the data. Column (II) presents the mean and standard deviation of our model estimates across 30 replications. The results show that we are able to recover the true parameters using our model specification and estimation strategy.

Table 3: Monte Carlo Simulation

	(I)	(II)	
	true coeff.	coeff.	s.e.
utility			
Constant	-3	-2.830	0.360
Price/100	-3	-3.028	0.101
Quality	2	1.976	0.106
# Previous Purchases: η	-0.05	-0.048	0.014
Baseline cost: k	1	0.985	0.028
Dirichlet Priors			
α_1	1	1.011	0.331
α_1	2	1.996	0.428
α_1	1	0.981	0.318
Observations	-	10,000	

Notes: Column (II) presents the mean and s.d. of the parameter estimates across 30 replications.

learning process fits our data much better than the normal learning model. More importantly, we feel that in our present context the assumption of consumers learning about the quality distribution of deals makes more sense.

5 Empirical Application

Table 4: Summary Statistics

	Restaurant				All Categories			
	Mean	Min.	Max.	S.D.	Mean	Min.	Max.	S.D.
Deal								
Price (USD)	13.23	2	1500	17.47	37.28	1	12500	109.64
Discount (%)	51	40	83	2.62	55.79	19	100	9.51
Original Price (USD)	25.99	3.33	3000	34.74	740.60	3	999999	25308
Quality[a]	6.14	1	9	1.95	5.10	1	9	2.44
User								
Tenure (days)	32.27	0	89	21.58	31	0	89	21.96
If Click	0.0457	0	1	0.2089	0.0483	0	1	0.2144
If Purchase	0.0049	0	1	0.0696	0.0038	0	1	0.0615
If Purchase (cond.) [b]	0.1065	0	1	0.3085	0.0785	0	1	0.2690
Cumu. Click	0.4001	0	20	0.8585	1.77	0	84	0.43
Observations	108,698				527,298			

[a] We used normalized total sales of each deal as proxy of quality. And divide the numbers to 9 categories with 1 indicate lowest and 9 the highest quality.

[b] Conditional on clicking.

We apply the model to a subset of the data on restaurant deals. We include new subscribers that have *ever clicked* restaurant deals over the sampling period for the estimation. Among all the 18 categories of deals predefined by Groupon, restaurant is the largest, accounting for about 20% of all the deals. Deals are differentiated products that may vary in different aspects. So focusing on one major category simplifies our analysis. We leave the extension of our model to multiple categories to future research.

The summary statistics of the data that we use for estimation is in Table 4 [21]. In total, we have 108,698 observations from 10,485 subscribers. The average price of restaurant deals is about 13.23 USD, which is after an average discount of 51%. The empirical distribution of the price after discount, the original price and discount rate for restaurant deals are shown in Figure 7a, b and c. Due to the lack of much variation in the discount rate, we do not include it in the estimation. The mean quality of the deal is above the quality bin 6. Figure 7d is the empirical distribution of the quality for restuarant deals. It shows a bimodal pattern which can be captured well by our Dirichlet assumption. In terms of the characteristics of individuals, on average, they have been with Groupon

[21] As comparison, we also include the summary statistics of the entire dataset in the table.

for about one month with some joining early in our data period while others join later. The click rate in the data is around 4.6% , the conditional and unconditional purchase rates are 11% and 0.48% respectively. [22] We also include the summary statistics of the variables from all categories in Table 4 as comparison.

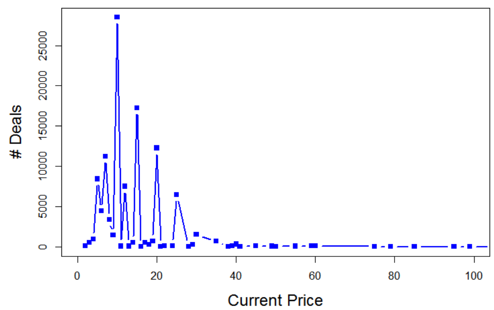

(a) Current Price Distribution

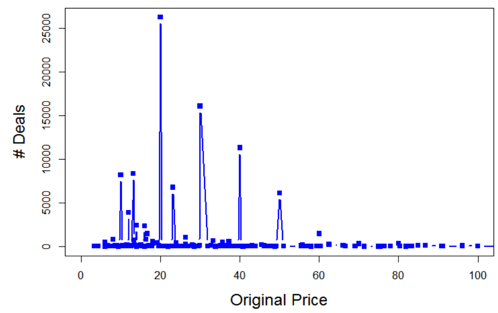

(b) Original Price Distribution

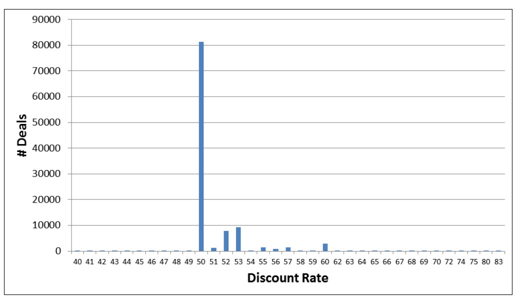

(c) Discount Rate Distribution

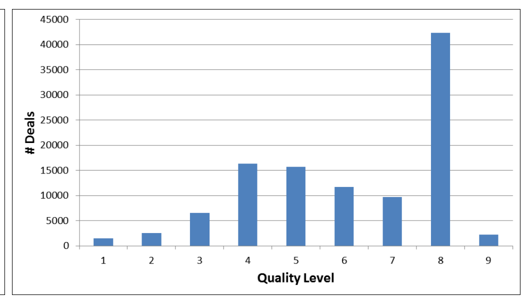

(d) Quality Level Distribution

Figure 7: Deal Characteristics: Restaurant Deals

5.1 Findings

We estimate the two stage dynamic model with Dirichlet learning using the data and present the results in Table 5. The results of the proposed model are intuitive. Lower price (coef. = -0.571, s.e.= 0.060) and better quality (coef. = 0.160, s.e. = 0.024) are related to higher purchase likelihood. The cost of search is positive and significant (coef. = 3.009, s.e. = 0.011). This indicates that consumers incur a cost when they click on the deal to go to the webpage. And this cost counterbalances the benefit of learning which would result in the decreasing trend of clicks. The effect of previous purchases is negative and significant (coef. = -0.034, s.e. = 0.013). This effect shows that "inven-

[22]A Pearson product-moment correlation coefficient was computed to assess the relationship between price and quality. There is a very weak negative correlation between the two variables for all category deals, $r = -0.178$, $df = 527300$, $p < .000$ and is even weaker for restaurant deals, $r = -0.017$, $df = 108700$, $p < .000$)

tory" discourages purchases on Groupon (Seiler, 2013). For the estimation of Dirichlet learning, we divided the priors to 9 grids from low to high values.The grids of 1 to 8 are not different from zero. Grid 9 is significant and has a large value (coef = 1.307, s.e. = 0.178). This means that in consumers' prior beliefs, merchants on Groupon are of high quality. They adjust this belief as they learn about the clicked deals. This makes sense as the data is about new subscribers. If they did not think the quality of deals are good, it is unlikely they would register to shop on Groupon at all.

Table 5: Estimation Results

	FL w/o learn [a]		**FL w/ Dirichlet learn**	
	est.	s.e.	est.	s.e.
utility				
Constant	-2.601***	0.670	-2.431***	0.891
Price/100	-0.569***	0.092	-0.571***	0.060
Quality	0.161***	0.023	0.160***	0.024
# Previous Purchases: η	-0.036***	0.014	-0.034***	0.013
Baseline cost	3.021***	0.088	3.009***	0.011
Dirichlet Priors				
α_1	-	-	4.003E-05	1.58E-04
α_2	-	-	3.999E-05	1.54E-04
α_3	-	-	3.999E-05	1.55E-04
α_4	-	-	3.995E-05	1.66E-04
α_5	-	-	3.999E-05	1.57E-04
α_6	-	-	3.999E-05	1.78E-04
α_7	-	-	3.999E-05	2.02E-04
α_8	-	-	3.999E-05	2.12E-04
α_9	-	-	1.307***	0.178
Observations	108,698		108,698	
-LL	22341.560		22262.826	
AIC	44693.120		44553.650	
BIC	44741.100		44688.000	

Notes: *** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

[a] We assume consumers are forward-looking without learning the quality distribution and the prior belief is what we estimated using forward-looking with Dirichlet learning model (i.e., a high prior).

As comparison, we use the two stage dynamic search model (without learning) as a benchmark (Seiler, 2013). We assume that consumers do not learn about the distribution of deals. Consumers are forward-looking and maximize their expected present-discounted sum of utilities. The click and purchase decisions are determined by the price and quality of current deal and the number of previous purchased deals. Consumers still face an intertemporal tradeoff - purchase now or wait for

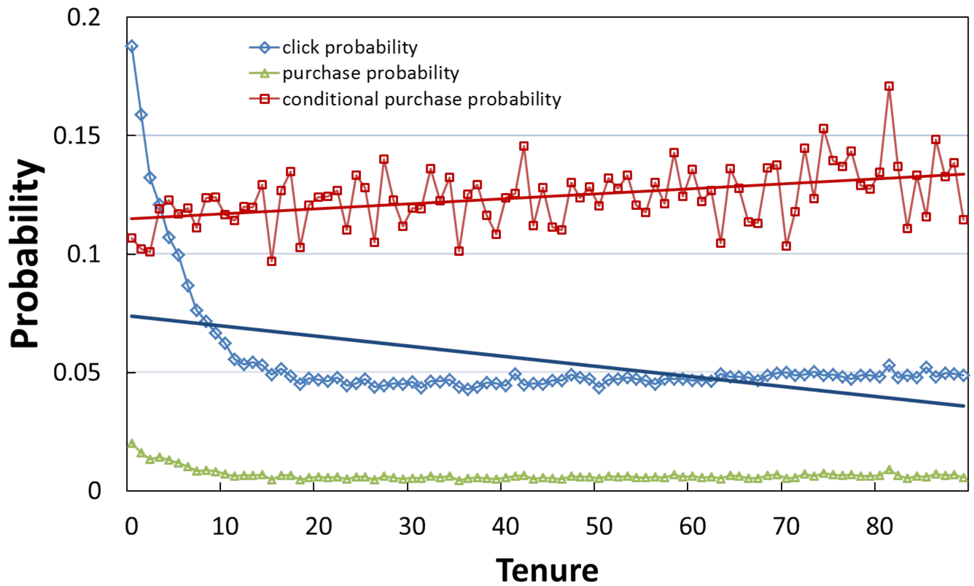

Figure 8: Simulated Data Trends

the next deal. Though consumers have knowledge of the price distribution and have beliefs about the quality distribution, they still face the uncertainty regarding the future deals and recognize that purchasing now lowers their utility from future purchases. Most state variables remains the same, only that now we do not have the state variables relating to learning (i.e., α_{it}^{pr} and α_{it}^{po}). To be specific, the click stage state variables are $x_{cs,it} = \{P_{it}, pp_{it}\}$ and the purchase stage state variables are $x_{ps,it} = \{P_{it}, Q_{it}, pp_{it}\}$. In this model, we assume consumers know the empirical distribution of deal prices, and they have a high prior belief (what we estimated using our proposed model as in Table 5) on the deal quality distribution. Since they are not learning, this belief will not be updated. As in the proposed model, both P_{it} and Q_{it} are observed by both consumers and researchers. P_{it} comes from a normal distribution with the empirical mean and empirical standard deviation. Q_{it} comes from the empirical categorical distribution. The estimates are comparable to the proposed model, with a negative effect of price and number of previous purchased deals, positive effect of quality and positive search cost.

The model with learning has higher likelihood and lower values of AIC and BIC, indicating a better fit to the data. Allowing for consumers to learn the quality distribution over time appears to better reflect the variation in the data than a model without learning. Based on the estimation results, we then simulate each individual's click and purchase behaviors. As in Figure 4, we plot the click probability and conditional purchase probability against the tenure with Groupon in Figure 8. We see that using the estimated model, we are able to successfully replicate the two observed data patterns.

Consumers click less since there is a reducing incentive to learn about quality. This must mean that the deals they click on over time should be increasing in quality as consumers get closer to and learn about the true quality distribution of deals on Groupon. Indeed this what the simulation from our proposed model indicates; specifically, with beliefs about quality getting closer to the true distribution, the quality of deals clicked on improves with tenure (Figure 9). This is a further validation of the model.

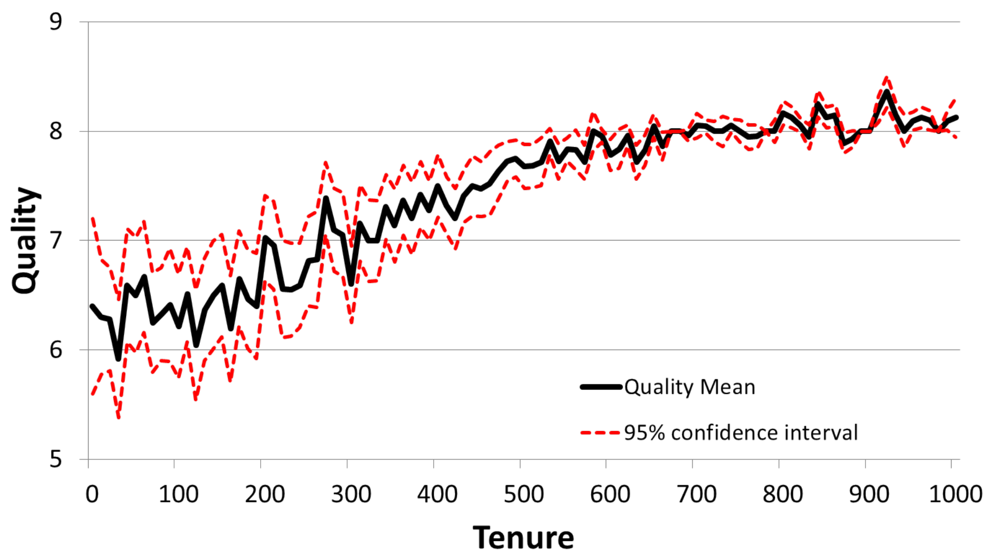

Figure 9: Mean Quality of Purchased Deals

5.2 Model without Learning

What is the benefit of incorporating learning in our context? As we have already shown, the model that does not incorporate subscriber learning does not fit the data as well as the proposed model. Further, simulating data from the model without learning reveals that this model is not able to recover the two observed data patterns, that is the declining click probability and the increasing conditional purchase probability. The reduced-form regession results are presented in Column (III) and (IV) in Table 6. We simulate the behavior of 500 agents for a 60-day period and repeat this process for 1000 times. Within each iteration, we generate data using the structural model specification [23] with the estimates of the proposed model and then test the trend using the same fix effect regression models in Section 3.4.

[23]In this case, proposed model without the learning part.

Without learning, consumers would not exhibit the declining tendency of clicking deals as their expectation or knowledge of the quality of future deals will not be updated so there is no benefit to searching as long as there is a search cost.[24] NOT EXACTLY SURE WHAT THIS MEANS? Furthermore, without learning, the subscribers are not able to choose better deals so purchase probabilities conditional on clicking will not change over time either. Although the estimates are similar, the proposed two stage dynamic search model with learning is able to capture the underlying patterns in the data, while the model without learning is unable to generate either of the observed data patterns.

Table 6: Robustness Check: Reduced Form Analysis of Simulation

	Benchmark Model		**Without Learning** [b]	
	(I) *Click*	(II) *Cond.Purchase* [a]	(III) *Click*	(IV) *Cond.Purchase*
Price/100	-0.0009	-0.042	-0.005	-0.041
	(0.003)	(0.014)	(0.003)	(0.007)
Quality	-0.0002	0.010	-0.0002	0.013
	(0.0000)	(0.0005)	(0.0001)	(0.0003)
pp	-0.037	-0.396	-0.029	-0.150
	(0.0009)	(0.005)	(0.0005)	(0.002)
Tenure	-0.0002	0.002	0.0006	0.003
	(0.0000)	(0.0000)	(0.0014)	(0.005)
Individual FE	Yes	Yes	Yes	Yes

Notes: Standard errors in parentheses. Benchmark model is our proposed model specification. Under each of these models, we simulate the behavior of 500 agents for a 60-day period and we repeat this process for 1000 times. Within each iteration, we generate data using the structural model specification with the estimates of the proposed model and then test the trend using the same fix effect regression models as in Column (I) and (III) of Table 2.

[a] Conditional on clicking.

[b] The Dirichlet updating part is shut down whereas we still preserve the estimated priors, which is the same with the benchmark model.

5.3 Alternative Measure of Quality

As a robustness check for our measure of quality, we consider several other alternatives. Our data contain information on seller ratings [25] provided by Groupon. There are two issues with this mea-

[24] Mathematically speaking, the total utility of click in Eq.5 is a convex function in quality, Q. So it is increasing in the variance of Q.With high uncertainty of quality in the early days of learning,the utility of click is also high which leads to the high click probability

[25] The ratings on the Groupon website are discrete integer values from 0 to 5.

sure. First, the ratings are just about the sellers of the coupon, i.e. the quality of the restaurant and may not reflect all aspects of the restaurant - deal combination. A deal from a highly rated restaurant can be viewed as bad due to reasons such as a short redemption period. The quality in our context also measures how good the deal is not just how good the restaurant is. Second, the Groupon ratings do not have much variation. The average is 4, with s.d. of 0.36. Among all the deals, 88% have a rating of 4. So we supplemented the ratings information with the number of referrals for each deal and came up with alternative measure of quality. A better deal is more likely to be recommended to other people through Facebook, Twitter or email. We combine the referral and rating information to come up with a composite metric of quality that correlates with the conditional purchase probability ($r = 0.61$, $p < .001$). We use this measure of quality to estimate the model and report the results in Table 7. Comparing these results to the results using the preferred operationalization, we find that the latter measure has a better fit, while the estimation results remains similar except for the point estimate of quality.

Table 7: Robustness Check: Alternative Quality Measure

	FL w/o learn [a]		**FL w/ Dirichlet learn**	
	est.	s.e.	est.	s.e.
utility				
Constant	-2.621***	0.611	-2.443***	0.732
Price/100	-0.561***	0.078	-0.569***	0.062
Quality	0.251**	0.116	0.248**	0.121
# Previous Purchases: η	-0.037***	0.016	-0.036***	0.014
Baseline cost	3.001***	0.121	3.000***	0.098
Dirichlet Priors [b]				
α_1	-	-	5.001E-04	1.78E-04
α_2	-	-	5.002E-04	1.82E-04
α_3	-	-	5.002E-04	1.69E-04
α_4	-	-	5.001E-04	1.80E-04
α_5	-	-	4.999E-04	1.85E-04
α_6	-	-	1.005***	0.237
Observations	108,698		108,698	
-LL	22455.120		22312.912	
AIC	44920.240		44647.820	
BIC	44968.220		44753.380	

Notes: *** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

[a] We assume consumers are forward-looking without learning the quality distribution and the prior belief is what we estimated using forward-looking with Dirichlet learning model (i.e., a high prior).

[b] The new quality measure use $(\log(Rating)+referral)$ to construct six quality bins. Therefore, Dirichlet prior vector $\vec{\alpha}$ has six parameters in this case; whereas our proposed quality measure has nine.

6 Managerial Implications

Using our proposed model, we test out two strategies related to consumers' search and learning behaviors on Groupon. Our objective is to increase the unconditional probability of purchase so as to improve the company's revenue.We assume that each successfully recommended deal purchased by the user would create 50% net revenue from the current deal price[26]. In that way, we can compute the total revenue for each method. The results are in Table 8.

First, we eliminate the search cost to examine its impact. When the subscriber clicks on the deals in the newsletter, currently (s)he will be redirected to the webpage of the deal on Groupon's website. More information is then revealed on the webpage. Subscribers therefore incur a search

[26]http://www.investopedia.com/articles/active-trading/080515/how-groupon-makes-money.asp

cost to navigate to the new page and to acquire the information therein - this includes the load time of web page as well as the information processing time. Researchers (Gu, 2017) and practitioners have long noted the negative correlation between web load time and conversion rate [27]. What if the information of deals is fully revealed in the newsletter and customers can make a purchase without going through the webpage of the deal? That is, what if we eliminate the search cost or at least a part of it? We consider the extreme case of no search costs and simulate the purchase decision using the proposed model under this policy change. Since the search step is now eliminated, we compare the click probability, unconditional purchase probability and revenue in Table 8. On average, we see an increase in the purchase probability from about 0.17% to 3.3%, with the average revenue increasing from 297 USD to 6245 USD.

We further investigate the role of learning by providing customers with different types of deals in the newsletter so as to induce different levels of uncertainty regarding the quality distribution (along with the high prior beliefs that we estimated from the data). Keeping the mean quality the same, We consider two conditions: 1) The real quality distribution from the data; 2) Groupon provides two types of deals with equal probability, either extremely good (quality rating is 9) or extremely bad (quality rating is 1) deals. The variance of the quality distribution of the second is larger than the first, indicating that consumers would face more uncertainty in the second situation. The results are shown in Table 8. We find that increasing the uncertainty of quality distribution will result in a higher unconditional purchase rate (0.17% vs. 0.25%) and more revenue (297 USD vs. 442 USD) . The reason is that with only extreme draws from the quality distribution, the learning process slows down as uncertainty is high. This leads to a higher click rate. Since subscribers start with high prior beliefs that they are not updating much, it leads to better purchase outcomes for the firm. A similar explanation of high variance obscuring learning is also provided in Sriram et al. (2015). On the flip side, one way to accelerate learning is to send newsletters more frequently. This will lead to the opposite effects as increasing the uncertainty.

To summarize, based on our counterfactual analysis, we recommend that Groupon 1) eliminate or lower the search cost by revealing as much information about the deal in the newsletter and allowing for direct purchases from there; 2) increase the variance of deal qualities in the newsletters it sends out (given the high prior beliefs of new subscribers).

[27] $https://www.youtube.com/watch?v=OpMfxzie2g$

Table 8: Counter-factual Analysis

	Click Rate	**Purchase Rate**	**Revenue (USD)**
Change Click Cost			
W/ Cost [a]	0.0536	0.0017	297
W/O Cost	0.4652	0.0330	6245
Change Distribution [b]			
Distribution: Empirical [a]	0.0536	0.0017	297
Distribution: Extreme	0.0726	0.0025	442

Notes: We simulate the behavior of 500 agents for a 60-day period and we repeat this process for 1000 times. The click (purchase) rate and revenues are the average of 30 simulations. Note that the revenue column stands for an average total revenue of 500 new subscribers for a 60-day period.

[a] Benchmark model using estimates in Table 5.

[b] The two distributions investigated here have the same mean (6.14) but different variance: empirical (3.76) vs. extreme (5.50)

7 Discussion

Analysts have been gloomy about the prospects for daily deal websites and have attributed this, at least in part, to what they call "daily deal fatigue" (Dholakia and Kimes, 2011). The empirical data presented in this paper do support the behavior underlying deal fatigue; at the same time we also find a reason for some optimism. In particular, we observe a declining probability that a consumer clicks on a merchant in the emailed newsletter over time but an increasing probability, over time, that the consumer makes a purchase conditional on clicking. The former trend is consistent with the notion of "fatigue". And the latter is a potential source of optimism for Groupon. Together, these patterns suggest that even though the consumer is getting more selective in terms of exploring the offers received, (s)he is more likely to yield the site revenue as time passes.

We propose a model of search and learning that tries to explain these observations. Based on the model estimates, we are able to replicate the two observed patterns observed in the data. So in daily deal websites, consumers who are uncertain about the quality of future deals constantly face the tradeoff between whether to click on the deal and purchase it now or to wait for the next one. Consumers tend to click on more deals in the beginning to learn about the quality distribution of deals. As their knowledge accumulates, the incentive to learn declines so they click fewer deals. Over time, since clicking is less about learning the quality of deals, the motivation for clicking is more likely to be to actually purchase the deal - hence the data pattern of increasing purchase probabilities. Through simulations, we show that forward looking and learning are two conditions that help us replicate the data patterns.Using the structural model allowing for the above aspects of subscriber behavior, we propose different strategies that Groupon may consider in order to increase revenue. We propose two strategies, eliminating or reducing search costs, and increasing the uncertainty in deal qualities to prolong the learning process.

From the methodological perspective, we combine a dynamic two stage search model with Dirichlet learning on the distribution of the entity that consumers are uncertain about. Compared to other parametric learning models, such non-parametric learning is flexible and maybe appropriate in contexts where consumers are not learning about the true quality of a specific product or service.

Due to the constraint on the data, we only look at consumer behavior on the Groupon website and that too with a focus on restaurant deals in the empirical analysis. As future work, we will extend to more categories of deals while allowing for the possibility that subscribers may be learning about specific categories as well as the site as a whole.

References

Berndt, Ernst R, Bronwyn H Hall, Robert E Hall, Jerry A Hausman. 1974. Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement, Volume 3, number 4*. NBER, 653–665.

Chen, Yuxin, Song Yao. 2015. Sequential search with refinement: Model and application with click-stream data. *Available at SSRN 2130646* .

Ching, Andrew T, Hyunwoo Lim. 2016. A structural model of correlated learning and late-mover advantages: The case of statins (october 18, 2016). *Rotman School of Management Working Paper No. 2662286.* .

Crawford, Gregory S, Matthew Shum. 2005. Uncertainty and learning in pharmaceutical demand. *Econometrica* **73**(4) 1137–1173.

De los Santos, Babur, Ali Hortaçsu, Matthijs R Wildenbeest. 2012. Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review* **102**(6) 2955–2980.

De los Santos, Babur, Ali Hortaçsu, Matthijs R Wildenbeest. 2013. Search with learning. *Available at SSRN 2163369* .

Dholakia, Utpal M. 2010. How effective are groupon promotions for businesses? *Available at SSRN 1696327* .

Dholakia, Utpal M. 2011. What makes groupon promotions profitable for businesses? *Available at SSRN 1790414* .

Dholakia, Utpal M. 2012. How businesses fare with daily deals as they gain experience: A multi-time period study of daily deal performance. *Available at SSRN 2091655* .

Dholakia, Utpal M, Sheryl E Kimes. 2011. Daily deal fatigue or unabated enthusiasm? a study of consumer perceptions of daily deal promotions. *A Study of Consumer Perceptions of Daily Deal Promotions (September 11, 2011)* .

Erdem, Tülin, Susumu Imai, Michael P Keane. 2003. Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics* **1**(1) 5–64.

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

Gu, Chris Naiqing. 2017. Consumer online search with partially revealed information. *Working Paper* .

Hartmann, Wesley R, Harikesh S Nair. 2010. Retail competition and the dynamics of demand for tied goods. *Marketing Science* **29**(2) 366–386.

Hendel, Igal, Aviv Nevo. 2006. Measuring the implications of sales and consumer inventory behavior. *Econometrica* **74**(6) 1637–1673.

Hong, Han, Matthew Shum. 2006. Using price distributions to estimate search costs. *The RAND Journal of Economics* **37**(2) 257–275.

Honka, Elisabeth. 2014. Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics* **45**(4) 847–884.

Honka, Elisabeth, Pradeep K Chintagunta. 2015. Simultaneous or sequential? search strategies in the us auto insurance industry. *Search Strategies in the US Auto Insurance Industry (November 20, 2015)* .

Hortaçsu, Ali, Chad Syverson. 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly Journal of Economics* **119**(2) 403–456.

Hu, Mantian Mandy, Russell S Winer. 2016. The tipping point feature of social coupons: An empirical investigation. *International Journal of Research in Marketing* .

Keane, Michael P, Kenneth I Wolpin. 1994. The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence. *The Review of Economics and Statistics* 648–672.

Kim, Jun B, Paulo Albuquerque, Bart J Bronnenberg. 2010. Online demand under limited consumer search. *Marketing science* **29**(6) 1001–1023.

Koulayev, Sergei. 2010. Estimating demand in online search markets, with application to hotel bookings. *Federal Reserve Bank of Boston Working Paper* (09-16).

Koulayev, Sergei. 2013. Search with dirichlet priors: estimation and implications for consumer demand. *Journal of Business & Economic Statistics* **31**(2) 226–239.

Luo, Xueming, Michelle Andrews, Yiping Song, Jaakko Aspara. 2014. Group-buying deal popularity. *Journal of Marketing* **78**(2) 20–33.

McCall, John Joseph. 1970. Economics of information and job search. *The Quarterly Journal of Economics* 113–126.

Orhun, A Yeşim, Sriram Venkataraman, Pradeep K Chintagunta. 2015. Impact of competition on product decisions: Movie choices of exhibitors. *Marketing Science* **35**(1) 73–92.

Rothschild, Michael. 1974. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy* **82**(4) 689–711.

Rust, John. 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.

Rust, John. 1994. Structural estimation of markov decision processes. *Handbook of econometrics* **4** 3081–3143.

Seiler, Stephan. 2013. The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics* **11**(2) 155–203.

Song, Inseong, Pradeep K Chintagunta. 2003. A micromodel of new product adoption with heterogeneous and forward-looking consumers: Application to the digital camera category. *Quantitative Marketing and Economics* **1**(4) 371–407.

Sriram, S, Pradeep K Chintagunta, Puneet Manchanda. 2015. Service quality variability and termination behavior. *Management Science* **61**(11) 2739–2759.

Stigler, George J. 1961. The economics of information. *The journal of political economy* 213–225.

Weitzman, Martin L. 1979. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society* 641–654.

Wu, Chunhua, Yitian Sky Liang, Xinlei Chen. 2015. Is daily deal a good deal for merchants? an empirical analysis of economic value in the daily deal market. *An Empirical Analysis of Economic Value in the Daily Deal Market (June 30, 2015)* .

Zhang, Juanjuan. 2010. The sound of silence: Observational learning in the us kidney market. *Marketing Science* **29**(2) 315–335.

Zhao, Yi, Sha Yang, Vishal Narayan, Ying Zhao. 2013. Modeling consumer learning from online product reviews. *Marketing Science* **32**(1) 153–169.

Appendix A: Value Functions

The choice-specific value functions at purchase stage, with regard to purchase ($d_{it}^P = 1$) or not purchase ($d_{it}^P = 0$), are:

$$\begin{aligned} v_{ps,it}^1 &= v_{ps}^1(x_{ps,it}, \tilde{\varepsilon}_{ps,it}) \\ &= \bar{u}_{ps,it}^1 + \varepsilon_{ps,it}^1 + \delta \cdot E\left[\max\{v_{cs}^0(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}), v_{cs}^1(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1})\} | x_{ps,it}, \tilde{\varepsilon}_{ps,it}, d_{it}^P = 1\right] \end{aligned} \tag{21}$$

$$\begin{aligned} v_{ps,it}^0 &= v_{ps}^0(x_{ps,it}, \tilde{\varepsilon}_{ps,it}) \\ &= \bar{u}_{ps,it}^0 + \varepsilon_{ps,it}^0 + \delta \cdot E\left[\max\{v_{cs}^0(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}), v_{cs}^1(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1})\} | x_{ps,it}, \tilde{\varepsilon}_{ps,it}, d_{it}^P = 0\right] \end{aligned} \tag{22}$$

where $\bar{u}_{ps,it}^j$ ($j \in \{0,1\}$) is the utility excluding the error term, and $v_{cs,it}^j$ and $v_{ps,it}^j$ are the choice-specific value functions for the click and purchase stage respectively. We denote the state variables at click and purchase stage as $x_{cs,it}$ and $x_{ps,it}$, respectively. The value function in the click stage next period is a function of the state variable $x_{cs,it+1}$ and error term $\tilde{\varepsilon}_{cs,it+1}$. The consumer forms expectations about these variables based on current state variables $x_{ps,it}$, error terms $\tilde{\varepsilon}_{ps,it}$, and decision d_{it}^P in the purchase stage.

The same rationale applies to the click stage. Upon opening the newsletter, if the consumer decides to click, (s)he receives the flow utility at this stage. And because (s)he forward-look on the purchase stage, (s)he also gets an expected maximum utility (s)he can get conditional on such a click decision.

$$
\begin{aligned}
v^1_{cs,it} &= v^1_{cs}(x_{cs,it}, \tilde{\varepsilon}_{cs,it}) \\
&= \bar{u}^1_{cs,it} + \varepsilon^1_{cs,it} + E\left[\max\{v^0_{ps}(x_{ps,it}, \tilde{\varepsilon}_{ps,it}), v^1_{ps}(x_{ps,it}, \tilde{\varepsilon}_{ps,it})\} | x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 1\right]
\end{aligned} \tag{23}
$$

Similarly, if (s)he does not click, (s)he knows for sure (s)he will enter into the click stage the next period. Therefore, (s)he forward-look on the click stage of next period. The utility (s)he gets comprises of the immediate flow utility plus the discounted expected maximum utility (s)he can get the next period.

$$
\begin{aligned}
v^0_{cs,it} &= v^0_{cs}(x_{cs,it}, \tilde{\varepsilon}_{cs,it}) \\
&= \bar{u}^0_{cs,it} + \varepsilon^0_{cs,it} + \delta \cdot E\left[\max\{v^0_{cs}(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}), v^1_{cs}(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1})\} | x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 0\right]
\end{aligned} \tag{24}
$$

Note that the expectations regarding the purchase-stage choice-specific value functions are not multiplied by the discounted factor. Despite the purchase decision happening shortly after the search decision, more information will be available to the consumer. This information structure makes including the *expected* utility from the purchase stage in the current time period into the click-stage function necessary.

Now we write the value functions, which are expressed in terms of their choice-specific counterparts.

$$
\begin{aligned}
V_{cs,it} &= V_{cs}(x_{cs,it}, \tilde{\varepsilon}_{cs,it}) \\
&= \max\{v^0_{cs,it}, v^1_{cs,it}\} \\
&= \max\{\bar{v}^0_{cs,it} + \varepsilon^0_{cs,it}, \bar{v}^1_{cs,it} + \varepsilon^1_{cs,it}\}
\end{aligned} \tag{25}
$$

$$
\begin{aligned}
V_{ps,it} &= V_{ps}(x_{ps,it}, \tilde{\varepsilon}_{ps,it}) \\
&= \max\{v^0_{ps,it}, v^1_{ps,it}\} \\
&= \max\{\bar{v}^0_{ps,it} + \varepsilon^0_{ps,it}, \bar{v}^1_{ps,it} + \varepsilon^1_{ps,it}\}
\end{aligned} \tag{26}
$$

As in Rust (1987), we define the expectation of value function at click (purchase) stage $EV_{cs,it}$ ($EV_{ps,it}$), integrated over the realization of $\tilde{\varepsilon}_{cs,it}$ ($\tilde{\varepsilon}_{ps,it}$) as

$$EV_{cs,it} = EV_{cs}(x_{cs,it}) = \int_{\tilde{\varepsilon}_{cs,it}} V_{cs,it}(x_{cs,it}, \tilde{\varepsilon}_{cs,it}) \cdot dF_{\tilde{\varepsilon}_{cs,it}} \tag{27}$$

$$EV_{ps,it} = EV_{ps}(x_{ps,it}) = \int_{\tilde{\varepsilon}_{ps,it}} V_{ps,it}(x_{ps,it}, \tilde{\varepsilon}_{ps,it}) \cdot dF_{\tilde{\varepsilon}_{ps,it}} \tag{28}$$

Also following Rust (1987), we make the following conditional independence assumptions:

$$\begin{aligned}
p(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}|x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 0) &= p(\tilde{\varepsilon}_{cs,it+1}|x_{cs,it+1}) \cdot p(x_{cs,it+1}|x_{cs,it}, d^C_{it} = 0) \\
p(x_{ps,it}, \tilde{\varepsilon}_{ps,it}|x_{cs,it}, \tilde{\varepsilon}_{cs,it}, d^C_{it} = 1) &= p(\tilde{\varepsilon}_{ps,it}|x_{ps,it}) \cdot p(x_{ps,it}|x_{cs,it}, d^C_{it} = 1) \\
p(x_{cs,it+1}, \tilde{\varepsilon}_{cs,it+1}|x_{ps,it}, \tilde{\varepsilon}_{ps,it}, d^P_{it} = j) &= p(\tilde{\varepsilon}_{cs,it+1}|x_{cs,it+1}) \cdot p(x_{cs,it+1}|x_{ps,it}, d^P_{it} = j), \quad j \in \{0, 1\}
\end{aligned} \tag{29}$$

Furthermore, by assuming the error terms ($\tilde{\varepsilon}_{cs,it}$, $\tilde{\varepsilon}_{ps,it}$) follow iid extreme value distribution, we yield closed-form expected value functions:

$$\begin{aligned}
EV_{cs,it} &= EV_{cs}(x_{cs,it}) \\
&= \log\{\exp(\bar{v}^0_{cs,it}) + \exp(\bar{v}^1_{cs,it})\} \\
&= \log\Big\{\exp\Big(\bar{u}^0_{cs,it} + \delta \cdot E\Big[EV_{cs}(x_{cs,it+1})|x_{cs,it}, d^C_{it} = 0\Big]\Big) + \\
&\qquad \exp\Big(\bar{u}^1_{cs,it} + E\Big[EV_{ps}(x_{ps,it})|x_{cs,it}, d^C_{it} = 1\Big]\Big)\Big\}
\end{aligned} \tag{30}$$

$$\begin{aligned}
EV_{ps,it} &= EV_{ps}(x_{ps,it}) \\
&= \log\{\exp(\bar{v}^0_{ps,it}) + \exp(\bar{v}^1_{ps,it})\} \\
&= \log\{\exp(\bar{u}^0_{ps,it} + \delta \cdot E[EV_{cs}(x_{cs,it+1})|x_{ps,it}, d^P_{it} = 0]) + \\
&\qquad \exp(\bar{u}^1_{ps,it} + \delta \cdot E[EV_{cs}(x_{cs,it+1})|x_{ps,it}, d^P_{it} = 1])\}
\end{aligned} \tag{31}$$

Following Eq. 12, Eq. 13 and the transition rules of respective state variables, we can explicitly write down the expected value functions as following [28].

[28]Note that the notation $\vec{\alpha}^{po}_{it}(Q_{it})$ means to apply the Dirichlet updating rule to the vector $\vec{\alpha}^{po}_{it}$ with the quality signal Q_{it}. Notation $\vec{\alpha}^{po}_{it}(q_n)$ has similar meaning.

$$
\begin{aligned}
EV_{cs,it} &= EV_{cs}(P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}) \\
&= \log\Big\{\exp\Big(\bar{u}_{cs,it}^{0} + \delta \cdot E\Big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr}, pp_{it+1})|P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}, d_{it}^{C} = 0\Big]\Big) + \\
&\qquad \exp\Big(\bar{u}_{cs,it}^{1} + E\Big[EV_{ps}(P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it})|P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}, d_{it}^{C} = 1\Big]\Big)\Big\} \\
&= \log\Big\{\exp\Big(\bar{u}_{cs,it}^{0} + \delta \cdot E\Big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr} = \vec{\alpha}_{it}^{pr}, pp_{it+1})|P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}, d_{it}^{C} = 0\Big]\Big) + \\
&\qquad \exp\Big(\bar{u}_{cs,it}^{1} + E\Big[EV_{ps}(P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}(Q_{it}), pp_{it}|P_{it}, \vec{\alpha}_{it}^{pr}, pp_{it}, d_{it}^{C} = 1\Big]\Big)\Big\} \\
&= \log\left\{\exp\left(\eta \cdot pp_{it} + \delta \cdot \int_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{pr}, pp_{it+1}) \cdot dF(P_{it+1})\right) + \right. \\
&\qquad \left. \exp\left(-c_{it} + \int_{Q_{it}} EV_{ps}(P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}(Q_{it}), pp_{it}) \cdot dF(Q_{it})\right)\right\} \\
&= \log\left\{\exp\left(\eta \cdot pp_{it} + \delta \cdot \int_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{pr}, pp_{it+1}) \cdot dF(P_{it+1})\right) + \right. \\
&\qquad \left. \exp\left(-c_{it} + \sum_{n=1}^{N} \left(EV_{ps}(P_{it}, \vec{\alpha}_{it}^{po}(q_n), pp_{it}) \cdot \rho_e^n\right)\right)\right\}
\end{aligned} \tag{32}
$$

$$
\begin{aligned}
EV_{ps,it} &= EV_{ps}(P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}) \\
&= \log\big\{\exp\big(\bar{u}_{ps,it}^{0} + \delta \cdot E\big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr}, pp_{it+1})|P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}, d_{it}^{P} = 0\big]\big) + \\
&\qquad \exp\big(\bar{u}_{ps,it}^{1} + \delta \cdot E\big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr}, pp_{it+1})|P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}, d_{it}^{P} = 1\big]\big)\big\} \\
&= \log\big\{\exp\big(\bar{u}_{ps,it}^{0} + \delta \cdot E\big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr} = \vec{\alpha}_{it}^{po}, pp_{it+1})|P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}, d_{it}^{P} = 0\big]\big) + \\
&\qquad \exp\big(\bar{u}_{ps,it}^{1} + \delta \cdot E\big[EV_{cs}(P_{it+1}, \vec{\alpha}_{it+1}^{pr} = \vec{\alpha}_{it}^{po}, pp_{it+1})|P_{it}, Q_{it}, \vec{\alpha}_{it}^{po}, pp_{it}, d_{it}^{P} = 1\big]\big)\big\} \\
&= \log\left\{\exp\left(\eta \cdot pp_{it} + \delta \cdot \int_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{po}, pp_{it+1}) \cdot dF(P_{it+1})\right) + \right. \\
&\qquad \left. \exp\left(\beta_0 + \beta_P \cdot P_{it} + \beta_Q \cdot Q_{it} + \eta \cdot pp_{it} + \delta \cdot \int_{P_{it+1}} EV_{cs}(P_{it+1}, \vec{\alpha}_{it}^{po}, pp_{it+1}) \cdot dF(P_{it+1})\right)\right\}
\end{aligned} \tag{33}
$$

Appendix B: Transition Rules

The detailed transition rule is illustrated in the following table. The first column is the possible transitions and the transit state variables. The second column is the corresponding transition rule.

Transition	Transition Rule
1. $x_{cs,it} \to x_{cs,it+1}$	
$P_{it} \to P_{it+1}$	$P_{it+1} \sim N(\mu_p, \sigma_p)$. Consumer‘s belief on price is always rational and follows the empirical distribution of price. μ_p (σ_p) is the mean (s.d.) of the empirical price distribution.
$\vec{\alpha}_{it}^{pr} \to \vec{\alpha}_{it+1}^{pr}$	$\vec{\alpha}_{it+1}^{pr} = \vec{\alpha}_{it}^{po} = \vec{\alpha}_{it}^{pr}$ Since there is no updating at time t, the knowledge on the quality distribution remains unchanged at the beginning of $t+1$.
$pp_{it} \to pp_{it+1}$	$pp_{it+1} = pp_{it} + \Delta pp_{it}$ [a]
2. $x_{cs,it} \to x_{ps,it}$	
$P_{it} \to P_{it}$	The purchase stage price value remains unchanged since it's the same deal.
$\vec{\alpha}_{it}^{pr} \to \vec{\alpha}_{it}^{po}$	The update of the "knowledge" about quality distribution upon receiving the quality signal (Q_{it}) follows Dirichlet learning rule which we will explain in detail in Section 4.2
Q_{it}	$Q_{it} \sim f_{emp}(Q)$ where $f_{emp}(Q)$ is the empirical distribution of quality levels on Groupon website.
$pp_{it} \to pp_{it}$	pp_{it} remain unchanged since it's within the same period
3. $x_{ps,it} \to x_{cs,it+1}$	
$P_{it} \to P_{it+1}$	$P_{it+1} \sim N(\mu_p, \sigma_p)$. Consumer's belief on price is always rational and follows the empirical distribution of price. μ_p (σ_p) is the mean (s.d.) of the empirical price distribution.
$\vec{\alpha}_{it}^{po} \to \vec{\alpha}_{it+1}^{pr}$	$\vec{\alpha}_{it+1}^{pr} = \vec{\alpha}_{it}^{po}$. That is, posterior at t becomes the prior at $t+1$.
$pp_{it} \to pp_{it+1}$	$pp_{it+1} = pp_{it} + \Delta pp_{it}$

[a] Δpp_{it} is the number of purchased deals during period t.

Appendix C: Dirichlet Learning Rule

Consumer i s expectation on the probability distribution of quality at time t can be fully represented by the vector:

$$\vec{\alpha}_{it} = \{\alpha_{it}^{1}, \alpha_{it}^{2}, ..., \alpha_{it}^{N}\} \tag{34}$$

More specifically, we denote the prior and posterior in period t as:

$$\vec{\alpha}_{it}^{pr} = \{\alpha_{it}^{pr,1}, \alpha_{it}^{pr,2}, ..., \alpha_{it}^{pr,N}\} \tag{35}$$

$$\vec{\alpha}_{it}^{po} = \{\alpha_{it}^{po,1}, \alpha_{it}^{po,2}, ..., \alpha_{it}^{po,N}\} \tag{36}$$

Now we use a naive example to illustrate the update rule. Suppose there are three quality levels $\{q_1, q_2, q_3\}$ and consumer i has a diffuse, non-informative prior at time t: $\vec{\alpha}_{it}^{pr} = \{\alpha_{it}^{pr,1}, \alpha_{it}^{pr,2}, \alpha_{it}^{pr,3}\} = \{1, 1, 1\}$. According to DLP, the prior expected values of the probability of re-sampling each level are given by $E[\vec{\rho}_t] = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$. If the consumer clicks the deal and is revealed that the deal quality level belongs to q_1, then the posterior is updated as follows: $\alpha_{it}^{po,1} = \alpha_{it}^{pr,1} + 1$, $\alpha_{it}^{po,2} = \alpha_{it}^{pr,2}$, $\alpha_{it}^{po,3} = \alpha_{it}^{pr,3}$. Therefore $\vec{\alpha}_{it}^{po} = \{2, 1, 1\}$ and posterior expected values $E[\vec{\rho}_t] = \{\frac{2}{4}, \frac{1}{4}, \frac{1}{4}\}$.

Appendix D: Literature

Table 9: Literature

		w/o dynamic FL	w/ dynamic FL	
		search model/ optimal stopping	**one-stage**	**two-stage**
no learning		Stigler(1961); McCall(1970); Weitzman(1979); De los Santos, Hortasu &Matthijs(2012); Honka(2014); Honka&Chintagunta(2016); Chen&Yao(2016);	Rust(1987); Erdem, Imai & Keane (2003); Song&Chintagunta(2003); Hendel&Nevo(2006)	Seiler(2013);
learning	**parametric learning**	Zhao, Yan., Narayan, & Zhao(2013); Ching&Lim(2016);	Erdem&Keane(1996); Crowford&Shum(2005); Zhang(2010); Chintagunta, Jiang & Jin(2009); Chintagunta, Goettlerz & Kim (2012);	
	non-parametric learning	Koulayev(2013); De los Santos, Hortasu & Wildenbeestet(2015);		**our model**