Household-Specific Regressions Using Clickstream Data[1]

Avi Goldfarb

University of Toronto

Qiang Lu

University of Sydney

December 20, 2005

[1]We thank the Guest editors Wolfgang Jank and Galit Shmueli, the referees, and participants at the first StatsChallenges in Ecommerce conference for helpful comments. Research support was provided by SSHRC Grant # 538-02-1013. Please address correspondence to agoldfarb@rotman.utoronto.ca. All opinions and errors are our own.

Abstract

This paper makes three contributions: (1) the paper provides a better understanding of online behavior by showing the main drivers of Internet portal choice, (2) the rich data allow for a deeper understanding of brand substitution patterns than previously possible, and (3) the paper introduces a wider statistics community to a new data opportunity and a recently developed method.

1 Introduction

The rapid growth of the Internet since the late 1990s have been an amazing opportunity for marketing and statistics research. We can now observe consumers in a new kind of environment where information search is relatively inexpensive. There are both direct and indirect reasons to be interested in online behavior. The direct reason is that the Internet has become an important element of daily life for millions of people. A growing body of evidence shows that online behavior differs from offline behavior (e.g. Park and Fader 2004, Chen and Hitt 2002, Danaher, Wilson, and Davis 2003). Therefore, it is essential for companies with an online presence to understand what drives online choices. The indirect reason is that Internet data and detail on online choices can inform researchers about more general economic questions such as the cost of search.[1]

Studies of online behavior have one important advantage over previous studies of behavior: rich, detailed data. This clickstream data tracks the website visit history for each online user, thereby enabling a detailed analysis of user choice. One advantage of clickstream data sets is that they are much larger than the data sets typically used to examine consumer behavior. A typical household may buy consumer goods like ketchup and detergent every month or two; however, this same household may

[1] Ellison and Ellison (2005) provide a review of how Internet research has informed a number of general economics and marketing questions.

visit a given website ten times each week or more. Clickstream data can provide us much more information about Internet users, and allow us to do more complicated analysis than traditional data. At the same time, this rich data means it is difficult for managers to use standard methods to quickly arrive at useful insights using their desktop computers.

This paper uses detailed clickstream data on Internet portals to better understand online choices.[2] The paper makes three contributions. First, it provides a better understanding of online behavior by showing the main drivers of Internet portal choice. This relates to the direct reason to study the Internet. In this paper, we show that the success of previous searches is a particularly important driver of website choice. The ability to provide deep searches with many results is less important. This may partially explain how Google, with its page-rank technology and "I feel lucky" button, grew to be dominant in search despite being a late entrant to the portal/search market.

Second, the rich data allow for a deeper understanding of brand substitution patterns than previously possible. This relates to the indirect reason to study the Internet. In our analysis, we show that consumers who rarely change portals often prefer either Yahoo, the most popular portal or the least popular portals such as Go2net and

[2]In this study, we define Internet portal as any site that classifies content and primarily presents itself as a one-stop point-of-entry to content on the web (Hargittai 2000).

Ask Jeeves. In contrast, those that prefer other popular portals such as Excite, MSN, and especially AOL are more likely to change the portal they visit. Many brands seem to have less loyal consumers. Both product features and habit formation are known to be important parts of brand building (e.g. Aaker 1991). Clickstream data provides a method for isolating the relationship between habit formation and brand strength that was not previously possible. In this context, habit formation is shown to be relatively uncorrelated with brand preferences.

Third, the paper introduces a wider statistics community to a new data opportunity and a recently developed method. We examine the choices of Internet portals by 2517 online households over three months. The average household makes 312 portal visits in this time.[3] This rich data set means that we can conduct household-specific regressions. By "household-specific regressions", we mean that we run a separate conditional logit regression on the portal choice of each household in the data. Following Goldfarb (2006), this method allows for more flexible (semiparametric) substitution patterns than the panel methods typically used to study offline consumer choice because all coefficients are allowed to vary at the household level without any assumed distribution. The method has the added advantage of being computationally inexpensive and relatively easy to understand. This means that managers can use it to

[3]The data are available to university-affiliated researchers. If interested in using this data, email Avi Goldfarb at agoldfarb@rotman.utoronto.ca .

quickly understand the behavior of their customers.

The next section details the data set used in the analysis. This is followed by a description of the methodology, the results, and a brief conclusion.

2 Data

The data set, provided by Plurimus Corporation, consists of 2,645,778 website visits by 2651 households from December 27, 1999 to March 31, 2000. For our analysis, we have chosen the top 15 Internet portals visited by 2517 different users with 784,882 observations. This data set is larger than those typically used in marketing and economics to analyze consumer choice. The raw data show the household identification, the website visited, the time of day the visit started and ended (in seconds), and the number of pages viewed at the website.

Despite its richness, the data set has four main limitations. First, it is not geographically representative. Most respondents are drawn from the Pennsylvania, Texas, Florida, and North Carolina. Second, it contains few users at-work. Third, the data are collected at the household level. And fourth, the data do not contain information on households the first time they go online. These limitations mean that results should be extended to different geographic distributions and at-work users with caution. Furthermore, the third limitation means that consumer loyalty may be underestimated if different household members have different preferences. The third

and fourth limitations together mean that it is difficult to study consumer learning.

From this data, measures of loyalty, past number of pages viewed at the website, and past search failure are constructed. Loyalty is measured as a dummy variable for whether that website was visited on the previous choice occasion: it is equal to one if the website was visited on the previous occasion and zero otherwise.[4] Past number of pages viewed is measured by the number of pages viewed at the website on the previous visit. Past search failure is measured by whether the previous search is repeated. In particular, if a household conducts two searches within five minutes for the same goal[5] then it is considered a repeated search. Furthermore, if a household visited two portal sites in a row, and there were less than five minutes between visits, then the first search is considered to have been repeated. It is pages viewed and repeated search during the previous visit to a portal that are relevant for the current choice of which website to visit. Therefore we include data on the experience of the household during its previous visit to the website in the analysis. The use of data on the previous experience at the website means that a correction is required for households that have not yet visited a given website. A *Missing Data* variable is used for this correction. It is equal to one if there is no information about a household's

[4]In the marketing literature, the variable measuring the previous choice is typically called "loyalty" (e.g. Guadagni and Little 1983). More formally, this variable measures "true state dependence".

[5]Over 100 possible goals were identified in the raw data.

previous visit to a website (because the household has not yet visited it) and zero otherwise. This variable serves as a control and has no clear economic interpretation. Table 1 provides univariate descriptive statistics. Table 2 provides pairwise correlation coefficients based on household-level aggregates. In particular, the average values for the variables were calculated at the household level. These are the correlations of the values across the 2517 households in the sample.

3 Method

The decision to visit an Internet portal is modeled as a standard economic discrete choice problem. Users visit the portal that they expect to give them the highest utility. The utility to household i of visiting portal j at time t is:

$$u_{ijt} = X_{ijt}\beta_i + \varepsilon_{ijt} \tag{1}$$

where X_{ijt} includes loyalty, whether the last search was repeated, the number of pages viewed during the last visit to the website (naturally logged), a missing data measure, and the brand dummy variables; β_i is a vector of household-specific coefficients; and ε_{ijt} is an idiosyncratic error term. Assuming that ε_{ijt} follows a type-II extreme value distribution over j and t, we can estimate equation 1 using a separate conditional logit regression (McFadden 1974) for each household in the data. Therefore the probability that household i visits portal j at time t is modeled as:

$$P_{ijt} = \frac{\exp(X_{ijt}\beta_i)}{\sum_{k=1}^{J_i} \exp(X_{ikt}\beta_i)} \tag{2}$$

Typically, the parameters β_i are fixed across households or assumed to follow a known (typically normal) distribution. Household-specific regressions allow the coefficients to vary non-parametrically across households and are computationally much less intensive.[6] In particular, a separate vector of coefficients is estimated for each household.[7] Furthermore, the discrete choice random coefficients models used in these settings may take weeks to converge with millions of observations (and only if the researchers have sufficient RAM in their computers). This matters if managers

[6] A number of previous studies have recommended running regressions on the time-varying dimension of a panel data set when there is sufficient data. Fischer & Nagin (1981) conducted experiments that showed coefficients vary across individuals. Pesaran & Smith (1995) examined employment in 38 industries and concluded (p. 102) that the "lesson for applied work is that when large T panels are available, the individual micro-relations should be estimated separately." Elrod & Haubl (1998), however, highlight the shortcomings of household-specific regressions: inefficiency, overestimation of populaiton variance, and an inability to predict out-of-sample. Therefore, household-specific regressions are best used to understand the underlying drivers of behavior in very large data sets.

[7] Numerous previous studies have shown the logit model to fit consumer choice behavior well (e.g. Guadagni and Little 1983, Jain, Vilcassim, and Chintagunta 1994). Still, these studies do not examine the fit at that household level. Goldfarb (2006) shows that taken together, household-specific regressions provide a better fit than aggregated methods. It is, however, likely that the fit is poor for some specific households in the data.

plan to use clickstream data analysis to make business decisions. Goldfarb (2006) shows that the coefficient estimates are consistent and directly comparable under a mild set of assumptions. Household-specific regressions require one key additional restriction because most households do not visit all 15 portals: only those websites that a particular household actually visits over the course of the sample are included in the analysis. For example, if household i only visits Yahoo, MSN, Excite, and Go.com then only these four portals are included in the regressions for that household. Implicitly, the coefficients on the dummies for the other portals approach negative infinity.

The likelihood function for household i is as follows:

$$\prod_{t=1}^{T_i} \prod_{j=1}^{J_i} (P_{ijt})^{d_{ijt}} \tag{3}$$

where d_{ijt} is equal to 1 if household i purchased brand j at time t, 0 otherwise. Maximum likelihood estimation is applied for each household. Thus, for 2517 households, 2517 sets of parameters are estimated.

4 Results

The results are presented in Tables 3 and 4 and in Figures 1 to 6. The regressions led to 2517 different vectors of parameter estimates. Table 3 presents a univariate description of these parameter estimates. It shows the mean of each coefficient across

households, the standard error of the mean, the standard deviation of the coefficients, and the percentage of coefficients that are significantly different from zero with 95% confidence (positive and negative respectively). Table 3 shows that loyalty and repeated search are important factors in portal choice. The result on repeated search, relative to pages viewed, is particularly interesting. It suggests that it is much more important for portals to direct people to the right website than for portals to provide many results. Information quality is much more important than information quantity. This may partially explain the rise of Google in the years since the data was collected.[8] Even before Google had a large data base, its technology was particularly good at ordering results. This likely reduced the frequency of repeated searches. The coefficients on the portal dummies are all presented relative to Yahoo.[9] Since the mean coefficients are all negative, this suggests that Yahoo is, on average, the preferred portal by a substantial margin (controlling for loyalty, search repeats, and pages viewed). MSN and Netscape also have substantially higher coefficients. The remaining twelve portals have relatively similar coefficient values. This characterization is subject to the important caveat that there is considerable heterogeneity across responses. The standard deviations of the coefficients are quite high relative to the

[8] Google is the 17th most popular portal in the data and is therefore not included in the study.

[9] The portal dummy coefficient distributions are based only on those 2,206 households that visited Yahoo at least once. This ensured that household-level coefficients would be comparable since they would all have the same base (Yahoo).

means.

Table 4 presents the pairwise correlation coefficients of the parameter estimates. It shows a number of interesting results. First, there is a strong negative correlation between Last Search Repeated and Ln(last # Pages), which is quite intuitive: the more a user cares about accuracy of the search results, the less he cares about the number of the search results. Figure 1 presents a scatter plot of these variables. Second, Loyalty does not seem to be correlated with Last Search Repeated or with Ln(last # Pages). Figures 2 and 3 present the scatter plots. Consumers who tend to return to the same portals are no more or less likely to care about search accuracy or search depth.

Third, Loyalty has an interesting relationship with preferences for the various brands. Consumers that rarely change portals have lower opinions of MSN, Excite, Altavista, and especially AOL relative to their opinions of Yahoo. Loyal households prefer Yahoo and Netscape relative to the other major brands. Figure 4 displays a scatter plot of the Loyalty coefficients relative to the coefficients of four popular portals. Interestingly, loyal households are also particularly likely to visit unpopular portals such as Go2Net, Goto, and Ask Jeeves. Figure 5 displays the scatter plot. Many strong brands seem to have less loyal consumers, suggesting that in this context brand strength relates to product features more than simply habit formation. While the extant literature acknowledges the roles of both habit formation and product

features in building brands (e.g. Aaker 1991), the rich clickstream data set used here allows for an assessment of the correlation between brand strength and loyalty that was not previously possible.

Fourth, brand preferences relative to Yahoo are positively correlated. This result emphasizes Yahoo's dominance and is shown in a scatter plot in Figure 6. Unsurprisingly, many brands that are owned by the same company or are linked to each other have highly correlated preference coefficients. For example, preferences for Lycos are highly correlated with Hotbot. Both had the same owner at the time. More interesting is a comparison of the MSN and Netscape correlation coefficients. Those brand coefficients that are highly correlated with Netscape are less correlated with MSN and vice versa. Overall, these results suggest that some portals are substitutes for each other in that they may be used for the same purposes. On the other hand, some portals are complements for each other. People may visit Lycos for searches about celebrities and then Ask Jeeves for science-related questions.

Last, there is a strong negative correlation between Last Search Repeated and Netscape, which means the more important the information accuracy to a household, the less the household visits Netscape; and there is a strong positive correlation between Ln(last # pages) and Netscape, which means the more a household cares about the number of search results, the better image he has on Netscape. At the time, Netscape ordered its search results by the search technology used rather than

by relevance. This provides an important reality check on the results suggesting. People who prefered many results to targeted results prefered a portal that provided many untargeted results.

In summary, examining the relationship between the parameter estimates across households provides a rich insight into consumer behavior that was not previously possible.

5 Conclusions

In summary, this paper provides a better understanding of Internet portal choice. It also shows that the household-level correlation of brand preferences and loyalty provides insights into brand building more generally. Finally, it introduces a wider audience to a rich clickstream data set and to household-specific regression. Clickstream data is an exciting opportunity for statisticians. Future research can proceed in a number of different directions. It can explore how behavior has changed as people get more comfortable online. It can explore how the online environment affects purchase behavior. It can use the data to develop new statistical tools for large data sets. And it can better inform our understanding of general economic and psychological issues. Going forward, there are a number of exciting research opportunities arising from the availability of clickstream data.

References

Aaker, David A. 1991. Managing Brand Equity: Capitalizing on the Value of a Brand Name. New York: The Free Press.

Chen, Pei-Yu and Lorin M. Hitt (2002). Measuring Switching Costs and Their Determinants in Internet Enabled Businesses: A Study of the Online Brokerage Industry. Information Systems Research 13, 255-276.

Danaher, Peter J., Isaac W. Wilson, & Robert A. Davis. 2003. A Comparison of Online and Offline Consumer Brand Loyalty. Marketing Science 22, 461-476.

Ellison, Glenn, and Sara Fisher Ellison. 2005. Lessons About Markets from the Internet. Journal of Economic Perspectives 19, 139-158.

Elrod T, & G. Haubl, 1998, An Extended Random Coefficients Model, With Application to Metric Conjoint Analysis, University of Alberta, Mimeographed.

Fischer, G. W., & D. Nagin, 1981, "Random versus Fixed Coefficient Quantal Choice Models," In C. F. Manski & D. McFadden, eds., Structural Analysis of Discrete Data with Econometric Applications. Cambridge, MA: The MIT Press, 273-304.

Goldfarb, A., 2006, "State Dependence at Internet Portals", Forthcoming Journal of Economics and Management Strategy.

Guadagni, Peter M. and John D. C. Little. 1983. "A Logit Model of Brand Choice Calibrated on Scanner Data." Marketing Science 2, 203-38.

Hargittai, Eszter. 2000. Open Portals or Closed Gates? Channeling Content on the World Wide Web. Poetics 27, 233-254.

Jain, Dipak C., Naufel J. Vilcassim, and Pradeep K. Chintagunta. 1994. A Random-Coefficients Logit Brand-Choice Model Applied to Panel Data. Journal of Business and Economic Statistics 12, 317-328.

McFadden, Daniel. 1974 "Conditional Logit Analysis of Qualitative Behavior." In P. Zarembka, ed., Frontiers of Econometrics. New York: The Academic Press, Inc, 105-142.

Park, Young-Hoon, and Peter S. Fader. 2004. Modeling Browsing Behavior at Multiple Websites. Marketing Science 23, 280-303.

Pesaran, M. H., & R. Smith, 1995, "Estimating long-run relationships from dynamic heterogeneous panels," Journal of Econometrics 68, 79-113.

Table 1: Basic Descriptive Statistics

Variable	Mean	Std. Dev.	Min.	Max.
Loyalty	0.1044	0.3058	0	1
Last Search Repeated	0.2528	0.4346	0	1
Ln(Last # Pages)	0.5873	0.8591	0	7.4565
Missing Data	0.2326	0.4225	0	1

Table 2: Correlation Coefficients of household-level data

	Total number pages visited	Average number pages visited	% searches repeated	% Yahoo	% MSN	% Netscape	% Excite	% AOL	% Altavista	% Iwon	% Lycos	% MyWay	% Go	% Hotbot	% Snap	% Go2Net	% Goto	% Ask Jeeves
Total # Pages	1																	
Average # Pages	0.0088	1																
Search Repeated	0.2502*	-0.0449*	1															
%Yahoo	0.0556*	-0.2341*	-0.0934*	1														
% MSN	-0.0614*	-0.0468*	-0.1759*	-0.3679*	1													
% Netscape	-0.0070	0.1159*	0.0141	-0.2661*	-0.2837	1												
% Excite	0.0732*	0.0665*	0.1335*	-0.1567*	-0.1368*	-0.0563*	1											
% AOL	-0.0182	0.1230*	-0.0318	-0.1535*	-0.0990*	-0.1125*	-0.0435	1										
% Altavista	-0.0156	0.1252*	0.0098	-0.1401*	-0.0864*	-0.1200*	-0.0581*	-0.0495*	1									
% Iwon	0.1051*	0.1598*	0.0140	-0.0961*	-0.0773*	-0.0639*	-0.0349	-0.0486*	-0.0282	1								
% Lycos	-0.0404*	-0.0683*	0.2115*	-0.1408*	-0.1019*	-0.0118	-0.0326	-0.0507*	-0.0389*	-0.0125	1							
% MyWay	-0.0055	0.0276	0.0000	-0.1293*	-0.0830*	-0.0793*	-0.0537*	-0.0468*	-0.0487*	-0.0312	-0.0314	1						
% Go	-0.0669*	0.0870*	-0.0034	-0.1105*	-0.0489*	-0.1006*	-0.0460*	-0.0290	-0.0117	-0.0355	-0.0164	-0.0406	1					
% Hotbot	-0.0283	-0.0442*	0.1836*	-0.1258*	-0.1218*	0.0148	-0.0159	-0.0426*	-0.0445*	-0.0275	0.3495*	-0.0331	-0.0444*	1				
% Snap	0.0597*	0.0113	0.0868*	-0.0890*	-0.0714*	-0.0221	-0.0857	-0.0305	-0.0123	-0.0052	-0.0174	-0.0218	-0.0268	-0.0202	1			
% Go2Net	-0.0155	-0.0834*	0.1186*	-0.0992*	-0.0707*	-0.0388*	-0.1294	-0.0167	-0.0363	-0.0259	-0.0274	-0.0361	-0.0377	-0.0310	-0.0218	1		
% Goto	-0.0564*	0.0036	0.2166*	-0.1449*	-0.0122	-0.0443*	-0.0318	-0.0282	-0.0291	-0.0258	0.0224	-0.0104	0.0311	-0.0098	-0.0286	0.2152*	1	
% Ask Jeeves	-0.0263	0.0231	0.0158	-0.0562*	-0.0757*	-0.0404*	-0.0266	-0.0076	-0.0035	-0.0155	-0.0234	-0.0164	0.0178	-0.0183	-0.0107	-0.0207	-0.0214	1

*significant at 95% confidence level

Table 3: Descriptive Statistics of the 2517 Household-Specific Coefficient Vectors

Variable	Mean Coefficient	Standard Error of the Mean	Standard. Deviation of the Coefficients	% Significantly Positive (95%)	% Significantly Negative (95%)
Variables					
Loyalty	0.979	0.0093	2.7756	69.62	4.41
Last Search Repeated	-0.3906	0.0117	1.4827	3.19	37.14
Ln(Last # Pages)	0.0267	0.0086	2.6014	12.58	6.73
Missing Data	-0.6261	0.0176	5.0659	7.35	21.96
Brand Dummies					
MSN	-0.8124	0.0120	4.8919	28.60	41.77
Netscape	-1.5678	0.0121	5.2511	23.53	48.08
Excite	-2.1220	0.0114	3.0820	8.30	59.43
AOL	-2.0665	0.0102	2.4683	14.40	61.80
Altavista	-2.6039	0.0119	2.3377	4.95	64.77
Iwon	-2.9994	0.0106	2.5799	5.13	69.25
Lycos	-1.9608	0.0114	4.2327	10.66	57.78
MyWay	-2.3776	0.0114	3.6333	6.90	62.99
Go	-2.8004	0.0119	3.4932	7.26	62.38
Hotbot	-2.2778	0.0107	2.1825	7.07	61.41
Snap	-2.8868	0.0106	3.2326	6.00	66.17
Go2Net	-2.1442	0.0114	3.1029	10.52	58.08
Goto	-2.6886	0.0095	2.1573	11.96	64.74
Ask Jeeves	-2.7133	0.0120	3.3165	17.63	57.93

Table 4: Correlation Coefficients of the Household-Specific Coefficients

	Loyalty	Last search repeated	Last # pages visited	Missing	MSN	Netscape	Excite	AOL	Altavista	Iwon	Lycos	MyWay	Go	Hotbot	Snap	Go2Net	Goto	Ask Jeeves
Total # Pages	1																	
Last search repeated	-0.0148	1																
Last # pages visited	-0.0227	**-0.8959***	1															
Missing	0.3339*	-0.7336*	0.7977*	1														
MSN	**-0.0830***	-0.1542*	0.2139*	0.0155	1													
Netscape	**0.0191**	**-0.9851***	**0.9724***	0.9625*	0.6125*	1												
Excite	**-0.1573***	-0.5250*	0.3481*	0.2771*	0.7321*	0.8843*	1											
AOL	**-0.7653***	-0.3874*	-0.7922*	-0.8148*	0.6707*	0.3591*	0.9440*	1										
Altavista	**-0.1562***	-0.4326*	-0.0322	0.0437	0.6978*	0.5695*	-0.1953*	0.5553*	1									
Iwon	**-0.1742***	-0.8082*	0.3413*	0.7965*	0.9085*	0.6732*	0.5979*	0.4937*	0.8876*	1								
Lycos	-0.4566*	-0.2967*	0.3027*	0.2204*	0.8574*	0.7772*	0.8322*	0.4652*	0.5761*	0.6966*	1							
MyWay	0.2769*	-0.4658*	-0.2918*	0.2264*	0.4525*	0.8453*	0.4510*	0.4522*	0.7308*	0.7189*	0.5939*	1						
Go	-0.1818*	0.1738*	-0.2581*	-0.5096*	0.7414*	0.8717*	0.7156*	0.2926*	0.7862*	0.5849*	0.7658*	0.5186*	1					
Hotbot	-0.8237*	-0.3971*	-0.5503*	-0.8181*	0.8194*	0.5073*	0.8579*	0.4938*	0.6192*	0.6519*	**0.8808***	0.7795*	0.8538*	1				
Snap	-0.1150*	-0.5247*	-0.4468*	-0.1919*	0.2648*	0.5674*	0.5718*	0.5072*	0.5350*	0.6014*	0.8250*	0.5646*	0.4597*	0.8258*	1			
Go2Net	**0.3960***	-0.0461	-0.8170*	0.1695*	0.5460*	0.8117*	0.6067*	0.9303*	0.6743*	0.5156*	0.7409*	0.4813*	0.4354*	0.5711*	0.6476*	1		
Goto	0.3616*	-0.4099*	0.5561*	0.8075*	0.7533*	0.8470*	0.3963*	0.4903*	0.8069*	0.5197*	0.3733*	0.6066*	0.5329*	0.5097*	0.3450*	0.4169*	1	
Ask Jeeves	**0.4283***	0.3633*	0.0573	0.4989*	0.5824*	0.4870*	0.4206*	0.4848*	0.3298*	0.5056*	0.4269*	0.4941*	0.6685*	0.5858*	0.3758*	0.4755*	0.4685*	1

*significant at 95% confidence level

Correlations mentioned in the text are written in **bold**.

Figure 1: Scatterplot of the Coefficients on Ln(Last # Pages) and Last Search Repeated

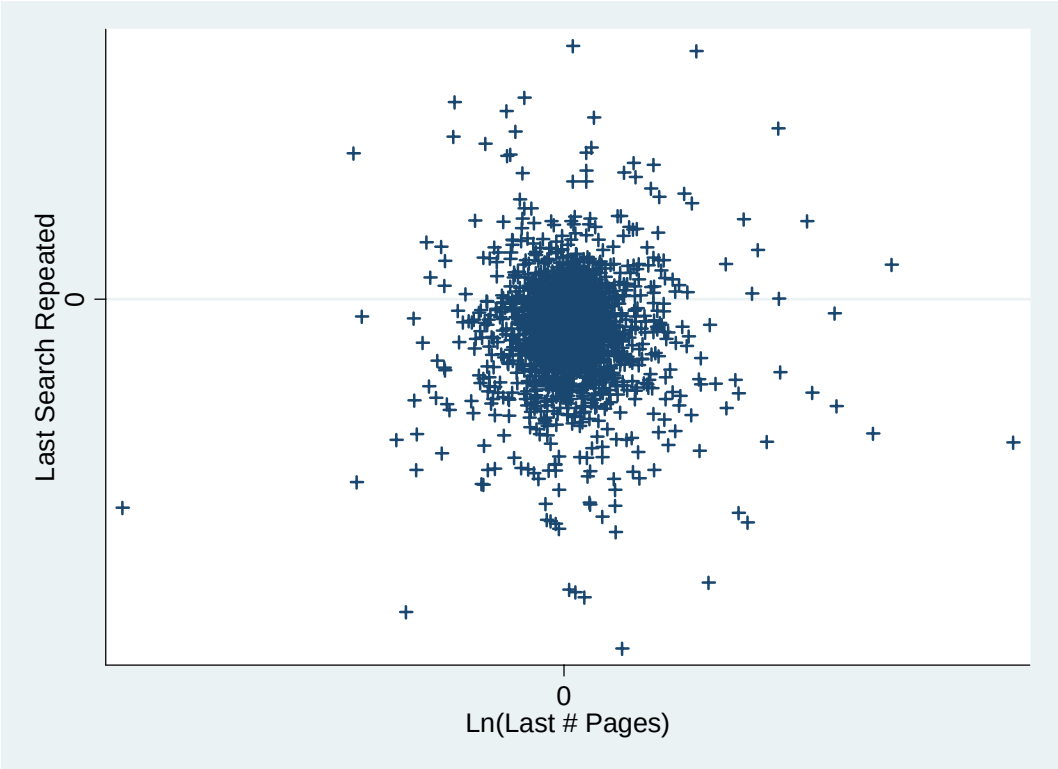

Figure 2: Scatterplot of the Coefficients on Loyalty and Last Search Repeated

Figure 3: Scatterplot of the Coefficients on Loyalty and Ln(Last # Pages)

Figure 4: Scatterplot of the coefficients of loyalty relative to popular portal dummies

Figure 5: Scatterplot of the coefficients of loyalty relative to less popular portal dummies

Figure 6: Scatterplot of the coefficients of the MSN dummy relative to five other portal dummies