Switching Websites After a Denial of Service Attack
Will they ever come back?

Avi Goldfarb[*]
Rotman School of Management
University of Toronto
105 St. George St.
Toronto, ON M5S 3E6
CANADA
(416) 946-8604
agoldfarb@rotman.utoronto.ca

January 30, 2003

Abstract

Dozens of Internet denial of service attacks are reported in the media each year and many more likely go unreported; however, there is little understanding of how these attacks affect customers and consequently little is known about how to mitigate the effects of these attacks.

In this paper, I examine the denial of service attacks of February 2000 to determine whether they had a negative impact on the attacked websites. Using clickstream data on every website visited by 2651 households from Dec. 27, 1999 to Mar. 31, 2000, I show that the attacks had a lasting impact on six of the seven websites attacked (CNN.com, Yahoo, ZDNet, Amazon, Buy.com, and EBay but not E*Trade). This shows that some damage control is necessary after denial of service attack.

I then explore whether this impact is due to developing loyalty to the new website or due to a lower opinion of the attacked website. I separately identify these effects. The identification is possible because the associated loyalty effect will only accrue at the websites that were visited during the attack and not at other websites that compete with the attacked site. Even though I find that free (to the user) websites and online shopping websites were all negatively affected by the attacks, the loyalty effect matters at free websites but not for online shopping websites.

Free websites should therefore focus on short-run promotional campaigns to bring back lost customers. Shopping websites, on the other hand, must emphasize that they have improved security against future attacks and that they are no more vulnerable than any other websites to an attack. In this way, they may be able to convince users that the bad experience at the website was a function of a one-time event that is unlikely to recur.

[*] Thanks to Plurimus Corporation for providing me with the data, to seminar participants at the University of Buffalo and Rutgers University, and to Michael Whinston for an offhanded remark that lead to this paper.

1. Introduction

On February 7, 2000 a teenaged hacker nicknamed ‘mafiaboy’ shut down the Yahoo website for approximately three hours in the first of a wave of ‘denial of service’ (DoS) attacks. Over the next two days, six other major websites would fall victim to mafiaboy’s attacks including Buy.com on the day of its initial public offering, Amazon, CNN.com, EBay, E*Trade, and ZDNet. Since the February 2000 attacks, DoS attacks have shut down dozens of websites including most Microsoft websites in January 2001, Weather.com in May 2001, and Foxnews.com and ESPN.com in June 2002.

In this paper, I show that these attacks had an impact on user behavior. Users that could not access a website due to a denial of service attack became less likely to visit that website in the future. Using a structural economic model, I explore the cause of this impact. Broadly speaking, users may be less likely to return to the website for two reasons. First, the user’s underlying opinion of the website may have changed. Since the user could not access the website, her view on the website’s reliability and overall quality may be affected. On the other hand, the user may have developed loyalty to the website visited instead of the attacked website. Consequently, the user may not switch back to the attacked website once the attack is over. This type of loyalty is often called true state dependence in the econometrics literature and short-run switching costs in the economics literature. It is distinct from the longer-run switching costs faced by customers who want to change banks or move residences.

I can separate out these two effects due to a detailed data set that contains every website visited by 2651 households from December 27, 1999 to March 31, 2000. The DoS attacks therefore occurred roughly in the middle of dataset. The data set allows me to see whether users visited a competing website to the attacked website. For example, the data show whether a user visited MSN.com during the Yahoo attack. This fact allows identification of loyalty to the rival website separately from a change in preferences for Yahoo. If the impact is solely due to changing preferences, then all competing websites should gain a proportionally equal amount. If, on the other hand, there is a loyalty effect then the website that is visited during the attack should gain more than other rival websites because the benefits of loyalty only accrue to that website. I discuss further identification issues in Section 4.

Understanding the impact of DoS attacks can help websites conduct damage control. First, knowing that users who experience the attacks are less likely to return than they otherwise would be shows that these attacks do damage the websites and that at least some damage control is necessary. Second, if the impact is mainly a function of loyalty developed to the new website then websites can perform short-run promotions to bring back old customers. Once the customer has returned, they will remain loyal. However, if the DoS attack changes a user's perception of a website's reliability then the website must focus promotional campaigns on the improved reliability of the website. Furthermore, the website must emphasize that they have improved security against DoS attacks and that they are no more vulnerable than any other websites to an attack. In this way, they may be able to convince users that the bad experience at the website was a function of a one-time event that is unlikely to recur.

I find that DoS attacks do not have the same impact on all websites. While all attacked websites except E*Trade appear to have lost traffic due to the attacks, there is only strong evidence of loyalty effects accruing to competitors of Yahoo. There is weaker evidence that the loyalty effects matter for CNN.com and ZDNet, and while I find little evidence in support of loyalty effects at EBay, I cannot reject this possibility. In general, loyalty effects appear more likely to develop at competitors to free websites (to the user) as a consequence of DoS attacks than to competitors to shopping websites.

There remains considerable disagreement as to whether this type of loyalty or switching costs matter in online business-to-consumer markets. On the one hand, Shapiro & Varian (1998) argue that the competition is just one click away and that consequently switching costs are negligable. Gandal (2001, p. 1105) claims that "there are little (if any) consumer switching costs" at Internet portals.

Alternatively, there is a considerable literature in marketing showing that consumers exhibit loyalty to particular brands, even when switching costs should be zero. In an online context, Johnson, Lohse, & Bellman (2000) label the cost of thinking involved in switching websites "cognitive switching costs."

I find evidence that grouping all online markets together is misleading. Loyalty develops differently for users at free websites than for users at online shopping websites. For shopping websites, Shapiro & Varian's argument appears to hold. When spending money, trust and reliability are more important that the website visited most recently. For free websites, being at the top of a user's mind is sufficient to bring a user to that website. The loyalty effects, however small, matter.

There are a number of different methods to identify loyalty in consumer markets. Typically, loyalty is identified by an individual's propensity to return to a website beyond the average propensity of that individual to visit the website. This method is used by Jones & Landwehr (1988), Keane (1997), Seetharaman, Ainslie, & Chintagunta (1999), Goldfarb (2002) and many others, and relies on assumptions about the influence of individual heterogeneity on serial correlation in the error term to identify the degree of true state dependence. Chen & Hitt (2000) exploit differences in the behavior of old and new customers to see whether the old customers are loyal. They rely on the assumption that all customers have the same overall preferences. I present a new way to identify loyalty that relies on the assumptions about the impact of loyalty of users to one firm on competing firms.

This method could be applied to measuring loyalty when a product is available in a grocery store for an exogenous reason. These 'stockouts' have been explored in other contexts by Jeuland (1979), Farquhar & Pratkanis (1993), Balachander & Farquhar (1994) and others. To my knowledge, they have not been used to measure the importance of loyalty in a product category.

The next section will give a brief description of the data set. Using a difference-in-difference econometric methodology, section 3 shows that the DoS attacks negatively impacted six of the seven attacked websites. Section 4 then shows that loyalty to the website visited during the attack played an important role for the free websites and little role for the shopping websites. Section 5 concludes that short-run promotional campaigns may overcome much of the negative impact of these attacks for free websites, but shopping websites will have to develop more comprehensive strategies aimed at winning back customers' trust.

2. Data

The raw data set, courtesy of Plurimus Corporation, consists of every website visited by 2651 households between December 27, 1999 and March 31, 2000 for a total of 3,228,595 observations. On average, therefore, there are 1217 observations per household. In addition, the data set contains the time of arrival at and departure from a website (to the second), the number of pages viewed at a website, the number of bytes downloaded from the website, and the number of bytes uploaded to the website.

The data set has a number of limitations. First, it is collected at the household level rather than at the individual level. One individual could be online during the DoS attack and never online again during the sample. All other observations could be another individual. If this is the case, being online for the DoS attack will have no effect. Second, I lack at-work data. It is possible that members of the control group attempted to access a website from work during a DoS attack. Both of these limitations, however, will bias the results toward finding no effect for the attacks and I find an effect. Also the data is not geographically representative and it does not include AOL users. These are unlikely to have an impact on the results in this study.

The data are divided into categories. In each category, are all major competitors of the attacked website. For example, competitors to Yahoo include MSN.com, Altavista, Lycos, Google, and dozens of other search engines and portals. The categories definitions were initially set by Plurimus Corporation.

I join this data set with a data set of ‘media mentions’ constructed from the Lexis-Nexis Academic Universe database. If one of the seven companies hit with a DoS attack is mentioned on network television news (ABC, CBS, or NBC) or in the New York Times then the media mentions variable is equal to one for that day. Also if a company is mentioned in the Pittsburgh Post-Gazette, the Tampa Tribune, the Dallas Observer, the Greensboro News & Record, or the Durham Herald-Sun then media mentions is equal to one for local residents on that day.

The other essential piece of data for this study is the identification of the timing of the DoS attacks. I found two sources that listed the exact time of the beginning and end of each attack: CNET online and New York Newsday. Unfortunately, the two sources provided different times. Therefore, I present many results under both definitions of the attack timing. I think that CNET’s definition, however, is more accurate in identifying household that were definitely hit by the DoS attack because no households in my sample access the affected websites during CNET’s definition of the attacks. Some households did access the affected websites under the Newsday definitions. Table 1 shows the exact times of the attacks on Amazon, Buy.com, CNN.com, EBay, E*Trade, Yahoo, and ZDNet according to both sources.

Following the language used in Manski (1995), those households that experienced the DoS attacks are called the 'treatment group' while those that did not are called the 'control group'. The treatment group received the stimulus of the DoS attack. The control group did not.

Since the websites were inaccessible, I cannot determine whether a household tried to access the website under attack and therefore cannot perfectly identify the treatment and control groups. Consequently, I estimate a probability for each household that it is in the treatment group of having experienced the denial of service attack. First, if the household was not online during the attack, then it is assigned zero probability of having experienced the attack. Second, for each household that was online during the attack, I estimate their prior propensity to visit the attacked website. For example, 41% of household 237's website visits prior to the attack are to Yahoo. From this propensity, I estimate the probability that the household visited the attacked website. Household 237 visited two websites during the DoS attack on Yahoo. Therefore the probability that household 237 tried to visit Yahoo during the attack is 0.41+(1-0.41)0.41=0.65. This is the probability that the first visit was to Yahoo plus the probability that the second visit was to Yahoo, given that the first visit was not to Yahoo. More generally:

$$\text{(1)} \qquad \Pr_{ij} = \sum\nolimits_{t=1}^{n_{ij}} (1 - \overline{P}_{ij})^{t-1} \overline{P}_{ij}$$

Where $\Pr_{ij}$ is the probability of household *i* of visiting the attacked website *j* during the attack, $\overline{P}_{ij}$ is the propensity of household *i* to visit attacked website *j* prior to the attack, and n_{ij} is the number of websites household i visits during the attack of website j.

Most of the paper uses this probabilistic definition of the treatment and control groups.[1] Table 2, however, uses a slightly looser definition of the treatment and control groups to show some general trends in the data. In this table, the treatment group is defined by households that visits another website in the same category as the attacked website. The control group consists of all other households. Table 2 shows that for five of the seven attacked websites, treatment group households were more negatively affected by the attacks than control group websites. The econometric analysis conducted in the next section provides evidence that, controlling for household-level differences and using probabilistic treatment groups, all websites except E*Trade were negatively affected by the attacks

3. The Total Effect of the Denial of Service Attacks

3. (a) Model and Identification

I assume that Internet users choose the website that they expect will give them the highest utility. For the shopping websites (Amazon, Buy.com, and EBay), this means that they go to the website that is most likely to have a product they want at an affordable price. For information services (CNN.com, ZDNet), this means users go to the website they expect will provide them with interesting information in an efficient manner. For search engines (Yahoo), users choose the website that will give them a high probability of finding what they seek in a small amount of time. For financial services (E*Trade), users choose the website that will allow them to conduct financial transactions efficiently and securely.

[1] Table A1 in the appendix shows the need for probabilistic modeling of the treatment group. Without these probabilities the signals become noisy and therefore the statistical results are insignificant.

The expected utility from visiting a website is then a function of past experience at the website, website characteristics, and an idiosyncratic error term. Formally, household *i* chooses website *j* on choice occasion *t* when

$$Eu_{ijt}\geq Eu_{ikt} \tag{2}$$

for all $k\neq j$. Where Eu_{ijt} is the expected utility and is defined by

$$Eu_{ijt}=X_{ijt}\beta+\mu_{ij}+\varepsilon_{ijt} \tag{3}$$

Here X_{ijt} are the covariates, β is the associated coefficient vector, μ_t is the household-specific effect, and ε_{ijt} is the idiosyncratic error term. In the estimated regressions, X_{ijt} includes the probability that a household has experienced a DoS attack. It is the coefficient on this variable that is the focus of all presented results.

I estimate both fixed effects and random effects specifications for the model. The fixed effect specification splits the panel into two parts: everything that happened before the DoS attack and everything that happened after. For each household in the data set, I calculate the average propensity to visit an attacked website relative to its competition both before and after the DoS attacks. The model now consists of the following two equations:

$$P_{1i}= X_{i1}\beta+Z_i\gamma+\mu_i+\varepsilon_{i1} \tag{4}$$

$$P_{2i}=Pr_i\alpha+ X_{i2}\beta +Z_i\gamma +\mu_i+\varepsilon_{i2} \tag{5}$$

Where P_{1i} is the probability that household *i* visits the attacked website before it is attacked, and P_{2i} is the probability that household *i* visits the attacked website after it is attacked.[2] Pr_i is the probability of experiencing the DoS attack. X_{it} are time-varying

[2] Any distribution can be assumed on ε_{it} if X_{it} does not vary before the attacks and if it also does not vary after. If this is the case then the probability that $Eu_{ijt}\geq Eu_{ikt}$ will be the same for all *t* before the attacks and again for all *t* after the attacks, and so $\Pr(Eu_{ijt}\geq Eu_{ikt})=\Pr(Eu_{ij\tau}\geq Eu_{ik\tau})$ for all *t* and τ. This implies that $(\Pr(Eu_{ijt}\geq Eu_{ikt})+\Pr(Eu_{ij\tau}\geq Eu_{ik\tau}))/2= \Pr(Eu_{ijt}\geq Eu_{ikt})$ and the above model can be applied. if X_{it} does vary, then if the ε_{it} are distributed uniformly, these equations can be derived from equations (2) and (3) by summing the vector for each variable in the before and after periods and then dividing by the number of observations.

covariates, Z_i are time-invariant covariates, μ_i is the household-specific fixed effect, and ε is the idiosyncratic error term. α, β, and γ are coefficients for Pr_i, X_{it}, and Z_i respectively.

Differencing these two equations gives the model I estimate:

$$P_{2i}-P_{1i} = Pr_i\alpha + (X_{i2}-X_{i1})\beta + \varepsilon_{i2} - \varepsilon_{i1} \tag{6}$$

The fixed effects and time invariant effects cancel out, and we are left with a model that can be consistently estimated by OLS. For most estimates, I assume there are no time-varying covariates. The effect of the denial of service attack is measured by the coefficient on the Pr_i variable, α. This fixed effects model can be viewed as an examination of correlations between the probability of being in the treatment group and the change in visit propensity without forcing a parametric form on any household-level characteristics or on serial correlation in the error terms. For these reasons, I emphasize the fixed effects model over the random effects model.

The random effects model estimates equation (3) assuming μ_i is distributed i.i.d. Normal. Some time-invariant covariates are included. The random effects models include variables for whether the household is in the treatment group, whether the DoS attack occurred before the observation, an interaction between these two, whether the company was mentioned in the media that day, whether website *j* was visited the previous time the household visited the category, how the previous experience at the website went in terms of bytes sent to the website by the user, and how the previous experience at an alternative website went (also in terms of bytes sent to the website by the user. For these last two variables pages viewed, time spent, and bytes sent by the website to the user give almost identical results, but with slightly less explanatory power in terms of log likelihood. Due to the skewness of these variables, the logarithms of pages viewed, time spent, and bytes received and sent are used in the estimation. The dependent variable, y_{it}, is equal to one if household *i* visits the website hit with a DoS attack that is the focus of that regression during visit *t* and zero otherwise. For example, when estimating the impact of a DoS attack on Amazon users, $y_{it}=1$ when a household visits Amazon, and $y_{it}=0$ otherwise.

The interaction between whether the household is in the treatment group and whether the DoS attack occurred before the observation is the key variable of interest and this is therefore the coefficient that is presented in all random effects models in this paper. There may have been systematic differences before and after the time of the attack. These are controlled for by the dummy variable for whether the DoS attack occurred after the observation. There may also be systematic differences between households that were online at the time of the attack and those that were not. These are controlled for by the dummy variable for whether the household is in the treatment group. The coefficient of interest therefore is on the interaction term: the effect of the DoS attack on the treatment group compared to the control group and to how the treatment group behaved before the attack.

In both models, therefore, I look at the impact on users who experienced the attack relative to their previous behavior and to changes in the behavior of others. This methodology can therefore be seen as a difference-in-difference approach. I look at the behavior of the treatment group (the group that experienced the DoS attacks) after the attacks occurred. I compare this behavior with that group's behavior before the attacks—the first difference—and with the other group's behavior after the attacks—the second difference. In this way, the econometric method borrows from Milyo & Waldfogel's (1999) study of the effects of advertising on prices, from Manski's (1995) discussion of identification in econometric models, and from other studies of treatment effects.

3. (b) Results

Table 3 shows the impact of mafiaboy's DoS attacks on each of the seven attacked websites. Unless otherwise specified, all results use CNET's definition of the timing of the attacks. Column 1 reports the fixed effects results of regressing the change in market share on the probability of having experienced the denial of service attack and a constant. Only E*Trade does not have a significantly negative result. Column 2 shows the marginal effects based on a one standard deviation increase in the probability of being in the treatment group based on the Column 1 model. For all six websites that are affected by the attacks, the effect is on the same order of magnitude (0.5% to 3.5%).

Column 3 shows that weighting by total number of households visited yields the same results. Column 4 shows results using the Newsday definition of the timing of the attacks instead of the CNET definition. As expected, the results generally become insignificant. The Newsday definition is much wider than the CNET definition and likely encompasses many households that did not experience the attack. Consequently, this added noise leads to insignificant results. The signs are generally as would be expected, suggesting noise in the treatment group definition.

Column 5 presents a model of logit demand as discussed in Berry (1994). He shows that the log of the market share of one good subtract the log of the market share of an outside good is a linear function of the covariates if the error term is assumed to be extreme value. Here "market share" is the propensity of a given individual to visit a website before and after the attack. This model is generally appealing as it is grounded in economic theory. It does not, however, work well when market shares are zero as is the case with the data used here. For the log of zero, I substitute negative one trillion. Consequently, these results have little meaning here and are only included to show robustness to alternative specifications.

Column 6 adds regressors to the base model. As stated above, this will only have economic meaning if we assume that errors are uniformly distributed or if we interpret these as unchanging average characteristics of the website for that person. Neither of these are appealing assumptions; however, I present this regression to show that adding information about individual interactions with websites does not change any of the qualitative results and the point estimates of the coefficients change little.

The random effects model of columns 7 and 8 shows two surprises. The DoS attacks seem to have had a positive effect on Amazon and ZDNet under this specification. All other results are negative and significant including E*Trade. I believe the positive correlation is a function of imperfect specification of this model in terms of household-level characteristics or serial correlation in the error term. However as a consequence of these random effects results, later in the paper I will only assert that there is weak evidence for loyalty effects accruing to ZDNet competitors as a consequence of the DoS attacks.

Since this is a probit model, the coefficients do not have economic meaning beyond statistical significance. Column 8 presents the impact on the probability of visiting the website if the probability of experiencing the attack goes from zero to one for the average household. For example, the average household will find its probability of visiting Yahoo fall roughly 2.5% from 33.1% to 30.6%.

In terms of significance, the results are strongest for Buy.com and Yahoo. Users of both websites clearly faced a negative fallout after mid-February 2000. It may seem like the Buy.com result is compromised by the fact that the Buy.com IPO was on the day of the attack. While this is likely to lead to behavioral differences before and after the attack, it seems unlikely that the effect on those in the treatment group will be different from those in the control group. CNN and EBay users also display negative effects across all models. Not surprisingly, significance is least common in the models with fewer observations. Nevertheless, the evidence strongly suggests that these DoS attacks negatively affected Buy.com, CNN, EBay, and Yahoo. There is also weaker evidence of an effect on Amazon and ZDNet. E*Trade does not appear to have been affected. Since E*Trade users tend to have accounts and a strong relationship with the company, it is not surprising that a short DoS attack was insufficient to cause users to abandon their discount broker. Similarly, the relatively large impact on CNN and Yahoo is not surprising as there are few long-term switching barriers associated with using these free websites.

Table 4 shows the impact of the DoS attacks on the websites that were visited during the attack instead of the attacked website. In particular, the regression looks at all households that visited a rival to the attacked website during the attack. A positive coefficient shows that of these households, those that were more likely to experience the attack were more likely to visit the rival. Sample sizes are much smaller here as only those users who visited a rival to the attacked website during the attack can be included. Even with small sample sizes, the results show CNN.com and especially Yahoo rivals who were visited during the attacks clearly benefited. The next section addresses whether these rivals benefited due to loyalty resulting from the visit during the attack, or whether they benefited due to a decreased preference for the attacked website as a consequence of a poor experience at that website.

4. Loyalty or Preference Changes?

4. (a) Model and Identification

In order to identify whether the competing websites visited during the attack benefited from loyalty, I use the fact that loyalty resulting from the attacks will only accrue to websites that users visit during the attack. All other websites competing with the attacked website can only benefit from a change in preferences. Formally, recall equation (3),

$$Eu_{ijt}=X_{ijt}\beta+\mu_i+\varepsilon_{ijt}$$

$X_{ijt}\beta$ can be decomposed into the effect of the denial of service attacks component, $Pr_i\alpha$, and all other components, $\overline{X}_{ijt}\overline{\beta}$. The key to the identification is that $Pr_i\alpha$ will have a different meaning for competing websites that were visited during the attack and those that were not. The utility of returning to a website that was visited during the attack will have a loyalty component. Other competing websites to the attacked website will not benefit from loyalty. They will only benefit from the reduced preference for the attacked website. Therefore the utility from visiting the attacked website for a household that experienced the attack is

$$Eu_{iat}= Pr_i\ \pi+\overline{X}_{iat}\overline{\beta}+\mu_{ij}+\varepsilon_{iat} \tag{7}$$

Where π is the preference change as a consequence of the attack. The utility from visiting a competing website that was visited during the attack is

$$Eu_{ict}= Pr_i\ \lambda+\overline{X}_{ict}\overline{\beta}+\mu_{ij}+\varepsilon_{ict} \tag{8}$$

Where λ is the added loyalty associated with having visited the website an extra time in the past due to the DoS attack. Finally, the utility from visiting another competing website that was not visited during the attack is

$$Eu_{iot}= \overline{X}_{iot}\overline{\beta}+\mu_{ij}+\varepsilon_{iot} \tag{9}$$

The DoS attack will not directly enter the utility function for a website that was neither attacked nor visited during the attack. The attack will only affect the probability of visiting these other websites through the impact on the attacked websites and the websites that were visited during the attack. Consequently, controlling for user behavior before the attacks, by exploring whether users are more likely to visit websites that they visited during the attack than other competing websites, λ, the coefficient on loyalty, is identified.

In particular, a household visits the website that was visited during the attack instead of another competing website if $Eu_{ict} \geq Eu_{iot}$. Rearranging terms, this means that the website visited during the attack is visited again if

$$Pr_i\ \lambda + (\overline{X}_{ict} - \overline{X}_{iot})\overline{\beta} + \mu_{ic} - \mu_{io} + \varepsilon_{ict} - \varepsilon_{iot} \geq 0 \tag{10}$$

Therefore, estimating a choice model to see whether competing firms that were visited during the attacks gained more than other competing firms will identify the effect of loyalty, λ.

This identification, however, is imperfect. Another possibility for the competing firm that was visited during the attack gaining more than other competing firms may be that the competing firm is a closer substitute for the attacked firm. Consequently, when households leave the attacked firm for preference reasons they will go to the firm visited during the attack. The fixed effects model largely overcomes this potential issue by differencing out all household-level preferences that are time-invariant. This is a further advantage of the fixed effects model over the random effects model. I also show that measures of substitutability based on user behavior at the websites do not change the results. Furthermore, there is little reason to believe this would only be relevant at free websites and not at shopping websites.

It is important to remember that this method identifies a particular kind of loyalty: the impact of a one-time switch. All that is required for the loyalty effect to exist is that the act of visiting a website once will increase the probability of visiting that website in the future, all else being equal including preferences for that website not based on this loyalty effect. Having visited a website at some point in the past must therefore have a lasting impact on the utility from visiting that website in the future.[3]

The results are presented in a similar format to those in section 3. I present results for both fixed and random effects models and the assumptions required are the same as those discussed above.

4. (b) Results

Table 5 shows whether loyalty (to the website visited instead of the attacked website) mattered in the DoS attacks for each of the seven firms attacked. Since only households that visited a competing website during the attack can be included, sample sizes are less than fifty for all fixed effects regressions using the CNET definition of loyalty except those involving Buy.com, CNN.com and Yahoo. Results for Amazon, EBay, E*Trade, and ZDNet should therefore be interpreted with caution.[4] Furthermore, Table 3 showed that E*Trade was not affected by the attacks and consequently, it is meaningless to determine how important loyalty effects were for E*Trade.

As in Table 4, a positive coefficient shows that of these households, those that were more likely to experience the attack were more likely to visit the website visited during that attack. Therefore this controls for any household-level preferences and even for any (however unlikely) systematic differences between households that were online at the time of the attacks and those that were not.

Column 1 presents results of the general fixed effects model with CNET's definition of the timing of the attacks. Column 2 presents the marginal effects of a one

[3] An example of this framework is Guadagni & Little's (1983) loyalty measure.

[4] Random effects results should also be interpreted with caution as they are derived from the same number of households as the fixed effects.

standard deviation increase in the probability of being in the treatment group, based on the model in Column 1. Column 3 presents this model weighted by the total number of visits. Column 4 adds a control for the closeness of the competing website visited during the attack to the attacked website relative to the closeness of other competing websites to the attacked website.[5] Technically, this control can only be added if it is viewed as only being relevant after the attacks occur. Column 5 presents the results using the noisier Newsday definition of the timing of the attacks. Column 6 presents a logit model of demand as described in Berry (1994). As in section 3, while usually conceptually appealing this model does not work well when market shares are often zero. Columns 7 and 8 present the base random effects model described in section 3 (without media mentions as these were difficult to obtain for the dozens of relevant websites) and the corresponding elasticities. Column 9 presents a random effects model that does not allow for systematic differences before the attacks between people who end up experiencing the attack and those that do not.

The most striking result of the table is that visitors who were likely to have wanted to visit Yahoo, but went elsewhere because of the DoS attack were much more likely to return to the website they visited than were other users who visited a website that competed with Yahoo at the time. The table therefore presents strong evidence, under many alternative specifications, that loyalty effects are an important factor in explaining the negative impact of the DoS attacks on Yahoo.

There is weaker evidence that competitors to CNN.com and ZDNet benefited from loyalty effects as a consequence of the attacks. Both of these websites have some significantly positive and no significantly negative results under different fixed effects

[5] This variable was defined as the $(PG_c-PG_a)^2-(PG_o-PG_a)^2$ where PG_j is average number of pages viewed at website j, c is the competing website visited during the attack, a is the attacked website, and o is other competing websites that were not visited during the attack.

models. The results for EBay are mixed. While the marginal effect seems large in the fixed effects case, it is not significant. Furthermore, the result is negative in the random effects case. For Amazon and Buy.com, there are cases where the coefficient is significantly negative. Since these models the reject the existence of loyalty effects, I conclude that the evidence suggests that competitors of Amazon and Buy.com did not benefit from loyalty effects as a consequence of the attacks. These statistical significance results suggest that loyalty mattered to the free websites, but not to the shopping websites. The magnitude of the marginal effects in Column 2 tells a similar story.

5. Conclusions

The denial of service attacks of February 2000 had an impact on user behavior. The attacks decreased the probabilities that users would visit Amazon, Buy.com, CNN.com, EBay, Yahoo, and ZDNet over the following seven weeks. E*Trade users, who face large costs of changing their accounts, did not change their behavior as a result of the attacks.

For Yahoo, this impact is a result of users developing loyalty to its competitors as a consequence of visiting a rival website during the attack. For CNN.com and ZDNet, it is also likely a result of loyalty effects accruing to their competitors. It is not only a function of a change in the user's underlying opinion the attacked website. Loyalty effects, however, do not seem to have played a major role in the impact of the DoS attacks on shopping websites.

Even though the competition to the free websites is just a click away, the potential benefits to switching seem to be dominated the loyalty effect generated by one forced visit to a free website. Despite no obvious impediments to switching websites and controlling for overall preferences, one visit generates a lasting effect.

For shopping websites, this is not the case. Perhaps the potentially larger perceived benefits surrounding the pricing and quality of a user's favorite shopping website mean that the loyalty benefits associated with a one-time switch are not large enough to be relevant to online shoppers.

Each of these DoS attacks lasted less than 4 hours. Therefore, I can only identify the impact of a brief one-time switch. It is likely that the results would be different if the attacks lasted for days. In this case, it is likely that some E*Trade customers may have switched and identifying a different types of loyalty effects would be possible. The purpose of this paper is to better understand the impact of DoS attacks on the websites that are attacked. I do not presume to understand the impact of a long-term shutdown of a website.

The identification method used here to identify loyalty effects can easily be applied to stockouts in grocery stores. Comparing the impact of the stockout on the brand that was bought instead with other brands that were not bought will allow identification of loyalty effects in these markets.

The reasons behind the impact of a DoS attack on a website have important strategic implications for websites. The above analysis suggests that free websites will benefit from a short promotional campaign aimed at those users who tried to access the website during the attack. The loyalty effects will then accrue to the promoting website again and the impact of the DoS attack will be minimized.

Shopping websites, on the other hand, are less likely to find this a useful strategy. They should focus on showing customers that their concerns that arose from the DoS attacks are no longer valid. Shopping websites that are victims of DoS attacks should emphasize to the customers that tried to access the website during that attack that they have improved security and are no more vulnerable than any other websites to an attack.

References

Balachander, Subramanian & Peter H. Farquhar. 1994. "Gaining More by Stocking Less: A Competitive Analysis of Product Availability." *Marketing Science* 13(1), 3-22

Berry, Steven T. 1994. "Estimating discrete-choice models of product differentiation." *RAND Journal of Economics* 25: 242-262.

Chen, Pei-Yu and Lorin M. Hitt. 2000. "Switching Cost and Brand Loyalty in Electronic Markets: Evidence from On-Line Retail Brokers." Wharton School of Business. Mimeographed.

Farquhar, Peter H., & Anthony R. Pratkanis. 1993. "Decision Structuring with Phantom Alternatives." *Management Science* 39(10), 1214-1226.

Heckman, James J. 1981. "Statistical Models for Discrete Panel Data." In *Structural Analysis of Discrete Data with Econometric Applications,* ed. Charles F. Manski and Daniel McFadden, 179-195. Cambridge, MA: The MIT Press.

Goldfarb, Avi. 2002. "State Dependence at Internet Portals." Mimeographed. University of Toronto.

Jeuland, Abel P. 1979. "The Interaction Effect of Preference and Availability on Brand Switching and Market Share." *Management Science* 25(10), 953-965.

Johnson, Eric J., Jerry Lohse, and Steve Bellman. 2000. "Cognitive Lock-In." Columbia University Business School. Mimeographed.

Jones, J. Morgan, and Jane T. Landwehr. 1988. "Removing Heterogeneity Bias from Logit Models Estimation." *Marketing Science* 7(1), 41-59.

Keane, Michael P. 1997. "Modeling Heterogeneity and State Dependence in Consumer Choice Behavior." *Journal of Business & Economic Statistics* 15(3), 310-327.

Manski, Charles. 1995. *Identification Problems in the Social Sciences.* Cambridge, MA: Harvard University Press.

Milyo, Jeffrey & Joel Waldfogel. 1999. "The Effect of Advertising on Prices: Evidence in the Wake of 44 Liquormart." *American Economic Review* 89(5), 1081-1096.

Seetharaman, P. B., Andrew Ainslie, & Pradeep Chintagunta. 1999. "Investigating Household State Dependence Effects Across Categories." *Journal of Marketing Research* 36(4), 488-500.

TABLE 1
Timing of the attacks

	Time of Attack (CNET)*	# users in category at time	# users online at time		Time of Attack (NY Newsday)*	# users in category at time	# users online at time
FREE							
CNN.com	Tues. Feb. 8: 7:00 PM–8:50 PM	56	587		Tues. Feb. 8: 7:15 PM–10:45 PM	99	833
Yahoo	Mon. Feb. 7: 1:20 PM–4:20 PM	401	650		Mon. Feb. 7: 1:15 PM–4:15 PM	400	656
ZDNet	Wed. Feb. 9: 6:45 AM–9:45 AM	16	397		Wed. Feb. 9: 7:15 AM–10:30 PM	19	448
SHOPPING							
Amazon	Tues. Feb. 8: 8:00 PM–9:00 PM	38	423		Tues. Feb. 8: 8:00 PM–11:45 PM	104	847
Buy.com	Tues. Feb. 8: 1:50 PM–4:50 PM	88	717		Tues. Feb. 8: 2:00 PM–6:00 PM	112	877
EBay	Tues. Feb. 8: 6:20 PM–7:50 PM	10	375		Tues. Feb. 8: 5:30 PM–10:30 PM	27	702
OTHER							
E*Trade	Wed. Feb. 9: 8:00 AM–9:30 AM	37	168		Wed. Feb. 9: 8:15 AM–11:00 AM	86	257

*All times EST

TABLE 2
Probability of visiting website before and after the attack for treatment and control groups (under CNET’s timing)

	Treatment Group				Control Group				
	Probability visit website before the attack if visit a similar website during the attack	Probability visit website after the attack if visit a similar website during the attack	% Change		Probability visit website before the attack if DO NOT visit a similar website during the attack	Probability visit website after the attack if DO NOT visit a similar website during the attack	% Change		% Treatment Change *subtract* % Control Change
FREE									
CNN.com	0.0311	0.0293	-5.8%		0.0955	0.0978	2.4%		-8.2%
Yahoo	0.286	0.281	-1.7%		0.361	0.358	-0.8%		-0.9%
ZDNet	0.0564	0.0228	-59.6%		0.293	0.320	9.2%		-68.8%
SHOPPING									
Amazon	0.0490	0.0245	-50.0%		0.168	0.136	-19.0%		-31.0%
Buy.com	0.00556	0.00833	49.8%		0.0151	0.0133	-11.9%		62.7%
EBay	0.537	0.677	26.1%		0.679	0.668	-1.6%		27.7%
OTHER									
E*Trade	0.00827	0.00874	5.7%		0.0254	0.0305	-20.1%		-14.4%

TABLE 3
Impact of Denial of Service Attacks on Attacked Websites[6] (standard errors in parentheses)

	Fixed Effects						Random Effects	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	General Model	Marginal Effects of 1 Std. Dev. Change	Weighted by total number of visits by household	Newsday reported time of attack	Logit Model of Demand (Berry 1994)	Additional regressors[7]	General Model	General Model Elasticity
FREE								
CNN.com	-0.135^ (0.101)	-0.00679^	-0.231*** (0.0804)	-0.0452 (0.0730)	3.08e+09 (2.64e+09)	-0.139^ (0.0913)	-0.132^ (0.102)	-0.0105^
Yahoo	-0.0632*** (0.018)	-0.0125***	-0.0723*** (0.00973)	-0.00744 (0.0182)	-6.08e+08^ (4.74e+08)	-0.0618*** (0.0182)	-0.0961*** (0.0159)	-0.0259***
ZDNet	-0.500^ (0.353)	-0.0177^	-0.377* (0.208)	0.0841 (0.216)	6.55e+09 (5.88e+09)	-0.395 (0.326)	0.314^ (0.242)	0.101^
SHOPPING								
Amazon	-0.456** (0.232)	-0.0130**	-0.646*** (0.216)	-0.102 (0.127)	-5.12e+09 (5.03e+09)	-0.455** (0.214)	0.654*** (0.222)	0.123***
Buy.com	-1.36*** (.116)	-0.0141***	-1.46*** (0.0885)	-0.539*** (0.0953)	-1.87e+10*** (6.10e+09)	-1.38*** (0.113)	-2.92*** (0.848)	-0.0126***
EBay	-0.231* (0.120)	-0.0322*	-0.173* (0.102)	-0.0663 (0.0790)	-6.44e+09*** (1.74e+09)	-0.194** (0.0984)	-0.106** (0.0470)	-0.0381**
OTHER								
E*Trade	-0.0291 (0.275)	-0.00113	-0.155 (0.124)	0.168 (0.155)	1.41e+07 (5.31e+09)	0.00678 (0.250)	-0.0421^ (0.0259)	-0.00145^

***significant with 99% confidence in a two-tailed test
**significant with 95% confidence in a two-tailed test
*significant with 90% confidence in a two-tailed test
^significant with 90% confidence in a one-tailed test

[6] Unless otherwise specified, numbers presented are coefficients, the treatment group is defined by the CNET definition and the variables included are the same as those in the general model.
CNET Fixed effects number of observations: Amazon=1932, Buy=1990, CNN=1544, EBay=572, E*Trade=432, Yahoo=2479 and ZDNet=944.
Random effects number of observations: Amazon=71,788, Buy=86,058, CNN=108,312, EBay=48,348, E*Trade=156,477, Yahoo=857,993 and ZDNet=39,242.

[7] Regressors include difference in average media mentions, difference in average bytes downloaded from the attacked website, and difference in average bytes downloaded from the attacked website's competitors. Bytes uploaded, pages viewed, and time spent give the same significance results, with less explanatory power.

TABLE 4

Impact of Denial of Service Attacks on Rival Websites that were Visited During the Attacks[8]

	Fixed Effects					Random Effects	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	General Model	Marginal Effects of 1 Std. Dev. Change	Weighted by total number of visits	With controls for substitutability	Newsday Definition	Coefficient	Elasticity
FREE							
CNN.com	0.704^ (0.433)	0.0826^	0.199** (0.0821)	0.702^ (0.437)	0.366^ (0.277)	-0.602* (0.313)	-0.143*
Yahoo	0.168** (0.0463)	0.0563**	0.122*** (0.0381)	0.168*** (0.0464)	0.159*** (0.0467)	0.00298 (0.0220)	0.00108
ZDNet	0.383 (1.52)	0.0245	0.124 (.567)	0.161 (1.56)	0.691 (0.547)	-0.230 (1.10)	-0.0589
SHOPPING							
Amazon	1.63 (1.91)	0.0486	-3.98** (1.74)	1.93 (1.96)	-0.0157 (0.343)	4.08** (1.81)	0.886**
Buy.com	1.13 (2.41)	0.0163	3.62*** (1.34)	1.13 (2.41)	1.35 (1.43)	-14.45*** (1.85)	-1.12***
EBay	-0.350 (0.357)	-0.125	0.00946 (0.160)	-0.540 (0.414)	0.0133 (0.195)	0.539** (0.225)	0.0834**
OTHER							
E*Trade	-0.539 (0.724)	-0.0344	-0.471 (0.406)	-0.551 (0.735)	-0.326 (0.364)	-2.19*** (0.330)	-0.870***

***significant with 99% confidence in a two-tailed test

**significant with 95% confidence in a two-tailed test

*significant with 90% confidence in a two-tailed test

^significant with 90% confidence in a one-tailed test

[8] *CNET* Fixed effects number of observations: Amazon=38, Buy=88, CNN=56, EBay=10, E*Trade=37, Yahoo=401 and ZDNet=16.
Random effects number of observations: Amazon=5,813, Buy=12,630, CNN=22,709, EBay=2,663, E*Trade=36,725, Yahoo=311,303 and ZDNet=2,549.
Newsday Fixed effects number of observations: Amazon=104, Buy=112, CNN=99, EBay=27, E*Trade=86, Yahoo=400 and ZDNet=19.

TABLE 5
Loyalty Effects: Impact of Denial of Service Attacks on Rival Websites that were Visited During the Attacks relative to Rival Websites that were not Visited During the Attacks[9]

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	General Model	Marginal Effects of 1 Std. Dev. Change	Weighted by total number of visits	With controls for substitutability	Newsday Definition	Logit Model of Demand (Berry 1994)	General Model	Elasticity	No controls for probability in treatment group
FREE									
CNN.com	0.487 (0.419)	0.0572	-0.0163 (0.0893)	0.484 (0.423)	0.400^ (0.279)	-2.09e+10 (4.33e+10)	0.478*** (0.0394)	0.190***	0.0611* (0.0360)
Yahoo	0.100** (0.0460)	0.0335**	0.105*** (0.0384)	0.0995** (0.0460)	0.0902* (0.0462)	-1.02e+10 (3.51e+10)	0.172*** (0.00695)	0.0683***	0.0433*** (0.00770)
ZDNet	2.35^ (1.41)	0.151^	2.91*** (0.481)	2.18^ (1.46)	0.193 (0.539)	-1.76e+11 (2.10e+11)	0.381 (1.13)	0.0641	0.170** (0.0861)
SHOPPING									
Amazon	-1.07 (1.94)	-0.0319	-5.14*** (1.59)	-0.819 (1.99)	-0.0351 (0.342)	9.38e+10 (3.01e+11)	4.53*** (1.70)	0.915***	0.201*** (0.0600)
Buy.com	1.43 (2.21)	0.0206	3.08** (1.29)	1.43 (2.22)	1.72 (1.43)	4.95e+11 (4.16e+11)	-13.74*** (1.81)	-1.88***	0.105** (0.0410)
EBay	0.509 (0.386)	0.181	1.11*** (0.167)	0.547 (0.473)	0.404* (0.218)	-1.01e+11 (3.66e+11)	-0.0836 (0.325)	-0.0300	0.0367 (0.117)
OTHER									
E*Trade	-0.638 (0.736)	-0.0407	-0.753* (0.417)	-0.653 (0.747)	-0.350 (0.364)	-2.17e+10 (6.07e+10)	-2.38*** (0.341)	-0.948***	0.219*** (0.0251)

***significant with 99% confidence in a two-tailed test
**significant with 95% confidence in a two-tailed test
*significant with 90% confidence in a two-tailed test
^significant with 90% confidence in a one-tailed test

[9] *CNET* Fixed effects number of observations: Amazon=38, Buy=88, CNN=56, EBay=10, E*Trade=37, Yahoo=401 and ZDNet=16.
Random effects number of observations: Amazon=5,813, Buy=12,630, CNN=22,709, EBay=2,663, E*Trade=36,725, Yahoo=311,303 and ZDNet=2,549.
Newsday Fixed effects number of observations: Amazon=104, Buy=112, CNN=99, EBay=27, E*Trade=86, Yahoo=400 and ZDNet=19.

Appendix

Table A1[10]

Impact of Denial of Service attacks: non-probabilistic treatment group definitions

	Treatment group= online at time of attack		Treatment group= visit category at time of attack	
	CNET reported time of attack	Newsday reported time of attack	CNET reported time of attack	Newsday reported time of attack
FREE				
CNN.com	-0.00248 (0.0111)	-0.00961 (0.0103)	0.00544 (0.0262)	-0.0141 (0.0205)
Yahoo	-0.00444 (0.00828)	-0.00424 (0.00826)	-0.00840 (0.00970)	0.00104 (0.00965)
ZDNet	0.0326 (0.0274)	0.00758 (0.0268)	-0.0538 (0.0792)	-0.0805 (0.0722)
SHOPPING				
Amazon	-0.000709 (0.0166)	-0.0185 (0.0137)	0.0280 (0.0478)	-0.0734** (0.0297)
Buy.com	-0.00157 (0.00264)	0.00245 (0.00253)	0.00121 (0.00615)	0.000610 (0.00551)
EBay	-0.00799 (0.0356)	0.0431 (0.0344)	-0.0522 (0.105)	0.0222 (0.06131)
OTHER				
E*Trade	0.0153 (0.0272)	0.0297 (0.0233)	0.00510 (0.0747)	0.0578 (0.0410)

***significant with 99% confidence in a two-tailed test
**significant with 95% confidence in a two-tailed test
*significant with 90% confidence in a two-tailed test
^significant with 90% confidence in a one-tailed test

[10] Unless otherwise specified, numbers presented are coefficients, the treatment group is defined by the CNET definition and the variables included are the same as those in the general model.
CNET Fixed effects number of observations: Amazon=1932, Buy=1990, CNN=1544, EBay=572, E*Trade=432, Yahoo=2479 and ZDNet=944.
Newsday Fixed effects number of observations: Amazon=1932, Buy=1990, CNN=1544, EBay=572, E*Trade=432, Yahoo=2479 and ZDNet=944.