

NBER WORKING PAPER SERIES

DOES THE INTERNET INCREASE TRADING?
EVIDENCE FROM INVESTOR BEHAVIOR IN 401(k) PLANS

James J. Choi
David Laibson
Andrew Metrick

Working Paper 7878
<http://www.nber.org/papers/w7878>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2000

We thank Lori Lucas, Jim McGhee, and Scott Peterson of Hewitt Associates, LLC for generously providing the data and extensive resources that made this project possible. We also acknowledge the technical help of Hewitt employees Lonnie Lee Buresh and Prema Palicharla. Outside of Hewitt, we benefited from the suggestions of Michael Kremer, Andrei Shleifer, Nick Souleles, Richard Zeckhauser, and seminar participants at Harvard and Wharton. Kunal Merchant, Parag Pathak, Kristy Piccinini, and Stephen Weinberg provided outstanding research assistance. Laibson acknowledges financial support from the Olin Foundation. Metrick acknowledges financial support from the Rodney White Center at the Wharton School of the University of Pennsylvania. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by James J. Choi, David Laibson, and Andrew Metrick. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Does the Internet Increase Trading? Evidence from Investor Behavior in 401(k) Plans

James J. Choi, David Laibson, and Andrew Metrick

NBER Working Paper No. 7878

September 2000

JEL No. D0, G0, Lo, O0

ABSTRACT

We analyze the impact of a Web-based trading channel on the trading activity in two corporate 401(k) plans. Using detailed data on about 100,000 participants, we compare trading growth in these firms to growth for a sample of firms without a Web channel. After 18 months of access, the inferred Web effect is very large: trading frequency doubles, and portfolio turnover rises by over 50 percent. We also document several patterns of Web-trading behavior. Young, male, and wealthy participants are more likely to try the Web channel. Frequent traders (before Web introduction) are less likely to try the Web. Participants who try the Web tend to stick with it. Web trades tend to be smaller than phone trades both in dollars and as a fraction of portfolio. “Short-term” trades make up a higher proportion of phone trades than of Web trades.

James J. Choi
Department of Economics
Littauer M-14
Harvard University
Cambridge, MA 02138

David Laibson
Department of Economics
Littauer M-14
Harvard University
Cambridge, MA 02138
and NBER
dlaibson@harvard.edu

Andrew Metrick
The Wharton School
University of Pennsylvania
3451 Walnut Street
Philadelphia, PA 19104
and NBER

There has been much speculation that the development of the Internet and the World Wide Web has had a significant impact on financial market behavior. Many commentators believe that the spread of online brokerage accounts and the rise of the day-trader have increased stock market volume and volatility.¹ Moreover, Web trading of financial instruments is only the tip of the iceberg. Thousands of new businesses have been founded on the belief that the Web will transform many types of economic activity.

The impact of the Web, however, is by no means clear. Even if the Web's emergence has coincided with a rise in trading activity, it is difficult to determine whether the Web's role is causal. It may simply be that online traders would have executed their trades on other channels. Similarly, one might wonder whether the existence of Amazon.com has raised aggregate book buying, or whether it has just appropriated the customers of other retailers.

To our knowledge, the only careful analysis of the behavior of online investors is Barber and Odean (1999), who focus on the trading behavior and investment performance of investors who switch from the phone to an online channel. While their data are useful for many purposes, the self-selected nature of discount-brokerage customers who choose to trade online makes it difficult to draw inferences about the impact of new trading technologies on the typical investor.

To evaluate the causal role of the Web, one needs a source of exogenous variation in access to Web-based transactions. We exploit such variation in this paper to examine the impact of the Web on the trading decisions of about 100,000 participants

¹For examples of press stories on these topics, see Dugan (1999) and Kunath (1999).

in two corporate 401(k) plans. Both of these plans opened a Web channel in August 1998, and we have about three years of detailed trading data for each plan. As a comparison, we also have a measure of trading activity for a set of large 401(k) plans that do not have a Web channel.²

As an illustration of our main results, Figure 1 plots a twenty-trading-day moving average of the daily trading frequency for one of our companies, code-named Alpha. At first glance, the Web effect appears dramatic. Within 18 months after the Web channel was opened, Web transactions represent approximately 60 percent of all transactions and the trading rate has quadrupled from its pre-Web level. But, as we emphasized above, all Web trading is not necessarily “new” trading. Participant trading is driven by many factors that have been trending up over our sample period. For example, stock price volatility has risen recently, and trading volume might be expected to rise as a result. When we control for such changes — including use of a trading index for a set of firms that do not have a Web channel — we continue to find a huge Web effect. After 18 months, the Web channel nearly doubles the daily trading frequency. Over the same period, daily turnover — the fraction of balances traded — increases by more than 50 percent.

In our analysis of Web trading we also document several other patterns. Young, male, and wealthy participants are more likely to try the Web for trading. Frequent traders (before Web introduction) are less likely to try the Web. Participants who try the Web tend to stick with it. Web trades tend to be smaller than phone trades both in dollars and as a fraction of the portfolio being traded. Lastly, “short-term”

²There is a substantial literature on 401(k) savings and average asset allocation choices, but only a few papers address trading behavior (Ameriks and Zeldes (2000) and Agnew, Balduzzi, and Sunden (2000)), and none focus on the determinants of trading frequency or the impact of trading technologies.

trades make up a higher proportion of phone trades than of Web trades.

The rest of the paper formally demonstrates these results. In Section I, we describe the dynamic economic forces that drive Web trading and sketch a simple model to help organize our empirical analysis. In Section II, we describe our dataset. In Section III, we present the main empirical results and show that the Web has increased trading. Section IV documents several specific features of Web trading. In Section V we conclude.

I. A framework for analyzing Web-based trading

In this section we describe the conceptual framework that we use to organize our analysis. The framework is based on two economic factors that are central to any analysis of the spreading impact of the Web: costs tend to fall over time and the evolution of costs contains some stochastic elements.

It is obvious that Web costs have *tended* to fall over the past several years. Numerous mechanisms drive this effect. For example, computer ownership has risen; Web access times have shortened; Web sites have become more user friendly; programming/interface bugs have been removed; and computer literacy has risen as more people use computers for work and leisure.

However, Web transaction costs have not fallen uniformly at all points in time. Indeed, the evolution of transactions costs is at least somewhat unpredictable. Unexpected events occasionally lead people to revise *upward* their estimates of the cost of Web transactions. For example, recent news about the lack of Web privacy may lead some users to reject the Web for more secure transaction technologies.³ Simi-

³For example, see Stillwell (2000).

larly, increasingly congested networks lead some users' access times to rise, effectively raising the cost of Web use.

We formalize these effects with a simple model which we sketch below. Before turning to the model we anticipate the predictions that the model makes. We present these predictions first, since they are not specific to our particular formalization.

First, the introduction of the Web channel will raise transaction frequency. This effect and other trading effects will tend to grow over time because Web usage costs fall over time. The Web channel will also raise turnover, measured as the traded fraction of portfolio value. This increase will be proportionally smaller than the increase in transaction frequency since Web access should lower average transaction size; Web traders will be willing to execute smaller (less valued) transactions, since Web transactions are less costly than phone transactions. Finally, Web traders will sometimes get discouraging signals about the Web and switch back to non-Web trading technologies. However, if a trader has repeatedly used the Web, such switching is relatively unlikely to occur in the near future.

We now sketch a formal model of Web use. We use this model to organize our analysis. Other economic models based on these principles would make similar predictions. Hence, readers who are not interested in the technical details of the model may wish to proceed immediately to the next section where we begin a discussion of the data.

In our formal model, trading opportunities have stochastic arrival times generated by a Poisson process with mean arrival rate λ . For simplicity, we assume that an investor can only exploit a trading opportunity at the moment that the opportunity arises. Each trading opportunity has an independent stochastic value η with density

function $f(\eta)$. We assume that the value of a trade is a proxy for the size of the trade.

Trades are executed either on the phone or on the Web. For phone trades, we subscript the associated variables with a p . If an individual is a phone trader, the individual trades if the value of the trading opportunity, η , is greater than the trading cost for phone trades, c_p . So the expected instantaneous flow payoff for a phone trader is given by

$$\pi_p = \lambda E \max \{0, \eta - c_p\}. \quad (1)$$

If an individual is a Web trader, the individual trades if the value of the trading opportunity is greater than the cost for Web trades, c_w . We assume $c_w < c_p$. The expected instantaneous flow payoff for a Web trader is given by

$$\pi_w = \lambda E \max \{0, \eta - c_w\}. \quad (2)$$

To be a Web trader, the individual must pay a flow cost of x , which represents the ongoing cost of using the Web.⁴ This flow cost will trend stochastically downward, converging to zero. It may be thought of as the cost of staying connected to and informed about the Web (e.g., access costs or ongoing human capital investment in Web-based skills). If an individual is paying x , then the expected instantaneous net flow payoff is

$$\pi_w - x = \lambda E \max \{0, \eta - c_w\} - x. \quad (3)$$

Over time, the instantaneous Web cost, x , falls to zero. The rate of decrease depends

⁴The individual may only execute a web trade that arrives at instant t if he has paid x continuously during an interval $(t - \varepsilon, t]$, for some $\varepsilon > 0$.

on whether the individual is using the phone or Web. When the individual is a phone trader we assume that

$$dx = -\mu_p x dt + \sigma_p x dz, \quad (4)$$

where dz is a Wiener process. Hence, x follows a geometric Brownian motion, with drift parameter μ_p and noise parameter σ_p . Likewise, when the individual is a Web trader we assume that

$$dx = -\mu_w x dt + \sigma_w x dz. \quad (5)$$

The drift and noise parameters capture ongoing changes in Web use costs. The drift parameters capture the tendency for Web costs to fall over time. The noise parameters capture the variability in Web costs. We believe that actively using the Web exaggerates both of these processes, implying that $\mu_p < \mu_w$ and $\sigma_p < \sigma_w$. However, the results that follow do not depend on this assumption.

There exists an endogenous threshold value x^* , below which the individual trades exclusively on the Web and above which the individual trades exclusively on the phone. In the appendix, we show how to solve for x^* , using a boundary condition at $x = 0$, a boundary condition at $x = \infty$, a value matching condition at x^* , a smooth pasting condition at x^* , and an optimality condition for x^* .

Figure 2 plots a simulated realization for the evolution of x for a single trader.⁵ At the beginning of the simulation, x lies far above x^* , implying that the individual does not pay Web flow costs and uses the phone for all transactions. At date 0.73, x falls below x^* . At this point the individual begins paying Web flow costs and conducting transactions on the Web. In this example, the initial Web use is short-lived, as a

⁵This simulation assumes the following parameter values: $\pi_v = .5$, $\pi_w = 1$, $\mu_v = .1$, $\mu_w = .2$, $\sigma_v = .25$, $\sigma_w = .5$, and $\gamma = .05$. See the appendix for details.

positive shock to Web costs quickly drives x back above x^* . The individual does not permanently switch to the Web until date 1.23, when x falls below x^* for the last time.

This model makes several predictions about trading dynamics. First, the introduction of the Web channel will raise transaction frequency. In the short-run, the trader does not use the Web at all, and the average instantaneous transaction frequency is $\lambda \Pr(\eta > c_p)$. In the long-run, the trader will use the Web with probability one, implying that the average instantaneous transaction frequency will be $\lambda \Pr(\eta > c_w)$. This is greater than $\lambda \Pr(\eta > c_p)$ since $c_w < c_p$.

The Web channel will also raise turnover. The instantaneous flow volume in the short-run is $\lambda \int_{\eta > c_p} \eta f(\eta) d\eta$. In the long-run, the instantaneous flow volume is $\lambda \int_{\eta > c_w} \eta f(\eta) d\eta$. This turnover increase will be proportionally smaller than the increase in trading frequency since Web access should lower average transaction size; Web traders will be willing to execute smaller (less valued) transactions, because Web transactions are less costly than non-Web transactions. Recall that η is a proxy for transaction size. For phone trades, the average transaction size is $E[\eta | \eta > c_p]$. For Web trades, the average transaction size is $E[\eta | \eta > c_w]$.

Because Web costs are uncertain, Web traders will sometimes get discouraging signals about the Web and switch back to non-Web trading technologies. Such switching is apparent in Figure 2. If a trader has repeatedly traded on the Web, switching is relatively unlikely to occur in the near future; conditional on observing a long string of trades on the Web, it is likely that the individual has drifted far below x^* . This implies that the individual is likely to continue trading on the Web (at least in the short-run).

II. Data

Our data is provided by Hewitt Associates LLC, a large provider of administrative and consulting services to firms with 401(k) plans. With their help, we identified two large companies that had recently introduced Web access to their 401(k) plans. In choosing these firms, we were careful to minimize any selection biases. Less than one half of the Hewitt's large-client firms offered Web trading as of January 2000. We asked Hewitt to identify the subset of these firms that had made the fewest changes to their plan rules for a wide window around the Web introduction. For example, changes to the menu of investment offerings, rules for matching contributions, or participant eligibility dates could all introduce noise into our attempts to identify a Web-trading effect. We also asked that Hewitt not calculate or prescreen the level of Web trading for any of these firms, so there was no chance of selecting firms based on unusual usage patterns. We also required that the selected firms have at least one year of data both before and after the Web introduction. Two firms' 401(k) plans – code-named Alpha and Omega – survived these filters; summary statistics for these plans are given in Table I. The sample period for Alpha begins in May 1997 and includes all of the data stored by Hewitt. Our sample period for Omega begins in January 1997. We ignore earlier data for Omega because participants were only allowed to trade once a month before this date.

As shown in the table, Omega has considerably more participants and investment options than Alpha. Omega has over 50,000 participants, who choose among 36 investment options covering every major asset class. Alpha offers their 10,000+ participants 11 investment options, but still includes several U.S. equity funds and one bond fund. Participants in both of these plans are able to transfer assets between

investment options through either a phone call or the Web. The phone call can be either through an automated menu system (the majority of calls) or with a live representative, which may entail a wait. All trades placed before 4 P.M. Eastern Time will be executed that day at closing prices.⁶ Trades in international funds are executed at their most recent (past) closing prices.

One interesting feature of many large 401(k) plans is the option to invest in company stock, which is available to participants in Omega. While many experts have pointed out the diversification costs of such own-company investment, company stock remains a popular choice among employees. Nationwide, participants in large (> 5,000 participants) plans invest more than 35 percent of their balances in company stock, and a significant portion of this is discretionary (Holden, VanDerhei, and Quick (2000)). In contrast to many other large plans, participants in Omega are not required to hold the company's matching contributions in company stock. This perhaps partially explains the relatively low holding of company stock – 6.6 percent of balances as of year-end 1999 – by Omega's participants. Across all forms of domestic equity – in company stock, equity mutual funds, and the equity portion of balanced or lifestyle funds – average allocations vary widely between the two plans, with Alpha at 78.2 percent and Omega at 40.4 percent.

The Web channel was opened in August 1998 by both plans. The channel introduction was announced by either a memo or a later article in the plan's newsletter. In no case was there any extra inducement to use the Web channel. The lack of any

⁶In some cases, a participant may place numerous trades in one day. Since such trades will be added together and executed at closing prices by the plan, we follow the same convention and treat the aggregate amount as one single trade. Some parts of our analysis require that we assign each trade a specific time and channel (Web vs. phone). When there are multiple trades aggregated into one, we assign the characteristics of the last trade to the aggregate.

special inducements is consistent with the overall focus of plan sponsors on long-term retirement planning and away from short-term trading. In discussions with representatives of these companies, we learned that the primary reasons for Web introduction were better communication with participants and a desire to give participants easier access to their account information.⁷

The last four rows of Table I suggest the same “Web effect” for Omega that is seen for Alpha in Figure 1. In both plans, the average monthly level of trading is higher after the Web channel is introduced than it is before, and this difference is approximately the same as the average number of Web trades made per month. In the next section, we show that these patterns are significant even after careful controls for other factors.

III. Does the Web affect Trading?

Do the patterns in Figure 1 and Table I demonstrate a “Web effect,” or are they caused by other factors? To answer this question we estimate regressions of the form

$$y_{it} = \alpha_i + \alpha_{iw}WEB_{it} + \beta_{iw} * WEB_{it} * TIME_{it} + \beta_i X_{it} + \varepsilon_{it}, \quad (6)$$

where y_{it} is a measure of trading activity in firm i on day t (described below), WEB_{it} is a dummy variable that takes on the value of 1 when Web trading is available and 0 otherwise, $TIME_{it}$ is the number of days since the Web channel was introduced at

⁷The preceding paragraph is based on private communications with an employee of Alpha and with Hewitt employees who administer these plans. Since both Alpha and Omega chose when to adopt this new technology, we cannot consider their Web introductions as purely random “natural” experiments. Nevertheless, our conversations with Hewitt and our specific findings of no immediate Web impact (discussed in Section III) suggest that the exact timing of the Web introduction was not motivated by participants’ trading demands.

company i , X_{it} is a vector of factors that influence and covary with trading activity, ε_{it} is a (possibly autocorrelated) error term, α_{iw} is the estimated level effect for Web trading, and β_{iw} is the estimated slope effect. If α_{iw} and β_{iw} are both zero, then there is no Web effect. If opening the Web channel causes an immediate increase in trading activity, then α_{iw} should be positive. If the Web channel causes trading to rise over time, then β_{iw} should be positive.

We consider two measures of trading activity on the left-hand-side of (6). Our first measure, $TRADES_{it}$, is the percent of participants that trade in plan i on day t . As shorthand, we refer to $TRADES_{it}$ as the “trading frequency” and express it in units of percent. Thus, $TRADES_{it} = 0.05$ means that 0.05 percent of all participants in plan i executed some trade on day t . Our second measure, $TURNOVER_{it}$, is the total dollars traded by participants in plan i on day t , divided by total balances for all participants in that plan on that day. Thus, if 0.05 percent of all participants in plan i each shift 20 percent of their portfolios on day t , then $TURNOVER_{it}$ would be 0.01 on that day.

The main difficulty in this analysis is in determining the elements of the X_{it} vector. What factors drive trading by participants in 401(k) plans? One obvious set of factors is day-of-the-week or day-of-the-month effects. Since participants cannot trade over the weekend, it is reasonable to expect heavier trading on the first and last (trading) day of the week. Also, since many financial decisions and transactions are made at month-end, participants may also engage in heavier trading around those times. Thus, we include dummy variables on the first and last trading days of the week and month. Our data series is too short to identify any end-of-the-year or tax-day effects, but we do include an overall time trend as part of the X vector. Then,

the coefficient β_{iw} in equation (6) can be interpreted as the additional time trend after the Web introduction.

Our control variables also include the past returns on each plan's investment options. Studies of individual trading behavior show that past returns on a portfolio's securities affect trading (Odean (1998), Barber and Odean (1999 and 2000), Grinblatt and Keloharju (1999)). Similarly, many studies of mutual-fund flows indicate that funds with high past returns attract high net flows (Sirri and Tufano (1993), Chevalier and Ellison (1997), Edelen (1999), Goetzmann, Massa and Rouwenhorst (2000), Bergstrasser and Poterba (2000)), with this relationship significantly nonlinear for funds with the highest past returns. Although each plan only offers a limited set of investment options, there are many possible choices of lags and powers, so that it is necessary to restrict this set to produce some interpretable results. Across both plans, an average of 60 percent of assets are invested either in equity mutual funds or in company stock. Even though participants in Alpha cannot invest in company stock, its past returns may still be salient when participants are forming expectations for future market returns. Thus, it seems reasonable to include the returns to both company stock and to a broad equity index, the S&P 500, as elements of X . Past studies, cited above, suggest that higher orders of returns may also affect trading, perhaps because more extreme returns are more salient for investors. Thus, for both asset classes, company stock and the S&P 500, we include the absolute value of the contemporaneous daily price return, the absolute value of the price return on the previous day, each of these daily returns squared, and, lastly, the standard deviation of daily price returns over the previous 20 trading days. Below, we discuss the robustness of our results to the inclusion of different return variables.

In the end, controlled experiments are the cleanest way to test for treatment effects. Ideally, we would have introduced Web trading for only a random sample of the participants at each firm, and then measured the differences between the Web and non-Web groups. Barring this possibility, we would like to identify some measure of trading activity that has been unaffected by the Web. Hewitt gathers data that allows us to construct such a measure. The “Hewitt 401(k) IndexTM” is designed to measure trading activity between asset classes on a daily basis. The index is constructed from the trading activity in 40 different large-company plans. For each plan, Hewitt calculates the aggregate net dollar amount traded between asset classes on each day.⁸ Individual trades between mutual funds in the same asset class are not counted, and trades of all individuals are added up and netted out to produce an aggregate figure for each firm. For example, if participant i transfers \$10,000 from a large U.S. equity fund to a bond fund, and participant j does the opposite, then Hewitt would cancel these transactions and show no aggregate activity from these two participants. By dividing this aggregate figure by the total assets in the plan, they calculate a Hewitt 401(k) IndexTM for each firm. Note that this index has the same denominator as but a different numerator than the $TURNOVER_{it}$ variable; the numerator of $TURNOVER_{it}$ is the sum of the dollar value of all transactions, irrespective of whether they are within or between asset classes, and without netting any offsetting trades. To construct our non-Web subsample of the Hewitt 401(k) IndexTM, we started with the same 40 plans as Hewitt, then eliminated the 13 firms who had a Web channel by the end of the sample and the ten firms that joined the

⁸The asset classes are money market, GIC/stable value, bond, balanced, lifestyle, large US equity, midsize US equity, small US equity, international, emerging markets, specialty sector, company stock, and self-directed window. Most plans, including Alpha and Omega, do not offer options in every class.

sample after August 4, 1997, which is the first day the index was calculated. We then averaged the aggregate net ratio for these 17 firms on each day to arrive at our *NON – WEB INDEX* variable, which is included in each X_i vector.

Overall, each X_i vector includes ten return-based variables, four timing dummies, a trend variable, and *NON – WEB INDEX*. In the regressions for Omega, we also include a dummy variable, *RULE CHANGE*, to reflect a change in trading rules made during 1999. This change prevented all transfers into one of the international funds; prior to this change, trades involving this fund constituted more than 15 percent of all trades. The dummy variable takes on the value of zero on all days before the rule change and a value of one after the rule change.

Tables II shows the results of estimating (6) for each firm with the variables described above and $y_{it} = TRADES_{it}$, the trading frequency, as the dependent variable.⁹ The table reports coefficient estimates and standard errors for all regressors, with the key test variables given in bold at the top. We use a Newey-West correction with maximum lag length of five trading days to estimate robust standard errors. The results demonstrate economically and statistically significant evidence of the Web’s effect on trading. The coefficient on $WEB_{it} * TIME_{it}$ is significant at the 1 percent level for Alpha and the 5 percent level for Omega. Neither of the level effects are statistically significant. Our calibrations, described below, indicate that all the point estimates for the level effects are economically small compared to the trend effects. From this evidence, we conclude that there is strong evidence that the Web’s effect on trading is growing over time, and no significant evidence of a jump at the time of introduction.

⁹Since the *NON – WEB INDEX* can only be calculated after August 4, 1997, the sample period for the regressions is truncated somewhat from the period listed in Table I.

To calibrate the economic significance of the level and trend Web effects, we can compare their estimated effects over the horizon of our sample to the trading frequency before the Web channel was open. For example, the estimated coefficient on $WEB * TIME$ for Alpha is 0.00072. Over one and a half years – 548 days – approximately the time the Web channel is open in our sample, this point estimate implies an increase in trading frequency of $548 * 0.00072 = 0.395$. If we subtract out the (insignificant) point estimate of the level effect, -0.095, we arrive at a total Web effect over the sample period of 0.300. In Table I, we report that the average monthly trades per participant before the Web was 0.0564; this translates into a daily trading frequency of $(0.0564/21) * 100 = 0.269$ percent. Thus, the total Web effect for Alpha is calibrated to be about $0.300/0.269 = 112$ percent of pre-Web trading. An analogous calculation for Omega yields an increase in trading frequency of $0.00064*548 - 0.024 = 0.327$ percent. In Table I, we report that the average monthly trades per participant was 0.0844 before the Web; this translates into a daily trading frequency of $(0.0844/21)*100 = 0.402$ percent. Thus, the total Web effect for Omega is calibrated to be about $0.327/0.402 = 81$ percent of pre-Web trading. The calibrated effects for both firms are consistent with back-of-the-envelope comparisons of trading before and after the Web, as shown in Table I. Averaging these two calibrations, we estimate that the Web nearly doubles trading at an 18-month horizon.

While our main focus is on the level and trend coefficients for the Web, it is worthwhile to examine the other coefficients and discuss the robustness of our results to alternative specifications. The coefficients on $NON - WEB INDEX$ are positive and significant for both plans, with t -statistics above four in both cases. Of the timing variables, only the first day of the (trading) week shows a significant effect

(1 percent level) in both plans, with point estimates implying effects on trading frequency of 0.12 percent. The qualitative results for the level and trend Web variables are unchanged if we use separate dummies for each day of the week.

Turning next to the past-returns variables, the most robust results are for the squared returns. The lagged, squared S&P 500 returns are positive and significant at the 1 percent level for both firms, and the squared company-stock returns, both lagged and contemporaneous, are positive and significant at the 1 percent level for Omega. These results suggest, not surprisingly, that investors react nonlinearly to market moves, with extreme moves causing disproportionately large reactions. In unreported regressions, we included up to five lags of daily absolute and squared returns for both company stock and the S&P 500. The influence of these variables was much smaller than for the first two days, and the effect on the Web variable coefficients and standard errors was negligible.

Overall, the qualitative results of Table II are robust to a broad range of control variables and alternative specifications. We conclude from this evidence that the pattern in Figure 1 is no illusion: the introduction of Web trading has a large effect on the trading frequency of plan participants. This result leads to a natural follow-up question: does the Web also affect the dollar volume of trade? It is possible, for example, that the large increase in trading frequency occurs because participants break up large trades into smaller pieces, with only a small or negligible increase in the total dollars traded. Also, if Web trading is predominantly an activity of young participants with small balances, then the Web's impact on dollar volume would be smaller than its impact on trading frequency. We analyze the Web's effect on dollar volume by using $TURNOVER_{it}$ as the dependent variable in (6). Recall that

$TURNOVER_{it}$ is the total dollars traded by participants in plan i on day t , divided by total balances for all participants in that plan on that day.

Table III summarizes the results. The coefficient on $WEB * TIME$ for Omega is positive and significant at the 1 percent level. The analogous coefficient for Alpha is positive and has a t -statistic of 1.88, implying a two-tailed p -value of 0.06. To evaluate the economic significance of these point estimates, we follow a procedure analogous to the one used to assess trading frequency. That is, we compute the total effect on $TURNOVER$ over 18 months that is implied by the point estimates, and then we compare this effect to the average turnover before the Web channel was opened. This computation yields an estimated increase of 45 percent for Alpha and 64 percent for Omega.¹⁰ The average effect across the two firms is about 55 percent, or a little more than half the estimated effect on trading frequency. Thus, the evidence suggests that the Web increased turnover, but not by as much as it increased trading frequency. Overall, the evidence from Tables II and III demonstrates that the Web had an economically large impact on investor behavior.

The patterns reported above are predicted by the simple model of Section I. First, the Web trading effects increase with time, a natural implication of falling Web usage costs (i.e., falling x). Second, the Web effects on trading frequency are proportionally larger than the Web effects on trading turnover, because Web trades are expected to be smaller, in dollars, than phone trades. As we show in the next section, this

¹⁰This computation uses the coefficients reported in the top two rows of Table III. The total Web effect for Alpha was $0.00038 * 548 - 0.059 = 0.149$ percent. The average daily pre-Web turnover in Alpha, not reported elsewhere in the paper, is 0.334 percent. Thus, the Web increased turnover by $0.149/0.334 = 45$ percent. The analogous calculation for Omega is $0.00026 * 548 - 0.021 = 0.122$ percent. The average daily pre-Web turnover at Omega was 0.192, implying an increase of $0.122/0.192 = 64$ percent.

prediction about relative trade sizes is also confirmed by the data.

IV. An Empirical Portrait of Web Trading

In this section, we analyze several elements of Web trading. Part A looks at participants' propensities to try the Web for trading. Part B studies whether participants who try the Web tend to stick with it. In Part C, we compare the sizes of trades made on the Web versus the phone. Part D analyzes whether the Web has increased "short-term" trading, which we define in two different ways.

A. Who trades on the Web?

Participants will trade on the Web if they have low Web usage costs. To identify who has low costs, we construct a sample of participants who executed at least one trade, either by phone or by Web, since the date that the Web channel was opened. Conditional on being in this sample, we then estimate the likelihood of executing at least one trade on the Web. As independent variables, we include age, tenure at the firm, salary, total balance in the 401(k) plan, length of time participating in the plan, contribution rates to the plan, monthly frequency of trading before the Web introduction, and dummy variables for sex, marital status, retirement status, and current employment status at the firm. All of the continuous variables, except age and trading frequency, are in logs.

Table IV summarizes the results of logit estimations for both firms. As might be expected, the coefficients on age are negative and significant for both firms. The coefficients on both salary and plan balances are positive and significant in both regressions. We only have gender data for Alpha: in that regression, the coefficient on the male dummy variable is positive and significant. Introspection and casual

empiricism from other contexts suggests that young, male, well-educated, wealthy people are more likely to be early adopters of computer technology, and these results are consistent with such beliefs.

The evidence on other demographic variables also demonstrates some interesting patterns. Retired participants are less likely to try the Web for trading at Omega. Participants coded as “terminated,” a mutually exclusive set from those who are retired, are also less likely to try the Web, with negative and significant coefficients at both firms. It is plausible that such participants are in less active information networks about plan changes and thus are less likely to know about plan changes. While they might receive the same formal documents as other participants, they are no longer able to hear about plan changes through word-of-mouth at the workplace. Finally, the coefficient on “pre-Web trades per month” is negative and significant at the one percent level for Alpha, and is negative and insignificant for Omega. This evidence suggests that traders who are already experienced and familiar with phone trading are less likely to try the Web.

B. Once a Web trader, always a Web trader?

The model sketched in Section I and illustrated in Figure 2 predicts that there will be some switching back and forth between Web and phone transactions. The specific assumptions of our model imply that, over time, the expected total costs of trading on the Web are falling and more traders will migrate to the Web. Over any short horizon, however, some Web traders receive bad news about their own costs and switch back to making phone transactions. The probability of switching back to the phone grows smaller the longer someone has been a Web trader because, on average, they have moved lower into the critical Web trading zone of Figure 2. In

this subsection, we examine the switching behavior of plan participants and analyze whether it is consistent with this framework.

Figure 3 shows the trade channel for participants starting with their *first* trade on the Web. In our two sample firms, 15,421 participants made at least one trade on the Web. By the analysis of Section I, at the time of this trade, these participants are within the critical range of Figure 2, with $x < x^*$. Because Web costs are uncertain, some of these traders will get discouraging signals about the Web and switch back to the phone. If a trader has repeatedly traded on the Web, such switching is relatively unlikely to occur in the near-future. Conditional on observing a long string of trades on the Web, it is likely that the individual has drifted far below x^* . Thus, we expect the switching probability to fall with each Web trade.

In the data, subsequent to the first Web trade, 10,413 participants make at least one more trade; of this group, 9,172 (88 percent) make their next trade on the Web, while 1,241 (12 percent) make their next trade by phone. With each subsequent trade on the Web, the probability of making the next trade on the Web increases, with an empirical frequency of 94 percent on the third trade and 96 percent on the fourth trade. Furthermore, the bottom branch of Figure 3 shows that the reverse is also true: the longer the sequence of trades by phone, the higher the likelihood that the next trade will be by phone. That is, conditional on the first trade being by the Web, the next trade is by phone 12 percent of the time; conditional on this second trade being by phone the third trade is by phone 57 percent of the time; conditional on the third trade being by phone the fourth trade is by phone 77 percent of the time.

The middle branches of Figure 3 are also consistent with the model described in Section I. In the model, the most recent trade provides the most information

about the position of x relative to the critical value x^* . A trader who executes in the order Web-Web-phone should be less likely to make his next trade on the Web than is a trader who has executed as Web-phone-Web. This difference is statistically significant in the data, with Web-Web-phone traders making their next trade on the Web 58 percent of the time, and Web-phone-Web traders making their next trade on the Web 81 percent of the time.

C. Trade Size

In the model discussed in Section I, Web traders have lower marginal costs per transaction and so should be willing to undertake trades with smaller expected benefits. A natural test of this prediction is to look at the average size of transactions executed on and off the Web. Figure 4 plots the 20-trading-day moving average of trade size before and after Web introduction for both companies. The size of a transaction is defined as the total amount transferred in one direction – there is no double-counting. For example, if \$10,000 is transferred out of company stock and into a bond mutual fund, then the transaction size would be \$10,000, not \$20,000. Recall that if a participant makes multiple transactions in one day, we use only the net change of all these transactions and count this net change as a single transaction. Panel A shows the results for Alpha and Panel B shows the results for Omega. Both panels show that the average size of Web transactions is substantially smaller than phone transactions. For example, in the period before the Web introduction at Omega, average transaction size varies between \$35,000 and \$55,000; at the time of the Web introduction in August 1998, this average stands at about \$45,000. After the Web introduction, the average size of phone trades trends upward and stands at \$70,000 by the end of the sample period. In contrast, the average size of Web

transactions starts at \$25,000 and never exceeds \$40,000. The patterns in Alpha are similar, with average phone transactions always larger than contemporaneous average Web transactions.

The patterns in Figure 4 result from some combination of two effects: (1) Web trades have lower turnover (per transaction) than do phone trades, or, (2) Web traders have lower balances than do phone traders. In fact, both of these effects seem to contribute in some part to the results of Figure 4. Figure 5 shows the average turnover per transaction separately for each channel. To compute average turnover, we compute the fraction of balances transferred in each transaction and then average across all transactions on each day. This series is relatively stable across the sample period, with Web trades always smaller than phone trades. Figure 5 also demonstrates that the average trade is large in relation to total balances: on average, participants in Alpha transfer more than half their portfolio with each transaction; participants in Omega transfer more than one-third of their portfolios.

Figure 5 implies that even if all participants had exactly the same balances, we would observe phone trades to be smaller (in dollars) than Web trades. In fact, Web traders and phone traders do not have the same balances. Figure 6 shows that phone traders are almost always “richer.” In Omega, the typical phone trader has balances more than \$50,000 higher than the typical Web trader. For Alpha, the gap is much smaller, but phone traders have higher balances for almost all of the sample period.

At first glance, there appears to be some tension between these findings and our earlier finding that wealthier participants are more likely to try the Web (Section IV.A). If wealthy participants are more likely to try the Web, why is it that the balances of phone traders are higher than the balances of Web traders? The resolu-

tion of this ostensible tension is in the distinction between traders and trades. The most frequent traders tend to be relatively wealthy participants who are engaging in short-term trades by phone (see Section IV.D below). The results of Figure 6 average across all *trades*, so that these frequent traders are counted many times. By contrast, the demographic results of Table I count each trader once. Thus the high-balance frequent phone traders do not dominate that earlier analysis.

D. Does the Web increase “short-term” trading?

The previous results established that the Web increased trading in these plans, but what kind of trading increased? We saw in Section IV.C that Web trades tend to be smaller than phone trades: does the size signify anything else about participants’ behavior? Does the fact that trading is only a “click” away lead participants to more short-term behavior? Since participants do not have to pay any capital gains taxes in 401(k) plans, they may be inclined to use them as vehicles to make short-term bets on the relative movements between asset classes. Furthermore, a popular strategy among participants in some plans is to take advantage of stale prices by buying (selling) funds with thinly or asynchronously traded securities on days with large market increases (decreases). This strategy has grown so popular that some plans and mutual fund companies have introduced trading restrictions to prevent it.¹¹

In this section, we use two different measures to classify trades as “short-term,” and then we study the pattern of speculative trading through each channel. Our first definition of a short-term trade is a trade that is at least partially reversed within five trading days. We will call these “reversed trades”. Under this definition, we would

¹¹For descriptions of this strategy, see Atchison, Butler, and Simonds (1987), Chalmers, Edelen and Kadlec (2000), and Goetzmann, Ivkovich, and Rouwenhorst (2000). This strategy is particularly profitable in international funds, but also works for domestic funds.

classify both the original trade and its reversal as a reversed trade. While not exactly day-trading – it is more like “week-trading” – such trades are likely to have been made with the intention of capturing some perceived short-term profit opportunity. Of course, there may be many trades that have this intention and are not reversed so quickly, so this filter is imperfect. We have developed this definition without reference to our model, since that model is mute on the motivation of trade. In our model, all trades are qualitatively the same, differing only in size.

Figure 7 plots the frequency of reversed trades by channel. Three interesting patterns are apparent. First, there are many reversed trades. By the end of the sample, these trades make up about one-half of the trading at Alpha and one-third at Omega. Second, there is an upward trend in reversed trades over the whole sample period for both firms. This trend begins before the Web channel is opened. Third, reversed trades constitute a smaller fraction of Web trades than of phone trades. At both firms, phone trades are more likely to be reversed than are Web trades. Thus, by this measure, the Web has proportionally decreased short-term trading.

Our second filter for speculation uses the time-of-day for a trade. Recall that participants’ trades may be placed at any time but are executed only once per day and use market closing prices. For both firms, this means that all trades placed before 4 P.M., Eastern Time, are executed at that day’s closing price, while any trades executed after that time must wait the close on the next day. Since any allocation decision would be more accurate later in the day (and closer to the closing prices), it seems likely that short-term traders – particularly the ones trying to profit from stale prices – would be influenced more by the time of day than would longer-term traders. The profitability of this strategy is highest if traders wait until the very end

of the day to execute, because that is when they will have the most information to exploit.

To examine the timing issue, we classify all trades by the time of day they were placed. Table V summarizes the results. All times are Eastern Time. Trades placed between 3 P.M. and 4 P.M. – the hour before the market closes – we define as “last-hour trades,” and are our second proxy for short-term trades. As shown in the table, a large fraction of trades are made in the last hour, and this fraction is significantly higher by phone than by Web: 46.4 percent of the phone trades and 26.7 percent of the Web trades for Alpha, and 51.2 percent of the phone trades and 29.1 percent of the Web trades for Omega. This reinforces the findings of Figure 7, which suggest that a large fraction of trades is driven by short-term motives. In fact, the categories overlap considerably: of all the last-hour trades, 81.3 percent for Alpha and 54.5 percent for Omega are also reversed trades. For trades made at all other times, only 39.3 percent for Alpha and 24.9 percent for Omega are reversed.

There is no way to know for sure if some participants are using the Web to gather information for their trades, but then use the phone to execute these trades. The fact that this trading goes on during the working day means that some participants may feel the need to hide this activity from their co-workers, and in the absence of a private office, it may be more discreet to use an automated menu on the phone than a computer screen.¹² Such concerns could explain why short-term traders seem to have a preference for phone transactions. On the other hand, the regression evidence in Table IV shows that frequent traders – who do most of the reverse and last-hour trading – are less likely to try the Web for even one trade. It may be that these

¹²Not all participants have computer access from their desks, but instead must use public computer kiosks.

active traders have very low costs for phone trades and see no need to switch to a new technology. Overall, the evidence suggests that it is the infrequent and longer-term traders who are the first to move to the Web, and that opening a Web channel does not increase the proportion of short-term trades.

V. Conclusion

We have identified the impact of a Web channel on the trading decisions of about 100,000 participants in two large 401(k) plans. To measure this Web effect, we control for numerous other sources of variability in trading activity. We find that, at a horizon of 18 months, a Web channel nearly doubles trading frequency. Turnover – the fraction of total portfolio value traded – increases by over 50 percent. Trading frequency increases by more than turnover, since Web trades tend to be smaller than phone trades both in dollars and as a portfolio fraction.

We also document several other findings about the behavior of Web trading. It appears that traders tend to use either the phone or the Web; very few individuals trade continuously on both channels. The Web traders are the ones with the lowest costs of using the new technology, like young, male, wealthy participants who may already be familiar with Web-based technologies. Finally, the new Web trading does *not* appear to be disproportionately short-term “day-trading” behavior.

Taken together, our results paint a detailed portrait of Web trading. This picture is broadly consistent with a theoretical framework where Web usage is driven by stochastically falling transaction costs. In this framework, the impact of the Web channel is not immediate, but as more and more traders migrate to the Web, the impact on total trading increases. In the data, after 18 months, this increase in total trading is economically large. We conclude that some of the “new economy” really is

new, after all.

Appendix: Model Solution

This appendix extends the analysis of Section II.

Assuming a discount rate of γ , and applying Ito's Lemma, we can represent the continuous-time Bellman Equation of the phone trader as

$$\gamma P(x) = \pi_p + (-\mu_p x)P'(x) + \frac{1}{2}\sigma_p^2 x^2 P''(x). \quad (7)$$

Likewise, we can represent the Bellman Equation of the Web trader as,

$$\gamma W(x) = \pi_w - x + (-\mu_w x)W'(x) + \frac{1}{2}\sigma_w^2 x^2 W''(x). \quad (8)$$

A particular solution for P is $P(x) = \frac{\pi_p}{\gamma}$. Note that this is equivalent to the value of the policy “always trade on the phone regardless of the value of x .” Two linearly independent solutions for the homogenous part of P take the general form Cx^{r_p} , where

$$r_p = \frac{\left(1 + \frac{2\mu_p}{\sigma_p^2}\right) \pm \sqrt{\left(1 + \frac{2\mu_p}{\sigma_p^2}\right)^2 + \frac{8\gamma}{\sigma_p^2}}}{2} \quad (9)$$

A particular solution for W is $W(x) = \frac{\pi_w}{\gamma} - \frac{x}{\gamma + \mu_w}$. Note that this is equivalent to the expected value of the policy “always trade on the Web regardless of the value of x .” Two linearly independent solutions for the homogenous part of W take the general form Cx^{r_w} , where r_w is defined analogously to r_p .

There exists a Web cost, x^* , below which the individual trades exclusively on the Web and above which the individual trades exclusively on the phone. To solve for

x^* , we apply a boundary condition at $x = 0$,

$$W(0) = \frac{\pi_w}{\gamma}, \quad (10)$$

a boundary condition at $x = \infty$,

$$\lim_{x \rightarrow \infty} P(x) = \frac{\pi_p}{\gamma}, \quad (11)$$

a value matching condition at x^* ,

$$W(x^*) = P(x^*), \quad (12)$$

a smooth pasting condition at x^* ,

$$W'(x^*) = P'(x^*), \quad (13)$$

and an optimality condition for x^* . To represent this optimality condition, define a class of x^* -contingent value functions.

$$G(x | x^* = y) = \begin{cases} W(x) & \text{if } x \leq y \\ P(x) & \text{if } x > y \end{cases} \quad (14)$$

Optimality requires that x^* be chosen to select the supremum value function:

$$F(\cdot) = \max_y G(\cdot | x^* = y) \quad (15)$$

Applying these constraints, and solving the system of two second-order differential

equations implies,

$$x^* = \frac{\left(\frac{\mu_w}{\gamma} + 1\right) (\pi_w - \pi_p)}{\left(1 - \frac{1}{r'_p}\right) \left(1 - \frac{1}{r'_w}\right)}, \quad (16)$$

where r'_p and r'_w are respectively negative and positive roots.

References

- Agnew, Julie, Pierluigi Balduzzi, and Annika Sunden, 2000, Portfolio choice, trading, and returns in a large 401(k) Plan, Working paper, Boston College.
- Ameriks, John, and Steve Zeldes, 2000, How do household portfolio shares vary with age?, Working paper, Columbia University.
- Atchison, Michael, Kirt Butler, and Richard Simonds, 1987, Nonsynchronous security trading and market index correlation, *Journal of Finance* 42, 111-118.
- Barber, Brad M., and Terrance Odean, 1999, Online investors: Do the slow die first?, Working paper, UC Davis.
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773-806.
- Bergstrasser, Daniel, and James Poterba, 2000, Do after-tax returns affect mutual fund flows, National Bureau of Economic Research, Working paper No. 7595.
- Chalmers, John M.R., Roger M. Edelen, and Gregory Kadlec, 1999, The wildcard option in transacting in mutual fund shares, Rodney L. White Center for Financial Research, Working paper No. 00-03.
- Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167-1200.
- Dugan, Ianthe Jeanne, 1999, Senate hearings to probe day trading; Impact on market, traders at issue, *Washington Post* (September 14), Financial Section, E01.
- Edelen, Roger M., 1999, Investor flows and the assessed performance of open-end mutual funds, *Journal of Financial Economics* 53 (3). 439-466.
- Goetzmann, William N., Zoran Ivkovich, and Geert Rouwenhorst, 2000, Day-trading international mutual funds: evidence and policy solutions, Working paper, Yale University.
- Goetzmann, William N., Massimo Massa, Geert Rouwenhorst, 2000, Behavioral factors in mutual-fund flows, Working paper, Yale University.
- Grinblatt, Mark and Matti Keloharju, 1999, What makes investors trade? Working paper, UCLA.

Holden, Sarah, Jack VanDerhei, and Carol Quick, 1999, 401(k) plan asset allocation, account balances, and loan activity in 1998, *Investment Company Institute Perspective* 6 (1).

Kunath, Brian D., 1999, Day trading: The 'Dealerization' of the Market, *Global Investment Magazine* (December), 51.

Odean, Terrance, 1998, Are investors reluctant to realize their losses, *Journal of Finance* 53, 1775-1798.

Sirri, Eric and Peter Tufano, 1993, Buying and selling mutual funds: funds, performance, fees, and service, Working paper, Harvard Business School.

Stillwell, Dennis, 2000, High-speed danger alert, *Journal of Commerce*, April 5, Global Commerce Section, 13.

Table I
401(k) Plan Characteristics

This table presents summary statistics for the 401(k) plans of firms Alpha and Omega. Because of data availability, demographic information is limited to participants who had positive plan balances or plan activity in 1998 or 1999.

	Alpha	Omega
Number of participants ¹	More than 10,000	More than 50,000
Data range	5/19/97 – 3/3/00	1/27/97 – 1/26/00
Number of investment options	11	36
Company stock available in plan?	No	Yes
Percent of plan assets in equity ^{2, 3}	78.2%	40.4%
Percent of plan assets in company stock ²	0.0%	6.6%
Average age ^{2, 4}	40.7	52.8
Average years since original hire ^{2, 4}	8.6	18.6
Average plan balance ^{2, 4}	\$68,202	\$112,456
Average contribution rate ^{2, 5}	6.49%	9.27%
Percent of participants who trade at least once in sample	41%	45%
Month of Web introduction	August 1998	August 1998
Average trades per month per participant before Web introduction ⁶	0.0564	0.0844
Average trades per month per participant after Web introduction ⁶	0.1285	0.1407
Average trades per month per participant on Web ⁶	0.0666	0.0597
Percent of participants who trade at least once on Web	24%	15%

¹ All participants in sample, including those who drop out of the plan before the end of the sample.

² At year-end 1999.

³ Includes all equity mutual fund and company stock balances.

⁴ Participants who had positive plan balances at year-end 1999 or who had plan activity in 1998 or 1999.

⁵ Current employees as of year-end 1999 only.

⁶ All sales and purchases on a given day by a participant are counted as one “trade.”

Table II
Determinants of Trading Frequency

The dependent variable, *TRADES*, is the percent of participants in each company who trade on each day. *WEB* is a dummy set to one if Web trading has been introduced. *WEB * TIME* is the interaction of *WEB* and *TIME*, the number of calendar days that have passed since Web trading was introduced. *NON-WEB INDEX* is the equally-weighted average of the daily percent of plan balances traded between asset classes for 17 companies without Web trading. *|S&P 500|* and *|LAG S&P 500|* are the absolute values of the S&P 500 return today and yesterday, respectively. $(S\&P\ 500)^2$ and $(LAG\ S\&P\ 500)^2$ are the squares of *|S&P 500|* and *|LAG S&P 500|*, respectively. *STD(S&P 500)* is the twenty-day lagged standard deviation of the S&P 500 price return. *|COMPANY STOCK|* and *|LAG COMPANY STOCK|* are the absolute values of the company stock's return today and yesterday, respectively. $(COMPANY\ STOCK)^2$ and $(LAG\ COMPANY\ STOCK)^2$ are the squares of *|COMPANY STOCK|* and *|LAG COMPANY STOCK|*, respectively. *STD(COMPANY STOCK)* is the twenty-day lagged standard deviation of the company stock price return. *START WEEK*, *END WEEK*, *START MONTH*, and *END MONTH* are dummies set to one if the day is the first trading day of the week, the last trading day of the week, the first trading day of the month, and the last trading day of the month, respectively. *RULE CHANGE* is a dummy set to one for Omega after the institution of a new rule restricting trading on an international fund. *TREND* is the number of calendar days that have elapsed since January 1, 1997. Newey-West robust standard errors (five lags) are reported in parentheses below the OLS point estimates.

SEE NEXT PAGE FOR TABLE

Independent Variables	Alpha	Omega
WEB	-0.0952 (0.0582)	-0.0241 (0.0518)
WEB * TIME	0.00072** (0.00020)	0.00064* (0.00028)
<i>NON-WEB INDEX</i>	140.0920** (22.3813)	81.6954** (18.6263)
<i>/S&P 500/</i>	0.2079 (1.8021)	3.4935* (1.7030)
<i>(S&P 500)²</i>	60.3637 (41.8465)	-1.0643 (32.3171)
<i>/LAG S&P 500/</i>	-1.5036 (1.7831)	-1.1995 (1.8917)
<i>(LAG S&P 500)²</i>	202.9359** (37.9670)	149.2298** (52.8017)
<i>STD(S&P 500)</i>	-2.7578 (3.9289)	-5.2980* (2.4152)
<i>/COMPANY STOCK/</i>	0.9426 (0.8968)	2.5601** (0.8864)
<i>(COMPANY STOCK)²</i>	-0.9552 (13.0412)	41.8645** (9.5403)
<i>/LAG COMPANY STOCK/</i>	0.2651 (0.9386)	0.2604 (1.1452)
<i>(LAG COMPANY STOCK)²</i>	-1.4287 (12.6027)	32.0787** (12.1468)
<i>STD(COMPANY STOCK)</i>	7.7854** (2.5549)	7.5359** (1.7749)
<i>START WEEK</i>	0.1208** (0.0173)	0.1287** (0.0172)
<i>END WEEK</i>	-0.0070 (0.0122)	0.0352* (0.0152)
<i>START MONTH</i>	-0.0488 (0.0410)	-0.0379 (0.0298)
<i>END MONTH</i>	-0.0149 (0.0308)	-0.0113 (0.0318)
<i>RULE CHANGE</i>		-0.0547 (0.0863)
<i>TREND</i>	0.00030 (0.00016)	0.00034** (0.00009)
Constant	-0.1894* (0.0953)	-0.0696 (0.0656)

* Significant at the 5 percent level

** Significant at the 1 percent level

Table III
Determinants of Turnover

The dependent variable, *TURNOVER*, is the daily dollar value of all sales as a percent of total balances on that day. *WEB* is a dummy set to one if Web trading has been introduced. *WEB * TIME* is the interaction of *WEB* and *TIME*, the number of calendar days that have passed since Web trading was introduced. *NON-WEB INDEX* is the equally-weighted average of the daily percent of plan balances traded between asset classes for 17 companies without Web trading. *|S&P 500|* and *|LAG S&P 500|* are the absolute values of the S&P 500 return today and yesterday, respectively. $(S\&P\ 500)^2$ and $(S\&P\ 500)^2$ are the squares of *|S&P 500|* and *|LAG S&P 500|*, respectively. *STD(S&P 500)* is the twenty-day lagged standard deviation of the S&P 500 price return. *|COMPANY STOCK|* and *|LAG COMPANY STOCK|* are the absolute values of the company stock's return *|LAG COMPANY STOCK|*, respectively. *STD(COMPANY STOCK)* is the twenty-day lagged standard deviation of the company stock price return. *START WEEK*, *END WEEK*, *START MONTH*, and *END MONTH* are dummies set to one if the day is the first trading day of the week, the last trading day of the week, the first trading day of the month, and the last trading day of the month, respectively. *RULE CHANGE* is a dummy set to one for Omega after March 19, 1999 to reflect a new rule instituted to restrict trading on an international fund. *TREND* is the number of calendar days that have elapsed since January 1, 1997. Newey-West robust standard errors (five lags) are reported in parentheses below the OLS point estimates.

SEE NEXT PAGE FOR TABLE

Independent Variables	Alpha	Omega
WEB	-0.0591 (0.0577)	-0.0205 (0.0119)
WEB * TIME	0.00038 (0.00020)	0.00026** (0.00006)
<i>NON-WEB INDEX</i>	194.0439** (25.9074)	31.6382** (4.3597)
<i>/S&P 500/</i>	0.1820 (2.6553)	0.1188 (0.3925)
<i>(S&P 500)²</i>	83.1802 (73.1748)	11.0200 (9.9516)
<i>/LAG S&P 500/</i>	-3.3204 (2.5653)	-0.4162 (0.3861)
<i>(LAG S&P 500)²</i>	302.6174** (64.2860)	44.1060** (8.7565)
<i>STD(S&P 500)</i>	-4.1020 (4.1555)	-0.3025 (0.4919)
<i>/COMPANY STOCK/</i>	0.5350 (1.2829)	-0.0198 (0.2006)
<i>(COMPANY STOCK)²</i>	9.5484 (20.8628)	1.8339 (1.8384)
<i>/LAG COMPANY STOCK/</i>	-0.4893 (1.1520)	-0.0768 (0.1800)
<i>(LAG COMPANY STOCK)²</i>	11.0976 (13.7542)	-0.2390 (1.7718)
<i>STD(COMPANY STOCK)</i>	8.5034** (2.4703)	1.3509** (0.3771)
<i>START WEEK</i>	0.1200** (0.0217)	0.0208** (0.0033)
<i>END WEEK</i>	0.0407* (0.0195)	0.0062* (0.0032)
<i>START MONTH</i>	-0.0442 (0.0487)	-0.0093 (0.0076)
<i>END MONTH</i>	-0.0222 (0.0408)	-0.0029 (0.0064)
<i>RULE CHANGE</i>		-0.0586** (0.0147)
<i>TREND</i>	0.00064** (0.00016)	0.00011** (0.00002)
Constant	-0.3360** (0.0956)	-0.0643** (0.0180)

* Significant at the 5 percent level

** Significant at the 1 percent level

Table IV
Demographics of Web Traders

This table presents the results of a binary logit regression of the likelihood of trading at least once on the Web in the sample, conditional upon trading at least once since Web trading was introduced. Participants must have been enrolled before Web introduction and had a positive plan balance or plan activity in 1998 or 1999 in order to have a full set of right-hand-side variables and be included in the regression. *MALE* and *MARRIED* are dummies set to one if the participant is male and married, respectively. *AGE* is the participant's age at December 31, 1999, and *TENURE* is the log of the number of years since the participant's original hire date, as of December 31, 1999. *SALARY* is the log of 1999 salary, and *BALANCES* is the log of total plan balance at year-end 1999. *PARTICIPATION LENGTH* is the log of the number of years since the participant originally enrolled in the plan. *PRE-WEB TRADES PER MONTH* is the number of trades per month the participant executed before the introduction of the Web. *CONTRIBUTION RATE* is the contribution rate effective at year-end 1999, in integers (e.g. "5" percent). *TERMINATED* and *RETIRED* are dummies set to one if the participant has been terminated or retired, respectively as of year-end 1999. Standard errors are given in parentheses below the point estimates.

Independent Variables	Alpha	Omega
<i>MALE</i>	0.4093** (0.0675)	
<i>MARRIED</i>	0.0564 (0.0639)	0.2783** (0.0406)
<i>AGE</i>	-0.0369** (0.0039)	-0.0479** (0.0028)
<i>TENURE</i>	-0.0059 (0.1496)	0.0674 (0.0578)
<i>SALARY</i>	0.1879** (0.0231)	0.0619** (0.0063)
<i>BALANCES</i>	0.3264** (0.0444)	0.2080** (0.0205)
<i>PARTICIPATION LENGTH</i>	-0.3683** (0.1360)	0.0287 (0.0567)
<i>PRE-WEB TRADES PER MONTH</i>	-0.3811** (0.0864)	-0.0535 (0.0275)
<i>CONTRIBUTION RATE</i>	0.0129 (0.0103)	-0.0076* (0.0032)
<i>TERMINATED</i>	-0.4862** (0.1615)	-0.1978** (0.0706)
<i>RETIRED</i>	-0.2244 (0.5898)	-0.3655** (0.0916)
Constant	-2.7538** (0.3831)	-0.7128** (0.1927)

* Significant at the 5 percent level

** Significant at the 1 percent level

Table V
Distribution of Trade Entry Times By Channel

This table presents the trades that have been entered through each channel at each hour (U.S. Eastern Time) since Web trading was introduced, as a percent of all trades that have gone through each channel since the introduction of Web trading.

Time	Alpha		Omega	
	Phone	Web	Phone	Web
12:00 A.M. – 8:59 A.M.	5.8%	12.0%	4.2%	13.6%
9:00 A.M. – 9:59 A.M.	2.6%	4.0%	3.1%	4.1%
10:00 A.M. – 10:59 A.M.	3.5%	4.9%	3.6%	4.7%
11:00 A.M. – 11:59 A.M.	4.2%	4.8%	3.7%	5.1%
12:00 P.M. – 12:59 P.M.	4.4%	5.7%	5.1%	7.0%
1:00 P.M. – 1:59 P.M.	5.2%	6.9%	6.1%	7.9%
2:00 P.M. – 2:59 P.M.	8.6%	8.6%	9.7%	10.5%
3:00 P.M. – 3:59 P.M.	46.4%	26.7%	51.2%	29.1%
4:00 P.M. – 4:59 P.M.	5.3%	4.6%	6.3%	2.9%
5:00 P.M. – 11:59 P.M.	13.9%	21.9%	7.0%	15.0%
Total	100.0%	100.0%	100.0%	100.0%




Figure 1. Alpha: Daily trading frequency, 20-trading-day moving average. On each trading day, the percent of participants enrolled in company Alpha's 401(k) plan who traded on that day is calculated.




Figure 2. A simulated realization of an x path for an individual trader. This figure plots a simulated path for the current flow cost of an individual's Web use. Whenever the flow cost lies above a critical threshold, x^* (see formula in the appendix), the individual does not pay the flow cost, stays off of the Web, and uses the phone to execute trades. When the flow cost lies below x^* , the individual actively uses the Web and is able to trade on the Web when sufficiently appealing trading opportunities arise. This simulation assumes the following model calibration: $\pi_v = .5$, $\pi_w = 1$, $\mu_v = .1$, $\mu_w = .2$, $\sigma_v = .25$, $\sigma_w = .5$, and $\gamma = .05$.

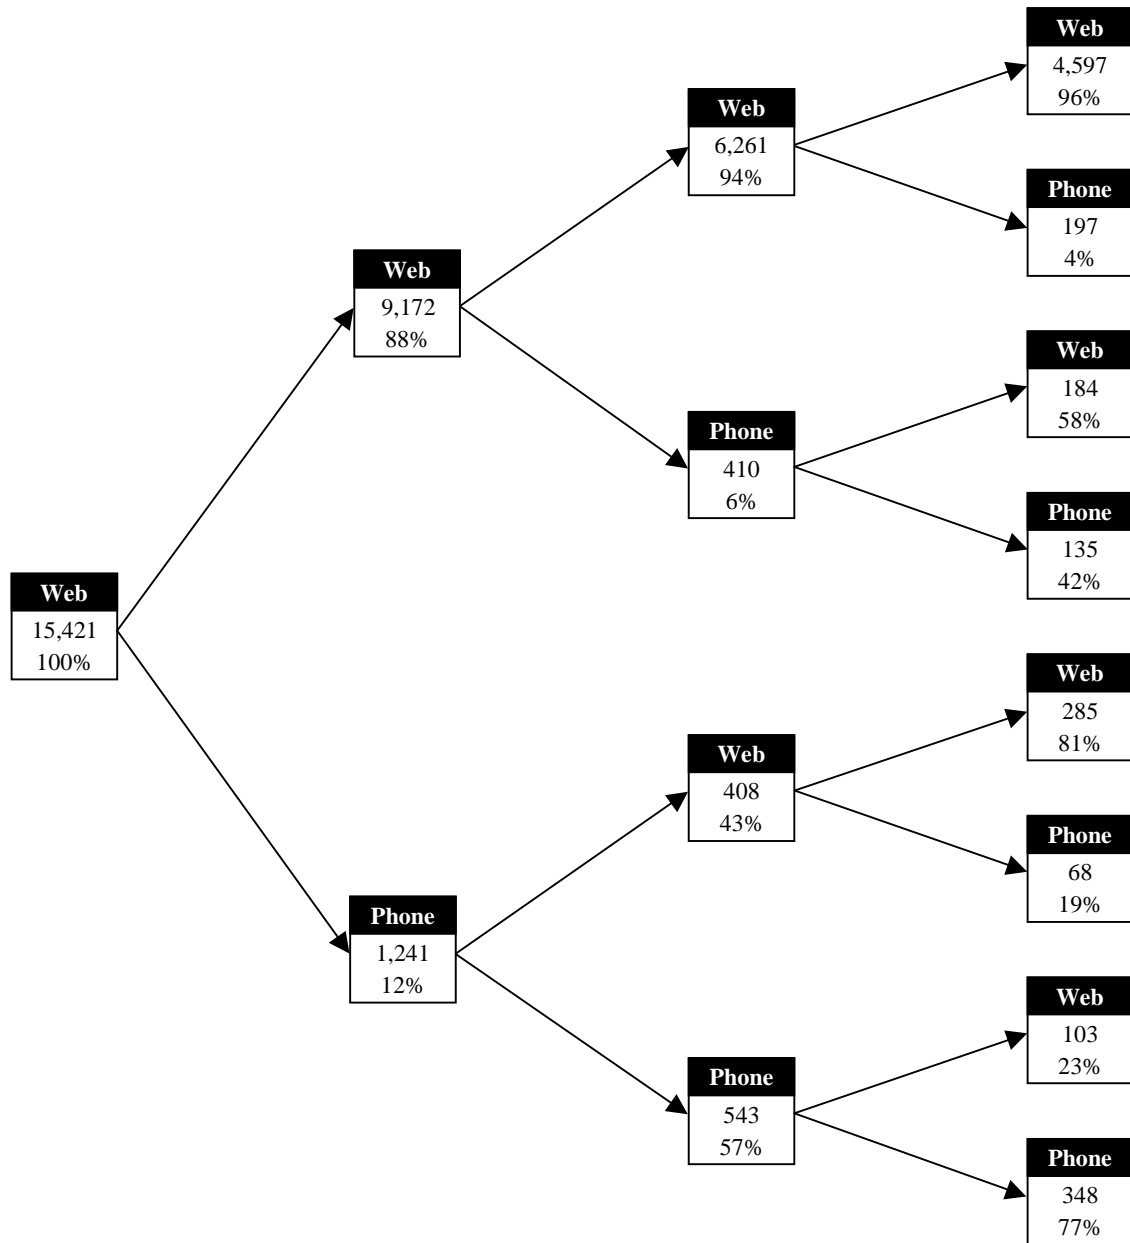


Figure 3. Trading channel tree. This figure traces the channels through which participants trade, starting with their first Web trade. The top number in each box is the number of participants who reach that particular node, across both companies. For example, 15,421 participants made at least one Web trade (column 1). Column 2 shows that $9,172 + 1,241 = 10,413$ participants traded again. 9,172 (88% of 10,412) made that second trade on the Web. 1,241 (12% of 10,413) made that second trade on the phone.




Figure 4. Average dollars per trade, 20-trading-day moving average. The dollar value of each day's trades on each channel is summed and divided by the number of participants who traded through that channel on that day. A moving average constructed using the most recent 20 trading days is displayed on the graphs.

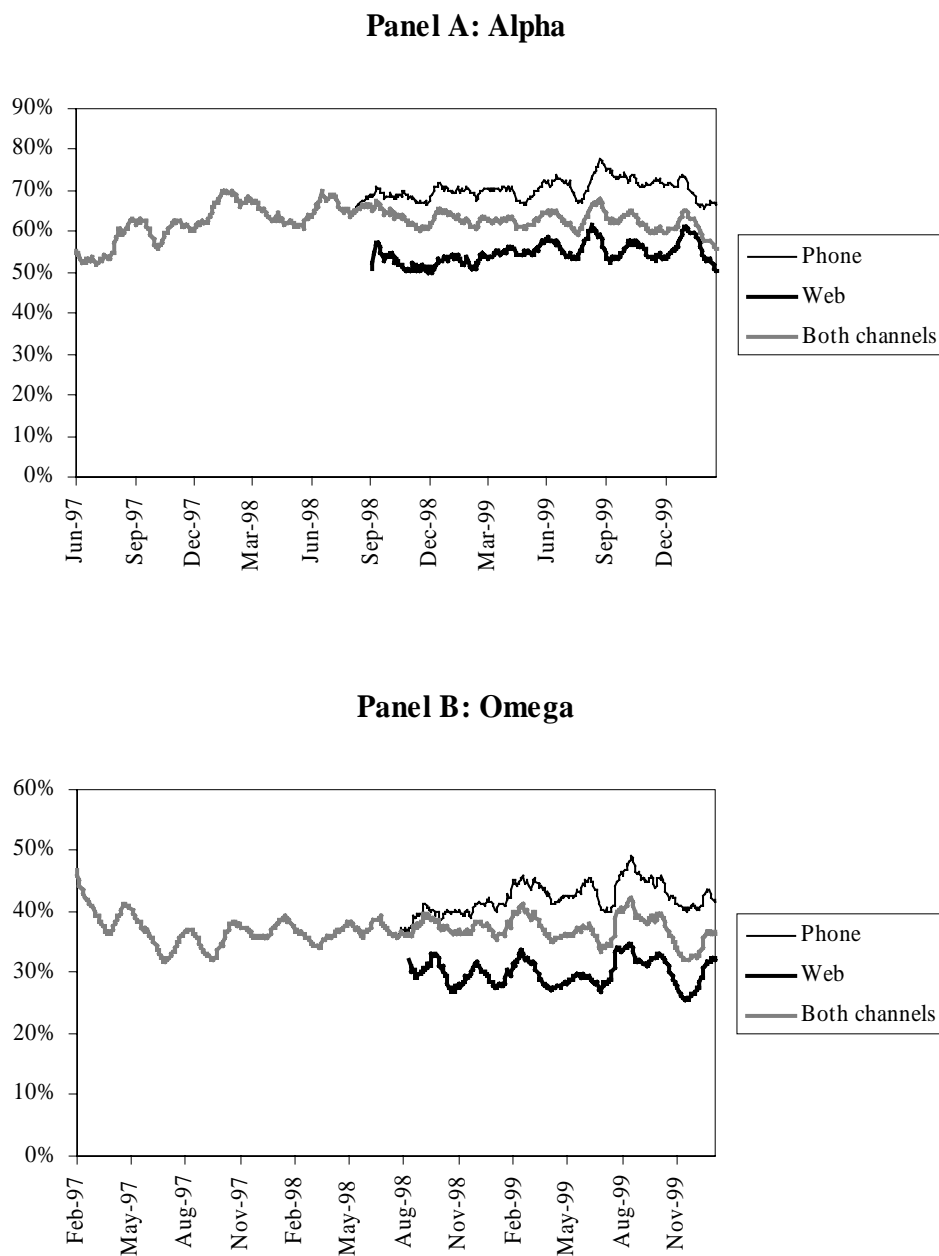


Figure 5. Average turnover per trader by channel, 20-trading-day moving averages. For each day we form a subsample of all of the participants who traded on that day, calculate their individual turnover rates (percent of day-end balance traded on that day), and then average these individual turnover rates by channel. A moving average constructed using the most recent 20 trading days is displayed on the graphs.

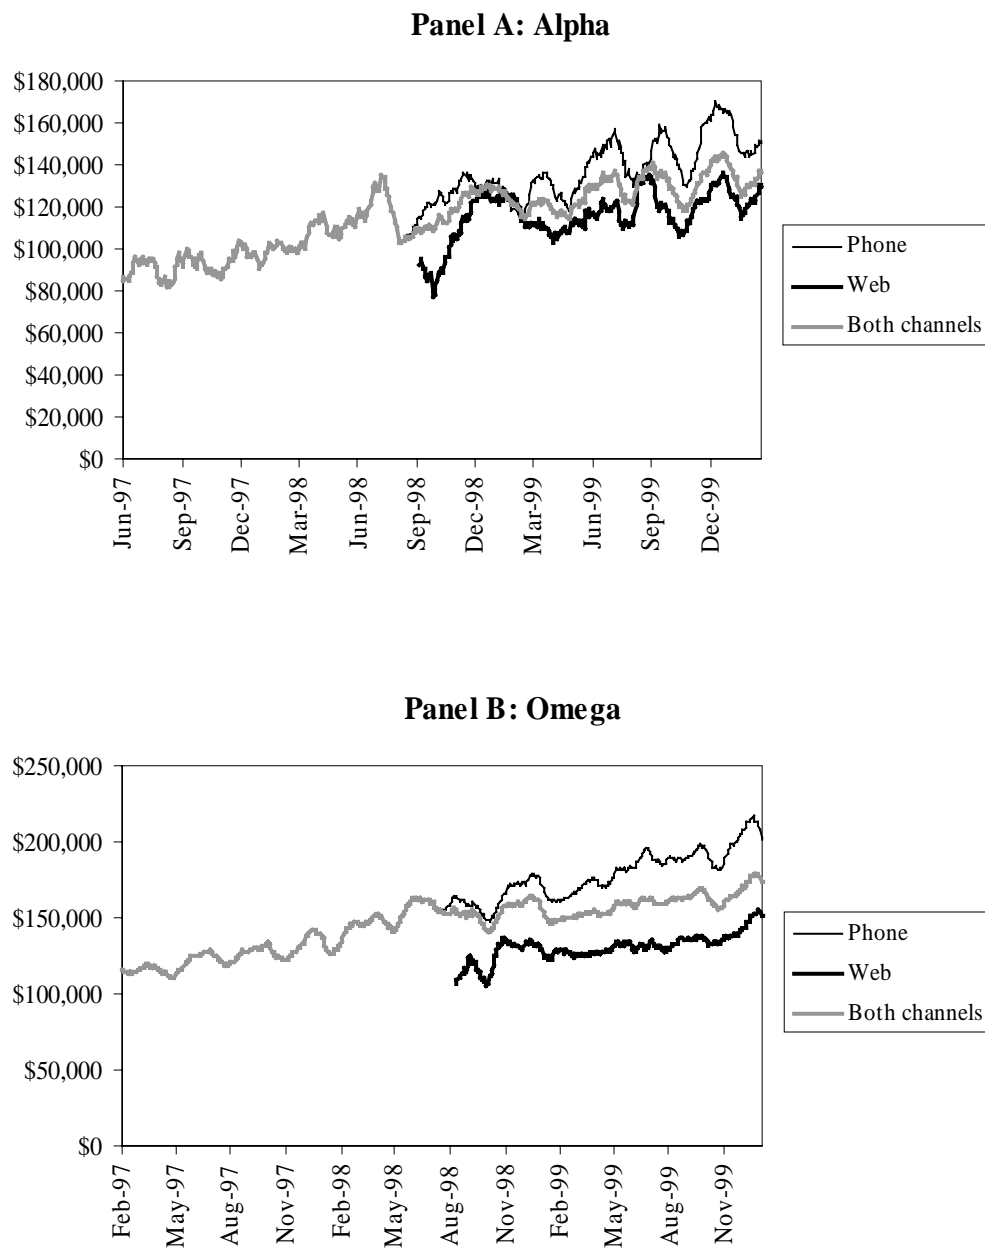


Figure 6. Average balance of traders, 20-trading-day moving average. The average day-end balances of participants who traded through each channel on each day are calculated. A moving average constructed using the most recent 20 trading days is displayed on the graphs.



Panel A: ABC**Panel B: Omega**

Figure 7. Percent of trades that are “reversed” trades, 20-trading-day moving average. “Reversed” trades are defined as trades that at least partially reverse a trade executed no more than five trading days prior, or trades that are at least partially reversed no more than five trading days hence. We divide the number of participants executing a reversed trade through each channel on each day by the total number of participants trading through that channel on that day. A moving average constructed using the most recent 20 trading days is displayed on the graphs.