

Do Investors Buy What They Know? Product Market Choices and Investment Decisions

Matti Keloharju

Aalto University School of Business and CEPR

Samuli Knüpfer

London Business School

Juhani Linnainmaa

University of Chicago Booth School of Business

This article shows that individuals' product market choices influence their investment decisions. Using microdata from the brokerage and automotive industries, we find a strong positive relation between customer relationship, ownership of a company, and size of the ownership stake. Investors are also more likely to purchase and less likely to sell shares of companies they frequent as customers. These effects are stronger for individuals with longer customer relationships. A merger-based natural experiment supports a causal interpretation of our results. We also find evidence of causality in the other direction: inheritances and gifts have an effect on individuals' patronage decisions. A setup in which customer-investors regard stocks as consumption goods, not just as investments, seems to best explain our results. (*JEL* G11, G24, D83)

If you like the store, chances are you'll love the stock.

(Lynch 1993, 152)

This article identifies a strong link between individuals' choices in the product market and those in the stock market. Analyzing daily panel data on stock

We would like to thank Amir Barnea (EFA discussant), John Cochrane, Lauren Cohen (AFA discussant), Gene Fama, Ron Goettler, Robin Greenwood, Bing Han, Dirk Jenter (NBER discussant), Brandon Julio, Joni Kokkonen, Mikko Leppämäki, Alexander Ljungqvist, Adair Morse, Toby Moskowitz, Mikko Niemenmaa, Peter Nyberg, Stavros Panageas, Lubos Pastor, Andrei Simonov, and Ilya Strebulaev, and seminar and conference participants at Aalto University, HEC Paris, Imperial College, London Business School, Norwegian School of Management, the University of Chicago, the University of Jyväskylä, the University of Lugano, the University of Nottingham, the University of Western Ontario, NBER Behavioral Economics working group meeting, the European Finance Association 2010 meeting, and the American Finance Association 2012 meeting for helpful comments and suggestions. Special thanks for detailed comments are due to two anonymous referees and the Editor, David Hirshleifer. We are also grateful to the OMX Nasdaq Foundation and OP-Pohjola Research Foundation for financial support and to Antti Lehtinen, Riikka Alatalo, and Rohit Kawathekar for superb research assistance. A previous draft of this article was entitled "From Customers to Shareholders: The Effect of Product Market Choices on Investment Decisions." Send correspondence to Juhani Linnainmaa, University of Chicago Booth School of Business, 5807 S Woodlawn Ave, Chicago, IL 60637; telephone: (773) 834 3176. E-mail: Juhani.Linnainmaa@chicagobooth.edu.

© The Author 2012. Published by Oxford University Press on behalf of The Society for Financial Studies. All rights reserved. For Permissions, please e-mail: journals.permissions@oup.com.

doi:10.1093/rfs/hhs090

Advance Access publication August 29, 2012

holdings, trades, and broker-customer relationships of Finnish investors, we find patrons of a given broker are more than twice as likely to invest in the corresponding broker stock, and have 13% larger ownership stakes conditional on investment. Our results also apply to the automotive industry, where owners of a given make are 51% to 83% more likely than owners of other makes to own shares in the respective car company. Taken together, our evidence suggests that a customer relationship is about as important to stock selection as the home bias.¹

Our analyses of the brokerage industry control for omitted company- and investor-level factors and take advantage of time-series variation in both patronage and ownership. The results also are robust to a wide array of alternative specifications and apply regardless of whether we study stock ownership, purchase, or sale decisions. Moreover, our key results hold for all sample brokers and different investor categories, including institutional investors.

We collect data on two types of shocks to measure causality between investment and consumption decisions. First, we study two events in which an insurance company merges with a brokerage firm; as a result of these mergers, customers of the (pre-merger) brokerage firm inadvertently become customers of the insurance company. We find the individuals experiencing this shock are more likely than the untreated individuals to purchase shares of the post-merger insurance company. Second, inheritances and gifts of broker stocks generate shocks to ownership status, allowing us to examine causality from investment decisions to consumption decisions. Although we find individuals who receive bank stocks are more likely to become these banks' customers, the causal effect from ownership to patronage is somewhat weaker than that from patronage to ownership.

The extraordinary detail in our data allows us to explore many potential explanations for these findings, including broker recommendations, home bias, awareness, private information, and beliefs on information precision and valuation. None of these explanations are fully satisfactory. The evidence against the explanation that brokers recommend their own stocks to customers comes from our analysis of the automotive industry; unlike stockbrokers, car dealers have little incentive to give stock recommendations to their customers. Yet the results from the automotive industry are similar to those from the brokerage industry. Home bias is also an unlikely explanation for our results, which remain unchanged even when we limit our sample to investors who live far away from the headquarters of any of the companies. Cross-sectional differences in awareness cannot easily explain why customers are less willing to *sell* their brokers' stocks; investors are, after all, aware of the stocks in their

¹ Grinblatt and Keloharju (2001a) use a subsample of our data set and document that households are 81% more likely to hold local stocks than domestic stocks at large. Hong et al. (2008) report U.S. households are 116% more likely to hold local stocks.

portfolios. We also find no evidence that investors have private information on the companies with which they do business, or that they behave as if they have private information: customers' investment performance in broker stocks and their behavior around earnings announcements is no different from that of other investors. Finally, an analysis of the prices at which investors place their limit orders suggests that customers value stocks similarly to other investors.

A theory in which individuals' product market choices alter the marginal utility they derive from different investments seems to best explain our results. A patron could derive utility from owning shares of a company she supports as a customer, just as a fan of a sports team could derive direct utility from her investment in the team.² Apart from explaining the effects patronage has on ownership and trades, we would expect such affection for a company to increase in the length of the customer relationship. The data support this expectation: individuals with a longer customer relationship with a broker are more likely to own and buy and less likely to sell the broker's stock. Investors who remain loyal to a broker thus display that same loyalty in their investment decisions.

We further evaluate this preference-based explanation by testing a prediction of Cao et al.'s (2011) model of familiarity. In this model, "familiarity-biased" investors are reluctant to invest in unfamiliar stocks and thus hold portfolios with fewer stocks than "unbiased" investors. We test this hypothesis by comparing diversification in customer-stockholders' (who reveal their preference for the familiar by investing in their own broker's stock) and non-customer-stockholders' portfolios. Consistent with the model's prediction, we find customer-stockholders invest in fewer stocks than non-customer stockholders.

Our article relates to four strands of literature. First, we contribute to the literature studying the effect of geographical and professional proximity on investing (see, e.g., Coval and Moskowitz 1999; Grinblatt and Keloharju 2001a; Benartzi 2001; Cohen 2009). The article most closely related to ours is that of Huberman (2001), who suggests that familiarity plays an important role in investors' portfolio choices. He finds that in almost every state, individuals overweight their local Regional Bell Operating Company relative to the six other RBOCs, in spite of the fact that all companies are equally accessible to investors. Huberman's findings are consistent with a portfolio bias induced by product market choices, but they can also reflect a preference for geographically proximate companies. The most important difference between his analyses and ours lies in our ability to attribute investment decisions directly to product market choices. Access to microdata allows us to observe heterogeneity in investment decisions even at the local level, and to exclude home bias as an

² The Green Bay Packers of the National Football League has been a publicly owned company since 1923 and has more than 100,000 shareholders today. Yet a dividend has never been paid, and the stock does not appreciate in value. Shares can only be sold back to the team, and then at a fraction of the original purchase price. See Sweet (2008).

explanation for the link between patronage and ownership. We also use broker–insurance company mergers and inheritances and gifts of bank stocks to explore the causal relation between investment and consumption decisions.

Second, our article is related to the literature analyzing the role of advertising and brand visibility on individual investor behavior (Grullon, Kanatas, and Weston 2004; Frieder and Subrahmanyam 2005; Billett, Jiang, and Rego 2010; Lou 2010). We focus on a different mechanism through which the product market could be related to investing: actual product market choices. Past product market choices are known to have a strong and persistent effect on consumer behavior, stronger and more persistent than advertising.³ Moreover, they can be observed at the level of an individual, and at high frequency, whereas advertising and brand visibility can be observed only in aggregate and generally at a much lower frequency.

Third, we add to the ongoing debate on whether geographically or professionally close investors enjoy an information advantage that allows them to beat the market. The literature can be divided into two broad camps: studies that find close investors outperform more distant investors, and those that do not.⁴ Our article falls into the second camp: although customers are likely to be more knowledgeable than non-customers about the company, we find they do not exhibit superior performance.

Finally, our analysis sheds light on the determinants of investors' behavioral biases. Although previous work has made significant progress in deciphering the underlying forces behind these biases (Cohen 2009 is worthy of special mention), we know of no other work that makes an effort to disentangle the relative contribution of preferences and beliefs to portfolio bias. Our results on loyalty, perceived information precision, private valuation, and diversification are consistent with the idea that preferences figure more importantly than beliefs in inducing customers to become and remain investors. As in Fama and French (2007), many customers appear to regard stocks as consumption goods, not just investments.

We organize the remainder of this article as follows. Section 1 describes the data and methods. Section 2 presents the results from our analyses. Section 3 evaluates possible explanations for our findings. Section 4 concludes.

1. Data and Methods

1.1 Data sources

We merge four data sets for our analysis.

³ See Bagwell (2007) for an excellent summary of the literature.

⁴ For example, Coval and Moskowitz (2001), Hau (2001), Choe, Kho, and Stulz (2005), Dvořák (2005), Ivković and Weisbenner (2005), Massa and Simonov (2006), and Baik, Kang, and Kim (2010) find that geographically close investors beat more distant investors, whereas Grinblatt and Keloharju (2000), Froot and Ramadorai (2008), and Seasholes and Zhu (2010) do not. Massa and Simonov (2006) and Døskeland and Hvide (2011) find no information edge for professionally close investors.

Finnish Central Securities Depository (FCSD) registry. This registry contains the daily portfolios and trades of all Finnish investors, both households and institutions, from January 1, 1995, through November 29, 2002. The electronic records we use are exact duplicates of the official certificates of ownership and trades and hence are reliable. Grinblatt and Keloharju (2000) report details on this data set, which includes information about each investor's attributes, date-stamped trades, holdings, and execution prices of publicly traded companies on the Helsinki Exchanges.

HEX transactions data. This data set is a record of all trades executed at the Helsinki Exchanges from January 1, 1995, through December 31, 2001. Each record indicates which brokerage firms the buyer and the seller used. We match these data against the FCSD registry to obtain information on investors' brokerage-firm choices as well as on individuals switching from one broker to another. Linnainmaa (2010) provides details of the match.

HEX microstructure data. This data set covers September 18, 1998, through October 23, 2001, and is a record of every order submitted to the fully electronic, consolidated limit order book of the Helsinki Exchanges. It tracks the life of each order submitted to the Exchanges, indicating when the order is executed, modified, or withdrawn. We first reconstruct second-by-second limit order books for all HEX-listed stocks, paying special attention to orders that are executed. Only executed orders contain certain markers that enable us to combine the limit order book with FCSD trading records to precisely identify the investor placing the executed order. Ultimately, we construct a data set that contains each investor's executed order type and what the limit order book (including unexecuted orders) looked like at any instant at, after, and prior to the moment of order execution.

Automobile Purchase Data. The Finnish Vehicle Administration (FVA) provided data on automobile purchases and ownership. The data set records the type of personal automobile each car-owning resident owned on June 10, 2002, for all purchases made prior to 2002, and the exact date of purchase. The data are comprehensive for residents in the provinces of Uusimaa and East Uusimaa, which contain Greater Helsinki and the most densely populated areas in Finland. Grinblatt, Keloharju, and Ikäheimo (2008) report details on this data set.

We match the automobile purchase data with stock-ownership records for Ford Finland (Ford Motor Company's Finnish subsidiary) and Metso (parent of the Finnish car maker Saab-Valmet, producer of more than 700,000 Saab cars between 1969 and 2003), both publicly listed companies. The automobile data are similar to the broker data in that we can identify changes in car- and stock-ownership statuses. These panel features in the data allow us to study whether an individual's decision to buy a Ford or a Saab car is associated with the same individual's decision to own or buy Ford or Metso shares.

1.2 Variables and methods

From the broker data set, we extract information on the key variables for all investors at the end of each year and on each day they trade. The unit of observation is an investor i -day t -broker j triplet, where the number of broker stocks listed on the stock exchange on day t determines the number of elements. In total, eight brokers were publicly listed for at least part of the sample period. Other, unlisted firms offering brokerage services do not enter our analyses, except when we record broker changes from listed brokers to unlisted ones, and vice versa.

We populate the investor-day-broker triplets with dummy variables that measure ownership (or purchases and sales) and patronage. The share-ownership dummy takes the value of one if an individual owns shares in a brokerage firm at the end of the year. We replace the ownership dummy with the log of the value of holdings when analyzing the determinants of the size of the holding in broker stock.

We measure changes in ownership status by identifying instances in which an investor either buys shares in a broker she does not currently hold or sells all her shares in a brokerage firm she currently holds. A share-purchase dummy takes the value of one in the first case, and a share-sale dummy takes the value of one in the second case.

We measure customer relationship with a dummy variable that takes the value of one if an investor has traded through a broker and zero otherwise. We restrict our sample to those investors who have traded at least once during the six months immediately preceding the ownership snapshot or the equity trade. This restriction reduces the likelihood that the recorded customer relationships are stale.

2. Results

2.1 Ownership decisions in the brokerage industry

Table 1, Panel A, reports descriptive statistics on the characteristics of the sample investors who own stock in at least one of the brokers. The median (mean) investor has a portfolio worth 4,000 euros (38,000 euros), of which broker stocks account for 11% (38%). The average investor is about 46 years old, and males account for almost two-thirds of the investors.

Panel B reports descriptive statistics on broker choice as a function of past broker choice. In a typical year, only 62% of traders use a broker they used in the previous year. This finding suggests that investors are surprisingly disloyal to their brokers. This disloyalty will be useful in later tests that use broker switches to estimate the effect of new and existing customer relationships on ownership.

Figure 1 reports descriptive analyses on individuals' patronage and ownership decisions. We aggregate the information on ownership and patronage to investor-year-broker (i, t, j) triplets by each year t in Panel A and by each

Table 1
Descriptive statistics on investors and their broker use

Panel A: Descriptive statistics						
Variable	<i>N</i>	Mean	Sd	25%	Median	75%
Financial wealth (1000 euros)	357,424	37.51	774.45	0.00	3.61	15.35
Holdings in broker stock (1000 euros)	357,424	2.16	23.26	0.00	0.00	0.34
Portfolio weight in broker stock (%)	357,424	37.66	43.49	0.00	10.97	100.00
Age	357,424	45.87	18.03	33.00	46.40	58.00
Female dummy (%)	357,424	36.88	48.25	0.00	0.00	100.00
Brokers used	357,424	0.77	0.62	0.00	1.00	1.00
Brokers used conditional on using at least one	228,387	1.21	0.44	1.00	1.00	1.25

Panel B: Broker use transition probabilities			
	Not current customer	Current customer	Probability of being current customer
Not past customer	784,278	99,574	11.27
Past customer	90,188	145,233	61.69
Prob. of being past customer	10.31	59.33	

Panel A reports descriptive statistics on the characteristics of the sample investors. Investor characteristics are averaged first across all years for each investor and then across all investors. An investor is selected in the sample if she has traded in any publicly listed stock in the 120 days prior to the end of that year. Financial wealth is measured as the sum of the value of all holdings in publicly listed stocks, holdings in broker stocks is the value of all holdings in publicly listed banks with a brokerage arm, and weight in broker stocks equals holdings in broker stocks divided by financial wealth. Brokers used calculates the number of brokers through which an investor has traded during a year. Panel B reports the transition probabilities between past and current customer states by counting the number of observations coming from investors who have had some trading activity in the current and previous years. Investors are divided into four categories: (i) those that continue to use the broker they used in the previous year, (ii) those that stop using the broker, (iii) those that continue not using the broker they did not use in the previous year, and (iv) those that start using the broker. The probability of being a current customer is calculated by dividing the number of current customer observations by the total number of observations in a row. The probability of being a past customer is the number of past customer observations divided by the total number of observations in a column. The total number of observations (1,119,273) is smaller than in Figure 1, as the sample is restricted to investors who have used some broker in the previous and current years.

broker j in Panel B. (Brokers 1 and 2 have the same number of observations in Panel B because (i) these brokers existed throughout the sample period and because (ii) an observation is generated for every investor-broker-year triplet.) This aggregation results in estimates of the probability that an individual owns shares in a brokerage firm at the end of the year conditional on whether the individual has been a customer of the same brokerage firm within the past six months.

Panel A shows that the probability of purchasing brokerage firm shares is higher for broker customers in all eight sample years. The difference is economically impressive: on average, a customer has an 18.2% likelihood of owning the broker stock, whereas the corresponding probability is only 4.6% for non-customers. The patronage-ownership relation also varies significantly over time. As new brokers get publicly listed, and thus enter our sample, the average patronage-ownership often changes significantly. For example, the drop in the relation from 1998 to 1999 arises from the listing of two brokers with a lower-than-average link between patronage and ownership; similarly, a listing of a

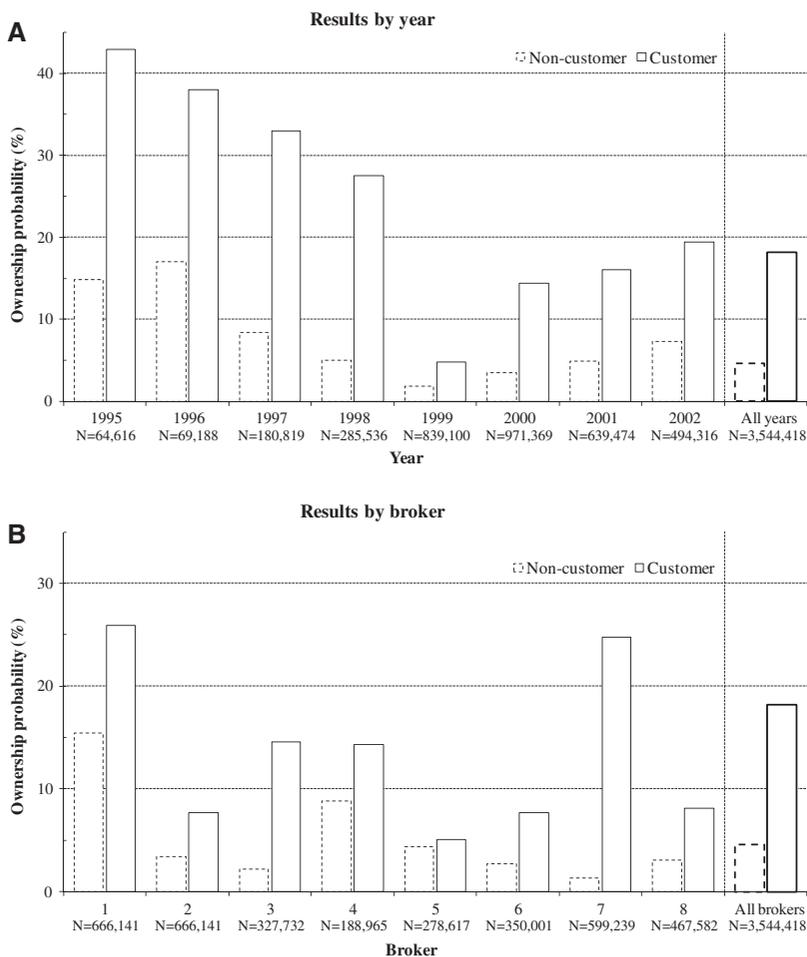


Figure 1

Descriptive statistics on ownership of brokerage firm stocks by customer status

Panels A and B report the probability that an investor owns shares in a brokerage firm. An investor is selected in the sample if she has traded in some stocks in the 120 days prior to the end of that year. An investor is classified as a customer of a brokerage firm if she has traded through it in the 120 days prior to the end of a year.

broker with an unusually strong patronage-ownership link explains the increase in the relation from 1999 to 2000.

Panel B shows that the probability of owning shares in a brokerage firm is significantly higher for broker customers for all eight brokers in the data. The strength of the customer-ownership relation, however, varies significantly across brokers. Whereas the customers of broker 7 are nineteen times more likely than non-customers to become owners, this ratio is only 1.7 for broker 1.

Table 2 builds on the descriptive analysis of Figure 1 by reporting a set of linear probability models. The first column runs a regression that examines

Table 2
Regressions of ownership

Dependent variable	Ownership dummy (0/1)			Ln (Value of holdings)		
	1	2	3	4	5	6
Patronage dummy	0.065 (5.15)	0.071 (6.49)	0.008 (3.24)	0.207 (3.19)	0.121 (3.66)	0.015 (0.49)
Ln (Wealth)	0.011 (2.34)	0.008 (1.75)	0.008 (1.06)	0.490 (10.31)	0.338 (6.42)	0.279 (3.42)
Female dummy	0.002 (0.36)		-0.059 (-3.62)			
Age	-0.0002 (-0.33)		0.008 (2.71)			
Age ²	0.00001 (1.15)		-0.00003 (-1.32)			
Fixed effects						
Investor	No	Yes	No	No	Yes	No
Investor-broker	No	No	Yes	No	No	Yes
Broker-year	Yes	Yes	No	Yes	Yes	No
Year	No	No	Yes	No	No	Yes
Reference prob. / value	0.069	0.069	0.069	7.910	7.910	7.910
Overall R ²	0.159	0.307	0.794	0.465	0.850	0.957
Number of investors	357,424	357,424	357,424	105,023	105,023	105,023
Number of observations	3,544,418	3,544,418	3,544,418	243,071	243,071	243,071

At the end of each sample year, the value of holdings in each brokerage firm stock is calculated for every investor. An investor is selected in the sample if she has traded in some stocks in the 120 days prior to the end of that year. Ownership dummy takes the value of one if an investor owns a brokerage firm stock, and zero otherwise. Ln (Value of holdings) is the value of an investor's holdings in a broker stock. Patronage dummy equals one if the investor has traded through a broker in the 120 days prior to the end of a year. Control variables include logged financial wealth measured as the sum of the value of all holdings in publicly listed stocks, a dummy for female investors, age/100, and age/100 squared. Columns 4 through 6 restrict the sample to cases in which an investor owns a broker stock. Columns 1 and 2 as well as 4 and 5 include broker-by-year fixed effects. Columns 2 and 5 drop the time-invariant investor characteristics and add investor fixed effects. Columns 3 and 6 replace the fixed effects with investor-broker and year fixed effects. The *t*-values, reported in parentheses below the coefficients, are robust for heteroscedasticity and are adjusted for clustering at the broker level.

the decision to own broker stock, controlling for broker-year fixed effects and observable investor characteristics, that is, financial wealth, gender, and age. The broker-year fixed effects control for any events that change the unconditional propensity to own a particular broker's stock in a given year, such as past returns on that stock. The investor characteristics address the possibility that customers are systematically different from other investors and that these differences make customers more likely to invest in any of the broker stocks.

The coefficient estimates and reference probabilities indicate that the probability of owning the stock of a given broker is 6.9% for non-customers and 13.4% (= 6.9% + 6.5%) for customers. Thus, same-firm customers are 1.94 times more likely than non-customers to purchase brokerage firm shares. This effect is highly significant (*t*-value = 5.15) even though we adjust the standard errors for broker-level clustering. This level of clustering assumes (conservatively) that we have just eight independent blocks of data. If we were to assume investor-by-investor independence, and hence clustered errors by investor, the *t*-value associated with the patronage dummy would increase to 104.9! We use broker-level clustering throughout the article to err on the side of caution and not oversell statistical significance.

The second column replaces the time-invariant investor characteristics with investor fixed effects. This specification captures any unobservable influences on ownership probability that go beyond gender and age. The fixed effects force the regression to identify the patronage-ownership link solely from within-broker and within-investor variation in the key variables. The results from this specification are even stronger, suggesting that a customer is more than twice as likely as a non-customer to own brokerage-firm shares. The effect of patronage dominates the effect of wealth and demographic variables on ownership propensity. In the investor fixed-effects regression, for example, a customer's ownership probability is 7.1% higher than that of a non-customer, whereas a one-standard-deviation change in logged financial wealth changes ownership probability only by 3.0% ($= 3.93 \times 0.008$).

Column 3 reports a more conservative regression specification with 2,179,278 investor-times-broker fixed effects. This specification uses information only from observations in which an investor either starts using a new broker or stops using an old broker. The slope on the customer dummy in this regression is 0.008 with a t -value of 3.24. The slope estimate is not comparable with the estimates reported in columns 1 and 2, because the specification excludes the information from investor-broker pairs in which the investor either remains the customer of the broker throughout the sample period or never uses the broker. We do not use this regression as our main specification, because the addition of investor-broker fixed effects absorbs important variation from the data and forces us to focus on observations in which the investor switches brokers (Angrist and Pischke 2009, p. 225–26). Nevertheless, it is comforting to know that our results survive even this more stringent identification strategy.

The association between patronage and ownership decisions is strong enough to influence the overall explanatory power of the regression. The first column shows that without the investor fixed effects, the R^2 is 0.159. The R^2 for the full regression, which includes the patronage dummy, is 0.307 in column 2. The model explains almost one-third of the variation in the ownership dummy, because of the investor and broker-year fixed effects. Omitting the patronage dummy from the full regression would decrease the explanatory power of the regression to 0.299, or by 2.6%. We consider the increase in R^2 nontrivial given that it comes solely from the patronage dummy's ability to explain within-investor variation in stock ownership.

Customer relationships correlate not only with the decision to invest in a broker stock, but also with the amount invested. The three rightmost columns in Table 2 replace the dependent variable by the log of the value of the holdings in the broker stock. Restricting the sample to investors with any holdings in broker stock, we gauge the extent to which the size of the holding differs across customers and non-customers. Given that our regressions control for the log of wealth, we can interpret the coefficient estimate on the patronage dummy as the effect a customer relationship has on portfolio weight. The coefficient estimate on the patronage dummy is 0.207 (t -value = 3.19) without investor

fixed effects (column 4) and 0.121 (t -value = 3.66) with these effects (column 5). The results are economically significant: conditional on investment, an average customer invests almost a quarter ($e^{0.207} - 1 = 23.0\%$) more than the average non-customer in the broker stock. The comparable effect is 12.9% in column 5's regression with investor fixed effects. Column 6 includes the same investor-times-broker fixed effects as the regression in column 3. The coefficient on the patronage dummy is statistically insignificant in this conservative specification.

2.2 Purchase and sale decisions in the brokerage industry

Table 3 analyzes investors' purchase and sale decisions of broker stocks using a set of linear probability models. The unit of observation is an investor-day-broker triplet (i, t, j) , and the vectors of explanatory and independent variables are populated with dummy variables measuring, first, which broker stock j investor i purchases on day t , and second, which broker j the investor has used in the 120 days preceding day t . Similarly to the ownership regressions, column 1 controls for observable investor characteristics, whereas column 2 replaces the time-invariant investor characteristics with investor fixed effects. Both regressions also control for past returns that are an important determinant of purchase and sell decisions.⁵

The coefficient on the patronage dummy in Table 3 measures the difference between customers and non-customers in the probability of purchasing brokerage-firm shares. The reference probability in column 2 indicates individuals who are not customers of a given broker have a 17.7% likelihood of investing in the broker's stock, but for customers this probability is 24.9% ($= 17.7\% + 7.2\%$). This 41% ($= 24.9\% / 17.7\% - 1$) increase in purchase probability is both economically and statistically (t -value = 5.36) significant.

The sell regressions reported in the two last columns of Table 3 follow the structure of the purchase regressions. The dependent variable is now a sale dummy, and the observations are generated for days on which an investor sells broker stocks. The sample comprises such investor-broker pairs in which the investor owns at least one share of the broker. These regressions assess how patronage of a broker changes one's likelihood of selling the shares of the broker. The slope estimate in column 4, which is statistically significant, indicates that customers are 3.6% less likely, or 0.96 times as likely, to part with their shares.

We also use an alternative methodology to study investors' sale decisions. In this robustness check, we estimate a Cox proportional hazards model to measure how an investor's patronage status influences the probability of a sale. We extract from the data all purchases of broker stocks for which an investor opens a new position, and then follow these positions until they are closed or the sample ends. We explain each position's survival with the patronage dummy,

⁵ See, for example, Grinblatt and Keloharju (2001b).

Table 3
Regressions of purchases and sales

Dependent variable	Purchase dummy (0/1)		Sale dummy (0/1)	
	1	2	3	4
Patronage dummy	0.043 (2.67)	0.072 (5.36)	-0.011 (-2.08)	-0.036 (-2.37)
Ln (Wealth)	-0.002 (-0.43)	-0.004 (-1.04)	-0.003 (-1.46)	-0.001 (-0.06)
Female dummy	0.007 (1.09)		-0.0004 (-0.37)	
Age	-0.009 (-0.07)		-0.028 (-0.50)	
Age ²	0.085 (0.91)		0.032 (0.60)	
Return [-1, -30]	-0.003 (-0.08)	-0.029 (-0.78)	0.076 (2.44)	0.245 (3.57)
Return [-31, -60]	0.029 (0.78)	0.029 (0.71)	0.047 (1.53)	0.090 (1.84)
Return [-61, -90]	-0.0004 (-0.01)	-0.016 (-0.87)	0.0229 (0.69)	0.051 (0.55)
Return [-91, -120]	0.002 (0.04)	0.003 (0.05)	-0.004 (-0.17)	-0.047 (-0.74)
Investor fixed effects	No	Yes	No	Yes
Broker fixed effects	Yes	Yes	Yes	Yes
Reference probability	0.177	0.177	0.863	0.863
Overall R ²	0.101	0.134	0.431	0.484
Number of investors	25,256	25,256	84,729	84,729
Number of observations	231,216	231,216	119,273	119,273

An observation is an investor-brokerage firm-day triplet. On each day an investor trades in a brokerage firm stock, the investor is assigned one observation for the stock she trades in, and observations for each brokerage firm stock she chooses not to trade in. Columns 1 and 2 restrict the sample to investor-days on which an investor purchases any of the brokerage firm stocks but does not own the purchased stock at the time of the purchase. The dependent variable takes the value of one for the purchased stock and zero for all the other broker stocks. Columns 3 and 4 restrict the sample to investor-days on which an investor sells all shares in any of the brokerage firm stocks and to brokerage firm stocks in which the investor has a position on the day. The dependent variable takes the value of one for the sold stock and zero for all other brokerage firm stocks in the investor's portfolio. Patronage dummy equals one if the investor has traded through a broker in the past 120 days. Control variables are logged financial wealth, measured as the sum of the value of all holdings in publicly listed stocks, a dummy for female investors, age/100, age/100 squared, and past returns on the broker stock, measured over four non-overlapping periods. Columns 3 and 4 also include unreported dummy variables for the number of broker stocks in an investor's portfolio, divided into dummy variables for one through four and more than five broker stocks. Columns 2 and 4 drop the time-invariant investor characteristics and add investor fixed effects. All the regressions include broker fixed effects. The *t*-values, reported in parentheses below coefficients, are robust for heteroscedasticity and are adjusted for clustering at the broker level.

logged portfolio value, gender, age, age squared, and the log-value of trades computed over the sample period preceding the purchase. The value of trades controls for the mechanical relation between trading activity and the average life span of a position in an investor's portfolio. The estimated coefficient on the patronage dummy variable is -0.149 (t -value = -2.37), suggesting that the hazard rate for customer-owners is 13.8% ($= 1 - e^{-0.149}$) lower than for non-customers-owners.⁶ These results are consistent with those reported in Table 3.

⁶ These results are reported in an Online Appendix, available online at <http://www.rfssfs.org/addenda.php>.

2.3 Patronage, ownership, and length of customer relationship

We have thus far measured patronage over a relatively short, six-month period. We now extend this time window to analyze how variation in the duration of a customer relationship changes the relation between patronage and investment decisions. We expect individuals who have been more loyal to a broker (i.e., by having a longer customer relationship) to show that same loyalty in their investment decisions. Specifically, we expect a freshly acquired customer to behave differently than an individual who has used the broker for a long time. Similarly, an investor who has not traded through a particular broker recently, but who used it in the past, is likely to differ from a customer who has been consistently loyal to the broker.

We extend the time window used to infer customer relations to twelve months and divide it into two six-month periods. We use these two windows to classify investors into four groups. The first group consists of investors who start using a new broker. These investors traded through a particular broker in months $m - 1$ through $m - 6$ (“recent past”) but not in months $m - 7$ through $m - 12$ (“distant past”). The other three groups consist of investors who (i) continue using an old broker from the distant-past period, (ii) stop using an old broker from the distant-past period, and (iii) did not use a particular broker in either six-month period.

Table 4 presents the decomposed results for the ownership probability, portfolio weight, purchase probability, and sale-probability regressions. These regressions replace the patronage dummy with dummy variables for the “continue using,” “start using,” and “stop using” investor groups. The omitted reference category consists of investors who did not use the broker in either six-month period.⁷ The sample sizes are smaller than in Tables 2 and 3 because we also require an investor to have traded through a particular broker in the distant-past period. This restriction ensures that the customer relationship is not stale in either period.

Column 1 shows the results of the ownership regressions. The reference probability of owning shares of a broker is 8.6% for the omitted group of investors (i.e., for those who did not use the broker in either period). The coefficient on the first row of the table shows that the ownership probability is more than twice as high (19.0%) for investors who used the broker in both periods. In other words, individuals who have remained loyal to a broker during the past year are more than twice as likely to own the stock of the broker in comparison to individuals who have not been customers in either six-month period. Investors who just started or stopped using a broker have ownership probabilities of 13.3% and 14.2%, respectively.

These estimates imply the following temporal pattern in ownership probabilities. Suppose an individual starts using a broker she has not used

⁷ The reference category accounts for 72.6% of the observations. The continuing, starting, and stopping investors account for 8.4%, 3.2%, and 15.8% of the sample, respectively.

Table 4
Decomposing the patronage effect

Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Continue using	0.104 (5.02)	0.178 (4.39)	0.098 (4.96)	-0.072 (-4.21)
Start using	0.047 (4.23)	0.109 (3.15)	0.036 (4.04)	-0.032 (-1.54)
Stop using	0.056 (5.00)	0.134 (5.29)	0.031 (6.28)	-0.043 (-1.92)
Ln (Wealth)	0.008 (2.18)	0.360 (7.77)	-0.003 (-0.91)	-0.008 (-0.89)
Return [-1, -30]			-0.043 (-1.02)	0.227 (3.70)
Return [-31, -60]			0.010 (0.21)	0.087 (1.50)
Return [-61, -90]			-0.015 (-0.85)	0.033 (0.55)
Return [-91, -120]			0.005 (0.09)	-0.022 (-0.43)
Investor fixed effects	Yes	Yes	Yes	Yes
Broker-year fixed effects	Yes	Yes	No	No
Broker fixed effects	No	No	Yes	Yes
Reference probability/value	0.086	8.018	0.173	0.762
Overall R^2	0.310	0.836	0.133	0.433
Number of investors	162,805	62,215	16,796	24,342
Number of observations	1,723,987	148,778	170,749	48,212

Columns 1 and 2 correspond to regressions in columns 2 and 5 in Table 2. Columns 3 and 4 report regressions similar to columns 2 and 4 in Table 3. All the regressions restrict the samples used in Tables 2 and 3 further by discarding investors who have not traded in 121 through 250 days preceding the day of the analysis. This restriction enables the calculation of a customer relationship from both the recent period of six months (recent) and the more distant period of six to twelve months (past). *Continue using* takes the value of one if an investor has traded recently through a broker she has also used in the past. *Start using* takes the value of one if an investor has recently started trading through a broker not used in the past. *Stop using* takes the value of one if an investor does not currently trade through a broker used in the past. The coefficient estimates of these variables tell the increase in ownership probabilities and portfolio weights, and purchase and sale probabilities, compared with the reference category of not having used a broker recently or in the past. Control variables and estimation methods are identical to those in Tables 2 and 3. The t -values, reported in parentheses below the coefficients, are robust for heteroscedasticity and are adjusted for clustering at the broker level.

before. During the first six months, the ownership probability increases from the baseline of 8.6% to 13.3%. If she stays with the broker for a full year, this probability increases to its peak of 19.0%. Now suppose the individual stops using the broker. The ownership probability drops to 14.2% during the first six months after the breakup. The probability then falls back to the baseline of 8.6% after the first year. This pattern is consistent with the hypothesis that the length of the customer relationship amplifies the relation between patronage and ownership. The estimates in column 2, which replaces the ownership dummy with the value of holdings, exhibit the same relation between the duration of the customer relationship and investment decisions as column 1.

The patterns are similar in the purchase and sale regressions in columns 3 and 4. For example, the probability of purchasing shares is 23% lower for new broker customers than for those who continue as customers of the same broker $((0.173 + 0.036)/(0.173 + 0.098) = 0.771)$. These results suggest that

new customers are less likely than customers who have been consistently loyal to their broker to purchase and retain shares of their new broker. The same applies to individuals who leave their broker for another.

2.4 Merger-based natural experiment

Although our analyses show that individuals' earlier patronage decisions predict later investment decisions, they do not rule out the possibility that some unobserved variable could drive both the patronage and investment decisions at different lags. We study two mergers of a publicly listed insurance company (with no brokerage business) with a publicly listed investment bank (with prominent brokerage business) to analyze causality.⁸ These mergers allow us to investigate the impact a plausibly exogenous shock to customer relationship has on portfolio choice. We test whether the individuals who are treated by the merger (to become customers of the insurance company) are more likely to buy shares of the post-merger insurance company than the untreated individuals.

The details of the natural experiment are as follows. The merger between insurance company Sampo and investment bank Mandatum was announced on December 4, 2000, and the merger between insurance company Pohjola and investment bank Conventum was announced on June 13, 2001. In both mergers, the insurance companies offered the banks' shareholders payment in either cash or shares. A small minority of the shareholders of both banks declined to tender their shares. As a consequence, Conventum retained its separate listing on the stock exchange for 259 calendar days after the announcement of the merger, and Mandatum for 361 calendar days. The names of the brokers did not change in these mergers and remained unchanged for the rest of our sample period.

The mergers were well publicized in the national media and created a transitory shock to the visibility of the companies involved. This shock itself may thus have increased share purchases of the insurance company (Barber and Odean 2008). We use a difference-in-difference design to isolate the pure effect of patronage from any market-wide developments surrounding the mergers. This design controls for the behavior of broker customers and other individuals before and after the merger. The "before" period covers the year (250 trading days) before the merger announcement, whereas the "after" period covers the year after the announcement.⁹ Although the insurance company's newly acquired customers could be less loyal than its original customers,¹⁰ such a difference in loyalty would only temper the effect of the merger treatment and would not matter for identification.

⁸ Several other studies, such as those by Garmaise and Moskowitz (2006) and Hong and Kacperczyk (2010), employ merger data to identify and measure causal effects.

⁹ For Conventum, the "after" period extends to the end of the company's listing as a separate company, that is, 259 calendar days.

¹⁰ See, for example, Parthasarathy (2007), Holliday (1995), and Morrall (1996).

We form our sample by imposing the following restrictions. First, we exclude investors who owned shares in the broker at the time of the merger. This exclusion ensures that broker shareholders investing the cash received in the tender offer in the shares of the insurance company do not affect our results. It also ensures that conversion of the broker shares to insurance company shares does not influence our results. Second, we omit investors who owned shares in the insurance company at the beginning of the before period. This restriction corresponds to the criteria we use in our previous share-purchase regressions to identify true changes in ownership. Finally, we narrow down the sample to investors who traded at least once during the one-year period immediately preceding the merger. We use these trades to determine the brokerage-customer relationships of the investors.

The unit of observation in our regression is an investor-merger-period (i, j, t) triplet. Each investor has four observations, representing two periods and two mergers. The first merger-specific observation corresponds to the before period, and the second observation corresponds to the after period. The dependent variable takes the value of one if investor i makes positive net purchases in the shares of the insurance company in merger j in period t ; that is, the number of shares bought exceeds the number of shares sold. A treatment dummy takes the value of one for investor i in merger j if investor i has traded through the broker involved in merger j during the one-year period immediately preceding the merger announcement. The regression also includes a dummy variable for the after period, an interaction variable for the product of the treatment and after dummies, and—to control for differences in observable characteristics of customers and other investors—age, gender, and past trading activity.

The after dummy measures the overall change in the probability that an individual purchases shares of the post-merger insurance company. This estimate is measured relative to the pre-merger share-purchase probability, and it captures, for example, the effect of the merger on the overall visibility of the insurance company. The interaction term measures differences in this post-merger purchase probability between customers and non-customers of the brokerage firm. This interaction term evaluates the effect of patronage on share ownership.

The estimates in Table 5 support the hypothesis that patronage increases the likelihood of ownership. The negative coefficient estimate for the after-dummy variable (-0.008), although not statistically significant, suggests that the merger decreased the likelihood that untreated individuals (i.e., those who are not customers of the merged broker) purchase shares of the insurance company. The positive coefficient for the treatment dummy, 0.019 (t -value = 4.36), indicates that customers of the broker were more likely than non-customers to purchase insurance-company shares before the merger. The interaction term, which measures the influence of the exogenous change in the ownership-patronage on individuals' portfolio choices, takes the value of 0.015 (t -value = 2.59). This positive estimate shows that unlike the untreated individuals, the treated individuals became more likely to purchase shares in

Table 5
Natural experiment with mergers

Dependent variable	Purchase dummy (0/1)
After	-0.008 (-0.55)
Treatment	0.019 (4.36)
After × Treatment	0.015 (2.59)
Ln (Trading activity)	0.011 (5.98)
Female	-0.003 (-4.17)
Age	0.00006 (5.17)
Age ²	-0.0000004 (-5.29)
Reference probability	0.018
Overall R^2	0.027
Number of observations	923,878

Two mergers of a publicly listed insurance company (without a brokerage arm) and a publicly listed brokerage firm provide a setting in which a link between ownership and patronage is exogenously established. “Before” and “After” refer to 250-day periods before and after the merger announcement (Sampo-Mandatum on December 4, 2000, and Pohjola-Conventum on June 13, 2001). The sample includes investors who have had some trading activity but do not own shares in either of the companies involved in the merger 250 days prior to the merger announcement. The dependent variable takes the value of one if an investor has made positive net purchases in the insurance company stock. Treatment status is defined as having traded through the brokerage firm in the 250 days preceding the merger. Control variables are logged trading activity, calculated as the combined value of all trades in the 250-day period prior to the merger, age/100, age/100 squared, and a dummy for female investors. Standard errors are adjusted for clustering at the merger level.

the post-merger insurance company. The incremental effect of patronage on the purchase probability is economically significant. Given that the unconditional purchase probability in this sample is 0.018, the estimate of 0.015 suggests that a shock to patronage status almost doubles the purchase probability over the merger event.¹¹

This natural experiment offers independent support for the hypothesis that individuals are more likely to invest in companies they frequent as customers. More importantly, it strengthens the argument that our results go beyond documenting a simple correlation between patronage and ownership.

2.5 Does investment predict patronage?

Tables 6 and 7 analyze the possibility that individuals’ investment decisions predict their patronage decisions—that is, that a causal effect also occurs in

¹¹ One can view the empirical estimates of the influence of patronage on ownership as a composite measure of all the factors that change when an individual becomes a customer. These factors include individuals’ own investments in information acquisition and firms’ marketing efforts toward their customers. This interpretation also applies to the natural experiment: the treated individuals became more likely marketing targets for the insurance company after the merger. We think these considerations are an advantage, as our estimates represent the true effect an individual would experience when becoming a customer of a company.

Table 6
Predicting patronage with investment

Dependent variable	Patronage dummy (0/1)	Ln (Value of trading activity through broker)	Start using broker dummy (0/1)	Stop using broker dummy (0/1)
Specification	1	2	3	4
Ownership dummy	0.125 (4.68)	0.445 (2.16)	0.069 (1.82)	-0.238 (-4.52)
Ln (Wealth)	0.011 (2.83)	0.049 (3.30)	-0.001 (-1.72)	0.004 (3.18)
Return [-1, -30]			0.003 (0.46)	-0.020 (-1.41)
Return [-31, -60]			0.0001 (0.01)	-0.032 (-1.53)
Return [-61, -90]			0.006 (1.21)	-0.0192 (-1.82)
Return [-91, -120]			0.0006 (0.22)	-0.006 (-0.55)
Investor fixed effects	Yes	Yes	Yes	Yes
Broker-year fixed effects	Yes	Yes	No	No
Broker fixed effects	No	No	Yes	Yes
Reference probability / value	0.197	8.649	0.051	0.551
Overall R ²	0.326	0.538	0.148	0.241
Number of investors	113,035	75,233	217,810	186,959
Number of observations	1,119,273	220,182	7,280,055	3,883,219

This table analyzes whether investment predicts patronage. Columns 1 and 2 are based on annual snapshots of ownership and patronage. Ownership and control variables are measured at the end of each year, and patronage is measured from the six months following the end of each year. Investors who do not trade in this period are dropped from the analysis. Columns 3 and 4 analyze the decisions to start and stop using a broker on a day when an investor trades through at least one of the sample brokers. Column 3 restricts the sample to cases where an investor has not traded through a particular broker in the preceding six months, and Column 4 to cases where an investor has traded through the broker. The *t*-values, reported in parentheses below the coefficients, are robust for heteroscedasticity and are adjusted for clustering at the broker level.

the opposite direction. Table 6 uses a framework similar to Tables 2 and 3 to measure how ownership of a broker stock today predicts patronage of the broker in the future. The dependent variable is a patronage dummy in the first column and the log-value of trades in the second column. (The sample in the second column is restricted to investors who trade at least once through the broker; that is, the value of trades exceeds zero.) The other two regressions explain individuals' decisions to either start using a new broker (column 3) or stop using their current broker (column 4).

The results in Table 6 are consistent with the idea that the patronage-ownership effect also runs in the opposite direction. Owners of a brokerage firm's shares (i) are more likely to be customers of the brokerage firm the next year, (ii) trade more through the broker whose shares they own, conditional on trading through this broker at all, (iii) are more likely to start using the broker whose shares they own, and (iv) are less likely to stop using the broker whose shares they own. Except for the start-using regression, the coefficients on the ownership dummy are statistically significant at the 5% level.

The economic significance of the results is comparable to what we observe for the causality in the opposite direction. The first column's regression suggests that owners of a broker stock are 63.4% more likely to be customers of the same

Table 7
Stock gifts, inheritances, and patronage

Dependent variable	Patronage dummy (0/1)		
	Inheritances	Gifts	Both
Specification	1	2	3
After	0.016 (1.20)	0.009 (0.55)	0.013 (1.22)
Treatment	0.106 (1.30)	0.132 (2.47)	0.115 (1.97)
After × Treatment	0.032 (1.20)	0.045 (1.95)	0.037 (2.18)
Ln (Portfolio value)	0.018 (4.45)	0.016 (3.19)	0.018 (4.34)
Female	-0.049 (-3.13)	-0.044 (-1.55)	-0.047 (-3.00)
Age	0.002 (0.53)	0.0003 (0.10)	0.0028 (2.48)
Age ²	-0.00003 (-1.00)	-0.000003 (-0.08)	-0.00004 (-3.00)
Reference probability	0.218	0.183	0.205
Overall R ²	0.310	0.258	0.290
Number of observations	4,312	2,502	6,814

This table analyzes how receiving broker stocks as gifts or inheritances influences patronage. For each gift (column 1) and inheritance (column 2) in broker stock, we compute patronage in the broker 250 days before and after the day of receiving the gift and inheritance, separately for each broker. We exclude trades in broker stocks in the computation of patronage. *Treatment* takes the value of one for the broker whose stock was received as a gift or inheritance and zero for all the other broker stocks. *After* takes the value of one for the 250-day period after the gift or inheritance and zero for the preceding 250-day period. The sample includes investors who (i) do not hold any broker stock on the day of receiving the gift or inheritance, and (ii) have some trading activity in the before and after periods. Control variables are logged portfolio value, age, age squared, and a dummy for female investors. The regressions also include dummies for each broker. Standard errors are adjusted for clustering at the broker level.

broker during the next six-month period. The second column finds that their trading volume is 44.5% higher than that of non-owners: share ownership thus helps bring not only more customers but also bigger business to the company. The rightmost column suggests that patron-owners are 43.2% less likely to switch to another broker. This finding is consistent with the idea that stock ownership makes customers more loyal to the company.

Table 7 examines causality from ownership to patronage by constructing a sample of events in which individuals, without any action on their part, receive bank shares. We study two kinds of events: inheritances (column 1) and gifts of bank shares (column 2).¹² We use the difference-in-difference methodology introduced in the merger experiment to study whether individuals who receive shares of a particular broker are more likely to become customers of that broker. We use the brokers whose shares the investor does not receive as the benchmark group. For example, if an investor receives broker 1's shares, the methodology measures the incremental probability of becoming a patron of broker 1 over the other brokers. In measuring post-event patronage, we ignore trades that take

¹² The exclusion restriction in this analysis is that inheritance and gift events are uncorrelated with the other unobserved determinants of the decision to become a firm customer.

place in the inherited or gifted shares. This restriction ensures that our sample investors use their brokers for more than just selling the shares they received using a newly opened brokerage-firm account.

The main coefficient of interest in Table 7's regression is that associated with the after-treatment interaction. This coefficient measures the incremental change in the probability of patronage for individuals who inherit shares of a particular broker. The coefficient estimate is 0.032 (t -value = 1.20) for inherited shares, 0.045 (t -value = 1.95) for shares received as gifts, and 0.037 (t -value = 2.18) for the pooled sample in column 3. Restricting our sample to these treatments leaves us with far fewer observations than in Table 5's test of the opposite causality. Although sample size can explain in part why the t -values are smaller in Table 7 than they are in Table 5, the point estimates suggest that the effect from investment to patronage is also smaller than the corresponding effect from patronage to investment. In the pooled sample, for example, a receipt of broker shares increases the patronage probability by 18% (= $0.037/0.205$). By contrast, the merger-based experiment suggests that for the causality in the opposite direction, patronage almost doubles the probability of purchasing broker shares. Thus, although Table 7 suggests that the causality also runs from shareholdings to consumption decisions, the effect is not as strong as in the other direction.

2.6 Evidence from the automotive industry

Table 8 uses data from the automotive industry to revisit the purchase decisions with another, independent data set. These data contain records on whether individuals own Ford or Saab cars and Ford or Metso shares. We use these data to examine whether the ownership of a Ford (Saab) car alters the probability of owning or purchasing Ford (Metso) shares. In the analysis of Metso, we exclude all employees who participated in the company's employee stock offering in 1996. Similar to our analysis of brokerage choices, we use time-series data to measure both how consumption decisions influence investment decisions and vice versa.

The descriptive analysis in Table 8, Panel A, suggests that the patronage-investment relation in the automotive industry is similar to that in the brokerage industry. The two leftmost columns of Panel A show that owners of Ford (Saab) cars are more likely than non-Ford-car (non-Saab-car) owners to own shares in Ford (Metso). The difference in ownership probabilities is 0.16% (2.54%), suggesting that Ford-car (Saab-car) owners are 1.84 (1.75) times more likely than non-owners to own Ford (Metso) shares. The two rightmost columns of Panel A repeat the analysis for open-market purchases of Ford or Metso shares, and we find qualitatively similar results. All differences are statistically significant at the 1% level.

Panel B measures the effect of owning car company shares on the probability of purchasing a car from the same company. The probability that a Ford shareholder purchases a Ford car is 14.1%, whereas this probability is 8.6%

Table 8
Car ownership and car-stock ownership and purchases

Car owner	Panel A: Stock ownership and purchases				Panel B: Car purchases			
	Ownership probability (%)		Stock purchase probability (%)		Stock owner		Car purchase probability (%)	
	Ford	Metso	Ford	Metso	Ford	Metso	Ford	Saab
Yes	0.34	5.92	0.30	6.86	Yes		14.13	5.83
No	0.19	3.38	0.07	4.84	No		8.60	2.91
Difference	0.16	2.54	0.23	2.02	Difference		5.53	2.92
z-value	2.92	7.88	3.57	3.02	z-value		1.89	6.90
N	89,204	89,151	26,743	26,714	N		62,986	62,965
	Panel C: Stock ownership and purchase regressions							
Dependent variable	Ownership dummy (0/1)				Share purchase dummy (0/1)			
	Ford	Metso	Ford	Metso	Ford	Metso	Ford	Saab
Ford-car owner	0.165 (2.42)		0.233 (2.02)				0.061 (1.67)	
Saab-car owner		1.789 (4.69)		1.586 (2.06)				0.022 (3.76)
Ln (Wealth)	0.206 (9.42)	4.073 (50.88)	0.014 (3.36)	0.461 (14.27)			-0.002 (-2.16)	0.002 (4.46)
Female dummy	-0.044 (-1.48)	0.166 (1.36)	-0.038 (-1.03)	-1.848 (-6.32)			0.010 (4.09)	-0.018 (-14.26)
Age	-1.711 (-2.2)	-29,220 (-10.05)	-0.486 (-0.48)	-0.492 (-0.08)			-0.224 (-4.54)	0.032 (1.04)
Age ²	1.728 (2.21)	31,020 (10.73)	0.287 (0.26)	3.339 (0.50)			0.184 (3.56)	-0.020 (-0.60)
Ref. prob. (%)	0.198	3.477	0.090	4.923			0.086	0.030
R ²	0.005	0.106	0.001	0.011			0.001	0.004
N	89,204	89,151	26,743	26,714			62,986	62,965

Panel A reports the association between Ford-car and Ford-stock ownership and Saab-car and Metso-stock (parent of Saab-Valmet, a major producer of Saab cars) ownership and purchases. Ownership of cars is measured from the period up to the end of 1999, ownership of stocks is measured in the beginning of 2000, and purchase of stocks from January 2000 to November 2002. Panel B reports the same association but changes the timing so that car purchases from January 2000 to June 2002 are explained by stock ownership at the end of 1999. Panel C reports regressions of the stock ownership and purchases that control for financial wealth, gender, age, and age squared. The regressions in Panel D change the timing of the variables and explain (future) car purchases with (past) stock ownership. The sample includes residents of Uusimaa and East Uusimaa provinces and includes the most recent purchase of a given car as of June 10, 2002. The *t*-values, reported in parentheses below the coefficients, are robust for heteroscedasticity. The table multiplies the original regression coefficients by 100. The Metso sample excludes individuals who participated in Metso's employee stock offering in 1996.

for individuals who do not own Ford shares. For the Metso-Saab pair, these probabilities are 5.8% and 2.9%, respectively, suggesting that portfolio holdings predict car-purchase decisions in the time series.

The two leftmost columns of Panel C report estimates of a model that examines the relationship between Ford- or Saab-car ownership and stock ownership.¹³ The regression estimates tell a story similar to the descriptive statistics in Panel A: Ford-car (Saab-car) owners are more likely than owners of other cars to purchase shares in Ford (Metso). Ownership of a Ford (Saab) car increases the share-ownership probability by 0.17% (1.79%), suggesting that Ford-car (Saab-car) owners are 1.83 (1.51) times more likely than others to own Ford (Metso) shares. The two rightmost columns repeat the analysis for purchases of Ford and Metso shares, and find qualitatively similar results. All regression results are statistically significant at the 5% level.¹⁴

We also repeat the regressions in Table 8 using a placebo specification that adds the “wrong” car dummy in every regression. For example, in an alternative to Panel C’s column 1, we regress the Ford-stock ownership dummy against a Ford-car ownership dummy and a Saab-car ownership dummy. This modification allows us to evaluate the possibility that the buyers of Ford cars are different from the buyers of other makes in some omitted dimension that also makes them more likely to buy Ford stock. None of the “wrong” car dummies are significant at conventional levels in any of the four alternative regressions. The “correct” car dummies, however, retain their significance in all regressions.

Table 8, Panel D, regresses the car-purchase decision against the stockowner status. The estimates are positive for both car companies. Metso stock ownership is associated with a $0.022/0.030 = 73\%$ increase ($t = 3.76$) in the probability of purchasing a Saab car. The corresponding estimate for Ford is of similar economic significance but statistically significant only at the 10% level.

The results from the automotive industry are invaluable for two reasons. First, they offer independent evidence on the consumption-investment relation from a second industry. A comparison of Table 8 with Tables 2 and 3 suggests that the patronage-ownership relation is about as strong in the brokerage and automotive industries, even though these industries are quite different.¹⁵ Second, the possibility of an omitted common shock is more remote in the

¹³ Unlike in the brokerage-firm data, the car regressions do not have enough within-investor variation for us to be able to control for individual-level fixed effects.

¹⁴ We also estimate Table 8, Panel B, using a logit model instead of a linear probability model. Although not reported formally, the t -values on the Ford- and Saab-car owner dummies in the four logit regressions are 3.3, 5.0, 3.3, and 2.2, respectively, and thus exceed those of the linear probability model. The economic significance of the car-ownership dummy is about the same in both models.

¹⁵ Issues such as taste and appeal may play a greater role in the car-purchase decision than in broker choice. Such an emotional response could conceivably be stronger for luxury or niche brands. Of the two car makes in our sample, at least Ford can be characterized as a mass-market brand. In the case of Saab, the fact that the company assembling the car (Metso) has a different name than the brand that is being sold to customers (Saab) weakens the possible emotional link between patronage and ownership.

car-firm than in the brokerage-firm data. The median waiting time between the consumption decision (car purchase) and the investment decision (stock purchase) is 1,493 calendar days, and the mean is over five years.¹⁶

2.7 Cross-sectional patterns

Do institutional investors, including mutual fund managers, exhibit patronage-ownership patterns similar to the patterns we document for individual investors? Individuals ultimately make the decisions for institutional investors, and these individuals may be subject to the same biases as individual investors trading on their own accounts. For example, Grinblatt and Keloharju (2001b) and Frazzini (2006) report that institutional investors sell winners too early and hold on to losers for too long (“the disposition effect”).

Table 9, Panel A, estimates the ownership, log-value of holdings, purchase, and sale regressions in a sample that includes both institutions and individual investors. The regressions interact the patronage dummy with an institutional-investor dummy to measure marginal differences between the investor types. The links between patronage and investment decisions are roughly of the same magnitude for institutions as they are for individual investors. The interaction term is statistically insignificant in all but one regression. The term is significant and positive in the log-value of holdings regression, indicating that the relation between patronage and investing is stronger for institutions in this dimension.

Table 9, Panel B, estimates the same regressions for individual investors but interacts the patronage dummy variable with five investor-specific variables and one broker-specific variable. These regressions measure the extent to which various investor characteristics either amplify or moderate the relation between patronage and investing. The investor-specific variables measure age, wealth, trading activity, and the use of multiple brokers.¹⁷ The “first-purchase” dummy variable is unique to the purchase regression. This variable (which takes the value of one for an investor’s first broker-stock purchase) allows the patronage-investment sensitivity to differ between investors’ “first” and “other” purchases. The broker-specific variable is a retail-broker dummy variable taking the value of one for the three retail brokers in the sample.

The estimated interaction terms indicate that the relation between patronage and investing varies significantly across investors. In the ownership regression (column 1), this relation is 14% weaker ($= -0.025/0.174$) for investors with above-median wealth, 20% weaker for those trading more actively, 25% weaker for investors using multiple brokers, and 43% weaker for customers of retail brokers. Moreover, it is 27% stronger for investors of above-median age.

¹⁶ Because of a high car tax, Finnish households tend to use their cars longer than households in many other industrialized countries.

¹⁷ Baumann, Burton, and Elliott (2005), Lambert-Pandraud, Laurent, and Lapersonne (2005), and Patterson (2007) show that older individuals tend to be more loyal customers. Vissing-Jørgensen (2003) suggests that wealth may moderate various trading patterns such as the disposition effect.

Table 9
Cross-sectional patterns

Panel A: Institutions				
Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Patronage dummy	0.072 (6.42)	0.123 (3.69)	0.068 (5.16)	0.016 (1.40)
× Institutional investor	-0.043 (-1.75)	0.272 (3.48)	0.037 (0.97)	0.003 (0.13)
Ln (Wealth)	0.008 (1.76)	0.338 (6.62)	-0.004 (-1.07)	-0.038 (-2.05)
Overall R^2	0.305	0.852	0.141	0.409
Number of observations	3,655,087	248,415	245,302	250,658
Panel B: Individual investors				
Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Patronage dummy	0.174 (5.83)	0.281 (1.84)	0.196 (5.29)	-0.045 (-0.53)
× Wealth above median	-0.025 (-1.98)	0.128 (4.41)	-0.056 (-1.77)	-0.009 (-0.29)
× Trading activity above median	-0.035 (-2.02)	-0.045 (-2.29)	0.023 (0.80)	0.011 (0.17)
× Uses multiple brokers	-0.044 (-7.44)	-0.078 (-5.12)	-0.123 (-21.45)	-0.212 (-10.00)
× Age ≥60 years	0.047 (4.22)	0.017 (1.94)	-0.005 (-0.49)	-0.015 (-0.57)
× Retail broker	-0.074 (-3.45)	-0.225 (-1.25)	-0.041 (-1.23)	0.090 (1.66)
× First purchase			0.041 (3.65)	
Ln (Wealth)	0.008 (1.75)	0.328 (6.38)	-0.002 (-0.68)	-0.051 (-1.53)
Overall R^2	0.311	0.850	0.140	0.433
Number of observations	3,544,418	243,071	231,216	119,273

Panel A adds institutions to the sample and interacts the patronage dummy with a dummy variable that takes the value of one for institutional investors and zero otherwise. Panel B runs the regression on individual investors and interacts the patronage dummy with dummies for an investor having wealth above the median, having value of trades in the previous six months above the median, using more than one broker, being 60 years or older, and making the first purchase in any of the broker stocks during the sample period. A dummy for the three retail brokers in the sample is also interacted with the patronage dummy. Column 1 in Panel B determines the median cutoffs for portfolio value and trading activity separately for investors who hold no broker stocks and for investors who hold at least one broker stock.

All of these effects are statistically significant at the 5% level. Many of the same interactions are important in the value-of-holdings (column 2), purchase (column 3), and sale (column 4) regressions as well, suggesting that the patronage-investing relation varies significantly across investors and brokers. The only other effect we emphasize here is that captured by the positive estimate on the “first purchase” dummy variable (t -value = 3.65) in column 3’s purchase regression. This estimate indicates that patronage decisions are more tightly linked to investment decisions for investors buying their first broker stock.

A first-time buyer is 21% ($= 0.041/0.196$) more likely than a non-first-time buyer to buy own-broker shares.¹⁸

3. Explanations

In this section, we discuss the channels through which product market experiences can influence investment decisions. We consider five channels and, where data permit, expose them to empirical scrutiny.

3.1 Institutional factors

Three institutional factors could explain why investors tilt their portfolios toward brokers they support as customers. First, brokers may recommend their firms' stocks to their customers. The evidence against this explanation comes from our analysis of the automotive industry, in which we find results similar to those in the brokerage firms.¹⁹

Second, bank employees may use their own companies' services and simultaneously invest in bank stock via an employee stock option or ownership plan. This explanation could influence the link between patronage and ownership but not between the patronage and open-market buy transactions our article analyzes. Nevertheless, we address this explanation by discarding finance professionals from the brokerage-firm analysis. Limiting our sample to those individuals for whom we have information on their profession gives us a sample size of about 1.3 million in the holdings regression similar to the one reported in column 2 of Table 2. The slope coefficient estimate is 0.065 (t -value = 5.81) in this alternative sample. This number, as well as the coefficients from the buy and sell regressions, are similar to our baseline estimates.

Third, shareholders could receive consumption perks from the firms in which they invest. Such practices are common in the United States and in the United Kingdom, where companies such as Berkshire Hathaway, Starbucks, Marks & Spencer, and British Airways provide shareholders with discounted products and services. We are not aware of any perks (e.g., free custodial services for the customers of brokers) our sample banks would have offered their shareholders. Similarly, neither Ford nor Saab dealers offered any discounts to shareholders of Ford and Metso. In the case of Saab-Metso, the lack of such incentives is particularly understandable given that Metso was only responsible for the assembly of Saabs.

¹⁸ The Online Appendix reports a specification that interacts the patronage dummy with a "previous-purchase" dummy. This variable takes the value of one if an investor has bought shares of the brokerage firm at any time in the past. The results from this analysis suggest that although historical purchase patterns are important—the coefficient estimate for the previous-purchase dummy is 0.108 (t -value = 3.66)—the patronage dummy remains significant. The coefficient estimate is 0.058 with a t -value of 6.03.

¹⁹ The results from the broker and car industries must be compared with caution because institutional factors vary across industries. For example, customers may have different emotional responses when choosing between cars than when choosing between brokers.

3.2 Home bias

This subsection investigates whether the link between patronage and ownership is driven by home bias: the tendency of consumers and investors to favor local companies. At the outset, home bias would seem an unlikely explanation for our results, because all sample brokers (except one) are headquartered in Helsinki, the capital city of Finland, and tend to serve customers in the whole country. As a consequence, there is almost no variation in “localness,” and thus very little room for home bias to explain why consumers and investors would favor one broker over another.

Table 10 repeats the regressions from Tables 2 and 3 for three subsamples in which we would expect home bias to have a negligible effect on the link between patronage and ownership. The first subsample (in Panel A) controls for home bias by excluding all investors who live within 250 kilometers of any sample broker. The second subsample (in Panel B) controls for home bias by excluding the broker headquartered outside Helsinki. The third subsample (in Panel C) controls for home bias by excluding all but three brokers with a nationwide branch network, and adds distance to the nearest branch as an additional control. If home bias were the cause of the link between patronage and ownership, we would expect the estimates on the patronage dummy in Table 10 to be significantly lower than the corresponding estimates in Tables 2 and 3.

We find, however, that the baseline estimates on the patronage dummy reported in Tables 2 and 3 are close to the restricted estimates reported in Table 10. For example, the baseline estimate in the ownership regression in Table 2, column 2, is 0.071—not much higher than the corresponding Table 10 estimate with a distance restriction (0.065), Helsinki restriction (0.066), and branch network restriction (0.062). The similarity of the restricted and full-sample estimates suggests that the effect of home bias on the patronage-ownership relation is small. A comparison of Table 3 against the restricted-sample estimates in Table 10 supports the same conclusion.

3.3 Private information

Customers may have an informational advantage over non-customers that could contribute to the customer-ownership link in two ways. First, customers (who receive private signals) would be more likely to invest in the company when constraints on short sales are present. Second, non-customers could avoid investing in the company because they have an informational disadvantage. Either of these scenarios could lead customers to own a disproportionate share of the company stock.

Using the calendar-time methodology of Jegadeesh and Titman (1993) to analyze the returns on customers’ and non-customers’ purchases and sales of broker stocks, we test whether customers have superior access to information. We use a one-month formation period and hold the portfolio for three months. (The results in this section are not sensitive to the choices of the lengths of

Table 10
Patronage effect controlling for home bias

Panel A: At least 250 km away from company headquarters				
Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Patronage dummy	0.065 (6.37)	0.108 (4.33)	0.072 (5.15)	-0.038 (-4.15)
Overall R^2	0.307	0.843	0.156	0.491
Number of investors	76,362	22,409	6,189	22,693
Number of observations	774,698	52,060	57,499	31,442
Panel B: Brokers with headquarters in Helsinki				
Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Patronage dummy	0.066 (6.49)	0.092 (3.98)	0.075 (5.17)	-0.027 (-2.60)
Overall R^2	0.334	0.857	0.120	0.493
Number of investors	357,424	103,187	25,256	84,072
Number of observations	2,945,179	232,847	192,438	116,634
Panel C: Distance to nearest bank branch				
Dependent variable	Ownership dummy (0/1)	Ln (Value of holdings)	Purchase dummy (0/1)	Sale dummy (0/1)
Specification	1	2	3	4
Patronage dummy	0.062 (9.85)	0.078 (2.08)	0.063 (5.87)	-0.030 (-2.34)
Ln (Distance to nearest branch)	-0.005 (-2.34)	-0.012 (-0.94)	0.002 (1.06)	0.003 (1.22)
Mean distance	6,444.94	5,259.04	6,484.80	5,176.52
Standard deviation of distance	13,724.11	10,942.45	13,984.70	10,096.93
Overall R^2	0.426	0.881	0.207	0.572
Number of investors	357,405	95,193	24,674	78,008
Number of observations	1,799,753	194,609	114,387	98,876

Columns 1 and 2 correspond to regressions in columns 2 and 5 in Table 2. Columns 3 and 4 report regressions similar to columns 2 and 4 in Table 3. Panel A reports results for the investors who live more than 250 kilometers away from the headquarters of any of the sample banks. Panel B reports regressions that leave out the bank not headquartered in Helsinki (the capital). Panel C runs the regressions for the three banks that have nationwide branch networks. The regressions include an additional explanatory variable that is the (log) distance from an investor's place of residence to the nearest branch of a bank. The distances are based on the population centroids of each zip code. Following Grinblatt and Keloharju (2001a), we define distance as one-quarter of the distance between the centroids of the zip code and the nearest zip code for investors living in the same zip code as the branch. The branch and investor locations are from the end of the year 2000.

the formation and holding periods.) We compute the returns as follows. First, for each month t , we compute the returns on three long-short portfolios: a portfolio formed in month $t-1$, a portfolio formed in month $t-2$, and a portfolio formed in month $t-3$. The month- t return on the one-month formation/three-month holding strategy is the average return over these three portfolio returns. This strategy in effect rebalances one-third of the portfolio each month. By computing averages of different strategies for each month, we avoid the overlap

in returns that would otherwise arise when studying three-month holding periods using monthly data.

We report the returns on six different long-short portfolios (and for their long and short components) in Table 11. Panel A reports on a strategy that buys the broker stocks broker customers buy and finances these purchases by selling the broker stocks non-customers buy; Panel B reports on a strategy that sells the broker stocks broker customers sell and buys the broker stocks non-customers sell; and Panel C first computes the buy-minus-sell returns for customers and non-customers and then examines the return on a strategy that takes a long position in the customers' buy-minus-sell portfolio and a short position in the non-customers' buy-minus-sell portfolio. Finally, we repeat each of these computations by defining the portfolio either in terms of the number of transactions (the left-hand side of the table) or in terms of the value of transactions (the right-hand side of the table).

The results in Table 11 do not support the information-asymmetry story. The returns on each of the three strategies are negligible. Although the relatively short sample period and the low number of broker stocks may contribute to

Table 11
Performance of purchases and sales in brokerage firm stocks

Panel A: Purchases						
	Number of transactions			Value of transactions		
	Customers	Others	Difference	Customers	Others	Difference
Mean return, %	1.27	1.23	0.04	1.39	1.28	0.11
Standard deviation	5.39	4.75	1.55	5.64	4.90	1.67
<i>t</i> -value	(2.12)	(2.32)	(0.24)	(2.22)	(2.36)	(0.59)
Panel B: Sales						
	Number of transactions			Value of transactions		
	Customers	Others	Difference	Customers	Others	Difference
Mean return, %	1.08	1.05	0.03	1.12	1.13	0.00
Standard deviation	5.19	4.64	1.32	5.48	5.03	1.18
<i>t</i> -value	(1.87)	(2.03)	(0.20)	(1.84)	(2.01)	(-0.04)
Panel C: Purchases less sales						
	Number of transactions			Value of transactions		
	Customers	Others	Difference	Customers	Others	Difference
Mean return, %	0.19	0.18	0.01	0.27	0.16	0.11
Standard deviation	1.12	1.05	0.93	1.23	1.06	1.17
<i>t</i> -value	(1.54)	(1.54)	(0.12)	(1.99)	(1.33)	(0.89)

Returns over 20-day periods are estimated separately for customers (investors who have traded through a broker in the past 120 days) and other investors. The sample includes investors who have traded during the 120-day period preceding a trade. Portfolios are formed in each 20-day period and held for three consecutive periods with equal weighting. Each portfolio has 81 observations. Panel A uses number and value of purchases, Panel B uses number and value of sales, and Panel C takes the difference between the purchase and sale portfolios. The table multiplies the original returns by 100.

this result, the mean estimates in the difference column are remarkably close to zero. For example, the return on the double-difference portfolio in Panel C is just 0.01% per month. These results suggest that customers' informational advantage over non-customers is unlikely to explain why customers are more likely to invest in the company stock and less likely to sell their holdings.

3.4 Information sets and beliefs

We have established so far that institutional factors and private information are unlikely explanations for the relation between product market experiences and investment. In what follows, we consider whether we can reconcile our results with systematic differences in the information sets and beliefs between customers and other investors.

3.4.1 Awareness hypothesis. Barber and Odean (2008) find that individual investors tend to be net buyers of attention-grabbing stocks, such as stocks in the news and stocks with extreme one-day returns. Attention-driven buying results from the difficulty investors have in searching the thousands of stocks they can potentially buy. An investor may end up buying a stock simply because a customer relationship with the company makes her aware of the stock. We call this explanation the awareness hypothesis.

The awareness hypothesis can reconcile an investor's tendency to own and purchase shares of a broker she frequents as a customer, but it cannot explain why a customer is less willing to part with the shares of her broker.²⁰ Given the paucity of short sales in our data, an investor cannot sell a stock unless she has first acquired and become aware of that stock. (An interesting extension here would be to measure whether customers are more or less likely to sell short their own brokerage firms. We leave this test for future research.) Moreover, awareness does not, by itself, explain why customers tend to have larger ownership stakes conditional on investment or why the patronage-ownership relationship is stronger for individuals with longer customer relationships. Thus, although the awareness hypothesis could amplify some of the effects we observe on the buy side, alone it cannot sufficiently explain all the patronage-ownership patterns we document.

3.4.2 Perceived information hypothesis. Although no evidence suggests that customers would have an informational advantage over non-customers, customers' trading decisions may nevertheless differ from those of other investors because of differences in perceived information. Frequent encounters

²⁰ "Awareness" could influence investors' purchase decisions by reducing search costs. Barber and Odean (2008), for example, find that individual investors are more likely to purchase stocks that "grab their attention." They motivate their study by noting that investors are more likely to take a shortcut (and avoid search costs) by buying stocks with abnormal visibility. They also note that attention is unlikely to influence investors' sell decisions, because investors need only to sift through the stocks they already own.

with the company may make customers perceive their private signals about the company as more precise than those of non-customers, in the spirit of Kyle and Wang (1997). We test this perceived-information hypothesis by studying customers' and non-customers' trading behavior before and after earnings announcements. If customers believe they are better informed about corporate matters than the rest of the market, we would expect them to be disproportionately present in the order flow in the days leading up to an earnings announcement, when asymmetries of information among investors tend to be particularly large.²¹

The first column in Table 12 estimates a regression that measures how customers and non-customers time their trades relative to earnings announcements. We use a three-week window from trading day $t-20$ to trading day $t+2$ as the main specification because some of the earnings-announcement information gets incorporated in prices only after the announcement.²² The regression's dependent variable takes the value of one if an investor trades in a broker stock during the $[-20, 2]$ earnings-announcement window. If the investor trades in the stock during the month after the announcement (that is, during a $[3, 20]$ window), the dependent variable takes the value of zero. As in our other tests, the key explanatory variable in the regression is the patronage dummy variable. A positive slope coefficient would indicate that customers trade more actively during periods of heightened information asymmetry. In this regression, as in our other regressions explaining trading decisions, we also control for investors' wealth and the past returns of the broker stock.

The results do not support the perceived-information hypothesis. The slope coefficient estimate for the patronage dummy variable is small (-0.004) and statistically insignificant (t -value= -0.67). The coefficient remains insignificant (-0.001 ; t -value= -0.20) if we shorten the pre-earnings-announcement window to include only the days before the announcement—that is, to a $[-20, -1]$ window. Because our sample consists of 131 quarterly earnings announcements and thousands of trades, we are confident about the robustness of our findings. Customers and non-customers do not appear to time their trades differently relative to earnings announcements.

3.4.3 Private valuation hypothesis. Research on social psychology has documented that people tend to think more positively about things merely because they are familiar with them.²³ It is therefore worthwhile to study whether product market experiences lead consumers to have a more favorable view of the companies they frequent as customers. Such optimism may spill

²¹ See, for example, Dubinsky and Johannes (2006).

²² See, for example, Lee, Mucklow, and Ready (1993), Kim and Verrecchia (1994), and Krinsky and Lee (1996).

²³ For a review of the literature of this so-called exposure effect, see Zajonc (1968), Bornstein (1989), and Zajonc (2001).

Table 12
Testing additional predictions of the belief-based hypotheses

Dependent variable	Trade before announcement (0/1)	Limit price relative to	
		ask	bid
		Purchases	Sales
Specification	1	2	3
Patronage dummy	-0.004 (-0.67)	-0.0001 (-0.10)	0.001 (1.01)
Ln (Wealth)	-0.005 (-3.26)	-0.0001 (-0.90)	-0.0003 (-0.84)
Return [-1, -30]	-0.301 (-17.82)		
Return [-31, -60]	-0.023 (-1.39)		
Return [-61, -90]	0.312 (17.21)		
Return [-91, -120]	-0.024 (-1.23)		
Ownership dummy		-0.001 (-1.62)	-0.001 (-1.35)
Spread over last 20 days		-0.897 (-13.98)	0.895 (10.95)
Investor fixed effects	Yes	Yes	Yes
Broker fixed effects	Yes	Yes	Yes
Broker-year fixed effects	No	No	No
Overall R^2	0.012	0.308	0.256
Number of investors	48,317	5,651	5,879
Number of observations	115,870	10,692	10,110

Column 1 tests the perceived-information hypothesis by investigating whether customers are more likely to trade before an earnings announcement. The dependent variable takes the value of one if an investor trades in a broker stock in trading days [-20,2] before a scheduled quarterly earnings announcement of the broker and zero if the investor trades in the stock in trading days [3,20] after the announcement. The analysis determines the value of the dependent variable and the independent variables (defined as in Table 2) each day an investor trades in the broker stock. The sample consists of 131 quarterly earnings announcements made by the sample companies. Columns 2 and 3 test the private valuation hypothesis by comparing the behavior of customers in the order book to that of investors at large. Column 2 uses purchase limit orders and divides the price of the limit offer by the ask price prevailing in the market at the time of entering the order, whereas column 3 uses sale limit orders and divides the limit price by the prevailing bid. The t -values, reported in parentheses below the coefficients, are robust for heteroscedasticity and are adjusted for clustering at the investor level.

over to stock return expectations, leading customers to consider the firms with which they do business to be undervalued in the stock market. We test the hypothesis that customers have higher private valuations of assets than investors at large by studying how customers and non-customers place their limit orders. If customers have higher private valuations, they should be standing ready to buy stocks at prices that are closer to the prevailing market price and sell at prices that are further away from the market price. In other words, if both a customer and a non-customer place a limit buy order for a stock, we would expect the customer's order to be closer to the market.

We base our test on the private valuation hypothesis on a three-year period (1998–2001) for which we have limit order data. We measure how investors place their limit orders by comparing the limit price with the bid and ask prices that prevailed at the time the order was initially submitted. In the results we report in Table 12, columns 2 and 3, we study purchases and sales separately.

For sell limit orders, we compute the percentage difference between the limit order price and the best bid price in the market. The larger this difference is, the farther away the limit order is placed from the current market price. The distance for buy limit orders is the difference between the best ask price and the price of the limit order. We use these distance variables as dependent variables in the buy and sell regressions. A negative coefficient estimate on the patronage dummy would indicate that customers place their limit orders closer than non-customers to the prevailing market price.

The estimation results in columns 2 and 3 do not support the private-valuation hypothesis. The economically and statistically insignificant patronage dummy estimates suggest that customers and non-customers do not make markedly different choices in the limit order book. This result is robust to how we measure distances. For example, the results remain unchanged if we measure distances not to the bid and ask price but to the midpoint of the bid-ask spread. Moreover, robustness tests suggest that unexecuted limit orders (which we do not observe) do not influence our estimates. These tests draw on the idea that sell (buy) limit orders on a stock tend to execute on the days the stock price rises (falls) significantly. We narrow our sample on days on which a particular stock (or the entire market) experiences either a large positive or negative return, ensuring that a greater fraction of (outstanding) limit orders executes in the conditional than in the unconditional sample. The patronage dummy is statistically and economically insignificant in every return-conditioned subsample. Potential differences in limit order execution rates thus do not appear to be significant enough to influence the outcome of the private-valuation test.

3.5 Investor preferences

The remaining explanation for the link between patronage and ownership involves preferences. A lasting preference for the products and services of a company may translate into a direct preference for its stock; a patron could derive direct utility from owning shares of a company she supports as a customer. One plausible reason for a patron becoming a shareholder is loyalty to the company. Our results are consistent with the loyalty explanation: the link between patronage and ownership is stronger for customers with a longer relationship with the company, and weaker for investors spreading their business across multiple companies. Also, Cohen (2009) finds that employees invest in employer stock partly because of loyalty.

The link between patronage and ownership could also stem from investors' preference for the familiar. We evaluate this hypothesis by testing a prediction based on Cao et al.'s (2011) model of familiarity. In this model, investor preferences include a familiarity bias: the larger the bias, the greater the investor's aversion toward stocks unknown to her. We would expect customer-stockholders (who have invested in a familiar stock) to have a larger bias than comparable non-customer-stockholders (who have invested in a less familiar stock).

In Cao et al.'s (2011, p. 184) model, a "familiarity-biased investor holds on to portfolios with a much fewer number of stocks." This assertion suggests that the portfolios of customer-stockholders should be less well diversified than those of non-customer-owners. This hypothesis also conforms to Heath and Tversky's (1991, p. 26) suggestion that "investors are sometimes willing to forego the advantage of diversification and concentrate on a small number of companies with which they are presumably familiar," a notion Huberman (2001) also refers to in his study of familiarity.

We test this prediction by estimating a zero-truncated negative binomial regression. The dependent variable is the number of stocks in the portfolio, and the main independent variable is the investor's customer-owner status. We include the same controls as elsewhere in the article and add one customer and one ownership dummy for each broker. We include these dummy variables to control for any omitted variables that could be correlated with an investor's choice to invest in or be a customer of a particular broker.²⁴ We restrict the sample to investors who own shares of at least one broker. Hence, the only difference between a customer-owner and a non-customer-owner in our sample is that the former invests in her own broker, whereas the latter invests in a broker other than her own.

Table 13 shows that our results are consistent with Cao et al.'s (2011) model: customer-owners' portfolios are less diversified than those of non-customers. The customer-owner dummy has a coefficient of -0.064 (t -value = -13.73), and the marginal effect associated with this estimate is -0.37 .²⁵ The latter number implies that whereas the average non-customer-owner holds a portfolio of 5.5 stocks, the average customer-owner holds 5.1 stocks. These results are consistent with the view that the link between consumption and ownership can be at least partly explained by investors' preference for familiarity (or, conversely, their fear of the unknown).

4. Conclusion

Consumer experiences shape investors' portfolio choices. We show that banks' brokerage customers are significantly more likely than non-customers to acquire shares in the bank and then less likely to sell them. Patronage also increases the size of the ownership stake, conditional on investment, and its effects on investment are stronger for individuals with longer customer

²⁴ For example, the average customer of one of the three retail brokers could be less sophisticated than the average customer of one of the remaining five (non-retail) brokers who compete for wealthier clients (Vissing-Jørgensen 2003). Such differences in sophistication could in turn correlate with the number of stocks in an investor's portfolio. At the same time, the patronage-ownership relation varies significantly across brokers (see Figure 1). The sixteen dummy variables resolve this issue by subsuming any systematic variation in investor characteristics across brokers.

²⁵ We also estimated this model using a sample without the restriction that non-customers also have to hold a broker stock. The coefficient estimate is -0.042 (t -value = -10.19) in this alternative specification.

Table 13
Customer-owners and portfolio diversification

Dependent variable	Number of stocks
Customer-owner dummy	-0.064 (-13.73)
Ln (Wealth)	0.312 (183.75)
Female dummy	-0.086 (-17.50)
Age	0.010 (16.70)
Age ²	-0.0001 (-18.88)
Year fixed effects	Yes
Ownership and patronage controls for each broker	Yes
Number of investors	104,840
Number of observations	193,786

This table explains the number of stocks held in an investor's portfolio at the end of each year with a dummy that takes the value of one if the investor holds shares in the broker whose customer she is and zero otherwise. The model is a zero-truncated negative binomial regression for count data. The regression also includes sixteen unreported dummy variables for an investor being either an owner (eight dummies) or a customer (eight dummies) of a broker. The sample is restricted to investors who own shares of at least one broker stock. The *t*-values, reported in parentheses below coefficients, are robust for heteroscedasticity and adjusted for clustering at the investor level.

relationships. A natural experiment using two insurance company-bank mergers yields estimates similar to those obtained in a full sample. Our evidence suggests a causal link not only from patronage to investing but also (a somewhat weaker effect) in the opposite direction.

Our results are difficult to explain by broker recommendations, home bias, customers' real or perceived information advantage, greater awareness of the stock, higher private valuation, or many other hypotheses. Instead, our results are most consistent with a preference-based setup in which customer-investors regard stocks as consumption goods, not just as investments. This setup could explain why patronage explains buy and sell decisions, why its effect is stronger for individuals with longer customer relationships, and why patrons tend to have larger stockholdings and less well-diversified portfolios than non-patrons.

Although data restrictions ultimately direct our focus to brokerage and automotive industries, both of these industries are well suited for an analysis of the patronage-ownership link. Both are prime examples of industries in which consumers are the end users of products and services in which one would hence expect the link between patronage and ownership to be the strongest. An additional benefit of having data on the brokerage industry is that it is natural to extend the examination of the patronage-ownership link to professional money managers.

Despite using data on only two industries, our results lead us to believe that individuals' consumer experiences and stock-selection decisions also interact

in other industries. In those industries in which individuals constitute the end users of the company products, we would expect to find an equally strong patronage-ownership link. By contrast, we would expect to find a weaker patronage-ownership relation in those industries in which companies do not service individuals directly. An interesting avenue for future research that would require more data would be to test whether the visibility of the company's brand affects the patronage-ownership link.

The patronage-ownership effect has both descriptive and normative implications for investor and firm behavior. First, it could account for some of the unexplained heterogeneity in individual investors' portfolios. It is well known that many individuals do not invest in the market portfolio (as recommended by theory), but their portfolio choices are also remarkably different in other ways. For example, the 46,791 active accounts in the Barber and Odean (2000) data set held 18,099 different securities in November 1996, and 7,979 of these securities were held by just one investor! Given that investors consume different product and service bundles, the effect of patronage on ownership may induce them to hold different portfolios as well.

Second, our findings can guide companies in their efforts to both acquire and retain customers and investors. The strong link between ownership and patronage decisions suggests that companies could benefit from integrating their ownership and customer relationship records. This integration would enable companies to target their products and services to their existing investors (to exploit the effect from ownership to patronage) or target equity issues to their existing customers (to exploit the effect from patronage to ownership).

Third, our results have implications for corporate control. Given that customers are less willing than other investors to sell their shares, a robust customer-owner base may allow a company to fend off takeover attempts. The link between patronage and ownership could even encourage managers to reduce their firm's vulnerability to a takeover by selling its shares to its customers, much in the same way managers might ally with employees to tip the scale toward management in a takeover attempt (Pagano and Volpin 2005).

References

- Angrist, J., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.
- Bagwell, K. 2007. The Economic Analysis of Advertising. In M. Armstrong and R. Porter (eds.), *Handbook of Industrial Organization*, vol. 3, pp. 1701–1844. Amsterdam: North-Holland.
- Baik, B., J.-K. Kang, and J.-M. Kim. 2010. Local Institutional Investors, Information Asymmetries, and Equity Returns. *Journal of Financial Economics* 97(1):81–106.
- Barber, B., and T. Odean. 2000. Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *Journal of Finance* 55(2):773–806.
- Barber, B., and T. Odean. 2008. All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies* 21(2):785–818.

- Baumann, C., S. Burton, and G. Elliott. 2005. Determinants of Customer Loyalty and Share of Wallet in Retail Banking. *Journal of Financial Services Marketing* 9:231–48.
- Benartzi, S. 2001. Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock. *Journal of Finance* 56(5):1747–64.
- Billett, M., Z. Jiang, and L. Rego. 2010. Does Customer Sentiment Contribute to Investor Sentiment: Glamour Brands and Glamour Stocks? Working Paper, University of Iowa.
- Bornstein, R. F. 1989. Exposure and Affect: Overview and Meta-Analysis of Research, 1968–1987. *Psychological Bulletin* 106(2):265–89.
- Brehm, J. 1956. Post-Decision Changes in Desirability of Alternatives. *Journal of Abnormal and Social Psychology* 52:384–89.
- Cao, H. H., B. Han, D. Hirshleifer, and H. H. Zhang. 2011. Fear of the Unknown: Familiarity and Economic Decisions. *Review of Finance* 15:173–206.
- Choe, H., B.-C. Kho, and R. Stulz. 2005. Do Domestic Investors Have an Edge? The Trading Experience of Foreign Investors in Korea. *Review of Financial Studies* 18(3):795–829.
- Cohen, L. 2009. Loyalty-Based Portfolio Choice. *Review of Financial Studies* 22(3):1213–45.
- Coval, J., and T. Moskowitz. 1999. Home Bias at Home: Local Equity Preference in Domestic Portfolios. *Journal of Finance* 54(6):2045–73.
- Coval, J., and T. Moskowitz. 2001. The Geography of Investment: Informed Trading and Asset Prices. *Journal of Political Economy* 109(4):811–41.
- Døskeland, T., and H. Hvide. 2011. Do Individual Investors Have Asymmetric Information Based on Work Experience? *Journal of Finance* 66(3):1011–1141.
- Dubinsky, A., and M. Johannes. 2006. Earnings Announcements and Equity Options. Working Paper, Columbia University.
- Dvořák, T. 2005. Do Domestic Investors Have an Information Advantage? Evidence from Indonesia. *Journal of Finance* 60(2):817–39.
- Fama, E. F., and K. R. French. 2007. Disagreement, Tastes, and Asset Prices. *Journal of Financial Economics* 83(3):667–89.
- Festinger, L. 1957. *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Frazzini, A. 2006. The Disposition Effect and Underreaction to News. *Journal of Finance* 61(4): 2017–46.
- Frieder, L., and A. Subrahmanyam. 2005. Brand Perceptions and the Market for Common Stock. *Journal of Financial and Quantitative Analysis* 40(1):65–85.
- Froot, K. A., and T. Ramadorai. 2008. Institutional Portfolio Flows and International Investments. *Review of Financial Studies* 21(2):937–71.
- Garmaise, M., and T. Moskowitz. 2006. Bank Mergers and Crime: The Real and Social Effects of Credit Market Competition. *Journal of Finance* 61(2):495–539.
- Grinblatt, M., and M. Keloharju. 2000. The Investment Behavior and Performance of Various Investor Types: A Study of Finland's Unique Data Set. *Journal of Financial Economics* 55(1):43–67.
- Grinblatt, M., and M. Keloharju. 2001a. How Distance, Language, and Culture Influence Stockholdings and Trades. *Journal of Finance* 56(3):1053–73.
- Grinblatt, M., and M. Keloharju. 2001b. What Makes Investors Trade? *Journal of Finance* 56(2): 589–616.
- Grinblatt, M., M. Keloharju, and S. Ikäheimo. 2008. Social Influence and Consumption: Evidence from the Automobile Purchases of Neighbors. *Review of Economics and Statistics* 90(4):735–53.

- Grullon, G., G. Kanatas, and J. P. Weston. 2004. Advertising, Breadth of Ownership, and Liquidity. *Review of Financial Studies* 17(2):439–61.
- Hau, H. 2001. Location Matters: An Examination of Trading Profits. *Journal of Finance* 56(5): 1959–83.
- Heath, C., and A. Tversky. 1991. Preferences and Beliefs: Ambiguity and Competence in Choice Under Uncertainty. *Journal of Risk and Uncertainty* 4:5–28.
- Holliday, K. 1995. Will Bank Mergers Destroy Customer Loyalty? *U.S. Banker* 105(10):30–37.
- Hong, H., and M. Kacperczyk. 2010. Competition and Bias. *Quarterly Journal of Economics* 125(4): 1683–1725.
- Hong, H., J. D. Kubik, and J. Stein. 2008. The Only Game in Town: Stock-Price Consequences of Local Bias. *Journal of Financial Economics* 90(1):20–37.
- Huberman, G. 2001. Familiarity Breeds Investment. *Review of Financial Studies* 14(3):659–80.
- Ivković, Z., and S. Weisbenner. 2005. Local Does as Local Is: Information Content of the Geography of Individual Investors' Common Stock Investments. *Journal of Finance* 60(1):267–306.
- Jegadeesh, N., and S. Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance* 48(1):65–91.
- Kim, O., and R. E. Verrecchia. 1994. Market Liquidity and Volume Around Earnings Announcements. *Journal of Accounting and Economics* 17(1–2):41–67.
- Krinsky, I., and J. Lee. 1996. Earnings Announcements and the Components of the Bid-Ask Spread. *Journal of Finance* 51(4):1523–35.
- Kyle, A. S., and F. A. Wang. 1997. Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test? *Journal of Finance* 52(5):2073–90.
- Lambert-Pandraud, R., G. Laurent, and E. Lapersonne. 2005. Repeat Purchasing of New Automobiles by Older Consumers: Empirical Evidence and Interpretations. *Journal of Marketing* 69(2):97–113.
- Lee, C. M. C., B. Mucklow, and M. J. Ready. 1993. Spreads, Depths, and the Impact of Earnings Information: An Intraday Analysis. *Review of Financial Studies* 6(2):345–74.
- Linnainmaa, J. 2010. Do Limit Orders Alter Inferences About Investor Behavior and Performance? *Journal of Finance* 65(4):1473–1506.
- Lou, D. 2010. Attracting Investor Attention Through Advertising. Working Paper, London School of Economics.
- Lynch, P. 1993. *Beating the Street*. New York: Simon and Schuster.
- Massa, M., and A. Simonov. 2006. Hedging, Familiarity and Portfolio Choice. *Review of Financial Studies* 19(2):633–85.
- Morrall, K. 1996. Managing a Merger Without Losing Customers. *Bank Marketing* 28(3):18–23.
- Pagano, M., and P. F. Volpin. 2005. The Political Economy of Corporate Governance. *American Economic Review* 95:1005–30.
- Parthasarathy, H. 2007. Impact of Mergers on Customer Retention and Acquisition. Working Paper, Harvard University.
- Patterson, P. 2007. Demographic Correlates of Loyalty in a Service Context. *Journal of Service Marketing* 21(2):112–21.
- Seasholes, M., and N. Zhu. 2010. Individual Investors and Local Bias. *Journal of Finance* 65(5): 1987–2010.
- Sweet, D. 2008. Packers Shareholders (All 122,000) Are Smiling. msnbc.com, January 16.

Vissing-Jørgensen, A. 2003. Perspectives on Behavioral Finance: Does “Irrationality” Disappear with Wealth? Evidence from Expectations and Actions. *NBER Macroeconomics Annual*: 139–94.

Zajonc, R. 1968. Attitudinal Effects of Mere Exposure. *Journal of Personality and Social Psychology* 9(2): 1–27.

Zajonc, R. 2001. Mere Exposure: A Gateway to the Subliminal. *Current Directions in Psychological Science* 10(6):224–28.

Copyright of Review of Financial Studies is the property of Oxford University Press / USA and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.