Who Buys the Crash? Household Buying and Price Reversals in Four Market Shocks, from the COVID-19 Crash to the Invasion of Ukraine[1]

P. Joakim Westerholm Abhishek Sehgal

November 2026

ABSTRACT

Finland's shareholder register spans four shocks, 2018-2022 (OMXH25: −17.0%, −36.6%, −8.8%, −23.9%). In the *crash window* of three shocks households were net buyers, and in two *foreign nominee accounts* net sellers, on the exchange. At *register prices* no five-day reversal is detected in the *most-traded shares*. In the other listed shares one is detected, and none on the midpoint. A third of crash-window traders can be measured: buyers after falls, or sellers of winners who keep losers, are net buyers more often. From 21 February 2022 returns fell with Russia exposure. That slope does not pass the family of tests.

Keywords: household finance, investor flows, market crashes, return reversal, retail investors, Finland.
JEL classification: G11, G12, G14, G41.

[1]Westerholm (corresponding author): Discipline of Finance, University of Sydney Business School, H69 Codrington Building, Sydney NSW 2006, Australia; joakim.westerholm@sydney.edu.au. Sehgal: independent researcher; abhishek@finsoeasy.com.

Introduction

When a stock market falls fast, every share that is sold is bought by someone. Whether the drop is news or price pressure is read from what prices do next; who bought says which investors took the other side. If sellers must sell and buyers ask a price concession, prices fall too far and recover. If the drop is news, prices stay down. A crash packs the selling into a few days and makes the price moves large, so the two readings make their sharpest predictions there. We ask who bought in four market shocks in Finland, and whether prices reversed afterward.

The data are the trades in the Finnish shareholder register from January 2018 to December 2022. The register holds every account that trades a listed Finnish share, and places nearly every account in one of seven investor classes (1,002 carry no class; Table I). In these years 499,166 household accounts made at least one exchange trade in a listed share. A custodian keeps *foreign nominee accounts* in its own name for clients abroad, so the register sees the custodian and not the owners. The register holds no orders, no quotes, no wealth and no opening positions. It does not record which side started a trade, so a net purchase does not say who paid for immediacy.

Flows are measured *on the exchange*: trades in listed shares that the register marks as exchange trades, dated by trade day (Section II.B). The measure leaves out tickets, single sides of a trade, of EUR 25 million or more, nearly all of them on foreign nominee accounts. Table IA.LIX repeats the flow results with those tickets counted. The class totals need not sum to zero.

A rule applied to weekly market returns finds four shocks: the 2018 sell-off, the COVID-19 crash, the 2021 scare and the invasion of Ukraine. Each shock has a *crash window*, the trading days of its flagged weeks: 15 days in the COVID-19 crash and five in each of the others. The invasion crash window is 28 February to 4 March 2022, the week after the invasion began on 24 February. Each shock also has a *fall*, the longer span from the peak of the OMXH25 index to its trough, and each result names the one it uses. The 2018 sell-off is the first of the four, and the two shocks in the title have the two largest falls (Table II). Baselines are *regular weeks* and regular days, which lie outside the falls (Section II.D).

Prices come from three series. *Register prices* are volume-weighted averages of each day's exchange trades. *Closes* are the last Compustat prices of the day. The *midpoint* is the last LSEG quote with a bid below the ask. The *most-traded shares* are the first 100 listed shares by median daily traded value, ranked before each window. Households did 90.2% of their crash-window trading in them and 4.2% in the other listed shares (Table IA.XXIII).

We test six hypotheses. Each verdict is the test Section I names, at the 5% level. Passing the family is a separate fact: q, the p-value adjusted across the 45 headline tests, below 5% (Section II.G).

Households were net buyers of directly held Finnish shares on the exchange in three crash windows. They bought EUR 162.6 million in the COVID-19 crash window and EUR 70.5 million in the invasion crash window (Table VI). In the crash window of the 2021 scare they were net sellers. Their purchase of EUR 9.6 million in the 2018 crash window changes sign when trades off the exchange are counted (Panel C of Table VI). Foreign nominee accounts were net sellers on the exchange in the crash windows of the 2018 sell-off and the COVID-19 crash, and net buyers in the 2021 crash window. In the invasion crash window the sign of their flow is not identified, because it rests on trades valued at closes. These sums are description and not a test.

A crash down day is one on which the OMXH25 returned −2% or less inside a crash window. Household net purchases on crash down days, and foreign nominee net purchases on the exchange, cannot be told from other down days (Table VII). Domestic financial institutions bought EUR 7.3 million a day on those days. Households bought on the day the index fell. That test is not the down-day comparison, and it passes the family (Table IA.LXXII). In Korea in 1997 foreign investors sold and domestic individuals bought, Korean institutions sold, and prices adjusted quickly (Choe, Kho, and Stulz (1999)). In the US COVID-19 crash institutions sold and retail investors bought (Glossner et al. (2025)). Studies of that crash often see one brokerage's clients, and the register is the whole market and every investor class (Giglio et al. (2021), Ortmann, Pelster, and Wengerek (2020), Welch (2022)).

In the most-traded shares the five-day slope at register prices is 0.008 (Table III). A slope

near zero is what news would leave. In the other listed shares, 4.2% of household crash-window trading, the slope from day 0 is −0.353. It passes the family. From day 1 that slope is not detected. The comparison with earnings days does not pass the family, and without one earnings-day move it no longer rejects (Table IA.XXIX). On the midpoint no five-day reversal is detected in either sample (Table IA.XXXI). Campbell, Grossman, and Wang read a reversal after non-informational selling as a concession for immediacy, and Nagel as the return to supplying liquidity (Campbell, Grossman, and Wang (1993), Nagel (2012)). Barber et al. find a later loss after the most extreme retail buying of small stocks (Barber et al. (2022)). Section III.E tests the three accounts, which predict when a reversal should be larger and are not themselves the five-day slope. The volume test and the VIX test do not pass the family, and the first-time-buyer result is not this slope.

A third of the accounts that traded in a crash window can be scored. Those that buy after falls, and those that sell winners and keep losers, are net buyers more often (Table IX). Both pass the family. Finnish households trade that way outside crashes too (Grinblatt and Keloharju (2000), Grinblatt and Keloharju (2001)). The register shows those trades in directly held shares, and not holdings through funds (Campbell (2006)). Purchases over 90 days are not detected to beat the index or the sales (Table XI).

From 21 February to 15 March 2022 a share lost 0.52% per point of the Russia-revenue score (Section VI.B; Table XII). The slope does not pass the family. On pricing before the invasion, the two p-values do not agree. Faster selling of the exposed shares is not detected.

The shocks are four observations in one market, so we describe them and make no claim about a typical crash. Two of the four, the 2018 sell-off and the 2021 scare, depend on the thresholds of the rule and on the series it is applied to. Household buying means net buying of directly held Finnish shares on household accounts, with positions traced only from January 2018.

Section I states the mechanisms and hypotheses, and Section II the data and design. Section III tests H1 and H4, Section IV H2 and H3, Section V H5 and Section VI H6. Section VII concludes.

I. Economic Mechanisms and Hypotheses

In a crash the market may have learned something about cash flows, and then a share's fall stays. Or sellers who must trade now may pay for immediacy, and then part of the move comes back once patient capital arrives. Six hypotheses follow, each with a rule. Unless it says otherwise, a rule is read at the 5% level, two-sided, on the unadjusted p-value of the test that it names. The q-value of a headline test says whether it survives adjustment across the family (Section II.G), and is reported beside it. Table IA.I in the Internet Appendix sets the six side by side.

A. H1, reversal (Section III). A seller who must trade now accepts a concession for immediacy, which reverses once the offsetting orders arrive (Grossman and Miller (1988)). H1 predicts that a crash-day move of 5% or more partly reverses within five trading days, and that an earnings-day move of the same size does not. It is judged at register prices on two five-day tests from day 0, in the most-traded shares and in the other listed shares (Table III). One is the crash slope of Panel A. The other is the interaction of Panel C, in the pooled regression that also has a crash indicator. Each is read on its wild bootstrap p-value by trading date. In a sample H1 holds if both are negative and can be told from zero. H1 holds in part if it holds in one sample of the two. The slope from day 1 and the ten-day slope do not enter the verdict.

The earnings-day benchmark is our expectation. Returns drift for weeks in the direction of an earnings surprise (Bernard and Thomas (1989)), in a sort on the surprise and not on the size of the announcement-day move. The reversal should be strongest after moves on high market volume, which mark non-informational selling (Campbell, Grossman, and Wang (1993)), and when the VIX is high (Nagel (2012)). It should also follow the most extreme retail buying of small stocks (Barber et al. (2022)). Section III.E tests each.

B. What the reversal slopes measure. A crash-day return has three parts. News stays, a price concession reverses, and bounce between the bid and the ask reverses the next day. Section III estimates three slopes of the forward return on the day-0 return: from day 0; from day 1, the next trading day; and net of the market. The first holds the concession and the bounce. The second drops

the bounce and the first day of the concession, and the third what the whole market shares. News predicts zero from day 1 and net of the market, and at or below zero from day 0. A concession predicts a negative slope that survives the later start. The midpoint drops the bounce and keeps the first day, and on it no five-day reversal is detected in either sample (Table IA.XXXI).

C. H2, who takes the other side (Section IV). The buyers in a crash should be the investors with cash and no obligation to sell, and the sellers those under mandates, redemptions, or funding pressure. In the model of Shleifer and Vishny (1997) investors withdraw delegated capital after losses. Funds facing redemptions sell at fire-sale prices (Coval and Stafford (2007)), and funding liquidity and market liquidity feed on each other (Brunnermeier and Pedersen (2009)). In Finland, across stocks, households buy past losers and foreign investors buy past winners in ordinary times (Grinblatt and Keloharju (2000)). In the Korean crisis of 1997 foreign investors were net sellers whatever the market return, and Korean individuals were net buyers (Choe, Kho, and Stulz (1999)).

H2 predicts that on crash down days households buy more, and *foreign nominee accounts* sell more on the exchange, than on other down days. A crash down day is a down day inside a crash window (Sections II.D and II.F). A part fails if the difference in mean daily net purchases, crash down days less other down days, is not of the predicted sign or cannot be told from zero. The test is Welch's t, which treats days as independent, at the 5% level on its unadjusted two-sided p-value (Table VII, Panel A). H2 holds if both parts hold, fails if neither does, and holds in part if one does. Every down day of the four falls, counted as a crash down day, is a check. A corollary: if mandates and risk limits bind in a crash, domestic financial institutions sell on the exchange with the foreign nominee accounts. It is judged on a sign, not at the 5% level: it fails if their mean net purchases on crash down days are positive. The register records no mandate (Section II.A), so the constraints are an interpretation.

D. H3, speed (Section IV.D). A price impact that comes from a demand for liquidity reverses slowly, because the capital that takes the other side arrives gradually (Duffie (2010), Mitchell, Pedersen, and Pulvino (2007)). Both papers are about price paths, and the step to flows is ours. Mitchell, Pedersen, and Pulvino find recoveries of months in convertible bonds and merger arbi-

trage. Duffie puts the delay at fractions of a second in an electronic order book for shares. Liquidity stays impaired for about two weeks after a large decline (Hameed, Kang, and Viswanathan (2010)). H3 predicts that the gap between the net purchases of households and of foreign nominee accounts closes more slowly after a crash than after an earnings day. H3 fails if the change in the gap, crash less earnings day (Table VIII), is not positive or cannot be told from zero on its permutation p-value.

E. H4, structure (Section III.G). When very large institutions trade in illiquid markets, returns and volume spike, and their extremes follow power laws (Gabaix et al. (2003), Gabaix et al. (2006)). On that account a crash is a large draw from the same distribution as other days, not a separate regime, which is the null. The alternative starts from scarce speculator capital, which makes price changes negatively skewed in a decline (Brunnermeier and Pedersen (2009)). The steps to volume and to the cross-section of shares on one day are ours. When the capital that takes the other side is scarce, prices move on little trading.

H4 predicts that inside a crash window, in the most-traded shares, the elasticity of a share's abnormal volume to its absolute return is lower than on other trading days. It also predicts that returns across those shares on one day are more negatively skewed, on quantile skewness. A linear slope of volume on the absolute return cannot separate H4 from its null. With a constant elasticity below one, that slope falls whenever moves are large. Each measure is taken at closes and at register prices and read on its permutation p-value (Table V). H4 holds if both are lower than on other trading days, and can be told from them, on each series. It fails if neither is on either series, and holds in part otherwise.

F. H5, which households (Section V). H5 predicts that an account's trading in regular weeks predicts whether it is a net buyer in a crash window. Accounts that trade against returns, and accounts that sell winners and keep losers, should be net buyers more often. We label such persistence habit, and the test does not set habit against skill. The rule is read on the pooled regression of Table IX, column 1, among the accounts with enough regular-week trades to score. A trait predicts if its coefficient can be told from zero on its HC1 t and has the predicted sign. That sign is negative

on the contrarian measure, which is lower for an account that buys after falls, and positive on the disposition measure. H5 holds if both traits predict, fails if neither does, and holds in part if one does.

Skill is a second reading with two predictions, each reported as detected or not detected. The shares bought in a crash should be followed by higher returns than the market and than the shares sold. This is read at 90 days, on every listed share, fall by fall (Table XI, Panel C). It is detected in a fall if purchases net of the index and purchases less sales are both positive and can be told from zero, on t clustered by security. And an account whose last crash purchase gained should come back in the next crash. This is detected if the gain coefficient of the 2018 to COVID-19 pair, the test in the family, is positive and can be told from zero (Table X). That prediction is a trade-level analog of the learning in Seru, Shumway, and Stoffman (2010).

If the buying is habit, we expect it to concentrate in positions below the price paid, the purchase-side counterpart of a reluctance to realize losses (Odean (1998)). In that study the repurchases of US brokerage clients favor small losses over large ones, so buying furthest below the price paid is not the US pattern. In the Finnish register the propensity to sell jumps at a zero return and is flat across losses (Kaustia (2010)). Prospect theory does not settle the direction: in some settings it predicts the opposite of the disposition effect (Barberis and Xiong (2009)).

Experience can cut either way. People who have lived through low stock returns take less financial risk (Malmendier and Nagel (2011)). Older US investors hold safer portfolios, but their stock selection worsens sharply around age 70 (Korniotis and Kumar (2011)). H5 therefore separates the owner's age from the account's time in the register. The register does not hold the returns an owner lived through, so lifetime experience is not tested. Estates and domestic financial institutions are two comparisons that bound the reading (Internet Appendix, below Table IA.I).

G. H6, the invasion (Section VI). The exposure score, from an unpublished manuscript of the same research project (Section VI.A), estimates the share of a firm's revenue that comes from Russia. If prices take in news as it arrives, exposed shares fall more than unexposed ones when the risk becomes public, and not before. The event window opens on 21 February 2022, when Russia

recognized the two breakaway regions of eastern Ukraine, three trading days before the invasion of 24 February.

H6 has three parts. Each is read on two p-values, and a part holds only if both reject with the predicted sign. The two are the HC3 p and the wild bootstrap p of a slope, and for the sale rates the p clustered by share and the wild bootstrap p. First, exposure is priced in the invasion weeks, 21 February to 15 March 2022: the slope of returns on exposure is negative (Table XII). Second, it is not priced in the weeks before, 3 January to 18 February. On one sample of shares, the slope before less the slope in the invasion weeks is positive, with the controls of Table XIII. Third, in the four weeks from 24 February the sale rate of household positions is higher in exposed than in unexposed shares (Table IA.CXI in the Internet Appendix). The slopes are at register prices. H6 holds if all three parts hold, fails if none does, and holds in part otherwise.

II. Data and Empirical Design

A. The register. The data are the trades in the shareholder register of Euroclear Finland, the central securities depository, from January 2018 to December 2022. The register records a trade as legs. A leg is one side of a trade on one account. It gives the security, the number of shares, the settlement date, a reference code and, for most legs, a price in euros. An account carries a class and a legal form. The account of a natural person also carries a birth year, a gender, a language and a postcode. One person can hold several accounts, so a count of accounts is not a count of people. No record of a single account is reported. Table IA.XIX in the Internet Appendix collects the terms defined here.

Table I gives the sample by class. Households are the accounts of the household class, nearly all of them natural persons or estates. A household account holds only the shares that its owner holds directly, and shares owned through a fund are on the account of the fund. The class of *foreign nominee accounts* holds the accounts that a custodian keeps in its own name for clients abroad. The register sees the custodian's account and not the owners behind it, so one account can carry the

Table I

Accounts, legs and gross trading on the exchange by investor class

investor class	accounts	legs	legs of EUR 25 million or more	gross, flow basis (EUR bn)
Households	499,166	17,458,758	5	57.9
Foreign nominee accounts	22	51,979,497	27,488	8,529.2
Domestic financial institutions	683	561,470	197	68.4
General government	147	33,482	154	38.3
Non-financial corporations	18,444	1,398,998	64	28.0
Non-profit institutions	2,200	129,474	8	6.4
Foreign accounts held directly	19,867	395,435	370	50.2
Accounts without a class	1,002	33,224	0	0.3
Every class		71,990,338	28,286	8,778.8

Notes. Exchange trades in listed shares with a trade day from January 2018 to December 2022. The second column counts the accounts with at least one such leg, and the third counts their legs of every size. The fourth counts the legs of EUR 25 million or more, which the flow basis of Section II.B leaves out. Gross is purchases plus sales. The household class holds 1,039,937 accounts, and all but 264 are those of natural persons or of estates. The accounts column is not summed.

trades of many owners.

Three things the register does not record bound what the paper can say. First, it has no opening positions. We therefore observe no holdings, no wealth and no ownership shares, and no result is grouped on them. Second, the register records an account's class but not its mandate or its funding. Third, it does not record which side initiated a trade. A net purchase says who ended the day with more shares, not who paid for immediacy.

An account has a *traced position* in a share when its purchases of that share on the exchange since January 2018 exceed its sales, in shares. Its *price paid* is the euros of all those purchases divided by the shares bought. It averages over every such purchase, shares since sold included, and is not a tax basis. Shares are counted with no adjustment for splits, which misstates both after a split (Internet Appendix, below Table IA.VII).

B. The flow basis. Every flow in the paper is computed on one basis unless its table says otherwise. Five rules define it, and the first part of the Internet Appendix holds the counts behind each.

First, a leg counts when the register marks it as an exchange trade. Two reference codes carry the mark: code 01, a marketplace trade, and code 02, which nearly all legs carry (Table IA.II in the Internet Appendix). The register names code 02 "other trade", which does not say where the leg was executed. We count it because the register marks it as an exchange trade. Transfers, inheritances, corporate actions, corrections and stock loans are left out.

Second, the security is a listed share. A *listed share* is a security with a daily close on the Helsinki exchange in Compustat Global, or a register security matched to such a listing. Nordea changed its identifier in October 2018 and is kept under both (Table IA.V in the Internet Appendix).

Third, a leg is dated by its *trade day*. The register records a trade date for 0.55% of exchange legs, and the recorded dates end on 2 May 2018. Every other leg is dated two settlement days before its settlement day. Legs traded in late December 2017 that settle in January 2018 count, so the weekly series starts in December 2017. A *settlement day* is a weekday on which the register settles legs: every weekday but the closing days of TARGET, the payment system of the euro area. Exchange holidays that are not closing days count toward the two, and a leg that the count puts on one is dated on the last trading day before it. The recorded trade dates and the register's settlement dates support the rule (Table IA.III and Table IA.IV in the Internet Appendix). We hold no calendar published by the depository to check the settlement days against.

Fourth, a leg is valued at the register's price for it. A leg without one is valued at the close of its trade day, and a leg with neither price is not counted.

Fifth, a ticket of EUR 25 million or more is left out. A *ticket* is one leg, sized by its euro value at the price that values it. The limit is that of an unpublished 2026 manuscript of the same research project, "Anatomy of a Black Swan: News, Panic, and the Reallocation of Risk". Its author line names the project and no person, so it has no entry in the references. That study set the limit to leave out tender and block crossings. We apply it unchanged. The limit leaves out 28,286 legs, 27,488 of them on foreign nominee accounts (Table I). Table IA.IX in the Internet Appendix counts those with a register price.

Table IA.LIX repeats the headline flow results with every ticket counted, and Table IA.X gives

the foreign nominee totals at four other limits.

A class buys or sells *on the exchange* when the statement rests on the legs these five rules count. The phrase means marked by the register as an exchange trade, not seen in the order book. We call purchases less sales in euros *net purchases*, and purchases plus sales *gross trading*. Net purchases do not sum to zero across the classes. The other side of a counted leg can be a leg that the rules leave out. The register also does not hold two valued sides of every trade: net purchases over every valued leg do not sum to zero either (Table IA.LXII).

Over 2018 to 2022, register gross divided by two is 7.26 times Compustat turnover in the same shares (Table IA.XI in the Internet Appendix). The register counts both sides of a trade, and Compustat counts a trade once. Most of the excess is in the exchange trades of foreign nominee accounts, on which a purchase and a sale of the same share often meet on one day. With such pairs set against each other on every account, the ratio is 0.93. Gross trading on those accounts is therefore not a measure of trading interest, and the flow results of the paper rest on net purchases. We cannot tell whether borrowed shares settled sales on the exchange (Table IA.XII).

C. Prices. The paper uses two price series for its central results, and the midpoint as a check. The first is *register prices*. For a share and a day, the register price is the volume-weighted average price of the day's exchange legs, on days with at least three priced legs. Two limits apply to them. First, the price recorded on a leg is not always the price of that trade. It equals the close for 67.2% of household legs with both prices and for 63.6% of those of foreign nominee accounts (Table IA.VI in the Internet Appendix). A register price is therefore nearer the close than an average over the day's trades would be. We never compare the price of one trade with the other prices of its day.

Second, the series dates a leg two weekdays before settlement, and not by the rule of Section II.B. The two rules differ only next to an exchange holiday. Without the moves whose day or five-day return lies next to one, the comparison of H1 in the other listed shares still rejects, unadjusted (Table IA.XXIX). The difference of the invasion slopes moves from a bootstrap p of 0.067 to 0.086, and its HC3 p rejects with and without them (Table IA.CV). Returns over 20 or 90 days have not been recomputed.

The second series is *closes*, the last prices of the day on the Helsinki exchange, in euros, from Compustat Global. Returns from closes include dividends and are adjusted for splits. Register prices are adjusted for neither, and Table IA.VII in the Internet Appendix drops the moves on split and ex-dividend days. Both series carry the bounce between the bid and the ask. A close carries it in full, and a register price averages it over the legs of the day. A reversal measured from the day of the move can therefore be bounce. A return from a *day-1 start* begins one trading day after the move and leaves the first day's bounce out. The midpoint is the last LSEG quote of the day with a bid below the ask, on the 149 shares with both prices. Five yearly quote files break off within the year: UPM's for 2018, 2019 and 2020, Stora Enso R's for 2019 and SSAB B's for 2021. We keep the complete days before each break, which for UPM cover the COVID-19 crash window (Table IA.XXXI).

The market is the OMXH25 index of Nasdaq Helsinki at its daily close. Our source for the index records no version. A coefficient of one on the dividend return of the largest shares is rejected ($p < 0.001$; Table IA.VIII in the Internet Appendix). We read it as the price index. A return from closes net of it is therefore too high by the market's dividend yield. The detection rule below uses the *register market factor*, a weighted daily return of Finnish shares at register prices. A share enters the factor on a count of its priced days over the whole sample period, so the factor uses later data. The rule was set on this factor in the manuscript named in Section II.B, and we apply it unchanged. Daily closes of the S&P 500 and the Cboe Volatility Index (VIX) measure global conditions; the VIX moves with risk aversion as well as with expected volatility (Bekaert and Hoerova (2014)). Earnings days are the announcement dates of earnings in the key developments file of S&P Capital IQ.

D. The four shocks. The shocks are found by a rule. For each week, Monday to Sunday, we compound the daily return of the register market factor. A week is flagged when its return is −6% or lower. It is also flagged when its return, divided by the standard deviation of weekly returns over the trailing 26 weeks, is −2.5 or lower. Flagged weeks no more than three weeks apart form one shock. The rule flags four shocks: the 2018 sell-off, the COVID-19 crash, the 2021 scare and

Table II
The four shocks

shock	peak	trough	fall	lowest weekly return	lowest return over trailing standard deviation	weeks
2018 sell-off	2018-08-28	2018-12-27	−17.0%	−4.7%	−2.85	18
COVID-19 crash	2020-02-11	2020-03-18	−36.6%	−16.4%	−3.96	6
2021 scare	2021-09-07	2021-10-06	−8.8%	−4.0%	−2.61	5
invasion of Ukraine	2022-01-05	2022-03-07	−23.9%	−8.8%	−3.53	10

Notes. Peak and trough are the dates of the highest and the lowest close of the OMXH25 within 60 trading days of the first flagged week. Fall is the return of the index from peak to trough. The fifth column is the lowest weekly return of the register market factor among the flagged weeks of the shock. The sixth is the lowest such return divided by the standard deviation of weekly returns over the trailing 26 weeks. Weeks counts the weeks from peak to trough.

the invasion of Ukraine, which we call the invasion. Table IA.XV applies 27 variants of the rule to the factor and to the OMXH25. Every variant finds the COVID-19 crash and the invasion, and the other two shocks depend on the series and on the thresholds.

A *crash window* is the set of trading days in the flagged weeks of one shock (Table IA.XIII in the Internet Appendix). It holds 15 trading days in the COVID-19 crash and five in each of the other three. The invasion began on Thursday 24 February 2022, and the rule flags the week after. The invasion crash window is therefore 28 February to 4 March 2022, five trading days. The *invasion weeks* are 21 February to 15 March 2022, 17 trading days, the event window of Section VI.

For comparisons over weeks we also use the *fall* of each shock. Its peak is the highest close of the OMXH25 in the 60 trading days up to the first flagged week. Its trough is the lowest close in the 60 trading days after. Table II dates the four falls. Dating a decline from peak to trough is in the spirit of Pagan and Sossounov (2003). Their bull and bear phases last at least four months on monthly prices, unless the price falls by 20% in one month. Three of the four falls here are shorter. The 2018 sell-off and the 2021 scare are flagged by the ratio alone.

The week of 26 October 2020 comes close to the rule and is not flagged (Table IA.XIV in the Internet Appendix). October 2020 is a robustness window of one week. It is pooled with the four

shocks only in one labeled row of Table IA.LXXIV and in column 5 of Table IA.XXXVIII. The fall around it runs from 12 to 29 October 2020 and is not one of the four falls.

We call *regular weeks* the weeks with no day inside one of the four falls or within 14 days of one. There are 208 of them (Table IA.XVI in the Internet Appendix). *Regular days* are the trading days outside the four falls. The regular-period beta and the account measures scored on trades use a subset, the days more than 14 days from every fall. A measure scored on regular weeks or regular days keeps the October 2020 fall.

The *regular weeks outside October 2020* are the regular weeks with no day inside the October 2020 fall or within 14 days of it. There are 201 of them, and they are the baseline of the weekly household comparisons of Section V. The *regular days outside October 2020* are the regular days less the 14 trading days of that fall. There are 1,067 of them, and they are the baseline of the daily flow comparisons of Section IV.A. We call *other trading days* every trading day outside the four crash windows. Down days are the narrower set in Section II.F. The days of a fall outside its crash window are among them, and so are the days of October 2020. They are the comparison of the tests of Section III.G and of the down-day tests of Section IV. Figure IA.1 places the four falls and the crash windows in the sample, and Table IA.XVIII in the Internet Appendix sets each fall against the S&P 500 and the VIX.

E. The most-traded shares. We form the list of shares for the price tests by a rule that uses only what was known before each window. On a formation date every listed share with a close that day is ranked by its median daily traded value over its last 250 priced days. Traded value is the close times the shares traded, from Compustat Global. A share needs volume on at least 120 of those days. The *most-traded shares* are the first 100 in that ranking. Days inside a crash window or the October 2020 window use the list formed on the last trading day before the window. Every other day uses the list formed on the last trading day of the month before. Our closes begin on 29 December 2017, so the first list forms at the end of June 2018 and the tests on either sample begin in July 2018. We call *other listed shares* the listed shares outside the list of that day. We keep *the fixed list of 55 shares* only as a labeled comparison; it was drawn up after the sample period

(Table IA.XXII). In the four crash windows 90.2% of household gross in exchange trades is in the most-traded shares (Table IA.XXIII in the Internet Appendix).

F. Filters. Three filters apply to the price tests. A large move is a daily return of 5% or more in absolute value. A daily return of 50% or more in absolute value is dropped as an error. A cell with fewer than 30 large moves is not estimated, and its table says so. A large move inside a crash window is a crash-day move. A large move on the day of a company's earnings announcement, outside the crash windows, is an earnings-day move. A return over five trading days runs to the share's fifth later priced day. For 47 of the 302 crash-day moves in the other listed shares that day is more than five trading days later (Internet Appendix). In the flow tests a down day is a day on which the OMXH25 returns −2% or less.

G. Estimation. Slopes are ordinary least squares with HC1 standard errors unless a table says otherwise. MacKinnon and White (1985) recommend HC3 in small samples, so the invasion cross-sections also report HC3 errors. Panel regressions of returns have firm fixed effects and standard errors clustered by date. A crash window holds few dates. Tests with clustered errors can reject too often when the clusters are few, and a bootstrap of the t-statistic by cluster reduces the problem (Cameron, Gelbach, and Miller (2008)). Tests inside a window therefore also carry p-values from a wild cluster bootstrap by date, with the null imposed and six-point weights (Webb (2023)).

Some tests have very few clusters: five dates in three of the crash windows, and four where the clusters are the four shocks. Wild cluster bootstrap tests perform poorly with fewer than eleven clusters, and six-point weights are shown to improve them, not to give correct size with four or five (Webb (2023)). We read a p-value from so few clusters as a guide. Time-series regressions use Newey-West standard errors (Newey and West (1987)), with five lags in the daily flow regressions and four in the weekly regressions on pandemic news. Comparisons of the weeks of a fall with regular weeks are Welch t-tests across weeks. Tables mark the 10%, 5% and 1% levels with stars, and their notes name the p-value that the stars use. Some regressors are themselves estimates: an account's traits, its same-day return slope, and a share's beta. The reported standard errors leave out their sampling error, and the direction of the bias is not known (Internet Appendix).

The headline tests of the paper form one family of 45, listed in Table IA.XX in the Internet Appendix. The list is declared in full, the tests that do not reject included, and no test enters or leaves it because of its p-value. It was not registered before estimation, and it does not hold one test per prediction. Fourteen of its tests belong to none of the six hypotheses, and 13 are named by a rule or a prediction of Section I (Table IA.XXI). The four reversal tests of H1 enter on their wild bootstrap p-values by trading date (Section I.A). The two slopes of H6 enter on the larger of their HC3 p and their wild bootstrap p (Section I.G).

Eight sets of tests that Section I names are outside the family. They are the tests of H3 and of H4, the down-day differences and the corollary of H2, and the sale-rate test of H6. In H5 they are the comparison of purchases with sales, the returns outside the COVID-19 fall and the other two pairs of the coming-back test. For H4 the family holds only the linear volume slope, which cannot separate H4 from its null (Section I.E). Replacing it once its result is known would be a choice made on the result.

We apply the step-up procedure of Benjamini and Hochberg (1995) at a 5% false discovery rate and write q for the adjusted value. Their proof covers independent tests. The tests here overlap, so q is a guide and not a guarantee. Benjamini and Yekutieli (2001) extend the result to tests that are positively dependent, and we do not claim that ours meet their condition. Of the 45 tests, 19 pass, four of them among the 13 that Section I names. A verdict that rests on one of those eight sets is marked as unadjusted where it is stated.

III. H1 and H4: Prices and Volume After a Crash-Day Move

In the most-traded shares none of the three five-day slopes of Section I.B can be told from zero. In the other listed shares the slope is negative for a day-0 start. Neither the elasticity of volume nor the skewness of returns differs detectably inside a crash window (Section III.G).

The test takes every stock-day inside a crash window with an absolute return of 5% or more. It regresses the return over the next five trading days on the day-0 return, with firm fixed effects. A

negative slope is a reversal: within a share, a more negative day-0 return went with a higher return over the next five days. Returns are at register prices. A crash window holds few trading dates, so a slope is judged by its wild bootstrap p-value by trading date. Five-day returns overlap from one date to the next, so that p-value can still overstate the evidence. Panel A also reports a bootstrap that treats the four shocks as clusters. With four clusters that bootstrap is a guide only (Section II.G).

A. The most-traded shares. Table III has three samples: the most-traded shares, the other listed shares (Section II), and the two together as all listed shares. The most-traded shares carry 90.2% of household gross trading in the four crash windows, and the other listed shares carry 4.2% (Table IA.XXIII). In the most-traded shares the five-day slope is 0.008, with a standard error of 0.094 (wild bootstrap p = 0.928). This headline test does not pass the family (q = 1.000; Table IA.XX).

Table III

Five-day slopes after crash-day and earnings-day moves, by sample and by start

	Most-traded shares: from day 0	Most-traded shares: from day 1	Other listed shares: from day 0	Other listed shares: from day 1	All listed shares: from day 0	All listed shares: from day 1
Panel A. Crash-day moves						
Crash slope	0.008	−0.047	−0.353**	−0.102	−0.232**	−0.081
(standard error)	(0.094)	(0.090)	(0.098)	(0.097)	(0.071)	(0.081)
[p, wild bootstrap]	[0.928]	[0.576]	[0.011]	[0.379]	[0.030]	[0.377]
[p, clustered by date]	[0.932]	[0.603]	[<0.001]	[0.293]	[0.001]	[0.319]
[p, wild bootstrap by shock]	[0.912]	[0.880]	[0.113]	[0.065]	[0.032]	[0.144]
Moves	289	289	302	302	591	591
Trading dates	29	29	29	29	30	30
Panel B. Earnings-day moves						
Earnings slope	0.315***	0.094	0.266	0.471**	0.263***	0.229**
(standard error)	(0.069)	(0.072)	(0.198)	(0.268)	(0.098)	(0.125)
[p, wild bootstrap]	[<0.001]	[0.233]	[0.191]	[0.024]	[0.004]	[0.032]
[p, clustered by date]	[<0.001]	[0.192]	[0.179]	[0.079]	[0.007]	[0.067]
Moves	149	149	94	94	259	259
Trading dates	110	110	82	82	175	175
Panel C. Crash slope less earnings slope						
Pooled, crash indicator (rule's test)	−0.259*	−0.160	−0.408**	−0.326	−0.407***	−0.252*
(standard error)	(0.105)	(0.106)	(0.182)	(0.227)	(0.112)	(0.134)

Table III (continued)

	Most-traded shares: from day 0	Most-traded shares: from day 1	Other listed shares: from day 0	Other listed shares: from day 1	All listed shares: from day 0	All listed shares: from day 1
[p, wild bootstrap]	[0.063]	[0.186]	[0.028]	[0.166]	[0.002]	[0.085]
[p, clustered by date]	[0.013]	[0.131]	[0.025]	[0.151]	[<0.001]	[0.060]
Pooled, no crash indicator	0.040	0.052	−0.462**	−0.369	−0.274**	−0.156
(standard error)	(0.125)	(0.117)	(0.198)	(0.238)	(0.123)	(0.134)
[p, wild bootstrap]	[0.750]	[0.645]	[0.022]	[0.130]	[0.046]	[0.267]
[p, clustered by date]	[0.747]	[0.655]	[0.020]	[0.122]	[0.026]	[0.244]
Difference of the two slopes	−0.307**	−0.141	−0.619***	−0.573**	−0.495***	−0.310**
(standard error)	(0.115)	(0.115)	(0.219)	(0.284)	(0.120)	(0.148)
[p, wild bootstrap]	[0.038]	[0.246]	[0.003]	[0.022]	[<0.001]	[0.040]
[p, clustered by date]	[0.008]	[0.218]	[0.005]	[0.044]	[<0.001]	[0.037]
Moves	438	438	396	396	850	850
Trading dates	139	139	111	111	205	205
Panel D. Crash slope on closes, raw and net of the market						
Firm effects	−0.006	−0.064	−0.258**	−0.112	−0.130*	−0.082
[p, wild bootstrap]	[0.926]	[0.418]	[0.032]	[0.365]	[0.095]	[0.427]
Firm and date effects	−0.101	−0.046	−0.414***	−0.148***	−0.296***	−0.126**
[p, wild bootstrap]	[0.598]	[0.640]	[<0.001]	[0.007]	[<0.001]	[0.024]
Market-model returns	−0.011	−0.043	−0.297***	−0.097	−0.159**	−0.063
[p, wild bootstrap]	[0.912]	[0.613]	[0.005]	[0.380]	[0.033]	[0.510]
Moves	384	384	326	326	710	710
Trading dates	26	26	29	29	30	30

Notes. Slopes of the five-day forward return on the day-0 return, across moves of 5% or more, with firm effects. Panel D uses closes, the others register prices. Day-1 returns start at the next day's price. A pooled interaction multiplies the day-0 return by a crash indicator. The rule's test adds the indicator; the last row is Panel A less B. Panel D adds date effects or a market beta. Errors cluster by trading date. Stars follow the wild bootstrap by date, not the clustered-by-date row. Earnings-day moves before July 2018 enter all listed shares only. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

Earnings days are the news benchmark. Returns after an earnings announcement drift in the direction of the surprise for weeks (Bernard and Thomas (1989)). That a large earnings-day move itself continues is this paper's expectation and not their finding. In the most-traded shares the earnings slope is 0.315 ($p < 0.001$), on 149 moves. A change of company guidance is a second news benchmark. In all listed shares the ten-day return falls after lowered guidance, and no drift is detected after raised guidance. The slope on the sign of the change is 1.82% ($p = 0.006$; Table IA.XLII), and this headline test passes the family ($q = 0.015$; Table IA.XX).

The rule's test of the second condition (Section I.A) is the pooled interaction in the regression that also has a crash indicator. In the most-traded shares it is −0.259 (wild bootstrap p = 0.063) and does not reject at 5%. This headline test does not pass the family (q = 0.118; Table IA.XX). Section III.F sets its p-value clustered by date beside the bootstrap p.

Without the indicator the interaction is 0.040, and that row is not the test. The difference of the two separately estimated slopes is −0.307 (p = 0.038). Without the one earnings-day move that changes the earnings slope most, that difference is −0.266 (p = 0.061; Table IA.XXVIII).

Table IA.XXXIII, Table IA.XXXVIII, and Table IA.XXXVII give the other horizons, the other windows, and the factor returns. None of them is the test Section I reads.

B. Other listed shares and the day-1 start. In the other listed shares the five-day slope is −0.353 (wild bootstrap p = 0.011). With the four shocks as clusters its p-value is 0.113, which we read as a guide only (Section II.G). The crash slope there is also below the earnings slope: the rule's test is −0.408 (p = 0.028).

The slope and the rule's test are headline tests. On its wild bootstrap p the slope passes the family of 45 (q = 0.025), and the rule's test does not (q = 0.063; Table IA.XX). On p-values clustered by date the q-values are 0.001 and 0.053, and the rule's test does not pass there either (Table IA.XXIX).

The comparison with earnings days is less secure than the slope. On three reduced samples the rule's test does not pass the family either (Table IA.XXIX). Two leave out moves next to an exchange holiday, where the register price series dates a leg by another rule (Section II.C). One leaves out the earnings-day move that carries much of the benchmark's slope. Without that move the rule's test is −0.297 and does not reject at 5% (p = 0.061; Table IA.XXVIII).

Most of this slope is the first day. From day 1 the slope is −0.102 (p = 0.379), and the rule's test is −0.326 (p = 0.166).

C. By shock and at ten days. Figure IA.2 traces the slope day by day on closes. In the most-traded shares the crash path is close to zero at day 5, and the earnings path lies above it from the first day. In all listed shares the crash path is below zero from the first day. The slopes by shock

and at ten days are in Table IA.XXV and Table IA.XXVI.

D. What the shares earned after a drop. A slope near zero says that no link between the size of a drop and the later return is detected. It does not say what the shares that dropped earned. On closes, in the most-traded shares, the five-day return net of the market after a drop of 5% or more is 0.00% (wild bootstrap p > 0.999; Table IA.XLIII). It is a headline test, and its q is 1.000 (Table IA.XX). At register prices the mean is 0.24% (p = 0.867; Table IA.XLV).

The longer horizons are in the same tables. Fire-sale discounts in US stocks take several quarters to reverse (Coval and Stafford (2007)), a horizon these tables do not reach.

E. Explanations of a reversal. Three explanations predict when a crash-day move should reverse. Table IV tests each in the most-traded shares, from closes, and marks the headline test of each.

Campbell, Grossman, and Wang (1993) interact the return with aggregate market volume. Llorente et al. (2002) read a positive interaction with a share's own volume as speculative trade, which this paper does not test. Nagel (2012) finds that a reversal book earns more after a high VIX. Panel C looks at the most extreme household buying, the comparison in Barber et al. (2022).

The volume test and the VIX test do not pass the family. The first-time-buyer test does, on 25 days, and it is a return after household buying, not the slope after a crash-day move.

The weekly contrarian strategy is in the Internet Appendix (Hameed, Kang, and Viswanathan (2010)), as is the split by flight-to-safety days (Baele et al. (2020)).

F. The verdict on H1. H1 sets two conditions (Section I.A): the five-day slope after a crash-day move must be below zero, and below the earnings-day slope. H1 holds in the other listed shares only. A part that fails is not shown to be zero.

Clustered by date, the p-value on the rule's test in the most-traded shares is 0.013 (Table III). On that p-value the test passes the same family with date-clustered p-values in place of the bootstrap (q = 0.030; Table IA.XXIX). The verdict follows the bootstrap p, as Section I.A sets it.

A concession predicts a negative slope that survives the day-1 start. With firm effects the other-listed slope is not detected from day 1, so the test cannot separate the first day of a concession from

Table IV

Three explanations of a reversal, most-traded shares at closes

Test	Estimate	t	p	p adjusted within the panel	Observations
Panel A. Volume: next-day return on the day's return times volume					
The share's own volume, all trading days	0.042***	6.01	<0.001		112,858 stock-days
Market volume, all trading days	−0.028	−0.75	0.452		112,859 stock-days
The share's own volume, crash-window days	0.065	1.08	0.331		3,000 stock-days
Market volume, crash-window days (headline)	−0.201	−0.81	0.475		3,000 stock-days
Market volume, October 2020 window days	−1.044	−1.57	0.141		500 stock-days
Panel B. The VIX: daily reversal-book return per ten VIX points					
Book formed on the previous five days (headline)	−0.013%	−0.63	0.530	0.530	54 months
Book formed two days back	0.040%	1.82	0.069	0.207	54 months
Book formed on the previous day	−0.054%	−0.89	0.371	0.530	54 months
Book formed two days back, with the lagged market return	0.045%	2.70	0.007		54 months
Panel C. Retail buying: twenty-day return of the tail less the rest					
Net buying, top 2%, crash-window days	1.90%	1.35	0.206	0.437	30 days
Net buying, top 5%, crash-window days	0.73%	0.75	0.474	0.710	30 days
First-time buyers from 2019, top 2%, crash-window days	−3.81%***	−3.01	0.005	0.020	25 days
First-time buyers from 2019, top 5%, crash-window days (headline)	−4.32%***	−4.14	<0.001	0.004	25 days
First-time buyers from 2019, top 5%, October 2020 window days	−0.44%	−0.12	0.903		5 days
First-time buyers from 2019, top 5%, all days	−0.07%	−0.23	0.815		987 days
Net buying, top 5%, all days	−0.69%**	−2.51	0.012		1,113 days

Notes. Panel A regresses the next-day return on the day's return and its product with log volume less its previous-year mean. Panel B regresses a reversal book's monthly mean daily return on the previous month-end VIX (Newey-West, five lags). It has no stars: its p depends on lag length. A tail is the day's top 2% or 5% of shares. All-days rows cluster by date (Panel A) or use 20 Newey-West lags (Panel C). Crash-window rows give the date-clustered t and bootstrap p; stars use that unadjusted p. The adjusted p is Benjamini-Hochberg over Panel B's three books and over Panel C's 36 cells (Table IA.XLIX). One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

bid-ask bounce. The comparison with earnings days does not pass the family, on the full sample or on any of three reduced samples (Section III.B).

In the most-traded shares the two prices agree. The midpoint slope is 0.032 (p = 0.471) and the close slope is 0.015 (p = 0.812; Table IA.XXXI).

In the other listed shares the ten-day close slope is −0.442 (p = 0.007), and the midpoint slope is −0.126 (p = 0.483; Table IA.XXXI).

G. H4: returns and volume in a crash. H4 predicts that inside a crash window the elasticity of a share's abnormal volume to its absolute return is lower, and returns across shares more negatively skewed (Section I.E). Under the null, one stable relation between volume and the size of a move holds in crashes as on other days (Gabaix et al. (2003), Gabaix et al. (2006)). Their relation holds within a stock over time, and the daily cross-section is this paper's form of it. Table V compares crash-window days with other trading days (Section II.D), one cross-section of the most-traded shares a day.

A difference in the elasticity is not detected. On closes it is 0.225 on crash-window days and 0.227 on other trading days (permutation p = 0.934). At register prices it is 0.140 and 0.124 (permutation p = 0.494; Table IA.LVI gives the other samples). With log volume in place of abnormal volume it is higher on crash-window days in all listed shares, against H4's sign (its Panel C).

Under the cubic tail law of Gabaix et al. (2006) the skewness of returns is borderline infinite. We therefore use a quantile measure of skewness, defined in the notes to Table V. It is lower on crash-window days on both price series, and neither difference is distinguishable from zero at 5%. The difference is −0.043 on closes (permutation p = 0.114) and −0.064 at register prices (permutation p = 0.061).

In all listed shares the difference in skewness has a permutation p below 0.05, at −0.057 on closes (p = 0.041) and −0.065 at register prices (p = 0.039; Table IA.LVII). By shock, in the most-traded shares on closes, it is negative in three shocks and positive in the COVID-19 crash, which holds half of the crash-window days. No single shock rejects at 5% (Table IA.LVIII).

Table V

Volume response and return skewness on crash-window days against other trading days, most-traded shares

	Crash-window days	Other trading days	Difference	t	Permutation p	Crash-window days (N)	Other trading days (N)
Panel A. Closes							
Volume slope	12.71	18.87	−6.16***	−5.61	<0.001	30	1,103
Volume elasticity	0.225	0.227	−0.002	−0.07	0.934	30	1,103
Quantile skewness	−0.002	0.041	−0.043	−1.50	0.114	30	1,103
Panel B. Register prices							
Volume slope	9.40	15.29	−5.90**	−4.45	0.013	30	1,084
Volume elasticity	0.140	0.124	0.016	0.54	0.494	30	1,084
Quantile skewness	−0.034	0.031	−0.064*	−1.92	0.061	30	1,084

Notes. Each trading day gives one cross-section of the most-traded shares. The volume slope regresses log abnormal volume on the absolute return, the elasticity on its logarithm. Abnormal volume is the day's volume over the share's median of the previous year. Quantile skewness is the 90th plus the 10th percentile of returns less twice the median, over the 90th less the 10th. Other trading days lie outside the four crash windows. The register price series dates no trade on 19 of them, each next to an exchange holiday (Section II.C). Stars use the permutation p, which reshuffles the crash-window days. Welch's t can differ from it. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

The headline family holds one test filed under H4, on a third measure: the linear slope of log abnormal volume on the absolute return. On closes the difference of −6.16 passes the family on its Welch p (Panel A of Table V; $q < 0.001$, Table IA.XX). Its permutation p is also below 0.001. The slope cannot separate H4 from its null, and it stays in the family for the reason given in Section II.G.

H. The verdict on H4. Each measure of Section I.E is read on its unadjusted permutation p, and neither test is in the family. H4 fails in the most-traded shares, the sample Section I names for it. Neither the elasticity nor the skewness is detectably lower there, on closes or at register prices (Table V). A part that fails is not shown to be zero: the tests rest on 30 crash-window days.

Of the eight tests of skewness over the four crash windows in Table IA.LVII, two reject at 5% unadjusted, both in all listed shares. Neither rejects on Welch's t, or with quartile or moment skewness on its permutation p (Panel D). On Welch's t, which the rule does not read, moment skewness at register prices rejects in both samples. We report the result as suggestive, and it is not

a verdict.

IV. H2 and H3: Who Takes the Other Side?

On the exchange, who buys in a crash differs from shock to shock. Households were net buyers on the exchange in three of the four crash windows. Foreign nominee accounts were net sellers on the exchange in the 2018 sell-off and the COVID-19 crash window, and net buyers in the 2021 scare. In the invasion crash window the sign of their flow is not identified (Section IV.A). Figure IA.4 draws each class's net purchases day by day through each window. Section IV.A describes the totals, Sections IV.B and IV.C test H2, and Section IV.D tests H3.

A. The crash windows in total. Table VI sums each class's net purchases over each crash window (Panel A) and over each fall (Panel B). These totals describe who traded; Section IV.B tests H2. In the COVID-19 crash window foreign nominee accounts sold EUR 914.8 million net on the exchange and households bought EUR 162.6 million.

Summed over every class, net purchases on the exchange are a sale of 55.3% of the foreign nominee sale (Table IA.LXIII). The other classes took up less than half of it. Only in the 2018 sell-off do foreign nominee accounts again sell on the exchange and households buy. Their sale on the exchange is under a fifth of that in the COVID-19 crash window. The largest net buyer of 2018 was foreign accounts held directly. In the 2021 scare the signs are reversed: on the exchange foreign nominee accounts bought and households sold. In the invasion crash window both totals are purchases on the exchange (qualified below). In October 2020, which the detection rule does not flag, households bought EUR 85.4 million on the exchange.

Over the whole fall two of these signs change (Panel B). Foreign nominee accounts were net sellers on the exchange in three of the four falls, the invasion fall among them. Households were net buyers on the exchange in two falls, and over the 2018 fall their net was a sale of EUR 1.6 million. Against the regular weeks before each fall, weekly household net purchases rise in the COVID-19 fall and the invasion fall, and neither rise is detected (Table IA.LXVII). Over the 2018

Table VI

Net purchases by investor class in each crash window and over each fall

	2018 sell-off	COVID-19 crash	2021 scare	invasion	October 2020 (robustness window)
Panel A. Crash windows					
Trading days	5	15	5	5	5
Households	9.6	162.6	−34.3	70.5	85.4
Foreign nominee accounts	−170.5	−914.8	148.2	130.8	−16.2
Domestic financial institutions	31.7	49.1	−15.5	43.5	15.9
General government	−3.7	123.8	−14.2	−60.3	−31.1
Non-financial corporations	14.6	74.4	−30.2	43.9	6.5
Non-profit institutions	−2.1	−25.2	−2.6	−25.1	−0.5
Foreign accounts held directly	62.8	22.0	−8.5	−19.6	−11.3
Accounts without a class	0.9	2.3	−0.5	3.0	0.1
All classes	−56.6	−505.7	42.3	186.7	48.7
Panel B. Falls, from the index peak to its trough					
Trading days	84	27	22	43	14
Households	−1.6	291.6	−43.1	189.2	109.8
Foreign nominee accounts	−937.9	−1,449.3	274.9	−1,017.5	162.2
Domestic financial institutions	433.7	66.4	46.5	292.2	31.1
General government	150.2	10.7	25.9	117.0	−96.4
Non-financial corporations	−0.2	89.9	−141.3	139.3	36.8
Non-profit institutions	−34.9	−6.9	−11.1	−50.9	−3.8
Foreign accounts held directly	89.0	157.0	−6.6	−243.4	−2.4
Accounts without a class	4.1	2.3	−0.9	5.8	−0.1
All classes	−297.5	−838.2	144.3	−568.2	237.2
Panel C. Off the exchange, shares and days of the crash windows					
Households	−47.4	66.1	−18.6	63.6	54.1
Foreign nominee accounts	−55.1	442.7	−32.9	109.8	−85.4

Notes. Net purchases (purchases less sales) of each investor class, EUR million, on the flow basis of Section II. Panel A sums the trading days of each crash window and Panel B the days of each fall, from the OMXH25 peak to its trough. The class rows add to 'All classes', which differs from zero (Section II.B; Table IA.LXII). Panel C gives two classes' net purchases in legs off the exchange with a register price, in the shares and days of Panel A. In the invasion crash window the foreign nominee total rests on legs valued at closes and its sign is not identified (Table IA.LX). October 2020 is a robustness window.

fall and the invasion fall the largest net buyer on the exchange was domestic financial institutions.

Three qualifications apply to these sums. The first is trades off the exchange in the same shares and days (Panel C). In the COVID-19 crash window foreign nominee accounts bought EUR 442.7 million net there, nearly half their sale on the exchange. With both counted they remain net sellers, and their flow also keeps its sign in the 2018 sell-off and the 2021 scare. The household flow does not in the 2018 sell-off: households sold EUR 47.4 million off the exchange, more than they bought on it.

The second is valuation. In the invasion crash window 26,699 legs carry no register price and are valued at closes, nearly all of them in one share, Nordea (Section VI.D). On legs with a register price foreign nominee accounts sold EUR 57.4 million on the exchange. Over all classes the legs valued at closes hold EUR 161.3 million more purchases than sales, so the flow basis does not hold all of their selling side (Table IA.LX). The sign of the foreign nominee flow in that window is therefore not identified. Their Nordea position rose under every count of legs (Table IA.LXI), but that is one share.

The third is the ticket limit. With every ticket counted foreign nominee accounts sold EUR 2,454.1 million on the exchange in the COVID-19 crash window and EUR 770.1 million in the 2018 sell-off (Table IA.LIX). No run of as many consecutive regular days holds a foreign nominee net sale on the exchange as large as that of the COVID-19 crash window (Table IA.LXIV). With every ticket counted, 2.4% of the runs do. The runs overlap, so this describes and does not test.

Day by day the split is common outside crashes. Foreign nominee accounts were net sellers on the exchange and households net buyers on 26.9% of regular days. In the COVID-19 crash window the share was 66.7% (Fisher's exact $p = 0.002$), a difference that passes the family of headline tests. In the other three windows the test cannot tell the share from that of regular days (Table IA.LXV). Outside the family, the 2018 sell-off also differs with every ticket counted ($p = 0.019$; Table IA.LIX), and the invasion over its whole fall ($p = 0.008$; Table IA.LXVI).

In Korea in 1997 foreign investors sold, domestic individuals bought and Korean institutions also sold (Choe, Kho, and Stulz (1999)). In US stocks in the COVID-19 crash institutions sold

and retail investors bought (Glossner et al. (2025)). In the 2018 sell-off and the COVID-19 crash window households bought and foreign nominee accounts sold on the exchange, as in Korea. Domestic financial institutions bought on the exchange in both, unlike institutions in either case.

B. Down days. H2 predicts that on a crash down day households buy more, and foreign nominee accounts sell more on the exchange, than on other down days. Table VII takes the down days of Section II.F and defines a crash down day in two ways. Panel A counts a down day as a crash down day when it lies inside a crash window, 11 days against 46. Six of the 11 are in the COVID-19 crash window and none is in the 2021 scare (Panel C). The rule of Section I is applied to this first definition. Panel B counts every down day of the four falls, 20 days against 37, and is a check. The tests are two-sided and their p-values unadjusted.

On down days households bought EUR 9.2 million a day inside the crash windows and EUR 11.1 million a day outside. The difference, EUR −1.8 million, is not detected (p = 0.704). On the second definition the difference is EUR −2.3 million (p = 0.621).

The definition matters inside a fall (Panel C). In the COVID-19 fall households bought EUR 10.0 million a day on the six down days inside the window and EUR 46.0 million a day on the two after it. Panel A counts those two as other down days.

The second part of H2 is the same comparison for foreign nominee accounts. On the first definition the difference in their daily net purchases on the exchange is EUR −33.1 million (p = 0.284). On the second it is EUR −37.5 million a day (p = 0.096). Three of the crash down days are in the invasion crash window, where their net rests in part on legs valued at closes. Without the days of the invasion fall the difference is EUR −73.1 million and is detected (p = 0.013; Table IA.LXIX). On legs with a register price it is detected on the second definition only. Both counts were chosen after the result, so the rule is not read on them.

C. Timing, and what the buying followed. Household buying comes on the day of the fall. In a regression of daily household net purchases on the OMXH25 return, the same-day slope is EUR −3.15 million per percentage point of return (t = −7.57). This result passes the family (Table IA.XX). It is absent when a trade is dated by its settlement day (Table IA.LXXII). The

Table VII

Net purchases on down days: crash down days against other down days

	inside	days inside	outside	days outside	difference	t	p
Panel A. Classes, inside the crash windows against outside them							
Households	9.2	11	11.1	46	−1.8	−0.39	0.704
Foreign nominee accounts	−33.0	11	0.1	46	−33.1	−1.12	0.284
Domestic financial institutions	7.3	11	0.8	46	6.4*	1.77	0.097
General government	3.3	11	−0.1	45	3.4	0.49	0.637
Non-financial corporations	5.8	11	−1.4	46	7.2	1.66	0.116
Non-profit institutions	−1.6	11	−0.3	46	−1.3	−0.77	0.457
Foreign accounts held directly	3.9	11	−3.1	46	7.0	1.40	0.175
Panel B. Classes, inside the four falls against outside them							
Households	9.3	20	11.5	37	−2.3	−0.50	0.621
Foreign nominee accounts	−30.7	20	6.9	37	−37.5*	−1.71	0.096
Domestic financial institutions	7.2	20	−0.7	37	7.8***	2.74	0.009
General government	3.8	20	−1.2	36	5.0	1.09	0.285
Non-financial corporations	1.9	20	−1.1	37	3.0	0.64	0.528
Non-profit institutions	−0.1	20	−0.8	37	0.7	0.54	0.596
Foreign accounts held directly	2.2	20	−4.0	37	6.2	1.24	0.220
Panel C. Households by shock, crash window against rest of the fall							
2018 sell-off	−2.1	2	8.3	2			
COVID-19 crash	10.0	6	46.0	2			
2021 scare		0	5.4	1			
invasion	15.3	3	−7.6	4			

Notes. Mean daily net purchases of each class, EUR million, on down days (the OMXH25 returns −2% or less) on which the class traded, with their number. Panel A sets down days inside the crash windows against other down days, the October 2020 fall among them. Panel B does the same for the falls. The difference is inside less outside, taken before rounding. Its Welch t-statistic and unadjusted two-sided p-value treat days as independent, though they cluster in four falls. Panel C splits each fall's household down days by its crash window, with no test. Flow basis as in Table VI. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

slope on the previous day's return is −1.20 (t = −2.20), a test outside the family that we report as suggestive. Joint tests cannot order returns and household flows: five daily lags of either series are indistinguishable from zero in the equation of the other (p = 0.528 and p = 0.240). The first is a family test and does not pass. Only a flow dated a day late shows returns leading (Table IA.LXXIII). We therefore do not claim that returns lead household flows, or that household flows predict the market.

Six family tests ask what the buying followed, and two pass. Over the four falls, household net purchases rise by EUR 0.819 million a day for each further point of loss in the shares households trade (t = 2.85). The slope passes the family, but without the COVID-19 crash it is −0.082 (t = −0.35; Table IA.LXXIV). Over the pandemic weeks, weekly household net purchases rise with news of world deaths (t = 3.21), which also passes (Table IA.LXXVI). The other four do not pass: the foreign nominee slope on the loss, and the household response to Finnish deaths, to dated pandemic events and to district cases. Over every trading day the foreign nominee slope on a drawdown is detected, outside the family (Table IA.LXXV).

H2 fails on both parts. By the rule of Section I each part is read on its difference from other down days on the first definition (Table VII, Panel A). The household difference has the wrong sign and is not detected at the 5% level (unadjusted p = 0.704). The foreign nominee difference has the predicted sign and is not detected (p = 0.284). Neither is detected on the second definition. A related test in the family, the change in the household return slope inside the crash windows, does not pass either (p = 0.196, q = 0.305; Table IA.XX, Table IA.LXVIII). A part that fails is not shown to be zero. The down-day tests are outside the family.

Other counts of legs do not all agree. The foreign nominee difference is detected without the invasion fall (Section IV.B). The change in the household slope is detected on two earlier counts, with opposite signs (Table IA.LXXI). With every ticket counted no verdict changes (Table IA.LIX).

The corollary fails. It was that domestic financial institutions sell on the exchange with foreign nominee accounts. By the rule of Section I it is judged on the sign of their mean net purchases on crash down days of the first definition, and not on a test. That mean is a purchase of EUR 7.3

million a day (Table VII, Panel A), and it is a purchase under every count of legs (Table IA.LXIX). Its difference from other down days is not detected on the first definition (p = 0.097) and is detected on the second (p = 0.009). One earlier count of legs detects it on the first definition too (p = 0.029; Table IA.LXXI). All these tests are outside the family and the days cluster in four falls, so we report them as suggestive.

D. H3: how fast the gap closes. H3 predicts that the gap between household and foreign nominee net purchases closes more slowly after a crash than after an earnings day. The gap is the absolute difference between the two classes' daily net purchases, divided by its median over all days of the sample and averaged over five-day weeks. It is unsigned, so it cannot tell a persisting imbalance from its reversal. In the week after a crash window the gap is 2.48 times that median, on average over the four shocks. One week later its excess over the median has more than halved (Table IA.LXXVII, Table IA.LXXVIII, Figure IA.3). Shock by shock the paths differ: after the 2021 scare and the invasion the largest week comes later.

In Table VIII each event contributes one change in its gap, week 6 less week 0, where weeks count from the index trough or the earnings day. Counted from the end of the crash window, as in Table IA.LXXVII, the mean change is also negative, with no test. The change is positive after two of the four troughs, and the invasion trough alone makes their mean, −0.158, negative (Panel B). At the 625 earnings days the mean change is 0.065. The benchmark is weak: an earnings day is a date on which any share has an earnings announcement, about every second trading day.

H3 fails. By the rule of Section I it fails if the change in the gap, crash less earnings day, is not positive or cannot be told from zero. The rule is read on the permutation p-value of Table VIII. The difference is −0.224, not positive, and neither test can tell it from zero (permutation p = 0.537, Welch p = 0.654). The invasion gap rests in part on legs valued at closes (Section IV.A), so its size is not identified. On legs with a register price the difference is −0.277 (permutation p = 0.447; Table IA.LXXX). The test is outside the family. Four crash events give it little power: the difference is not shown to be zero, and the gap is not shown to close as fast after a crash.

Households buy on the day prices fall, and the tests detect no more of it in a crash than on other

Table VIII

The gap between household and foreign nominee flows after crash troughs and after earnings days

	Crash troughs	Earnings days	Crash less earnings
Panel A. Mean over events			
Gap in week 0	1.63	1.25	
Gap in week 6	1.48	1.31	
Change in the gap, week 6 less week 0	−0.158	0.065	−0.224
			[0.537]
t, Welch			(−0.49)
p, Welch			0.654
Events	4	625	
Panel B. Change in the gap at each crash trough			
2018 sell-off, 27 December 2018	0.785		
COVID-19 crash, 18 March 2020	−0.141		
2021 scare, 6 October 2021	0.098		
invasion, 7 March 2022	−1.376		

Notes. The gap is the absolute difference between household and foreign nominee daily net purchases, divided by its full-sample median and averaged within five-day weeks from each event. Each event contributes week 6 less week 0; Panel A gives means. Crash events are the four OMXH25 troughs. An earnings day is a date with an earnings announcement in S&P Capital IQ. Twelve fall on a weekend, where week 0 starts on the next weekday, and four in late December 2017, where it holds only trades settled from 2018. Brackets hold the two-sided permutation p-value and parentheses the Welch t-statistic; both treat events as independent, though their spans overlap. H3 predicts a positive difference. Flow basis as in Table VI.

down days. Section V asks which households bought.

V. H5: Which Households Buy in a Crash?

H5 predicts that the way an account trades in regular weeks tells whether it is a net buyer in a crash window. We label that persistence habit. Skill is a second reading, tested separately, with two predictions. The purchases of a crash earn more than the market and than the sales. An account whose last crash purchase gained comes back in the next crash, which we call learning.

Four limits of the register apply throughout the section (Sections II.A and II.C). An account is not a person, and a traced position is not a holding or wealth. On about two thirds of the household exchange legs with both prices the recorded price is the day's close, so the price paid is not always what the account paid. The register does not show holdings through funds or of foreign stocks (Campbell (2006)). Household buying here means net buying of directly held Finnish shares on the exchange.

A. How many accounts buy. Table IA.C counts, in each crash window, the household accounts that traded and the share that bought more euros than they sold. The benchmark is the same share in each of the 201 regular weeks outside October 2020 (Section II.D), with a mean of 67.6%.

In the COVID-19 crash 75.2% of the accounts that traded were net buyers, and in the invasion crash window 77.8%. Both shares are above the regular-week mean, and 19 and 9 of the regular weeks reach them. These are counts with no test attached, outside the family of headline tests (Table IA.XX). The share was not above the mean in the 2018 sell-off or in the 2021 scare. It was highest, 81.2%, in October 2020, the robustness window, which the detection rule does not flag.

Week by week, the share in the COVID-19 crash window was 62.8% in the first week, below the regular-week mean, and 81.1% and 83.0% in the next two (Table IA.LXXXII). In that window the largest tenth of net buyers made 69.5% of the net purchases of all net buyers, and the median net buyer bought EUR 1,850 (Table IA.LXXXIII). A count of accounts therefore describes many small purchases, and the euro totals of Section IV rest on fewer and larger accounts.

B. Who buys. Table IX regresses being a net buyer in a crash window on account traits scored in regular weeks (Section II.D). The contrarian measure is the correlation of an account's weekly

net purchases with the market return in the week before, over regular weeks. A lower value is an account that buys after falls. The disposition measure follows Odean (1998): the proportion of gains realized minus the proportion of losses realized, against the price paid.

The first column pools the four crash windows and nets each account over them. Its coefficient on the contrarian measure is −0.041 per standard deviation (t = −17.83), and on the disposition measure 0.026 (t = 11.59). Both pass the family of headline tests (Table IA.XX), and so does age: older owners were net buyers less often. The model covers 39,314 accounts, one in three of the accounts that traded in a crash window. Grinblatt and Keloharju (2001) document a reluctance to realize losses among Finnish households, and about half of that effect disappears without sales from limit orders (Linnainmaa (2010)).

The strongest single trait is the same-day return slope: the slope of an account's buy indicator on the share's return on the trade day. A limit order executes only when the price moves against it, so its trade is contrarian (Linnainmaa (2010)). An account with a negative slope buys on days the share falls, as resting limit orders do. The register records no order type, so the slope may reflect limit orders, and we cannot tell. In the pooled column its coefficient is −0.099 (t = −42.27), outside the family of headline tests. Its correlation with the contrarian measure is 0.11.

The traits are scored over the whole sample period, so weeks after a crash enter them. The third column therefore scores every trait before 28 January 2020, two weeks before the index peak of the COVID-19 crash. The contrarian, disposition and slope coefficients keep their signs there, and the age coefficient cannot be told from zero. In columns (2) and (3), accounts that bought in the 2018 sell-off were net buyers more often in the COVID-19 crash, and accounts that sold in it less often. These columns and rows are outside the family.

Table IX

Who buys in a crash: account traits and the same-day return slope

trait	(1) every window	(2) COVID-19 crash	(3) COVID-19 crash, traits from before it	(4) October 2020 (robustness window)
Contrarian measure	−0.041***	−0.044***	−0.032***	−0.027***
	(−17.83)	(−16.99)	(−6.98)	(−9.72)
Disposition measure	0.026***	0.024***	0.050***	0.012***
	(11.59)	(10.17)	(10.91)	(4.83)
Same-day return slope	−0.099***	−0.096***	−0.126***	−0.082***
	(−42.27)	(−35.88)	(−27.86)	(−26.37)
Age	−0.028***	−0.011***	0.001	−0.022***
	(−11.29)	(−3.82)	(0.28)	(−7.22)
Woman	0.005**	0.012***	0.009**	−0.002
	(2.10)	(4.62)	(2.06)	(−0.68)
Years in the register	−0.013***	−0.023***	−0.003	0.001
	(−4.00)	(−8.41)	(−0.59)	(0.35)
Traded in early or late 2018 (flag)	−0.015***			
	(−4.48)			
Bought in the 2018 sell-off		0.015***	0.041***	0.006**
		(5.89)	(8.94)	(2.00)
Sold in the 2018 sell-off		−0.032***	−0.020***	−0.020***
		(−11.15)	(−3.97)	(−6.13)
Log postcode income	−0.003	0.002	0.008	−0.005
	(−0.86)	(0.41)	(1.09)	(−1.12)
Postcode higher-education share	0.003	0.003	0.004	0.004
	(0.72)	(0.72)	(0.60)	(1.03)
Accounts that traded, owners aged 18 to 100	117,429	71,339	71,339	36,333
Share with a contrarian measure	40.1%	47.6%	17.0%	64.8%
Share with a disposition measure	48.7%	54.4%	27.3%	67.8%
Accounts in the model	39,314	28,737	10,029	20,323
Share of accounts that traded in the model	33.5%	40.3%	14.1%	55.9%
R-squared	0.069	0.081	0.123	0.057
Share of accounts with a negative slope	70.7%	71.6%	74.7%	75.2%
Correlation of slope and contrarian measure	0.11	0.12	0.12	0.08

Notes. Linear probability models of net buying, on household accounts with all three traits. Coefficients are per standard deviation; HC1 t-statistics in parentheses. A lower contrarian measure buys after falls. The slope regresses the buy indicator on the share's same-day return. Traits are scored in regular weeks; column (3) scores them before 28 January 2020. Those weeks include the window of column (4), which is therefore no test. A score needs 20 regular weeks with a trade (contrarian), a gain and a loss on sale days (disposition), or 10 trades (slope). Column (1) nets the four windows and flags a trade in early or late 2018. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

Owners aged 70 or over were less likely to be net buyers in the COVID-19 crash: in the family's test the coefficient is −0.045 (t = −7.59), and it passes (Table IA.LXXXVI). On the accounts with every trait measured it cannot be told from zero (Table IA.LXXXV). The cut follows Korniotis and Kumar (2011), who find that stock selection worsens around the age of 70. Their finding is about stock selection and not about buying in a crash.

Below that age there is no single direction, and the age result of the family is a pooled one. In the COVID-19 crash alone the linear age term changes sign with the model (Table IA.LXXXIV). About 2,000 surveyed Vanguard clients, with a median age of about 60, cut their equity shares by 0.6 to 1.1 points in the same crash (Giglio et al. (2021)). Most of them did not trade. Their measure is the equity share of the portfolio held at Vanguard, and ours is net euros in directly held Finnish shares.

Estates of deceased owners test whether beliefs are needed (Section I): the owner of record holds none, though an executor still decides. Estates buy less often at all times. Against their own regular weeks, their weekly net-buyer share rose 6.3 points more than that of living owners in the COVID-19 crash (z = 5.05; Table IA.LXXXVII). That difference passes the family. The pattern is not the same in every shock: the difference is negative in the 2021 scare and cannot be told from zero in the invasion crash window. Estates therefore bound the role of the owner's beliefs in one crash and do not rule beliefs out.

The add rate is the share of traced positions with a purchase in a week. Figure IA.5 sorts it by the last register price before the week against the price paid. In the weeks of the four crash windows the add rate is highest, 5.82%, just below the price paid, and it falls with distance on both sides (Table IA.LXXXVIII). US brokerage clients are more likely to buy a stock they already own on the day after a large negative return (Barber and Odean (2008)). That sort is a one-day return and not a gap to the price paid.

The family's test is the step in the add rate at the price paid, from linear fits on each side within 10% of it (Table IA.XC). In the weeks of the COVID-19 crash window the add rate is 0.72 points lower just above the price paid than just below. Its difference from the step in regular weeks, −0.80

points (t = −6.92), passes the family. The step is not a clean break: it shrinks as the fit narrows, and placebo cut-offs, where no step is predicted, have steps too. We read the test as a change in the add rate near the price paid, and not as a step at it.

Section I set a second comparison: if domestic financial institutions add below the price paid as households do, the pattern may be market-wide. Institutions also add more near the price paid than far below it, the two groups of the right panel of Figure IA.5 (Table IA.LXXXIX). Inside each window their gap is +4.23 points (t = 2.18), against +3.98 for households, outside the family. In the COVID-19 crash alone it cannot be told from zero. The comparison is thin: it rests on 368 accounts, and the institutions' add rate is nearly flat (Table IA.LXXXVIII). It does not support reading the ordering as a household trait, and it leaves open whether the change near the price paid is market-wide.

Rebalancing toward a target weight would also buy after a fall, and Swedish households rebalance in part (Calvet, Campbell, and Sodini (2009)). The register holds no portfolio weights, so rebalancing is not tested here. The gap to the price paid is also the share's return since the purchase, so the profile is not a test against skill.

A reference to the price paid could also come from taxes: an account that sells for tax reasons realizes losses before the year ends. The share of sales below the price paid is higher in December in four of five years, by 12.3 points on average (Table IA.XCI). With five yearly differences the test does not pass the family (q = 0.212; Table IA.XX), and a tax reading of the price paid stays open.

C. Was the buying learned? Kaustia and Knüpfer (2008) find a pattern analogous to learning among Finnish IPO subscribers: those whose earlier subscriptions gained subscribe again more often. They read it as reinforcement. A pass would be consistent with skill and with reinforcement alike; a failure counts against both. Table X takes the accounts that bought in the 2018 sell-off. A purchase gained when its return over the next three months was positive. The family's test is a logit of trading in the COVID-19 crash window on that gain, with errors clustered by security. The other two pairs of consecutive crashes are outside the family, and in neither can the coefficient be

Table X

Did a gain on purchases of the 2018 sell-off bring a household back in the COVID-19 crash?

outcome in the COVID-19 crash window and model	bought and gained	sold	accounts	securities
Traded in the COVID-19 crash: average marginal effect	−0.024***	−0.123***	16,464	
	(−2.60)	(−15.84)		
Traded, accounts the register still saw: average marginal effect	−0.024**	−0.088***	15,456	
	(−2.52)	(−10.43)		
Traded, with log trades between the two windows as a control: average marginal effect	−0.030***	−0.068***	16,464	
	(−3.47)	(−9.02)		
Net buyer, accounts that traded: average marginal effect	−0.025**	−0.215***	9,395	
	(−2.09)	(−22.93)		
Traded: logit coefficient, clustered by largest 2018 purchase	−0.138*		11,129	153
	(−1.94)			
Net buyer: logit coefficient, clustered by largest 2018 purchase	−0.142		6,919	149
	(−1.17)			

Notes. Logit models on household accounts that traded in the crash window of the 2018 sell-off. The base group bought in it and its purchases, weighted by euros, were not up from the end of the window to 84 to 91 days later. Controls are log euros traded in it, log exchange trades before it, age and woman. The first four models report average marginal effects, with HC1 z-statistics in parentheses. The last two keep the 2018 buyers and report the logit coefficient, with errors clustered by the security of the account's largest 2018 purchase. The first of the two is the test in the family. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

told from zero (Table IA.XCIII).

The coefficient is −0.138 (t = −1.94), with p = 0.052 and a family q-value of 0.102 (Table IA.XX), and it does not pass. Its sign is also the opposite of learning. Of the accounts whose 2018 purchases gained, 59.7% traded in the COVID-19 crash window, against 63.2% of those whose purchases did not gain (Table IA.XCII). With errors that ignore the common return of accounts in one share, the effect is negative and can be told from zero in each of the first four rows. Those rows are outside the family, and we do not claim that a gain deters.

Seru, Shumway, and Stoffman (2010) find that Finnish investors learn mainly by ceasing to trade after poor results. The test here finds neither that pattern nor reinforcement on the margin of taking part, and attrition is not the reason (Table IA.XCII).

Danish investors who lost money in a bank failure cut their risk taking, and a failure they

only witnessed had a negligible effect (Andersen, Hanspal, and Nielsen (2019)). Table IA.XCIV asks whether a loss in one crash changes trading in the next. It compares accounts whose traced positions, priced at each share's lowest register price in the COVID-19 crash, stood more than 30% below the price paid with all other accounts. They are three accounts in five. The outcome is trading in the 12 trading days from the invasion, net of the account's trading in 10 trading days of November 2021.

Net of that benchmark the coefficient is −0.017 (t = −8.03), and it passes the family. The sign is the benchmark's: before netting the coefficient is positive. Against the first 12 trading days of each month of 2021 in regular weeks, it is negative in five of ten periods and positive in five. We therefore do not read the test as an effect of a marked loss.

D. What the shares returned after household trades. Table XI tests the first prediction of skill: it follows every household purchase and sale for 5 and 20 trading days and for 90 days, net of the OMXH25. Each return is on the share from the close of the trade day, not what the account earned. The table uses the fall of each shock from index peak to trough (Table II), not the crash window. It splits the trades into the most-traded shares and other listed shares (Table IA.XXIII).

Two return tests belong to the family, and neither passes. The first is the return in the 90 days after purchases in the COVID-19 fall, net of the index. It is +0.2% and cannot be told from zero (t = 0.08, q = 1.000; Panel C of Table XI). The index is a price index, so this level is too high by the market's dividend yield (Section II.C).

The second sets purchases by owners aged 70 or over against those of younger owners, weighted by euros (Table IA.XCIX). Its prices are closes for the fixed list of 55 shares and register prices for the rest. The older owners' purchases returned 0.3 points less, with p = 0.736 on the t-statistic clustered by security (q = 0.872; Table IA.XX). A wild bootstrap by trade date gives p = 0.415. We read every returns test on the first, as Section I.F does, because returns measured from the days of one fall overlap.

Table XI

Returns 90 days after household purchases and sales in each fall: most-traded shares, other listed shares and every listed share

	Purchases	Minus sales (pp)	Purchases, beta-adjusted	Minus sales, beta-adjusted (pp)	Accounts buying	Accounts selling	Securities
Panel A. Most-traded shares, 90 days							
2018 sell-off	−0.2%	−0.6	−0.1%	−0.2	74,173	52,717	98
	(−0.11)	(−0.70)	(−0.07)	(−0.20)			
COVID-19 crash	+0.1%	+0.4	−2.4%	−0.2	78,726	32,957	96
	(0.03)	(0.29)	(−0.90)	(−0.08)			
2021 scare	+0.6%	+1.0*	+2.2%	+0.8	50,104	28,400	97
	(0.47)	(1.75)	(1.66)	(1.47)			
invasion	−1.1%	+1.3	+1.1%	+2.0	111,081	52,682	97
	(−0.50)	(1.11)	(0.48)	(1.55)			
October 2020 fall	+14.2%***	+7.9*	+10.4%**	+7.9*	43,462	18,400	95
	(2.95)	(1.92)	(2.24)	(1.87)			
Panel B. Other listed shares, 90 days							
2018 sell-off	−5.5%*	−1.5	−5.5%*	−1.6	7,446	5,651	62
	(−1.84)	(−0.85)	(−1.82)	(−0.89)			
COVID-19 crash	+2.3%	+0.3	+2.2%	+1.8	6,658	4,348	60
	(0.54)	(0.20)	(0.53)	(1.63)			
2021 scare	−0.6%	−1.1	+0.6%	−1.1	11,842	5,595	71
	(−0.28)	(−1.35)	(0.25)	(−1.42)			
invasion	+9.3%	+3.5**	+8.9%	+3.4**	16,441	9,166	81
	(1.47)	(2.43)	(1.39)	(2.30)			
October 2020 fall	+18.5%***	−0.5	+19.6%***	−0.5	6,691	3,873	59
	(5.19)	(−0.34)	(7.25)	(−0.34)			
Panel C. Every listed share, 90 days							
2018 sell-off	−0.5%	−0.7	−0.4%	−0.3	75,846	54,278	152
	(−0.26)	(−0.81)	(−0.21)	(−0.30)			
COVID-19 crash	+0.2%	+0.3	−2.3%	−0.2	80,053	34,220	152
	(0.08)	(0.23)	(−0.87)	(−0.09)			
2021 scare	+0.5%	+0.8	+2.0%*	+0.6	54,568	30,299	163
	(0.41)	(1.52)	(1.69)	(1.20)			
invasion	−0.4%	+1.4	+1.7%	+2.1*	114,317	55,174	175

	(−0.19)	(1.23)	(0.75)	(1.68)			
October 2020 fall	+14.4%***	+6.8*	+11.0%**	+6.6	45,302	20,080	152
	(3.23)	(1.74)	(2.54)	(1.64)			

Notes. Household exchange trades of listed shares in each fall (Table II); the October 2020 rows cover the fall around the robustness window, 12 to 29 October 2020. Returns run from the trade-day close to the close 90 calendar days later; Table IA.XCV gives 5 and 20 trading days. Purchases is the mean over accounts of the euro-weighted return less the OMXH25; minus sales is purchases less sales, in percentage points (pp). A beta-adjusted return subtracts the share's prior-year beta times the register market factor's return. t-statistics in parentheses are clustered by security. The COVID-19 Purchases cell of Panel C is a test in the family. One, two, and three stars mark significance at the 10%, 5%, and 1% levels (two-sided).

With no index subtracted, purchases in the COVID-19 fall were followed over 90 days by returns 11.5 points above those of sales (t = 4.20, Table IA.XCVI). Net of the index the gap is 0.3 points, so nearly all of it is the path of the index, and one fall has one path. The gap is therefore not a test of timing skill, nor evidence that it is absent. Net of the index, the returns compare which shares were bought and do not credit the day an account chose.

Over 90 days and net of the index, neither purchases nor purchases less sales can be told from zero in any of the four falls (Panel C). The levels lie between −0.5% and +0.5%, and the differences between −0.7 and +1.4 points. Nor can either for trades inside the crash windows; beta-adjusted, the COVID-19 level is negative, and at shorter horizons cells can be told from zero in both directions (Table IA.XCVII).

Table XI and Table IA.XCV hold 40 differences for the four shocks, all outside the family and not independent of one another. Four can be told from zero: three positive, all in the fall of the invasion, and one negative, in the COVID-19 fall over 20 trading days (Table IA.XCV). Five levels of purchases can, all at 5 or 20 trading days. Four are positive, in the 2021 scare and the fall of the invasion, and one is negative, −3.6% beta-adjusted in the COVID-19 fall.

In the October 2020 fall, which is not one of the four shocks, the ninety-day return on purchases was 14.4% net of the index (t = 3.23), outside the family. Purchases less sales cannot be told from zero there (Panel C).

Among French retail clients of one European online broker, most of the return to retail liquidity provision is dissipated after about 20 days (Barrot, Kaniel, and Sraer (2016)). Net of the index,

purchases in the 2021 scare returned +2.1% over 20 trading days and +0.5% over 90 days (Table IA.XCV, Panel B, and Table XI, Panel C). The difference is not tested. A portfolio that weights each stock by its number of Robinhood holders had good timing and good alpha from mid-2018 to mid-2020 (Welch (2022)). That is a portfolio of holdings over two years, which purchases followed for 90 days cannot test.

E. The verdict on H5. H5 holds if both the contrarian and the disposition measure predict net buying in column (1) of Table IX, and holds in part if one does (Section I.F). Both do, with the predicted signs, on the unadjusted p at the 5% level, and both pass the family (Table IA.XX). H5 holds, for the one account in three that the model scores. The family's tests score the traits over the whole sample period. On traits scored before the COVID-19 crash the signs hold in column (3), for one account in seven, outside the family. The traits explain a small part of who buys (R-squared 0.069), and the tests do not separate habit from skill.

The first prediction of skill is detected in a fall if two cells are positive and can be told from zero (Section I.F). The two are purchases net of the index and purchases less sales. In Panel C of Table XI neither can, in any of the four falls, so it is not detected. Other cells can be told from zero, for purchases and against them (Section V.D). On the other weighting and split, eight more cells of the four shocks can, seven of them for purchases (Table IA.XCVIII). None of the eight is a level of purchases net of the index.

The second prediction, learning, is not detected ($p = 0.052$, $q = 0.102$; Section V.C), and its estimate is against learning. Section V holds eleven tests of the family, and seven pass (Panel C of Table IA.XX). None of the other four is read as a zero effect. The 90-day returns are imprecise, and four shocks are four observations.

VI. H6: Was Russia Exposure Priced at the Invasion?

The invasion of Ukraine is the one shock in the sample for which we hold a measure of each firm's exposure: the share of its revenue that came from Russia. Stock indexes fell in most markets at the

invasion, and against a world benchmark Finland fell on the event day and further over the next five days (Boubaker et al. (2022)). The response was strongest in countries bordering Russia or Ukraine (Boungou and Yatié (2022)). H6 has three parts: exposure is priced when the risk becomes public, it is not priced before, and households sell exposed shares more than unexposed shares.

A. The exposure score. A Russia exposure score exists for 150 listed shares. The source study states it as the "share of revenue attributable to Russia, extracted from annual-report segment disclosures", in percentage points. That study is an unpublished 2026 manuscript of the same research project, "Exposure and Exit: How the Finnish Market Traded the Russian Invasion of Ukraine". Its author line names the project, not a person, which is why the references have no entry for it. Neither it nor our files record who compiled the scores or for which fiscal year. We use the score unchanged. A share is exposed at a score of two points or more, the source study's cut: 44 are, and the other scored shares are unexposed.

Every listed share is read under the same rule, the exposed names and the rest alike. A figure is kept only when the report states Russia's share of net sales, or states Russian net sales and the group total on the same page. A 2022 or 2023 report is kept only when it states the 2021 share. Nokia's 2022 report states Russia's 2021 net sales as under 2%. Fiskars states approximately 1% of 2021 net sales. Huhtamaki's later reports state 2021 Russia sales as under 3%, beside the share already kept. Those bounds stay out of the disclosed file. For 2021 the disclosed file has YIT, Harvia, Huhtamaki and Rapala. For 2020 it has Raisio and Tikkurila. YIT states that Housing Russia was 7% of revenue.

Nokian's 20% is Russia and Asia together. Fortum states Russia segment sales of EUR 906 million and no share. Metsä Board's Russia line is the origin of its wood. The disclosed set is still shorter than eight shares, so the regressions use the copied score. A report we could not open is not recorded as a firm with no Russian sales.

The source study also sorts the 44 by the confirmation it found, and we did not check its sorting (Internet Appendix, Table IA.CI; Table IA.CII gives each score). On the reports we read, that published figure is not always a Russia-only share of net sales. For nine the score agrees

closely with a figure on Russia that the firm itself published. Seven are confirmed in kind only, by news of business in Russia or by a later exit announcement, which is not evidence on the level of the score. One is out of date and 27 are unconfirmed. The main regressions keep the 148 scored shares with a register price and so include whatever error the score has. If the score was compiled after the invasion, that error need not be unrelated to the returns.

B. When exposure was priced. Table XII regresses each share's return on its score in three windows. The window dates are the source study's. Its event window, which we call the invasion weeks, runs from 21 February to 15 March 2022. It opens on the day Russia recognized the two breakaway regions of eastern Ukraine, three trading days before the invasion of 24 February. The crash window of Section II.D, 28 February to 4 March, lies inside the invasion weeks.

The largest hat value is 0.28, so one share can move a slope. Two standard answers are HC3 errors and a wild bootstrap of the t-statistic. They disagree here, and the choice was not fixed before the slopes were read. So each slope is read on both, and a part of H6 holds only if both reject at the 5% level (Section I.G). The 34 traded shares with no score, 19.0% of gross trading, are left out (Table IA.CIII).

On register prices the slope in the invasion weeks is −0.52% per exposure point and rejects on both p-values, 0.031 on HC3 errors and 0.022 on the wild bootstrap. It enters the family on the larger and does not pass (q = 0.063; Table IA.XX). It rejects under every error estimate on file for this regression (Table IA.CIV). With the source study's count of Russia news as a control the slope is −0.43% and falls short of rejection (HC1 t = −1.93). The slopes before and after cannot be told from zero.

At closes of the 86 most-traded shares the slope in the invasion weeks is −0.76%, which rejects on HC3 errors (p = 0.006) and not on the wild bootstrap (p = 0.137).

The scored shares of the fixed list (Section II.E), the source study's sample at closes, give the larger slope before the invasion (Table IA.CVII in the Internet Appendix). On those 40 shares it is −0.83%, which rejects on the wild bootstrap (p = 0.023) and not on HC3 errors (p = 0.129).

H6 also says "not before", which needs a test of the difference between the two slopes. Table

Table XII

Return on Russia exposure by window under three error estimates

	Closes, before	Closes, invasion weeks	Closes, after	Register prices, before	Register prices, invasion weeks	Register prices, after
Return per exposure point	−0.25%	−0.76%	0.39%	−0.10%	−0.52%	−0.23%
	(−1.37)	(−2.75)	(0.62)	(−0.44)	(−2.16)	(−0.75)
p, HC1 errors	0.143	0.001	0.452	0.589	0.008	0.412
p, HC3 errors	0.170	0.006	0.535	0.663	0.031	0.452
p, wild bootstrap	[0.158]	[0.137]	[0.478]	[0.625]	[0.022]	[0.440]
Largest weight of one share in its own fitted value	0.29	0.29	0.29	0.28	0.28	0.28
Shares	86	86	86	148	148	148
Shares with exposure of two points or more	26	26	26	42	42	42

Notes. Each column regresses shares' window returns on the Russia exposure score. The windows are 3 January to 18 February 2022, 21 February to 15 March, and 16 March to 30 June. A window return is the last price over the first, winsorized at the 1st and 99th percentiles. Parentheses hold HC3 t-statistics. The bracketed wild bootstrap p imposes the null and redraws the HC1 t across shares with six-point weights. The largest weight is the largest hat value. Closes cover the most-traded scored shares of 3 January 2022. Register prices cover 148 of the 150 scored shares.

XIII tests the slope before the invasion less the slope in the invasion weeks on the same shares. Its controls are momentum and volatility of late 2021. Neither the source study nor our code records a reason for them, and this is the specification in the family. On register prices the difference is 0.46% per exposure point. HC3 errors reject (p = 0.020) and the wild bootstrap does not (p = 0.067), so the second part does not hold (q = 0.120; Table IA.XX).

Without the controls the register-price difference is 0.42% and rejects on neither p (0.061 and 0.116; Table IA.CIV). On closes the difference with controls is 0.54% and rejects on neither (0.066 and 0.145). The paper does not claim that exposure was unpriced before the invasion, and it does not claim the two slopes are equal.

Two tables in the Internet Appendix change one choice at a time at register prices: Table IA.CV the sample and Table IA.CVI the window dates. The slope in the invasion weeks rejects on both p-values in every column but two. On the 86 most-traded shares it is −0.73% and rejects on HC3 errors only (wild bootstrap p = 0.150). Without the 27 exposed shares in the sample whose score

Table XIII

Exposure slope before the invasion less the slope in the invasion weeks on one sample of shares

	Closes	Register prices
Slope before the invasion	−0.27%	−0.07%
	(−1.14)	(−0.35)
Slope in the invasion weeks	−0.81%	−0.53%
	(−3.25)	(−2.11)
Before less invasion weeks	0.54%	0.46%
	(1.84)	(2.33)
p, HC1 errors	0.027	0.008
p, HC3 errors	0.066	0.020
p, wild bootstrap	[0.145]	[0.067]
Shares	86	148

Notes. Both slopes use the shares with a return in both windows, 3 January to 18 February 2022 and 21 February to 15 March, with momentum and volatility from October to December 2021 as controls. The row 'Before less invasion weeks' regresses each share's return before the invasion less its return in the invasion weeks on the same regressors. Parentheses hold t-statistics with HC3 errors. Brackets hold the wild bootstrap p of Table XII.

is unconfirmed or out of date it is −0.62% and rejects on both.

H6 dates the news to 21 February, and the main windows start from that day's price. Started from the last price of 18 February, the slope in the invasion weeks is −0.59% and the difference with controls 0.53%. Each rejects on both p-values, the difference narrowly (wild bootstrap p = 0.048) and not on the other two bootstrap forms. The difference also rejects on both without the unconfirmed scores and with the window ended at the trough, and in no other column.

A split at the invasion, which ends the first window on 23 February, leaves a rejection in the first window only, and on the wild bootstrap alone (Section VI.H). No window slope changes without the days that the price series and the flows date differently (Section II.B).

Caldara and Iacoviello (2022) separate threats of war from acts of war. Had exposure been priced as the threat grew, exposed shares should have fallen on days the threats index rose. The regressions use 55 trading days between 1 December 2021 and 23 February 2022. Of 24 coefficients on the threats index one has a p below 5%, on closes and of the wrong sign, which is what chance gives. So short a sample would miss a small relation (Internet Appendix, Table IA.CVIII).

C. Who traded the exposed shares. Table IA.CIX sums each investor class's net purchases on the exchange, per trading day, in exposed and unexposed shares. The cells are sums of trades with no test attached, reported as description. In the invasion weeks households were net sellers of exposed shares, EUR 1.7 million a day, and net buyers of unexposed shares, EUR 3.1 million a day.

The main net buyers of exposed shares in the invasion weeks were general government and non-financial corporations. Foreign nominee accounts were net sellers on the exchange in both groups, EUR 4.3 million a day of exposed shares and EUR 22.9 million of unexposed shares. The unexposed group is larger, so the table does not say that they left exposed shares in particular.

Household selling of exposed shares fell mostly in the weeks of 14 and 21 February (Table IA.CX). The sign of foreign nominee accounts' net sales of exposed shares on the exchange over the invasion weeks rests on the last two days. Over the three whole weeks before those days they were net buyers under the ticket limit. Their tickets over the limit in the invasion weeks, which the tables do not split by week, were net sales on the exchange of EUR 102.6 million (Table IA.CXIII).

The third part of H6 has a test, which is outside the family. Table IA.CXI in the Internet Appendix compares the share of household positions with a sale on the exchange in exposed and in unexposed shares. Its window follows Section I: the four weeks from the invasion day, three trading days after the invasion weeks open. The gap of 0.322 percentage points rejects on neither p (0.056 clustered by share; 0.074 on the wild bootstrap).

Other forms of the test, also outside the family, do not all agree. With the exposed group cut at five points, 15 exposed shares, its gap rejects at the 5% level on both the clustered p (0.024) and the wild bootstrap p (0.0496). With one sale rate per share the gap rejects in the four weeks from the invasion (Welch p = 0.019) and in the four weeks before it (p = 0.032). That comparison therefore does not date the selling to the invasion. No gap rejects on its clustered p once the exposed share with most sales is left out.

D. Legs the register does not price. The 26,699 exchange legs without a register price in the invasion crash window are all but two in Nordea (Section IV.A; Internet Appendix, Table IA.LX).

Nordea has no exposure score, so its legs do not enter the flows of this section.

E. Trading ahead of the invasion. Table IA.CXI also tests whether households' sale rate was higher in exposed shares in the four weeks to 23 February 2022, whose last three days fall in the invasion weeks. The position-weighted gap of 0.105 percentage points cannot be told from zero (clustered p = 0.261; wild bootstrap p = 0.332). With one rate per share, or with positions in one share treated as independent, the gap rejects (Welch p = 0.032). Neither rejection is evidence that households kept the exposed shares.

General government was the largest buyer of exposed shares in the invasion weeks, so we ask whether it had bought them early. In the eight weeks from 27 December 2021 to 20 February 2022 it bought EUR 28.6 million net, more than in 71.7% of earlier eight-week windows (p = 0.566; q = 0.749 in Table IA.XX; Table IA.CXII). A regression over share-weeks cannot tell its net purchases over gross trading in exposed shares from those in unexposed shares (Table IA.CXIV). Neither test rejects, alone or in the family: under the ticket limit, its trading in exposed shares in those weeks cannot be told from its usual trading.

One purchase of EUR 28.5 million lies over the limit and is in neither test (Table IA.CXIII). The other classes are not tested.

F. Which households sold. From 21 February to 22 April 2022, 4.7% of 806,771 household positions in exposed shares have a sale on the exchange (Internet Appendix, Table IA.CXV).

A linear probability model regresses the sale on six account traits, each in standard deviations, with errors clustered by postcode area and share (Internet Appendix, Table IA.CXVI). The six tests are in the family; they belong to none of the six hypotheses (Table IA.XXI). Four pass. Older account holders sold more, by 0.81 percentage points per standard deviation of age (q = 0.012; Table IA.XX). Women, Swedish speakers and accounts in postcodes with more higher education sold less. Accounts first seen longer ago sold more, but a q of 0.063 makes this suggestive only, and postcode income cannot be told from zero.

The sign on age turns on the source study's positions and traits, where it is negative and rejects (Table IA.CXVII). The traits describe who sold, not why.

G. Exit announcements. Over the five trading days from twelve first announcements of an exit from Russia the mean cumulative abnormal return is −0.58%. Its 90% bootstrap interval runs from −2.56% to 1.46% (p = 0.655; q = 0.819 in Table IA.XX; Table IA.CXVIII in the Internet Appendix). Twelve events cannot tell a loss of a few percent from zero, and the test does not pass the family.

H. The verdict on H6. H6 holds in its first part only: by the rule of Section I.G it holds if all three parts hold and fails if none does. The first part holds. At register prices the slope in the invasion weeks is −0.52% per exposure point, negative, and rejects on both p-values (HC3 0.031; wild bootstrap 0.022; Table XII). On the larger of the two it does not pass the family (q = 0.063; Table IA.XX). On the most-traded shares neither price series rejects on the wild bootstrap, and with the source study's news control the slope falls short of rejection (Section VI.B).

The first part does not rest on the 27 exposed shares in the sample whose score is unconfirmed or out of date. Without them the slope still rejects on both p-values (Table IA.CV). It still rests on the score: some of the 15 exposed shares kept, the highest score among them, are confirmed in kind only, and no unexposed score is confirmed.

The second part, "not before", fails. The slope before the invasion less the slope in the invasion weeks is 0.46%, positive. Its HC3 p of 0.020 rejects and its wild bootstrap p of 0.067 does not (Table XIII; q = 0.120 in Table IA.XX). Without the controls it rejects on neither.

Outside the family, the difference rejects on both p-values in three columns (Table IA.CV and Table IA.CVI; Section VI.B). With the split at the invasion the slope of the window that ends on 23 February rejects on the wild bootstrap (p = 0.046) and not on HC3 errors (p = 0.114). The two slopes cannot be told apart (p = 0.766). That window includes the first three days of the invasion weeks, so the split does not speak to pricing before them.

The third part fails on its named test. In the four weeks from the invasion the gap in household sale rates, exposed less unexposed, is 0.322 percentage points, positive. It rejects on neither p (clustered 0.056; wild bootstrap 0.074; Table IA.CXI). With the cut at five points, and with one rate per share, the gap rejects (Section VI.C). Neither the difference in slopes of the second part

nor the gap in sale rates of the third is shown to be zero. Households were net sellers of exposed shares in the invasion weeks (Table IA.CIX), which is a sum of flows. The verdict holds for this one shock and for the shares that have a score.

VII. Conclusion

Who buys in a crash has no single answer across the four shocks. Households were net buyers on the exchange in three of the four crash windows, in one of them only while trades off the exchange are left out (Table VI). Foreign nominee accounts were net sellers on the exchange in two. The foreign nominee sale on the exchange in the COVID-19 crash window is larger than in every run of regular days of the same length (Table IA.LXIV). With tickets of EUR 25 million or more counted it no longer is (Table IA.LIX). In the other three windows their net is not larger, and no other class's euros are placed among the runs. These sums are description and not tests.

By the rules of Section I, H5 holds, H6 holds in its first part only, and H2, H3 and H4 fail. H1 holds in the other listed shares only (Table III), and that slope is not detected from day 1 or on the midpoint (Table IA.XXXI). A part that fails is not shown to be zero.

Purchases in a fall, on every listed share, are not detected to earn more than the index or than sales over 90 days (Table XI). The second part of H6 fails because the two p-values disagree: the HC3 p of 0.020 rejects and the wild bootstrap p of 0.067 does not (Table XIII).

On the midpoint no five-day reversal is detected (Table IA.XXXI). Nor is a further loss at 20, 60 or 90 days, so that price does not point to a slow concession. The register does not record which investor class started each trade (Section II.A). Data on the owners behind foreign nominee accounts would say whether the sellers on the exchange were funds under redemption. A longer register would add shocks, and would follow accounts through more than one earlier crash. In the most-traded shares no five-day reversal is detected, and household buying on crash down days cannot be told from other down days.

Disclosures

Conflict of interest. P. Joakim Westerholm: I have nothing to disclose. Abhishek Sehgal: I have nothing to disclose.

Funding. The authors received no financial support for this research.

Data. The Euroclear Finland shareholder register is proprietary; it was accessed for research and cannot be redistributed. Account identifiers are never reported and every output is an aggregate. Euroclear Finland had no right to review the paper before circulation.

Ethics review. No ethics review was required. The study uses existing register records, involves no contact with any person, and reports only aggregates in which no account can be identified.

Code. The code that produces every statistic, table, and figure, and a replication notebook that recomputes them from the register, accompany the paper.

Artificial intelligence. Generative AI tools wrote and ran analysis code and drafted and edited the text, tables, and figures under the authors' direction. No AI model classified data or produced a reported estimate: every number is computed from the source data by the replication code. The authors reviewed the results and are solely responsible for the content.

References

Andersen, Steffen, Tobin Hanspal, and Kasper Meisner Nielsen, 2019, Once bitten, twice shy: The power of personal experiences in risk taking, *Journal of Financial Economics* 132, 97–117.

Baele, Lieven, Geert Bekaert, Koen Inghelbrecht, and Min Wei, 2020, Flights to safety, *Review of Financial Studies* 33, 689–746.

Barber, Brad M., Xing Huang, Terrance Odean, and Christopher Schwarz, 2022, Attention-induced trading and returns: Evidence from robinhood users, *Journal of Finance* 77, 3141–3190.

Barber, Brad M., and Terrance 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,

785–818.

Barberis, Nicholas, and Wei Xiong, 2009, What drives the disposition effect? An analysis of a long-standing preference-based explanation, *Journal of Finance* 64, 751–784.

Barrot, Jean-Noël, Ron Kaniel, and David Sraer, 2016, Are retail traders compensated for providing liquidity? *Journal of Financial Economics* 120, 146–168.

Bekaert, Geert, and Marie Hoerova, 2014, The VIX, the variance premium and stock market volatility, *Journal of Econometrics* 183, 181–192.

Benjamini, Yoav, and Yosef Hochberg, 1995, Controlling the false discovery rate: A practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society: Series B (Methodological)* 57, 289–300.

Benjamini, Yoav, and Daniel Yekutieli, 2001, The control of the false discovery rate in multiple testing under dependency, *Annals of Statistics* 29, 1165–1188.

Bernard, Victor L., and Jacob K. Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research* 27, 1–36.

Boubaker, Sabri, John W. Goodell, Dharen Kumar Pandey, and Vineeta Kumari, 2022, Heterogeneous impacts of wars on global equity markets: Evidence from the invasion of Ukraine, *Finance Research Letters* 48, 102934.

Boungou, Whelsy, and Alhonita Yatié, 2022, The impact of the Ukraine–Russia war on world stock market returns, *Economics Letters* 215, 110516.

Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Caldara, Dario, and Matteo Iacoviello, 2022, Measuring geopolitical risk, *American Economic Review* 112, 1194–1225.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009, Fight or flight? Portfolio rebalancing by individual investors, *Quarterly Journal of Economics* 124, 301–348.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2008, Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics* 90, 414–427.

Campbell, John Y., 2006, Household finance, *Journal of Finance* 61, 1553–1604.

Campbell, John Y., Sanford J. Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *Quarterly Journal of Economics* 108, 905–939.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Choe, Hyuk, Bong-Chan Kho, and René M. Stulz, 1999, Do foreign investors destabilize stock markets? The Korean experience in 1997, *Journal of Financial Economics* 54, 227–264.

Coval, Joshua, and Erik Stafford, 2007, Asset fire sales (and purchases) in equity markets, *Journal of Financial Economics* 86, 479–512.

Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *Journal of Finance* 65, 1237–1267.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

French, Kenneth R., 2026, US and developed-Europe five factors and momentum factor, daily returns.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2003, A theory of power-law distributions in financial market fluctuations, *Nature* 423, 267–270.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2006, Institutional investors and stock market volatility, *Quarterly Journal of Economics* 121, 461–504.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, The joint dynamics of investor beliefs and trading during the COVID−19 crash, *Proceedings of the National Academy of Sciences* 118, e2010316118.

Glossner, Simon, Pedro Matos, Stefano Ramelli, and Alexander F. Wagner, 2025, Do institutional investors stabilize equity markets in crisis periods? Evidence from COVID−19, *Management Science* 71, 8623–8641.

Grinblatt, Mark, and Matti Keloharju, 2000, The investment behavior and performance of various investor types: A study of Finland's unique data set, *Journal of Financial Economics* 55, 43–67.

Grinblatt, Mark, and Matti Keloharju, 2001, What makes investors trade? *Journal of Finance* 56, 589–616.

Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.

Hameed, Allaudeen, Wenjin Kang, and S. Viswanathan, 2010, Stock market declines and liquidity, *Journal of Finance* 65, 257–293.

Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual investor trading and stock returns, *Journal of Finance* 63, 273–310.

Kaustia, Markku, 2010, Prospect theory and the disposition effect, *Journal of Financial and Quantitative Analysis* 45, 791–812.

Kaustia, Markku, and Samuli Knüpfer, 2008, Do investors overweight personal experience? Evidence from IPO subscriptions, *Journal of Finance* 63, 2679–2702.

Kelley, Eric K., and Paul C. Tetlock, 2013, How wise are crowds? Insights from retail orders and stock returns, *Journal of Finance* 68, 1229–1265.

Korniotis, George M., and Alok Kumar, 2011, Do older investors make better investment decisions? *Review of Economics and Statistics* 93, 244–265.

Linnainmaa, Juhani T., 2010, Do limit orders alter inferences about investor performance and behavior? *Journal of Finance* 65, 1473–1506.

Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang, 2002, Dynamic volume-return relation of individual stocks, *Review of Financial Studies* 15, 1005–1047.

MacKinnon, James G., and Halbert White, 1985, Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties, *Journal of Econometrics* 29, 305–325.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics* 126, 373–416.

Mitchell, Mark, Lasse Heje Pedersen, and Todd Pulvino, 2007, Slow moving capital, *American Economic Review* 97, 215–220.

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

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

Ortmann, Regina, Matthias Pelster, and Sascha Tobias Wengerek, 2020, COVID−19 and investor behavior, *Finance Research Letters* 37, 101717.

Pagan, Adrian R., and Kirill A. Sossounov, 2003, A simple framework for analysing bull and bear markets, *Journal of Applied Econometrics* 18, 23–46.

Seru, Amit, Tyler Shumway, and Noah Stoffman, 2010, Learning by trading, *Review of Financial Studies* 23, 705–739.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Webb, Matthew D., 2023, Reworking wild bootstrap-based inference for clustered errors, *Canadian Journal of Economics* 56, 839–858.

Welch, Ivo, 2022, The wisdom of the Robinhood crowd, *Journal of Finance* 77, 1489–1527.