NBER WORKING PAPER SERIES

SELLING FAST AND BUYING SLOW: HEURISTICS AND TRADING PERFORMANCE OF INSTITUTIONAL INVESTORS

Klakow Akepanidtaworn
Rick Di Mascio
Alex Imas
Lawrence Schmidt

Working Paper 29076
http://www.nber.org/papers/w29076

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

We thank Bronson Argyle, Ned Augenblick, Aislinn Bohren, Pedro Bordalo, John Campbell, Colin Camerer, Cary Frydman, Samuel Hartzmark, David Hirshleifer, Zwetelina Iliewa, Lawrence Jin, David Laibson, Ulrike Malmendier, Joshua Miller, Terry Odean, Alex Rees-Jones, Andrei Shleifer, Richard Thaler, Tuomo Vuolteenaho, Russ Wermers, conference and seminar participants at the NBER Behavioral Finance Meeting, Society of Judgment and Decision Making Conference, WFA Annual Conference, Booth School of Business, UC Berkeley, briq Institute, Caltech, ETH Zurich, Florida State, Harvard Business School, Higher School of Economics, New School of Economics, University of Amsterdam, and University of Frankfurt for helpful comments. We thank Kushan Tyagi, and especially Brice Green, for outstanding research assistance. The views expressed in the paper are solely those of the authors and do not necessarily represent the views of the institutions to which the authors are affiliated. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Selling Fast and Buying Slow: Heuristics and Trading Performance of Institutional Investors
Klakow Akepanidtaworn, Rick Di Mascio, Alex Imas, and Lawrence Schmidt
NBER Working Paper No. 29076
July 2021
JEL No. D91,G02,G12,G4,G41

ABSTRACT

Are market experts prone to heuristics, and if so, do they transfer across closely related domains—buying and selling? We investigate this question using a unique dataset of institutional investors with portfolios averaging $573 million. A striking finding emerges: while there is clear evidence of skill in buying, selling decisions underperform substantially—even relative to random selling strategies. This holds despite the similarity between the two decisions in frequency, substance and consequences for performance. Evidence suggests that an asymmetric allocation of cognitive resources such as attention can explain the discrepancy: we document a systematic, costly heuristic process when selling but not when buying.

Klakow Akepanidtaworn
University of Chicago
5757 S University Ave
Chicago, IL 60637
klakowa@uchicago.edu

Rick Di Mascio
Inalytics Ltd
12th Floor, Emerald House
15 Lansdowne Road
CR0 2BX, Croydon, UK
rick@inalytics.com

Alex Imas
Booth School of Business
University of Chicago
5897 S. Woodlawn Avenue
Chicago, IL 60637
and NBER
alex.imas@chicagobooth.edu

Lawrence Schmidt
Sloan School of Management
Massachusetts Institute of Technology
100 Main Street
Cambridge, MA
ldws@mit.edu

A data appendix is available at http://www.nber.org/data-appendix/w29076

1 Introduction

A large literature has demonstrated that market participants use heuristics and are prone to systematic biases. However, the majority of evidence comes from unsophisticated traders such as retail investors, who have been shown to be overconfident (Barber and Odean 2001), loss-averse (Larson, List, and Metcalfe 2016) and to exhibit limited attention in their trade decisions (Barber and Odean 2008; Hartzmark 2014). Comparatively little is known about the decision-making of market experts and it remains important to better understand whether they are prone to behavioral biases and, if so, the extent to which these biases affect performance.

This paper examines the decisions of sophisticated market participants—experienced institutional portfolio managers (PMs)—using a rich data set containing the entirety of their *daily* holdings and trades. Our data is comprised of 783 portfolios, with an average portfolio valued at approximately $573 million. The modal PM is hired by a single institutional client such as a pension fund in a separately managed account to generate excess returns by forming concentrated portfolios that depart substantially from their benchmarks. More than 89 million fund-security-trading dates and 4.4 million high-stakes trades (2.0 and 2.4 million sells and buys, respectively) are observed between 2000 and 2016. We evaluate performance by constructing counterfactual portfolios, and compare PMs' actual decisions to returns of the counterfactual strategy. Our data set uniquely allows us to evaluate selling decisions relative to a conservative counterfactual that assumes no skill: *randomly* selling an alternative position that was not traded on the same date.[1]

We document a striking pattern: While the investors display clear skill in buying, their selling decisions underperform substantially. Positions added to the portfolio outperform both the benchmark and a strategy which randomly buys more shares of

[1]Throughout the paper we construct counterfactuals based on *current* holdings. In a setting such as ours with active managers forming concentrated portfolios subject to short-sales constraints, evaluating a selling decision relative to a counterfactual which is unrelated to existing holdings (e.g., a benchmark index) is not an appropriate comparison. First, because portfolios depart quite a bit from the benchmark, selling the benchmark is often not feasible on the margin. Additionally, PMs may be selling in order to raise capital to buy, or because their opinion about a security has changed. An asset sold may outperform (underpeform) a benchmark index, but the sale may be optimal (suboptimal) depending on what is bought with that capital and what other assets *could* have been sold (e.g. an alternative may have gone up even more).

assets already held in the portfolio by over 100 basis points annually per dollar of purchase volume. In contrast, selling decisions not only fail to beat a no-skill *random* selling strategy, they consistently underperform it by substantial amounts. In our preferred specification, PMs forgo 80 basis points per year relative to a factor neutral, random-selling strategy.[2] We perform a wide array of robustness checks, constructing counterfactuals that match assets bought and sold on size, value, idiosyncratic volatility, prior returns, and momentum, and 'characteristic selectivity' (Daniel, Grinblatt, Titman, and Wermers 1997). We also perform sample splits to account for the potential impact of outflow pressure and price impact. Our results are robust to these alternative specifications both in terms of statistical significance and economic magnitude. Heterogeneity analyses reveal that underperformance in selling appears most prominently amongst fundamentals-oriented managers who hold more active, concentrated portfolios with higher tracking error; PMs who rely on momentum strategies exhibit the least underperformance in selling. Finally, we show that the decomposed buying and selling measures have a meaningful relationship with overall portfolio returns: outperforming (underperforming) the counterfactual when buying (selling) is strongly predictive of higher (lower) excess portfolio returns.[3]

Why would a majority of portfolio managers appear to exhibit skill in buying while at the same time underperforming substantially in selling? At face value, the fundamentals of buying and selling to optimize performance are essentially the same: Both require incorporating information to forecast the distribution of future returns of an asset.[4] Skill in both decisions requires the investor to look for relevant information and integrate it into the forecast. However, there is reason to suspect that selling and buying decisions involve different psychological processes (Barber and Odean 2013). Recent

[2]As a benchmark, active managers of mutual funds charge between 20 to 50 basis points per year in fees, depending on the size of the portfolio. The foregone returns due to poor selling are also substantially larger than the average cost differential (and net-of-fee performance differential) between mutual funds and institutional seperately managed accounts (SMAs) of the sort we have in our sample, which is around 10-35 bp per annum (Chen, Chen, Johnson, and Sardarli 2017; Elton, Gruber, and Blake 2013).

[3]Moreover, we show that poor selling affects overall performance through two channels: the first is through the stocks sold outperforming the counterfatual (direct channel), and the second is through 'fast' selling leading PMs to prematurely disgard viable investment ideas (indirect channel). As evidence for the latter, PMs rarely re-purchase assets that have left their portfolio, suggesting that sales eliminate equities from future consideration sets.

[4]As noted in Barber and Odean (2008), 'in formal models, the decisions to buy and to sell often differ only by a minus sign.'

work from the lab is consistent with this discrepancy: Buying decisions appear to be more forward-looking and belief-driven than selling decisions in an experimental asset market (Grosshans, Langnickel, and Zeisberger 2018). Indeed, anecdotal evidence from our sample points to PMs thinking differently about and allocating different amounts of effort towards the two decisions; extensive interviews suggest that they appear to focus primarily on finding the next great idea to add to their portfolio and view selling largely as a way to raise cash for purchases.[5]

We argue that the stark discrepancy in performance between buys and sells is consistent with an asymmetric allocation of limited cognitive resources towards buying and away from selling. As a result of this asymmetry, we propose that while buying will resemble decisions made through standard portfolio optimization, selling decisions are driven by a two-stage process that has elements of bounded rationality entering at both stages. First, limited attention leads PMs to constrain their consideration set to assets with extreme attributes on a salient dimension—prior returns.[6] Next, from this constrained consideration set, PMs choose to unload positions in which they have the least conviction. The latter effect can generate systematic underperformance if these positions happen to include neglected, but still viable investment ideas.[7]

As a first piece of evidence, we examine the performance trades that occur contemporaneously with exogenous events that draw investors' attention to current holdings. The release of salient and portfolio-relevant information though company earnings announcements has been exploited to study limited attention in asset markets (Menkveld 2013). Earnings announcements, like other salient news releases, not only draw attention to specific assets or asset classes, they also provide new decision-relevant information (Ball and Brown 1968) on which skilled traders are able to capitalize (Ben-Rephael,

[5]The following quotes are illustrative of this attitude: "When I sell, I'm done with it. In fact, after I sell, I go through and delete the name of the position from the entire research universe." "Selling is simply a cash raising exercise for the next buying idea." "Buying is an investment decision, selling is something else." In section 4.1, we also discuss why certain institutional features of our setting might lead PMs to prioritize buying decisions.

[6]Limited attention has been argued to generate a higher propensity to buy and sell assets with extreme returns (Hartzmark 2014; Ungeheuer 2017). Various measures of prior returns are among the most readily available pieces of information about a security – trading terminals and research platforms all highlight asset-specific prior returns – and research has shown that the salience of these return measures can affect investment decisions above and beyond the information they provide about future performance (Frydman and Rangel 2014; Frydman and Wang 2018).

[7]Barber and Odean (2008) argue for a similar two-stage trading process, writing that "preferences determine choices after attention has determined the choice set."

Da, and Israelsen 2017; Easley, Engle, O'Hara, and Wu 2008; Fedyk 2018; Hendershott, Livdan, and Schürhoff 2015). We exploit variation in earnings announcements as predetermined shifters of attention which may lead PMs to think more deliberately about positions that they would have otherwise not considered selling. Accordingly, we predict that earnings announcements broaden the PMs' consideration sets and, as a result, contemporaneous selling decisions will outperform those made on non-announcement days. Since attention is already being channeled towards purchases, the performance of buying decisions will not change on announcement days. In contrast, if the difference in buying and selling performance is driven by some fundamental discrepancy between the two decisions, e.g. differences in skill, then trades should look similar on announcement and non-announcement days.

We find that selling decisions on earnings announcement days outperform those on non-announcement days by more than 150 basis points annually. Whereas sell decisions on non-announcement days substantially underperform (similar to the overall result), on average, stocks sold on announcement dates substantially *outperform* the random-sell counterfactual. Consistent with PMs focusing on buys throughout, we do not detect a systematic difference in performance of buying decisions on announcement versus non-announcement days. These results suggest that investors do not lack the fundamental skill to sell well—in fact, the point estimates of buying and selling performance on announcement days are similar—it is just not transferred.

As further evidence for the asymmetric allocation of cognitive resources, we document that PMs are prone to use a heuristic process when selling but not when buying. PMs in our sample have substantially greater propensities to sell positions that are extreme on the salient dimension of prior returns: both the worst and best performing assets in the portfolio are sold at rates more than *50 percent* higher than assets that just under- or over-performed. Non-psychological instrumental motives do not seem to explain this pattern: results are robust to controlling for position size and holding length, and are unlikely to be explained by risk management and tax motives. The pattern persists even after the inclusion of stock-date fixed effects which absorb a number of time-varying, stock-specific unobservables. On any given day, the *same* asset is more likely to be sold from a portfolio where it exhibits relatively extreme returns than from a portfolio where its recent performance stands out less compared to other positions held. In contrast, we observe *no* similar tendency to focus on extremes on the buying side—unlike with selling, buying behavior correlates little with past returns.[8] Prior returns appear to guide the PMs' consideration sets of what assets to sell but have little effect on decisions of what to buy.

Next, we show that given the consideration set of assets with extreme returns, PMs systematically choose to sell those that they have the least conviction in. We define 'conviction' as the extent to which the PM has developed a position as an integral, active part of her portfolio. This can be measured by examining the asset's weight relative to the benchmark: assets that score low on this dimension are most likely to be ideas that the PM began to develop but neglected to do so further. We find that these 'neglected ideas' are associated with much of the underperformance in selling; moreover, they are most likely to be sold and their sales exhibit the most pronounced relationship with extreme returns.[9] In contrast, sales of 'high conviction' assets are not associated with any systematic underperformance.

Finally, we present evidence that the heuristic process outlined above is costly. Our analysis looks at within-manager variation in the propensity to sell assets with extreme returns. When the same manager becomes more likely to sell extremes (top quartile), they forgo almost 180 foregone basis points annually; in contrast, sales do not underperform when managers are least likely to sell extremes (bottom quartile). These results point to an empirical link between heuristic thinking and overall underperformance in selling. Moreover, we show that selling performance is further degraded during periods when attention devoted towards sales is likely to be stretched thin, such as when PMs are stressed (during periods when the overall portfolio is underwater) or selling in order to raise cash for buying decisions (attending to their selling choices even less). Importantly, the link between heuristic use and poor selling does not seem

[8] Since prior returns may reflect changes in relative valuations, it is not unreasonable to see a correlation between extreme prior returns and trading behavior. However, the revealed preferences of PMs' buying decisions suggest that the public signal provided by recent relative returns has little effect on PMs' beliefs about future expected returns. The lack of this correlation for buying decisions suggests that such instrumental motives are not the primary driver of behavior on the selling side.

[9] Note that the sale of assets with extreme returns may be sufficient to generate systematic underperformance if a large enough number of investors share assets that are categorized as extreme within their portfolios An (2015); An and Argyle (2016); Coval and Stafford (2007). However, as discussed in Section 2.1, this channel is unlikely to be driving the underperformance in selling within our sample as PMs tend to hold concentrated portfolios that depart quite a bit from the benchmark. To that end, we show that the correlation between assets sold in the cross-section is essentially zero.

to be driven by a persistent trait or lack of skill: consistent with the proposed attentional mechanism, differences in selling performance are associated with variation in covariates within-manager rather than in the cross-section.

Our paper contributes to the literature documenting biased decision-making in financial markets. The majority of research on investor behavior has focused on retail traders for whom daily holdings and trade data has been been more readily available (see Barber and Odean (2011) for review). The selling pattern we document is most related to the rank effect described in Hartzmark (2014). There, retail investors appear to exhibit a similar pattern in selling *and* buying behavior—unloading and purchasing assets with more extreme returns. However, it is not clear from the data whether these trading strategies are costly: This set of investors have been found to underperform the market net of fees and display a host of heuristics and biases.[10]

While prior work has documented biases amongst experts in corporate finance settings, e.g. CEOs in charge of mergers (Malmendier, Tate, and Yan 2011) or other restructuring decisions (Camerer and Malmendier 2007), substantially less research exists on the behavioral biases of expert institutional investors. A number of papers have used data on aggregate returns to demonstrate slow and inefficient incorporation of certain types of signals into asset prices (Chang, Hartzmark, Solomon, and Soltes 2016a; Giglio and Shue 2014; Hartzmark and Shue 2017; Hong, Torous, and Valkanov 2007). Although these findings highlight inefficiencies in the overall market, they cannot identify bias in expert investors per se. Other work has used the mandated release of quarterly holdings data to show the tendency of mutual funds to herd on the decisions of others (Wermers 1999), follow past prices (Griffin, Harris, and Topaloglu 2003), and display a 'reverse' disposition effect (Chang, Solomon, and Westerfield 2016b), but the coarseness of the data makes it difficult to identify biases from instrumental motives. As a result, the behavioral finance literature has mostly assumed unbiased institutional investors exploiting the behavioral biases of retail investors (Malmendier 2018). Our findings suggest that such an assumption may not be a valid one.

Our findings speak to the literature examining the performance of institutional investors. Although much of the research has argued that actively managed funds un-

[10]These biases include the disposition effect (Odean 1998), overconfidence (Odean 1999), and narrow bracketing (Frydman, Hartzmark, and Solomon 2017).

derperform the market after fees (Gruber 1996; Jensen 1968), a number of studies have presented evidence for some skill in managers' ability to pick stocks (see Wermers 2011, for a review). Using quarterly holdings data, Wermers (2000) shows that stocks held by mutual funds outperform the market on average, while net returns underperform. The majority of the difference can be explained by fees and transaction costs, suggesting skill in managers' ability to pick stocks. Kosowski, Timmermann, Wermers, and White (2006) employ a bootstrap analysis that identifies a sizable minority of funds who persistently beat the market even net of costs. Puckett and Yan (2011) argue that quarterly data may mask skill since interim trades (i.e. trades that are initiated and reversed within-quarter) can add considerable value to the fund. Finally, other work has looked at how fund characteristics are associated with performance; for example, Chen, Hong, Huang, and Kubik (2004) study the effect of fund size on returns and document a negative relationship.[11]

Our results also complement the analysis of Di Mascio, Lines, and Naik (2017) and von Beschwitz, Lunghi, and Schmidt (2017), who used the Inalytics Ltd database.[12] They study the relation between purchase/sale volume (aggregated across PMs) and market conditions to test predictions of theoretical models of optimal strategic trading with private information. They find evidence that both opening and closing trades tend to earn positive risk-adjusted returns. Crucially, the analysis of closing trades in both papers does not consider feasible alternatives based on existing holdings. As such, while the papers can speak to skill in security selection, the question of whether or not PMs over- or under-perform with respect to feasible strategies and whether decisions are potentially subject to behavioral biases is outside the scope of their investigation.

Finally, our results also contribute to the literature demonstrating heuristics and biases amongst experts in domains such as sports (Green and Daniels 2017; Massey and Thaler 2013; Pope and Schweitzer 2011; Romer 2006), judges (Chen, Moskowitz, and Shue 2016), professional forecasters (Coibion and Gorodnichenko 2015), monitoring of corporate actions (Kempf, Manconi, and Spalt 2016), and retail markets (DellaVigna and Gentzkow 2017). This line of work highlights the persistence of behavioral biases

[11] Papers have also looked at how the presence of institutional investors affects prices (Gompers and Metrick 2001), market stability (Brunnermeier and Nagel 2004), and and corporate governance (Ferreira and Matos 2008).

[12] Whereas the former also studies a sample of long only portfolios similar to ours, the latter restricts attention to a smaller set of 21 long-short hedge funds.

despite significant experience and exposure to market forces.

The paper proceeds as follows. Section 2 describes the data. Section 3 presents results on performance of buying and selling decisions, while Sections 4 and 5 present results on the use of heuristics in trading strategies and how those strategies affect performance. Section 6 concludes.

2 Institutional Sample

This section reviews the data sources which are assembled for our analysis and presents some relevant background and summary statistics on the data we observe in our sample of institutional portfolios. Further details are in Appendix sections A.1 and A.2.

2.1 Empirical Setting: Institutional Portfolio Managers

Our primary source of data is compiled by Inalytics Ltd. These data include information on the portfolio holdings and trading activities of institutional investors. Inalytics acquires this information as part of one of its major lines of business, which is to offer portfolio monitoring services for institutional investors that analyze the investment decisions of portfolio managers.[13] The majority of portfolios in our sample are sourced from asset owners—institutional investors such as pension funds who provide capital to PMs to allocate on their behalf. In these cases, we see holdings and trades related to the specific assets owned by the client. The remainder of the portfolios are submitted by PMs themselves who seek to benchmark their own performance; in these cases, data will frequently correspond with holdings and trades aggregated over multiple clients. These data are associated with a single strategy, so we do not observe assets managed by the same PMs using alternative strategies.

For purposes of this study, Inalytics assembled a dataset of long-only equity portfolios spanning from January 2000 through March of 2016. These portfolios are almost always tax-exempt, hold limited cash, and are prohibited from using leverage or taking short positions. The portfolios are internationally diversified, including data from a large number of global equity markets. Data are only available during periods for which Inalytics' monitoring service is performed. After applying a sequence of filters

[13] We will use the terms fund and portfolio interchangeably throughout our discussion.

described further in Appendix A.1, our final sample includes daily information about holdings and trades covering about 51 thousand portfolio-months of data, which are compiled from a set of 783 institutional portfolios. During this period, we have an average of just over 5 years (65 months) of data per portfolio and observe 89 million fund-security-trading date observations, 2.4 million buy trades, and 2 million sell trades. We convert all market values to US dollars at the end of each trading day.[14]

Our sample consists entirely of active managers, and, based on our review of companies' publicly available marketing materials, the vast majority of these PMs identify mispriced securities through a mix of balance-sheet analysis, conversations with management, and qualitative judgements. There are only 22 systematic, quantitative portfolios in our sample, and our results are insensitive to excluding them. Section A.2 in the Appendix provides a summary of several qualitative characteristics of our sample. Portfolios are mostly beholden to a single specific institutional client—the modal portfolio is managed by a PM on behalf of a pension fund.

Relative to the well-studied universe of mutual funds, the portfolios in our sample are highly concentrated and have large tracking error budgets.[15] These types of high-conviction PMs are typically hired by pension funds or endowments to generate alpha as a complement to other managers with more passive strategies. As a consequence, the PMs in our sample are expected to be concentrated stock-pickers with low correlation to other assets. Goyal, Wahal, and Yavuz (2020) describe the client's the search process for managers in detail, where an institutional plan sponsor puts out a request for proposals, hears presentations from three to five finalists, and finally selects a manager from amongst this group. The authors find two important factors predict which managers are selected: past returns and previous interactions.[16] Speaking loosely, the selection of highly active managers requires a record of high past returns and favorable qualitative judgements about their quality. [17]

[14]We compile data on exchange rates from three sources: Datastream, Compustat Global, and Inalytics' internal database, with Datastream being our primary source. In the vast majority of cases, at least two of these sources have identical exchange rates.

[15]In results presented in Tables 3, A.4, and A.5, we show that tracking error is not likley to be be a binding constraint.

[16]Yale, for example, invites managers for a weekend of tennis and golf to discuss investment process and philosophy (https://www.institutionalinvestor.com/article/b1k52pp0nq4f8b/The-Easily-Misunderstood-Yale-Model).

[17]The CFA Institute makes a similar claim, saying that selecting a manager is based primarily on, "a

This institutional arrangement has several important consequences. First, because the sample is made up almost entirely of institutional separately managed accounts (SMAs), the PMs are subject an entirely different flow structure compared to mutual funds. Mutual funds invest in a common set of securities on behalf of all of their shareholders, each of whom has essentially been promised daily liquidity. In contrast, managers of an SMA are typically beholden to a single client, and as a result, direct outflows are unlikely to cause forced sales. Flows in or out of portfolios tend to either be small withdrawals for cash flow needs of a pension, which are entirely predictable by the PMs, or large block withdrawals that come from firing a manager. In these latter cases, asset owners engage third party transition managers to assist with transfer of assets to new PMs. As a result, the prospect of unexpected changes in assets under management (AUM) is unlikely to significantly affect PMs' decision-making and flow-induced sales tend to be rare (relative to mutual funds).[18]

2.2 Descriptive Statistics

For each portfolio, we have a complete history of holdings and trades at the daily level throughout the sample period. Analytics collects portfolio data on a monthly basis and extends them to a daily basis by adjusting quantities using daily trades data. As a result, we observe the *complete* equity holdings of the portfolio at the end of each trading day (quantities, prices, and securities held), as well as a daily record of buy and sell trades (quantities bought/sold and prices) and daily portfolio returns, though we do not observe cash balances. Each portfolio is associated with a specific benchmark (usually a broad market index) against which its performance is evaluated. Our dataset includes an unbalanced panel of both active and inactive portfolios, with the vast majority of the data collected in real-time, suggesting that incubation and survivorship biases are not a substantial concern for our analysis.[19]

quantitative analysis of the manager's performance track record, and a qualitative analysis of the manager's investment process."

[18]Our results support this conjecture: we show that buying and selling performance is not qualitatively affected by proxies for flows. This is in contrast to the results of Alexander, Cici, and Gibson (2007) who report that flow-induced buys and sells substantially underperform discretionary trades in a sample of mutual funds. See also Chen et al. (2017); Edelen (1999).

[19]Furthermore, given that the majority of our analyses involve comparisons of stocks held with stocks traded, a number of common portfolio-specific factors which could potentially be associated with incubation/survivorship biases are differenced out via our methodology.

Table 1. Portfolio level summary statistics

This table reports summary statistics for various monthly variables from the analysis dataset of 783 portfolios. See Appendix Table A.2, Appendix A.1-A.2, and main text for additional details on variable construction and summary statistics.

Variable	Count	Mean	Std	25th	50th	75th
Assets under management ($million)	51228	573.6	1169.3	71.70	201.8	499.0
Number of stocks	51229	78.49	68.46	40.95	58.60	86.58
Turnover(%)	51223	4.10	5.76	0.93	2.54	5.03
Fraction of distinct stocks sold over all holdings (%)	51221	10.14	12.13	1.923	5.695	13.70
Fraction of distinct stocks bought over all holdings (%)	51221	14.86	17.68	3.788	8.820	19.23
Monthly benchmark-adjusted returns (%)	48786	0.22	1.77	-0.60	0.17	1.01
SD of daily benchmark-adjusted returns (%)	48041	0.35	0.21	0.21	0.29	0.43
Loading on Market	48705	0.97	0.26	0.81	0.94	1.12
Loading on SMB	48705	0.01	0.50	-0.32	-0.06	0.27
Loading on HML	48705	-0.06	0.50	-0.36	-0.07	0.22
Loading on Momentum	48705	0.05	0.34	-0.13	0.04	0.22

To complement these data we merge in external information on past and future returns (including periods before and/or after we have portfolio data). When possible, we use external price and return series from CRSP; otherwise, we use price data from Datastream. When neither of these sources are available, Inalytics provided us with the remaining price series which are sourced (in order of priority) from MSCI Inc. and the portfolio managers themselves. Since PMs in our sample have global mandates and tend to hold concentrated portfolios that depart quite a bit from the benchmark, there is a fairly low overlap between stocks held contemporaneously by different portfolios in our sample.

Using these data we construct a wide array of measures at the portfolio-time and portfolio-stock-time (position) level. Table 1 summarizes a number of key monthly fund characteristics. For brevity, we emphasize a few key attributes of our sample here and relegate discussion of summary statistics for many of these variables to Appendix Table A.2 and section A.2, respectively. All portfolios are large, and there is considerable heterogeneity in portfolio size.[20] In addition, funds differ noticeably in terms of their

[20]Note that given the size of the portolios (average AUM $533 million), trades in the 1 to 5 million dollar range are unlikely to face issues related to price impact, such as difficulty exiting from a position. These magnitudes are in sharp contrast with trades of big mutual funds (e.g. between $25 and $200 billion dollars) which have been examined in prior work (Coval and Stafford 2007; Lou

trading activity levels. Average monthly turnover is about 4 percent of assets under management, but some funds are considerably more active in their trading behavior than others (the standard deviation is 5.7 percent).

While holding fairly diversified portfolios (average number of stocks is about 78 with a standard deviation of 68), funds in our sample remain active, with positions that deviate substantially from their benchmarks. This is more concentrated than the average mutual fund, which holds a median of 92 stocks during our sample period Wermers, Yao, and Zhao (2012). The average tracking error—the standard deviation of the difference between the daily portfolio return and the benchmark—is about 0.35 percent per day, or about 5.7 percent on an annualized basis. As discussed in the previous subsection, this corresponds to portfolios that are substantially more concentrated and have larger tracking error budgets than the typical passive or retail mutual fund. On average, a manager will initiate a sell trade for about 10 percent and a buy trade for about 15 percent of the stocks in his/her portfolio each month. We also characterize fund portfolios in terms of factor exposures by computing rolling Carhart 4-factor regressions (using the prior 1 year of daily data with the Fama-French international factors), adjusted for asynchronous trading.[21] The average market beta is about 1, and average exposures to the SMB, HML, and Momentum factors are fairly close to zero.

The average fund in our sample beats its respective benchmark by about 0.22 percent per month, or 2.6 percent per year. This, in conjunction with the fact that funds' average betas are close to 1 and have little average exposure to the three other priced risk factors, suggests that these managers are highly skilled.[22] We view the positive selection of managers in our sample as an advantage when studying expertise and heuristic use: The population we examine is clearly skilled, and thus identifying biased behavior is likely a lower bound when generalizing the results.

In Appendix A.2, we summarize a variety of additional attributes of our sample. Relative to US mutual funds, PMs in our sample have a larger share of international holdings. Managers have concentrated holdings and long holding periods. Specifically,

---

2012).

[21]Following Dimson (1979), we adjust for asynchronicity by including one lag and one forward return of each factor. We use these adjustments throughout when estimating daily factor loadings.

[22]Prior work has demonstrated that a subset of insitutional investors do persistently outperform the market and generate alpha (Kosowski et al. 2006).

the average holding length is at least 485 calendar days (or about 15 months).[23]

Differences from other datasets This sample offers some unique opportunities to study expert decision-making relative to other datasets in the literature. First, in contrast to the Large Discount Brokerage dataset of Barber and Odean (2000), which features portfolio holdings and trades of individual retail investors, our data include complete portfolio and trade-level detail for a population of professional investors managing large pools of assets.[24] Illustrative of this distinction, Barber and Odean (2000) report that the value of the average portfolio is $26,000 and that the *top quintile* of investors by wealth had account sizes around $150,000—the average portfolio in our sample is almost four *thousand* times larger. Second, unlike other datasets which characterize institutional portfolios such as mutual fund portfolio holdings reports and 13-F filings (e.g. Chang et al. (2016b)), we are able to observe portfolio holdings and changes to those holdings on a *daily* level. This allows us to test hypotheses on individual decision-making that are infeasible with quarterly data. Additionally, in the other most widely used database with institutional trading information—the Abel Noser/ANcerno database (for an overview, see Hu, Jo, Wang, and Xie 2018)—researchers often do not observe all trades made by a given institutional investor and tend to lack timely information on portfolio holdings. Finally, unlike passive and retail mutual funds, our sample is comprised of professional managers who are mostly beholden to a single client (e.g. pension fund) and are paid to generate alpha through concentrated portfolios that depart from the respective benchmark. These differences are highlighted in Appendix A.2, which compares characteristics of portfolios in our sample to those in other datasets. As discussed further in Section 3, one notable consequence of this is that PMs are not very constrained by client-mandated flows or tight tracking error budgets (in contrast to mutual funds, e.g. Alexander et al. (2007) or earlier studies of pension funds, e.g. Del Guercio and Tkac (2002)). This suggests that trading decisions are largely up to the managers' discretions.

[23] This is an underestimate: our measure is right censored since we can only measure holding length beginning with the first portfolio snapshot.

[24] See Barber and Odean (2011) for a survey of studies using this and other similar datasets.

3 Overall Trading Performance

Having described the basic properties of our dataset, we now examine performance of PMs' decisions. We begin by discussing our methodology for computing counterfactual portfolio returns and, accordingly, value-added measures. We then present the first of our empirical results, which correspond to the average value-added (or lost) associated with managers' active buying and selling decisions.

3.1 Constructing counterfactuals

Given that PMs in our sample tend to hold limited cash positions and are not generally permitted to use leverage, the primary mechanism for raising money to purchase new assets is selling existing ones. Since the portfolios already include stocks that are carefully selected to outperform their respective benchmarks, the choice of which asset to sell is far from innocuous. Precisely if managers' use of information makes them skilled at picking stocks, biased selling strategies have the potential to cannibalize existing, still viable investment ideas and to reduce the potential value for executing new ones. It is therefore important to construct the appropriate benchmark to serve as the counterfactual for evaluating buying and selling decisions. In contrast, we would expect unskilled investors neither to gain nor lose money (on a risk-adjusted basis) by relying on a simple rule of thumb for selling existing positions.

The fact that we observe daily transactions allows us to compare observed buy and sell decisions to counterfactual strategies constructed using concurrent portfolio holdings data. Our measures correspond to the relative payoffs from two hypothetical experiments: one for evaluating buying decisions, and one for evaluating selling decisions. For evaluating buys, suppose that we learned that a manager was planning to invest \$1 to purchase a stock tomorrow and to hold it for a fixed period of time. We then suggest that instead of executing the proposed idea, the PM invests that money in a randomly selected stock from her other holdings. Likewise, we can suggest that instead of selling a particular stock, the PM randomly sells one of her other positions to raise the same amount of cash, holding the stock that was to be sold for the same period.

Since our conditioning information was also available to the manager and our strategies are always feasible on the margin, one would expect decisions of a skilled PM to outperform our counterfactual.[25] Note that the expected payoff from the counterfactual strategy (integrating out uncertainty about which stock is randomly selected) corresponds to the equal-weighted mean of realized returns across stocks held in the portfolio, which we denote by R_{hold}. Similar results obtain if we use lagged portfolio weights to construct a value-weighted mean instead. The manager's decision adds value relative to the random counterfactual if $R_{buy} - R_{hold} > 0$ in the first experiment and if $R_{hold} - R_{sell} > 0$ in the second experiment. Following this logic, we compute $R_{buy} - R_{hold}$ and $R_{hold} - R_{sell}$ over horizons ranging from 1 week to 1 year for all buy and sell trades, respectively, to characterize the value-added associated with each.

If the return measure of interest is a cumulative return over the relevant horizon, these measures capture the impact on benchmark-adjusted returns associated with switching from the counterfactual to the actual trade, per dollar transacted. According to our discussions with clients and managers these relative returns are the primary metric by which our PMs are evaluated.[26] That said, given that stocks should earn ex-ante compensation for systematic risk exposures, our preferred measures will use "factor-neutral" stock returns, all of which should earn the same expected return per period. We estimate stock-level exposures to the Fama-French/Carhart 4 factors using pre-trade data, then use them to adjust stock-level returns to hedge ex-ante differences in exposures. Further details about this method are presented in Appendix A.3.

We aggregate across trades and conduct inference as follows. If multiple stocks are bought or sold on a given day, we average these measures for buy and sell trades separately. Since not all funds trade every day and are not necessarily present throughout our sample period, this averaging procedure yields a portfolio-day unbalanced panel. Because some funds trade much more frequently than others—see the dispersion in

[25] In contrast, selling the benchmark to finance a purchase, which implicitly corresponds to the counterfactual in measuring benchmark-adjusted returns of stocks sold, is likely infeasible for a long-only manager who, similar those in our sample, holds a portfolio with a small (relative to the number of assets in the benchmark) number of high active share positions and thus deviates substantially from the benchmark. Purchasing the benchmark is feasible on the other hand. We discuss concerns about transaction costs and price impact below.

[26] These measures also have an alternative interpretation to the extent that buy and sell trades are not motivated by a desire to change a portfolio's systematic risk exposures. In that case, we would expect loadings on priced factors of the assets being traded and the hold portfolio to be similar and these measures would also correspond to differences in risk-adjusted returns (i.e., "alpha"). Generally, we find similar results regardless of our approach for correcting for systematic risk.

monthly turnover in Table A.1—we weight observations inversely to a measure of trading frequency.[27] We compute double-clustered standard errors using a panel estimator similar to Hansen and Hodrick (1980). It allows individual fund time series to be serially correlated and additionally allows these value-added measures to cross-sectionally correlated across PMs who place trades at similar times. While the autocorrelation correction is likely required since our use of long horizon returns potentially introduces an overlapping structure in the error term of each fund's value-added time series, the latter adjustment turns out to be small. For further details, see Appendix A.4.

3.2 Overall performance relative to counterfactuals

Figure 1, Panel A shows average factor neutral counterfactual returns for buying decisions. As will turn out to be the case across all of our specifications, we find strong evidence that purchases add value relative to the random buy counterfactual, $R_{buy} - R_{hold}$. The average stock bought outperforms the counterfactual by nearly 120 basis points over a one year horizon.

Figure 1, Panel B presents average value-added, $R_{hold} - R_{sell}$, for sell trades relative to a factor neutral counterfactual. Recall that our measure is already signed so that positive values indicate that a trade helps portfolio performance relative to the counterfactual, and negative values point to a trade hurting performance. In stark contrast to Panel A, these estimates suggest that managers' actual sell trades *underperform* a simple random selling strategy. Magnitudes are quite substantial: The value lost from an average sell trade is indistinguishable from zero at a 1 month horizon but on the order of 80 basis points at a 1 year horizon relative to a simple counterfactual which randomly sells other stocks held on the same day.

To help assess magnitudes associated with these estimates, Table 2 links our performance measures with portfolios' benchmark-adjusted returns. Specifically, Inalytics provides a daily estimate of the benchmark-adjusted return on each portfolio, and, for each portfolio, we compute the average of this return compounded over the next

[27] We weight observations inversely to the number of trading days in a calendar year that the fund buys and sells a stock. This measure allows for an easier comparison across buys and sells, since we use the same weights across both types of trades. We obtain similar results when we instead weight inversely to the number of days with trades (buys or sells), which ends up assigning a higher weight to funds with higher turnover.

Figure 1. Post-trade returns relative to counterfactual

This figure presents differences between average factor-neutral returns of stocks bought/sold and those of random buy/sell counterfactual strategies for buy and sell trades. The bracket at the top of each bar is the 95% confidence interval of the point estimate at each horizon. Confidence intervals in brackets are computed using double-clustered standard errors, calculated as described in Appendix A.4.

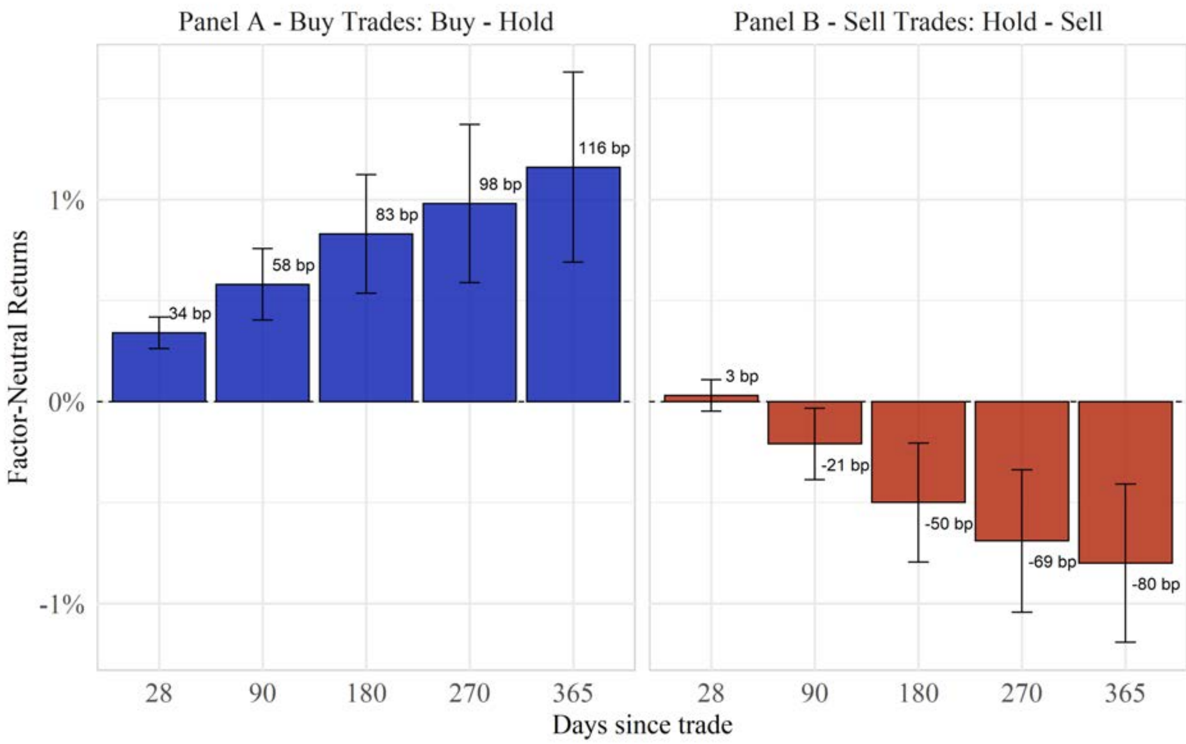

month, quarter, or year. We regress these measures on averages of our trade-based buy and sell performance measures over the same horizons.[28] Importantly, both the buying and selling performance measures are strongly linked with benchmark-adjusted performance. This is consistent with poor selling imposing nontrivial opportunity costs on the PM's potential performance. Given that PMs in our sample have fairly long holding periods, coefficients on our performance measures increase as the horizon increases; the coefficients on both measures range between 0.22-0.25 over a 1 year horizon. In the right panel, we see that results are unchanged if we control for the average difference between the hold portfolio and the benchmark, which is somewhat similar to an overall attribution measure that can be computed with a single snapshot of holdings.

[28]Similar results obtain if we weight portfolios by the length of time for which we observe data or use value-weighted performance measures in place of our baseline.

Table 2. Linking buying and selling-based performance measures with fund benchmark-adjusted performance

This table presents the association between factor-adjusted portfolio returns in excess of the benchmark and our measures of buying and selling performance. A fund's average performance at holding, buying, and selling (by our measures) is highly associated with the magnitude by which they beat their benchmark in the cross-section, even after adjusting for exposures to value, size, market, and momentum factors. Buy, sell, and hold portfolios are adjusted for factor exposures to value, momentum, size, and market, calculated for the forward-looking horizons, the Buy - Hold, Hold - Sell, and Hold - Benchmarks are computed for each fund date and then averaged across the full sample by fund. Fund returns are similarly calculated on a rolling basis and then averaged over the full period. Fund returns are not adjusted for factor exposures, and neither are benchmarks. Benchmarks are assigned based on what the fund manager has chosen for use with the Inalytics service.

	Dependent variable: Average benchmark-adjusted portfolio return over					
	1 Month	1 Quarter	1 Year	1 Month	1 Quarter	1 Year
	(1)	(2)	(3)	(4)	(5)	(6)
Hold - Benchmark				0.123***	0.131***	0.110***
				(0.031)	(0.029)	(0.026)
Buy - Hold	0.061**	0.120*	0.250***	0.057**	0.100*	0.245***
	(0.022)	(0.049)	(0.074)	(0.021)	(0.046)	(0.069)
Hold - Sell	0.094***	0.086**	0.227**	0.082***	0.079**	0.239***
	(0.021)	(0.029)	(0.072)	(0.021)	(0.029)	(0.072)
Observations	750	750	749	750	750	749
R^2	0.053	0.041	0.073	0.105	0.096	0.107

Note: *p<0.05; **p<0.01; ***p<0.001

Intriguingly, estimated coefficients on our trade-based performance measures are larger than the Hold-Benchmark coefficient at longer horizons.

As with any performance evaluation exercise, a number of questions emerge about the robustness and interpretation of our main results. Accordingly, we have conducted a battery of additional tests to address several of potential concerns. Table 3 reports results from a subset of these tests in which we consider a variety of alternative ways of constructing our performance measures, either by restricting the set of holdings/trades used in the analysis or by changing the way in which we construct the relevant long-short portfolios. See Appendix A.5 for a more detailed discussion and additional tests.

First, we illustrate the implications of altering the approach for selecting securi-

Table 3. Post-trade returns relative to counterfactual, overall and robustness checks

This table presents the average value added measures (post-trade returns relative to a random sell counterfactual) for buy and sell trades for a variety of specifications. See text and Appendix A.5 for additional details.

	Performance	Panel A: Buy					Panel B: Sell				
	Measure	28	90	180	270	365	28	90	180	270	365
(1)	Factor-Neutral	0.34	0.58	0.83	0.98	1.16	0.03	-0.21	-0.50	-0.69	-0.80
		(0.04)	(0.09)	(0.15)	(0.20)	(0.24)	(0.04)	(0.09)	(0.15)	(0.18)	(0.20)
(2)	Value-weighted*	0.30	0.63	0.99	1.23	1.45	-0.02	-0.31	-0.71	-0.92	-1.04
		(0.05)	(0.11)	(0.23)	(0.32)	(0.39)	(0.04)	(0.10)	(0.21)	(0.26)	(0.29)
(3)	Omit recently added stocks	0.35	0.63	1.94	1.14	1.39	0.01	-0.26	-0.61	-0.78	-0.91
	from hold portfolio*	(0.04)	(0.09)	(0.15)	(0.21)	(0.27)	(0.04)	(0.09)	(0.14)	(0.17)	(0.20)
(4)	Unadjusted Return	0.39	0.71	0.92	1.08	1.27	0.06	-0.12	-0.42	-0.59	-0.70
		(0.05)	(0.10)	(0.18)	(0.22)	(0.27)	(0.04)	(0.10)	(0.17)	(0.20)	(0.23)
(5)	DGTW Adjusted	0.30	0.45	0.43	0.48	0.55	0.04	-0.07	-0.31	-0.48	-0.49
		(0.04)	(0.09)	(0.15)	(0.19)	(0.24)	(0.04)	(0.09)	(0.13)	(0.17)	(0.20)
	Match counterfactual by:										
(6)	Benchmark-adjusted return quintile*	0.32	0.57	0.85	1.09	1.32	0.03	-0.19	-0.45	-0.66	-0.79
		(0.05)	(0.10)	(0.16)	(0.22)	(0.28)	(0.05)	(0.09)	(0.12)	(0.17)	(0.20)
(7)	Idiosyncratic volatility quintile*	0.35	0.59	0.87	1.02	1.10	-0.01	-0.29	-0.61	-0.86	-0.98
		(0.05)	(0.08)	(0.15)	(0.20)	(0.24)	(0.05)	(0.09)	(0.15)	(0.18)	(0.21)
(8)	Position size quintile*	0.26	0.41	0.56	0.72	0.93	0.05	-0.21	-0.52	-0.71	-0.86
		(0.04)	(0.09)	(0.13)	(0.19)	(0.25)	(0.04)	(0.09)	(0.13)	(0.16)	(0.19)
(9)	Omit 5% Smallest*	0.19	0.36	0.35	0.46	0.46	-0.08	-0.19	-0.69	-0.71	-0.86
		(0.04)	(0.09)	(0.14)	(0.19)	(0.24)	(0.04)	(0.10)	(0.14)	(0.18)	(0.21)
(10)	Omit 5% Most Illiquid*	0.35	0.61	0.88	1.06	1.28	0.02	-0.25	-0.50	-0.57	-0.73
		(0.05)	(0.10)	(0.15)	(0.20)	(0.25)	(0.05)	(0.10)	(0.15)	(0.20)	(0.23)
(11)	Omit returns for 7 days post-trade*	0.18	0.40	0.67	0.81	0.91	-0.05	-0.31	-0.61	-0.80	-0.90
		(0.03)	(0.09)	(0.15)	(0.19)	(0.24)	(0.04)	(0.09)	(0.15)	(0.19)	(0.22)

* Indicates that results are factor-neutral (where not otherwise specified).

ties for our no-skill random selling counterfactual. The first approach in row 2 uses weights proportional to lagged portfolio values in place of equal weights for the hold portfolio—analogous to selecting a security at random with a probability proportional to size or selling all holdings pro-rata—and weights multiple buy/sell trades on the same day proportionally to transaction size. The second approach in row 3 uses additional available information to construct the counterfactual. It is rare for PMs to sell stocks that were recently added to the portfolio (see Appendix Table A.8). Thus, we construct a counterfactual hold portfolio which excludes stocks which are in the bottom quintile of the distribution of holding length at each date. Both adjustments yield modest increases in the buying performance measure and modest decreases in the selling performance measure.

Second, a natural concern is that stocks traded tend to have above average exposures to systematic factors linked with expected returns relative to stocks held, meaning that our estimates could be driven by risk compensation rather than skill. Indeed, any performance evaluation exercise involves a joint hypothesis about whether the benchmark is correct. While our long-short construction differences out many systematic exposures which are common between stocks traded and held, there is always a possibility that we are sorting on characteristics which are already known to drive expected returns, presumably because they proxy for priced systematic factors. Specification 4 of Table 3 also reports estimated return measures from repeating the analysis in Figure 1 using raw cumulative returns. Results are quite similar between our baseline and this unadjusted specification, so the explicit role of the Carhart risk adjustment procedure on our overall estimates is small or nonexistent. We also leverage financial information from the Worldscope database to construct characteristics following the approach of Daniel et al. (1997) (row 5), matching counterfactual trades based on joint quintiles by size, book to market, and momentum characteristics. Rows 6-7 match counterfactuals by quintiles of characteristics which are potentially correlated with propensity to sell assets (as we show below)—prior returns over the previous quarter and idiosyncratic volatility.[29] While each exercise changes magnitudes to some modest extent, results

[29]Further, to allow for the possibility that PMs adjust position sizes partly to reflect ex-ante expectations about benchmark-adjusted returns from systematic factors, we also demonstrate that results are quite similar if we select the counterfactuals within the same quintile by position size (row 8). We discuss the role of conviction further in section 4.5 below. See also Appendix A.5 for many additional tests. Appendix A.9 shows that we find similar results if we express long-short returns

are quite consistent across all specifications: both the annualized outperformance of buys and underperformance of sells are always statistically significant and economically meaningful in magnitude.

Third, we consider the potential role of unobserved transaction costs in driving our results. Perhaps the hold portfolio includes many illiquid stocks which would be too costly (whether due to direct trading costs or indirect costs from price impact) to sell. In this case, our estimates of value-added relative to a feasible counterfactual potentially neglect any differences in these transaction costs between the security traded and our proposed alternative. To address this concern, we have conducted a battery of tests which drop the smallest (by market capitalization) or least liquid (by Amihud liquidity measures) stocks in our sample (see e.g., rows 9-10) and/or form counterfactuals matched within-portfolio by quintiles of these measures. Similar results obtain if we omit the smallest or least liquid stocks from counterfactual portfolios only.

In addition, it is possible that PMs' actual buy and sell trades could be sufficiently large so as to generate price impacts. In that case, we could measure a change in the price of the security actually traded but fail to incorporate a corresponding change which would have been induced in the counterfactual portfolio if it was traded instead. Mean reversion in prices would tend to bias both performance measures downward. Our approach to addressing this concern is simple: since these direct price impacts are largely transitory, we can compute performance starting several days after the transaction takes place, which allows for transitory effects to dissipate (row 11).[30] Our results are not meaningfully affected across these specifications.[31]

in the modal currency of each portfolio rather than USD. In addition, Appendix A.10 conducts a complementary analysis in which we form long-short calendar time portfolios on the basis of signals which are constructed using our trading data. These portfolios, which *buy* stocks recently sold by PMs and short sell stocks recently held, earn very high Sharpe ratios due to substantial benefits of diversifying across (largely uncorrelated) positions traded by different PMs. For several reasons which we discuss further there, we prefer our baseline approach to the calendar time portfolio method. However, these results are consistent with our main analysis, and these portfolios earn Sharpe ratios which are consistent with the degree of diversification implied by standard errors on our baseline point estimates.

[30] Further note that the timing over which underperformace of selling, which is most pronounced at longer horizons, manifests is inconsistent with a simple price impact explanation.

[31] While these analyses help to rule out a number of concerns related to transaction costs and potential price impacts, we acknowledge some potential limitations of these tests. First, we do not observe intraday data, so our analysis cannot capture profits which were captured at a very high frequency (see, e.g., Puckett and Yan 2011). As a result, while the daily frequency of our data does allow us to rule out favorable intraquarter trading as driving our results, we are unable to directly test

Finally, Appendix A.5 also presents several empirical tests which address the potential role for position size limits and tracking error budgets in driving our main results. Once again, we find little evidence consistent with PMs facing constraints which "force" sales which underperform a random selection of existing holdings. Perhaps most convincingly, we find almost no correlation between the "Buy - Hold" portfolio and the "Sell - Hold" portfolio (between 1.5% and 2.5%) for buys and sells on the same days, suggesting that there are not systematic factor exposures matched between buys and sells within portfolios. If a PM were trying to maximize return within a tracking error constraint, we would expect trades on each side to have common factor exposures, so as to not accidentally tilt the fund away from its target benchmark.

3.3 Heterogeneity

Thus far we have presented results based on averages across all PMs. In this section we consider potential sources of heterogeneity in average buying and selling performance across PMs in our sample. Figure 2 presents performance results based on a number of covariates such as turnover, tracking error, and investment style. We find that underperformance in selling appears most prominently amongst fundamentals-oriented managers who hold more active, highly concentrated portfolios with higher tracking error. PMs in this space may have fewer procedures in place to systematically process information about existing holdings (e.g., one might expect these PMs to prioritize soft information relative to quantitative relative valuation signals). In other words, if buys are less related to systematic information, it might be less natural to have effective systems to assist with selling decisions.

whether sales are executed at particularly favorable intraday transaction prices. While this could be the case, we do observe that both our buying and selling measures predict benchmark-adjusted portfolio performance across short and longer horizons. Moreover, if unobserved trading costs were driving our results, then basic comparative statics logic would imply that matching on observable liquidity proxies and timing adjustments should significantly dampen the effect. We do not find evidence for this prediction. In this sense, our analysis parallels an approach taken in the literature which tries to test for adverse selection in insurance markets (see, e.g., Einav and Finkelstein 2011, for a survey) by testing for selection on observables.Thus, price impacts of the average individual trade that are large in magnitude and persistent over substantial periods of time (e.g. longer than 1 week) could potentially change the interpretation of our selling results. Given the size of the portfolios (which is still fairly small relative to total institutional trading volume), the average trade size, and the results which omit illiquid securities and microcaps, we consider this possibility unlikely, but it is difficult to rule out completely.

Figure 2. Variation in Buying and Selling Skill by Manager Characteristics

The end of the bar is the 95% confidence interval of the point estimate at each horizon. Confidence intervals are computed using double-clustered standard errors, calculated as described in Appendix A.4. Details on characteristics construction can be fund in Appendices A.2 and A.5.

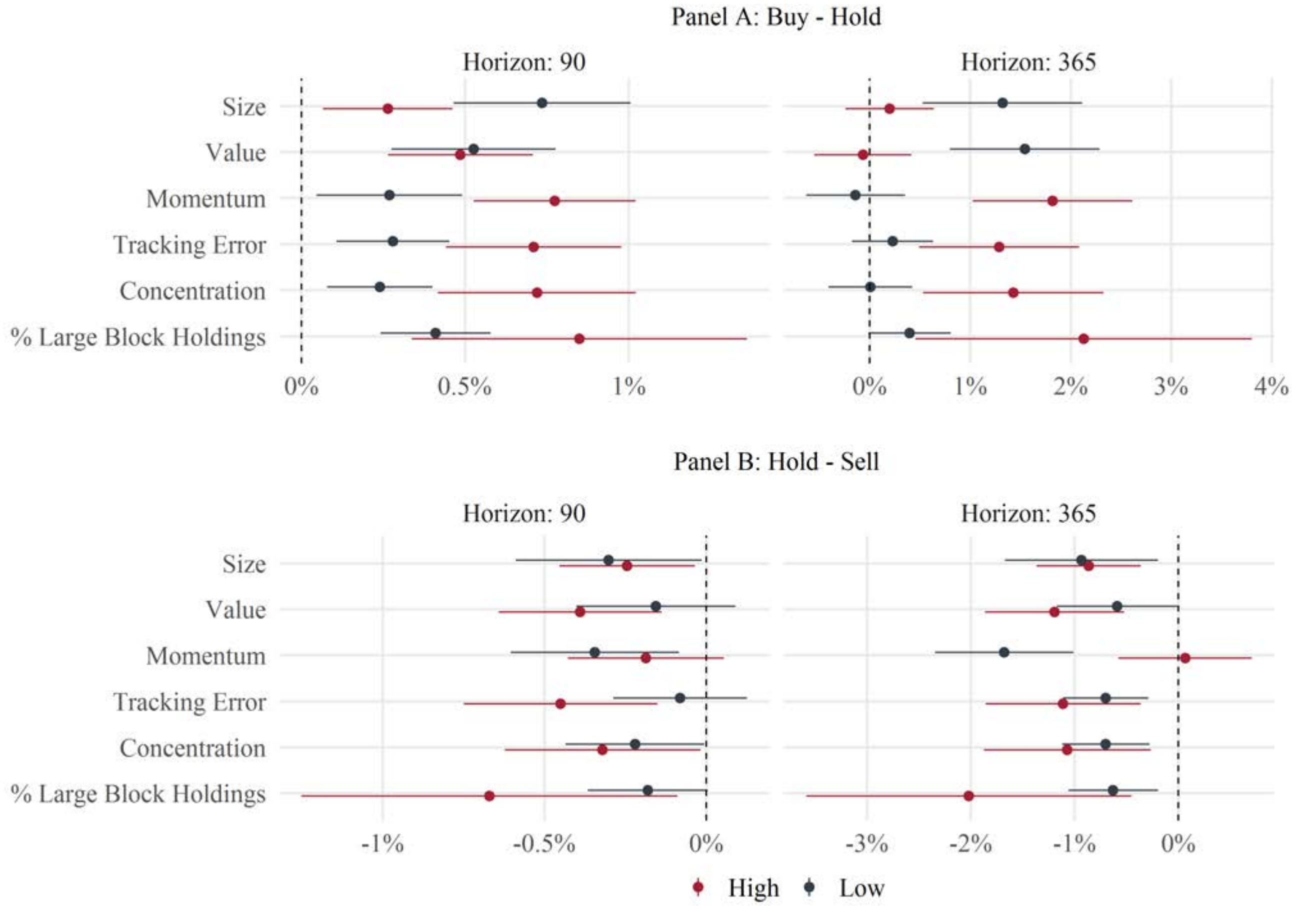

With respect to style, funds that score higher on momentum appear to perform better in selling; funds that score higher on value appear to underperform most in selling. Whereas selling is a key component of the momentum strategy, value plays—at least in terms of how they are described heuristically—are often more about identifying undervalued assets, which seems psychologically more closely aligned with buying. For further details, see Appendix A.2.

4 Explaining Underperformance

In this section, we propose a potential mechanism linking the use of heuristics to systematic underperformance of selling strategies relative to a feasible counterfactual. We then provide evidence for the mechanism by exploiting the panel nature of our database to ask whether patterns in funds' actual trading strategies are associated with predictable differences in performance.

We argue that our results can be explained by expertise with the asymmetric allocation of cognitive resources: PMs focus more on buying and are prone to 'fast', heuristic decisions when selling. This conjecture implies that 1) shifts in cognitive resources towards the latter has the potential to improve performance and 2) heuristics associated with limited attention are more likely to be observed for selling than buying.

4.1 Potential Reasons for an Asymmetric Allocation of Attention

Our interviews with PMs suggest that decisions are overwhelmingly focused on the buying domain relative to the selling domain. Why might this be the case? Several institutional features offer partial explanations. Perhaps the most obvious reason is that, as noted above, virtually all managers in our sample are fundamental rather than quantitatively oriented and have very long holding periods. Research, especially in-depth research that involves constructing custom valuation models and qualitative judgements via conversations with management of various companies, is costly and difficult. The fact that search processes emphasize large blocks of unstructured interactions between PMs and clients (see, e.g., Goyal et al. 2020) suggests that being able to present new investment ideas in a compelling narrative is an important qualitative signal which helps PMs to win and keep existing business. As such, PMs may devote the lion's share of their cognitive resources towards finding the next winner to add to the portfolio.

At the same time, a manager who is buying based on private relative valuation judgements can also sell based on similar valuations. In an "information-ratio" centric world, where managers are evaluated based on bets relative to a benchmark, it is straightforward to map short-selling type positions to an equivalent underweighted position relative to a benchmark. For owned positions that make up a small proportion of the benchmark, relative value judgements still inform which stocks are best to eliminate from a portfolio.

In turn, even taking as given that the client-manager relationship may lead to an initial asymmetric allocation of attention, the question remains of why professional PMs have not learned that their selling decisions are underperforming simple no-skill strategies. Indeed, Table 2 suggests that poor selling has a substantial drag on benchmark-adjusted returns, one of the key metrics which affects PMs' ability to attract clients (Del Guercio and Tkac 2002). The environment in which fund managers make decisions offers several clues. As Hogarth (2001) notes, learning requires frequent and consistent feedback. While it is feasible to generate this type of feedback for both buy and sell decisions, Appendix A.6 discusses how common reporting standards are much better suited to identify underperformance in buying decisions than selling decisions. Thus, PMs and their clients are more likely to receive frequent, valid feedback about their purchases than their sales, which can explain the failure to learn about underperformance in the latter domain.

Importantly, a PM may have two potential reasons to sell a stock: first, she could receive negative information about the stock and subsequently decide that its current price is not justified, and second she could need to raise cash to finance a new buy. Our results point to the second type of selling decision as being worse relative to a counterfactual strategy of simply trimming equal amounts from every other position (on either a value-weighted or equal-weighted basis). The first type of trade seems to perform well; as we discuss in section 4.3, sales on earnings announcement days, when PMs are likely to be paying more attention to their trades, perform significantly better than on other days.

4.2 Bounded Rationality in Selling

How might bounded rationality affect PMs' trading decisions? Many models of decision-making in psychology (Hauser and Wernerfelt 1990) and economics (Lleras, Masatlioglu, Nakajima, and Ozbay 2017) split choices between multiple alternatives—in our case, choosing what asset(s) to buy and sell—into two stages: generating a consideration set and then selecting an option from that set. Prior work has shown that cognitive constraints can lead to the use of heuristics in both stages of the process (Hauser 2014).

For example, Barber and Odean (2008) posit a two-stage process for both buying and selling decisions, where limited attention constrains the consideration set to assets with salient attributes and biases in preferences lead to potentially suboptimal choices from that consideration set.[32] We outline how a two-stage decision-model would operate in our setting and provide evidence for bounded rationality in both stages of the selling process but, importantly, not the buying process. We then demonstrate how the specific heuristics we document can potentially explain the results presented in the preceding sections, including the underperformance of sales relative to a random-sell counterfactual.

In the first stage, rather than considering the entire portfolio, we posit that limited attention places bounds on the consideration set (Hirshleifer and Teoh 2003). Research in psychology and economics has found that these consideration sets are often determined by the ranking and filtering of objects based on salient attributes.[33] Information on prior returns is ubiquitous, and according to models of salience (Bordalo, Gennaioli, and Shleifer 2013), the high variation around average returns should make this attribute particularly top-of-mind for a fast-thinking PM.[34] In turn, extreme deviations in relative returns in either the positive or negative direction naturally emerge as candidates for filtering assets to be included in the consideration set.

In the second stage, the PM must choose which assets from the consideration set to trade. According to the attribute substitution framework of (Kahneman and Frederick 2002), people making 'fast,' heuristic decisions may replace the more difficult question of "which asset in this set is least likely to outperform in the future" with an easier question to answer, such as "how much conviction do I have in this position?" or "how well do I understand this company?" This attribute substitution process can potentially generate systematic underperformance if it leads PMs to, for example, sell still-viable investment ideas.

[32] Also see Sakaguchi, Stewart, and Walasek (2017) for how a two-stage model explains the disposition effect.

[33] See Lleras et al. (2017) for an overview of such attention-based filtering in decision-making. For example, in consumer choice Gourville and Soman (2007) find that people faced with options that differ along several attributes end up only considering those that rank on the extreme ends of those dimensions.

[34] Consistent with this, former investment banker and Bloomberg columnist Matt Levine writes "The rule of thumb wisdom for buying is about fundamentals, but for selling it's usually about price action." https://www.bloomberg.com/opinion/articles/2019-01-10/investors-have-to-sell-stocks-too.

Consistent with attention-based filtering in the first stage, we first show that exogenous events which potentially *expand* PMs' consideration sets by drawing attention to current holdings—earnings announcement days—are associated with substantially better selling performance. We then demonstrate that positions with extreme returns are indeed over-represented in PMs' consideration sets: assets that are in the top or bottom 5 percent based on prior returns are nearly 50 percent more likely to be sold relative to those that just over- or underperformed. Importantly, consistent with the conjecture that more attention is channeled towards buying than selling, earnings announcement days are not associated with changes in buying performance and prior returns have no detectable association with purchase decisions.[35] Finally, we show that low conviction assets are significantly more likely to be sold from the PMs' consideration sets and that the systematic sale of these assets can potentially explain the associated underperformance.

Note that, at first glance, these results might appear to contrast with those of Barber and Odean (2008), who find that retail traders' buying decisions are more prone to attention-based heuristics than their selling decisions. We argue that the two sets of results are complementary once one considers important differences in the groups' decision environments and levels of expertise. Barber and Odean (2008) argue that salient cues are more likely to affect buying decisions because it is much more difficult to attend to the entire set of potential buying opportunities than the limited number of stocks in one's portfolio. On the buying side, PMs are less likely to display attention-driven heuristics because of the process involved in purchasing an asset. As is described in Section 2.1, managers spend a great deal of time and resources cultivating an investment universe that includes the set of stocks that they are considering at any given time. Active research identifies potential mispricing and investment opportunities, at which point the PM will add the respective assets to their portfolio.[36] Moreover, institutional investors are likely to possess expertise that retail traders do not; as a result, when PMs' cognitive resources are devoted to a decision, those deci-

[35]The results from Hartzmark (2014) offer additional support for our proposed mechanism of expertise with asymmetric allocation of cognitive resources: retail investors, who tend to be less sophisticated overall, appear to make both their selling *and* buying decisions based on extreme returns.

[36]This process is described at length in David Swensen's bestselling text titled "Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment" (Swensen 2009), and is consistent with interviews of PMs in our sample.

sions are considerably more likely to outperform those of retail investors.[37] Thus PMs' purchases are less likely to be heuristic-driven because of the combination of expertise and resources devoted to the decision. On the other hand, if PMs are devoting few resources to their selling decisions, those choices are more likely to be driven by attentional heuristics. Notably, while the retail traders in Barber and Odean (2008) hold 4.3 assets on average, PMs in our sample hold 78.5. While it is possible even for an attention-constrained individual to consider the entire portfolio when it only includes 4.3 stocks, this is less likely to be the case for a portfolio that is almost 20 times larger. In turn, we would expect that the consequence of asymmetric resource allocation towards buying decisions would manifest as attention-based heuristics on the selling side.

4.3 Performance on announcement days

We gather earnings announcement dates from the I/B/E/S database and recompute our counterfactual return strategies for stocks which are bought/sold on those days, relative to all other trading days.[38] Managers have a strong incentive to pay close attention to stocks in their portfolios on these dates for several reasons. As discussed in Section 1, the information in financial statements, associated press releases, and conference calls (which even offer opportunities for managers to directly address questions to the company) provide a wealth of new pieces of hard and soft information that are decision-relevant and can potentially improve trading performance (Easley et al. 2008). This information is both (relatively) easily available and salient, since earnings announcement dates are known in advance, and results are heavily covered by the financial press. In turn, we conjecture that earnings announcements prompt PMs to broaden their consideration sets of what to sell, potentially mitigating the constraints imposted by attentional heuristics.

Panel II of Table 4 presents the difference in average performance of trades on announcement versus non-announcement days. Panel A reports the difference be-

[37] This is reflected in the stark difference in overall performance. While PMs in our sample largely outperform their respective benchmarks, retail traders actually underperform the market (Odean 1999).

[38] Our results do not change if we look at performance of trades within a 1, 2, 3, or 4 day window of the announcement.

Table 4. Average Trading Performance Differential: Earnings vs. Other Days

This table presents the the difference between averages of our value added measures for trades of stocks on their earnings announcement days versus all other days (I), and we report the difference between average performance of buys and sells for trades on announcement dates vs non-announcement dates using the baseline measure (II). Double-clustered standard errors, computed using the method described in Appendix A.4, are reported in parentheses.

	Panel A: Buy					**Panel B: Sell**				
Horizon	28	90	180	270	365	28	90	180	270	365
I. Average difference earnings vs. other days:										
Factor-neutral	-0.04	0.14	0.24	0.70	-0.32	0.37	0.29	1.08	2.05	1.54
	(0.21)	(0.39)	(0.62)	(0.81)	(1.02)	(0.20)	(0.37)	(0.58)	(0.77)	(0.80)
Unadjusted	-0.14	0.08	0.16	0.62	-0.25	0.28	0.40	1.16	2.13	1.60
	(0.21)	(0.40)	(0.65)	(0.85)	(0.98)	(0.19)	(0.39)	(0.62)	(0.80)	(0.95)
II. Average performance difference of buys and sells:										
Non-announcement trades						0.31	0.79	1.33	1.68	1.97
Announcement trades						-0.12	0.56	0.41	0.20	0.05

tween average value-added of buy trades executed on earnings announcement days compared with average value-added from all other buy trades.[39] There is little systematic difference in performance, and whatever differences exist are not statistically significant. This is consistent with attentional resources already being devoted towards purchase decisions; information released on earnings announcement days is carefully incorporated into purchase decisions just like other forms of information is used on non-announcement days. Panel B demonstrates the stark contrast in the performance of selling decisions on announcement versus non-announcement days. Selling decisions on announcement days add substantially more value than those sold on non-announcement days. Our point estimate of this difference is 154 basis points over a one year horizon in the baseline specification, with all estimates at the 180 day horizon and onwards being significant at the 5% level. These results hold for raw returns as well.

[39]Given the much smaller number of observations associated with stocks sold on earnings announcement dates, the average performance of sells on earnings announcement dates is positive but imprecisely estimated. Accordingly, we emphasize and report differences between average returns on non announcement days rather than levels. Point estimates for non-announcement days are virtually identical to the overall numbers in section I of Table 3. When computing a standard error for the difference between the two estimates, we impose the assumption that the covariance between the two estimates is zero. This is likely conservative, given that most likely the two estimates are positively correlated (e.g., because stocks sold on earnings announcement days might also be sold several days later as well), which would have the effect of reducing the standard error on the difference.

Finally, Panel II of Table 4 compares the difference between average performance of buys and sells on earnings announcement days versus non-announcement days. In the baseline specification, our point estimate of selling performance on announcement days is +79 bp at a 1 year horizon, which is only 5 bp lower than the corresponding estimate for buying performance on those days. In contrast, buys outperform sells by nearly 200 bp at a 1 year horizon on all other dates, and differences in relative magnitudes are similarly large at other horizons as well.

These findings suggest that when contemporaneous predetermined events shift PMs' attention towards existing positions—leading them consider a wider set of assets and information than they otherwise would when deciding what to sell—their selling performance improves substantially. In Appendix A.8, we present complementary evidence consistent with better selling performance when PMs are more likely to attend to these trades. In particular, we argue that larger transactions are more likely to be 'slow' decisions that are attended to. Indeed, we show that the largest quintile of sales based on transaction size does not underperform the counterfactual at longer horizons, while all other quintiles underperform. Together, this evidence suggests that the overall poor selling performance is does not seem to be due to a fundamental lack of skill in selling.

4.4 Predicting Buying and Selling Decisions

We now attempt to provide more direct evidence that attention constrains PMs' consideration sets in what to sell—but not in what to buy—by looking at the relationship between trade decisions and prior returns.

4.4.1 Measuring association between prior returns and trade decisions

For each portfolio-date, we identify a set of stocks (a subset of holdings in the prior day's portfolio) potentially under consideration to be bought or sold, rank existing holdings according to past benchmark-adjusted returns, and then ask whether managers are more likely to trade based on these ranks. Given the size of our dataset, we adopt a flexible, non-parametric approach to measuring managers' tendency to buy and sell positions based on past returns. Specifically, for the set of prior holdings which are included in the analysis, we compute a measure of returns, usually relative to the benchmark over the same horizon. For each portfolio and trading date, we sort prior positions into 20 bins using these relative rankings, where we set the breakpoint between bins 10 and 11 equal to zero, so that all stocks in bin 10 and below have declined relative to the benchmark. We also emphasize within-manager rankings, rather than absolute levels of these measures, since the definition of "extreme returns" depends on the types of assets in a given PM's investment opportunity set. [40]

While this approach is straightforward for a long-only PM's selling decisions given that the consideration set for sell trades is composed of the current holdings, constructing the consideration set for buying decisions requires taking additional factors into account. Our first approach looks at purchases of assets that already exist in the portfolio; this captures the majority of buys, including large ones (adding up to 99 percent to existing holdings). Our second approach, described in Appendix A.7, includes all purchase decisions—including the opening of brand new positions—and calculates relative prior returns by broadening the consideration set to assets that are likely being considered for purchase.

For this exercise, our preferred measure of prior returns is computed as follows. For positions which were opened more than 1 quarter (90 days) prior to the date of interest, we use the benchmark-adjusted return of the stock from 90 calendar days prior through the trading day before the date of interest. For positions with shorter holding periods, we change the starting point for computing the benchmark-adjusted return to the opening date.[41] We use this as our preferred measure because performance is often reported to clients at a quarterly frequency, and, from a more pragmatic perspective, this construction is less sensitive to fact that position opening dates are left-censored. However, as we show in Section 4.4.2, results are robust to alternative definitions of past returns.[42]

[40] For details on constructing the bins, see Appendix A.7

[41] Results are robust to a wide variety of horizons and whether of not gains and losses are computed since purchase or over a fixed horizon. For example, in Appendix Table A.7, we report probabilities of buying using a prior return measure which does not depend on the time of initial purchase and does not impose the restriction on holding length. In that specification, our main results on the relationship between average buying probabilities and prior returns maintain.

[42] We find nearly identical results if we restrict attention to stocks with opening dates that are observed during our sample. To avoid a fairly mechanical relationship between our prior return measure and the probability that a manager will add to/reduce an existing position from splitting trades across adjacent days (see Appendix A.7), we exclude stocks that were bought in the very recent past. Related to this concern, our regression analyses below always control for the holding period since the position was opened and the holding period since last buy, as well as squared terms of each.

4.4.2 Buying and selling based on past returns

We present results as fractions of positions that are bought or sold within each of the prior return bins. These fractions, which can be interpreted as probabilities, are computed by first calculating the proportion of stocks sold within each bin at the fund-date level, then averaging across all fund-dates in the sample. Figure 3 depicts the results for selling and buying decisions of assets that are already held graphically using a variety of different prior return measures, with 20 bins formed on each measure. Bins are sorted from left to right according to prior returns. We begin with the buying probabilities. The probability of purchasing a stock already held is quite flat across the bins of prior returns. These results hold across all prior return measures considered and no pronounced patterns appear as we move towards more extreme bins in all cases.

A very different picture emerges for the selling probabilities. Assets with more extreme relative returns are substantially more likely to be sold relative to stocks in the central bins. An asset with a prior return in one of the most extreme bins is more than *50 percent* more likely to be sold than an asset with a less extreme return. Moreover, assets in these most extreme bins (1 and 20) have much higher selling probabilities than adjacent bins; such discrete jumps are altogether absent for buying probabilities. Despite the fact different specifications use prior return measures calculated over a variety of horizons, a very pronounced U-shape appears across all.

In Appendix A.7 and Appendix Table A.7, we report two additional results on the relationship between prior returns and buying/selling strategies. First, we extend the analysis to allow for a wider consideration set for purchases; this enables us to incorporate additions of new stocks into the analysis. As above, we find essentially no relationship between prior returns and the likelihood that a stock is purchased. In addition, we look at the relationship between making large purchases or sales (which more than double or halve the existing portfolio weight) and prior returns. Again, we find a strong relationship between extremeness of prior returns and selling probabilities, while seeing no such relationship for purchases.

The U-shaped selling pattern is similar to the rank effect in Hartzmark (2014), where retail traders are more likely to buy *and* sell assets with extreme returns. Hartzmark (2014) documents a similar rank effect in mutual funds as well. However, due to

Figure 3. Probability of buying and selling based on past returns

This set of figures reports daily buying (blue) and selling (red) probabilities for stocks in the portfolio sorted into 20 bins by various past return measures. Panel A sorts on cumulative past benchmark-adjusted returns since the purchase date or one quarter/year, whichever is shortest. Panel B sorts on past benchmark- adjusted returns of a position over one quarter and one year. Panel C sorts on past raw returns of a position over one week and one day. The ten bins on the left are positions with negative returns and the ten bins on the right are positions with positive returns.

Panel A: Cumulative benchmark-adjusted returns capped at 1-quarter and 1-year

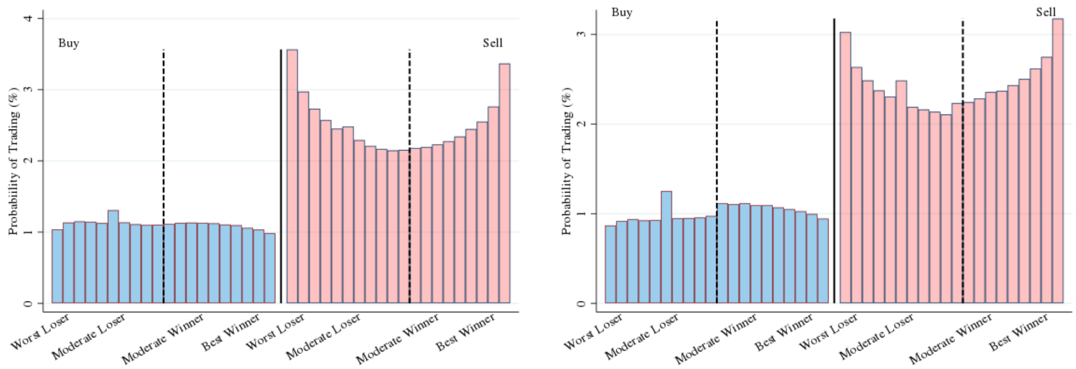

Panel B: Past benchmark-adjusted 1-quarter and 1-year returns of a position

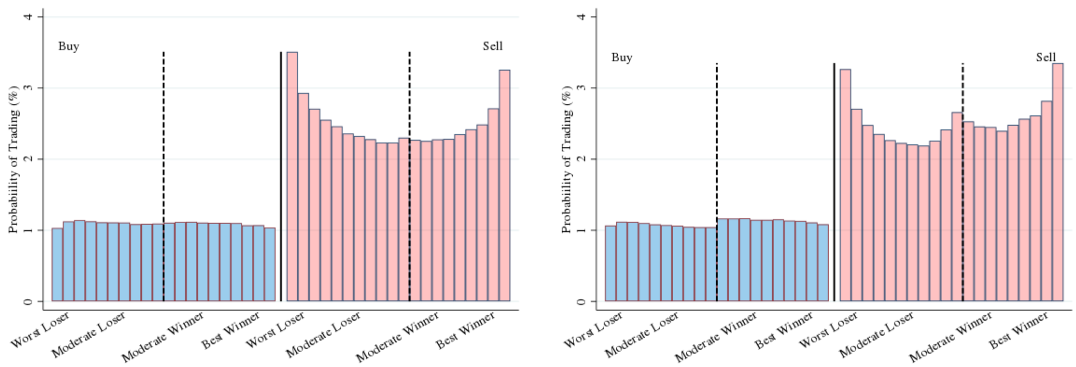

Panel C: Short-horizon 1-week and 1-day returns

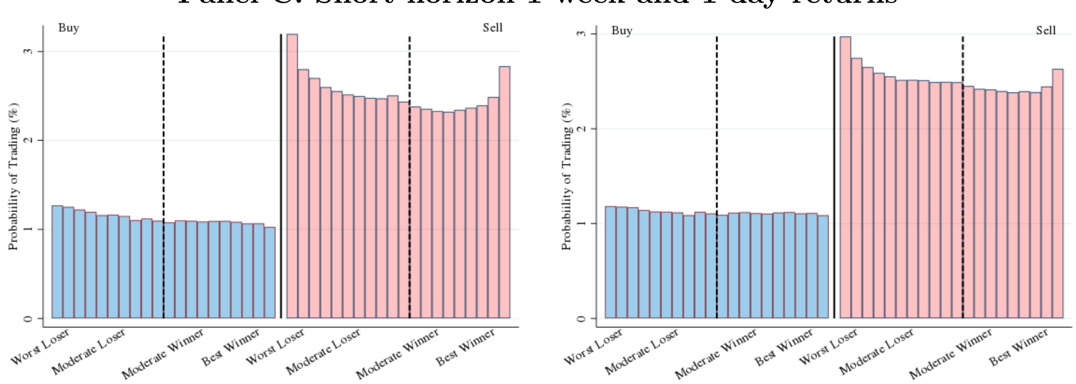

the use of quarterly data, he notes that the behavior can be driven by strategic concerns in response to investor preferences. We show that in contrast to these findings, expert PMs display this heuristic when selling but not when buying. Additionally, as outlined in Section 5, our ability we see each manager's entire portfolio and trades over time allows us to study whether this heuristic is costly for institutional investors. The richness of the dataset also permits exploring the underlying psychological mechanism, yielding evidence that the heuristic is driven by an asymmetric allocation of attention and that its costs are not driven by the unloading of extreme positions *per se*, but rather by the types of positions sold from the constrained consideration set.[43]

4.4.3 Alternative explanations

We now consider several instrumental reasons that could potentially explain our results. As discussed in Section 2, the vast majority of portfolios in our sample are tax-exempt, so the U-shaped selling pattern cannot be rationalized with tax concerns. Our finding that sales are more likely for positions with extreme returns over very long (1 year) and very short (1 week) horizons makes agency-based explanations—where PMs are reluctant to report realized losses to their clients—unlikely. Agency-based explanations also seem unlikely to explain the large jumps in probabilities observed between the 19^{th} and 20^{th} (1^{st} and 2^{nd}) bins relative to the 18^{th} and 19^{th} (2^{nd} and 3^{rd}) bins. These jumps are consistent with limited attention constraining the consideration set of what to sell, as the top and bottom 5 percent of returns are much more likely to be displayed and made salient to PMs (see Ungeheuer (2017) for direct evidence). This also mitigates concerns about risk management motives, since the relative risk of assets in extreme bins is likely to be fairly comparable to less extreme adjacent bins.[44]

We also examine whether our observed pattern can be explained by other variables

[43] Ben-David and Hirshleifer (2012) document a related phenomenon where retail investors are more likely to sell securities with high or low prior returns; An (2015) shows that stocks with such returns are more likely to outperform in the following month. This V-shaped trading pattern is distinct from ours in that the high or low returns are calculated at the level of the individual security rather than with respect to overall portfolio returns. Moreover, the pattern is documented both for buying and selling decisions.

[44] In subsequent regression analyses, we will include controls for idiosyncratic volatility, systematic factor exposures, and position size, all of which are potentially relevant for risk management. Inclusion of these controls generally has a very limited impact on estimates analogous to the nonparametric statistics presented above.

which could be correlated with our prior return measures: holding length and position size. We find that both the U-shaped selling pattern and the flat buying pattern is observed across the holding length and position size variable sorts. For more detailed results, see Appendix A.7.

Finally, Table 5 reports estimates from a series of linear probability models for the likelihood of selling or buying, which allow us to control for a number of time-varying fund characteristics (either via controls, fund fixed effects, or fund-date fixed effects), calendar time effects, as well as other position characteristics. All specifications include linear and quadratic controls for holding length since the position was opened, holding length since last buy, and position-level portfolio weight (as a fraction of total portfolio assets under management). The key regressors of interest are dummies for each of the prior return categories, which have the same interpretation as the bins used in the preceding analyses, where the omitted category remains bin 10 (slight loser positions). Results are similar with different prior return measures and different numbers of bins.

We begin with the right panel which characterizes selling probabilities. Across all of these specifications, the difference in the predicted probability of selling a stock in bin 20 is at least 50 percent higher than the probability of selling a stock in bins 6 through 15, and always considerably higher than bin 19. Likewise, we observe similar strong nonlinearities for stocks in bins 1 through 2 relative to more central bins. Column 4 includes stock-date fixed effects as in Hartzmark (2014), so the main coefficients of interest are identified via variation in the relative return categories across portfolio managers who hold the same stock on the same date. Even then, we find that positions in the most extreme returns bins are substantially more likely to be sold.

Turning to the left panel, the relationship between buying probabilities and prior return measures is much more muted. In the loss domain, most of the coefficients are insignificant despite being estimated on a sample of over 50 million observations. Even the significant coefficients are substantially smaller in magnitude than the coefficients associated with selling probabilities. In the saturated specification presented in column 5, only the coefficient on bin 1 is statistically distinguishable from zero. Looking at large positive returns, we observe a greater number of significant coefficients, but the differences between central and extreme bins (e.g., bins 16 through 18 and bin 20 or bins 19 and 20) are very small relative to the same differences for sales.

Table 5. Probability of trading based on prior returns

This table presents position-level estimates of a linear probability model (in percentage points) for the likelihood of trading a given stock. The key explanatory variables of interest are indicators corresponding to 20 bins of past benchmark-adjusted returns capped at one year, where the tenth bin is the omitted category. We control for fund characteristics including lagged log(AUM), prior-month turnover, the annual volatility of a funds benchmark-adjusted returns, and prior month loadings on Fama-French Cahart regressions (calculated using the Dimson (1979) procedure using 1 year of prior daily returns). We adjust for position-level characteristics including linear and quadratic terms in holding lengths (overall and since last buy) and position sizes(% AUM) at the beginning of the day. The coefficients and t-statistics are reported for the estimates on each bin. We present versions of these estimates which condition on trading (buys for buying probability, sells for selling probability) in Table A.9 in Appendix A.7.

	Buying Probability				Selling Probability			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Bin 1	-0.041	-0.047*	-0.074**	-0.144*	1.389***	1.379***	1.237***	0.796***
	(-1.789)	(-2.068)	(-3.178)	(-2.192)	(14.828)	(14.592)	(13.032)	(6.343)
Bin 2	0.040*	0.033	0.008	-0.058	0.806***	0.799***	0.666***	0.543***
	(2.071)	(1.685)	(0.410)	(-1.217)	(12.580)	(12.330)	(10.471)	(6.212)
Bin 3 to 5	0.046**	0.041**	0.025	0.001	0.432***	0.428***	0.331***	0.308***
	(3.275)	(2.817)	(1.606)	(0.031)	(10.998)	(10.731)	(8.980)	(5.686)
Bin 6 to 9	0.025**	0.022*	0.011	0.012	0.109***	0.107***	0.067***	0.079**
	(2.856)	(2.337)	(1.164)	(0.763)	(6.750)	(6.328)	(4.549)	(3.262)
Bin 11 to 15	-0.029**	-0.031**	-0.032***	-0.014	0.028	0.037*	0.005	0.105***
	(-2.701)	(-2.679)	(-3.568)	(-0.771)	(1.560)	(1.966)	(0.250)	(4.213)
Bin 16 to 18	-0.088***	-0.092***	-0.109***	-0.152***	0.312***	0.318***	0.234***	0.559***
	(-4.699)	(-4.668)	(-6.291)	(-3.679)	(8.004)	(7.945)	(6.063)	(8.760)
Bin 19	-0.129***	-0.133***	-0.155***	-0.220***	0.578***	0.582***	0.485***	0.834***
	(-5.810)	(-5.769)	(-7.276)	(-4.071)	(10.849)	(10.725)	(9.125)	(9.505)
Bin 20	-0.169***	-0.175***	-0.212***	-0.286***	1.186***	1.186***	1.071***	1.132***
	(-6.700)	(-6.665)	(-8.366)	(-4.179)	(15.869)	(15.638)	(14.303)	(9.581)
Fund Control	Yes	Yes	No	Yes	Yes	Yes	No	Yes
Fixed Effects	None	Date	Fund × Date	Stock × Date	None	Date	Fund × Date	Stock × Date
R^2	0.022	0.028	0.283	0.281	0.005	0.009	0.179	0.317
N	54.2M	54.2M	56.2M	45.5M	54.2M	54.2M	56.2M	45.5M

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Taking stock, the regression specifications, in conjunction with the nonparametric evidence in Appendix Table A.8, suggest that the considered sources of omitted variable bias are unlikely to explain our results. Additionally, Table A.9 calculates the same model, conditional on a trade (at least one buy on a given day for a given fund when analyzing buys, and at least one sell when analyzing sells). We find that the results are qualitatively similar, but with much larger magnitudes for the selling decisions. Together these results are consistent with non-instrumental motives driving selling but not buying decisions.

4.5 Selling and Conviction

We now examine which assets from the consideration set PMs choose to sell. As discussed above, the attribute substitution framework of Kahneman and Frederick (2002) predicts that people making heuristic decisions may replace the difficult question of determining which asset is least likely to outperform with an easier question related to conviction or psychological attachment. To capture this, we order positions based on their *active share*, or how much they are overweighted relative to the benchmark (Cremers and Petajisto 2009). This measure captures how much the PM stands to gain if the stock beats the benchmark.[45] Assets with high active share typically correspond to positions that the manager has spent a good deal of effort building up over time, becoming familiar and attached to the firm in the process. This costly procedure likely generates greater conviction in the position for non-instrumental reasons, such as sunk costs or psychological ownership (Anagol, Balasubramaniam, and Ramadorai 2018; Heath 1995; Kahneman, Knetsch, and Thaler 1990).

Positions with low active share can manifest for three main reasons: 1) a position had a high active share but has done very poorly, 2) the PM has added a position to the portfolio but has not built it into a larger one, or 3) the PM chooses to hold a position close to the benchmark weight of a stock which is large in the benchmark (or underweight, though negative active shares are fairly uncommon).[46] The PM may

[45] Active share is calculated by taking the difference between a PM's weight on a stock in the fund and subtracting the corresponding weight, if any, of the same stock in the client-provided benchmark, a measure which is provided to us by Inalytics. Since performance is evaluated based on benchmark-adjusted returns, an asset that is overweighted generates excess returns when it goes up and excess losses when it goes down.

[46] A fourth alternative is that the PM has actively reduced a formerly large position through prior

Table 6. Probability of selling by active share and past returns

This tables reports differences in probabilities, in percentage points, of selling by bins of past benchmark-adjusted returns double sorted with bins of position-level active share, relative to a baseline category (the tenth bin of past benchmark-adjusted returns, within each active share quartile). Columns represent different active share bins, along with the difference across rows between the smallest active share bin and the average across the other bins (active share bins 2-4). Calculations for 8 categories of prior returns, formed from 20 bins of past position returns, are reported in rows. We report the baseline selling probability for the 10th bin in the last row.

	Active share Bins				
Prior Return Bins	Lowest	Low	Higher	Highest	Lowest - Others
1	2.264	0.724	0.452	0.414	1.734
2	1.501	0.445	0.320	0.223	1.171
3-5	0.828	0.280	0.150	0.105	0.650
6-9	0.254	0.122	0.051	0.013	0.192
11-15	-0.067	0.093	0.115	0.146	-0.185
16-18	0.159	0.387	0.424	0.488	-0.274
19	0.580	0.689	0.790	0.859	-0.199
20	1.426	1.338	1.360	1.410	0.057
Baseline Level: Bin 10	3.329	1.851	1.691	1.779	1.556

still have conviction in a stock in the first category as she had exerted time and effort in building it up in the past. Rather than reflecting a particular view about future returns, assets in the third category may be in place to minimize exposure of a fund's relative returns to the idiosyncratic returns of large assets in the benchmark.

In contrast, assets in the second category are most likely to include the PM's 'neglected ideas.' The manager has gathered enough favorable information on each asset to add a position to the portfolio, but has not elected to build it up further. In turn, heuristic thinking would generate fewer reasons to keep a low active share asset in this category. Moreover, time-constrained PMs may monitor them less closely than higher conviction positions, especially if other factors draw attention elsewhere (e.g. attractive buying opportunities). Thus, in discarding these positions, PMs may be throwing out still-viable investment ideas—ones potentially capable of outperforming many existing holdings—especially if high conviction positions are more likely to have already realized their anticipated upside potential.

sells.

Table 6 documents the PMs' propensity to sell based on active share, both overall and within each of the 20 bins of prior returns. To construct this table, assets within each portfolio are sorted into four bins based on their active share. We then construct a measure capturing the propensity to sell an asset based on its prior returns; specifically, the difference in the probability of selling a stock in a given bin of prior returns relative to the middle one (bin 10, Slight Loser). The last column reports the difference between the lowest active share bin and the average across the other three active share bins in the same row of prior returns. We report baseline probabilities for the omitted bin in the last row.

Results are consistent with PMs being most prone to selling neglected ideas, as measured by low active share, from the consideration set of extreme returns.[47] First, examining the baseline probabilities, we note that low active share positions are substantially more likely to be sold regardless of the level of prior returns. Second, we find that stocks in the lowest active share bin are much more likely to be sold when they exhibit prior returns below the benchmark, especially extreme ones, relative to high active share assets. The probability of selling a stock with the lowest active share and lowest prior return bin is 5.6, or 140 percent larger than the baseline probability of selling, which is 2.3. Assets in these bins are also 155 percent more likely to be sold than those which experienced similar levels of underperformance (bin 1) but have the highest active share. Thus, low active share positions are particularly likely to be discarded when they are in the consideration set of extreme underperformance. Selling probabilities in the lowest active share bin are relatively less affected by moderate gains, and responses to the most extreme gains in bin 20 are similar regardless of active share.

We then examine whether sales of low active share assets tend to underperform relative to a random selling counterfactual. Panel A of Table 7 depicts the performance of sales relative to a random counterfactual by bins of stocks' active share for our baseline value-added measure. As above, we sort positions into four bins based on a prior day estimate of active share, then separately form counterfactual performance measures for these different subsamples of trades. We see a stark contrast in performance: assets in the lowest three active share bins underperform substantially more

[47] Consistent with results in the prior section, we find that buying probabilities do not exhibit a significant relationship with prior returns.

Table 7. Post-trade sell returns relative to counterfactual by active share

This table presents the average factor-neutral returns relative to random sell counterfactuals for sell portfolios sorted by active share. We compute average returns of stock held minus returns of stocks sold. We rank the active share measures within funds at a daily level and sort them into four bins from Lowest, Low-Med, Med-High to Highest sizes. For the Lowest Active Share bin, we further split into two halves by a position's weight: Smaller positions (below the 50^{th} percentile) and Larger positions (above the 50^{th} percentile), after sorting by active share. Columns represent sell performance measures at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. We report point estimates of average counterfactual returns for each portfolio at different horizons. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.4, are reported in parentheses.

Active Share	Horizon				
Bins	28	90	180	270	365
Lowest,	-0.46	-0.96	-1.46	-1.80	-1.62
Smaller Positions	(0.10)	(0.32)	(0.66)	(0.91)	(0.96)
Lowest,	0.10	0.01	-0.05	0.14	0.55
Larger Positions	(0.07)	(0.13)	(0.21)	(0.29)	(0.33)
Low-Med	0.20	-0.15	-0.47	-0.86	-1.19
	(0.07)	(0.18)	(0.35)	(0.41)	(0.58)
Med-High	0.16	0.13	-0.15	-0.24	-0.84
	(0.07)	(0.13)	(0.19)	(0.26)	(0.34)
Highest	0.32	0.29	0.23	0.13	-0.35
	(0.08)	(0.21)	(0.39)	(0.50)	(0.56)

than sales of assets in the top active share bin, where the latter performs much closer to the the counterfactual. This is consistent with PMs holding on to high active share assets when thinking fast; thus, when sales of those positions are observed, they are more likely to be informed ones.

As discussed above, very low active share positions may be held to hedge a fund's exposure to idiosyncratic returns of large stocks in the benchmark. These positions would appear in the data as having a very low active share and a high portfolio weight.[48] In turn, we separate stocks in the lowest active share bin into two categories based on absolute position size, and recompute our counterfactual performance measures for these subsamples. Whereas large positions with low active share tend to have positive

[48] As an example, if Apple is 3% of a PM's benchmark index we might observe a position in Apple of 3% which would have an active share of zero, despite the fact that 3% would be a quite large absolute position size.

counterfactual performance measures (consistent with these being somewhat passive positions and good candidates for sales), we find that sales of assets from the second 'neglected ideas' category—low active share and low position size—to substantially underperform.

It is tempting to conclude from these results that since the underperformance of selling strategies is associated with low active share (usually smaller) positions, the costs in terms of overall portfolio performance associated with these transactions is likely to be small. However, this reasoning is incorrect provided that changes in portfolio weights induced by selling smaller initial positions are similar to those from selling larger ones. Holding trade size as a fraction of portfolio market value constant, the cost in foregone profits from a suboptimal trade are independent of the initial size of the position.[49] Indeed, we find that average trade sizes for sells are quite similar across both active share and initial position size bins. Further, recall from Table 2 above that our measures of average selling performance are highly predictive of overall benchmark-adjusted performance.

In addition, we note that poor selling is likely associated with performance in two ways. The first is a direct effect: it changes the weights in the current portfolio. The second is an indirect effect: 'fast' selling of low conviction positions may lead those stocks to be discarded from the consideration set of future buys. Consistent with the potential importance of this latter effect, we find that, once sold, assets are substantially less likely to be purchased again, suggesting that attention-constrained elimination of small positions may interfere with a PM's security selection process.[50]

Note that it is possible that our results may be explained by an alternative mechanism which can generate underperformance of sales as a function of only the first stage outlined above. Specifically, if extreme positions are ranked similarly across a large enough number of institutional portfolios and sold for reasons unrelated to funda-

[49]Further, since the effect of an idiosyncratic stock return on overall portfolio variance is a convex function of the weight, one could argue that the effect on measures of performance that adjust for idiosyncratic risk exposures such as the information ratio are larger for small positions.

[50]Before a security has been sold once, there is a 76% chance that the portfolio will purchase a security at least once more. After the first time a portfolio sells a security, there is only a 40% chance the portfolio will ever increase its position in that security again before liquidating it. Once a stock is removed from the portfolio, there is only a 40% chance that the position is ever added back to the portfolio. On the intensive margin, the daily probability that a stock is purchased again declines from 3.3% to 2.4% once a stock is sold for the first time.

mental value, then the increased propensity to sell them can generate downward price pressure on current prices. Although the decision to unload an asset by an individual investor may have been driven by contextual factors specific to their portfolios, if the asset is categorized as extreme in enough portfolios, then this would lead to performance being a function of prior returns at the market level. An (2015) and An and Argyle (2016) provide evidence from holdings of individuals and mutual funds suggesting that, as selling pressure is reduced over the medium term, subsequent mean reversion in prices can lead these stocks to systematically outperform on a risk-adjusted basis going forward.[51]

5 Heuristic use and Performance

In this section, we exploit the panel nature of our dataset in order to illustrate a more direct link between the performance of selling strategies and fund-level characteristics, such as the propensity to sell assets with extreme returns. As in Section 3, we compare the returns of the actual stocks traded with counterfactual random selling strategies. We then ask whether patterns in funds' actual trading strategies are associated with predictable differences in performance. To operationalize this, we compute several fund-level characteristics and sort fund-weeks into categories based on these characteristics, then compute the average value-added associated with PMs' trades in each bin. Before proceeding, we note that this analysis is only able to identify correlations in the data, so it is not feasible to rule out other time-varying fund characteristics which simultaneously drive performance and observable properties of trading behavior.

We begin by considering the impact of heuristic use on performance. Based on the mechanism outlined in Section 4.1, we use the greater propensity to sell assets with extreme returns as a proxy for heuristic use.[52] To capture what we term 'heuristic intensity,' positions are sorted into 4 bins based on the fraction of stocks sold that are located in the extreme bins (Worst Loser and Best Winner) for each fund-week.[53]

[51] At the same time, as we discuss further in Appendix A.4, we find very little correlation between stocks sold at the same time in our dataset. These correlations are slightly positive for stocks sold on the same day and slightly negative at other leads and lags, which provides some evidence against this channel in our data.

[52] This greater propensity is only a proxy for heuristic use because, as demonstrated in Section 4.1, PMs do not randomly sell assets from the consideration set of extreme returns.

[53] For instance, the mean of this heuristics intensity measure is around 0.4 on a monthly basis, which

In order to reduce noise in the measures related to discreteness, we only compute the heuristic intensity measure for funds which hold at least 24 stocks (i.e., there are 4 stocks on average per bin) and only include fund-weeks in which at least 3 stocks are traded.[54] We then rank fund-weeks into four categories according to this measure to calculate relative performance of the associated selling decisions. Our primary rationale for a weekly frequency is that it provides a satisfactory balance between reducing potential noise in the sorting variable (by averaging over multiple trades) while still operating at a high enough frequency to capture variation in heuristic-use.

Panel A of Table 8 presents sample averages of counterfactual returns sorted into four bins based on heuristic intensity within-fund. Over the course of each PM's time series, each week is sorted into one of four categories based on its level of heuristic intensity. We find that our proxy for heuristic intensity is positively associated with significant underperformance. The highest levels of heuristic intensity are associated with the worst performance, especially at the longer horizons. Magnitudes are quite substantial: at a 1 year frequency, the highest level of heuristic intensity predicts an average of around 220 foregone basis points relative to a random-sell counterfactual. At the same time, average performance of sales which occur when PMs are selling fewer extreme positions is statistically indistinguishable from the counterfactual. Appendix A.8 demonstrates the robustness of our results when using the unadjusted performance measures. In contrast to these results, we find no similar relationship when looking at differences in average heuristic use *between*-managers, a result which further supports our assertion that underperformance in selling, rather than being driven by persistent differences in skill, is due to the asymmetric allocation of attention.[55]

In the preceding section we argued that both stages of the selling process are prone to heuristic thinking—limiting the consideration set to assets with salient attributes and then choosing to sell those that the PM has least conviction in. The literature on heuristics and biases documents that people are more likely to rely on heuristics during

would imply (through a simple application of Bayes' rule) that the likelihood of a stock being sold in the extreme bin is 4/3 the likelihood of a stock being sold in one of the central bins. In Appendix Table A.12, we show, perhaps surprisingly, our measure of heuristics intensity is largely uncorrelated with a variety of observable fund characteristics.

[54] We lose about one quarter of fund-week observations due to these restrictions.

[55] More specifically, during periods of time when the PM is attending to sales and less reliant on heuristics, the performance of her selling decisions increases substantially.

Table 8. Post-trade sell returns relative to counterfactual by fund behavior

This table presents average returns relative to random sell counterfactuals for sell portfolios sorted by heuristics intensity, cumulative benchmark-adjusted fund returns since the beginning of a quarter, and a proxy for 'cash raising' episodes. We divide these measures into four bins from Lowest, Low-Med, Med-High and Highest, based on their rankings. Columns represent sell performance measures at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. Double-clustered standard errors, computed using the method described in the section A.4, are reported in parentheses.

Fund Characteristics	Bins	Horizon 28	90	180	270	365
Panel A: Heuristics Intensity Fraction of extreme stocks sold weekly (sorted across funds)	Lowest	-0.07	-0.21	-0.29	-0.42	-0.28
		(0.07)	(0.14)	(0.20)	(0.24)	(0.33)
	Low-Med	0.00	-0.07	-0.16	-0.12	0.03
		(0.06)	(0.11)	(0.15)	(0.21)	(0.24)
	Med-High	0.05	0.05	-0.13	-0.20	-0.31
		(0.06)	(0.14)	(0.21)	(0.22)	(0.26)
	Highest	0.11	-0.59	-1.33	-1.57	-2.20
		(0.09)	(0.19)	(0.31)	(0.39)	(0.47)
Panel B: Cumulative Benchmark-adjusted Fund Return since the beginning of a quarter (sorted across funds)	Lowest	-0.22	-0.96	-1.43	-1.81	-1.95
		(0.10)	(0.24)	(0.38)	(0.48)	(0.56)
	Low-Med	0.01	-0.07	-0.46	-0.50	-0.69
		(0.06)	(0.13)	(0.19)	(0.28)	(0.30)
	Med-High	-0.02	-0.15	-0.42	-0.69	-0.59
		(0.07)	(0.12)	(0.21)	(0.26)	(0.31)
	Highest	0.16	0.12	0.21	0.03	0.06
		(0.09)	(0.17)	(0.29)	(0.40)	(0.43)
Panel C: Net Buy Weekly number of stocks bought minus number of stocks sold (sorted within fund)	Lowest	-0.01	-0.47	-0.73	-1.25	-1.44
		(0.08)	(0.18)	(0.26)	(0.36)	(0.46)
	Low-Med	0.04	-0.20	-0.73	-0.97	-1.15
		(0.08)	(0.17)	(0.23)	(0.36)	(0.40)
	Med-High	0.04	0.11	-0.01	-0.08	-0.35
		(0.07)	(0.13)	(0.19)	(0.24)	(0.29)
	Highest	0.08	-0.03	-0.26	0.01	0.23
		(0.06)	(0.13)	(0.20)	(0.25)	(0.28)
Panel D: Flows Weekly Assets entering portfolio as % of AUM (sorted within fund)	Most outflows	-0.06	-0.53	-0.86	-1.10	-0.99
		(0.09)	(0.18)	(0.30)	(0.38)	(0.46)
	Low-Med	0.05	-0.12	-0.47	-0.56	-0.73
		(0.06)	(0.12)	(0.18)	(0.24)	(0.29)
	High-Med	0.08	-0.12	-0.27	-0.53	-0.79
		(0.06)	(0.12)	(0.18)	(0.23)	(0.27)
	Most inflows	0.07	-0.09	-0.42	-0.54	-0.71
		(0.07)	(0.13)	(0.20)	(0.28)	(0.33)

situations when cognitive resources are in higher demand, such as in times of stress or when attention is otherwise occupied (see Kahneman (2003) for review). Panels B and C of Table 8 consider two empirical proxies intended to capture periods emblematic of such episodes. As in Panel A, these measures are computed on a weekly basis and sort fund-weeks into four categories to capture either between or within-manager variation.

The first aims to capture performance when the PM is likely to be stressed. Institutional investors are known to take stock of their own performance based on calendar time, e.g. on a quarterly or yearly basis. Based on the conjecture that the PMs are more likely to be stressed when their overall portfolio is underperforming, we construct a measure that captures portfolio performance relative to the beginning of the preceding quarter. Table 8, Panel B demonstrates that selling quality is worst (relative to a random-sell counterfactual) when the PM's overall portfolio is underperforming the most. Panel C considers a measure that proxies for sales that are likely to be driven by cash raising considerations rather than forecasts of relevant performance metrics. We posit that observing larger bundles of assets being sold (relative to being bought) is emblematic of the manager being in "cash-raising mode." We compute the difference between the number of stocks bought and the number of stocks sold, where both measures are expressed as fractions of the number of stocks in the portfolio. We find that the difference between the number of stocks bought and sold predicts greater underperformance of the selling decisions.[56]

As discussed above, some institutions such as mutual funds face direct pressure from performance-driven fund flows, creating the possibility that flows induce PMs to engage in large sales under unfavorable conditions (Alexander et al. 2007). Following Coval and Stafford (2007), we construct daily flows as the portion of assets under management which are not explained by the previous day's returns, then aggregate them to a weekly frequency. We test the above hypothesis in Panel D by creating proxies for flows and sorting fund-weeks into bins based on these proxies. Consistent with our discussion in Section 2.1 that flow concerns are less pronounced in our institutional

[56] In similar spirit, we present additional evidence in Appendix A.8 which relates performance to turnover. There, we sort each portfolio's time series into periods with low and high turnover. We find that performance of sales deteriorates during high turnover periods (which are presumably characterized by better trading opportunities), while performance of buys remains unchanged. This is consistent with a mechanism where better buying opportunities attract the already scarce attentional resources allocated to sales, further degrading performance in that domain.

context, the performance results—that buys outperform the counterfactual while sells underperform—are unchanged across the flow measures.

6 Conclusion

We utilize a unique dataset and find evidence that financial market experts—institutional investors managing portfolios averaging $573 million—display costly, systematic heuristics. A striking finding emerges: While investors display skill in buying, their selling decisions underperform substantially—even relative to random-sell strategies. We provide evidence that investors use heuristics when selling but not when buying, and that these heuristic strategies are empirically linked to the documented difference in performance.

As shown in Section 3, the comparison of trades on earnings announcement versus non-announcement days suggests that PMs do not lack fundamental skills in selling; rather, results are consistent with PMs devoting more cognitive resources to buying than selling. When decision relevant information is salient and readily available—as it is on announcement days—PMs' selling performance improves substantially. We propose a mechanism through which overall underperformance in selling can be explained by a heuristic two-stage selling process, where PMs limit their consideration set to assets with salient characteristics (extreme prior returns) and sell those they have the least conviction in. A proxy for this heuristic strategy is associated with substantial losses relative to a no-skill random selling strategy.

In light of the imbalance in feedback discussed in Section 4.1 and Appendix A.6 , the theoretical framework of Gagnon-Bartsch, Rabin, and Schwartzstein (2018) suggests that PMs may fail to recognize their underperformance in selling even in the long-run. The authors show that a mistaken theory such as the favorability of selling positions with extreme returns may persist in the long run because people channel their attention through the lens of this theory. As in Schwartzstein (2014), errors persist due to the person ignoring information that seems irrelevant and only updating her beliefs based on information that is attended to. Our findings imply significant benefits to creating environments where learning can occur more effectively, such as through amended reporting standards that emphasize counterfactual sales performance. Moreover, our empirical results on a link between heuristic use and underperformance of selling strategies suggest that PMs adoption of decision aids and/or simple alternative selling strategies may substantially improve performance.

References

Alexander, G. J., G. Cici, and S. Gibson (2007): "Does motivation matter when assessing trade performance? An analysis of mutual funds," *The Review of Financial Studies*, 20, 125–150.

An, L. (2015): "Asset pricing when traders sell extreme winners and losers," *The Review of Financial Studies*, 29, 823–861.

An, L. and B. Argyle (2016): "Overselling Winners and Losers: How Mutual Funds Trading Affects Asset Prices," *PBCSF-NIFR Research Paper*.

Anagol, S., V. Balasubramaniam, and T. Ramadorai (2018): "Endowment effects in the field: Evidence from India's ipo lotteries," *The Review of Economic Studies*, 85, 1971–2004.

Ball, R. and P. Brown (1968): "An empirical evaluation of accounting income numbers," *Journal of accounting research*, 159–178.

Barber, B. and T. Odean (2000): "Trading is hazardous to your wealth: The common stock investment performance of individual investors," *The journal of Finance*, 55, 773–806.

——— (2011): "Chapter 22. The Behavior of Individual Investors *," 2.

Barber, B. M. and T. Odean (2001): "Boys Will be Boys: Gender, Overconfidence, and Common Stock Investment," *The Quarterly Journal of Economics*, 116, 261–292.

——— (2008): "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors," *The Review of Financial Studies*, 21, 785–818.

——— (2013): "The behavior of individual investors," in *Handbook of the Economics of Finance*, Elsevier, vol. 2, 1533–1570.

Ben-David, I. and D. Hirshleifer (2012): "Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect," *The Review of Financial Studies*, 25, 2485–2532.

Ben-Rephael, A., Z. Da, and R. D. Israelsen (2017): "It depends on where you search: Institutional investor attention and underreaction to news," *The Review of Financial Studies*, 30, 3009–3047.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013): "Salience and consumer choice," *Journal of Political Economy*, 121, 803–843.

Brunnermeier, M. and S. Nagel (2004): "Hedge funds and the technology bubble," *The Journal of Finance*, 59, 2013–2040.

Camerer, C. F. and U. Malmendier (2007): "Behavioral Organizational Economics," in *In Peter Diamond and Hannu Vartiainen, eds., Behavioral Economics and Its Applications. Princeton and*, Princeton University Press.

Cameron, A. C. and D. L. Miller (2015): "A practitioner's guide to cluster-robust inference," *Journal of Human Resources*, 50, 317–372.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001): "Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk," *The journal of finance*, 56, 1–43.

Chang, T. Y., S. M. Hartzmark, D. H. Solomon, and E. F. Soltes (2016a): "Being surprised by the unsurprising: Earnings seasonality and stock returns," *The Review of Financial Studies*, 30, 281–323.

Chang, T. Y., D. H. Solomon, and M. M. Westerfield (2016b): "Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect," *The Journal of Finance*, 71, 267–302.

Chen, D. L., T. J. Moskowitz, and K. Shue (2016): "Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *The Quarterly Journal of Economics*, 131, 1181–1242.

Chen, F., L.-W. Chen, H. Johnson, and S. Sardarli (2017): "Tailored versus Mass Produced: Portfolio Managers Concurrently Managing Separately Managed Accounts and Mutual Funds," *Financial Review*, 52, 531–561.

Chen, J., H. Hong, M. Huang, and J. D. Kubik (2004): "Does fund size erode mutual fund performance? The role of liquidity and organization," *American Economic Review*, 94, 1276–1302.

Coibion, O. and Y. Gorodnichenko (2015): "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105, 2644–78.

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

Cremers, K. M. and A. Petajisto (2009): "How active is your fund manager? A new measure that predicts performance," *The review of financial studies*, 22, 3329–3365.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers (1997): “Measuring mutual fund performance with characteristic-based benchmarks,” *The Journal of finance*, 52, 1035–1058.

Del Guercio, D. and P. A. Tkac (2002): “The determinants of the flow of funds of managed portfolios: Mutual funds vs. pension funds,” *Journal of Financial and Quantitative Analysis*, 523–557.

DellaVigna, S. and M. Gentzkow (2017): “Uniform Pricing in US Retail Chains,” Working Paper 23996, National Bureau of Economic Research.

Di Mascio, R., A. Lines, and N. Y. Naik (2017): “Alpha decay,” Working paper.

Dimson, E. (1979): “Risk measurement when shares are subject to infrequent trading,” *Journal of Financial Economics*, 7, 197–226.

Easley, D., R. F. Engle, M. O’Hara, and L. Wu (2008): “Time-varying arrival rates of informed and uninformed trades,” *Journal of Financial Econometrics*, 6, 171–207.

Edelen, R. M. (1999): “Investor flows and the assessed performance of open-end mutual funds,” *Journal of Financial Economics*, 53, 439–466.

Einav, L. and A. Finkelstein (2011): “Selection in insurance markets: Theory and empirics in pictures,” *Journal of Economic Perspectives*, 25, 115–38.

Elton, E. J., M. J. Gruber, and C. R. Blake (2013): “The Performance of Separate Accounts and Collective Investment Trusts*,” *Review of Finance*, 18, 1717–1742.

Fama, E. F. and K. R. French (2012): “Size, value, and momentum in international stock returns,” *Journal of Financial Economics*, 105, 457–472.

Fedyk, A. (2018): “Front page news: The effect of news positioning on financial markets,” Tech. rep., working paper.

Ferreira, M. A. and P. Matos (2008): “The colors of investors? money: The role of institutional investors around the world,” *Journal of Financial Economics*, 88, 499–533.

Frydman, C., S. M. Hartzmark, and D. H. Solomon (2017): “Rolling Mental Accounts,” *The Review of Financial Studies*, forthcoming.

Frydman, C. and A. Rangel (2014): “Debiasing the disposition effect by reducing the saliency of information about a stock’s purchase price,” *Journal of economic behavior & organization*, 107, 541–552.

Frydman, C. and B. Wang (2018): "The Impact of Salience on Investor Behavior: Evidence from a Natural Experiment," *Journal of Finance*, forthcoming.

Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2018): "Channeled Attention and Stable Errors," *HBS Working Paper 18-108*.

Giglio, S. and K. Shue (2014): "No news is news: do markets underreact to nothing?" *The Review of Financial Studies*, 27, 3389–3440.

Gompers, P. A. and A. Metrick (2001): "Institutional investors and equity prices," *The quarterly journal of Economics*, 116, 229–259.

Gourville, J. T. and D. Soman (2007): "Extremeness seeking: When and why consumers prefer the extremes," *HBS Marketing Research Paper*.

Goyal, A., S. Wahal, and M. D. Yavuz (2020): "Choosing Investment Managers," *Available at SSRN 3651476*.

Green, E. and D. Daniels (2017): "Bayesian Instinct," *mimeo*.

Griffin, J. M. (2002): "Are the Fama and French factors global or country specific?" *The Review of Financial Studies*, 15, 783–803.

Griffin, J. M., J. H. Harris, and S. Topaloglu (2003): "The dynamics of institutional and individual trading," *The Journal of Finance*, 58, 2285–2320.

Grosshans, D., F. Langnickel, and S. Zeisberger (2018): "The Role of Beliefs in Trading Decisions," .

Gruber, M. J. (1996): "Another puzzle: The growth in actively managed mutual funds," *The journal of finance*, 51, 783–810.

Hansen, L. P. and R. J. Hodrick (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88, 829–853.

Hartzmark, S. M. (2014): "The worst, the best, ignoring all the rest: The rank effect and trading behavior," *The Review of Financial Studies*, 28, 1024–1059.

Hartzmark, S. M. and K. Shue (2017): "A tough act to follow: Contrast effects in financial markets," *The Journal of Finance*.

Hauser, J. R. (2014): "Consideration-set heuristics," *Journal of Business Research*, 67, 1688–1699.

Hauser, J. R. and B. Wernerfelt (1990): "An evaluation cost model of consideration sets," *Journal of consumer research*, 16, 393–408.

Heath, C. (1995): "Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting," *Organizational Behavior and Human Decision Processes*, 62, 38–54.

Hendershott, T., D. Livdan, and N. Schürhoff (2015): "Are institutions informed about news?" *Journal of Financial Economics*, 117, 249–287.

Hirshleifer, D. and S. H. Teoh (2003): "Limited attention, information disclosure, and financial reporting," *Journal of accounting and economics*, 36, 337–386.

Hogarth, R. M. (2001): *Educating intuition*, University of Chicago Press.

Hong, H., W. Torous, and R. Valkanov (2007): "Do industries lead stock markets?" *Journal of Financial Economics*, 83, 367–396.

Hu, G., K. M. Jo, Y. A. Wang, and J. Xie (2018): "Institutional trading and Abel Noser data," *Journal of Corporate Finance*, 52, 143–167.

Jensen, M. C. (1968): "The performance of mutual funds in the period 1945–1964," *The Journal of finance*, 23, 389–416.

Kahneman, D. (2003): "Maps of Bounded Rationality: Psychology for Behavioral Economics," *American Economic Review*, 93, 1449–1475.

Kahneman, D. and S. Frederick (2002): "Representativeness revisited: Attribute substitution in intuitive judgment," *Heuristics and biases: The psychology of intuitive judgment*, 49, 81.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1990): "Experimental tests of the endowment effect and the Coase theorem," *Journal of political Economy*, 98, 1325–1348.

Kempf, E., A. Manconi, and O. Spalt (2016): "Distracted shareholders and corporate actions," *The Review of Financial Studies*, 30, 1660–1695.

Kosowski, R., A. Timmermann, R. Wermers, and H. White (2006): "Can mutual fund ?stars? really pick stocks? New evidence from a bootstrap analysis," *The Journal of finance*, 61, 2551–2595.

Larson, F., J. A. List, and R. D. Metcalfe (2016): "Can Myopic Loss Aversion Explain the Equity Premium Puzzle? Evidence from a Natural Field Experiment with Professional Traders," Working Paper 22605, National Bureau of Economic Research.

Lleras, J. S., Y. Masatlioglu, D. Nakajima, and E. Y. Ozbay (2017): "When more is less: Limited consideration," *Journal of Economic Theory*, 170, 70–85.

Lou, D. (2012): "A flow-based explanation for return predictability," *The Review of Financial Studies*, 25, 3457–3489.

Malmendier, U. (2018): "Behavioral Corporate Finance," in *In Douglas Bernheim, Stefano DellaVigna and David Laibson, eds., Handbook of Behavioral Economics.*, Elsevier.

Malmendier, U., G. Tate, and J. Yan (2011): "Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies," *The Journal of Finance*, 66, 1687–1733.

Massey, C. and R. H. Thaler (2013): "The Loser's Curse: Decision Making and Market Efficiency in the National Football League Draft," *Management Science*, 59, 1479–1495.

Menkveld, A. J. (2013): "High frequency trading and the new market makers," *Journal of Financial Markets*, 16, 712–740.

Odean, T. (1998): "Are investors reluctant to realize their losses?" *The Journal of finance*, 53, 1775–1798.

——— (1999): "Do investors trade too much?" *American Economic Review*, 89, 1279–1298.

Pástor, L., R. F. Stambaugh, and L. A. Taylor (2017): "Do funds make more when they trade more?" *The Journal of Finance*, 72, 1483–1528.

Pope, D. G. and M. E. Schweitzer (2011): "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes," *American Economic Review*, 101, 129–57.

Puckett, A. and X. Yan (2011): "The interim trading skills of institutional investors," *The Journal of Finance*, 66, 601–633.

Romer, D. (2006): "Do Firms Maximize? Evidence from Professional Football," *Journal of Political Economy*, 114, 340–365.

Sakaguchi, H., N. Stewart, and L. Walasek (2017): "Selling Winners or Losers: Two-Stage Decision Making and the Disposition Effect in Stock Trading," *Available at SSRN 3053331*.

Schwartzstein, J. (2014): "Selective attention and learning," *Journal of the European Economic Association*, 12, 1423–1452.

Swensen, D. F. (2009): *Pioneering portfolio management: An unconventional approach to institutional investment, fully revised and updated*, Simon and Schuster.

UNGEHEUER, M. (2017): "Stock returns and the cross-section of investor attention," .

VON BESCHWITZ, B., S. LUNGHI, AND D. SCHMIDT (2017): "Limits of arbitrage under the microscope: evidence from detailed hedge fund transaction data," .

WERMERS, R. (1999): "Mutual fund herding and the impact on stock prices," *the Journal of Finance*, 54, 581–622.

——— (2000): "Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses," *The Journal of Finance*, 55, 1655–1695.

——— (2011): "Performance measurement of mutual funds, hedge funds, and institutional accounts," .

WERMERS, R., T. YAO, AND J. ZHAO (2012): "Forecasting stock returns through an efficient aggregation of mutual fund holdings," *The Review of Financial Studies*, 25, 3490–3529.