Selling Fast and Buying Slow: Heuristics and Trading Performance of Institutional Investors[*]

Klakow Akepanidtaworn[†] Rick Di Mascio[‡]
Alex Imas[§] Lawrence Schmidt [¶]

November 2018

Most research on heuristics and biases in financial decision-making has focused on non-experts, such as retail investors who hold modest portfolios. We use a unique data set to show that financial market experts – institutional investors with portfolios averaging $573 million – exhibit costly, systematic biases. A striking finding emerges: while investors display clear skill in buying, their selling decisions underperform substantially – even relative to strategies involving no skill such as random selling. We present evidence for limited attention as the driver of this discrepancy, with investors devoting more attentional resources to buy decisions than sell decisions. When attentional resources are more likely to be equally distributed between prospective purchases and sales – around company earnings announcement days – sells outperform similar to buys. Moreover, a salience heuristic explains much of the underperformance in selling: investors are prone to sell assets with extreme returns across all specifications. This strategy is a mistake, resulting in substantial losses relative to randomly selling assets to raise the same amount of money. In contrast to selling decisions, the salience heuristic does not appear to drive buying decisions, which are not affected by prior returns.

KEYWORDS: Heuristics, Behavioral Finance, Expert Decision-Making, Asset Pricing
JEL: D91, G11, G12, G14, G41

THIS IS A PRELIMINARY DRAFT. ALL COMMENTS VERY WELCOME!

[*]We thank Aislinn Bohren, Pedro Bordalo, Cary Frydman, Samuel Hartzmark, Zwetelina Iliewa, David Laibson, Ulrike Malmendier, Joshua Miller, Alex Rees-Jones, Richard Thaler and seminar participants at Booth School of Business, Florida State, University of Amsterdam and University of Frankfurt for helpful comments.

[†]Booth School of Business, University of Chicago: klakowa@uchicago.edu

[‡]Inalytics Ltd: rick@inalytics.com

[§]Social and Decision Sciences, Carnegie Mellon University: aimas@andrew.cmu.edu

[¶]Sloan School of Management, Massachusetts Institute of Technology: ldws@mit.edu

1 Introduction

In developing the efficient market hypothesis, Fama (1970) argued that prices reflect the true fundamental values of assets because rational investors trade away any opportunities to purchase undervalued stocks or sell stocks at inflated prices. Even if some investors do trade irrationally, sophisticated *experts* would trade against them and eliminate any temporary mispricing. Since then, a large literature has demonstrated that investors do indeed use heuristics and are prone to systematic biases. Individual investors have been shown to be overconfident (Barber and Odean 2001), sensation-seeking (Grinblatt and Keloharju 2009), and to have limited attention (Barber and Odean 2008). However, the majority of evidence documenting biased behavior of individual investors comes from data on retail investors (Barber and Odean 2011) or day traders (Barber, Lee, Liu, and Odean 2014), who generally hold modest portfolios.[1] As a result, it remains important to demonstrate the extent to which the decisions of market experts are prone to behavioral biases, and, if so, the effect of the resulting biases on performance.

This paper examines the trade decisions of sophisticated market participants – experienced institutional portfolio managers (PMs) – using a rich data set containing their daily holdings and trades. Our data is comprised of 783 portfolios, with an average portfolio (managed on behalf of a single institutional client) valued at approximately $573 million. More than 89 million fund-security-trading dates and 4.4 million 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. Since PMs often need to raise capital by selling existing positions in order to buy, evaluating a selling decision relative to a counterfactual which is unrelated to existing holdings (e.g., a benchmark index) is not an appropriate comparison.[2] Instead, we 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.

We document a striking pattern: while the investors display clear skill in buying, their selling decisions underperform substantially. Positions added to the portfolio outperform

[1]There are several notable exceptions: Frazzini (2006) and Jin and Scherbina (2010) present evidence for the disposition effect using individual-level data from mutual fund holdings. Coval and Shumway (2005) and Liu, Tsai, Wang, and Zhu (2010) present evidence for history-dependent risk-taking from market makers on the Chicago Board of Trade and the Taiwan Futures Exchange, respectively. Work has also documented behavioral biases amongst experts in corporate finance settings (see Malmendier (2018) for review).

[2]An asset sold may outperform a benchmark index, but the sale may still be optimal 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). In turn, a counterfactual must consider current holdings.

both the benchmark and a strategy which randomly buys more shares of assets already held in the portfolio. This result holds both in terms of raw returns and adjusted for risk. In contrast, selling decisions not only fail to beat a no-skill strategy of selling another randomly chosen asset from the portfolio, they consistently underperform it by substantial amounts. PMs forgo between 50 and 110 basis points over a 1 year horizon relative to this random selling strategy, depending on the specification.[3] As with buys, the selling result is robust when adjusting for systematic risk of the assets sold. The divergence in performance and skill across buy and sell decisions is surprising given the similarity between the two choices: optimizing over both involves forecasting the expected returns of the respective asset.

We present evidence that the discrepancy in performance between buy and sell decisions is driven by an asymmetric allocation of cognitive resources, particularly attention. When selling decisions coincide with salient information releases that draw attention to current holdings – company earnings announcements – sales outperform the counterfactual. It does not appear that the investors lack fundamental skills for selling: when relevant earnings announcements draw attention to particular holdings, selling decisions appear to capitalize on this information and outperform the counterfactual to a similar extent as when buying. Sales on non-announcement days consistently trail the random selling strategy. In contrast, the performance of purchases around announcement days does not meaningfully differ than those on non-announcement days. This is consistent with attention being generally allocated to buying decisions; the relevant information shock does not add much to the attention already devoted to making purchases.

Moreover, we find that much of the underperformance in selling can be explained by a salience heuristic associated with limited attention, which allocates more attentional resources to assets whose attributes deviate most from the average values within the choice set (Bordalo, Gennaioli, and Shleifer 2013). PMs in our sample have substantially greater propensities to sell positions with extreme 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. This strategy is a mistake. Systematically selling assets with extreme positions forgoes substantial earnings relative to a random selling strategy. Importantly, no such pattern is found on the buying side – unlike with selling, buying behavior correlates little with past returns and other observables. This suggests that PMs are purchasing assets based on private information not available to the researchers.

The tendency to sell positions with extreme returns is robust to numerous alternative

[3] As a benchmark, 100 basis points is often double the management fees charged to clients.

specifications. Using cumulative, annual and quarterly returns yields the same U-shaped selling pattern. The pattern emerges regardless of how the assets are grouped; PMs are more likely to sell assets with extreme returns when considering 6 bins of past returns (1st and 6th bins are most likely to be sold), or 20 bins of past returns (1st and 20th bins are most likely to be sold). Moreover, assets with extreme returns are more likely to be sold when conditioning on other, less salient, attributes. We find the same U-shaped pattern regardless of the assets' weights in the portfolio or how long they have been held. Similar to a test proposed in Hartzmark (2014), the pattern persists even after the inclusion of stock-date fixed effects and thus absorbs a number of time-varying stock-specific unobservables.

In addition to documenting the use of heuristics in selling decisions but not buying decisions, we also demonstrate that the salience heuristic is costly. A large literature has shown that heuristics often lead to systematic deviations from optimal behavior – costly biases that result in worse outcomes than counterfactuals (Kahneman 2003). We find that the tendency to sell extreme positions leads to substantially lower returns than the counterfactual of a random selling strategy both *between* and *within* manager. Comparing managers based on their proclivity to sell positions with extreme returns, PMs who have a more pronounced U-shaped selling pattern have substantially worse selling outcomes than those with a less pronounced pattern; PMs in the top quartile of heuristic use forgo more than 140 basis points per year relative to a random selling strategy. Changes in the proclivity to sell extremes within manager also yield similar magnitudes, explaining substantial variation in poor performance. For PMs prone to using the heuristic, simply adopting a random selling strategy would generate greater earnings than the average management fees charged to clients.

Lastly, prior work has shown that reliance on heuristics increases when cognitive resources are in greater demand (Kahneman 2003). We provide evidence that, consistent with these findings, sell decisions are particularly poor when attentional resources are likely to be stretched. We hypothesized the PMs would devote fewer attentional resources to specific sell decisions when 1) he or she is selling a large number of distinct assets and 2) during periods when cash is likely being raised for a purchase. We found that episodes involving more unique assets being sold are associated with worse selling performance compared to selling episodes involving fewer unique assets (relative to the number of assets in the portfolio). We also found that more pronounced selling episodes emblematic of raising cash – measured within-manager, relative to their average selling rate – are associated with greater heuristic use and lower selling performance.

Taking stock, the robustness of the U-shaped selling pattern to alternative specifications, the costs associated with this behavior, and the decrease in selling performance when attentional resources are likely to be stretched all suggest that pecuniary motives such as agency concerns are unlikely to be driving the selling of positions with extreme returns. Instead, these results point to a salience heuristic as the driver of poor selling performance.

A key innovation of our analysis is examining buying strategies separately from selling strategies. The findings shed light on the nature of expertise in financial markets. While a large fraction of fund managers display prominent biases in selling that hurt performance, we find no evidence of such biases in their buying decisions. This is surprising given that the two decisions are similar operationally – the outcomes of both are a function of future asset returns of the asset. They also have equivalent consequences for trading profits: the 'quality' of a buy decision is evaluated relative to alternatives in the choice set (not purchasing, purchasing other assets), as is a sell decision (not selling, selling other assets). Yet we do not observe traces of a salience heuristic in buy decisions, which unlike sell decisions, outperform counterfactuals such as the benchmark.

Our results suggest that PMs systematically fail in porting their expertise in buying to selling decisions, consistent with the concept of fractionation of expertise (Kahneman and Klein 2009). Fractionation refers to individuals who attain expertise in one domain failing to successfully port these skills to other domains, even if those domains are nearly identical. Moreover, people often fail to recognize this lack of portability, and make decisions as if expertise attained in one domain makes them experts in other domains as well. This lack of awareness can lead to overconfidence in decision-making and a reluctance to embrace learning opportunities that would be helpful for attaining expertise in the related domains.[4] In Section 6 we discuss the potential role of learning environments in the development of expertise for buying assets in our setting, and strategies for porting the expertise to selling decisions.

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. While it is not clear from the data whether these trading strategies are particularly maladaptive, the investors from this data set tend to underperform the market in general, and have been shown to

[4]In a recent demonstration of this concept, Green, Rao, and Rothschild (2017) asked expert decision makers, ESPN analysts, to make judgments on outcomes that their profession is focused on – forecasting NBA playoff games. This population of experts has been shown to outperform novices on similar predictions both in terms of accuracy and consistency of forecasts. The same experts were then given an isomorphic task that kept the fundamental structure of playoff forecasting, but changed the contextual cues. Despite the tasks sharing identical properties in terms of the underlying uncertainty, expert forecasts were substantially worse on the second task than the first.

display a host of heuristics and bias such as the disposition effect (Odean 1998), overconfidence (Odean 1999), and narrow bracketing (Frydman, Hartzmark, and Solomon 2017).[5]

Our findings contribute to the literature in finance documenting biased decision-making in individual investors (see Barber and Odean (2011) for review). While prior work has documented biases amongst experts in corporate finance settings, e.g. CEOs in charge of merger (Malmendier, Tate, and Yan 2011) or other restructuring decisions (Camerer and Malmendier 2007), substantially less research exists on the biases of expert investors.[6] In fact, for the most part the behavioral finance literature has assumed unbiased institutional investors exploiting the behavioral biases of retail investors (Malmendier 2018). Our documented findings suggest that this assumption may not be a valid one. Lastly, our results 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), and retail markets (DellaVigna and Gentzkow 2017). This line of work highlights the persistence of behavioral biases despite significant experience and exposure to market forces. As our findings demonstrate, in the presence of cognitive constraints, the extent to which expert market participants are prone to behavioral biases depends on the context in which these decisions take place.

The paper proceeds as follows. Section 2 describes the data. In Section 3 we outline the role of salience in investment decisions and outline hypotheses for our setting. Section 4 presents results on performance of buying and selling decisions, while 5 presents results on heuristics and trading strategies, and how those strategies affect performance. Section 6 discusses our results and concludes.

2 Data and Methodology

This section discusses the data sources which are assembled for our analysis, presents descriptive statistics, and discusses a number of portfolio and position-specific variables which we use throughout the analysis. Section 2.3 outlines our methodology for computing counterfactual portfolio returns and value-added measures, while Section 2.4 describes our measures

[5]Though Hartzmark (2014) focuses on the behavior of retail investors, he also present evidence that mutual funds are prone to such behavior as well. However, due to the limitations of the data, which comes from quarterly holdings reports, he notes that the behavior can be driven by strategic concerns in response to investor preferences.

[6]One exception to this is a literature which emphasizes slow/inefficient incorporation of certain types of *aggregate* signals into asset prices; see, e.g., Chang, Hartzmark, Solomon, and Soltes (2016); Giglio and Shue (2014); Hartzmark and Shue (2017); Hong, Torous, and Valkanov (2007).

of PMs' use of a salience heuristic.

2.1 Data sources and sample selection

Our primary source of data for this analysis is compiled by Inalytics Ltd. These data, which were first introduced to the literature and are discussed in greater detail by Di Mascio, Lines, and Naik (2017), include information on the portfolio holdings and trading activities of institutional investors.[7] Inalytics acquires this information as part of one of its major lines of business, which is to offer delegated portfolio monitoring services for institutional clients that analyze the investment decisions of portfolio managers.[8] Our dataset includes both active and inactive portfolios, and the vast majority of the data are collected essentially in real-time, suggesting that incubation and survivorship biases are unlikely to be a substantial concern for our analysis.

For purposes of this study, Inalytics assembled an extract of data of long-only equity portfolios, spanning from January 2000 through March of 2016. In our data, the names of funds and managers are anonymized – only a numerical identifier for each fund is provided. These portfolios are internationally diversified, including data from a large number of global equity markets. Data are only collected during periods for which Inalytics' monitoring service is performed, leaving us with an unbalanced panel.

For each portfolio, we have a complete history of holdings and trades at daily level throughout the time period that the fund subscribes to Inalytics' service. Inalytics 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. Further, each portfolio is associated with a specific benchmark (usually a broad market index) against which its performance is evaluated – a feature we exploit heavily throughout our analysis.

We apply two primary filters to select the set of portfolios to include in our analysis. First, daily trading data are unavailable for a subset of portfolios or appear to be incom-

[7] While Di Mascio et al. (2017) also use Inalytics data and compute several performance measures, their focus is quite distinct and complementary to ours. Whereas our primary emphasis is on understanding differences between PMs' buying and selling strategies, their primary emphasis is on buying behavior. To this end, they find evidence that PMs' buying strategies outperform and are consistent with predictions of theoretical models of optimal strategic trading with private information.

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

Table 1. Summary statistics of the analysis dataset

This table reports the summary statistics of the analysis dataset for 783 portfolios at various levels of aggregation. The position level summary statistics include various holding lengths, portfolio weights, future return measures and the number of trades (indicator for buy and sell trades). Future returns are reported in percentage points over specified horizons. The fund-level and position-level summary statistics are reported at monthly and daily frequencies, respectively. See table 2 and text for additional details on variable construction.

Variable	Count	Mean	Std	25th	50th	75th
Panel A: Fund level Summary (monthly)						
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	0.927	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
Fraction of distinct stocks bought minus fraction of distinct stocks sold over all holdings (%)	51221	4.675	16.87	-0.691	1.852	7.030
Monthly benchmark-adjusted returns (%)	48786	0.217	1.767	-0.599	0.165	1.010
SD of daily benchmark-adjusted returns (%)	48041	0.348	0.208	0.205	0.293	0.431
Loading on Market	48705	0.971	0.259	0.807	0.943	1.121
Loading on SMB	48705	0.00669	0.497	-0.320	-0.0624	0.271
Loading on HML	48705	-0.0636	0.503	-0.358	-0.0655	0.215
Loading on Momentum	48705	0.0447	0.336	-0.133	0.0430	0.221
Heuristics Intensity	47335	0.404	0.240	0.267	0.385	0.522
Panel B : Position Level Summary (daily)						
Buying indicator	89.8M	0.0264	0.160	0	0	0
Selling indicator	89.8M	0.0226	0.149	0	0	0
Holding length since position open (days)	89.8M	484.4	512.9	119	314	679
Holding length since last trade (days)	89.8M	73.36	113.5	10	32	88
Holding length since last buy (days)	89.8M	112.3	152.4	18	57	144
Portfolio weight(%)	89.7M	1.2	1.61	.24	.79	1.65
1-day return (%)	82.1M	0.0511	4.15	-1.11	0.0115	1.17
Future 7-day return (%)	82.9M	0.205	5.830	-2.454	0.179	2.833
Future 28-day return (%)	82.8M	0.781	11.04	-4.634	0.810	6.181
Future 90-day return (%)	82.6M	2.561	20.16	-7.711	2.308	12.30
Future 180-day return (%)	81.5M	5.315	30.51	-10.46	4.164	18.88
Future 270-day return (%)	80.3M	7.873	38.54	-13.10	5.562	24.47
Future 365-day return (%)	78.9M	10.37	44.84	-15.08	7.241	29.73
Future 485-day return (%)	76.9M	13.43	51.12	-16.81	9.006	35.60
Future 605-day return (%)	74.9M	16.73	58.82	-18.73	9.871	41.01
Future 665-day return (%)	73.9M	18.53	62.94	-19.55	10.32	43.66
Future 730-day return (%)	72.7M	20.40	66.82	-20.13	10.86	46.43
Earning announcement day indicator	49.3M	0.007	0.08	0	0	0

plete.[9] Second, we exclude funds that do not have a sufficient fraction (at least 80%) of portfolio holdings which could be reliably matched with CRSP or Datastream, which we discuss further below. After applying these screening procedures, our final sample includes about 51 thousand portfolio-months of data, which are compiled from a set of 783 institutional portfolios. Summary statistics are presented in Table 1. We have an average of just over 5 years (65 months) of data per portfolio. During this time frame, we observe 89 million fund-security-trading date observations and 4.4 million (2.4 million buy and 2 million sell) trades. We convert all market values to US dollars at the end of each trading day.[10]

This sample offers some unique opportunities for the study of 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 and has been used in numerous studies[11], our data include complete portfolio and trade-level detail for a population of professional investors who are tasked with managing large pools of assets. For instance, 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 of roughly $150,000, whereas 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, we are able to observe changes in portfolio holdings at a *daily* level. This facilitates the testing of hypotheses that is infeasible with quarterly data. Most other datasets with institutional trading information often lack timely information on portfolio holdings.

To complement these data, which characterize portfolios and trades at specific points in time, 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.

[9]Trades are sometimes imputed at month-end because Inalytics receives portfolio snapshots in adjacent months which do not fully match with the portfolio which would be expected from aggregating the trade data, which necessitates a reconciliation process. We exclude funds that have a large fraction of trades occurring at the end of each month.

[10]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.

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

Table 2. Characteristics summary

This table reports the construction of characteristics in our analysis. The first column reports the variables, the second column reports the frequency that we compute the variables and the type of sorting methods (across-fund or within-fund) used in the analysis. The third column reports the formula or the description on the construction of sorting variables.

Characteristics	Sorting	Construction
Cumulative Returns capped at 1-year	Within Fund-date across stocks	$r^{cum}_{s,f,t} = \Pi^{i=t}_{i=t-\min\{365,d\}}(1 + r_{s,f,t}) - 1$, where d is the time since a position is open.
Position past k day returns	Within Fund-date across stocks	$r^{past\ k}_{s,f,t} = \Pi^{i=t-1}_{i=t-k}(1 + r_{s,f,t}) - 1$.
Fund past k day returns	Across funds on daily basis	$r^{k}_{f,t} = \Pi^{i=t-1}_{i=t=k-1-1}(1 + r_{,f,t}) - 1$.
Heuristics Intensity	Across/Within funds on weekly/monthly basis	$\frac{\text{Total \# of Position sold in Bin 1 or Bin 6 of past returns}}{\text{Total \# of Position Sold}}$.
Gross Sell	Within funds on weekly basis	# of Positions sold.
Net Buy	Within funds on weekly basis	# of stocks bought - # of stocks sold.
Monthly Turnover	Across funds on monthly basis	$turnover_{f,m} = \frac{\min\{total\ MarketValue^{buy}_{f,m} total\ MarketVvalue^{sell}_{f,m}\}}{\overline{MarketValue_{f,m}}}$.
Position Size	Within Fund-date across stocks	$PositionSize_{s,f,t} = \frac{Quantity^{beginiing\ t}_{s,f,t} \times P_{s,f,t}}{Fund\ AUM_{s,f,t}}$.
Holding length last buy	Within Fund-date across stocks	# of trading days from last day on which a position was bought

2.2 Fund and position-level characteristics

With these data in hand, we construct a wide array of measures at the portfolio-time and portfolio-stock-time (position) level. Formulas for many of these variables are presented in Table 2. We begin by discussing some characteristics of fund portfolios in our sample; these are summarized in Panel A of Table 1 on a monthly basis. All portfolios are large, and there is considerable heterogeneity in portfolio size. In addition, funds differ noticeably in terms of their levels of trading activity. Average monthly turnover is about 4% of assets under management, but some funds are considerably more active in their trading behavior than others (the standard deviation is 5.7%).

While holding fairly diversified portfolios (average number of stocks is about 78 with a standard deviation of 68), funds in our sample remain active, holding positions that deviate substantially from their benchmarks. The average tracking error–the standard deviation of the difference between the daily portfolio return and the benchmark–is about 0.35% per day, or about 5.7% on an annualized basis. On average, a manager will initiate a sell trade for about 10% and a buy trade for about 15% 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.[12] The average market beta is about 1, and average exposures to the SMB, HML, and Momentum factors are fairly close to zero.

Panel A also reports the average benchmark-adjusted return, where we use each portfolio-specific return series. The average fund in our sample beats its respective benchmark by about 0.22% per month, or 2.6% per year. This, in conjunction with the fact that funds' average betas are close to 1 and have little average exposure to the 3 other priced risk factors, suggests that these managers are highly skilled, earning returns above and beyond exposure to known risk factors. In turn, demonstrating systematic biases and the resultant negative effects on performance could be interpreted as a lower bound for extending the results to broader populations.

Next, we turn to our position-level data. Our simplest position-level variable is an indicator variable which equals 1 if the manager buys or sells a given stock on a given date. Of the 89 million position-date combinations in our sample where a stock was in the portfolio at either the start or end of the day, about 2.4 million of them involved an active purchase decision on that same day and 2 million of them involved active sell decisions, or about 2.6% and 2.2% of the time, respectively.

We compute three other primary measures at the position level. First, we construct several different measures of the holding length associated with a given position. Specifically, we consider the length of time (in calendar days) elapsed since the position was first added to the portfolio. In many case, this measure will be censored because a stock may have been in the portfolio since it was first added to our sample. The average holding length is 485 calendar days (or about 15 months), though this measure is downward-biased. As such, we also examine holding length measures which consider the time elapsed since a stock was most recently bought (or traded). The average position was last purchased about 112 calendar days (a bit less than 4 months) ago and was last traded about 10 weeks ago. In much of the analysis that follows, we will exclude stocks which were very recently bought to avoid having our results being driven by predictable buying (and lack of selling) behavior as managers split trades over several days while building up positions over time. Second, we compute the portfolio weight as a fraction of market value associated with each position on each date. The average stock has a weight of about 1.2% with a standard deviation of 1.6%. Finally, we compute a number of measures of backward or forward-looking returns at the position

[12] Following Dimson (1979), we adjust for asynchronicity by including 1 lag and 1 forward returns of each factor.

level over various horizons, both overall and relative to the benchmark return. With the exception of 1-day measures (which refer to the prior trading day), we measure horizons in calendar days.[13] For brevity, we only report summary statistics for forward-looking returns that are not adjusted for the benchmark. Volatilities of individual stocks are quite large, with a standard deviation of 45% at a 1 year horizon. As we discuss further below, we also consider several measures of prior position performance that are computed using periods of time which depend on holding period length.

2.3 Constructing counterfactuals

This section outlines how we construct counterfactual strategies in order to evaluate trade performance, which is greatly facilitated by the availability of daily holdings information.

Given that the portfolio managers 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 may be far from innocuous. Thus, precisely if managers have useful private information that 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. Note that this issue is less important when considering unskilled investors; in that case, we would expect them to neither to gain or lose money (on a risk-adjusted basis) by relying on a simple rule of thumb for selling existing positions.

The ability to observe daily transactions allows us to compare observed buy and sell decisions to counterfactual strategies constructed using portfolio holdings data. Our measures correspond to the relative payoffs from two hypothetical experiments: one for evaluating buy decisions and one for evaluating sell 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 manager invests that money in a randomly selected stock from his other holdings. For evaluating sells, suppose that we learned that the manager was planning to sell a given stock tomorrow and hold the rest of the portfolio for a fixed period of time. We then suggest that

[13]This choice is, in part, motivated by the fact that trading calendars differ slightly across exchanges. We take a number of precautions to reduce the potential influence of measurement errors in prices, including winsorizing 0.1% of returns in either tail by date. These steps are discussed at greater length in the Appendix.

instead of executing this trade, the manager randomly sells one of his/her other positions to raise the same amount of cash, holding the stock that was to be sold for the same period.

Since the information being used by us was also available to the manager, we would expect the decisions of a skilled PM to outperform our suggested strategies; this is due to the fact that, on the margin, our strategies are always feasible.[14] Note that the expected payoff from the counterfactual strategy (integrating out uncertainty about which stock is randomly selected) simply corresponds to the equal-weighted mean of realized returns across stocks held in the portfolio, which we denote by R_{hold}. The manager's selection adds value relative to the random counterfactual if $R_{buy} - R_{hold} > 0$ in the first example and if $R_{hold} - R_{sell} > 0$ in the second example.[15] Following this logic, we compute $R_{buy} - R_{hold}$ and $R_{hold} - R_{sell}$ over horizons ranging from 1 week to 2 years for all buy and sell trades, respectively, to characterize the value-added associated from each.

We can also use additional information to construct a "more intelligent" counterfactual. As we show in Figure 4 below, very few managers elect to sell stocks that were very recently purchased. In constructing the counterfactual, we exclude stocks which are in the bottom quintile of the distribution of holding length since last purchase; across all specifications, results are similar if we include them as well. 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 on every day and are not necessarily present throughout our sample period, this averaging procedure yields a portfolio-day unbalanced panel. Since some funds trade much more frequently than others–see the dispersion in monthly turnover in Table 1– we weight observations inversely to a measure of trading frequency.[16]

Note that one disadvantage of this value-added measure is that our use of long horizon returns introduces an overlapping structure in the error term of each fund's value-added time series. To address this concern, we compute heteroskedasticity and autocorrelation robust standard errors using a panel version of the Hansen-Hodrick (1980) correction using a lag of

[14]In contrast, selling the benchmark to finance a new idea, 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 conviction positions and thus deviates substantially from the benchmark. Purchasing the benchmark is feasible on the other hand.

[15]These measures also directly correspond to changes in benchmark-adjusted returns associated with different trading strategies. To the extent that buy and sell trades are not motivated by a desire to change a portfolio's systematic risk exposures, we would expect factor loadings of the assets being traded and the hold portfolio to be similar. If so, these measures would also correspond to differences in risk-adjusted returns (i.e., alpha).

[16]In our baseline analysis, we weight observations inversely to the number of trading days in a calendar year that the fund buys a stock. For easier comparison across buys and sells, we use the same weights across buys and sell trades. Unweighted results are reported in the Appendix. They are qualitatively similar in terms of direction and statistical significance as the weighted results.

the horizon minus 1.[17] This allows for individual fund time series to be serially correlated but assumes that these value-added measures are cross-sectionally independent across funds and across non-overlapping periods of time within funds.

2.4 Measuring heuristic use

We construct the following measure to examine whether past returns drive trading decisions in line with salience heuristic. 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 an empirical proxy for salience (past returns), then ask whether managers are more likely to trade the more salient holdings.

Given the size of our dataset, we adopt a fairly 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. Then, on each trading date, we sort stocks into N_{bin} bins using these relative rankings. We always choose an even number of bins and always set the breakpoint between bins $N_{bin}/2$ and $N_{bin}/2+1$ equal to zero. This ensures that all stocks in bins $N_{bin}/2$ have declined relative to the benchmark. We choose all remaining breakpoints so that (ignoring issues related to discreteness) there are equal numbers of stocks in bins $1,\ldots,N_{bin}/2$ and bins $N_{bin}/2+1,\ldots,N_{bin}$. As a baseline, we consider $N_{bin}=6$ and name the first three bins "Worst loser", "Loser", "Slight Loser", respectively, and adopt analogous naming conventions for bins three through six.

Our preferred measure of prior returns is computed as follows. For positions which were opened more than 1 year prior to the date of interest, we use the benchmark-adjusted return of the stock from 365 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. We use this as our preferred measure because it is unclear whether large returns relative to benchmarks more than 1 year in the past are likely to be salient attributes for the PMs. From a more pragmatic perspective, this construction is less sensitive to the censoring issues for holding length discussed above. However, as we show in Section 5, results are robust to alternative definitions of past returns.[18]

We make one substantive restriction on the sample of stocks which are under consideration

[17] We compute these standard errors using the ivreg2 package in Stata.

[18] We find nearly identifcal results if we restrict attention to stocks with opening dates that are observed during our sample.

for this analysis. In predicting the probability that a manager will add to/reduce an existing position, we exclude stocks that were bought in the very recent past. Specifically, we sort positions into 5 bins based on the holding length since the last buy trade, and exclude the bottom bin (shortest time elapsed since last purchase) from our calculations. We elect to do this to avoid a fairly mechanical relationship between our prior return measure, which has a variance which shrinks with the holding period, and the probability of buying/selling that can be generated if managers build up positions by splitting buy trades over short windows of time in order to minimize price impacts.[19] Such trades likely originate from a single purchase decision being executed over time, and so we construct our measures to treat them as such. Further, to ensure meaningful distinctions between bins, we exclude fund-dates which include fewer than 40 stocks in the portfolio throughout the analysis in this section, though results do not meaningfully change without such a restriction.

3 Hypotheses

In this section, we discuss buying decisions separately from selling decisions, and apply the insights from the literature on limited attention to develop hypotheses on the trading strategies in our setting.

A skilled manager of a long-only (short-sale constrained) portfolio can add value in two ways: 1) by identifying and increasing positions in undervalued assets and 2) by identifying assets within his/her portfolio which offer less attractive upside potential, reducing these positions in favor of more attractive alternatives. Based on extensive interviews with PMs, the first emerges as central: identifying new investment opportunities is seen as perhaps the most critical aspect of a PM's role. Moreover, the decision to add an asset to the portfolio or substantially increase the size of an existing position often follows lengthy periods of research and deliberation. In contrast, there was substantially less emphasis on decisions of what to sell, which involves research into which existing positions are likely to underperform going forward – i.e., finding stocks whose "investment thesis has changed."[20]

If PMs are subject to constraints on their ability to acquire and process information – they have limited cognitive resources such as attention – it would follow that more of those

[19] This phenomenon mechanically tends to increase the likelihood that positions with non-extreme returns are bought and decrease the likelihood that they are sold, since a manager is unlikely to sell an asset immediately after or while actively building a position in it. Related to this concern, in addition to imposing this selection criterion, our analyses 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.

[20] 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."

cognitive resources are likely to be devoted to buying decisions than selling decisions. This asymmetric allocation of resources lends itself to several testable predictions.

The first is that the performance of buying decisions will surpass that of selling decisions relative to a counterfactual. To the extent that attention and costly information processing yields positive returns, activities associated with higher allocations of those cognitive resources (buys) will result in better performance than those associated with lower allocations (sells). Moreover, if limits on attentional resources drive performance differences, we should observe a mitigation of performance discrepancies between buy and sell decisions during periods when attention to portfolio holdings is likely to increase (e.g. days with earnings announcements).

Hypothesis 1. *Selling decisions underperform relative to buying decisions. Differences in performance are mitigated during periods where more attention shifts more towards selling.*

Second, if limited attention is responsible for the discrepancy in performance between buying and selling decisions, then selling strategies will be more prone to heuristics that typify a lack of attentional resources, such as those that overweigh salient features of the choice environment. Bordalo, Gennaioli, and Shleifer (2012) and Bordalo et al. (2013) develop a theoretical framework where individuals attach disproportionately high weights to salient attributes of a good or lottery. An attribute's salience is a function of the availability of the relevant information, either from the environment or from memory (Bordalo, Gennaioli, and Shleifer 2017), and the extent to which the values of this attribute deviate from the attribute's average value in the choice set.

Relative to other forms of information relevant to the decision problem, such as forecasted returns, data on past returns is ubiquitously available to PMs in our setting. This information is prominently featured on trading terminals, which typically break down past returns by year, quarter, month, day and since last purchase. Most news programs and popular webpages that cover financial markets include a segment which covers the stocks which experienced the largest moves on a given (both positive and negative).[21] The availability of this information, as well as the large range of values past returns take relative to the portfolio average (as captured by their standard deviation, 51% over the average holding period in our sample), makes it highly likely that past returns are a very salient attribute of a given asset. For investors prone to the salience heuristic, assets with extreme past returns will be more likely to be present in the consideration set of positions to be sold or not. In contrast, because

[21] See Kumar, Ruenzi, and Ungeheur (2018) for discussion of media focus on past returns.

attentional resources are hypothesized to be less constrained for buying decisions, past returns should not predict them.

Hypothesis 2. *Past returns predict selling decisions but not buying decisions.*

The extent to which such a salience heuristic is adaptive, optimally factoring past returns into a PM's decision to sell an asset, is unclear. For instance, the strong form of the Efficient Markets Hypothesis (Fama 1970) posits that past price movements should not predict future risk-adjusted returns across stocks. The present value identity (Campbell 1991; Campbell and Shiller 1988) and a vast literature on return predictability of the aggregate stock market suggests that a substantial fraction of broad market returns is driven by time-variation in expected returns (discount rate news). However, this result may not apply when considering relative valuations across stocks. To this point, Vuolteenaho (2002) performs an extensive variance decomposition exercise using data on individual stocks and finds that the majority of firm-level stock market variance is driven by cash-flow news – that is, news about valuations which do not affect future expected returns – where these shocks are largely idiosyncratic. Further, the fraction of variance explained by discount rate news is highly correlated across stocks. This suggests that the majority of the variation in relative returns is likely to be driven by idiosyncratic cash flow news (and, accordingly, unrelated to the manager's objective function) rather than information captured by past returns of individual stocks.

If the reliance on past returns in selling decisions is a bias – a systematic deviation from optimal behavior – then the outcomes of decisions driven by the bias should be worse than those less driven by the bias.

Hypothesis 3. *Greater propensity to sell positions with extreme returns will result in worse outcomes.*

This prediction should hold both when comparing *individuals* – those who are more prone to the bias will perform worse than those who are less prone to it – and when comparing *decisions* of the same individual – decisions that are more prone to the bias will result in worse outcomes than decisions less prone to the bias.

Lastly, as the literature on heuristics and biases has documented (Kahneman 2003), episodes that engage greater cognitive resources will lead to a greater reliance on heuristics. In our setting, this predicts that episodes which divert further attention from individual selling decisions – such as periods where more unique assets are being sold or larger amounts of capital are needed for purchasing stocks – will be associated with worse performance.

Hypothesis 4. *Factors that further engage attentional resources away from selling decisions will lead to lower relative returns on those decisions.*

We now proceed to present evidence for these predictions.

4 Overall Trading Performance

Having described the basic properties of our dataset and variable construction procedures, we now turn to our tests of the hypotheses developed above. Here, we present the first of our empirical results, which calculate the average value-added (or lost) associated with managers' active buying and selling decisions.[22]

Figure 1, Panel A shows the average counterfactual returns for buying decisions. As will turn out to be the case across the vast majority of our specifications, we find very strong evidence that buy trades add value relative to our random buy counterfactual, $R_{buy} - R_{hold}$. The average stock bought outperforms stocks held by 0.6% over one year and 0.87% over two years.

Figure 1, Panel B presents the average value-added, $R_{hold} - R_{sell}$, for sell trades. 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, consistent with Hypothesis 1. Magnitudes are quite substantial: the value lost from an average sell trade is on the order of 100 basis points at a 1 year horizon relative to a simple counterfactual which randomly sells other stocks held on the same day.

In Section 3 we hypothesized that the results in Panel B are due to managers' allocation of cognitive resources rather than a fundamental difference between buying and selling decisions. The latter explanation seems quite unlikely given that payoffs from buying and selling are mirror images of one another. Indeed, Figure 2 – which depicts the outcomes of buy and sell decisions on days when attention is more likely to be devoted to current holdings – provides initial evidence that poor selling performance is likely due to a lack of attentional resources devoted to the task rather than a fundamental inability to sell. Specifically, we show that managers' sell decisions actually do add value on those days when more attention is likely devoted to them. 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 earnings

[22] We will return to this analysis in more depth in Section 5.2 below, which will link other position and fund-characteristics with predictable differences in trading performance.

Figure 1. Overall average post-trade returns relative to counterfactual

This figure presents weighted average returns relative to random buy/sell counterfactuals for buy and sell portfolios. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. We then compute the average of these performance measures across all portfolios and dates, weighted inversely to funds' trading activity. Each bar represents average counterfactual returns in percentage over specified horizons on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Panel A: Buy Trades, weighted

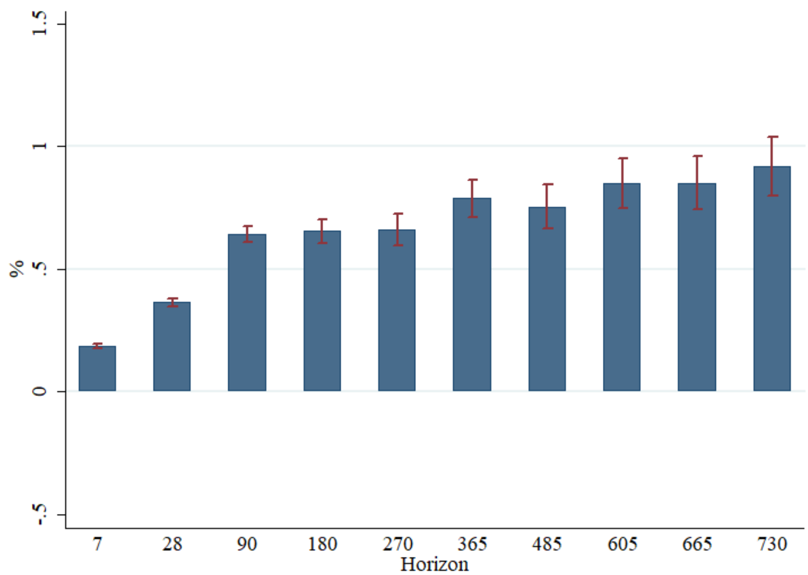

Panel B: Sell Trades, weighted

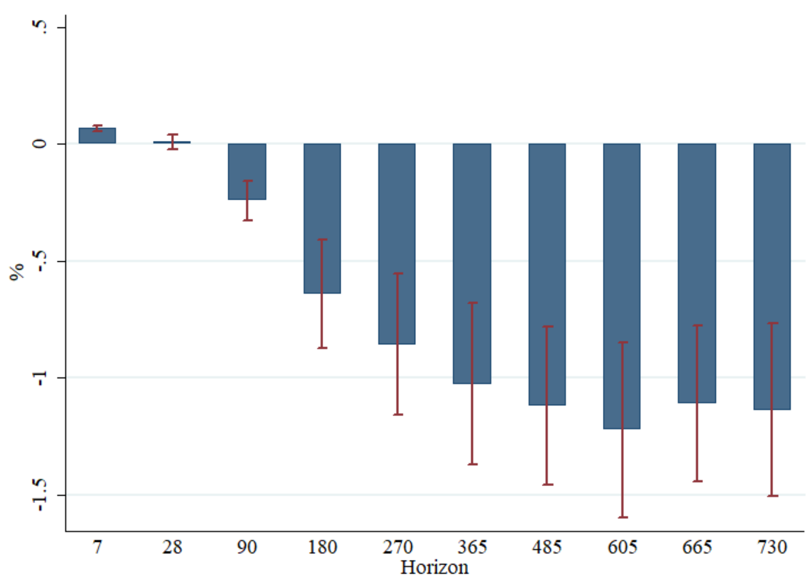

Figure 2. Trading performance on earnings announcement days vs other days

This figure presents weighted average returns relative to random buy/sell counterfactuals for overall buy and sell trades that take place on firm's earning announcement days (red bars) vs trades that are executed on all other days (blue bars). Earning announcement dates are taken from the I/B/E/S database. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. We then compute the average of these performance measures across all portfolios and dates, weighted inversely to funds' trading activity. Each bar represents average counterfactual returns in percentage over specified horizons on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Panel A: Buy Trades

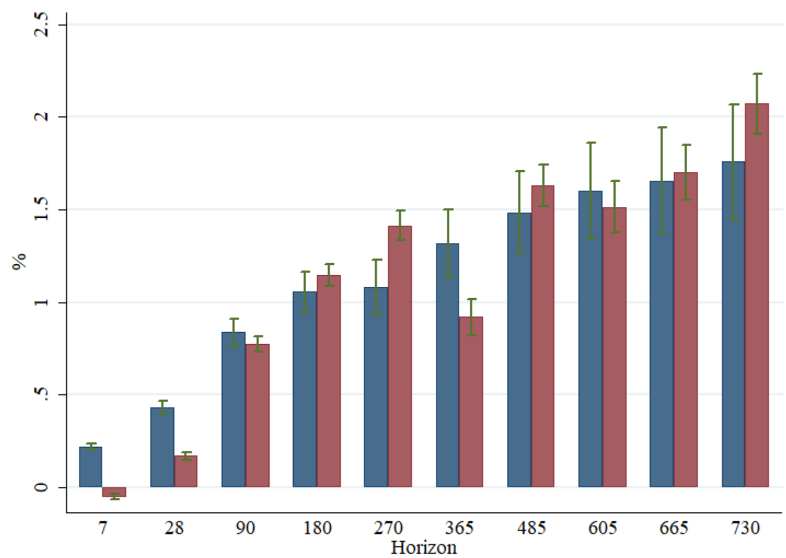

Panel B: Sell Trades

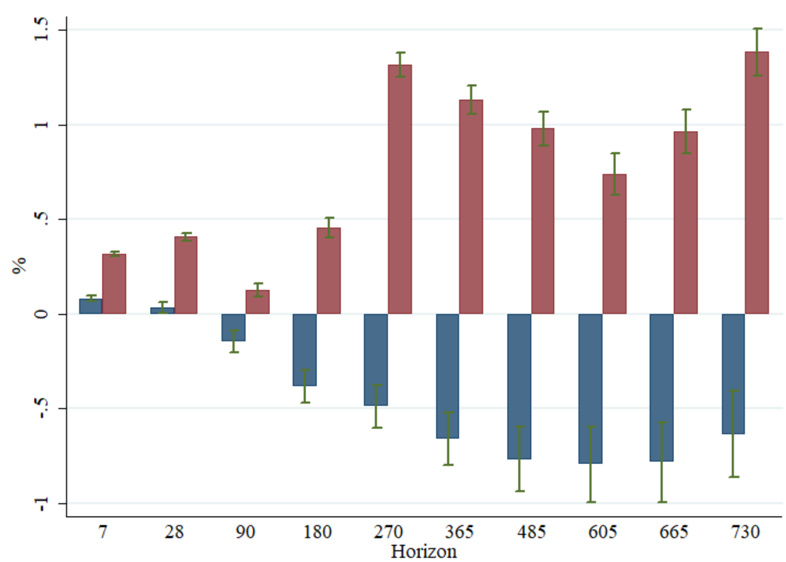

announcement days, relative to all other trading days. Managers have a strong incentive to pay close attention to stocks in their portfolios on these dates for several reasons. The information in financial statements, associated press releases and conference calls (which even offer opportunities for managers to directly ask questions to the company), provide a wealth of new pieces of hard and soft information that are likely to be relevant for firm valuations. This information is both freely available and salient, since earnings announcements are heavily covered by the financial press.

Figure 2, Panel A looks at the value-added of buy trades executed on earnings announcement days compared to other days. In both cases, averages are positive, and, consistent with attentional resources *already* being devoted towards purchase decisions, magnitudes are similar on both types of days. Panel B performs the same comparison for sell trades. On non-announcement days, results are similar to Panel B of Figure 1: observed sells substantially underperform a random sell strategy.[23] However, stocks sold on announcement days are associated with substantial value-added, especially at longer horizons. In sum, when the manager has access to a cheap, valuable source of information and is likely to be paying close attention, we observe no asymmetry between buy and sell trades, providing further evidence for Hypothesis 1.

5 Heuristic Use and Trading Performance

In this section, we provide evidence for the use of a salience heuristic based on past returns, then establish an empirical link between the use of this heuristic and performance.

5.1 Past Returns and Heuristic Use

Table 3 summarizes our primary result on PMs' heuristic use. Each row reports the fraction of existing positions that are bought or sold within each of the six bins formed on prior position returns, so that positions with the least salient attributes according to our measure appear in the center of the table and the most salient ones appear at the edges.[24] The fraction of assets bought in each bin is the first row, and the fraction of assets sold in each bin is depicted in the second row.

We begin with the buying probabilities. Consistent with the discussion above and the second part of Hypothesis 2, the probability of purchasing a stock already held is quite flat

[23] Magnitudes differ somewhat because the sample composition is limited to stocks that can be linked with I/B/E/S for this exercise.

[24] 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.

Table 3. Probability of buying and selling based on past position returns

This table reports the buying and selling probability (in percentage points) by six bins of past benchmark-adjusted returns capped at one year. We create bins based on past cumulative returns of a position capped at 1-year. The first three bins are positions with negative benchmark adjusted returns and the last three bins are positions with positive benchmark-adjusted returns The selling (buying) probability is computed by the number of stock sold (bought) in a particular bin divided by the total number of stocks in that bin. We exclude recently bought stocks by sorting based on the holding length from last buy on each day within a portfolio and dropping the bottom quintile of holding length since last buy. For buying probability, we only consider stocks that a portfolio manager has already held as of the prior day when computing the probability in order to avoid mechanical zero returns for newly bought stocks. The first row reports buying probabilities and the second row represents selling probabilities.

Trade	Bins of Cumulative Returns capped at 1-year					
	Worst Loser	Moderate Loser	Slight Loser	Slight winner	Moderate Winner	Best Winner
Buy	0.91	1.04	0.96	1.11	1.08	1.00
Sell	2.68	2.36	2.15	2.27	2.41	2.83

across the bins of prior returns. Figure 3, which we discuss further below, depicts this result graphically using a variety of different prior return measures with 20 bins formed on each measure, where bins are sorted from left to right according to prior returns. These results hold across prior return measures and no pronounced patterns appear as we move towards extreme bins regardless of the measure considered.

A very different picture emerges for the selling probabilities. In Table 3, which sorts past returns into 6 bins, a U-shaped pattern appears. Consistent with Hypothesis 2, stocks with more extreme relative returns are substantially more likely to be sold relative to stocks in the central bins. A stock with a more extreme return (bins 1 and 6) is 25% more likely to be sold than a stock with a less extreme return (bins 3 and 4). As depicted in Figure 3, this result is even more pronounced when sorting into 20 bins, with probabilities of selling stocks with more extreme returns nearly 50% higher than stocks with less extreme returns. Despite the fact that each specification is associated with prior return measures that are calculated over a variety of horizons, a very pronounced U-shape appears across all specifications.

Panel A of Figure 3 considers our baseline measure and an analogous one that caps relative returns at the shorter horizon of 90 calendar days instead of 1 year. In this second specification, the difference between central and extreme bins is even more pronounced than when using the baseline measure. Panels B and C look at benchmark-adjusted returns over fixed horizons of 1 year, 90 days, and returns over 1 week and 1 day, respectively. Across all horizons, there is a strong increase in selling probabilities as one moves from intermediate to more extreme bins. This is in stark contrast to buying probabilities which remain relatively flat both for intermediate and extreme returns.

Figure 3. Probability of buying and selling for various measures of past position returns

This set of figures reports buying and selling probability for stocks in the portfolio sorted into 20 bins of various past return measures. Panel A sorts on cumulative past benchmark-adjusted returns since the purchase date or one year/quarter, whichever is shortest. Panel B sorts on past benchmark-adjusted returns of a position over one year and one quarter. Panel C sorts on past raw returns of a position over one week and one day. The first ten bins are positions with negative returns and the last ten bins are positions with negative returns. The selling (buying) probability is computed as the number of stock sold (bought) in a particular bin divided by the total number of stocks in that bin. We exclude recently bought stocks by sorting based on the holding length from last buy on each day within a portfolio and dropping the bottom quintile of holding length since last buy. For buying probability, we only consider stocks that a portfolio manager has already held before when computing the probability in order to avoid mechanical zero returns for newly bought stocks. Blue bars represent selling probabilities and the red bars represent buying probabilities.

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

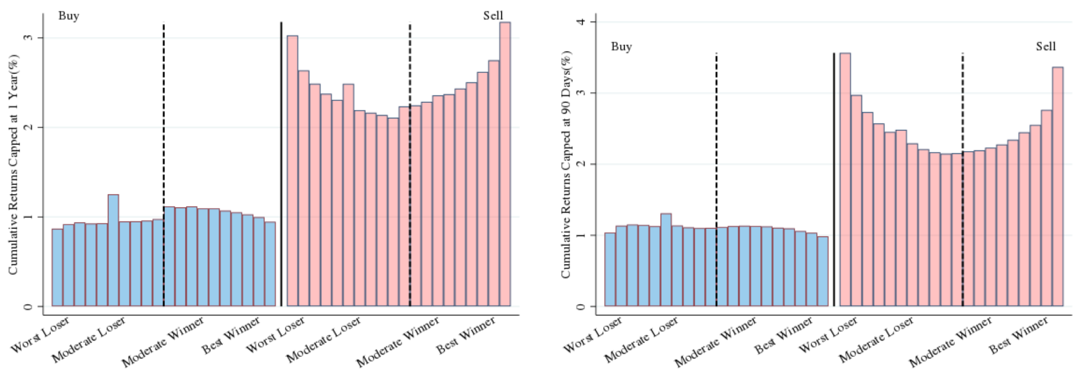

Panel B: Past benchmark-adjusted 1-year and 1-quarter returns of a position

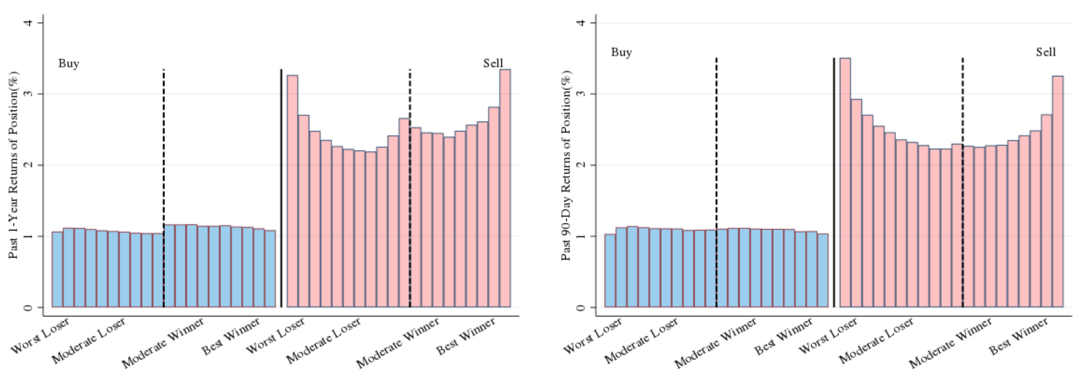

Panel C: Short-horizon 1-week and 1-day returns

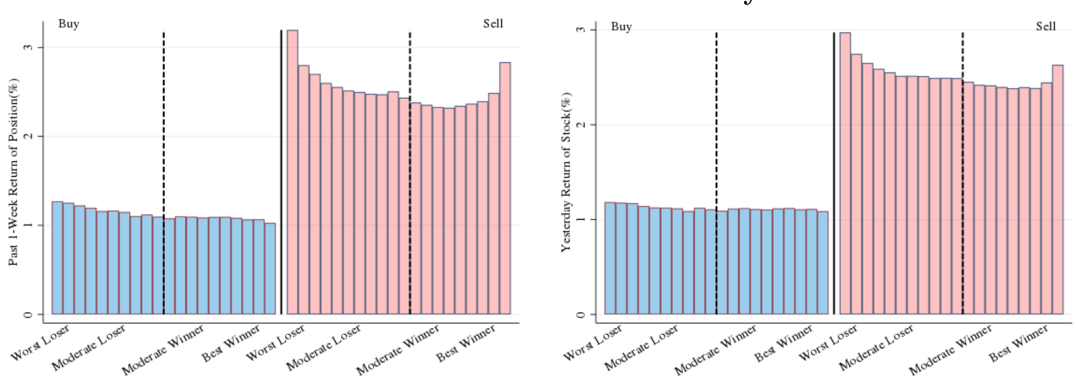

Figure 4. Probability of selling by prior returns and characteristics of holdings

This set of figures reports probabilities, in percentage points, of buying/selling by 6 bins of past benchmark-adjusted returns double sorted with bins of holding characteristics including position sizes and holding lengths. The left panel plots probabilities of buying and the right panel plots probabilities of selling. The x axis represents different position sizes in panel A and holding length in panel B. 6 bins of past position returns are plotted within each section on the horizontal axis. The selling (buying) probability is computed by the number of stock sold (bought) in a particular bin divided by the total number of stocks in that bin. For Panel A, we exclude recently bought stocks by sorting based on the holding length from last buy on each day within a portfolio and dropping the bottom quintile of holding length since last buy. For Panel B, we do not exclude the bottom quintile of holding length since last buy when computing buying probabilities.

Panel A: Position Size

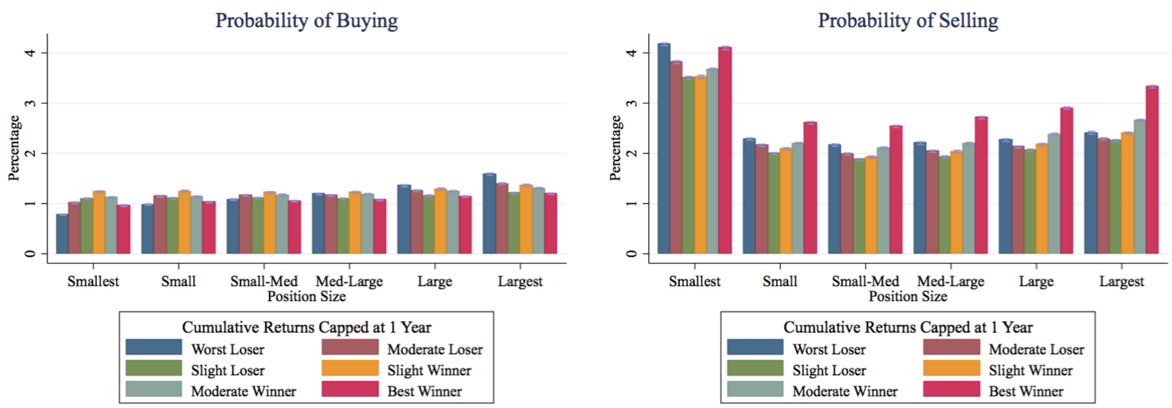

Panel B: Holding Length

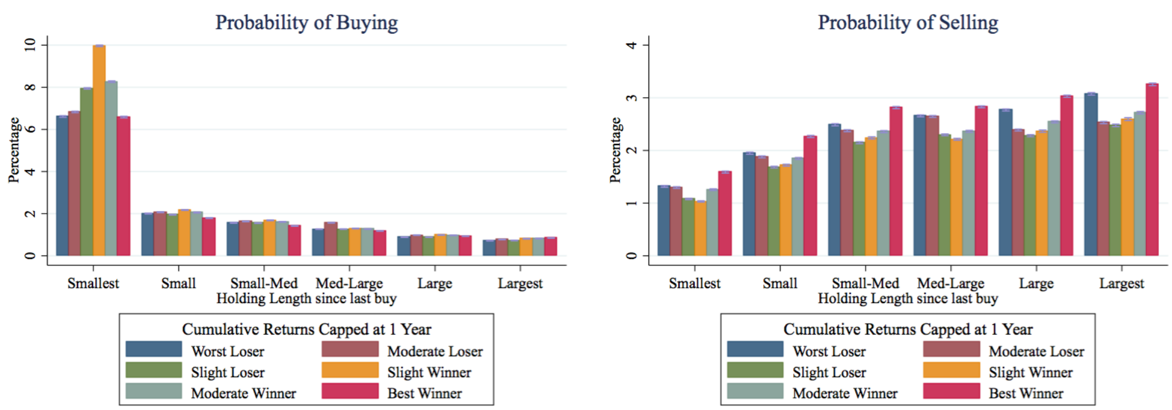

Figure 4 considers the extent to which our observed pattern can be explained by two potential omitted variables which may be correlated with our prior return measures: position size and holding length. We use the same prior return sorting procedure as Table 3 for the remaining analyses in this section, though results are similar with different numbers of bins. As a step towards addressing these concerns, we conduct simple double-sorting analyses. We assign each stock into one of 6 bins based on prior returns and the other sorting variable, respectively. Since the breakpoints used for the second characteristic are the same regardless of the bin associated with the first characteristic, there will be unequal numbers of observations in each bin. We then compute the buying (left panel) or selling (right panel) probabilities within each group.

First, even if initial positions all begin at the same size, portfolio drift will imply that stocks that experience extreme relative returns will tend to have very large or very small portfolio weights in the absence of trading. Therefore, simple rebalancing motives (e.g., to reduce portfolio exposures to idiosyncratic risk) could motivate managers to sell positions with extreme positive returns that have become too large.[25] As shown in Panel A, we observe that buying probabilities are relatively flat in position size, while selling probabilities feature a pronounced U-shape for all position sizes.

Second, as discussed above, positions which have only been held for a short period of time will tend to have less dispersion in returns and also be more likely to be bought and less likely to be sold. Panel B double sorts on 6 bins based on time elapsed since last buy (the variable we filter on) and prior returns. For this analysis only, we do not discard any stocks from the analysis based on this holding period measure. One can observe the mechanical patterns discussed in Section 2.4 when looking at the buying probabilities of assets in the bin with the shortest holding length; buying probabilities are flat in prior returns for all other holding periods. In contrast, the U-shaped pattern of selling probabilities persists across all holding lengths.

Finally, Tables 4 and 5 report 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

[25] Note, however, that similar logic would potentially imply that we would see more buying of positions that have became small due to portfolio drift, which we do not observe.

Table 4. Regressions of selling indicators on various fixed effects

This table presents position-level estimates of a linear probability model (in percentage points) for the likelihood of selling a given stock, where the key explanatory variables of interest are indicators six bins of past benchmark-adjusted returns capped at one year under different models, where the Slight Loser bin is the omitted category. We control for fund characteristics including log(yesterday's assets under management), prior-month turnover, the volatility of a fund benchmark-adjusted returns over the past year, and prior month loadings on Fama-French Cahart regressions (calculated using the Dimson (1979) procedure using 1 year of prior daily returns). We control 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 models in different columns differ predominantly by fixed effects. We consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. We exclude recently bought stocks by dropping the bottom quintile of holding length since last buy from the analysis. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at fund level. * denotes statistical significance at 5% level , ** denotes statistical significance at 1% level and *** denotes statistical significance at 0.1% level.

	(1) No FE	(2) Fund FE	(3) date FE	(4) Fund x Date	(5) Stock x Date
Worst Loser	0.493***	0.418***	0.497***	0.359***	0.203***
	(9.705)	(8.732)	(9.733)	(7.684)	(5.289)
Loser	0.108***	0.065**	0.114***	0.022	0.128***
	(4.042)	(2.973)	(4.298)	(1.076)	(5.406)
Slight Loser	0.000				
	(.)				
Slight Winner	-0.009	-0.081***	0.013	-0.083***	-0.037
	(-0.338)	(-3.992)	(0.423)	(-4.722)	(-1.551)
Winner	0.151***	0.061*	0.176***	0.003	0.169***
	(4.026)	(2.115)	(4.361)	(0.117)	(5.057)
Best Winner	0.631***	0.529***	0.648***	0.450***	0.405***
	(12.790)	(12.263)	(12.643)	(10.471)	(8.553)
Fund Control	Yes	Yes	Yes	No	Yes
FE	None	Fund	Date	Fund x Date	Stock x Date
r2	0.005***	0.018***	0.009***	0.179***	0.317***
N	54.2M	54.2M	54.2M	56.2M	45.5M

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

same interpretation as differences across rows in Table 3, where the omitted category is bin 3 (slight loser positions). Again, results are similar with different prior return measures and different numbers of bins.

We begin with Table 4, which characterizes selling probabilities. Coefficients are quite similar across columns 1-4, which include different types of fixed effects. Across all of these specifications, the difference in the predicted probability of selling a stock in bin 1 is at least 0.44% higher than the probability of selling a stock in either bin 3 or 4. The final column includes stock-date fixed effects, so the main coefficients of interest are identified off

Table 5. Regressions of buying indicators on various fixed effects

This table presents position-level estimates of a linear probability model (in percentage points) for the likelihood of buying a given stock, where the key explanatory variables of interest are indicators six bins of past benchmark-adjusted returns capped at one year under different models, where the Slight Loser bin is the omitted category. We control for fund characteristics including log(yesterday's assets under management), prior-month turnover, the volatility of a fund benchmark-adjusted returns over the past year, and prior month loadings on Fama-French Cahart regressions (calculated using the Dimson (1979) procedure using 1 year of prior daily returns). We control 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 models in different columns differ predominantly by fixed effects. We consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. We exclude recently bought stocks by dropping the bottom quintile of holding length since last buy from the analysis. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at fund level. * denotes statistical significance at 5% level , ** denotes statistical significance at 1% level and *** denotes statistical significance at 0.1% level.

	(1) No FE	(2) Fund FE	(3) date FE	(4) Fund x Date	(5) Stock x Date
Worst Loser	0.033*	0.050**	0.031*	0.034*	0.021
	(2.150)	(3.119)	(1.980)	(2.118)	(1.004)
Loser	0.012	0.031***	0.012	0.019*	0.009
	(1.563)	(3.946)	(1.470)	(2.561)	(0.859)
Slight Loser	0.000				
	(.)				
Slight Winner	-0.026	-0.004	-0.022	0.008	-0.036*
	(-1.831)	(-0.335)	(-1.742)	(1.085)	(-2.390)
Winner	-0.073***	-0.038**	-0.069***	-0.045***	-0.055**
	(-3.853)	(-2.704)	(-4.047)	(-3.998)	(-2.876)
Best Winner	-0.149***	-0.116***	-0.146***	-0.131***	-0.135***
	(-6.199)	(-5.483)	(-6.516)	(-6.938)	(-4.827)
Fund Control	Yes	Yes	Yes	No	Yes
Fixed Effect	None	Fund	Date	Fund x Date	Stock x Date
r2	0.022***	0.028***	0.028***	0.283***	0.281***
N	54.2M	54.2M	54.2M	56.2M	45.5M

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

of variation in the relative return categories across portfolio managers who hold the same stock on the same date. Even when coefficients are only identified using this narrow source of variation, we find that positions in the most extreme returns are substantially more likely to be sold.

Turning to Table 5, the relationship between buying probabilities and prior return measures is much more muted. Most of the coefficients are insignificant despite being estimated on a sample of over 50 million observations. Even the significant coefficients are substantially less economically meaningful than the coefficients associated with selling probabilities. In the saturated specification presented in column 5, none of the coefficients in the Loser categories are statistically distinguishable. Taking stock, the regression specifications, in conjunction with the nonparametric evidence in Figure 4, suggest that the considered sources of omitted variable bias are unlikely to explain our results. Together, these results are consistent with Hypothesis 2, suggesting that asymmetric allocation of attentional resources leads to a salience heuristic being used for selling decisions but not buying decisions.

5.2 Linking heuristic use and trading performance

In this section, we consider the potential implications (or lack thereof) of the heuristic strategies documented in Section 5.1. We exploit the panel nature of our dataset in order to illustrate a more direct link between the performance of selling strategies associated with heuristic use. To do so, as in Section 4, we compare the returns of the actual stocks traded with counterfactual random selling strategies. Here, we ask whether patterns in funds' actual trading strategies are associated predictable differences in performance. To operationalize this, we compute several fund and position-level characteristics and sort trades into categories based on relative levels of these characteristics, then compute the average value-added associated with each bin.

To capture heuristic intensity, we calculate the fraction of stocks sold that are located in the extreme bins (worst loser and best winner) for each fund-week.[26] We then rank fund-weeks into 4 categories according to this measure to calculate relative performance of the associated selling decisions.

Table 6 presents performance results as a function of heuristic intensity. Different panels correspond to three alternative ranking schemes. Panel A corresponds to a between-manager measure of heuristics intensity. In each week, we sort each manager into one of four categories based on its level of heuristic intensity. In Panel B, we use a within-manager measure, comparing trades that a manager makes during weeks in which heuristic intensity is relatively high or relatively low. Panel C repeats this analysis using a monthly measure of heuristics instead. The left panel plots average performance of buy trades, while the right panel plots average performance of sell trades.

To the extent that increased reliance on heuristics is suboptimal, we would expect trades

[26]For instance, the mean of this heuristics intensity measure is 0.4 on a monthly basis, which 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 C, we use a variety of fund sorts to show that, perhaps surprisingly, our measure of heuristics intensity is nearly uncorrelated with a variety of observable fund characteristics.

Table 6. Average post-trade returns relative to counterfactual by heuristics intensity

This table presents the average returns relative to random buy/sell counterfactuals for buy and sell portfolios sorted by heuristics intensity. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. The heuristics intensity is computed by measuring the fraction of sells in the lowest and highest of 6 bins of cumulative returns capped at 1-year at weekly or monthly horizons. We rank the heuristics intensity both in the cross section of funds and within-fund time series and sort funds into four bins from Lowest, Low-Med, Med-High to Highest heuristics use. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, 1 year and 2 years. We report point estimates of weighted average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate), where we weigh observations inversely to the number of trades per year of a fund. Standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Heuristics Intensity	Bin	Buy				Sell			
		Horizon				Horizon			
		28 days	90 days	1 year	2 year	28 days	90 days	1 year	2 year
Panel A: Across-fund weekly	Lowest	0.39	0.74	0.92	1.31	0.10	0.12	-0.10	0.08
		(0.009)	(0.015)	(0.035)	(0.055)	(0.01)	(0.02)	(0.05)	(0.06)
	Low-Med	0.40	0.62	0.68	0.47	-0.11	-0.30	-0.92	-0.67
		(0.007)	(0.013)	(0.030)	(0.046)	(0.02)	(0.03)	(0.08)	(0.08)
	Med-High	0.39	0.58	0.65	0.36	-0.02	-0.19	-0.64	-0.79
		(0.007)	(0.013)	(0.029)	(0.047)	(0.01)	(0.02)	(0.05)	(0.07)
	Highest	0.41	0.66	0.67	0.83	0.04	-0.30	-1.63	-2.52
		(0.008)	(0.015)	(0.033)	(0.053)	(0.02)	(0.03)	(0.09)	(0.13)
Panel B: Within-fund weekly	Lowest	0.38	0.72	0.82	1.05	0.05	0.09	0.02	0.48
		(0.008)	(0.015)	(0.033)	(0.052)	(0.01)	(0.02)	(0.04)	(0.06)
	Low-Med	0.38	0.57	0.55	-0.03	-0.09	-0.36	-1.08	-1.08
		(0.008)	(0.014)	(0.030)	(0.046)	(0.02)	(0.04)	(0.10)	(0.10)
	Med-High	0.40	0.63	0.82	1.03	0.02	-0.14	-0.79	-1.05
		(0.007)	(0.014)	(0.030)	(0.049)	(0.01)	(0.02)	(0.06)	(0.08)
	Highest	0.43	0.66	0.68	0.77	0.02	-0.29	-1.48	-2.28
		(0.008)	(0.015)	(0.032)	(0.051)	(0.01)	(0.03)	(0.07)	(0.11)
Panel C Within-fund monthly	Lowest	0.35	0.61	0.69	0.52	0.02	0.05	0.32	0.65
		(0.009)	(0.016)	(0.036)	(0.054)	(0.01)	(0.02)	(0.09)	(0.12)
	Low-Med	0.38	0.68	0.77	0.51	-0.07	-0.31	-1.27	-0.83
		(0.009)	(0.016)	(0.035)	(0.055)	(0.02)	(0.04)	(0.17)	(0.15)
	Med-High	0.37	0.59	0.77	1.66	0.06	-0.03	-0.28	-0.42
		(0.009)	(0.015)	(0.034)	(0.054)	(0.01)	(0.02)	(0.04)	(0.06)
	Highest	0.39	0.65	0.81	0.61	0.04	-0.53	-2.22	-3.45
		(0.008)	(0.016)	(0.037)	(0.058)	(0.02)	(0.06)	(0.21)	(0.27)

Figure 5. Average post-trade returns relative to counterfactual by past position returns

This figure presents weighted average returns relative to random buy/sell counterfactuals, where trades are separated into 6 bins based on past cumulative returns capped at 1 year (similar to Table 3). For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. We then compute the average of these performance measures across all portfolios and dates, weighted inversely to funds' trading activity. Each bar represents average counterfactual returns in percentage over specified horizons on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Panel A: Buy Trades

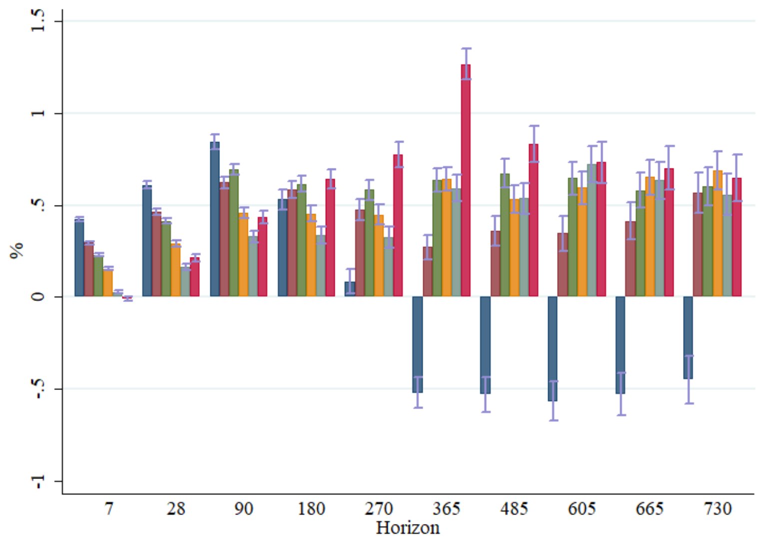

Panel B: Sell Trades

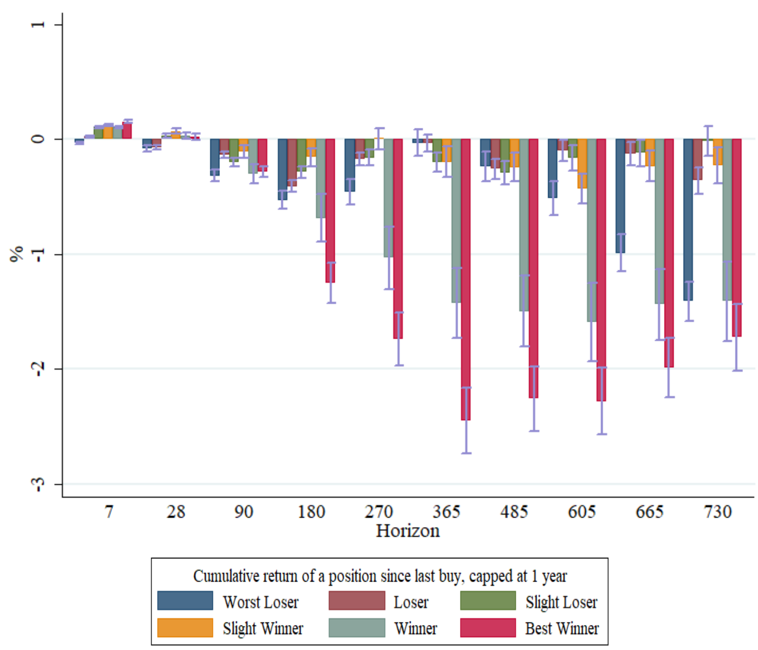

occurring during periods with high heuristic use to add the least value. Consistent with Hypothesis 3, across the majority of the specifications, the highest levels of heuristic intensity are associated with the worst performance. This effect is most stark at 1 and 2 year horizons, where the highest level of heuristic use predicts an average of 200 forgone basis points relative to a random selling strategy.

Figure 5 constructs counterfactual measures for trades associated with different levels of prior returns (the 6 bins used in Section 5.1) at various horizons. Panel A suggests that buy trades outperform the counterfactual at all horizons up to 270 days. At longer horizons, buy trades add value for all bins except the lowest one. In contrast, Panel B suggests that sell trades underperform at all except very short horizons. At longer horizons, stocks in the most extreme bins are associated with the *worst* performance relative to the counterfactual. Both Table 6 and Figure Figure 5 suggest that a systematic tendency to sell positions with extreme – as explained by use of a salience heuristic – explains the underperformance of PMs' selling decisions.

To test Hypothesis 4, Table 7 considers two proxies associated with periods that engage additional attentional resources. As in Panels B and C of Table 6, both measures are computed on a weekly basis and sort fund-weeks into 4 categories to capture within-manager variation. In Table 7, Panel A considers the number of distinct names being sold in a given week relative to the total number of names in the portfolio. Consistent with the hypothesis that attentional resources will be stretched further when a greater variety of assets is being sold, we find that sell trades underperform most during weeks when a larger number of distinct names are being sold.

Panel B considers a proxy for a manager being in "cash-raising mode." We hypothesize that when managers are most focused on raising money for a purchase, they will allocate even fewer attentional resources to selling decisions – leading to greater underperformance. To capture this, we compute the difference between the number stocks bought and the number of stocks sold, where both measures are expressed as fractions of the number of stocks in the portfolio. Intuitively, executing on a new idea likely involves reducing the size of a number of positions in order to free up capital to prepare to invest in a small number of new positions. Therefore, low numbers would likely be consistent with periods of time where the manager is engaging in "cash-raising" activities. Consistent with our hypothesis, we find that the number of sell trades relative to buy trades predicts greater underperformance of the selling decisions.

Together, the results of this section paint a stark picture: reliance on a salience heuristic

Table 7. Average post-trade returns relative to counterfactual by fund trading behavior

This table presents the average returns relative to random buy/sell counterfactuals for buy and sell portfolios sorted by by weekly trading activities (Gross Sell and Net Buy). For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. The weekly gross sell is computed by counting the number of unique positions sold within a week. Weekly net buy is computed by the unique number of positions bought per week minus the unique number of positions sold per week. For each sorting variable, we then rank these measures within portfolios across all weeks in the sample and divide them into four bins from Lowest, Low-Med, Med-High and Highest, across funds. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, 1 year and 2 years. We report point estimates of weighted average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate), where we weigh observations inversely to funds' trading activity. Standard errors are computed using Hansen-Hodrick standard errors with a lag equal to the number of horizon -1.

Fund Characteristics	Bin	Buy				Sell			
		Horizon				Horizon			
		28 days	90 days	1 year	2 year	28 days	90 days	1 year	2 year
Panel A: Gross Sell Weekly # of distinct stocks sold (sorted within fund)	Lowest	0.45	0.75	1.20	1.25	-0.23	-0.05	-0.78	0.08
		(0.008)	(0.015)	(0.035)	(0.053)	(0.009)	(0.015)	(0.052)	(0.063)
	Low-Med	0.34	0.68	0.92	1.33	0.20	0.13	-0.32	-0.77
		(0.008)	(0.015)	(0.033)	(0.054)	(0.008)	(0.015)	(0.037)	(0.058)
	Med-High	0.47	0.73	0.72	0.68	0.05	-0.09	-0.37	-0.24
		(0.008)	(0.014)	(0.031)	(0.048)	(0.008)	(0.015)	(0.036)	(0.053)
	Highest	0.26	0.49	0.54	0.64	-0.05	-0.51	-1.73	-2.02
		(0.008)	(0.014)	(0.031)	(0.050)	(0.025)	(0.063)	(0.260)	(0.277)
Panel B: Net Buy Weekly # of stocks bought - # of stocks sold (sorted within fund)	Lowest	0.42	0.66	0.45	0.05	-0.05	-0.67	-2.43	-3.26
		(0.008)	(0.014)	(0.031)	(0.050)	(0.029)	(0.077)	(0.325)	(0.343)
	Low-Med	0.46	0.96	1.14	0.67	0.13	0.05	-0.63	-0.61
		(0.008)	(0.014)	(0.033)	(0.049)	(0.008)	(0.015)	(0.038)	(0.056)
	Med-High	0.40	0.67	0.79	1.33	-0.04	-0.02	-0.32	0.04
		(0.008)	(0.014)	(0.033)	(0.051)	(0.008)	(0.014)	(0.036)	(0.052)
	Highest	0.28	0.49	0.80	1.17	0.03	0.02	0.28	0.76
		(0.008)	(0.014)	(0.033)	(0.052)	(0.008)	(0.014)	(0.032)	(0.050)

amongst expert investors leads to substantial underperformance in their selling decisions. For PMs who are most prone to this heuristic, adopting a random selling strategy would add substantial value to their portfolios.

6 Discussion and Conclusion

We use a unique data set to show that financial market experts – institutional investors managing portfolios averaging \$573 million – display costly, systematic biases. A striking finding emerges: while investors display skill in buying, their selling decisions underperform substantially – even relative to random sell strategies. A salience heuristic explains the underperformance: investors are prone to sell assets with extreme returns. This strategy is a mistake, resulting in substantial losses relative to randomly selling assets to raise the same amount of money.

The question remains of why professionals would develop expertise in one domain but not the other; selling decisions are essentially buying decisions with a minus sign. The environment in which fund managers make decisions offers several clues. As Hogarth (2001) notes, the development of expertise requires frequent and consistent feedback. While it is feasible to generate this type of feedback for both buy and sell decisions, in practice the environment in which fund managers make decisions is overwhelmingly focused on one domain over the other. The vast majority of the investors' research resources are devoted to finding the next winner to add to the portfolio. Purchased assets are tracked, providing salient and frequent feedback on the outcomes of buying decisions. This process appears successful in producing expertise – purchased assets consistently outperform the benchmark. In comparison, paltry resources are devoted to decisions of what to sell. Importantly, the relevant feedback is largely lacking: assets sold are rarely, if ever, tracked to quantify returns relative to potential alternatives such as our random sell counterfactual.

Given this imbalance in feedback, it is perhaps not surprising that fund managers display skill in buying while simultaneously rely on costly heuristics in selling. The disparity in material resources devoted to the two decision domains also suggests that fund managers likely use their limited cognitive resources to focus on buying, while neglecting to devote similar resources to improve their selling decisions. This would only exacerbate any effects generated by the differential feedback mechanisms. Our findings imply significant benefits to creating environments where learning can occur more effectively. Given the heterogeneity in selling skills – managers who do not use the extreme-selling heuristic outperform those who do – fund managers who are underperforming can adopt learning tools and simple alternative selling strategies to substantially improve performance.

Perhaps more surprising than the fact that sell trades appear to add less value than buy trades is our empirical finding that sell trades also substantially underperform a random selling strategy, which requires no skill. While formal modeling of this question is beyond the scope of this paper, we suggest one potential explanation here. All else constant, PMs' highest conviction ideas – those for which managers' ex-ante estimates of expected risk-adjusted returns are the largest – may also be more easily accessible in PMs' minds. Moreover, high conviction ideas which recently experienced large price movements may be particularly salient relative to ideas about which the manager was less confident prior to observing these signals. If so, he or she may be especially likely to select these stocks when using the salience heuristic documented above. If this were indeed the case, periods with large fractions of extreme positions sold might actually be expected to underperform more neutral strategies such as our random sell counterfactual, consistent with our empirical results. We leave the further exploration of such interactions for future research.

References

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., Y.-T. Lee, Y.-J. Liu, and T. Odean (2014): "The cross-section of speculator skill: Evidence from day trading," *Journal of Financial Markets*, 18, 1 – 24.

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.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012): "Salience Theory of Choice Under Risk," *The Quarterly Journal of Economics*, 127, 1243–1285.

——— (2013): "Salience and Consumer Choice," *Journal of Political Economy*, 121, 803–843.

——— (2017): "Memory, attention, and choice," Tech. rep., National Bureau of Economic Research.

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.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *The Economic Journal*, 101, 157–179.

Campbell, J. Y. and R. J. Shiller (1988): "The dividend-price ratio and expectations of future dividends and discount factors," *The Review of Financial Studies*, 1, 195–228.

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

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.

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. D. and T. Shumway (2005): "Do Behavioral Biases Affect Prices ?" *Journal of Finance*, 60, 1–34.

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.

Fama, E. F. (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, 25, 383–417.

Frazzini, A. (2006): "The Disposition Effect and Underreaction to News," *The Journal of Finance*, 61, 2017–2046.

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

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

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

Green, E., J. M. Rao, and D. M. Rothschild (2017): "A Sharp Test of the Portability of Expertise," *mimeo*.

Grinblatt, M. and M. Keloharju (2009): "Sensation Seeking, Overconfidence, and Trading Activity," *The Journal of Finance*, 64, 549–578.

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*.

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.

Jin, L. and A. Scherbina (2010): "Inheriting losers," *The Review of Financial Studies*, 24, 786–820.

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

Kahneman, D. and G. Klein (2009): "Conditions for intuitive expertise: a failure to disagree." *American psychologist*, 64, 515.

Kumar, A., S. Ruenzi, and M. Ungeheur (2018): "Daily Winners and Lowers," Working paper.

Liu, Y. J., C. L. Tsai, M. C. Wang, and N. Zhu (2010): "Prior Consequences and Subsequent Risk Taking: New Field Evidence from the Taiwan Futures Exchange," *Management Science*, 56, 606–620.

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.

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.

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.

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

Vuolteenaho, T. (2002): "What drives firm-level stock returns?" *The Journal of Finance*, 57, 233–264.

Table A.1. Average post-trade returns relative to counterfactual by heuristics intensity, unweighted

This table presents the average returns relative to random buy/sell counterfactuals for buy and sell portfolios sorted by heuristics intensity. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. The heuristics intensity is computed by measuring the fraction of sells in the lowest and highest of 6 bins of cumulative returns capped at 1-year at weekly or monthly horizons. We rank the heuristics intensity both in the cross section of funds and within-fund time series and sort funds into four bins from Lowest, Low-Med, Med-High to Highest heuristics use. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, 1 year and 2 years. We report point estimates of unweighted average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate). Standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Heuristics Intensity	Bin	Buy				Sell			
		Horizon				Horizon			
		28 days	90 days	1 year	2 year	28 days	90 days	1 year	2 year
Panel A: fraction, across-fund weekly	Lowest	0.39	0.69	0.91	1.31	0.14	0.20	0.20	0.37
		(0.010)	(0.018)	(0.044)	(0.070)	(0.010)	(0.019)	(0.045)	(0.075)
	Low-Med	0.41	0.58	0.48	0.11	0.01	-0.04	-0.10	0.23
		(0.009)	(0.016)	(0.040)	(0.063)	(0.009)	(0.017)	(0.039)	(0.062)
	Med-High	0.45	0.57	0.51	0.01	0.03	-0.08	-0.20	-0.46
		(0.009)	(0.016)	(0.039)	(0.062)	(0.009)	(0.017)	(0.042)	(0.069)
	Highest	0.47	0.62	0.57	0.51	-0.06	-0.20	-0.49	-0.89
		(0.010)	(0.018)	(0.043)	(0.067)	(0.013)	(0.024)	(0.059)	(0.093)
Panel B: fraction, within-fund weekly	Lowest	0.38	0.63	0.68	0.79	0.07	0.05	-0.10	0.04
		(0.009)	(0.017)	(0.042)	(0.066)	(0.011)	(0.021)	(0.049)	(0.080)
	Low-Med	0.42	0.56	0.55	0.00	0.06	0.10	-0.08	0.32
		(0.009)	(0.016)	(0.039)	(0.060)	(0.011)	(0.020)	(0.047)	(0.077)
	Med-High	0.46	0.64	0.58	0.41	0.04	-0.09	-0.05	-0.27
		(0.009)	(0.017)	(0.039)	(0.064)	(0.012)	(0.022)	(0.052)	(0.084)
	Highest	0.46	0.60	0.59	0.47	-0.05	-0.20	-0.41	-0.90
		(0.009)	(0.017)	(0.041)	(0.065)	(0.013)	(0.026)	(0.065)	(0.103)
Panel C fraction, within-fund monthly	Lowest	0.37	0.62	0.68	0.61	0.09	0.12	0.21	0.47
		(0.011)	(0.021)	(0.051)	(0.080)	(0.009)	(0.018)	(0.041)	(0.068)
	Low-Med	0.41	0.56	0.68	0.26	0.03	-0.03	-0.14	0.18
		(0.011)	(0.021)	(0.051)	(0.078)	(0.009)	(0.016)	(0.038)	(0.062)
	Med-High	0.43	0.62	0.63	1.04	0.05	-0.04	-0.19	-0.56
		(0.011)	(0.021)	(0.049)	(0.079)	(0.010)	(0.018)	(0.044)	(0.071)
	Highest	0.45	0.59	0.48	0.15	-0.07	-0.20	-0.50	-0.79
		(0.011)	(0.021)	(0.052)	(0.081)	(0.012)	(0.023)	(0.056)	(0.089)

Table A.2. Average post-trade returns relative to counterfactual by fund trading behavior, unweighted

This table presents the average returns relative to random buy/sell counterfactuals for buy and sell portfolios sorted by by weekly trading activities (Gross Sell and Net Buy). For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. The weekly gross sell is computed by counting the number of unique positions sold within a week. Weekly net buy is computed by the unique number of positions bought per week minus the unique number of positions sold per week. For each sorting variable, we then rank these measures within portfolios across all weeks in the sample and divide them into four bins from Lowest, Low-Med, Med-High and Highest, across funds. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, 1 year and 2 years. We report point estimates of unweighted average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate). Standard errors are computed using Hansen-Hodrick standard errors with a lag equal to the number of horizon -1.

Fund Characteristics	Bin	Buy				Sell			
		Horizon				Horizon			
		28 days	90 days	1 year	2 year	28 days	90 days	1 year	2 year
Panel A: Gross Sell	Lowest	0.44	0.68	0.99	0.83	0.05	0.21	0.82	1.65
Weekly		(0.009)	(0.017)	(0.042)	(0.063)	(0.009)	(0.015)	(0.038)	(0.059)
Number of distinct	Low-Med	0.38	0.60	0.80	1.23	0.09	0.10	-0.05	-0.40
stocks sold		(0.009)	(0.016)	(0.039)	(0.064)	(0.009)	(0.017)	(0.041)	(0.066)
(sorted within fund)	Med-High	0.47	0.64	0.65	0.44	0.03	-0.02	-0.23	-0.20
		(0.008)	(0.016)	(0.038)	(0.059)	(0.009)	(0.017)	(0.042)	(0.065)
	Highest	0.37	0.51	0.31	0.07	-0.01	-0.18	-0.42	-0.58
		(0.009)	(0.016)	(0.040)	(0.063)	(0.010)	(0.019)	(0.046)	(0.075)
Panel B: Net Buy	Lowest	0.43	0.63	0.35	-0.25	0.01	-0.25	-0.79	-1.26
Weekly		(0.009)	(0.016)	(0.038)	(0.060)	(0.010)	(0.019)	(0.045)	(0.074)
# of stocks bought -	Low-Med	0.47	0.81	0.75	0.40	0.07	0.03	-0.42	-0.69
# of stocks sold		(0.008)	(0.016)	(0.037)	(0.057)	(0.009)	(0.017)	(0.042)	(0.065)
(sorted within fund)	Med-High	0.43	0.58	0.54	0.68	0.06	0.13	0.29	0.53
		(0.008)	(0.016)	(0.038)	(0.060)	(0.009)	(0.016)	(0.037)	(0.057)
	Highest	0.36	0.48	0.82	1.07	-0.01	0.02	0.46	0.90
		(0.009)	(0.017)	(0.042)	(0.067)	(0.009)	(0.017)	(0.038)	(0.062)

Table A.3. Average heuristics intensity by bins of fund characteristics

This table reports the average heuristic intensity by four bins of various fund characteristics. We measure heuristics intensity by the fraction of positions sold in extreme bins of past position returns. We report this for a variety of fund characteristics, sorted in ascending order. For each bin of fund characteristics denoted by b, we measure heuristics intensity by fraction of position sold by computing :

$$HI_b^{frac} = \frac{\text{\#position sold in past return bin 1 or 6 given bin of fund characteristics b}}{\text{\# positions sold in bin of fund characteristics b}}.$$

Fund Characteristics	Lowest	Low-Medium	Medium-High	Highest
Panel A: Trading Style				
Weekly Gross sell	41.067	40.397	40.333	38.529
Monthly Turnover	39.354	38.892	39.995	39.224
Median Holding Length	38.926	39.601	39.86	38.978
Panel B: Past Fund Returns				
Fund past 2-day return	39.985	39.843	39.821	40.396
Fund past 7-day return	40.121	39.536	39.819	40.513
Fund past 30-day return	39.74	39.677	39.642	40.972
Fund past 60-day return	39.681	39.745	39.59	40.971
Fund past 90-day return	39.672	39.616	39.678	41.001
Fund past-year return	39.407	39.719	39.286	40.715
Fund past 2 year returns	39.985	39.843	39.821	40.396

Figure A.1. Average post-trade returns relative to counterfactual, unweighted

This figure presents weighted average returns relative to random buy/sell counterfactuals for (a) all buy and sell trades and (b) buy and sell trades sorted by prior performance, where trades are separated into 6 bins based on past cumulative returns capped at 1 year (similar to Table 3). For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The hold portfolio excludes recently bought stocks, defined as stocks in the lowest quintile when sorted by holding length since last buy on each day of a fund. We then compute the average of these performance measures across all portfolios and dates. Each bar represents average counterfactual returns in percentage over specified horizons on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using Hansen-Hodrick standard errors with number of lags equal to the horizon -1.

Panel A: Overall Counterfactual Returns

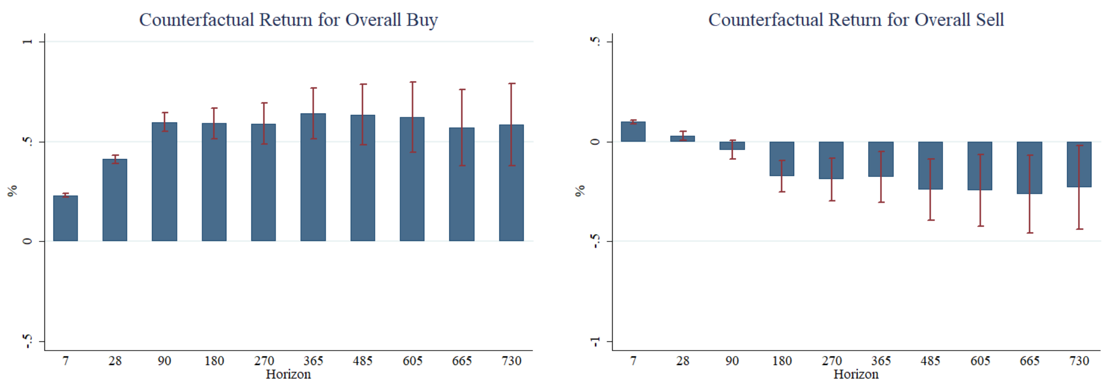

Panel B: By Bin of Past Returns