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

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

May 2019

Is expertise portable across closely related domains? We investigate this question empirically using a unique dataset of market experts—institutional investors with portfolios averaging $573 million. A striking finding emerges: while there is clear evidence of skill in buying, the investors' selling decisions underperform substantially–even relative to *random* selling strategies. This holds despite the similarity between the two decisions in frequency, substance and consequences for portfolio performance. We present evidence that an asymmetric allocation of cognitive resources such as attention can explain the discrepancy: investors are prone to a systematic, costly heuristic process when selling but not when buying.

Keywords: Heuristics, Behavioral Finance, Expert Decision-Making, Limited Attention
JEL: D91

*We thank Bronson Argyle, Ned Augenblick, Aislinn Bohren, Pedro Bordalo, John Campbell, Colin Camerer, Cary Frydman, Samuel Hartzmark, David Hirshleifer, Zwetelina Iliewa, Lawrence Jin, David Laibson, Ulrike Malmendier, Joshua Miller, Terry Odean, Alex Rees-Jones, Andrei Shleifer, Richard Thaler, Tuomo Vuolteenaho, Russ Wermers, conference and seminar participants at the NBER Behavioral Finance Meeting, Society of Judgment and Decision Making Conference, Booth School of Business, UC Berkeley, Caltech, ETH Zurich, Florida State, University of Amsterdam, and University of Frankfurt for helpful comments. We thank Kushan Tyagi for outstanding research assistance.

[†]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

Most economic models assume that skill in zero-sum environments such as financial markets is closely linked to differences in *information* across agents. As an example, when trading in frictionless markets, sophisticated investors who are more informed exploit their informational advantages to make more profitable trading decisions against less informed market participants. The nature of the optimal trade (e.g., buying versus selling) required to exploit such an advantage is irrelevant for performance. In contrast, models in behavioral economics predict that seemingly irrelevant changes in the decision-making environment, such as whether to part with or acquire the same good (Kahneman, Knetsch, and Thaler 1990), can have a large impact on choices. If contextual factors prompt heuristic-use in one decision environment and not the other, performance may differ even if the choices overlap in their fundamentals; expertise and skill that manifests in one domain will not transfer to the other.

In this paper, we study whether expertise—here, the ability to identify outperforming assets—is transferred across fundamentally similar decision contexts: buying and selling stocks in a portfolio. We examine this question using a unique data set containing the *daily* holdings and trades of sophisticated market experts—experienced institutional portfolio managers (PMs). Our data is comprised of 783 portfolios, with an average portfolio valued at approximately $573 million. More than 89 million fund-security-trading dates and 4.4 million high-stakes trades (2.0 and 2.4 million sells and buys, respectively) are observed between 2000 and 2016. We evaluate performance by constructing counterfactual portfolios, and compare PMs' actual decisions to returns of the counterfactual strategy. Our data set uniquely allows us to evaluate selling decisions relative to a conservative counterfactual that assumes no skill: *randomly* selling an alternative position that was not traded on the same date.[1]

We document a striking pattern: While the investors display clear skill in buying, their selling decisions underperform substantially. Positions added to the portfolio outperform both the benchmark and a strategy which randomly buys more shares of assets already held in the portfolio by over 100 basis points per year. In contrast, selling decisions not only fail to beat a no-skill *random* selling strategy, they consistently underperform it by substantial amounts. In our preferred specification, PMs forgo 70 basis points per year in

[1] In a setting such as ours with short-sales constraints, evaluating a selling decision relative to a counterfactual which is unrelated to existing holdings (e.g., a benchmark index) is not an appropriate comparison because PMs may be selling in order to raise capital to buy, or because their opinion about a security has changed. An asset sold may outperform 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).

raw returns.[2] Restricting the sample to only developed markets leads to a similar result of over 70 basis points in forgone returns per year. One potential alternative explanation is that stocks sold have above average exposure to systemic risk relative to the counterfactual strategy. If this is the case, measures of raw returns would overstate the performance of buys and understate the performance of sells. To address this, we replace raw counterfactual returns with those of factor-neutral strategies that take out exposure to the Carhart (1997) risk factors. Correcting for risk exposure does little to change the results: Buys continue to outperform the counterfactual by over 100 basis points while sales forgo 80 basis points a year relative to a random selling strategy.

Why would a majority of portfolio managers appear to exhibit skill in buying while at the same time underperforming substantially in selling? At face value, the fundamentals of buying and selling to optimize portfolio performance are similar: Both require incorporating information to forecast the distribution of future returns of an asset. Skill in both decisions requires the investor to look for relevant information and integrate it into the forecast. However, there is reason to suspect that selling and buying decisions involve different psychological processes (Barber and Odean 2013). Recent work from the lab is consistent with this discrepancy: Buying decisions appear to be more forward-looking and belief-driven than selling decisions in an experimental asset market (Grosshans, Langnickel, and Zeisberger 2018). And indeed, anecdotal evidence from our sample points to PMs thinking differently about the two decisions; extensive interviews suggest that they appear to focus primarily on finding the next great idea to add to their portfolio and view selling largely as a way to raise cash for purchases.[3]

We argue that the stark discrepancy in performance between buys and sells is consistent with an asymmetric allocation of limited cognitive resources such as attention towards buying and away from selling. As a first piece of evidence, we examine performance of trades that occur contemporaneously with the release of salient and portfolio-relevant information. Company earnings announcements have been used to study limited attention in asset markets (DellaVigna and Pollet 2009; Hirshleifer, Lim, and Teoh 2009) and exploited as exogenous events that draw investors' attention to assets in their portfolio (Menkveld 2013). Earnings announcements not only draw attention to specific assets or asset classes, they also provide new decision-relevant information (Ball and Brown 1968) on which skilled traders are able to capitalize (Easley, Engle, O'Hara, and Wu 2008). We exploit the variation in earnings

[2] As a benchmark, active managers of mutual funds charge between 20 to 40 basis points per year in fees.

[3] 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."

announcements as predetermined shifters of attention which may lead PMs to think more deliberately about positions that they would have otherwise not considered selling. Accordingly, we predict that contemporaneous sales are more likely to be informed and, as a result, perform better than those made on non-announcement days. In contrast, if the difference in buying and selling performance is driven by some fundamental discrepancy between the two decisions, then trades should look similar on announcement and non-announcement days.

We find that selling decisions on respective earnings announcement days outperform those on non-announcement days by about 160 basis points at a yearly horizon. Whereas sell decisions on non-announcement days substantially underperform (similar to the overall result), on average, stocks sold on announcement dates substantially *outperform* the random sell counterfactual. Consistent with PMs focusing on buys throughout, we do not detect a systematic difference in performance of buying decisions on announcement versus non-announcement days. These results suggest that investors do not lack the fundamental skill to sell well—in fact, the point estimates of buying and selling performance on announcement days are similar—it is just not transferred. Rather, the findings on overall performance are consistent with an asymmetric allocation of cognitive resources between buying and selling decisions.

As evidence for this hypothesis, we document that PMs are prone to use a heuristic process when selling but not when buying. Specifically, we propose and provide evidence for the following two-stage trading process. First, limited attention leads PMs to constrain their consideration set of what to sell to assets with extreme attributes on a salient dimension—prior returns.[4] From this set, PMs choose to unload positions to which they are least attached. The latter effect can generate systematic underperformance if the positions to which they are least attached to happen to be their newest ideas.[5]

Providing evidence for the first stage, we find that PMs in our sample have substantially greater propensities to sell positions from the set of holdings 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. Instrumental motives do not seem to explain this pattern: results are robust to controlling for position size and holding length and are unlikely to be explained by risk management motives. The pattern persists even

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

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

after the inclusion of stock-date fixed effects which absorb a number of time-varying, stock-specific unobservables. On any given day, the *same* asset is more likely to be sold from a portfolio where it exhibits relatively extreme returns than from a portfolio where its recent performance stands out less compared to other positions held. Moreover, the vast majority of portfolios in our sample are tax-exempt, meaning that tax considerations are unlikely to explain the selling of extreme performers. In contrast, we observe *no* similar tendency to focus on extremes on the buying side—unlike with selling, buying behavior correlates little with past returns. This suggests that PMs are purchasing assets based on factors that are not available to the researchers. According to the revealed preferences of their buying decisions the public signal provided by recent relative returns does not tend to change PMs' beliefs about future expected returns.[6] Rather, prior returns appear to guide the PMs' consideration sets of what assets to sell but have little effect on decisions of what to buy.

Next, we show that from the consideration set of assets with extreme returns, PMs systematically choose those to which they appear to be least attached. We define 'attachment' as the extent to which the PM has developed a position as an integral active part of their portfolio. This can be measured by examining the asset's weight relative to the benchmark and its initial position size: assets with low relative weights and initial position sizes are most likely to be new ideas that the PM is just beginning to develop. We find that the underperformance in selling is largely associated with these 'new ideas'; in fact, sales of more developed assets are not associated with any systematic underperformance. Additionally, these 'new idea' assets are also most likely to be sold and demonstrate the most pronounced relationship between extreme returns and selling probability. This systematic selling of assets to which the PM appears to be least attached is consistent with behavioral evidence on sunk costs and psychological ownership effects (Anagol, Balasubramaniam, and Ramadorai 2018; Heath 1995; Kahneman et al. 1990).

Lastly, we provide suggestive evidence that the documented heuristic process is costly. To do this, we examine the correlation between the tendency to sell assets with extreme returns—which is taken as an empirical proxy for heuristic thinking—and selling performance. Sell trades executed during periods of time in which PMs are most prone to this behavior (top quartile) forgo nearly 200 basis points annually relative to a random selling strategy, whereas sell decisions do not underperform when PMs score low on our measure of heuristic thinking. These results point to an empirical link between heuristic thinking and overall

[6] Since prior returns may reflect changes in relative valuations, it is not unreasonable to see a correlation between extreme prior returns and trading behavior. The lack of this correlation for buying decisions suggests that such motives are not the primary driver of behavior on the selling side.

underperformance in selling. Moreover, we show that selling performance is further degraded when PMs are likely to be experiencing stress (during periods when the overall portfolio is underwater) or selling in order to raise cash for buying decisions (attending to their selling choices even less).

2 Related Literature.

Our results suggest that PMs systematically fail in transferring their expertise in buying to selling decisions. Prior work has documented the fractionation of expertise (Kahneman and Klein 2009), where individuals who attain expertise in one domain fail to successfully port these skills to other related domains (Green, Rao, and Rothschild 2017). Our setting differs from these results in that investors buy and sell at approximately the same rate and are likely to have been doing so since they started in the field. It also does not seem that the PMs lack a fundamental ability to sell well—rather, they appear to allocate fewer cognitive resources to the decision compared to buying.

The paper contributes to the literature documenting biased decision-making in financial markets. The majority of research on investor behavior has focused on non-expert retail traders for whom daily holdings and trade data has been been more readily available (see Barber and Odean (2011) for review).[7] 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 behavioral biases – or lack thereof – of expert investors.[8] Haigh and List (2005) recruited professional traders from the Chicago Board of Trade and used an experimental paradigm from the lab (Gneezy and Potters 1997) to show that the investors exhibited significant myopic loss aversion – to an even greater extent than a student population (Benartzi and Thaler 1995). Papers have also used the mandated release of quarterly holdings data to study the decisions of mutual funds and other institutional investors. Funds have been found to herd on the decisions of others (Wermers 1999) and display a 'reverse' disposition effect (Chang, Solomon, and Westerfield 2016b), but the coarseness of the data makes it difficult to identify biases from instrumental motives. Due to the relative dearth

[7] A large literature has shown that non-expert market participants use heuristics and are prone to systematic biases, such as overconfidence (Barber and Odean 2001), loss-aversion (Larson, List, and Metcalfe 2016) and limited attention in their trade decisions (Barber and Odean 2008). The majority of this evidence comes from data on retail investors (Barber and Odean 2011) or day traders (Barber, Lee, Liu, and Odean 2014), who generally hold modest portfolios.

[8] 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 (2016a); Giglio and Shue (2014); Hartzmark and Shue (2017); Hong, Torous, and Valkanov (2007).

of field data, the behavioral finance literature has mostly assumed unbiased institutional investors exploiting the behavioral biases of retail investors (Malmendier 2018). The findings documented in this paper suggest that such an assumption may not be a valid one.

Our results also contribute to the literature demonstrating heuristics and biases amongst experts in domains such as sports (Green and Daniels 2017; Massey and Thaler 2013; Pope and Schweitzer 2011; Romer 2006), judges (Chen, Moskowitz, and Shue 2016), professional forecasters (Coibion and Gorodnichenko 2015), 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.

The selling pattern we document is most related to the rank effect described in Hartzmark (2014).[9] There, retail investors appear to exhibit a similar pattern in selling *and* buying behavior—unloading and purchasing assets with more extreme returns. However, it is not clear from the data whether these trading strategies are particularly maladaptive: This set of investors have been found to underperform the market in general and display a host of heuristics and biases such as the disposition effect (Odean 1998), overconfidence (Odean 1999), and narrow bracketing (Frydman, Hartzmark, and Solomon 2017).[10] Our results also relate to the analysis of Di Mascio, Lines, and Naik (2017), who used the Inalytics Ltd. dataset of institutional investors to test theoretical models of optimal strategic trading with private information. Most of their analyses aggregate information across managers to examine the speed at which managers trade and, in turn, the rate at which private information is incorporated into prices. The authors argue that the results support models of optimal trading strategies: stocks with above average buying and selling volume tend to outperform the benchmark. Given the different focus of their paper (aggregate metrics rather than individual decision-making), they do not explore individual-level determinants of trading behavior nor use existing holdings to compare performance of strategies to feasible alternatives (e.g. evaluating quality of actual selling strategies relative to counterfactual strategies).

The paper proceeds as follows. Section 3 describes the data. Section 4 presents results on performance of buying and selling decisions, while Sections 5 and 6 present results on the use of heuristics in trading strategies and how those strategies affect performance, respectively. Section 7 discusses implications for learning and the porting of expertise.

[9] Ben-David and Hirshleifer (2012) also demonstrate a V-shaped trading pattern with respect to prior returns, but at the level of the individual security rather than with respect to overall portfolio returns.

[10] Though Hartzmark (2014) focuses on the behavior of retail investors, he also presents 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.

3 Data

This section reviews 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.

3.1 Data sources and sample selection

Our primary source of data for this analysis is compiled by Inalytics Ltd. These data include information on the portfolio holdings and trading activities of institutional investors. Inalytics acquires this information as part of one of its major lines of business, which is to offer portfolio monitoring services for institutional investors that analyze the investment decisions of portfolio managers.[11] The majority of portfolios in our sample are sourced from asset owners—institutional investors such as pension funds who provide capital to PMs to allocate on their behalf. In these cases, we see holdings and trades related to the specific assets owned by the client. The remainder of the portfolios are submitted by PMs themselves who seek to benchmark their own performance; in these cases, data will frequently correspond with holdings and trades aggregated over multiple clients. These data are associated with a single strategy, so we do not observe assets managed by the same PMs using alternative strategies. For purposes of this study, Inalytics assembled a dataset of long-only equity portfolios spanning from January 2000 through March of 2016. These portfolios are almost always tax-exempt, hold limited cash, and are prohibited from using leverage or shorting positions. 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.

For each portfolio, we have a complete history of holdings and trades at the daily level throughout the sample period. 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. Our dataset includes an unbalanced panel of both active and inactive portfolios, with the

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

vast majority of the data collected essentially in real-time, suggesting that incubation and survivorship biases are unlikely to be a substantial concern for our analysis.[12]

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.

We apply two primary filters to select the set of portfolios to include in our analysis. First, we apply the filter when daily trading data are unavailable for a subset of portfolios or appear to be incomplete.[13] Second, we exclude funds that do not have a sufficient fraction (at least 80 percent) of portfolio holdings which could be reliably matched with CRSP or Datastream. In demonstrating the robustness of our results, we perform the analyses using data from developed markets only; these markets arguably have better price discovery and higher match rates with CRSP/Datastream. 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.[14]

Differences from other datasets This sample offers some unique opportunities to study expert decision-making relative to other datasets in the literature. First, in contrast to the Large Discount Brokerage dataset of Barber and Odean (2000), which features portfolio holdings and trades of individual retail investors and has been used in numerous studies[15], our data include complete portfolio and trade-level detail for a population of professional investors managing large pools of assets. Illustrative of this distinction, Barber and Odean

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

[13] 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.

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

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

Table 1. Summary statistics

This table reports 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	47114	0.414	0.227	0.289	0.397	0.526
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
Earnings announcement day indicator	49.3M	0.007	0.08	0	0	0
Active share	89.8M	0.86	1.27	0.11	0.55	1.28

(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—the average portfolio in our sample is almost four *thousand* times larger. Second, unlike other datasets which characterize institutional portfolios such as mutual fund portfolio holdings reports and 13-F filings (e.g. Chang et al. (2016b)), we are able to observe portfolio holdings and changes to those holdings on a *daily* level. This facilitates the testing of hypotheses on individual decision-making that is infeasible with quarterly data. Additionally, in the other most widely used database with institutional trading information—the Abel Noser/ANcerno database (for an overview, see Hu, Jo, Wang, and Xie 2018)—researchers often do not observe all trades made by a given institutional investor and tend to lack timely information on portfolio holdings.

3.2 Fund and position-level characteristics

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

While holding fairly diversified portfolios (average number of stocks is about 78 with a standard deviation of 68), funds in our sample remain active, with positions that deviate substantially from their benchmarks. On an asset level, deviation from the benchmark is captured by an asset-specific measured called *active share*, which corresponds to the asset's weight in the portfolio relative to its weight in the benchmark. The average tracking error—the standard deviation of the difference between the daily portfolio return and the benchmark—is about 0.35 percent per day, or about 5.7 percent on an annualized basis. On average, a manager will initiate a sell trade for about 10 percent and a buy trade for about 15 percent of the stocks in his/her portfolio each month. We also characterize fund portfolios in terms of factor exposures by computing rolling Carhart 4-factor regressions (using the prior 1 year of daily data with the Fama-French international factors), adjusted for asynchronous trading.[16] The average market beta is about 1, and average exposures to the SMB, HML, and Momentum factors are fairly close to zero.

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

Table 2. Summary of characteristics

This table describes how we construct several characteristics for use 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 a description of the sorting variable construction.

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

Panel A also reports the average benchmark-adjusted return that uses each portfolio-specific return series. The average fund in our sample beats its respective benchmark by about 0.22 percent per month, or 2.6 percent per year. This, in conjunction with the fact that funds' average betas are close to 1 and have little average exposure to the three other priced risk factors, suggests that these managers are highly skilled, earning returns above and beyond exposure to known risk factors.[17] We view the positive selection of managers in our sample as an advantage when studying expertise and heuristic use: The population we examine is clearly skilled, and thus identifying biased behavior is likely a lower bound when generalizing the results.

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

[17] Prior work has demonstrated that a subset of insitutional investors do persistently outperform the market and generate alpha (Kosowski, Timmermann, Wermers, and White 2006).

percent and 2.2 percent 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 four 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 percent with a standard deviation of 1.6 percent. Inalytics also provides us with a measure of active share.

Finally, we compute a number of measures of backward or forward-looking returns at the position 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.[18] 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 percent 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.

4 Overall Trading Performance

Having described the basic properties of our dataset and variable construction procedures, we now begin to analyze performance of PMs' decisions. We begin by discussing our methodology for computing counterfactual portfolio returns and, accordingly, value-added measures. We then present the first of our empirical results, which calculates the average value-added (or

[18] 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 percent of returns in either tail by date. These steps are discussed at greater length in the Appendix.

lost) associated with managers' active buying and selling decisions.[19]

4.1 Constructing counterfactuals

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

Given that PMs in our sample tend to hold limited cash positions and are not generally permitted to use leverage, the primary mechanism for raising money to purchase new assets is selling existing ones. Since the portfolios already include stocks that are carefully selected to outperform their respective benchmarks, the choice of which asset to sell is far from innocuous. Precisely if managers' use of information 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; there, we would expect them neither to gain nor 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 buying decisions, and one for evaluating selling decisions. For evaluating buys, suppose that we learned that a manager was planning to invest $1 to purchase a stock tomorrow and to hold it for a fixed period of time. We then suggest that instead of executing the proposed idea, the PM invests that money in a randomly selected stock from his other holdings. For evaluating sells, suppose that we learned that the PM 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 instead of executing this trade, the PM 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.[20] Note that the expected payoff

[19] We will return to this analysis in more depth in Section 6, which will link other position and fund-characteristics with predictable differences in trading performance.

[20] In contrast, selling the benchmark to finance a purchase, which implicitly corresponds to the counterfactual in measuring benchmark-adjusted returns of stocks sold, is likely infeasible for a long-only manager who, similar those in our sample, holds a portfolio with a small (relative to the number of assets in the benchmark) number of high active share positions and thus deviates substantially from the benchmark. Purchasing the

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 decision 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. Following this logic, we compute $R_{buy} - R_{hold}$ and $R_{hold} - R_{sell}$ over horizons ranging from 1 week to 485 days (the average holding period) for all buy and sell trades, respectively, to characterize the value-added associated from each.

Note that these measures can be interpreted as changes in benchmark-adjusted returns associated with different trading strategies. According to our discussions with clients and managers this is the primary manner in which these managers are evaluated. That said, they also have an alternative interpretation to the extent that buy and sell trades are not motivated by a desire to change a portfolio's systematic risk exposures. In that case, we would expect loadings on priced factors of the assets being traded and the hold portfolio to be similar and these measures would also correspond to differences in risk-adjusted returns (i.e., "alpha"). However, a natural concern is that stocks traded tend to have above average exposures to systematic risk, meaning that our estimates could be driven by risk compensation rather than skill. If this were the case, we would tend to overstate positive performance of buy trades and understate performance of sells.

To address this concern, we also construct counterfactuals to form "factor-neutral" portfolios. Specifically, we estimate stock-level exposures to the Fama-French/Carhart 4 factors using data from prior to the trade, then use these estimates to adjust our long short portfolios for ex-ante differences in these exposures.[21] For each stock-date, we subtract off the inner product of factor loadings and factor realizations, so

$$R_{i,t}^{FN} \equiv R_{i,t} - \Lambda'_{i,q(t)-1} F_t,$$

where $R_{i,t}$ is stock i's excess return on date t and F_t is a (4×1) vector of factor realizations. $\Lambda_{i,q(t)-1}$ is a (4×1) vector of factor loadings which are estimated 1 year of daily data using data up to the end of the previous calendar quarter. $R_{i,t}^{FN}$ thus captures return of a self-financing portfolio which, if factor loadings are estimated correctly and are stable, has zero exposure to the priced risk factors on each date. Thus, if the asset pricing model holds, all $R_{i,t}$ should earn zero excess return in expectation, and, accordingly, randomly

benchmark is feasible on the other hand.

[21] The four factors are the market excess return, the Fama-French (2012) international size and value factors, as well as the Carhart momentum factor. As above, we compute loadings using data for the global factors from Ken French's website.

sold portfolios should have the same factor-neutral returns period-by-period as actual stocks sold. Next, we compute value-added as before, by compounding factor neutral returns and compare cumulative factor-neutral returns of stocks traded with the average of cumulative factor-neutral returns of stocks held.[22]

Lastly, to address potential issues about measurement errors (e.g., stale prices) and/or liquidity, we re-run our main counterfactual analyses excluding stocks which are traded in developing and emerging markets.[23]

We aggregate across trades in the following manner. If multiple stocks are bought or sold on a given day, we average these measures for buy and sell trades separately. Since not all funds trade every day and are not necessarily present throughout our sample period, this averaging procedure yields a portfolio-day unbalanced panel. Because some funds trade much more frequently than others—see the dispersion in monthly turnover in Table 1—we weight observations inversely to a measure of trading frequency.[24]

We conduct statistical inference by computing standard errors using a panel heteroskedasticity and autocorrelation estimator which is similar to Hansen and Hodrick (1980). Our estimator 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. The autocorrelation correction is likely required since our use of long horizon returns potentially introduces an overlapping structure in the error term of each fund's value-added time series. Further details for these calculations are described in Appendix A.2.[25]

[22] We have used additional information to construct a potentially "more intelligent" counterfactual. As we show in Figure 6 below, very few PMs elect to sell stocks that were very recently purchased. Thus, we have also considered a counterfactual which exclude stocks which are in the bottom quintile of the distribution of holding length since last purchase. Since results are similar between the two approaches, we elected to use the simpler of the two. Results from alternative counterfactual specifications are available upon request.

[23] Similar to Fama and French (2015), we re-run our analyses restricting attention to developed countries in four regions: (i) North America (NA), including the United States and Canada; (ii) Japan; (iii) Asia Pacific, including Australia, New Zealand, Hong Kong, and Singapore; and (iv) Europe, including Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

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

[25] As an alternative, we also computed standard errors using a simple Monte Carlo ("placebo") approach which is quite similar in spirit to the manner in which we construct the counterfactual portfolios themselves. Specifically, rather than use the actual positions traded, we randomly allocate (without replacement) the same number stocks from the portfolio to be bought/sold as we observe in the data. We then form counterfactuals and aggregate across funds and time as we do in the data. Results, which are available upon request, are quite similar to those with our HAC standard error estimator.

4.2 Overall performance relative to counterfactuals

Figure 1, Panel A shows 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 the random buy counterfactual, $R_{buy} - R_{hold}$. The average stock bought outperforms the counterfactual by more than 120 basis points over a one year horizon.

Figure 1, Panel B presents 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. Magnitudes are quite substantial: The value lost from an average sell trade is on the order of 70 basis points at a 1 year horizon relative to a simple counterfactual which randomly sells other stocks held on the same day.

Table 3, Panel 1 reports estimated return measures from the analysis in Figure 1 for our baseline specification as well as alternatives that adjust for risk and restrict the sample to developed markets. Our first alternative is the factor-neutral performance measure described immediately above. Our second two alternatives recompute baseline and factor-neutral performance measures for the subsample of developed countries only. In all cases, results are fairly similar between the baseline model and the three alternatives. Consistent with trading activity not being concentrated among stocks with above-average systematic risk exposure, we find fairly similar estimates of value-added for factor-neutral portfolios compared to the baseline estimates. Results are also quite similar in the developed only sample.[26]

The results thus far have examined performance of all buying and selling decisions together. However, both buys and sells differ in the extent to which they add or subtract from the portfolio. Some buys add a little bit to an existing position while others introduce a substantial amount of shares or start a whole new position in the portfolio; similarly, some sells cut a bit from existing positions while others unload substantial shares or cut the asset altogether. We refer to buy decisions that add 50 percent or more of an asset to the portfolio as 'marriages' (which includes opening a new position) and sell decisions that cut 50 percent or more from an asset as 'divorces' (100 percent corresponds to cutting a position completely). Table 4 below presents performance of marriages and divorces relative to the

[26] We also conducted additional tests in which we benchmark performance against a weighted random-sell counterfactual, where the weights in the hold portfolio (i.e., the probability of being selected) are proportional to prior-day portfolio weights and obtained similar results.

Figure 1. Post-trade returns relative to counterfactual

This figure presents differences between average returns of stocks bought/sold and those of random buy/sell counterfactual strategies for buy and sell trades. Bars present average return spreads in percentage points over horizons specified on the x axis. The bracket at the top of each bar is the 95% confidence interval of the point estimate at each horizon. Confidence intervals in brackets are computed using heteroskedasticity robust standard errors, calculated as described in Appendix A.2.

Panel A: Buy Trades

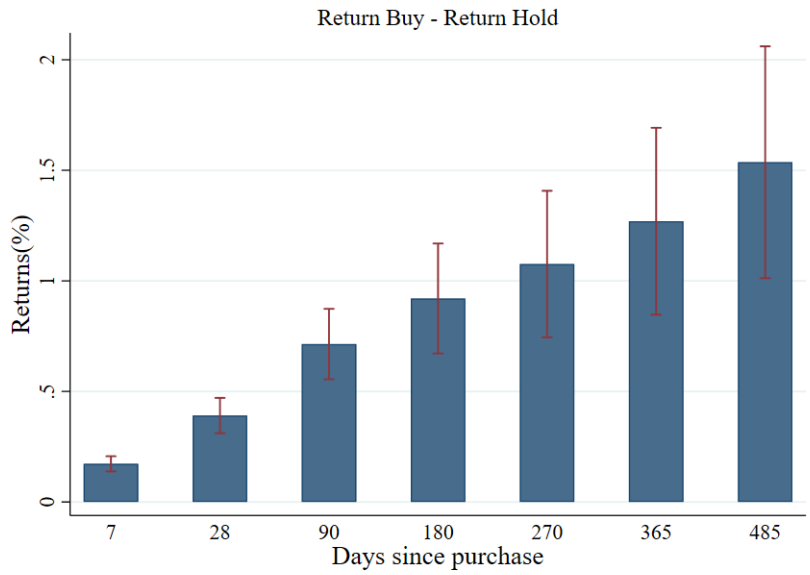

Panel B: Sell Trades

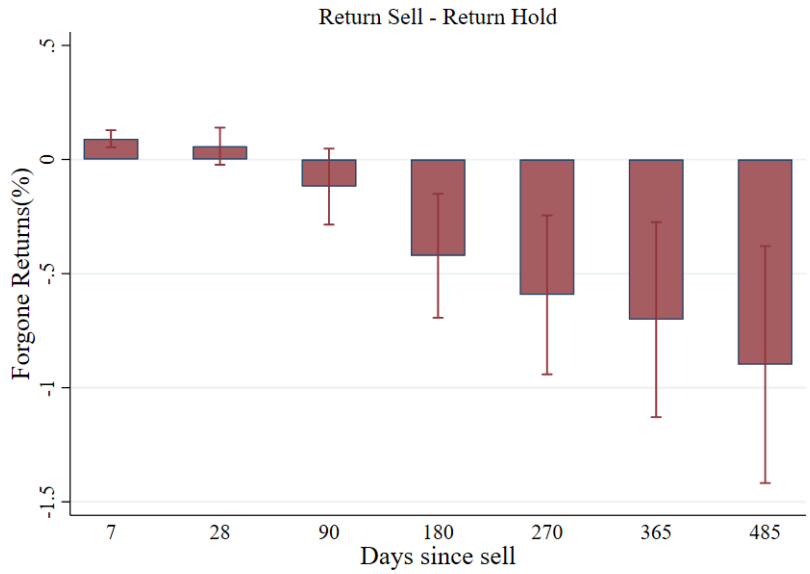

Table 3. Post-trade returns relative to counterfactual, overall and robustness checks

This table presents the average value added measures (post-trade returns relative to a random sell counterfactual) for buy and sell trades under two measures of returns 1) raw returns and 2) factor-neutral returns, for the whole sample and the subsample of stocks from developed markets (see text for further details). We first present the overall average returns relative to random buy/sell counterfactuals (I), and then report the difference between averages of these measures for trades of stocks on their earnings announcement days versus all other days (II). Last, in part III, we report the difference between average performance of buys and sells for trades on announcement dates vs non-announcement dates using the baseline measure. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in Appendix A.2, are reported in parentheses.

Performance	Specification	**Panel A: Buy**					**Panel B: Sell**				
Measure	Horizon	28	90	180	270	365	28	90	180	270	365
	Baseline	0.39	0.71	0.92	1.08	1.27	0.06	-0.12	-0.42	-0.59	-0.70
		(0.04)	(0.08)	(0.13)	(0.17)	(0.22)	(0.04)	(0.08)	(0.14)	(0.18)	(0.22)
I.	Baseline	0.34	0.63	0.73	0.70	0.81	-0.01	-0.11	-0.44	-0.58	-0.71
Overall	(Developed)	(0.04)	(0.08)	(0.12)	(0.16)	(0.19)	(0.05)	(0.09)	(0.14)	(0.18)	(0.23)
	Factor-neutral	0.34	0.58	0.83	0.98	1.16	0.03	-0.21	-0.50	-0.69	-0.80
		(0.04)	(0.08)	(0.13)	(0.17)	(0.22)	(0.04)	(0.08)	(0.14)	(0.18)	(0.22)
	Factor-neutral	0.28	0.47	0.57	0.53	0.57	0.00	-0.19	-0.49	-0.58	-0.69
	(Developed)	(0.04)	(0.09)	(0.12)	(0.16)	(0.19)	(0.05)	(0.09)	(0.14)	(0.18)	(0.22)
II.	Baseline	-0.14	0.08	0.16	0.62	-0.25	0.28	0.40	1.16	2.13	1.60
Average Trading		(0.18)	(0.31)	(0.47)	(0.58)	(0.72)	(0.18)	(0.33)	(0.51)	(0.65)	(0.77)
Performance	Baseline	-0.21	0.55	0.24	0.47	0.07	0.64	0.46	1.48	2.24	1.52
Differential:	(Developed)	(0.20)	(0.34)	(0.51)	(0.60)	(0.72)	(0.24)	(0.41)	(0.55)	(0.63)	(0.79)
Earnings	Factor-neutral	-0.04	0.14	0.24	0.70	-0.32	0.37	0.29	1.08	2.05	1.54
vs Other Days		(0.17)	(0.31)	(0.46)	(0.62)	(0.77)	(0.17)	(0.31)	(0.49)	(0.63)	(0.78)
	Factor-neutral	-0.16	0.32	0.03	0.19	-0.19	0.68	0.36	1.55	2.48	2.01
	(Developed)	(0.19)	(0.32)	(0.52)	(0.62)	(0.77)	(0.25)	(0.41)	(0.55)	(0.68)	(0.85)
III.											
Average performance difference of buys and sells: Non-announcement trades							0.33	0.84	1.37	1.72	2.04
Average performance difference of buys and sells: Announcement trades							-0.09	0.51	0.37	0.21	0.19

same counterfactual used in Table 3. Results are largely the same as for overall trade performance: Marriages outperform the counterfactual while divorces underperform it. We see larger effects in both directions for marriages and divorces compared to other trades.[27]

4.3 Performance on announcement days

We conjecture that the discrepancy in performance depicted in Panel A versus Panel B of Tables 3-4 is driven by the asymmetric allocation of limited cognitive resources such as attention towards buying and away from selling. To provide evidence for this claim, we examine performance on days when decision-relevant information is salient and readily available—earnings

[27] The greater performance of marriages compared to other buy trades is consistent with PMs trading on information when opening or increasing a position by a significant amount.

Table 4. Post-trade returns relative to counterfactual, marriage/divorce vs other trades

This table presents the average value added measures (post-trade returns relative to a random buy/sell counterfactual) for large (marriage and divorce) trades under two measures of returns 1) raw returns and 2) factor-neutral returns, for the whole sample and the subsample of stocks from developed markets (see text for further details). Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.2, are reported in parentheses.

Performance Measures	Bins\Horizon	Panel A: Regular Buy / Marriage					Panel B: Regular Sell / Divorce				
		28	90	180	270	365	28	90	180	270	365
Baseline	Normal Trade	0.28	0.52	0.67	0.81	0.94	0.20	0.06	-0.19	-0.32	-0.46
		(0.04)	(0.08)	(0.13)	(0.18)	(0.23)	(0.04)	(0.09)	(0.14)	(0.18)	(0.23)
	Marriage /	0.60	1.03	1.29	1.47	1.74	-0.43	-0.67	-1.05	-1.27	-1.30
	Divorce	(0.06)	(0.11)	(0.16)	(0.20)	(0.26)	(0.06)	(0.12)	(0.19)	(0.26)	(0.30)
Baseline	Normal Trade	0.23	0.45	0.45	0.38	0.44	0.14	0.11	-0.14	-0.21	-0.32
(Developed)		(0.04)	(0.08)	(0.12)	(0.15)	(0.18)	(0.05)	(0.09)	(0.15)	(0.19)	(0.23)
	Marriage /	0.58	0.98	1.20	1.30	1.53	-0.49	-0.73	-1.15	-1.40	-1.57
	Divorce	(0.07)	(0.12)	(0.17)	(0.22)	(0.27)	(0.07)	(0.13)	(0.21)	(0.28)	(0.33)
Factor-neutral	Normal Trade	0.24	0.46	0.68	0.82	0.91	0.17	-0.05	-0.26	-0.46	-0.63
		(0.04)	(0.09)	(0.14)	(0.19)	(0.24)	(0.04)	(0.09)	(0.14)	(0.18)	(0.22)
	Marriage /	0.53	0.80	1.04	1.19	1.46	-0.45	-0.70	-1.10	-1.20	-1.14
	Divorce	(0.06)	(0.11)	(0.16)	(0.21)	(0.25)	(0.06)	(0.12)	(0.19)	(0.26)	(0.30)
Factor-neutral	Normal Trade	0.18	0.35	0.40	0.34	0.32	0.15	0.00	-0.20	-0.27	-0.43
(Developed)		(0.04)	(0.09)	(0.12)	(0.16)	(0.19)	(0.05)	(0.09)	(0.15)	(0.19)	(0.22)
	Marriage /	0.49	0.69	0.84	0.88	1.04	-0.49	-0.74	-1.17	-1.24	-1.22
	Divorce	(0.06)	(0.12)	(0.18)	(0.22)	(0.26)	(0.07)	(0.13)	(0.21)	(0.28)	(0.33)

announcement days. We gather earnings announcement dates from the I/B/E/S database and recompute our counterfactual return strategies for stocks which are bought/sold on those days, relative to all other trading days.[28] Managers have a strong incentive to pay close attention to stocks in their portfolios on these dates for several reasons. As discussed in Section 1, the information in financial statements, associated press releases, and conference calls (which even offer opportunities for managers to directly address questions to the company) provide a wealth of new pieces of hard and soft information that are decision-relevant and can potentially improve trading performance (Easley et al. 2008). This information is both (relatively) easily available and salient, since earnings announcement dates are known in advance, and results are heavily covered by the financial press. In turn, we conjecture that earnings announcements prompt PMs to broaden their consideration sets of what to sell, potentially mitigating the limitations imposed by heuristic processes related to attention.

Panel 2 of Table 3 depicts the difference in average performance of trades on announcement versus non-announcement days. Panel A reports the difference between average value-added

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

of buy trades executed on earnings announcement days compared with average value-added from all other buy trades.[29] There is little systematic difference in performance, and whatever differences exist are not statistically significant. This is consistent with attentional resources already being devoted towards purchase decisions; information released on earnings announcement days is carefully incorporated into purchase decisions just like other forms of information are incorporated on non-announcement days. Panel B demonstrates the stark contrast in the performance of selling decisions on announcement versus non-announcement days. Selling decisions on announcement days add substantially more value than those sold on non-announcement days. Our point estimate of this difference for the baseline specification, while somewhat noisily estimated due to the smaller number of earnings-related trades, is significant at the 5% level about 160 basis points over a one year horizon, and all estimates are statistically significant at horizons of 180 and 270 days. The remaining rows Table 3, Panel 2 shows that these results hold when adjusting for risk and restricting the sample to developed markets.

Finally, Panel 3 of Table 3 compares the difference between average performance of buys and sells on earnings announcement dates versus all other dates. In the baseline specification, our point estimate of selling performance on earnings dates is +86 bp at a 1 year horizon, which is only 19 bp lower than the corresponding estimate for buying performance on earnings dates. In contrast, buys outperform sells by more than 200 bp at a 1 year horizon on all other dates, and differences in relative magnitudes are similar at other horizons as well.

These findings suggest that when contemporaneous predetermined events shift PMs' attention towards existing positions—and to potentially consider a wider set of assets and information that they would otherwise ignore when selling—performance of selling decisions improves substantially. This also provides evidence suggesting that the overall poor selling performance is not necessarily due to a fundamental lack of skill in selling.

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

5 Predicting Buying and Selling Decisions

Why might the performance of buying and selling decisions diverge? The preceding section provides some initial evidence that asymmetric allocation of limited cognitive resources such as attention can potentially explain the difference in performance. Here we provide additional, more direct evidence for this mechanism by documenting a heuristic process in selling—but not in buying—that has been previously linked to limited attention.

Work in psychology and economics suggests that limited attention prevents a person making 'fast' choices from considering the entire portfolio of assets. Rather, she may consider a narrower subset of potential choices usually comprised of assets that rank particularly high or low on some salient dimensions (Gourville and Soman 2007; Lleras, Masatlioglu, Nakajima, and Ozbay 2017). We examine whether assets with extreme prior returns—one of the most salient attributes available to traders—are more likely to be traded. Consistent with an asymmetric allocation of attention, we find that extreme returns help to explain decision-making in the domain where we predict the PMs are making 'fast' choices (selling) but not where they are making deliberative 'slow' ones (buying).

5.1 Measuring relation between prior returns on buying and selling decisions

For each portfolio-date, we identify a set of stocks (a subset of holdings in the prior day's portfolio) potentially under consideration to be bought or sold, rank existing holdings according to past benchmark-adjusted returns, and then ask whether managers are more likely to trade the holdings based on these ranks.

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. We also emphasize within-manager rankings, rather than absolute levels of these measures, since the definition of "extreme returns" may depend on the types of assets in a given PM's investment opportunity set. 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 $N_{bin} = 20$. Some specifications collapse

across bins to fit more conveniently in tabular format—the results are always robust to the number of bins considered.

While this approach is straightforward for selling decisions since the consideration set of what to sell is composed of the current holdings, constructing the consideration set for buying decisions is a bit more challenging. Our first approach considers purchases of assets that already exist in the portfolio; this approach captures the majority of buys and includes most 'marriages' (adding up to 99 percent to existing holdings). Our second approach includes all purchase decisions—including the opening of brand new positions—and calculates relative prior returns by broadening the consideration set to assets that are likely being considered for purchase. Specifically, because our dataset contains not only current and past holdings for each PM but future holdings as well, we can include assets that the PM is likely considering by looking at what he ended up buying within 12 months of the current date. We include those assets in the portfolio when computing the prior return bins to examine whether new positions are more or less likely to be bought depending on prior returns relative to the larger consideration set.

For the first approach, our preferred measure of prior returns is computed as follows. For positions which were opened more than 1 quarter (90 days) prior to the date of interest, we use the benchmark-adjusted return of the stock from 90 calendar days prior through the trading day before the date of interest. For positions with shorter holding periods, we change the starting point for computing the benchmark adjusted return to the opening date.[30] When we use the wider consideration set approach, we use the prior 1 quarter benchmark-adjusted return. We use this as our preferred measure because performance is often reported to clients at a quarterly frequency, and, from a more pragmatic perspective, this construction is less sensitive to the censoring issues for holding length discussed above. However, as we show in Section 5.2, results are robust to alternative definitions of past returns.[31]

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

[30] For buying specifications which use the wider "consideration set" approach, we use prior returns over a fixed period of calendar time (90 days), though results are robust to a wide variety of horizons.

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

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 impact.[32] 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 for predicted selling probabilities do not meaningfully change without such a restriction.[33]

5.2 Buying and selling based on past returns

We present results as fractions of positions that are bought or sold within each of the prior return bins. These fractions, which can be interpreted as probabilities, are computed by first calculating the proportion of stocks sold within each bin at the fund-date level, then averaging across all fund-dates in the sample. Figure 2 depicts the results for selling and buying decisions of assets that are already held graphically using a variety of different prior return measures, with 20 bins formed on each measure. Bins are sorted from left to right according to prior returns. We begin with the buying probabilities. The probability of purchasing a stock already held is quite flat across the bins of prior returns. These results hold across all prior return measures considered and no pronounced patterns appear as we move towards more extreme bins in all cases.

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

Panel A of Figure 2 considers our baseline measure and an analogous one that caps relative returns at the longer horizon of 1 year instead of 90 days. In this second specification, the

[32] 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 regression analyses below always control for the holding period since the position was opened and the holding period since last buy, as well as squared terms of each.

[33] Further, in Table 5, we report probabilities of buying using a prior return measure which does not depend on the time of initial purchase and do not impose the restriction on holding length. In that specification, our main results on the relationship between average buying probabilities and prior returns maintain.

difference between central and extreme bins is fairly similar, though slightly smaller, than estimates with the baseline measure. Panels B and C look at benchmark-adjusted returns over fixed horizons of 1 quarter, 1 year, and returns over 1 week, 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.

Table 5, Panel A replicates the buying decisions presented in Figure 2 but includes new buys using our second approach which expands the consideration set. Specifically, we report differences in probabilities relative to a baseline category (bin 10, stocks which barely underperformed the benchmark) of buying across categories of prior returns. For ease of comparison, the top row reports our estimate from Panel B of Figure 2, which uses the 1 quarter prior benchmark-adjusted return measure as the sorting variable. We average probabilities across several intermediate bins for brevity, and report the baseline probability associated with the omitted category in the final column. Then, the second row uses the same sorting variable but also includes stocks in the broader consideration set (as defined above) and eliminates our restriction which excluded stocks in the bottom bin of holding length since last buy. We see that the probabilities of purchasing an asset remain quite flat with respect to prior returns.

Panel B of Table 5 depicts the propensity to engage in 'divorces'—selling more than 50 percent of an asset—as a function of prior returns. We see a similar U-shape emerge as when we consider the all sales together, where changes in probabilities of engaging in 'divorces' are particularly likely to increase in response to extreme losses. Again, this pattern is not matched for the probability of engaging in 'marriages,' which remain quite flat across bins of prior returns. Together, these results demonstrate that we can predict selling decisions based on observables from the PM's current holdings with some confidence; in contrast, these observables—nor any others that we have considered—do not predict buying decisions.

5.2.1 Alternative explanations

We now consider several instrumental reasons that could potentially explain our results. As discussed in Section 3, the vast majority of portfolios in our sample are tax-exempt, so the U-shaped selling pattern cannot rationalized with tax concerns. Our finding that positions with extreme returns in terms of both very long (1 year) and very short (1 week) horizons makes agency-based explanations—where PMs are reluctant to report realized losses to their clients—unlikely. Agency-based explanations also seem unlikely to explain the large jumps

Figure 2. Probability of buying and selling based on past returns

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

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

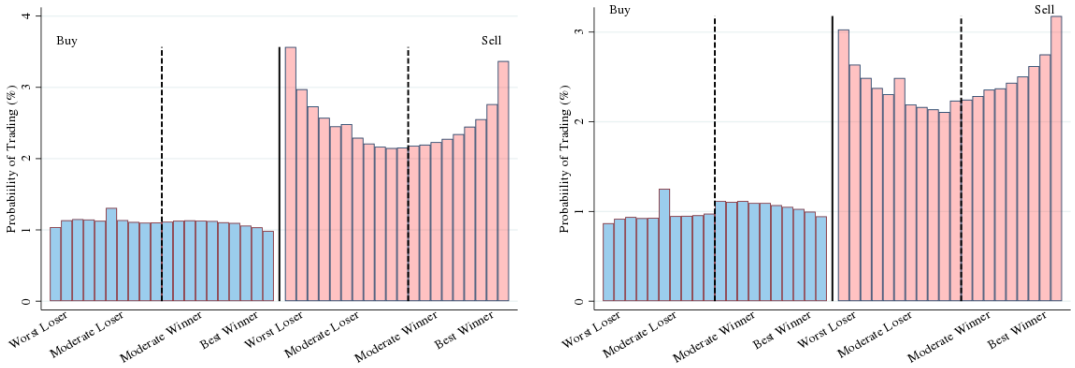

Panel B: Past benchmark-adjusted 1-quarter and 1-year returns of a position

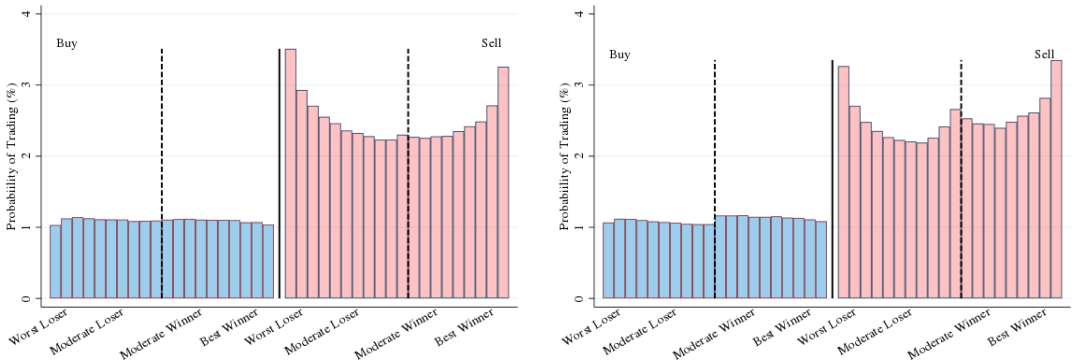

Panel C: Short-horizon 1-week and 1-day returns

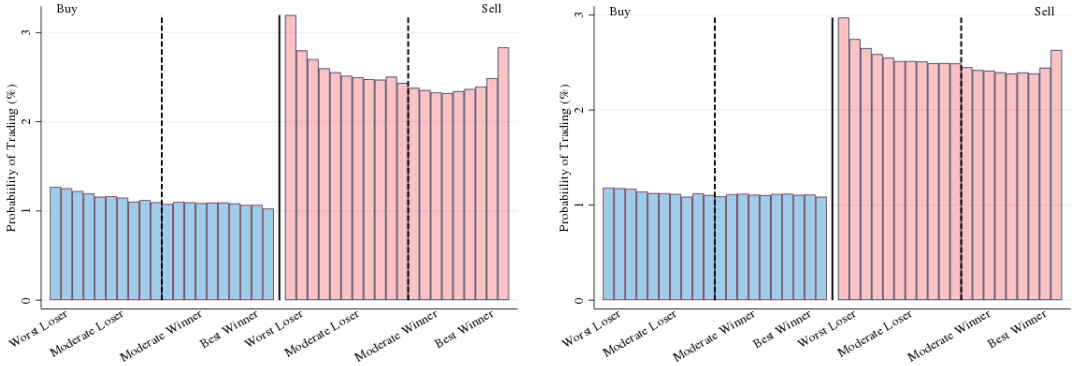

Table 5. Probability of buying and selling: larger consideration set and Marriage/Divorces

This table presents the probability of buying for security in the investors holdings and consideration sets, as well as the probability of marriage and divorces as a function of prior returns. Panel A presents reports differences in probabilities, in percentage points, of buying selling by bins of past benchmark-adjusted returns for stocks in the larger consideration set (current holdings plus stocks added to the portfolio in the next year) by 20 bins of positions past benchmark-adjusted returns. The baseline probability of trading a stock in the omitted category, bin 10, is reported in the rightmost column. The first row applies the same filters as Figure 2, while the second row considers a wider set of stocks following the approach described in the main text. Panel B presents the relative probability of marriage/divorce by 20 bins of past positions benchmark-adjusted returns capped at 90 days.

Panel A: Buying probability by bins of prior 1 quarter benchmark-adjusted returns

	Differences relative to bin 10 for prior return bins								Baseline
Buying prob for	1	2	3-5	6-9	11-15	16-18	19	20	prob (bin 10)
Current holdings	0.03%	0.05%	0.03%	0.00%	0.02%	0.00%	-0.02%	-0.06%	1.10%
Consideration set	0.09%	0.07%	0.05%	0.00%	0.00%	0.07%	0.10%	0.06%	1.62%

Panel B: Trading probability by bins of prior benchmark-adjusted returns, capped 1 quarter

	Differences relative to bin 10 for prior return bins								Baseline
Trading Probability	1	2	3-5	6-9	11-15	16-18	19	20	prob (bin 10)
Marriage	0.02%	0.02%	0.01%	0.02%	0.01%	0.01%	0.01%	0.01%	0.16%
Divorce	0.38%	0.27%	0.21%	0.07%	0.01%	0.01%	0.03%	0.12%	0.37%

in probabilities observed between the 19th and 20th (1st and 2nd) bins relative to the 18th and 19th (2nd and 3rd) bins. These jumps are consistent with limited attention constraining the consideration set of what to sell, as the top and bottom 5 percent of returns are much more likely to be displayed and made salient to PMs (see Ungeheuer (2017) for direct evidence). This observation also mitigates concerns about risk management motives, since the relative risk of assets in extreme bins is likely to be fairly comparable to less extreme adjacent bins.[34]

Table 6 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: holding length and position size. As a step towards addressing these concerns, we conduct simple double-sorting analyses. As above, we assign each stock into one of 20 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 report the buying (top panel) or selling (bottom panel) probabilities within each group relative to the middle, least extreme

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

Table 6. Probability of trading by prior returns and position characteristics

This set of tables reports differences in probabilities, in percentage points, of buying/selling by bins of past benchmark-adjusted returns double sorted with bins of position characteristics holding length and position sizes relative to bin 10 of past benchmark-adjusted returns within each category. The top section of each panel reports relative probabilities of buying and the bottom section reports relative probabilities of selling. Baseline probabilities for the omitted category are reported below. Columns represents different holding lengths in Panel A and position sizes in Panel B. Different bins of past position returns are reported in rows, together with the baseline probability of the omitted category.

Panel A: Holding Length

Trade	Past Return\Holding Length	Shortest	Short	Short-Med	Med-Long	Long	Longest
Buy	1	-2.34	-0.20	-0.15	-0.07	-0.03	0.03
	2	-1.99	-0.05	-0.03	0.02	0.05	0.07
	3-5	-2.08	-0.02	-0.01	0.05	0.04	0.05
	6-9	-1.20	0.05	0.05	0.13	0.07	0.04
	11-15	0.28	-0.03	-0.03	0.01	0.04	0.05
	16-18	-1.59	-0.19	-0.12	-0.03	0.02	0.08
	19	-2.27	-0.33	-0.21	-0.06	0.02	0.09
	20	-2.93	-0.48	-0.29	-0.10	-0.02	0.09
	Baseline: 10	8.78	2.03	1.56	1.22	0.86	0.68
Sell	1	0.45	0.78	1.17	1.14	1.52	1.70
	2	0.31	0.40	0.59	0.64	0.93	1.05
	3-5	0.24	0.24	0.35	0.34	0.42	0.58
	6-9	0.16	0.14	0.19	0.18	0.12	0.16
	11-15	0.07	0.08	0.08	0.02	0.05	0.01
	16-18	0.32	0.27	0.32	0.32	0.35	0.34
	19	0.49	0.47	0.61	0.56	0.63	0.58
	20	0.75	0.94	1.12	1.15	1.25	1.21
	Baseline: 10	1.02	1.63	2.07	2.15	2.17	2.33

Panel B: Position Size

Trade	Past Return\Position Size	Smallest	Small	Small-Med	Med-Large	Large	Largest
Buy	1	-0.21	-0.07	0.11	0.19	0.31	0.49
	2	-0.12	-0.01	0.11	0.15	0.25	0.39
	3-5	-0.11	-0.03	0.05	0.12	0.18	0.27
	6-9	0.05	0.02	0.10	0.06	0.07	0.13
	11-15	0.08	0.01	0.05	0.03	-0.02	-0.02
	16-18	0.02	-0.05	-0.01	-0.02	-0.05	-0.08
	19	-0.07	-0.09	-0.04	-0.04	-0.07	-0.14
	20	-0.10	-0.13	-0.09	-0.09	-0.13	-0.20
	Baseline: 10	0.98	1.07	1.03	1.07	1.17	1.28
Sell	1	1.33	0.90	0.93	1.02	1.02	0.90
	2	0.92	0.54	0.55	0.65	0.64	0.53
	3-5	0.55	0.30	0.31	0.35	0.37	0.30
	6-9	0.17	0.13	0.12	0.12	0.12	0.09
	11-15	-0.06	0.04	0.03	0.10	0.13	0.14
	16-18	0.07	0.23	0.30	0.42	0.51	0.59
	19	0.22	0.41	0.51	0.71	0.83	0.96
	20	0.77	0.97	1.13	1.28	1.40	1.66
	Baseline: 10	3.38	1.93	1.81	1.85	1.95	2.18

bin (bin 10). As in Figure 5, we average across several intermediate categories and separately report the probability of trading for the omitted category.

First, 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 A of Table 6 double sorts on six 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 the holding period measure. One can observe the mechanical patterns discussed in Section 5.1 when looking at the buying probabilities of assets in the bin with the shortest holding length; buying probabilities are quite flat in prior returns for longer holding periods. In contrast, assets in extreme bins are much more likely to be sold across all holding lengths.

Second, 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 larger or smaller portfolio weights in the absence of trading. Therefore, 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.[35] As shown in Panel A of Table 6, we observe that selling probabilities feature a pronounced U-shape for all position sizes, a pattern that holds robustly within all position size bins. For larger positions, we see some evidence of PMs adding to their biggest positions following losses. However, even within these categories, buying probabilities increase gradually with losses and decrease gradually with gains, whereas corresponding sell probabilities increase much more dramatically for the more extreme return categories. Additionally, the magnitudes for buys are generally much smaller compared to the respective selling probabilities. We discuss position size more in Section 6.1 below.

Finally, Tables 7 and 8 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 same interpretation as the bins used in the preceding analyses, where the omitted

[35] Note, however, that similar logic would potentially imply that we would see less selling of positions that have became small due to portfolio drift, which we do not observe. Also, from the univariate evidence above, we do not see large increases in buying for positions that declined in value, as would be predicted by this channel. In regressions below, we will always include controls for position size.

Table 7. Probability of selling based on prior returns

This table presents position-level estimates of a linear probability model (in percentage points) for the likelihood of selling a given stock. The key explanatory variables of interest are indicators corresponding to 20 bins of past benchmark-adjusted returns capped at one year, where the tenth bin is the omitted category. We control for fund characteristics including log(yesterdays assets under management), prior-month turnover, the volatility of a funds 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. Columns consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at the fund level. $*$ denotes statistical significance at 5% level , $**$ denotes statistical significance at 1% level and $***$ denotes statistical significance at 0.1% level. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at the fund level. * denotes statistical significance at 5% level , ** denotes statistical significance at 1% level and *** denotes statistical significance at 0.1% level.

	(1)	(2)	(3)	(4)	(5)
Bin 1	1.389***	1.329***	1.379***	1.237***	0.796***
	(14.828)	(13.961)	(14.734)	(13.126)	(6.812)
Bin 2	0.806***	0.748***	0.799***	0.666***	0.543***
	(12.580)	(11.685)	(12.498)	(10.561)	(6.749)
Bin 3 to 5	0.432***	0.403***	0.428***	0.331***	0.308***
	(10.998)	(10.729)	(10.916)	(9.094)	(6.216)
Bin 6 to 9	0.109***	0.102***	0.107***	0.067***	0.079***
	(6.750)	(6.642)	(6.644)	(4.797)	(3.546)
Bin 11 to 15	0.028	0.016	0.037*	0.005	0.105***
	(1.560)	(0.929)	(2.064)	(0.260)	(4.572)
Bin 16 to 18	0.312***	0.295***	0.318***	0.234***	0.559***
	(8.004)	(7.798)	(8.100)	(6.147)	(9.630)
Bin 19	0.578***	0.552***	0.582***	0.485***	0.834***
	(10.849)	(10.601)	(10.884)	(9.222)	(10.429)
Bin 20	1.186***	1.139***	1.186***	1.071***	1.132***
	(15.869)	(15.398)	(15.810)	(14.410)	(10.261)
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

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

category remains bin 10 (slight loser positions). Results are similar with different prior return measures and different numbers of bins.

We begin with Table 7, 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 20 is at least 1 percent higher than the probability of selling a stock in bins 6 through 15, and always

Table 8. Probability of buying based on prior returns

This table presents position-level estimates of a linear probability model (in percentage points) for the likelihood of buying a given stock. The key explanatory variables of interest are indicators corresponding to 20 bins of past benchmark-adjusted returns capped at 90 days, where the tenth bin is the omitted category. We control for fund characteristics including log(yesterdays assets under management), prior-month turnover, the volatility of a funds 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. Columns consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at the fund level. $*$ denotes statistical significance at 5% level , $**$ denotes statistical significance at 1% level and $***$ denotes statistical significance at 0.1% level. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at the fund level. * denotes statistical significance at 5% level , ** denotes statistical significance at 1% level and *** denotes statistical significance at 0.1% level.

	(1)	(2)	(3)	(4)	(5)
Bin 1	-0.041	-0.038	-0.047*	-0.074**	-0.144*
	(-1.789)	(-1.652)	(-2.110)	(-3.238)	(-2.284)
Bin 2	0.040*	0.045*	0.033	0.008	-0.058
	(2.071)	(2.241)	(1.735)	(0.421)	(-1.278)
Bin 3 to 5	0.046**	0.056***	0.041**	0.025	0.001
	(3.275)	(3.719)	(2.924)	(1.655)	(0.034)
Bin 6 to 9	0.025**	0.032***	0.022*	0.011	0.012
	(2.856)	(3.584)	(2.538)	(1.252)	(0.839)
Bin 11 to 15	-0.029**	-0.011	-0.031**	-0.032***	-0.014
	(-2.701)	(-1.247)	(-2.853)	(-3.844)	(-0.823)
Bin 16 to 18	-0.088***	-0.068***	-0.092***	-0.109***	-0.152***
	(-4.699)	(-3.876)	(-4.809)	(-6.450)	(-3.911)
Bin 19	-0.129***	-0.112***	-0.133***	-0.155***	-0.220***
	(-5.810)	(-5.177)	(-5.905)	(-7.434)	(-4.280)
Bin 20	-0.169***	-0.163***	-0.175***	-0.212***	-0.286***
	(-6.700)	(-6.388)	(-6.798)	(-8.492)	(-4.376)
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

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

considerably higher than bin 19. Likewise, we observe similar strong nonlinearities for stocks in bins 1 through 2 relative to more central bins. The final column includes stock-date fixed effects, so the main coefficients of interest are identified off 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 8, the relationship between buying probabilities and prior return measures is much more muted. In the loss domain, most of the coefficients are insignificant despite being estimated on a sample of over 50 million observations. Even the significant coefficients are substantially smaller in magnitude than the coefficients associated with selling probabilities. In the saturated specification presented in column 5, only the coefficient on bin 1 is statistically distinguishable from zero. Turning to extreme gains, we observe a greater number of significant coefficients, but the differences between central and extreme bins (e.g., bins 16 through 18 and bin 20 or bins 19 and 20) are much more muted relative to selling decisions. Taking stock, the regression specifications, in conjunction with the nonparametric evidence in Table 6, suggest that the considered sources of omitted variable bias are unlikely to explain our results.[36] Together these results are consistent with non-instrumental motives driving selling but not buying decisions.

6 Explaining Underperformance

In this section, we propose a potential mechanism linking the use of heuristics to systematic underperformance of selling strategies relative to a feasible counterfactual. We then provide evidence for the mechanism by exploiting the panel nature of our database to ask whether patterns in funds' actual trading strategies are associated 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 average value-added associated with sells in each bin. Differences in expected returns across these categories consistently point to a link between reliance on heuristics and selling underperformance. Importantly, we observe little evidence of heuristic use in buying strategies.

6.1 Potential mechanisms linking heuristics and underperformance

Many models of decision-making in psychology (Hauser and Wernerfelt 1990) and economics (Lleras et al. 2017) split choices between multiple alternatives—in our case, choosing what asset(s) to sell—into two stages: generating a consideration set and then selecting an option from that set. Prior work has shown that cognitive constraints can lead to the use of heuristics in both stages of the process (Hauser 2014).[37] Barber and Odean (2008) posit this type of two-stage process for trading decisions, where limited attention constraints the consideration

[36] Increases in selling probabilities for very extreme bins are even larger when considering other measures of prior return rankings. We elected not to report these estimates since magnitudes are quite similar to Figure 2. These pronounced increases in probabilities of selling extremes are not matched for buys.

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

set to assets with salient attributes and biases in preferences lead to potentially suboptimal choices from that consideration set. We outline how this process would operate in our setting and provide evidence for biases in both stages of the decision. We then demonstrate how the heuristic process can explain the results presented in the preceding sections, including underperformance relative to random sell counterfactuals.

In the first stage, rather than considering the entire portfolio, limited attention places bounds on the consideration set (Hirshleifer and Teoh 2003). Research in psychology and economics has found that these consideration sets are often determined by the ranking and filtering assets on salient attributes.[38] Information on prior returns is ubiquitous, and according to theories of salience (Bordalo, Gennaioli, and Shleifer 2013), the high variation around average returns should make this attribute particularly top-of-mind for a fast-thinking PM.[39] In turn, extreme deviations in relative returns in either the positive or negative direction naturally emerge as drivers for the construction of consideration sets. A focus on assets with extreme returns can also be rationalized using common investment maxims: It is easy to argue that assets with extreme gains have already realized their anticipated upside potential and will mean revert, while extreme losses suggest the investment thesis has changed, or that prices will become even more volatile.[40] Our results from Section 5.2 provide strong evidence that extreme returns at least partially govern the consideration set of potential sales: assets that are in the top or bottom 5 percent based on prior returns are nearly 50 percent more likely to be sold relative to those that just over- or underperformed, a pattern not observed for buy decisions.[41]

Next, the PM must choose which asset(s) from the consideration set to sell. According to the *attribute substitution* framework of Kahneman and Frederick (2002), people making 'fast,' heuristic decisions replace the more difficult question of "which asset in this set is least likely to outperform in the future" with an easier question to answer, such as "how

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

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

[40] While these these may sound like contrasting reasons, they are consistent with prominent investing advice. While buying advice is mostly about the fundamentals, the popular investment publication Barron's instructs, "If you double your money, sell and take profit" as part of the '5 Rules of Options Trading.' At the same time, the similarly prominent platform Investopedia advices traders to cut their losses: "Taking corrective action before losses worsen is always a good strategy...Selling these 'dogs' has another advantage: You will not be reminded of your past mistakes."

[41] The results from Hartzmark (2014) offer additional support for this mechanism—retail investors, who tend to be less sophisticated overall, appear to make both their selling *and* buying decisions based on extreme returns.

attached am I to this stock?" Positions in an actively managed fund can be ordered based on how much they are overweighted relative to the benchmark. This measure, known as *active share*, captures how much the PM stands to gain if the stock beats the benchmark.[42] Assets with high active share typically correspond to positions that the manager has spent a good deal of effort building up over time, likely becoming familiar and attached to the firm in the process. This costly process likely generates greater attachment to the position for non-instrumental reasons, such as sunk costs or psychological ownership (Anagol et al. 2018; Heath 1995; Kahneman et al. 1990).

Positions with low active share can manifest for three main reasons: 1) a position which had a high active share but has done very poorly, 2) the PM has added a position to the portfolio but has not yet built it up over time, or 3) the PM chooses to hold a position close to the benchmark weight (or underweight, though negative active shares are fairly uncommon) of a stock which is large in the benchmark.[43] The PM may still be attached to a stock in the first category as he had exerted time and effort in building it up in the past. Rather than due to a particular view about future returns, assets in the third category may be in place to minimize exposure of a fund's relative returns to the idiosyncratic returns of large stocks in the benchmark. In contrast, assets in the second category are most likely to be the PM's 'new ideas.' The manager has gathered enough information on the asset to add it to the portfolio, but has not yet put in the effort to build up the position over time and become attached to it. In turn, heuristic thinking would generate fewer reasons to keep a low active share asset from the latter category, while at the same time, experiencing an extreme return produces a salient reason to sell. However, selling newer, less entrenched ideas may be exactly the wrong thing to do since the information which drove the initial buying decision is likely to still be fresh, leading the asset to outperform a random holding from the portfolio. As we now proceed to demonstrate, this process appears to explain the underperformance of the PMs' selling decisions.

Panel A of Table 9 documents the PMs' propensity to sell based on active share, both overall and within each of the 20 bins of prior returns. To construct this table, assets within each portfolio are sorted into four bins based on their active share. We then construct a measure capturing the propensity to sell an asset based on its prior returns; specifically, the difference in the probability of selling a stock in a given bin of prior returns relative to the

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

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

Table 9. Probability of selling by active share and past returns

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

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

middle one (bin 10, Slight Loser). The last column reports the difference between the lowest active share bin and the average across the other three active share bins in the same row of prior returns. We report baseline probabilities for the omitted bin in the last row.

Results are consistent with PMs being more prone to sell assets they are least attached to (as measured by low active share) from the consideration set of extreme returns.[44] First, examining the baseline probabilities, we note that low active share positions are substantially more likely to be sold regardless of the level of prior returns. Second, we find that stocks in the lowest active share bin are much more likely to be sold when they exhibit prior returns below the benchmark, especially extreme ones, relative to high active share assets. The probability of selling a stock with the lowest active share and lowest prior return bin is 5.6 percent, or 140 percent larger than the baseline probability of selling (which is 2.3 percent). Assets in these bins are also 155 percent more likely to be sold than those which experienced similar levels of underperformance (bin 1) but have the highest active share. Thus, low active share positions are particularly likely to be discarded when they are in the consideration set of extreme underperformance. Selling probabilities in the lowest active share bin are relatively less effected by moderate gains, and responses to the most extreme gains in bin 20 are similar regardless of active share.

[44] Consistent with results in the prior section, we find that buying probabilities do not exhibit a significant relationship with prior returns. We do not report these results for brevity.

Table 10. Post-trade sell returns relative to counterfactual by active share

This table presents the average returns relative to random sell counterfactuals for sell portfolios sorted by active share. We compute average returns of stock held minus returns of stocks sold. We rank the active share measures within funds at a daily level and sort them into four bins from Lowest, Low-Med, Med-High to Highest sizes. For the Lowest Active Share bin, we further split into two halves by a position's weight: Smaller positions (below the 50th percentile) and Larger positions (above the 50th percentile), after sorting by active share. Columns represent sell performance measures at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. We report point estimates of average counterfactual returns for each portfolio at different horizons. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.2, are reported in parentheses.

Performance Measure	Active Share Bins	Horizon				
		28	90	180	270	365
Baseline	Lowest,	-0.41	-0.78	-1.2	-1.47	-1.38
	Smaller Positions	(0.07)	(0.14)	(0.22)	(0.30)	(0.36)
	Lowest,	0.21	0.28	0.09	0.2	0.48
	Larger Positions	(0.06)	(0.12)	(0.17)	(0.23)	(0.28)
	Low-Med	0.27	-0.08	-0.47	-0.97	-1.56
		(0.07)	(0.12)	(0.21)	(0.26)	(0.34)
	Med-High	0.10	-0.03	-0.37	-0.39	-0.84
		(0.06)	(0.12)	(0.21)	(0.26)	(0.32)
	Highest	0.24	0.27	0.18	0.13	-0.14
		(0.06)	(0.12)	(0.19)	(0.25)	(0.30)
Factor-neutral	Lowest,	-0.46	-0.97	-1.46	-1.81	-1.63
	Smaller Positions	(0.07)	(0.14)	(0.22)	(0.31)	(0.36)
	Lowest,	0.10	0.01	-0.05	0.14	0.54
	Larger Positions	(0.06)	(0.11)	(0.17)	(0.23)	(0.28)
	Low-Med	0.20	-0.15	-0.47	-0.86	-1.19
		(0.07)	(0.13)	(0.21)	(0.27)	(0.33)
	Med-High	0.16	0.13	-0.15	-0.24	-0.84
		(0.07)	(0.12)	(0.22)	(0.27)	(0.32)
	Highest	0.32	0.29	0.23	0.13	-0.35
		(0.06)	(0.12)	(0.19)	(0.25)	(0.30)

We then examine whether sales of low active share assets tend to underperform relative to a random selling counterfactual. Panel A of Table 10 depicts the performance of sales relative to a random counterfactual by bins of stocks' active share for our baseline value-added measure. As above, we sort positions into four bins based on a prior day estimate of active share, then separately form counterfactual performance measures for these different subsamples of trades. We see a stark contrast in performance: assets in the lowest three active share bins underperform substantially more than sales of assets in the top active share bin, where the latter actually tend to outperform the counterfactual at shorter horizons, and perform in line with the counterfactual at longer ones. The latter result is consistent with

PMs holding on to high active share assets when thinking fast; so when a sale is observed, it is more likely to be an informed one.

As discussed above, very low active share positions may be held to hedge a fund's exposure to idiosyncratic returns of large stocks in the benchmark. These positions would appear in the data as having a very low active share and a high portfolio weight.[45] In turn, we separate stocks in the lowest active share bin into two categories based on absolute position size, and recompute our counterfactual performance measures for these subsamples. Whereas large positions with low active share tend to have positive counterfactual performance measures (consistent with these being somewhat passive positions and good candidates for sales), we find that sales of assets from the second 'new ideas' category – low active share and low position size – to substantially underperform. Panel B repeats this analysis for factor-neutral portfolios and obtains similar results.[46]

As additional evidence that the underperformance in sales is driven by PMs letting go of positions they are least attached to, we present performance results based on initial position size. Stocks with low initial position sizes were sufficiently interesting to put into the portfolio but the PM has not yet put in the time to justify making a large bet. To investigate whether assets with these characteristics are indeed associated with underperformance in selling, we construct a measure of PMs' initial choices of portfolio weights when a position was added to the portfolio and rank stocks into four bins according to the measure.[47] Table 11 reports average performance of sales in each bin for the baseline and factor-neutral performance measures, respectively. Consistent with a mechanism of potentially promising ideas being discarded "too early," we find a strong empirical link between underperformance of selling strategies and a measure of PMs' initial position sizes. Specifically, sales of stocks in the smallest bin of initial position size have particularly poor average performance relative to the counterfactual. Results are quite similar when we correct for systematic risk.

It is tempting to conclude that since the underperformance of selling strategies is driven by smaller initial positions, the costs in terms of overall portfolio performance associated with these transactions is likely to be small. However, this reasoning is incorrect provided that changes in portfolio weights induced by selling smaller initial positions are similar to those from selling larger ones. Holding trade size as a fraction of portfolio market value constant,

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

[46] In this section, we omit results for the subsample of developed markets, which are similar, for brevity.

[47] We also construct an alternative measure which updates when a "marriage"—defined as above according to whether the PM's buying activity on a given day increases a stock's weight by 50 percent—takes place. This measure obtains similar results, which are omitted for brevity.

Table 11. Post-trade sell returns relative to counterfactual by initial position size

This table presents the average returns relative to random sell counterfactuals for sell portfolios sorted by initial position size. We compute average returns of stock held minus returns of stocks sold. Initial position size is defined as the portfolio weight of a stock on the day of its first buy. We rank the initial position size measure within funds at a daily level and sort them into four bins from Smallest, Small-Med, Med-Large to Largest sizes. Columns represent sell performance measures at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. We report point estimates of average counterfactual returns for each portfolio at different horizons. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.2, are reported in parentheses.

Performance Measure	Initial Position Size Bins	Horizon				
		28	90	180	270	365
Baseline	Smallest	0.00	-0.31	-1.45	-2.10	-2.88
		(0.08)	(0.21)	(0.26)	(0.35)	(0.46)
	Small-Med	-0.01	-0.14	-0.32	-0.28	-0.50
		(0.07)	(0.15)	(0.24)	(0.28)	(0.38)
	Med-Large	0.07	-0.06	-0.05	-0.38	-0.29
		(0.07)	(0.15)	(0.23)	(0.30)	(0.37)
	Largest	0.15	0.06	0.04	0.25	0.28
		(0.07)	(0.14)	(0.24)	(0.29)	(0.37)
Factor-neutral	Smallest	-0.14	-0.70	-1.60	-2.29	-2.88
		(0.07)	(0.16)	(0.26)	(0.37)	(0.46)
	Small-Med	0.03	-0.12	-0.25	-0.15	-0.30
		(0.07)	(0.15)	(0.23)	(0.29)	(0.38)
	Med-Large	0.07	-0.02	0.09	-0.23	-0.09
		(0.07)	(0.15)	(0.24)	(0.30)	(0.36)
	Largest	0.14	-0.01	-0.09	0.08	-0.01
		(0.07)	(0.15)	(0.25)	(0.31)	(0.39)

the cost in foregone profits from a suboptimal trade are independent of the initial size of the position.[48] Indeed, we find that average trade sizes for sells are quite similar across both active share and initial position size bins.

6.2 Heuristic use and overall fund selling performance

In this section, we exploit the panel nature of our dataset in order to illustrate a more direct link between the performance of selling strategies and fund-level characteristics, such as the propensity to sell assets with extreme returns. 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

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

in performance. To operationalize this, we compute several fund-level characteristics and sort fund-weeks into categories based on these characteristics, then compute the average value-added associated with PMs' trades in each bin. Before proceeding, we note that this analysis is only able to identify correlations in the data, so it is not feasible to rule out other types of time-varying fund characteristics which simultaneously drive performance and observable properties of trading behavior.

We begin by considering the impact of heuristic use on performance. Based on the mechanism outlined in Section 6.1, we use the greater propensity to sell assets with extreme returns as a proxy for heuristic use.[49] To capture what we term 'heuristic intensity,' positions are sorted into 4 bins based on the fraction of stocks sold that are located in the extreme bins (Worst Loser and Best Winner) for each fund-week.[50] In order to reduce noise in the measures related to discreteness, we only compute the heuristic intensity measure for funds which hold at least 24 stocks (i.e., there are 4 stocks on average per bin) and only include fund-weeks in which at least 3 stocks are traded.[51] We then rank fund-weeks into four categories according to this measure to calculate relative performance of the associated selling decisions. Our primary rationale for a weekly frequency is that it provides a satisfactory balance between reducing potential noise in the sorting variable (by averaging over multiple trades) while still operating at a high enough frequency to capture between-manager variation in heuristic-use.

Panel A of Table 12 presents sample averages of counterfactual returns where funds are sorted into four bins based on heuristic intensity. In each week, we sort each portfolio into one of four categories based on its level of heuristic intensity. We find that our proxy for heuristic intensity is positively associated with significant underperformance. The highest levels of heuristic intensity are associated with the worst performance, especially at the longer horizons. Magnitudes are quite substantial: at a 1 year frequency, the highest level of heuristic intensity predicts an average of around 200 foregone basis points relative to a random-sell counterfactual. At the same time, average performance of sales which occur when PMs are selling fewer extreme positions is statistically indistinguishable from the counterfactual.

Appendix Table A.2 demonstrates the robustness of our results relating funds' high heuristic intensity and performance to using the factor-neutral performance measures. We consis-

[49] This greater propensity is a proxy for heuristic use because, as demonstrated in Section 6.1, PMs do not randomly sell assets from the consideration set of extreme returns.

[50] For instance, the mean of this heuristics intensity measure is around 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 Table A.1, 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.

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

Table 12. Post-trade sell returns relative to counterfactual by fund behavior

This table presents average returns relative to random sell counterfactuals for sell portfolios sorted by heuristics intensity, cumulative benchmark-adjusted fund returns since the beginning of a quarter, and a proxy for 'cash raising' episodes. We divide these measures into four bins from Lowest, Low-Med, Med-High and Highest, based on their rankings. Columns represent sell performance measures at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.2, are reported in parentheses.

Fund Characteristics	Bins	Horizon				
		28	90	180	270	365
Panel A: Heuristics Intensity Fraction of extreme stocks sold weekly (sorted across funds)	Lowest	-0.06	-0.07	-0.22	-0.31	-0.31
		(0.08)	(0.17)	(0.25)	(0.32)	(0.39)
	Low-Med	-0.02	-0.14	-0.23	-0.20	-0.04
		(0.06)	(0.11)	(0.17)	(0.21)	(0.26)
	Med-High	0.03	0.05	-0.20	-0.21	-0.34
		(0.05)	(0.11)	(0.17)	(0.21)	(0.29)
	Highest	0.09	-0.48	-1.34	-1.38	-2.00
		(0.09)	(0.18)	(0.29)	(0.38)	(0.46)
Panel B: Cumulative Benchmark-adjusted Fund Return since the beginning of a quarter (sorted across funds)	Lowest	-0.09	-0.62	-1.29	-1.62	-1.80
		(0.11)	(0.21)	(0.36)	(0.45)	(0.58)
	Low-Med	0.02	-0.04	-0.42	-0.44	-0.51
		(0.06)	(0.11)	(0.17)	(0.22)	(0.28)
	Med-High	0.01	-0.13	-0.30	-0.50	-0.44
		(0.06)	(0.11)	(0.17)	(0.22)	(0.28)
	Highest	0.12	0.05	0.09	-0.12	-0.14
		(0.07)	(0.15)	(0.24)	(0.33)	(0.38)
Panel C: Net Buy Weekly Number of stocks bought minus Number of stocks sold (sorted within fund)	Lowest	0.04	-0.30	-0.67	-1.10	-1.31
		(0.08)	(0.17)	(0.29)	(0.37)	(0.46)
	Low-Med	0.04	-0.19	-0.61	-0.95	-1.29
		(0.08)	(0.15)	(0.23)	(0.32)	(0.38)
	Med-High	0.04	0.11	-0.14	-0.25	-0.43
		(0.06)	(0.12)	(0.17)	(0.22)	(0.27)
	Highest	0.11	0.06	-0.11	0.24	0.46
		(0.06)	(0.12)	(0.17)	(0.22)	(0.26)

tently find that the highest heuristic intensity is associated with the worst performance is quite similar across all of the specifications. We also report results where bins are formed using within-manager variation (i.e., we compare the same PM's trades at different points in time, sorting time periods into bins based on our heuristics intensity estimates).[52]

In the preceding section we argued that both stages of the selling process are prone to heuristic thinking—limiting the consideration set to assets with salient attributes and then choosing to sell those that the PM is least attached to. The literature on heuristics and biases documents that people are more likely to rely on heuristics during situations when cognitive resources are in higher demand, such as in times of stress or when attention is

[52]Conclusions are similar if we focus on the developed subsample or sort on a monthly basis.

otherwise occupied (see Kahneman (2003) for review). Panels B and C of Table 12 consider two empirical proxies intended to capture periods emblematic of such episodes. As in Panel A, these measures are computed on a weekly basis and sort fund-weeks into four categories to capture either between or within-manager variation.

The first aims to capture performance when the PM is likely to be stressed. Institutional investors are known to take stock of their own performance based on calendar time, e.g. on a quarterly or yearly basis. Based on the conjecture that the PMs are more likely to be stressed when their overall portfolio is underperforming, we construct a measure that captures portfolio performance relative to the beginning of the preceding quarter. Table 12, Panel B demonstrates that selling quality is worst (relative to a random-sell counterfactual) when the PM's overall portfolio is underperforming the most. Panel C considers a measure aimed to proxy for sales that are more driven by cash raising considerations rather than forecasts of relevant performance metrics. We posit that observing larger bundles of assets being sold (relative to being bought) is emblematic of the manager being in "cash-raising mode." We compute the difference between the number of stocks bought and the number of stocks sold, where both measures are expressed as fractions of the number of stocks in the portfolio. We find that the difference between the number of stocks bought and sold predicts greater underperformance of the selling decisions.

7 Conclusion

We utilize a unique dataset and find evidence 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. We provide evidence that investors use heuristics when selling but not when buying, and that these heuristic strategies are empirically linked to the documented difference in performance.

As shown in Section 4, the comparison of trades on earnings announcement versus non-announcement days suggests that PMs do not lack fundamental skills in selling; rather, results are consistent with PMs devoting more cognitive resources to buying than selling. When decision relevant information is salient and readily available—as it is on announcement days—PMs' selling performance improves substantially. We propose a mechanism through which overall underperformance in selling can be explained by a heuristic two-stage selling process, where PMs limit their consideration set to assets with salient characteristics (extreme prior returns) and sell those they are least attached to (low active share assets). A proxy

for this heuristic strategy is associated with substantial losses relative to a no-skill random selling strategy.

The question remains of why professional PMs have not learned that their selling decisions are underperforming simple no-skill strategies. While we can only speculate, 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, anecdotal evidence from our interviews with PMs suggests that decisions are overwhelmingly focused on one domain over the other. In terms of time allocations, our understanding is that the vast majority of the investors' research resources are devoted to finding the next winner to add to the portfolio. Moreover, standard reporting practices are well-suited for evaluating performance of buying decisions: 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, and 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, the theoretical framework of Gagnon-Bartsch, Rabin, and Schwartzstein (2018) suggests that PMs may fail to recognize their underperformance in selling even in the long-run. The authors show that a mistaken theory such as the favorability of selling positions with extreme returns may persist in the long run because people channel their attention through the lens of this theory. As in Schwartzstein (2014), errors persist due to the person ignoring information that seems irrelevant and only updating her beliefs based on information that is attended to. Our findings imply significant benefits to creating environments where learning can occur more effectively. Moreover, our empirical results on a link between heuristic use and underperformance of selling strategies suggest that PMs adoption of decision aids and/or simple alternative selling strategies may substantially improve performance.

References

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

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

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

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

Barber, B. M., 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.

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

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

Benartzi, S. and R. H. Thaler (1995): "Myopic loss aversion and the equity premium puzzle," *The quarterly journal of Economics*, 110, 73–92.

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

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

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

Carhart, M. M. (1997): "On persistence in mutual fund performance," *The Journal of Finance*, 52, 57–82.

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

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

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

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.

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

DellaVigna, S. and J. M. Pollet (2009): "Investor inattention and Friday earnings announcements," *The Journal of Finance*, 64, 709–749.

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

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

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

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

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116, 1–22.

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

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

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

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

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

Gneezy, U. and J. Potters (1997): "An experiment on risk taking and evaluation periods," *The Quarterly Journal of Economics*, 112, 631–645.

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

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*.

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

Haigh, M. S. and J. A. List (2005): "Do professional traders exhibit myopic loss aversion? An experimental analysis," *The Journal of Finance*, 60, 523–534.

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

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

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

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

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

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

Hirshleifer, D., S. S. Lim, and S. H. Teoh (2009): "Driven to distraction: Extraneous events and underreaction to earnings news," *The Journal of Finance*, 64, 2289–2325.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

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

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

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

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

A Appendix

A.1 Data construction: additional details

This appendix provides additional detail about how we construct and clean our dataset, then presents some supplemental tables/figures referenced in the main text. The primary source of our analysis is Inalytics' holdings data and changes in holdings. After we clean the holdings data, we convert all the prices into USD using exchange rates mainly from Datastream. To ensure accuracy in exchange rates, we compare the exchange rate in Datastream with two other sources of exchange rates from Compustat and Inalytics. In the event of a discrepancy, we pick the two out of three that are the same, and this procedure takes care of discrepancy in all cases. We then augment the holdings data by merging in external prices series and forward and backward returns from CRSP (US stocks), Datastream (International stocks) and Inalytics' provided price series in this order. The external price series allow us to compute the market value of each holding precisely. There are instances where the market value of a stock (likely due to a measurement error in price/quantity) seems implausibly high, so we employ an iterative weight cleaning algorithm to eliminate these positions from the analysis. We provide additional details about these steps below.

We begin by outlining the key steps of our data cleaning procedure:

1. **Cleaning identifiers**: Inalytics has four main types of identifiers for stocks: SEDOL, ISIN CUSIP, and LOCAL. For the first three types of identifiers, they are distinguishable by the number of digits. SEDOL has 6-7 digits, CUSIP has 8-9 digits, and ISIN has 12 digits. In a few instances one type of identifier is mislabeled by the clients, so we correct them according to the number of digits.

2. **Merging in liquidated stocks with holdings data**: There are instances when a fund completely closes a position, so a stock disappears from the holdings data. Since our main trade measure is computed from the change in stock's holding, a position-closing trade will not be observed in the holdings. To do so, we first measure the minimum date of a fund and maximum date. Then, we compute tag the instance when the stocks disappear on some date between the minimum and maximum dates of each fund. We then append those stocks back to the holdings data in order to measure trading activities, from the changes in holdings accurately.

3. **Dropping portfolios without daily trades**: Some of the portfolios in the dataset do not receive daily time-stamped trade data. In these cases, only monthly holdings are reported and trades are imputed at the end of the month. To filter out these portfolios, we count the fraction of trades after the 27th of the month for each fund. If a fraction of trades after the 27th for a fund is over 50% or missing (in case of no trades observed), we drop the portfolio from the analysis sample. In addition, Inalytics independently provided a list of these portfolios from their internal records, essentially all of which were filtered out by this criterion. We also remove these manually flagged portfolios.

Next, we discuss some potential issues related to measurement errors in the price data. We use external price series from CRSP and Datastream, and we additionally have data provided by Inalytics. Inalytics relies on multiple data vendors such as MSCI or Thompson, as well as clients themselves, for price series in the holdings data. Since these prices are

collected for thousands of unique securities, they inevitably will be occasionally subject to measurement errors. In some cases, reported prices may be overstated, which may lead us to incorrectly characterize portfolio weights and potentially introduce measurement errors in various counterfactual return calculations. We rely on our external price series as the primary measure for a price when computing returns and portfolio weights throughout the analysis, though we take precautions to limit the potential influence of outliers.

When we compute cumulative returns for purposes of evaluating trading performance, we winsorize extremely small and large return realizations, some of which may be due to measurement errors in the price data. To mitigate the effect of the extreme returns when computing the average returns, we winsorize returns in the holdings dataset across all measures (raw,betaneutral) before forming portfolios. In our baseline results that we present here, we employ two winsorizing thresholds. First, we winsorize the cumulative return measures on each date across all positions at 0.1% on either tail. As an additional precaution, we winsorize large positive returns in the whole sample at the 99.99% threshold on the right tail of the distribution for raw returns and 0.01% on either tail for beta-neutral returns. The rationale is that beta-neutral returns can have also have extreme negative returns after adjusting for risks, so it is necessary to winsorize on both tails for risk-adjusted returns measures. We have also considered larger thresholds for winsorizing such as 0.3%, 0.5%, and 1% and obtain similar results.

In a handful of cases (e.g., because a stock split has led to an incorrectly high market value), the market value of a single position appears to be extremely large relative to the rest of the portfolio, which is indicative of a likely measurement issue. In order to flag situations when one errant price could cause our estimates of portfolio weights to be substantially biased, we employ an iterative procedure to drop potentially problematic positions. The idea of the procedure is to look for situations where the entire portfolio is concentrated in a single, extremely large position. For these purposes, we compute the market value of a position as the minimum of raw Inalytics price and raw external prices times the quantity of stocks. Then, we compute the position-level weight by dividing through by the dollar value of all positions. With these weights in hand, the procedure proceeds as follows. First, we compute the first three largest weights at a portfolio-date level. We then compute two measures 1) the difference between weights of the largest and second largest-held stocks and 2) the Difference between weights of the second and third largest-held stocks. If the first difference minus the second difference is over 15%, the largest weight is over 10% and the second difference is less than 5%, we flag the stock with the largest weight to exclude from the analysis and the weight calculation. We then recompute stocks' weights after the largest-held stocks are dropped and repeat the procedure to flag other stocks with unusually high weight in a portfolio. We repeat this algorithm until there is no stock with an unusually large weight in portfolios. This iterative weight-dropping algorithm finishes in 5 runs. There are 57,982 stock-date observations to be excluded from weight calculation. 84.3M observations (94.12%) in the holdings data have no weight errors. The first run of this algorithm cleans up weights for 4.1M observations (4.62%) in the holdings data. After five runs of this algorithm, whereby we exclude five stocks at most, 99.86% of holding observations have no weight problems. There are two portfolios for which this procedure still indicates the presence of a handful of extremely concentrated positions but their total number of associated observations is only 39,128 out of 89M holdings observations.

A.2 Standard error computation for counterfactual return measures

In sections 4.2 and 6.2, we estimate a number of average performance measures which are computed as a weighted average of cumulative long-short portfolio returns ($R_{buy} - R_{hold}$ and $R_{hold} - R_{sell}$). In this section, we briefly describe our methodology for estimating heteroskedasticity and autocorrelation robust standard errors associated with these mean calculations. We use a clustered standard error estimator which allows for autocorrelation and heteroskedasticity at the fund level following a similar approach to autocorrelation as Hansen and Hodrick (1980), who propose a standard error estimator for a general moving average process with a finite number of lags.

For each fund, let $\tau \in \{1, \ldots, T_i\}$ index dates for which we have return measures for fund i, and $w_{i,\tau}$ be the weight associated with the τth observation for fund i and $t_i(\tau)$ be the mapping from the index τ to calendar time t for fund i. Our weighted mean calculation takes the following form for various choices of horizon H:

$$\bar{R} = \sum_{i=1}^{N} \sum_{\tau=1}^{T_i} w_{i,\tau} R_{i,t_i(\tau):t_i(\tau)+H}, \tag{1}$$

where $\sum_{i=1}^{N} \sum_{\tau=1}^{T_i} w_{i,\tau} = 1$ and $R_{i,t_i(\tau):t_i(\tau)+H}$ is the cumulative return measure for the associated portfolio for fund i from $t_i(\tau)$ to $t_i(\tau) + H$. Note that (1) also corresponds with the OLS formula for the following regression:

$$\sqrt{w_{i,\tau}}\ R_{i,t_i(\tau):t_i(\tau)+H} = \sqrt{w_{i,\tau}} \left(\bar{R} + \varepsilon_{i,t_i(\tau):t_i(\tau)+H} \right). \tag{2}$$

We assume that $\varepsilon_{i,t_i(\tau):t_i(\tau)+H}$ has mean zero in the cross-section and in the time series, but $Cov[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H}, \varepsilon_{i,t_i(\tau_2):t_i(\tau_2)+H}]$ may be larger than zero for $|\tau_1 - \tau_2| < H$ due to the overlapping structure in calendar time. Then, a general formula for the clustered standard error for equation (2) in this case (see, e.g., Cameron and Miller 2015) is

$$V_{clu}[\bar{R}] = \sum_{i=1}^{N} \sum_{\tau_1=1}^{T_i} \sum_{\tau_2=1}^{T_i} w_{i,\tau_1} w_{i,\tau_2} Cov[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H} \cdot \varepsilon_{i,t_i(\tau_2):t_i(\tau_2)+H}] \mathbf{1}[t_i(\tau_1), t_i(\tau_2) \text{ in same cluster}]. \tag{3}$$

Because of the highly unbalanced panel structure of the daily counterfactual returns data, many fund-specific autocorrelation estimates are informed by a very small number of observations. To improve the precision of our estimates, we pool information across funds to obtain more accurate estimates of autocorrelations while still allowing for heteroskedasticity. Specifically, exploiting the identity

$$Cov[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H} \cdot \varepsilon_{i,t_i(\tau_2):t_i(\tau_2)+H}] = Corr[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H} \cdot \varepsilon_{i,t_i(\tau_2):t_i(\tau_2)+H}] \cdot Var[\varepsilon_{i,t:t+H}], \tag{4}$$

we apply a common correlation term across funds but allow variances to vary at the fund level. Specifically, we estimate a different autocorrelation for each lag length, which is measured as the number of weekdays between any two calendar time dates. We do this in three stages. First, at the fund level we estimate $Var[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H}] \approx \sum_{\tau=1}^{T_i} \varepsilon^2_{i,t_i(\tau_1):t_i(\tau_1)+H}$. Next, on

a fund-by-fund basis we estimate $Corr[\varepsilon_{i,t_i(\tau_1):t_i(\tau_1)+H} \cdot \varepsilon_{i,t(\tau_2):t(\tau_2)+H}]$ by taking the ratio of the sample mean of the cross terms at a given lag length to our fund-specific variance estimate. Then, finally, we aggregate to a single correlation number by taking a weighted average of the individual fund-level autocorrelation estimates and compute the standard error using equations (3-4) with the common correlation and fund-specific variance estimates, respectively.

Inspecting our autocorrelation estimates, it appears that the primary source of autocorrelation in our counterfactual performance measures comes from repeated trading of the same asset. For instance, if a PM sells a stock today, she is also somewhat more likely to sell it tomorrow and in the near future. Since we use cumulative return measures, consecutive trades of the same asset will mechanically be autocorrelated due to a moving average structure in the error term which comes from overlapping calendar time periods for cumulative returns.

Autocorrelation in trading appears to a very high frequency phenomenon. As a result, our estimated autocorrelation for our 1 year counterfactual sell measures is about 50% at a 1 trading day frequency and falls to 23% at a 5 trading day horizon. Within 1 and 3 months, these autocorrelations fall to about 7% and 3%, respectively, and decay rapidly towards zero thereafter. At long horizons, we find little to no evidence of autocorrelation even for periods which overlap in calendar time, as might be expected if stocks bought/sold tend to have different exposures to systematic risk relative to stocks not sold. In light of these estimates, since the computational cost of computing a full set of daily autocorrelations is quite high, we impose an estimate of zero after 100 business days, though overall estimates were quite similar when we allowed all autocorrelations for overlapping calendar time period to be nonzero.

Table A.1. Average heuristics intensity by bins of fund characteristics

This table reports averages of our measure of heuristics intensity at the fund-level, where funds are sorted into four bins according to various fund characteristics. We measure heuristics intensity by the fraction of positions sold in extreme bins of past position returns formed using each position's benchmark-adjusted return since time of purchase, capped at 90 days. We report this for subsamples formed by sorting on 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 :

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

Fund Characteristics	Lowest	Low-Medium	Medium-High	Highest
Panel A: Trading Style				
Weekly Net Buy	37.62%	40.61%	40.51%	40.78%
Monthly Turnover	39.37%	40.01%	40.17%	38.47%
Median Holding Length	38.89%	39.87%	39.36%	38.81%
Panel B: Past Fund Returns				
Fund past 2-day return	39.99%	39.26%	39.32%	40.22%
Fund past 7-day return	39.75%	39.02%	39.35%	40.47%
Fund past 30-day return	39.47%	39.09%	39.22%	40.91%
Fund past 60-day return	39.52%	38.96%	39.24%	41.01%
Fund past 90-day return	39.50%	39.03%	39.31%	40.91%
Fund past-year return	39.32%	39.24%	39.21%	40.65%
Fund past 2 year returns	39.10%	38.96%	38.95%	40.69%

Table A.2. Post-trade sell returns relative to counterfactual by heuristics intensity, robustness checks

This table presents average returns relative to random sell counterfactuals for sell portfolios sorted by heuristics intensity, where the heuristics intensity measure is sorted across funds by week (panel I) and within funds over time (panel II). We divide these measures into four bins from Lowest, Low-Med, Med-High and Highest, based on their rankings. Columns represent sell performance measures, with the baseline measure in Panel A and the factor-neutral measure in Panel B, at the following horizons: 1 month, 3 months, 6 months, 9 months, and 1 year. Heteroskedasticity and autocorrelation robust standard errors, computed using the method described in the section A.2, are reported in parentheses.

Heuristics	Bins	**Panel A: Baseline**					**Panel B: Factor-neutral**				
Intensity (weekly)	Horizon	28	90	180	270	365	28	90	180	270	365
I. Across Funds,	Lowest	-0.06	-0.07	-0.22	-0.31	-0.31	-0.07	-0.21	-0.29	-0.42	-0.30
by week		(0.08)	(0.17)	(0.25)	(0.32)	(0.39)	(0.08)	(0.17)	(0.25)	(0.32)	(0.38)
	Low-Med	-0.02	-0.14	-0.23	-0.20	-0.04	0.00	-0.07	-0.16	-0.11	0.03
		(0.06)	(0.11)	(0.17)	(0.21)	(0.26)	(0.06)	(0.11)	(0.17)	(0.21)	(0.25)
	Med-High	0.03	0.05	-0.20	-0.21	-0.34	0.04	0.04	-0.14	-0.20	-0.30
		(0.05)	(0.11)	(0.17)	(0.21)	(0.29)	(0.05)	(0.10)	(0.16)	(0.21)	(0.25)
	Highest	0.09	-0.48	-1.34	-1.38	-2.00	0.11	-0.59	-1.33	-1.58	-2.21
		(0.09)	(0.18)	(0.29)	(0.38)	(0.46)	(0.09)	(0.18)	(0.30)	(0.38)	(0.46)
II. Within Funds,	Lowest	-0.08	-0.19	-0.22	-0.30	-0.27	-0.10	-0.30	-0.27	-0.32	-0.13
by week		(0.08)	(0.15)	(0.22)	(0.29)	(0.35)	(0.07)	(0.15)	(0.21)	(0.28)	(0.34)
	Low-Med	-0.05	-0.02	-0.29	-0.32	-0.13	-0.04	-0.05	-0.27	-0.38	-0.24
		(0.06)	(0.12)	(0.19)	(0.23)	(0.29)	(0.06)	(0.12)	(0.20)	(0.24)	(0.29)
	Med-High	0.15	0.02	-0.36	-0.32	-0.63	0.15	-0.03	-0.31	-0.37	-0.62
		(0.07)	(0.13)	(0.21)	(0.26)	(0.32)	(0.06)	(0.12)	(0.20)	(0.26)	(0.30)
	Highest	0.02	-0.48	-1.19	-1.23	-1.76	0.07	-0.51	-1.15	-1.34	-1.90
		(0.09)	(0.17)	(0.27)	(0.35)	(0.42)	(0.09)	(0.17)	(0.27)	(0.36)	(0.43)