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

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

April 2019

Most research on heuristics and biases in financial decision-making has focused on non-experts, such as retail investors who hold modest portfolios. We use a unique data set to show that financial market experts – institutional investors with portfolios averaging $573 million – exhibit costly, systematic biases. A striking finding emerges: while investors display clear skill in buying, their selling decisions underperform substantially – even relative to strategies involving no skill such as *randomly* selling existing positions – in terms of both benchmark-adjusted and risk-adjusted returns. Across many specifications, foregone profits from underperformance in selling relative to a random-sell strategy are of a similar order of magnitude (though somewhat smaller) as the gains accrued from buying. We present evidence that an asymmetric allocation of cognitive resources such as attention towards buying relative to selling can explain this discrepancy. Looking at events when attention is more likely to be evenly split between prospective buys and sells – earning announcement days – we find that stocks bought *and* sold both outperform counterfactual strategies. We show that a heuristic process associated with limited attention can explain selling but not buying decisions. Assets with salient features in the form of extreme past returns are 50% more likely to be sold than those with zero benchmark-adjusted returns. Past returns have little predictive power for buying decisions. Lastly, we show that such heuristics are costly; funds that are more prone to use heuristic strategies exhibit the most underperformance in selling.

KEYWORDS: Heuristics, Behavioral Finance, Expert Decision-Making, Asset Pricing
JEL: D91

THIS IS A PRELIMINARY DRAFT. ALL COMMENTS VERY WELCOME!

*We thank Ned Augenblick, Aislinn Bohren, Pedro Bordalo, 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 and seminar participants at Booth School of Business, Caltech, Florida State, University of Amsterdam and University of Frankfurt for helpful comments.

†Booth School of Business, University of Chicago: klakowa@uchicago.edu
‡Inalytics Ltd: rick@inalytics.com
§Social and Decision Sciences, Carnegie Mellon University: aimas@andrew.cmu.edu
¶Sloan School of Management, Massachusetts Institute of Technology: ldws@mit.edu

Electronic copy available at: https://ssrn.com/abstract=3301277

1 Introduction

A large literature has demonstrated that market participants use heuristics and are prone to systematic biases. Individual investors have been shown to be overconfident (Barber and Odean 2001), sensation-seeking (Grinblatt and Keloharju 2009), and to have limited attention (Barber and Odean 2008). However, the majority of evidence documenting biased behavior of individual investors comes from data on retail investors (Barber and Odean 2011) or day traders (Barber, Lee, Liu, and Odean 2014), who generally hold modest portfolios.[1] It remains important to better understand the extent to which the decisions of market experts are prone to behavioral biases, and, if so, their effects on performance and real outcomes.

This paper examines the trade decisions of sophisticated market participants – experienced institutional portfolio managers (PMs) – using a rich data set containing their *daily* holdings and trades. Our data is comprised of 783 portfolios, with an average portfolio valued at approximately $573 million. More than 89 million fund-security-trading dates and 4.4 million trades (2.0 and 2.4 million sells and buys, respectively) are observed between 2000 and 2016. We evaluate performance by constructing counterfactual portfolios, and compare PMs' actual decisions to returns of the counterfactual strategy. Since PMs often need to raise capital by selling existing positions in order to buy, evaluating a selling decision relative to a counterfactual which is unrelated to existing holdings (e.g., a benchmark index) is not an appropriate comparison.[2] Instead, we evaluate selling decisions relative to a conservative counterfactual that assumes no skill: *randomly* selling an alternative position that was not traded on the same date.

We document a striking pattern: while the investors display clear skill in buying, their selling decisions underperform substantially. Positions added to the portfolio outperform 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 raw returns.[3] Restricting the sample to only developed markets increases forgone returns to over 90 basis

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

[2] An asset sold may outperform a benchmark index, but the sale may still be optimal depending on what is bought with that capital and what other assets could have been sold (e.g. an alternative may have gone up even more). In turn, a counterfactual for selling in a long-only portfolio must consider current holdings.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

points. 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.

We argue below that the stark discrepancy in performance between buying and selling appears to be driven by cognitive constraints, with the asymmetric allocation of limited attention towards buys and away from sells being the likely culprit. 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 likely able to capitalize (Easley, Engle, O'Hara, and Wu 2008). We exploit the variation in earnings announcements as predetermined shifters of attention that 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 (e.g. skill), then trades should look similar on announcement and non-announcement days.

We find that the performance of stocks sold on their earnings announcement days increases by more than 200 basis points over a yearly horizon relative to stocks sold on non-announcement days. 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 difference in the 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 fairly similar – but they are consistent with asymmetric overall allocation of attention between buying and selling decisions.

Electronic copy available at: https://ssrn.com/abstract=3301277

We also provide evidence that PMs are prone to use a heuristic that has been associated with limited attention when selling but not when buying. Prior work has shown that the salience of prior returns can affect investment decisions above and beyond the information they provide about future performance (Frydman and Rangel 2014; Frydman and Wang 2018). Indeed, various measures of prior returns are among the most readily available pieces of information about assets; trading terminals and research platforms all highlight past returns as amongst the first pieces of information available to an investor. And the propensity to buy and sell assets with extreme returns has been argued to stem from limited attention (Hartzmark 2014; Ungeheuer 2017). We find that PMs in our sample have substantially greater propensities to sell positions with extreme returns: both the worst and best performing assets in the portfolio are sold at rates more than *50 percent* higher than assets that just under or over performed. Instrumental motives do not seem to explain this pattern: it is robust to controlling for position size and holding length, and is unlikely to be explained by risk management motives. The pattern persists even after the inclusion of stock-date fixed effects which absorb a number of time-varying stock-specific unobservables. On any given day, the *same* asset is more likely to sold from a portfolio where it exhibited relatively extreme returns than from a portfolio where its recent performance stands out less relative 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.

Since prior returns may reflect changes in relative valuations, it is not unreasonable to see a correlation between extreme prior returns and trading behavior. Indeed, large price movements likely accompany arrival of new information, and models with heterogeneous beliefs point to instances in which information releases may generate more dispersion in beliefs and more trade (Kondor 2012). We argue against such a channel by exploiting the richness of our data to provide a natural comparison: PMs' decisions of which assets to buy. Importantly, 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 and implies that the public signal provided by recent relative returns tends to change PMs' ex-ante beliefs about future expected returns fairly little. Rather, as discussed further below, prior returns appear to guide the PMs' consideration sets of what assets to sell.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 points to PMs in our sample 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.[4] If this is the case, PMs likely spend relatively little time choosing between alternatives to sell in order to focus on buying decisions.

In Section 6.1 we propose and provide evidence for the conjecture that limited attention leads PMs to constrain their consideration set of what to sell to assets with extreme attributes on a salient dimension – for example, prior returns. From this set, PMs then choose to unload positions that they are least attached to. Importantly, the latter effect can generate systematic underperformance if the positions they are least attached to happen to be their newest ideas. We document significant evidence for this process. The more than 50% greater propensity to sell assets with extreme returns suggests that PMs are not considering their entire portfolio when choosing what to sell. We then identify new ideas by their active share in the portfolio, which is defined as the weight relative to the respective benchmark. Low active share positions can be observed for two reasons: 1) an asset with a large active share has lost value and not been replenished or 2) the PM is initializing a new idea but has not yet overweighted it relative to the benchmark. We can distinguish between the two types of assets by looking at prior position sizes: the latter would have never been a large position in the portfolio. We find that underperformance of selling strategies is particularly pronounced for these 'new ideas' – positions with low active share that have never occupied a large proportion of the portfolio. Sells of larger positions with higher active share are not associated with systematic underperformance.[5] Importantly, 'new idea' assets are also most likely to be sold and have the most pronounced relationship between extreme returns and

[4]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."

[5]Note that while a small position can only have a small impact on a portfolio's overall performance, the *incremental* impact on expected portfolio returns of a decision to reduce the size of a position that outperforms depends on the associated *change* in the stock's portfolio weight. We find that the average amount sold as a fraction of total portfolio market value is fairly similar regardless of the size of the position, suggesting that cumulative foregone returns associated with suboptimally selling small positions can be quite large.

Electronic copy available at: https://ssrn.com/abstract=3301277

selling probability. As we discuss further in Section 6.1, 'new ideas' also represent assets in which the PM has invested the least amount of time and effort to research and build up compared to high conviction positions that constitute a larger active share. The systematic selling of these low conviction assets is consistent with behavioral evidence on investors' decisions being affected by sunk costs and ownership effects (Anagol, Balasubramaniam, and Ramadorai 2018; Arkes and Blumer 1985; Heath 1995).

To provide further evidence that heuristic thinking is costly, we examine the correlation between the tendency to sell assets with extreme returns – which we view as an empirical proxy for heuristic thinking – and selling performance. We find that our proxy for heuristic thinking is indeed costly and at least partially explains the overall underperformance in selling. PMs who are most prone to this behavior (top quartile) forgo nearly 150 basis points annually relative to a random selling strategy. There appears to be substantial heterogeneity, with a significant minority of managers neither exhibiting this tendency nor underperforming in their selling decisions.

Additionally, if the use of heuristics is driven by limited cognitive resources, then its predicted effects will be exacerbated when these resources are limited further (Kahneman 2003). To examine this, we look at periods when PMs are likely to be experiencing periods of stress or selling in order in order to raise cash for buying decisions (i.e. attending to their selling choices even less). Because PMs primarily focus on their benchmark-adjusted returns, as a proxy for stress we study the quality of trades when the PM's overall portfolio is underwater relative to the benchmark in any given quarter. We document a strong relationship: the worse the overall portfolio is doing the lower the quality of sells relative to a random-sell counterfactual. There does not seem to be a similar relationship on the buying side, where the quality of decisions does not depend on portfolio performance. We then examine performance when PMs are likely to be selling in order to raise cash rather than focusing on sales as investment decisions. We proxy for these episodes by examining performance during periods when a large number of assets are being sold relative to buys. Based on this conjecture, we find that these periods are associated with a marked decrease in selling performance.

The paper proceeds as follows. Section 2 relates our findings to the existing literature. 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 concludes.

Electronic copy available at: https://ssrn.com/abstract=3301277

2 Related Literature.

Our findings contribute to the literature in finance documenting biased decision-making in individual investors (see Barber and Odean (2011) for review). While prior work has documented biases amongst experts in corporate finance settings, e.g. CEOs in charge of merger (Malmendier, Tate, and Yan 2011) or other restructuring decisions (Camerer and Malmendier 2007), substantially less research exists on the biases of expert investors.[6] In fact, for the most part the behavioral finance literature has assumed unbiased institutional investors exploiting the behavioral biases of retail investors (Malmendier 2018). Our documented findings suggest that this assumption may not be a valid one. Lastly, our results contribute to the literature demonstrating heuristics and biases amongst experts in domains such as sports (Green and Daniels 2017; Massey and Thaler 2013; Pope and Schweitzer 2011; Romer 2006), judges (Chen, Moskowitz, and Shue 2016), professional forecasters (Coibion and Gorodnichenko 2015), and retail markets (DellaVigna and Gentzkow 2017). This line of work highlights the persistence of behavioral biases despite significant experience and exposure to market forces.

The selling pattern we document is most related to the rank effect described in Hartzmark (2014). There, retail investors appear to exhibit a similar pattern in selling and buying behavior – unloading and purchasing assets with more extreme returns. However, it is not clear from the data whether these trading strategies are particularly maladaptive: this set of investors have been found to underperform the market in general and display a host of heuristics and bias such as the disposition effect (Odean 1998), overconfidence (Odean 1999), and narrow bracketing (Frydman, Hartzmark, and Solomon 2017).[7] Our results also relate to the analysis of Di Mascio, Lines, and Naik (2017), who used the Inalytics 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

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

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

Electronic copy available at: https://ssrn.com/abstract=3301277

compare performance of strategies to feasible alternatives (e.g. evaluating quality of actual selling strategies relative to counterfactual strategies).

Our results suggest that PMs systematically fail in porting 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 appear as if the PMs do not know how to sell – they seem to just not attend to the decision. Given the substantial foregone earnings even relative to a no-skill selling strategy, it is natural to ask why the investors do not appear to recognize their underperformance and adjust their behavior. While a full analysis of the learning environment is beyond the scope of the paper, recent theoretical work by Gagnon-Bartsch, Rabin, and Schwartzstein (2018) provides insight for this question. 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. Anecdotal evidence suggests that PMs extensively track both absolute and relative portfolio returns (required to evaluate buys) but rarely, if ever, calculate foregone returns from selling decisions. In Section 7 we further discuss the potential role of learning environments in the development of expertise for buying assets in our setting, and strategies for porting the expertise to selling decisions.

3 Data

This section discusses the data sources which are assembled for our analysis, presents descriptive statistics, and discusses a number of portfolio and position-specific variables which we use throughout the analysis.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

of portfolio managers.[8] 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 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. This leaves us with an unbalanced panel.

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 both active and inactive portfolios, with the 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.[9]

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.

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

[9] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

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	47335	0.404	0.240	0.267	0.385	0.522
Panel B : Position Level Summary (daily)						
Buying indicator	89.8M	0.0264	0.160	0	0	0
Selling indicator	89.8M	0.0226	0.149	0	0	0
Holding length since position open (days)	89.8M	484.4	512.9	119	314	679
Holding length since last trade (days)	89.8M	73.36	113.5	10	32	88
Holding length since last buy (days)	89.8M	112.3	152.4	18	57	144
Portfolio weight(%)	89.7M	1.2	1.61	.24	.79	1.65
1-day return (%)	82.1M	0.0511	4.15	-1.11	0.0115	1.17
Future 7-day return (%)	82.9M	0.205	5.830	-2.454	0.179	2.833
Future 28-day return (%)	82.8M	0.781	11.04	-4.634	0.810	6.181
Future 90-day return (%)	82.6M	2.561	20.16	-7.711	2.308	12.30
Future 180-day return (%)	81.5M	5.315	30.51	-10.46	4.164	18.88
Future 270-day return (%)	80.3M	7.873	38.54	-13.10	5.562	24.47
Future 365-day return (%)	78.9M	10.37	44.84	-15.08	7.241	29.73
Future 485-day return (%)	76.9M	13.43	51.12	-16.81	9.006	35.60
Future 605-day return (%)	74.9M	16.73	58.82	-18.73	9.871	41.01
Future 665-day return (%)	73.9M	18.53	62.94	-19.55	10.32	43.66
Future 730-day return (%)	72.7M	20.40	66.82	-20.13	10.86	46.43
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

Electronic copy available at: https://ssrn.com/abstract=3301277

We apply two primary filters to select the set of portfolios to include in our analysis. First, daily trading data are unavailable for a subset of portfolios or appear to be incomplete.[10] Second, we exclude funds that do not have a sufficient fraction (at least 80%) of portfolio holdings which could be reliably matched with CRSP or Datastream.In demonstrating the robustness of our results, we perform the analyses using data from developed markets only; these data have reliably high match rates with CRSP. 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.[11]

Differences from other datasets This sample offers some unique opportunities for the study of expert decision-making relative to other datasets in the literature. First, in contrast to the Large Discount Brokerage dataset of Barber and Odean (2000), which features portfolio holdings and trades of individual retail investors and has been used in numerous studies[12], 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 (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. Frazzini (2006)), 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 may not necessarily observe all trades made by a given institutional investor and often lack timely information on portfolio holdings.

[10] 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.

[11] 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.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 the description of the sorting variable construction.

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

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% of assets under management, but some funds are considerably more active in their trading behavior than others (the standard deviation is 5.7%).

While holding fairly diversified portfolios (average number of stocks is about 78 with a standard deviation of 68), funds in our sample remain active, 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

Electronic copy available at: https://ssrn.com/abstract=3301277

– is about 0.35% per day, or about 5.7% on an annualized basis. On average, a manager will initiate a sell trade for about 10% and a buy trade for about 15% of the stocks in his/her portfolio each month. We also characterize fund portfolios in terms of factor exposures by computing rolling Carhart 4-factor regressions (using the prior 1 year of daily data with the Fama-French international factors), adjusted for asynchronous trading.[13] The average market beta is about 1, and average exposures to the SMB, HML, and Momentum factors are fairly close to zero.

Panel A also reports the average benchmark-adjusted return that uses each portfolio-specific return series. The average fund in our sample beats its respective benchmark by about 0.22% per month, or 2.6% per year. This, in conjunction with the fact that funds' average betas are close to 1 and have little average exposure to the 3 other priced risk factors, suggests that these managers are highly skilled, earning returns above and beyond exposure to known risk factors. 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% and 2.2% of the time, respectively.

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

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

Electronic copy available at: https://ssrn.com/abstract=3301277

portfolio weight as a fraction of market value associated with each position on each date. The average stock has a weight of about 1.2% with a standard deviation of 1.6%.

Finally, we compute a number of measures of backward or forward-looking returns at the position 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.[14] For brevity, we only report summary statistics for forward-looking returns that are not adjusted for the benchmark. Volatilities of individual stocks are quite large, with a standard deviation of 45% at a 1 year horizon. As we discuss further below, we also consider several measures of prior position performance that are computed using periods of time which depend on holding period length.

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 lost) associated with managers' active buying and selling decisions.[15]

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's 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,

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

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 manager invests that money in a randomly selected stock from his other holdings. For evaluating sells, suppose that we learned that the manager was planning to sell a given stock tomorrow and hold the rest of the portfolio for a fixed period of time. We then suggest that instead of executing this trade, the manager randomly sells one of his/her other positions to raise the same amount of cash, holding the stock that was to be sold for the same period.

Since the information being used by us was also available to the manager, we would expect the decisions of a skilled PM to outperform our suggested strategies; this is due to the fact that, on the margin, our strategies are always feasible.[16] Note that the expected payoff from the counterfactual strategy (integrating out uncertainty about which stock is randomly selected) simply corresponds to the equal-weighted mean of realized returns across stocks held in the portfolio, which we denote by R_{hold}. The manager's selection adds value relative to the random counterfactual if $R_{buy} - R_{hold} > 0$ in the first example and if $R_{hold} - R_{sell} > 0$ in the second example. Following this logic, we compute $R_{buy} - R_{hold}$ and $R_{hold} - R_{sell}$ over horizons ranging from 1 week to 2 years for all buy and sell trades, respectively, to characterize the value-added associated from each.

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

[16] 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 conviction positions and thus deviates substantially from the benchmark. Purchasing the benchmark is feasible on the other hand.

Electronic copy available at: https://ssrn.com/abstract=3301277

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.[17] 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 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.[18]

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.[19]

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

[17] The four factors are the market excess return, the Fama-French (1993) 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.

[18] In additional unreported results, we have used additional information to construct a potentially "more intelligent" counterfactual. As we show in Figure 6 below, very few managers 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 one.

[19] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

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.[20]

We compute 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. When stocks traded are separated into multiple categories, our approach is similar except that we randomly allocate stocks to the different categories.

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 our random buy counterfactual, $R_{buy} - R_{hold}$. The average stock bought outperforms stocks held by more than 120 basis points over a one year horizon.

Figure 1, Panel B presents the average value-added, $R_{hold} - R_{sell}$, for sell trades. Recall that our measure is already signed so that positive values indicate that a trade helps portfolio performance relative to the counterfactual and negative values point to a trade hurting performance. In stark contrast to Panel A, these estimates suggest that managers' actual sell trades underperform a simple random selling strategy. 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 Figures 1-2 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, magnitudes are fairly similar between the baseline model and the three alternatives. Consistent with

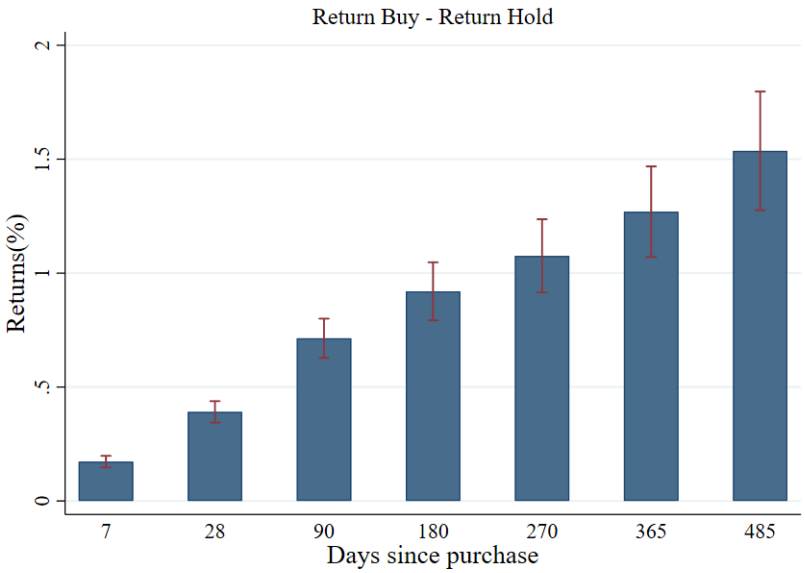
[20] We weight observations inversely to the number of trading days in a calendar year that the fund buys and sells a stock. This measure allows for an easier comparison across buys and sells, since we use the same weights across both types of trades. We obtain similar results when we instead weight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

Electronic copy available at: https://ssrn.com/abstract=3301277

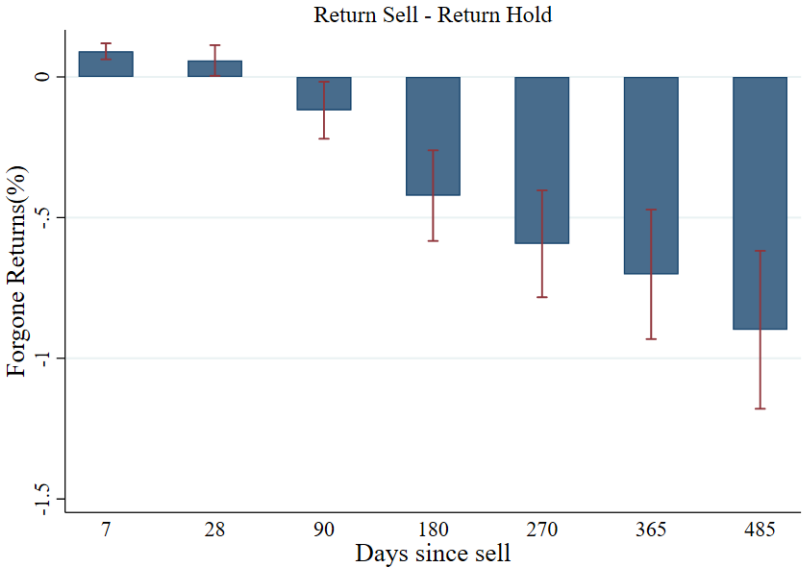
Figure 1. Post-trade returns relative to counterfactual

This figure presents average returns relative to random buy/sell counterfactuals for buy and sell trades. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. We then compute the average of these performance measures across all portfolios and dates, weighted inversely to funds' trading activity. Each bar represents average counterfactual returns in percentage over specified horizons on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using the Placebo method.

Panel A: Buy Trades, weighted



Panel B: Sell Trades, weighted



Electronic copy available at: https://ssrn.com/abstract=3301277

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) 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, and then report the difference between these measures for trades of stocks on their earnings announcement days versus all other days, where we weigh observations inversely to a fund's trading activity. Standard errors are computed using the Placebo method.

Sorting	Return Measures	**Panel A: Buy**			**Panel B: Sell**		
	Horizon	28	90	365	28	90	365
	Baseline	0.39	0.71	1.27	0.06	-0.12	-0.70
		(0.02)	(0.04)	(0.10)	(0.03)	(0.05)	(0.12)
I.	Baseline (Developed)	0.39	0.80	1.42	0.02	-0.25	-0.94
		(0.03)	(0.05)	(0.11)	(0.03)	(0.06)	(0.12)
Overall	Factor-neutral	0.34	0.58	1.16	0.03	-0.21	-0.80
		(0.03)	0.05	(0.12)	(0.03)	(0.06)	(0.12)
	Factor-neutral (Developed)	0.28	0.47	0.57	0.00	-0.19	-0.69
		(0.03)	(0.05)	(0.12)	(0.03)	(0.06)	(0.12)
	Baseline	-0.22	-0.06	-0.47	0.40	0.67	2.20
II.		(0.20)	(0.35)	(0.83)	(0.16)	(0.28)	(0.59)
Earnings	Baseline(Developed)	-0.26	0.40	-0.35	0.76	0.74	2.07
Announcement		(0.20)	(0.35)	(0.83)	(0.16)	(0.28)	(0.59)
(Difference in	Factor-neutral	-0.14	-0.09	-0.90	0.46	0.55	2.47
Counterfactual)		(0.20)	(0.35)	(0.83)	(0.16)	(0.28)	(0.59)
	Factor-neutral (Developed)	-0.25	0.03	-1.11	0.79	0.66	2.92
		(0.20)	(0.35)	(0.83)	(0.16)	(0.28)	(0.59)

much 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 as the baseline estimates. Results also are quite similar in the developed only sample as in the full sample.

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% or more of an asset to the portfolio as 'marriages' (100% corresponds to a opening a new position), and sell decisions that cut 50% or more from an asset as 'divorces' (100% corresponds to cutting a position completely). Table 4 below presents performance of marriages and divorces relative to the same counterfactual used in Table 3. Results are largely the same as for overall trade performance: marriages outperform the counterfactual while divorces underperform it. Across all specifications, es-

Electronic copy available at: https://ssrn.com/abstract=3301277

Table 4. Post-trade returns relative to counterfactual, large trades (marriage/divorce)

This table presents the average value added measures (post-trade returns relative to a random sell counterfactual) for large (marriage and divorce) trades under two measures of returns 1) raw cumulative returns and 2) factor-neutral cumulative returns, for the whole sample and the subsample of stocks from developed markets (see text for further details). Marriage is defined as the buy trades whose trade sizes exceed half of the existing beginning-of-the-day position sizes. Divorce is defined as sell trades whose sizes exceed half of the beginning-of-the-day position's sizes. We first present the overall average counterfactual returns for marriage and then divorce, where we weigh observations inversely to a fund's trading activity. Standard errors are computed using the Placebo method.

Return Measures	Bins	**Panel A: Marriage (Buy)**			**Panel B: Divorce (Sell)**		
Horizon		28	90	365	28	90	365
Baseline	Normal Trade	0.37	0.65	1.24	0.20	0.06	-0.46
		(0.03)	(0.05)	(0.11)	(0.03)	(0.05)	(0.12)
	Large Trade	0.39	0.81	1.08	-0.43	-0.67	-1.30
		(0.05)	(0.10)	(0.24)	(0.05)	(0.09)	(0.21)
Baseline (Developed)	Normal Trade	0.32	0.58	0.79	0.14	0.11	-0.32
		(0.03)	(0.05)	(0.11)	(0.03)	(0.06)	(0.12)
	Large Trade	0.36	0.71	0.75	-0.49	-0.73	-1.57
		(0.06)	(0.11)	(0.26)	(0.06)	(0.11)	(0.23)
Factor-neutral	Normal Trade	0.32	0.54	1.14	0.17	-0.05	-0.63
		(0.03)	(0.05)	(0.12)	(0.03)	(0.06)	(0.13)
	Large Trade	0.35	0.65	0.84	-0.45	-0.70	-1.14
		(0.06)	(0.11)	(0.28)	(0.06)	(0.10)	(0.23)
Factor-neutral (Developed)	Normal Trade	0.25	0.42	0.58	0.15	0.00	-0.43
		(0.03)	(0.05)	(0.11)	(0.04)	(0.06)	(0.13)
	Large Trade	0.31	0.52	0.38	-0.49	-0.74	-1.22
		(0.06)	(0.11)	(0.29)	(0.06)	(0.11)	(0.23)

timated foregone returns associated with divorces are considerably larger relative to those associated with other sell trades.

4.3 Performance on announcement days

We conjecture that the discrepancy in performance depicted in Panel A versus Panel B is driven by the asymmetric allocation of limited attention towards buying rather than selling. To provide evidence that discrepancy is due to attention rather than a fundamental difference in skill, we examine performance on days when decision-relevant information is salient and readily available – earnings 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.[21] 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

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

Electronic copy available at: https://ssrn.com/abstract=3301277

releases and conference calls (which even offer opportunities for managers to directly ask questions to the company), provide a wealth of new pieces of hard and soft information that are decision-relevant and improve trading performance (Easley et al. 2008). This information is both (relatively) easily available and salient, since earnings announcements are heavily covered by the financial press.

Figure 2 depicts the difference in performance of trades on announcement versus non-announcement days. Figure 2, Panel A looks at the difference in value-added of buy trades executed on earnings announcement days compared to other days.[22] 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. Assets sold on announcement days add substantially more value than those sold on non-announcement days: more than 200 basis points over a one year horizon. Table 3, Panel 2 shows that these results hold when adjusting for risk and restricting the sample to developed markets.

These findings suggest that when contemporaneous predetermined events are likely lead PMs to allocate more attention to selling, performance of those decisions improves substantially. This also provides evidence suggesting that the overall poor selling performance not necessarily due to a fundamental inability to sell.

5 Trading and Past Returns

Why might the performance of buying and selling decisions diverge? The preceding section provides some initial evidence that asymmetric allocation of limited attention can potentially explain the difference in performance. Such an allocation of attention would naturally lead to a prediction that buys would outperform sells on average, but not necessarily that sells would underperform feasible alternatives. Here we provide additional, more direct evidence consistent with an asymmetric attention allocation and defer our discussion of potential mechanisms to explain such underperformance until the next section.

Work in psychology and economics suggests that limited attention prevents a person

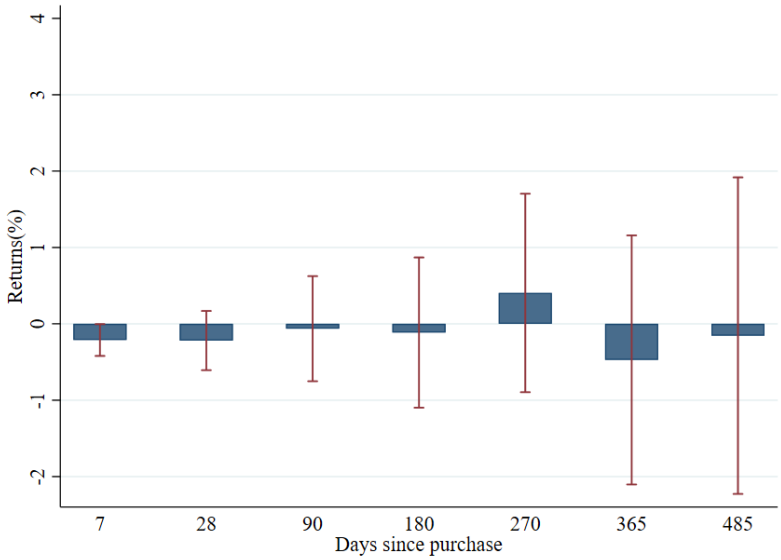
[22] Given the much smaller number of observations associated with stocks sold on earnings announcement dates, the average performance of sells on earnings announcement dates is positive but imprecisely estimated. Accordingly, we emphasize and report differences between average returns on non announcement days rather than levels.

Electronic copy available at: https://ssrn.com/abstract=3301277

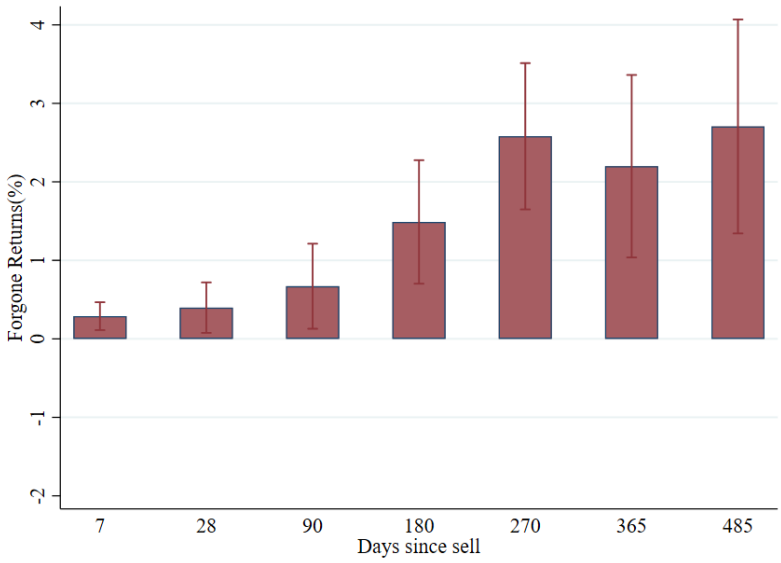
Figure 2. Difference in trading performance on earnings announcement vs other days

This figure presents the difference between average returns relative to random buy/sell counterfactuals for overall buy and sell trades that take place on firm's earnings announcement days (red bars) vs trades that are executed on all other days (blue bars). Earnings announcement dates are taken from the I/B/E/S database. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stocks held minus returns of stocks sold. We then compute the average of these performance measures across all portfolios and dates, weighted inversely to funds' trading activity and report the mean of the difference between the measure for stocks traded on there earnings announcement dates versus all other trades of the same type. Each bar represents the difference between the two average performance measures in percentage points over specified horizons indicated on the x axis. The range on the top of each bar is the confidence interval of the average returns of a portfolio at each horizon. The standard errors are computed using the Placebo method.

Panel A: Buy Trades



Panel B: Sell Trades



Electronic copy available at: https://ssrn.com/abstract=3301277

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). In this section, we explore when assets with extreme prior returns – one of the most salient attributes available to traders – are more likely to be traded. Consistent with our hypothesis about 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).

We emphasize prior returns because, relative to other forms of information relevant to the decision problem such as forecasted returns, data on past returns are ubiquitously available to PMs in our setting. This information is prominently featured on trading terminals, which typically break down past returns by year, quarter, month, day and since last purchase. Most news programs and popular webpages that cover financial markets include a segment which covers the stocks which experienced the largest moves on a given (both positive and negative).[23] The availability of this information and its close theoretical connection with the task at hand (forecasting changes in relative valuations), as well as the substantial amount of dispersion in idiosyncratic returns (as captured by their standard deviation, 51% over the average holding period in our sample), make it highly likely that past return is one of the most salient attributes of a given asset.

5.1 Measuring effect of prior returns on buying and selling

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, 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 $N_{bin}/2$

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 is a bit more challenging for buying decisions. Our first approach considers purchases of assets that already exist in the portfolio; this approach captures the majority of buys and includes most 'marriages' (up to 99%). 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 both approaches, our preferred measure of prior returns is computed as follows. For positions which were opened more than 1 year prior to the date of interest, we use the benchmark-adjusted return of the stock from 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.[24] We use this as our preferred measure because performance is often reported to clients at a quarterly frequency, and because, 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.[25]

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 5 bins based on the holding length since the last buy trade, and exclude the

[24] 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.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

bottom bin (shortest time elapsed since last purchase) from our calculations. We elect to do this to avoid a fairly mechanical relationship between our prior return measure, which has a variance which shrinks with the holding period, and the probability of buying/selling that can be generated if managers build up positions by splitting buy trades over short windows of time in order to minimize price impact.[26] 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.[27]

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 3 depicts the results for selling and buying decisions of assets that are already held graphically using a variety of different prior return measures, with 20 bins formed on each measure. Bins are sorted from left to right according to prior returns.

We begin with the buying probabilities. The probability of purchasing a stock already held is quite flat across the bins of prior returns. These results hold across all prior return measures considered and no pronounced patterns appear as we move towards more extreme bins in all cases.

A very different picture emerges for the selling probabilities. Assets with more extreme relative returns are substantially more likely to be sold relative to stocks in the central bins. An asset with a more extreme return is more than 50% more likely to be sold than assets with less extreme returns. Moreover, assets in these most extreme bins (1 and 20) have much higher selling probabilities than adjacent bins, increases which are not matched 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.

[26] 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.

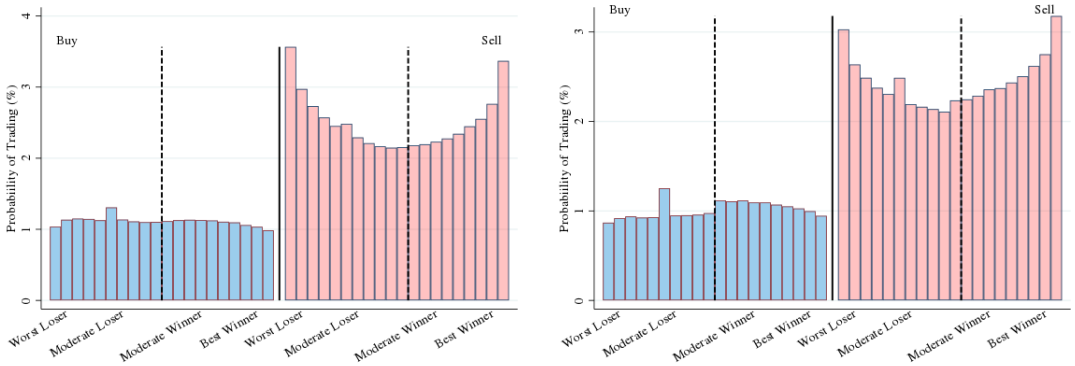
[27] Further, in Table 5, we report probabilities of buying using a prior return measure which does not depend on the time of initial purchase and do not impose the restriction on holding length. In that specification, our main results on the relationship between average buying probabilities and prior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

Electronic copy available at: https://ssrn.com/abstract=3301277

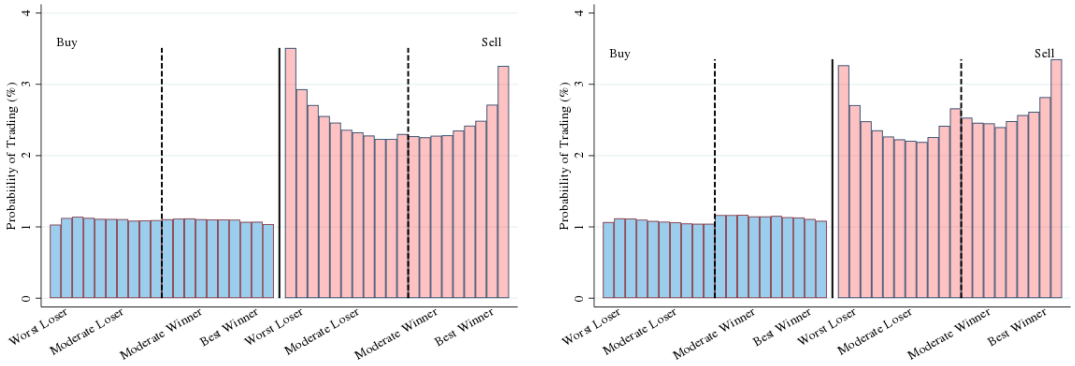
Figure 3. Probability of buying and selling based on past returns

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

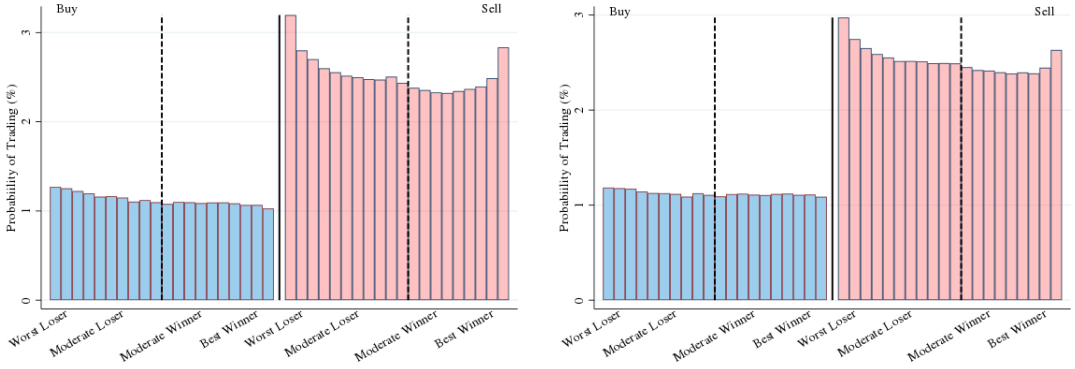
Panel A: Cumulative benchmark-adjusted returns capped at 1-quarter and 1-year



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



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



Electronic copy available at: https://ssrn.com/abstract=3301277

Table 5. Probability of buying and selling: larger consideration set and large trades

This table presents the probability of buying for security in the investor's holdings and consideration sets, as well as the probability of marriage (buying a large position) and the probability of divorce (selling a large position) for sell trades. The consideration set for each manager includes positions that investors buy within the next year for each month. Marriage (divorce) is defined as a buy (sell) trade whose size exceeds half of the beginning-of-the-day position's size. Panel A presents reports differences in probabilities, in percentage points, of buying selling by bins of past benchmark-adjusted returns for stocks in the consideration set by 20 bins of position's 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 3, 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 position's 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 (large buy)	0.02%	0.02%	0.01%	0.02%	0.01%	0.01%	0.01%	0.01%	0.16%
Divorce (large sell)	0.38%	0.27%	0.21%	0.07%	0.01%	0.01%	0.03%	0.12%	0.37%

Panel A of Figure 3 considers our baseline measure and an analogous one that caps relative returns at the longer horizon of 1 year calendar days instead of 90 days. In this second specification, the 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 3 but includes new buys using our expanded consideration set approach. 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 3, 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

Electronic copy available at: https://ssrn.com/abstract=3301277

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% 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 large sales are particularly likely to increase in response to extreme losses. Again, this pattern is not matched for the probability of engaging in large buys, which remain quite flat across bins of prior returns.

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 extremes in terms of both very long (1 year) and very short (1 week) horizon returns makes agency-based explanations where PMs are reluctant to report realized losses to their clients unlikely. Agency-based explanations also seem unlikely to explain the large jumps in probabilities observed between the 19^{th} and 20^{th} (1^{st} and 2^{nd}) bins relative to the 18^{th} and 19^{th} (2^{nd} and 3^{rd}) bins; in contrast, the top and bottom 5% of returns are much more likely to be displayed and salient to PMs while they are making selling decisions than adjacent bins (see Ungeheuer (2017) for direct evidence of this). This observation also mitigates concerns about risk management motives, since the relative risk of assets in extreme bins should be comparable to less extreme adjacent bins.[28]

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

[28] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

Table 6. Probability of Trading by Prior returns and Holding 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 holdings characteristics – holding lengths and position sizes, respectively – relative to the 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 and baseline probabilities of the omitted category 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. The selling (buying) probability is computed by the number of stocks sold (bought) in a particular bin divided by the total number of stocks in that bin. For Panel A, we do not exclude the bottom quintile of holding length since last buy when computing buying probabilities. For Panel B, we exclude recently bought stocks by excluding the bottom quintile of holding length since last buy.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 6 bins based on time elapsed since last buy (the variable we filter on) and prior returns. For this analysis only, we do not discard any stocks from the analysis based on this holding period measure. One can observe the mechanical patterns discussed in Section 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 all other holding periods. In contrast, assets in extreme bins are much more likely to be sold for 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, simple rebalancing motives (e.g., to reduce portfolio exposures to idiosyncratic risk) could motivate managers to sell positions with extreme positive returns that have become too large.[29] 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. 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 category remains bin 10 (slight loser positions). Results are similar with different prior return

[29] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

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(yesterday's assets under management), prior-month turnover, the volatility of a fund benchmark-adjusted returns over the past year, and prior month loadings on Fama-French Cahart regressions (calculated using the Dimson (1979) procedure using 1 year of prior daily returns). We control for position-level characteristics including linear and quadratic terms in holding lengths (overall and since last buy) and position sizes(% AUM) at the beginning of the day. Columns consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. We exclude recently bought stocks by dropping the bottom quintile of holding length since last buy from the analysis. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at 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

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

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% higher than the probability of selling a stock in bins 6-15, and always considerably higher than bin 19. 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 man-

Electronic copy available at: https://ssrn.com/abstract=3301277

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(yesterday's assets under management), prior-month turnover, the volatility of a fund benchmark-adjusted returns over the past year, and prior month loadings on Fama-French Cahart regressions (calculated using the Dimson (1979) procedure using 1 year of prior daily returns). We control for position-level characteristics including linear and quadratic terms in holding lengths (overall and since last buy) and position sizes(% AUM) at the beginning of the day. Columns consider various fixed effects including Fund, Date, Fund x Date and Stock x Date for different comparisons. We exclude recently bought stocks by dropping the bottom quintile of holding length since last buy from the analysis. The coefficients and t-statistics are reported for the variables included for each model. The standard errors for each model are clustered at 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

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

agers 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. On the downside, we see a similar effect: stocks in bin 1 are considerably more likely to be sold than those in bin 2.

Turning to Table 8, the relationship between buying probabilities and prior return measures is much more muted. Most of the coefficients are insignificant despite being estimated on a sample of over 50 million observations. Even the significant coefficients are substan-

Electronic copy available at: https://ssrn.com/abstract=3301277

tially 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 0. Turning to extreme gains, we observe many significant coefficients, but the differences between central and extreme bins (e.g., bins 16-18 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.[30] Together, these results are consistent with non-instrumental motives stemming from cognitive constraints driving selling but not buying decisions.

6 Heuristics and performance

In this section, we propose a potential mechanism which can link the use of heuristics to systematic underperformance of selling strategies relative to feasible counterfactuals. 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 the average value-added associated with 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 lead to the use of heuristics in both stages of the process (Hauser 2014).[31] Next, we outline and provide evidence for heuristics in both stages of the selling decision that can explain the reported above, including underperformance relative to random sell counterfactuals.

[30] Increases in selling probabilities for very extreme bins are even larger in additional unreported results other measures of prior return rankings. We elected not to report these estimates since magnitudes are quite similar to Figure 3. These pronounced increases in probabilities of selling extremes are not matched for buys.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 ranking and filtering assets on some salient attributes.[32] 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.[33] In turn, extreme deviations in relative returns in either the positive or negative direction naturally emerge as candidate characteristics 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.[34] Our results from Section 5.2 provide strong evidence that extreme returns at least partially govern the consideration of sales: assets that are in the top or bottom 5% based on prior returns are nearly 50% more likely to be sold relative to those that just over or underperformed, a pattern not observed for buy decisions.[35]

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 "do I have a reason to keep (or discard) 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.[36]

[32] 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.

[33] 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.

[34] 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 that "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."

[35] 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.

[36] 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,

Electronic copy available at: https://ssrn.com/abstract=3301277

Assets with high active share are typically referred to as the PM's high conviction positions: greater deviation from the benchmark implies that the manager has spent a good deal of effort building up the position over time and becoming familiar with the firm. This costly process likely generates abundant reasons to retain these high active share assets.

Positions with low active share can manifest for two main reasons: 1) a position which was initially large has done very poorly, or 2) the PM has added a position to the portfolio but has not yet built it up over time.[37] A stock in the former category was previously high conviction one – and may still well be. In contrast, assets in the latter 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 make it a high conviction one. In turn, heuristic thinking would generate fewer reasons to to keep it, while at the same time, experiencing an extreme return produces a salient reason to sell. Consistent with behavioral evidence on sunk cost and ownership effects, the greater time and effort that managers have spent on developing high active share positions may also lead them to be more psychologically attached to them when deciding what to sell (Anagol et al. 2018; Arkes and Blumer 1985; Heath 1995). But selling newer, less entrenched ideas may be exactly the wrong thing to do since the information that drove the initial buying decision is likely to still be fresh, leading the asset to outperform a random holding from the portfolio. We now proceed to demonstrate that this process appears to explain the underperformance of the PM's 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 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 below.

Results are consistent with sales of low conviction assets being particularly easy to justify.[38] First, examining the baseline probabilities, we note that low active share positions are

an asset that is overweighted generates excess returns when it goes up and excess losses when it goes down.

[37] A third alternative is that the PM has actively reduced a formerly large position through prior sells. Given that the majority of our portfolios are quite concentrated, negative active shares are observed very infrequently.

[38] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

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 (bin 1) and the average across the other bins. Calculations for 8 categories of prior returns, formed from 20 bins of past position returns, are reported in rows. Below, we also report the baseline selling probability for the 10th bin. The selling probability is computed by the number of stocks sold in a particular bin divided by the total number of stocks in that bin. We exclude recently bought stocks by sorting based on the holding length from last buy on each day within a portfolio and dropping the bottom quintile of holding length since last buy.

	Active share bins				
Prior return bin	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

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 conviction assets. The probability of selling a stock with the lowest active share and lowest prior return bin is 5.6%, or 140% larger than the baseline probability of selling (which is 2.3%). Assets in these bins are also 155% 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 change relatively less in response to moderate gains and responses to the most extreme gains in bin 20 is similar regardless of active share.

We then examine whether sales of low conviction assets tend to underperform relative to a counterfactual (i.e., the stocks actually sold end up outperforming the sale of randomly chosen holdings). Panel A of Table 10 depicts the performance of sales relative to a random counterfactual by bins of stocks' active share. We see a stark contrast in performance: low conviction assets underperform substantially more than sales of high conviction assets in

Electronic copy available at: https://ssrn.com/abstract=3301277

Table 10. Post-trade sell returns relative to counterfactual by Initial Position Size and Active share

This table presents the average returns relative to random sell counterfactuals for sell portfolios sorted by Active share and Initial Position Size. For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. We compute average returns of stock held minus returns of stocks sold. Initial position size is defined as portfolio weights of stocks on the day of their first buy. We rank the Active share and Initial Position Size within fund at a daily level and sort them into four bins from Lowest, Low-Med, Med-High to Highest sizes. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, and 1 year. We report point estimates of average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate), where we weigh observations inversely to the number of trades per year of a fund. Standard errors are computed using the Placebo method.

Return Measures	Bins	**Panel A: Active share**			**Panel B: Initial Position Size**		
Horizon		28	90	365	28	90	365
Baseline	Lowest	-0.12	-0.28	-0.53	0.00	-0.31	-2.88
		(0.03)	(0.06)	(0.13)	(0.05)	(0.08)	(0.19)
	Low-Med	0.27	-0.08	-1.56	-0.01	-0.14	-0.50
		(0.05)	(0.09)	(0.21)	(0.05)	(0.08)	(0.22)
	Med-High	0.10	-0.03	-0.84	0.07	-0.06	-0.29
		(0.05)	(0.09)	(0.19)	(0.05)	(0.08)	(0.21)
	Highest	0.24	0.27	-0.14	0.15	0.06	0.28
		(0.05)	(0.08)	(0.18)	(0.04)	(0.08)	(0.20)
Baseline (Developed)	Lowest	-0.13	-0.20	-0.38	-0.13	-0.20	-0.38
	Low-Med	0.15	-0.11	-1.68	0.15	-0.11	-1.68
	Med-High	0.10	-0.12	-0.82	0.10	-0.12	-0.82
	Highest	0.19	0.31	0.01	0.19	0.31	0.01
Factor-neutral	Lowest	-0.21	-0.55	-0.65	-0.21	-0.55	-0.65
	Low-Med	0.20	-0.15	-1.19	0.20	-0.15	-1.19
	Med-High	0.16	0.13	-0.84	0.16	0.13	-0.84
	Highest	0.32	0.29	-0.35	0.32	0.29	-0.35
Factor-neutral (Developed)	Lowest	-0.20	-0.50	-0.55	-0.20	-0.50	-0.55
	Low-Med	0.14	-0.22	-1.25	0.14	-0.22	-1.25
	Med-High	0.19	0.08	-0.65	0.19	0.08	-0.65
	Highest	0.30	0.39	0.00	0.30	0.39	0.00

Electronic copy available at: https://ssrn.com/abstract=3301277

the top active share bin, where the latter actually tend to outperform the counterfactual at shorter horizons and have performance in line with the counterfactual at longer ones. The latter result is consistent with PMs holding on to high conviction assets when thinking fast, so when a sale is observe, it is more likely to be an informed one.

As noted above, one reason why low active share positions could be particularly bad sales is that they may represent the PM's newer ideas. These stocks were sufficiently interesting to put into the portfolio but the investment thesis has not developed enough to justify making a large bet. To investigate whether stocks with these characteristics are indeed associated with underperformance of selling strategies, 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.[39] Panel B of Table 10 reports average performance of sales in each bin. 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 and hold in the developed market subsample.

While it is tempting to conclude that since the underperformance of selling strategies is driven by smaller initial positions, that the costs in terms of overall portfolio performance associated with these transactions is likely to be small. As we noted above, this reasoning is incorrect provided that changes in portfolio weights induced by selling smaller initial positions are similar to those from selling larger ones. Holding trade size as a fraction of portfolio market value constant, the cost in foregone profits from a suboptimal trade are independent of the initial size of the position.[40] 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/position-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

[39] 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% – takes place. This measure obtains similar results.

[40] 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

ask whether patterns in funds' actual trading strategies are associated predictable differences in performance. To operationalize this, we compute several fund and position-level characteristics and sort trades into categories based on relative levels of these characteristics, then compute the average value-added associated with each bin. Before proceeding, we note that this analysis is only able to identify correlations in the data, so it is not feasible via these designs 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 potential implications for performance (or lack thereof) of selling positions with extreme returns. 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.[41] To capture what we term 'heuristic intensity,' we calculate the fraction of stocks sold that are located in the extreme bins (worst loser and best winner) for each fund-week.[42] We then rank fund-weeks into 4 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 attention allocation.

Panel A of Table 11 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. The left panel plots average performance of buy trades, while the right panel plots average performance of sell trades. 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 150 foregone basis points relative to a random-sell counterfactual. At the same time, sales of managers that appear less prone to heuristic thinking do not underperform the counterfactual – directionally, the coefficient is actually positive.

Appendix Table A.2 demonstrates the robustness of our results relating funds' heuristic

[41] 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. They sell the low conviction ones, which tend to increase in value compared to a randomly chosen holding.

[42] For instance, the mean of this heuristics intensity measure is 0.4 on a monthly basis, which would imply (through a simple application of Bayes' rule) that the likelihood of a stock being sold in the extreme bin is 4/3 the likelihood of a stock being sold in one of the central bins. In Appendix 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

Table 11. Post-trade returns relative to counterfactual by fund behavior

This table presents average returns relative to random buy/sell counterfactuals for buy and sell portfolios sorted by heuristics intensity, cumulative benchmark-adjusted fund returns since the beginning of a quarter, and a proxy for buying episodes (net buy). For buy trades, we compute average returns of stocks bought minus returns of stocks held on each day. For sell trades, we compute average returns of stock held minus returns of stocks sold. The heuristics intensity is computed by measuring the fraction of sells in the lowest and highest of 6 bins of cumulative returns capped at 1-year at weekly horizons. Weekly net buy is computed by subtracting the unique number of positions bought per week from the unique number of positions sold per week. In panel C, we compute each portfolio's cumulative benchmark-adjusted return since the beginning of the quarter, then sort funds into bins based on this measure for each week in the sample. Since the units of the net buy measure differ substantially across funds, we rank these measures within portfolios across all weeks in the sample. We divide these measures into four bins from Lowest, Low-Med, Med-High and Highest, based on their rankings. Columns represent buy or sell performance measures at the following horizons: 1 month, 3 months, and 1 year. We report point estimates of average counterfactual returns for each portfolio at different horizon as well as their standard errors in parenthesis (below the point estimate), where we weigh observations inversely to funds' trading activity. Standard errors are computed using the Placebo method.

Fund Characteristics	Bin	Buy			Sell		
		Horizon			Horizon		
		28 days	90 days	1 year	28 days	90 days	1 year
Panel A: Heuristics Intensity fraction, across-fund weekly	Lowest	0.42	0.85	0.92	0.10	0.19	0.10
		(0.07)	(0.12)	(0.25)	(0.07)	(0.12)	(0.25)
	Low-Med	0.33	0.53	0.69	-0.01	-0.11	-0.60
		(0.05)	(0.08)	(0.21)	(0.05)	(0.08)	(0.21)
	Med-High	0.32	0.54	0.68	0.05	-0.12	-0.30
		(0.04)	(0.08)	(0.17)	(0.04)	(0.08)	(0.17)
	Highest	0.44	0.63	0.98	0.02	-0.32	-1.49
		(0.06)	(0.11)	(0.24)	(0.06)	(0.11)	(0.24)
Panel B: Cumulative Benchmark-adjusted Fund Return since the beginning of a quarter (sorted across funds)	Lowest	0.40	0.58	1.70	-0.09	-0.62	-1.80
		(0.14)	(0.18)	(0.48)	(0.14)	(0.18)	(0.48)
	Low-Med	0.36	0.80	1.04	0.02	-0.04	-0.51
		(0.06)	(0.12)	(0.28)	(0.06)	(0.12)	(0.28)
	Med-High	0.43	0.74	0.96	0.01	-0.13	-0.44
		(0.04)	(0.07)	(0.15)	(0.04)	(0.07)	(0.15)
	Highest	0.44	0.72	1.88	0.12	0.05	-0.14
		(0.05)	(0.09)	(0.19)	(0.05)	(0.09)	(0.19)
Panel C: Net Buy Weekly Number of stocks bought minus Number of stocks sold (sorted within fund)	Lowest	0.37	0.60	0.13	0.04	-0.30	-1.31
		(0.14)	(0.18)	(0.48)	(0.14)	(0.18)	(0.48)
	Low-Med	0.54	1.19	2.24	0.04	-0.19	-1.29
		(0.06)	(0.12)	(0.28)	(0.06)	(0.12)	(0.28)
	Med-High	0.36	0.72	1.23	0.04	0.11	-0.43
		(0.04)	(0.07)	(0.15)	(0.04)	(0.07)	(0.15)
	Highest	0.36	0.60	1.40	0.11	0.06	0.46
		(0.05)	(0.09)	(0.19)	(0.05)	(0.09)	(0.19)

Electronic copy available at: https://ssrn.com/abstract=3301277

intensity and performance to counterfactuals that account for risk factors and restrict our sample to developed markets. Although magnitudes vary somewhat, the main result that the highest heuristic intensity is associated with the worst performance is quite similar across all of the specifications. If anything, magnitudes are somewhat larger for the alternative specifications relative to the baseline.[43]

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 lack a readily available reason to keep them. 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 otherwise occupied (see Kahneman (2003) for review). Panels B and C of Table 11 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 4 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 11, Panel B demonstrates that selling quality is worst (relative to a random-sell counterfactual) when the PM's overall portfolio is underperforming the most – consistent with the notion that stress exacerbates suboptimal decision-making. We do not observe a similar relationship between portfolio performance and quality of buying decisions. 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 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 number of sell trades relative to buy trades predicts greater underperformance of the selling decisions. Consistent with attentional resources being allocated away from selling decisions, we do not find that the quality of buying decisions is affected by these measures.

[43] In additional unreported results, we find a similar result – the highest heuristic intensity is linked with substantially worse selling performance – when we form bins 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 heuristics intensity). Results are also similar if we sort on a monthly basis.

Electronic copy available at: https://ssrn.com/abstract=3301277

7 Conclusion

We use a unique data set to show that financial market experts – institutional investors managing portfolios averaging $573 million – display costly, systematic biases. A striking finding emerges: while investors display skill in buying, their selling decisions underperform substantially – even relative to random sell strategies. We provide evidence that investors use heuristics when selling but not when buying, and that these heuristic strategies can explain the difference in performance.

As shown in Section 4, the performance of buy decisions is substantially superior to the performance of sell decisions. Moreover, the PMs do not lack the fundamental skills to sell well; they appear to be devoting more cognitive resources to buying than selling. When relevant information is salient and readily available – on earnings announcement days – the investors are able to effectively incorporate it into their selling decisions which end up outperforming the counterfactual. We find that the overall underperformance in selling can be explained by a two-stage heuristic selling process where PMs limit their consideration set to assets with salient characteristics – extreme prior returns – and sell those that lack a readily-available reason to keep – low conviction assets. A proxy of this 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 here, the environment in which fund managers make decisions offers several clues. As Hogarth (2001) notes, the development of expertise requires frequent and consistent feedback. While it is feasible to generate this type of feedback for both buy and sell decisions, in practice the environment in which fund managers make decisions is overwhelmingly focused on one domain over the other. The vast majority of the investors' research resources are devoted to finding the next winner to add to the portfolio. Purchased assets are tracked, providing salient and frequent feedback on the outcomes of buying decisions. This process appears successful in producing expertise – purchased assets consistently outperform the benchmark. In comparison, paltry resources are devoted to decisions of what to sell. From our interviews with the PMs, 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 et al. (2018) suggests that PMs may fail to recognize their underperformance in selling even in the long run. Our findings imply significant benefits to creating environments where learning can occur more effectively. Given the heterogeneity in selling skills – managers who do not use the

Electronic copy available at: https://ssrn.com/abstract=3301277

extreme-selling heuristic outperform those who do – fund managers who are underperforming can adopt learning tools and simple alternative selling strategies to substantially improve performance.

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.

Arkes, H. R. and C. Blumer (1985): "The psychology of sunk cost," *Organizational behavior and human decision processes*, 35, 124–140.

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.

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.

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 (2016): "Being surprised by the unsurprising: Earnings seasonality and stock returns," *The Review of Financial Studies*, 30, 281–323.

Chen, D. L., T. J. Moskowitz, and K. Shue (2016): "Decision making under the gambler?s fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *The Quarterly Journal of Economics*, 131, 1181–1242.

Electronic copy available at: https://ssrn.com/abstract=3301277

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

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

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

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.

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

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

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.

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*.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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

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.

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

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

Kahneman, D. and 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.

Kondor, P. (2012): "The more we know about the fundamental, the less we agree on the price," *The Review of Economic Studies*, 79, 1175–1207.

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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

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," .

Electronic copy available at: https://ssrn.com/abstract=3301277

A Appendix

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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 stock?s 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.

Electronic copy available at: https://ssrn.com/abstract=3301277

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

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

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

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

Electronic copy available at: https://ssrn.com/abstract=3301277

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

This table presents the average counterfactual returns for buy and sell trades under two return measures (raw, factor-neutral) for the whole sample and the subsample of developed markets, sorted into four bins based on our measure of heuristics intensity. Heuristics intensity is computed by measuring the fraction of sell trades in the lowest and highest of 6 bins of cumulative returns capped at 1-year, sorted weekly across funds. See text for further details on variable definitions. Panel A and B report mean counterfactual returns of each measure for buy and sell trades respectively, weighted by a fund's trading activity, as well as their standard errors in parenthesis (below the point estimate). Each cell represents the average counterfactual returns in percentage over specified horizons and levels of heuristics intensity. Standard errors are computed using the Placebo method.

Return Measures	Bins	**Panel A: Buy**			**Panel B: Sell**		
Horizon		28	90	365	28	90	365
Baseline	Lowest	0.42	0.85	0.92	0.10	0.19	0.10
		(0.05)	(0.10)	(0.20)	(0.07)	(0.12)	(0.27)
	Low-Med	0.33	0.53	0.69	-0.01	-0.11	-0.60
		(0.04)	(0.07)	(0.16)	(0.04)	(0.07)	(0.17)
	Med-High	0.32	0.54	0.68	0.05	-0.12	-0.30
		(0.03)	(0.06)	(0.13)	(0.04)	(0.07)	(0.17)
	Highest	0.44	0.63	0.98	0.02	-0.32	-1.49
		(0.05)	(0.09)	(0.20)	(0.06)	(0.11)	(0.24)
Baseline (Developed)	Lowest	0.41	0.79	0.71	-0.08	0.09	0.11
		(0.05)	(0.09)	(0.20)	(0.08)	(0.15)	(0.29)
	Low-Med	0.34	0.60	0.71	0.00	-0.14	-0.81
		(0.04)	(0.06)	(0.15)	(0.05)	(0.08)	(0.20)
	Med-High	0.31	0.53	0.70	0.05	-0.08	-0.40
		(0.04)	(0.07)	(0.16)	(0.05)	(0.08)	(0.20)
	Highest	0.36	0.53	0.85	-0.03	-0.37	-1.82
		(0.05)	(0.09)	(0.20)	(0.06)	(0.12)	(0.24)
Factor-neutral	Lowest	0.33	0.59	0.89	0.10	0.12	0.04
		(0.05)	(0.09)	(0.21)	(0.07)	(0.13)	(0.27)
	Low-Med	0.26	0.43	0.60	0.04	-0.10	-0.39
		(0.04)	(0.06)	(0.16)	(0.04)	(0.07)	(0.16)
	Med-High	0.29	0.48	0.40	0.03	-0.15	-0.30
		(0.03)	(0.06)	(0.13)	(0.04)	(0.07)	(0.16)
	Highest	0.38	0.52	0.88	-0.07	-0.59	-1.90
		(0.05)	(0.09)	(0.20)	(0.07)	(0.12)	(0.25)
Factor-neutral (Developed)	Lowest	0.34	0.53	0.52	-0.02	0.06	0.16
		(0.05)	(0.10)	(0.22)	(0.08)	(0.14)	(0.31)
	Low-Med	0.26	0.49	0.66	0.04	-0.15	-0.47
		(0.04)	(0.07)	(0.17)	(0.05)	(0.08)	(0.19)
	Med-High	0.28	0.44	0.38	0.05	-0.12	-0.34
		(0.03)	(0.07)	(0.15)	(0.05)	(0.09)	(0.19)
	Highest	0.29	0.42	0.72	-0.09	-0.66	-2.11
		(0.05)	(0.10)	(0.21)	(0.06)	(0.14)	(0.27)

Electronic copy available at: https://ssrn.com/abstract=3301277