Trading Costs of Asset Pricing Anomalies

ANDREA FRAZZINI, RONEN ISRAEL, AND TOBIAS J. MOSKOWITZ *

First draft: October 23, 2012
This draft: September 9, 2015

Abstract

Using over a trillion dollars of live trading data from a large institutional money manager across 21 developed equity markets over a 16-year period, we measure the real-world transactions costs and price impact function facing an arbitrageur and apply them to trading strategies based on empirical asset pricing anomalies. We find that actual trading costs are an order of magnitude smaller than previous studies suggest. In addition, we show that small portfolio changes to reduce transactions costs can increase the net returns and break-even capacities of these strategies substantially, with little tracking error. Use of live trading data from a real arbitrageur and portfolios designed to address trading costs give a vastly different portrayal of implementation costs than previous studies suggest. We conclude that the main capital market anomalies – size, value, and momentum – are robust, implementable, and sizeable in the face of transactions costs.

*Andrea Frazzini is at AQR Capital Management, e-mail: andrea.frazzini@aqr.com. Ronen Israel is at AQR Capital Management, e-mail: ronen.israel@aqr.com. Tobias Moskowitz is at the Booth School of Business, University of Chicago, NBER, and AQR Capital Management, email: tobias.moskowitz@chicagobooth.edu. We thank Cliff Asness, Malcolm Baker, John Campbell, Josh Coval, Darrell Duffie, Eugene Fama, John Heaton, Bryan Kelly, John Liew, Gregor Matvos, Michael Mendelson, Hitesh Mitthal, Stefan Nagel, Lasse Pedersen, Amit Seru, Amir Sufi, Richard Thaler, Annette Vissing-Jorgensen, Brian Weller, and seminar participants at the University of Chicago, Harvard Business School, Stanford University, and University of California Berkeley for helpful comments. We also thank Laura Serban for outstanding research assistance. Moskowitz thanks the Center for Research in Security Prices and the Fama-Miller Center at the University of Chicago for financial support.

Empirical asset pricing largely focuses on the expected gross returns of assets. For investors, however, the net of transaction cost returns are the critical input for investment decisions. A large literature documents several strong predictors for the cross-section of average returns, which have been thrust into the efficient markets debate as challenges to standard asset pricing models. However, an understanding of the net of transaction cost returns and capacity limits of strategies based on these predictors is crucial to this debate. Determining whether such strategies are robust, implementable, and sizeable or whether they face significant practical impediments or "arbitrage limits" that prevent traders from profiting from them is a key and open question.

We explore this issue by examining the cross-section of net of trading cost returns using a unique dataset of $1.05 trillion dollars of live trades from a large institutional money manager from August 1998 to September 2013, across 21 developed equity markets. The data offer a singular look into the real-time trading costs of an investor who resembles the theoretical "arbitrageur" for most of these anomalies. Our institutional investor implements strategies very similar to those in the literature. We evaluate the robustness of the three most prominent capital market anomalies — size, value, and momentum[1] — to trading costs and assess their implied capacity limits.

The results help shed light on the market efficiency debate surrounding these return predictors. If some predictors do not survive real-world trading costs, or only survive at small dollar investments (low capacity), then limits to arbitrage may simply prevent them from being exploited and disappearing. On the other hand, if returns after trading costs are significantly positive at very large fund sizes, then strategies based on these predictors may offer large profit opportunities for an arbitrageur or perhaps represent a risk factor in the economy that a significant fraction of the market is exposed to. While we remain agnostic on risk versus non-risk based explanations for these predictors, we estimate the frictions and real-world costs of implementing portfolios based on these factor/anomalies, which are relevant to both sets of theories.

Our live trading data contains some unique features not previously studied in the literature. First, we can identify the actual real-time price impact of a trade at various trade sizes. We see whether the trade was a buy or sell, the market price at trade initiation, the exact amount traded, and the exact execution price for each share traded. These data allow us to calculate the actual price impact of over one trillion dollars of equity trades over a 16-year period. In addition, our dataset

[1] These three anomalies dominate the cross-sectional return landscape. See Fama and French (1996, 2008, and 2012), Asness, Moskowitz, and Pedersen (2012), and Ilmanen (2013) for a description of cross-sectional stock return predictors that appear most robust in the data. While other anomalies can be examined with our data, too, we focus on size, value, and momentum for brevity and because they are the focus of current asset pricing models.

contains the *intended* trade, where the difference between the executed trade and intended (theoretical) trade can be used to measure “implementation shortfall” in order to capture the opportunity cost of a trade. We can also differentiate between trades to speculate versus cover a position, and sell versus short-sell trades, all providing a more accurate picture of the real trading costs to long-short portfolios common in the literature. In addition to estimating these costs for U.S. stocks, our data also covers 20 other developed equity markets internationally, providing the first look at the trading costs in different markets, for similar strategies deployed markets simultaneously.[2]

Armed with real trading costs from live trading data, we evaluate the robustness of size, value, and momentum strategies to these costs. Assessing standard long-short strategies commonly used in the literature – such as those of Fama and French (1993) pertaining to SMB (a portfolio of small minus big stocks), HML (a portfolio of high minus low book-to-market equity stocks), and UMD (a portfolio of up minus down momentum stocks) – we find that size, value, and momentum each survive transactions costs at fairly substantial fund sizes.

We begin by using the exact quantities traded by our manager facing the exact actual trading costs the manager faced in real time, to reconstruct SMB, HML, and UMD. The resulting portfolios have correlations greater than 0.96 with the original Fama-French factors and represent actual traded quantities of those factor portfolios in real time. The average volume traded by our manager in these long-short factor portfolios is 9.7, 6.0, and 6.2 billion dollars per month in the U.S. (11.8, 7.2, and 8.4 billion dollars per month internationally), for SMB, HML, and UMD, respectively. At these trade sizes, average realized costs per year are 1.47, 1.35, and 3.03 percent, which given historical return premia of 3, 5, and 8.2 percent per year for SMB, HML, and UMD, respectively, safely rejects that net of trading costs returns to these strategies are non-positive.

To give some perspective, the implied fund sizes of SMB, HML, and UMD traded by our manager are 18.2, 9.4, and 5.2 billion dollars in U.S. equities, which by themselves are larger than the break-even market-wide capacities proposed in the literature (Chen, Stanzl, and Watanabe (2002), Lesmond, Schill, and Zhou (2003), and Korajczyk and Sadka (2004)). This suggests that

[2]While a few other studies have looked at proprietary trade data from a large institution or set of institutions (e.g., Keim (1995), Keim and Madhavan (1997), Engle, Ferstenberg, and Russell (2008)), the data used in these studies is much more limited in breadth and depth, covering only a couple of years of U.S.-only data and a small cross-section of stocks. Our sample spans 16 years of trading activity across more than 20,000 stocks in 21 different equity markets. In addition, the data used in these other studies only reflects *executed* trades, containing no information about the process generating the trades. Finally, none of these other studies address the question we pose in this paper on the net of trading cost returns to asset pricing anomalies.

estimated trading costs and capacity limits from the literature are grossly different than the realized costs and *actual* traded quantities of a large institutional manager.

We then use the live trading database to estimate a trading cost model based on 1) market conditions (e.g., the VIX, time trends), 2) stock characteristics (e.g., market cap, idiosyncratic volatility), and most importantly, 3) trade size (the fraction of daily volume traded), where we scale up the fraction of daily trading volume (DTV) traded to compute the break-even fund size or capacity of each strategy, which is the fund size where trading costs would entirely offset the return premium to a strategy. We find that break-even fund sizes for long-short factors of size, value, and momentum are more than an order of magnitude greater than capacity sizes suggested in the literature. In the U.S. equity market alone, the break-even capacities of SMB, HML, and UMD are $275, $214, and $56 billion, respectively.

Our trading cost and capacity estimates differ substantially from those in the literature. Previous papers generally find a huge effect from transactions costs on the viability of trading strategies. For instance, Knez and Ready (1996) find that trading costs swamp the profitability of lead-lag strategies between large and small firms. Mitchell and Pulvino (2001) estimate commissions and price-impact costs for a merger arbitrage portfolio and find that trading costs account for much of the returns. Chen, Stanzl, and Watanabe (2002) estimate that only small maximal fund sizes are attainable before costs eliminate profits on size, value, and momentum strategies. Lesmond, Schill, and Zhou (2003) find that trading costs eliminate the profits to momentum strategies and Korajczyk and Sadka (2004) find break-even fund sizes for momentum portfolios are in the two to five billion dollar range. We find that our real-world estimates of transaction costs are an order of magnitude smaller than those in the literature, and therefore that the break-even capacity of these strategies is more than an order of magnitude larger.

One key difference between our study and the literature is that previous studies compute trading costs for the *average* investor, which are an order of magnitude higher than the costs realized by our large institutional manager. Hence, the average investor costs estimated in the literature may be a very poor proxy for the marginal investor's trading cost in these strategies.

Our unique approach has some advantages and disadvantages relative to other studies on trading costs. While explicit trading costs, such as commissions and bid/ask spreads, can be measured relatively easily and have been studied at length,[3] the full costs of trading are dominated by price impact (Kyle (1985), Easley and O'Hara (1987), Glosten and Harris (1988), Hasbrouck (1991a,

[3] Schultz (1983) and Stoll and Whaley (1983) examine the effect of commissions and spreads on size portfolios.

1991b), Huberman and Stanzl (2000), Breen, Hodrick, and Korajczyk (2002), Loeb (1983), Keim and Madhavan (1996, 1997), and Knez and Ready (1996)), where transaction costs are a function of the size of a trade and its ability to move the market price of an asset. Indeed, this is certainly true of our manager. For instance, the effective bid-ask spread across all trades in our database averages only 0.015% per year in the U.S., with many trades facing zero or near-zero spreads. These costs are tiny relative to the price impact costs facing a large institutional trader.

Previous studies attempting to estimate price impact costs rely on theoretical models and typically use one of two types of data: daily spread and price information or intra-day transaction level data aggregated across all traders over fixed time intervals.[4] These data provide estimates of transactions costs for the average trader, reflecting aggregated volume of traders with different intentions, objectives, and information. Our trading costs, on the other hand, are those facing a real arbitrageur investing in portfolios based on anomalies discovered in the academic literature in real time, which we argue is a better proxy for the marginal investor's trading costs in these strategies. From a market efficiency standpoint, we are interested in the costs facing the marginal arbitrageur.

The drawback of our data is that we can only measure costs for one particular money manager, albeit a large manager who invests in many of the anomalies we investigate. The ability to generalize our cost estimates to other strategies studied in the academic literature or perhaps to other investors depends on how exogenous the trading costs are to the portfolios being traded by our manager. For instance, if the portfolios themselves are a function of trading costs, then the cost estimates we obtain would be most relevant to the specific portfolios run by the manager, but less relevant for portfolios not identical to the manager's. We argue and show that the trading costs we estimate are fairly independent from the portfolios being traded. First, we only examine the live trades of our manager's longer-term strategies, where the portfolio formation process generating the desired trade is *separate* from the trading process executing it. This separation exists because the set of intended trades is created from longer-term models (such as value and momentum effects, which we study here) and from client-specific mandates that often adhere to a benchmark subject to a tight tracking error constraint. This separation of portfolio formation from trading execution is a feature only of the

[4] Examples of trading cost measures using daily price and spread data include Roll (1984), Huang and Stoll (1996), Chordia, Roll, and Subrahmanyam (2000), Amihud (2002), Acharya and Pedersen (2005), Pastor and Stambaugh (2003), Watanabe and Watanabe (2006), Fujimoto (2003), Korajczyk and Sadka (2008), Hasbrouck (2009), and Bekaert, Harvey, and Lundblad (2007). Examples using transaction level data from the trade and quote (TAQ) data from the NYSE, Rule 605 data (e.g., Goyenko, Holden, and Trzcinka (2009)), or proprietary broker data (e.g., Engle, Ferstenberg, and Russell (2008)), include Goyenko (2006), Sadka (2006), Holden (2009) and for a comparison of daily versus intra-daily measures, see Lesmond, Ogden, and Trzcinka (1999), Lesmond (2005), Engle, Ferstenberg and Russell (2008), Lehmann (2003), Werner (2003), Hasbrouck (2009), and Goyenko, Holden, and Trzcinka (2009).

longer-term models and would not apply, for instance, to intra-day or higher frequency trading strategies. This is precisely why we focus only on the long-term trades and exclude all short-term intra-day strategies from our data. Given a portfolio, the manager then trades to its intended positions in the cheapest possible manner using a proprietary trading algorithm that importantly, cannot make any buy or sell decisions. The algorithm merely determines how best to execute the desired trades. For example, it can decide how patient to trade (minutes versus days), but not what to trade (or not to trade). In our sample, 99.9% of all intended trades are completed, but not necessarily immediately. The average realized trade horizon for completion is slightly less than one day and no more than three days, indicating very low tracking error to the intended trade. However, trading over several days versus instantaneously has huge trading cost consequences.

To rule out concerns over the endogeneity of portfolio weights to (expected) transactions costs, we also examine trading costs using only the first trade from new inflows of long-only mandates that specifically adhere to a benchmark. The initial trades to a benchmark from new inflows are essentially exogenous to trading costs since there is little scope for deviation. We find that the trading costs from these "forced" initial trades are identical to those from all other trades.

The only endogenous choice in security selection with respect to trading costs our manager makes is the universe of stocks the manager considers in forming portfolios. Our manager does not trade in extremely small and illiquid microcap or penny stocks or stocks with limited DTV. We show, however, that portfolios created from the stocks the manager does trade have extremely high correlation to those used in academia created from the entire universe of securities.

Finally, since the standard academic portfolios pertaining to these anomalies are not designed to address or consider transactions costs in any way, it may be premature to conclude what the efficacy or capacity of these strategies might be without recognizing what a real-world arbitrageur might do to respond to transactions costs. For example, Garleanu and Pedersen (2012) show theoretically how a strategy's net performance can be significantly improved by taking transaction costs into account when computing optimal portfolios. We therefore construct portfolios that seek to minimize trading costs, subject to maintaining similar exposure to the original factor. Using a static portfolio optimization and our trading cost model, we minimize expected trading costs subject to a tracking error constraint that seeks to maintain the "style" or exposure of the strategy to the factor being studied. We analyze how much after-transaction costs returns can be improved across different anomalies through portfolio optimization and assess the tradeoff between reducing trading costs and introducing tracking error across the different strategies. We find that for all strategies net alphas are improved significantly without incurring large tracking error or significant reductions in gross alphas. The optimizer achieves this by both lowering turnover and changing the types of stocks traded.

The optimized portfolios further increase the capacity of the anomalies by several times. At less than one percent tracking error, optimized versions of size, value, and momentum can reach break-even sizes of $1.9 trillion, $1 trillion, and $160 billion among U.S. securities, respectively. These sizes are three to six times larger than our base case estimates and several orders of magnitude higher than estimates in the literature.

Our findings indicate that the three main capital market anomalies related to size, value, and momentum, are robust, implementable, and sizeable. These conclusions differ significantly from the literature because 1) we use live trading costs of a real-world arbitrageur, which are many times smaller than the average investor demanding liquidity and 2) we examine portfolios that consider transactions costs in their design. Both innovations result in trading costs (break-even fund sizes) that are several orders of magnitude smaller (larger) than previous studies suggest. We believe our estimates are much closer to the real-world transactions costs facing an arbitrageur, who is closer to the theoretical marginal investor in these strategies, and therefore the relevant costs for discussion of market efficiency surrounding these asset pricing anomalies.

The paper proceeds as follows. Section I describes the live trading data and anomaly-based strategies we examine and discusses the manager's portfolio formation and trading processes. Section II outlines our methodology for measuring transactions costs and examines realized trading costs over time, across markets, and for different trade types. Section III presents the cross-section of net after-cost returns, where we use the actual costs from the actual quantities traded by our manager to assess the net returns of commonly used factor portfolios. Section IV then uses the live trading data to estimate a price impact function, which we use to calculate break-even capacities of the factors. Section V uses the model for trading cost optimization and assesses the tradeoff between minimizing trading costs and tracking error across anomalies. Section VI concludes.

I. Live Trading Data, Trading Process, and Equity Portfolios

We describe our live trading data, including a description of how our institutional manager generates desired trades and executes and manages trading, and present summary statistics on the live trading data. We also detail the long-short portfolios we examine in the face of trading costs, that are commonly used to capture various anomalies from the literature.

A. Trade execution data

Our trading costs data are drawn from the internal trade execution database maintained by our institutional manager, who manages global strategies ranging from benchmark-oriented long-only strategies (including mutual funds) to absolute return alternative strategies in a variety of asset classes. The database is compiled by the execution desk and covers all trades executed algorithmically (described below) in any of the firm's funds since inception, excluding trades associated with intra-day or higher frequency models.[5]

The data contains information about orders, execution prices, and quantities. We exclude all non-equity and emerging markets trades, restricting the sample to developed market equity transactions, which includes cash equities and equity swap transactions. For each individual order we collect the trade-order-set identifier, stock identifier, order size, trade horizon, order type, region and country, size and portfolio type, benchmark, ex-ante predicted beta (described below), model-implied price P^{theory}, arrival price P^{start}, execution time, execution price P^{ex} and execution quantity Q^{ex}.

The trade-order-set has a unique identifier that allows the mapping of individual trades to corresponding portfolios. A trade-order-set is simply a basket of trades submitted together. Usually a trade-order-set contains trades from a single fund, but in some cases trades from multiple funds are bundled in the same set. In a given day, the same stock can appear in multiple order sets. Therefore, we use a trade-order-set stock pair as our cross sectional unit of observation for each date. Hereafter, we will use the term "order" to indicate a trade-order-set stock pair submission date triplet.

The order size is the desired number of shares to trade when orders are submitted. For our data, which excludes higher frequency statistical arbitrage trades, the desired trade amount will almost always equal the sum of the actual execution quantities (desired quantities are filled 99.9% of the time). The only times the desired amount of shares fails to fill is due to trading suspensions in a stock or other rare events. While the quantity being traded is not a choice variable for the trading algorithm, the duration (or patience) of the trade is. The trade horizon is the target trade duration (in days) at the time of order submission, but since the actual total trade duration is endogenous and depends on realized fill rates, it may, and often does, differ from the target duration. Our estimate of trading costs, therefore, includes how our manager executes trades, which is designed to minimize costs of a given desired quantity, but importantly where the trade quantity is essentially fixed.

---

[5] Our manager does not engage in high-frequency trading, but does run some portfolios on short-term intra-day price pressure and statistical arbitrage. The trades from these models are specifically excluded here since they are necessarily and by design endogenous to trading costs – i.e., trades for these strategies are only executed when costs are expected to be low. We focus exclusively on trades from the long-term models, where the trading algorithm is not allowed to make buy or sell decisions, but is allowed to be patient making the trade, as discussed below.

The order type indicator identifies the transaction as "buy-long", "buy-to-cover", "sell-long", and "sell-short". The size and portfolio indicators classify each portfolio as "large" or "small" in terms of market capitalization (based on their relevant benchmark, such as Russell 1000 or 2000 in U.S. equities, for example) and "long/short" or "long-only" based on the portfolio's mandate.[6] The benchmark identifies the relevant benchmark, while the ex-ante predicted beta is a forecasted beta for each stock to the benchmark at the time of order submission, which is based on a regression using the past year of daily stock returns on the corresponding benchmark. The model price is defined as the stock price at portfolio formation. For portfolios constructed during market hours this is equal to the current price. For portfolios constructed off market hours this is the latest closing price. The arrival price is defined as the current price at the time the first order is submitted to the market, which is recorded by the trading algorithm when orders are submitted. The execution prices are time-stamped and incorporate commissions. Since the trading algorithms break orders into smaller pieces (and typically dynamically submit and cancel a range of limit orders as discussed in the next subsection), there are many executions per order.

The final dataset contains 32,158,831 orders covering 9,128 stocks between August 1998 and September 2013 and totaling $1,046,940,000,000 in trades.[7] Panel A of Table I reports the amount traded (in $billions) for each year from 1998 to 2013 in the U.S. equity market, across the 20 international markets, for large cap and small cap stocks separately, and for long-short and long-only portfolios separately. Not surprisingly, the amounts traded have grown substantially over time, with faster growth in the U.S. from $1.3 billion traded in 1998 to $121 billion in 2012 (the last full year in our data). Equities traded in international markets have grown from $1.7 to $56 billion. Large cap (as defined by the Russell 1000 or MSCI) portfolios comprise the bulk of trades, accounting for $1.013 trillion out of the total $1.046 trillion traded over the sample period. Nevertheless, $33.5 billion worth of trades were placed among small cap stocks. Finally, $775.46 billion, or roughly 74%, of trades occur for long-short strategies similar to those constructed in the literature.

Panel B of Table I reports time-series means, medians, standard deviations, minimums, and maximums for the number of stocks traded, number of countries traded, and number of exchanges traded on per year. A minimum of 386 stocks across eight countries on 11 exchanges are traded each year with a maximum of over 5,000 stocks in 21 markets across 34 exchanges. Panel C of Table I

[6] We classify relaxed constraint portfolios, such as 130-30 or 140-40, as long-only. These represent only a small fraction of trades and excluding them or reclassifying them does not alter any of our results.

[7] Throughout our analysis we exclude netting trades (which are trades settled internally among accounts). Including netting trades does not affect any of our results.

reports year-by-year time series averages in the style of Fama and MacBeth (1973) of the average trade size (in $thousands) and fraction of daily volume traded. The average (median) trade over time is $657,600 ($344,900), but there is substantial heterogeneity in the size of trades. The standard deviation of trade size is about $1 million and ranges from $52,600 to $5,993,200 over our sample period. As a fraction of daily volume, the average (median) trade represents 1.1% (0.5%) with a standard deviation of 2% of total daily volume that ranges from 0.1% to 13.1%. This variation allows us to measure the price impact function across a wide range of trade sizes.

B. Portfolio formation and trading process

Understanding how portfolios and trades are created and executed is important for interpreting the trading cost measures we obtain from these data. Trades are executed using proprietary, automated trading algorithms designed and built by the manager. Before describing how these algorithms are designed, we first briefly describe the portfolio generation process.

Importantly, we only examine trades where the portfolio generation process is separate from the trading process. Specifically, we examine the live trades of our manager's longer-term strategies, where the portfolio formation process generating the desired trade is separate from the trading process executing it. This is because the set of intended trades is primarily created from longer-term models (such as value and momentum effects, which we study here) and from specific client mandates that often adhere to a benchmark subject to a tight tracking error constraint. This separation of portfolio formation from the trading process is a feature of the longer-term trades we examine and would not apply, for instance, to high frequency or intra-day trading strategies like statistical arbitrage. Hence, we exclude all such trades from our trade data.

For the lower frequency trades we focus on, a set of desired trades is first obtained for a portfolio from a set of optimal holdings based on a particular factor or model, often using historical data. While the models employed are often more complex, a lot of the portfolio generation process is based on factors similar to those from the academic literature and those we investigate in this paper and describe in the next subsection. Once a theoretical portfolio is determined, an optimization process generates portfolio weights subject to various constraints which can include risk constraints and other client mandated constraints (e.g., tracking error to a benchmark, shorting constraints, industry and position limits, etc.). This portfolio then implies a set of trades that moves the current portfolio to the desired portfolio. These trades are then loaded into the trading algorithm which is responsible for implementing or executing the desired trades at the lowest cost they can, but where the algorithms do not make any explicit aggregate buy or sell decisions. Those decisions are made solely by the portfolio generation process. The trading algorithms, however, determine the trade horizon—how long it will take to buy or sell those positions. The average realized trade horizon for completion is less than one day and no more than three days.

While the quantity traded is not altered by the trading algorithm, and hence not endogenous to expected trading costs, there is one endogenous choice the manager makes with regard to stock selection in relation to trading costs, which is to only trade stocks of a certain size and ex ante liquidity (roughly the Russell 3000 or higher). Thus, we have no live trading data for microcap stocks, for instance.

The trading algorithm directly and anonymously accesses market liquidity through electronic exchanges and, in order to minimize market impact, is designed to provide rather than demand liquidity by using a system of limit orders (with prices generally set to buy at the bid or below and sell at the ask or above) that dynamically break up total orders (parent orders) into smaller orders (child orders), with the sizes of child orders and the time in which they are sent being randomized.[8] We include all of these features as part of the trading costs facing our manager. That is, the costs our manager faces are a function of these execution choices, which are designed to lower the costs of trading a given fixed quantity of shares. These costs would not represent the cost of trading those shares immediately/instantaneously from a market order, for instance. However, we believe these features are very similar to the way most large institutions trade (or could trade), and therefore the costs we estimate are not representative of the average investor in the market, but rather represent those facing an arbitrageur engaged in these strategies and hence much closer to the marginal investor we have in mind theoretically.

We treat our trading costs as (mostly) exogenous to the portfolio weights and the set of securities being traded. There are several reasons to justify this assumption. First, as described above, the portfolio determination process is separate from the trading process. Second, we throw out higher frequency trades where this is not true. Third, many of the client mandates force the portfolio to trade to a benchmark with low tracking error, leaving little scope for deviation. Fourth, the trading algorithm controls trade execution, not stock selection, and since most trades are completed within a day, the intended trades for the long-term strategies we study have very little tracking error. Finally, to rule out any remaining concerns over the endogeneity of portfolio weights to transactions costs, we also estimate our trading costs using only the first trade from new inflows from long-only mandates that specifically adhere to a benchmark. The initial trades to a benchmark from new

[8] There are additional features the trading algorithm employs to handle corporate announcements, special events, etc. that are beyond the scope of this paper. Results are robust to excluding such events from the data.

inflows are exogenous to trading costs since there is no scope for deviation. In the next section we show that the trading costs from these exogenous, initial trades from new inflows are identical to those from all other trades.

C. Equity anomaly/factor portfolios

We apply our trading cost data to the preeminent anomaly/factor portfolios studied in the literature, which also happen to comprise (not coincidentally) a significant fraction of the models and trades made by our manager. Specifically, we examine long-short equity portfolios pertaining to size, value, and momentum, as well as combinations of these portfolios. We focus on these equity styles because research shows they capture much of the cross-sectional variation in returns (Fama and French [1996, 2008] and Asness, Moskowitz and Pedersen [2013]), and are also the focus of attention in the investment management industry.

We examine equity strategies in the U.S. and developed international markets. We obtain stock returns and accounting data from the union of the CRSP tapes and the XpressFeed/Compustat Global database. Our U.S. equity data include all available common stocks on the merged CRSP/Compustat data between July 1926 and September 2013. Our global equity data include all available common stocks on the XpressFeed/Compustat Global database for 21 developed markets. The international data run from January 1986 to September 2013. We assign individual stocks to the corresponding market based on the location of the primary exchange. For international companies with securities traded in multiple markets, we use the primary trading vehicle identified by XpressFeed/Compustat. Table A1 in the Appendix presents the summary statistics for the global data, including number of stocks and average market cap for stocks in each country.

Our portfolio construction closely follows Fama and French (1996 and 2012) and Asness and Frazzini (2013). Our global portfolios are country neutral: we form long-short portfolios within each country and then compute a global factor by weighting each country's long-short portfolio by the country's total (lagged) market capitalization. The size and value factors are constructed using six value-weighted portfolios formed on size and book-to-market.[9] At the end of each month stocks are

[9] To obtain shareholders' equity we use Stockholders' Equity (SEQ), but if not available, we use the sum of Common Equity (CEQ) and Preferred Stock (PSTK). If both SEQ and CEQ are unavailable, we will proxy shareholders' equity by Total Assets (TA) minus the sum of Total Liabilities (LT) and Minority Interest (MIB). To obtain book equity, we subtract from shareholders' equity the preferred stock value (PSTKRV, PSTKL, or PSTK depending on availability). Finally, to compute book value per share (B) we divide by common shares outstanding (CSHPRI). If CSHPRI is missing, we compute company-level total shares outstanding by summing issue-level shares (CSHOI) at fiscal yearend for securities with an earnings participation flag in the security pricing file. Following Fama and French (1992) we assume that accounting variables are known with a minimum six-month lag

assigned to two size-sorted portfolios based on their market capitalization. For the U.S., the size breakpoint is the median NYSE market equity. For the international sample the size breakpoint is the 80th percentile by country (to roughly match the U.S. size portfolios). We use conditional sorts that first sort on size, then on book-to-price.

The size factor, SMB (Small minus Big), is the average return on the three small portfolios minus the average return on the three big portfolios = 1/3 (Small Value + Small Neutral + Small Growth) - 1/3 (Big Value + Big Neutral + Big Growth). SMB portfolios are refreshed in June of each year and rebalanced every calendar month to maintain value weights within each of the six portfolios and equal weights across the three small and three big portfolios.

The value factor, HML (High minus Low), is the average return of the two value portfolios minus the average return of the two growth portfolios = 1/2 (Small Value + Big Value) - 1/2 (Small Growth + Big Growth). HML portfolios are refreshed every calendar month using the most recent price following the methodology of Asness and Frazzini (2013), and rebalanced every calendar month to maintain value weights within and equal weights across the portfolios.

The momentum portfolios UMD (Up minus Down) are constructed in a similar manner: we use six value-weighted portfolios formed on size and prior returns (the cumulative $US return from months *t-12* to *t-2*). The portfolios are the intersections of two portfolios formed on size and three portfolios formed on prior returns. The momentum factor UMD is constructed as UMD = 1/2 (Small High + Big High) - 1/2(Small Low + Big Low). UMD portfolios are refreshed and rebalanced every calendar month. All returns are in $US and excess returns are relative to the one-month U.S. Treasury bill rate. We include delisting returns when available in CRSP.

Table A2 in the Appendix reports the fraction of firms covered and the fraction of total market cap covered by our trading database for all stocks on CRSP and Xpressfeed/Compustat Global. Panel A of Table A2 reports the year-by-year coverage, which reaches as high as 60 percent of names and 97 percent of market cap. Panel B reports the coverage by region and year. Finally, Panel C reports the fraction of the anomaly/factor portfolios our data covers that were actually traded (fraction of absolute value of portfolio weights) by our manager. For the U.S., the trade data covers about 62% of the portfolio weights of the factors we study – SMB, HML, and UMD. Internationally, we cover about 34%. Recall, that through its trade execution – patient trading, limit orders, not

---

and align book value of the firm at the end of the firm's fiscal year, ending anywhere in calendar year *t-1* to June of calendar year *t*. In order to be included in any of our tests we require a firm to have a non-negative book value and non-missing price at of the prior month.

trading tiny, illiquid stocks, and news events – our manager creates tracking error with respect to the academic portfolios. However, as Panel C of Table A2 also reports, the correlation of returns of the factor portfolios we generate from the trade data with respect to those from Ken French's website are very high. For HML and UMD, the return correlations are 0.96 and 0.97 in the U.S. and 0.91 and 0.94 internationally. Hence, despite the deviations the trading algorithm imposes on the factor portfolios to lower trading costs, their returns are highly correlated. For SMB, the correlations are lower (0.78 in the U.S. and 0.53 internationally) which makes sense since our manager screens out anything below the Russell 3000, which the Fama-French SMB overweights.

Finally, the academic portfolios standard to the literature are not designed in any way to address transactions costs, and as such, may not be proper benchmarks to evaluate the after-cost efficacy or capacity of an anomaly, which may depend critically on the design of the portfolios, including simple steps taken to limit trading costs. Therefore, we also consider in Section V a set of portfolios specifically designed to optimize for trading costs ex ante to maximize after-cost returns.

II. Realized Actual Trading Costs

We first describe our methodology for calculating trading costs. Then, using the live trading data we present summary statistics of trading costs over time, for different types of stocks, and across markets for various types of trades.

A. Measuring transactions costs: price impact and implementation shortfall

To measure transactions costs, we use an *implementation shortfall* (*IS*) methodology as defined in Perold (1988). Implementation shortfall measures the difference between a theoretical or benchmark price (e.g., a model price) and an actual traded price, scaled by the amount traded. We also define *market impact* (*MI*) in a similar way by looking at the difference between an arrival price (i.e., the price that exists when a trade begins in the market) and an actual traded price, scaled by the amount traded. The difference between these two measures represents pre-trade moves that might occur from the time a model portfolio is generated and the time trading begins. The following equations illustrate the two costs:

$$\begin{aligned} IS &= ret_{p,theory} - ret_{p,actual} \\ &= ret_{p,theory} - \text{cost}_{execution} - \text{cost}_{opportunity} \end{aligned} \tag{1}$$

where the opportunity cost is the difference in returns before trading costs between the theoretical and actual portfolio. The cost of execution is measured as

$$\text{cost}_{execution} = Q_{+}^{ex}(P^{ex} - P^{theory}) + Q_{-}^{ex}(-P^{ex} + P^{theory}) \tag{2}$$

with Q_{+}^{ex} and Q_{-}^{ex} representing the quantity of shares bought and sold, respectively. Equation (2) takes into account bid-ask spread, market impact, and commissions. The largest part of the cost of execution for a large institutional trader such as our manager is market impact costs:

$$MI = Q_{+}^{ex}(P^{ex} - P^{start}) + Q_{-}^{ex}(-P^{ex} + P^{start}). \tag{3}$$

Other costs of execution such as commissions or effective bid-ask spreads are small by comparison, especially at large trade sizes since they do not scale up with size. We show below that for the size our manager traded, these costs are only a few basis points and will be inconsequential for determining break-even fund capacities.

Both *IS* and *MI* are effective measures of the costs of trading. Implementation shortfall measures the total amount of slippage a strategy might experience from its theoretical returns. For the strategies we examine, the theoretical or benchmark price is defined as the closing price at the time the strategy's desired holdings and trades are generated (e.g., prior day's closing price). Market impact is similarly defined as the opening price on the day the strategy begins trading to the new desired holdings (e.g., current day's opening price).[10] Our unique data allow us to calculate implementation shortfall since we have the underlying theoretical model prices our institutional manager was attempting to trade to, which is absent from previous studies since data on the model generating the portfolios is typically unavailable to the researcher.

However, in expectation, *IS* and *MI* should be the same, since pre-trade moves for the strategies we examine (or the difference between the two) should be random with expected value of zero. In other words, the strategies we examine based on long-term return premia related to size, value, and momentum are not forming opinions about the overnight moves in the stocks they desire to hold such that there would be an expected difference in the two measures. Empirically, we estimate *IS* and *MI* separately and find little difference.

There are many ways to think about transactions costs and many different objectives when doing so. Our notion of trading costs computes the difference between the results of a theoretical portfolio which has zero transactions costs and the results of a practical portfolio which attempts to track the theoretical portfolio but is subject to actual traded prices. Effectively, our cost computation captures how much of the theoretical returns to a strategy are actually achieved in practice. This measure seems to be the most relevant for addressing market efficiency questions like: does an anomaly

[10] Our manager does not engage in any significant overnight trading.

survive trading costs? Or, at what equilibrium size would an anomaly's returns be completely eaten up by transactions costs? Both are questions we aim to answer in this paper.

There are other ways to measure transactions costs, such as to compare actual traded prices to other possible traded prices that existed during the trading period. For example, one could compare the actual trade price to the average trade price achieved in the market, such as the volume weighted average price (VWAP) or the price that would have been achieved using market orders over the same period. Defining market impact costs relative to the VWAP,

$$MI^{relative} = Q_{+}^{ex}(P^{ex} - VWAP) + Q_{-}^{ex}(-P^{ex} + VWAP). \tag{4}$$

This can measure the effectiveness of a trade relative to other traders in the market at the same time. However, this measure does not help us understand whether the theoretical returns of a strategy are achievable or what the breakeven capacity is to being profitable. For comparison, we show both market impact measures, recognizing that the relative market impact measure in equation (4) will generally be a lot lower than our estimates of trading costs from the theoretical portfolio in equation (3), where the latter is most relevant for thinking about capacity or the returns available to an investor implementing such a strategy.

B. Trading costs

Table II reports the mean, median, and value-weighted mean (weighted by dollar value of trades) of the market impact (*MI*) and implementation shortfall (*IS*) estimates we get from the live trading data, following equations (1) through (3). Specifically, the cross-sectional mean, median, and weighted mean are computed each month and the time-series average of these monthly measures are reported, where each month is weighted by the number of stocks traded in that month. Standard errors of the monthly estimates are reported in the style of Fama and MacBeth (1973) in the bottom half of each panel. Panel A reports results for the full sample period from 1998 to 2013. The first column reports the summary statistics for all trades, where the mean market impact measure is 11.25 basis points and the mean implementation shortfall measure is 11.89 basis points. The 0.64 basis point difference between the two represents the difference between the intended model and the actual portfolio at the time of trading, which, as expected, is negligible. The median *MI* and *IS* costs are quite a bit lower at 7.36 and 9.42 basis points, respectively, suggesting that trading costs are positively skewed by a few much more expensive trades. Weighting trades by their dollar value, the value-weighted means are quite a bit higher at 17.06 basis points for market impact and 18.34 basis points for implementation shortfall, which indicates that the largest trades are the most expensive trades, consistent with trading cost models that argue that trading costs increase substantially with trade size – a fact we confirm in the data as well. The standard errors for these estimates are small and typically close to one basis point, so the difference between the *MI* and *IS* measures are usually within two standard errors of zero.

Event-time market impact. Figure 1 plots the average price movement in the 24 hours following a trade for the average trade from our live trading data. We restrict the sample to U.S. stock trades on the NYSE/AMEX with a trading horizon of less than one trading day and compute average price impact during trading and the average overnight and intraday return over the subsequent 24 hours. We estimate the average price impact during the day and split it equally across 30 minute trading windows. For overnight returns and for the next day's returns, we split the sample into two-hour windows for simplicity. The plot in Figure 1 shows the evolution of price impact for the average trade in our data. Since trades are broken up into smaller child orders, we estimate market impact as the integral of the function plotted in Figure 1 over the trade horizon – from the beginning of trading at the open to the last child order filled for that trade (when the trade horizon exceeds one day we compute market impact over multiple days). From Figure 1, our estimate of market impact is just under 9 basis points on average for all trades completed within a day, which matches the number in Table II (8.81 bps).

In measuring the price impact of a trade, as defined in equation (3), it is interesting to assess how much of the impact is permanent versus temporary. The permanent component, while still a cost relative to the theoretical price, is driven by outside forces such as market demand, while the temporary component should reflect the trader's demand for immediacy, which once alleviated should result in price reversion. As Figure 1 shows, most (around 85%) of the price impact in our trade execution data appears to be permanent. Of the average 9 basis point market impact, only about 1.5 basis points are reversed over the next 24 hours. While some may consider the 1.5 basis points as the true price impact of our manager coming from immediacy, for questions regarding the efficiency and after-cost returns to a strategy, as well as break-even fund size calculations, we use the entire price impact (permanent + temporary), since it is the total cost the strategy incurs from implementation, which is relevant for these issues. In other words, the investor pays both costs – temporary and permanent – when investing in these strategies.

By exchange. The next three columns of Panel A of Table II report the same statistics for NYSE-AMEX and NASDAQ traded stocks as well as for all stocks traded on international exchanges separately. Whether measured by market impact or implementation shortfall, the average costs of trading on the NYSE appear smaller than on NASDAQ or in international markets, which makes sense since the NYSE contains larger and more liquid firms.[11]

By size. The fifth and sixth columns of Panel A of Table II report results for large and small cap stocks separately. The average large cap stock trade generates 10.25 basis points of market impact costs compared to 21 basis points for small cap stocks. Of course, this does not control for the size of the trade, which we will address in the next section.

By portfolio type. The last two columns of Panel A of Table II report trading costs for trades made within the context of long-short portfolios and long-only portfolios separately. The trading costs for long only portfolios are larger than those made in long-short portfolios. The average long-only trade faces nearly 15.5 basis points of market impact costs compared to only 9.3 basis points for the average long-short trade. However, on a value-weighted basis, long-short portfolios face about the same or slightly larger costs, since long-short portfolios trade more extreme positions.

Panel B of Table II repeats the estimates over the more recent sample period from 2003 to 2013. As shown in Table I, trading activity increased substantially starting in 2003, which also coincides (but not coincidentally) with the adoption of an electronic trading algorithm. Since the time-series averages we report in Table II weight each month by the number of stocks traded, which overweights the more recent periods, the results from 2003 to 2013 look very similar to those from the full sample from 1998 to 2013.

By trade type. Panel C of Table II reports trading costs for different types of trades: buy long, buy to cover, sell long, and sell short. The first two columns report the percentage of trades, in both dollar terms and numbers, of each trade type. Buying long and selling long account for two-thirds of all trades, with the remaining one third split evenly between short selling and covered buys. Over the full sample, buying long generates about 12.5 basis points of price impact, but buying to cover has 16.2 basis points of price impact, though the difference between them is statistically insignificant. Short selling is slightly more expensive by 2.7 basis points on average than selling long, but the difference is also not statistically significant.[12] Although a large literature discusses the additional costs associated with short-selling, conditional on actually shorting, we see no marked difference in

[11] Figure A1 in the Appendix plots the mean and median market impact costs separately for each of the 21 equity markets we examine: Australia, Austria, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, U.K., Hong Kong, Israel, Italy, Japan, Netherlands, Norway, Portugal, Singapore, Sweden, and the U.S. Singapore, Hong Kong, and Japan have some of the highest price impact costs, along with Australia. The U.S. and most of the European countries have the lowest price impact costs.

[12] Shorting costs and shorting revenues from lending fees are not included in these costs, which just capture price impact. The majority (around 99%) of short positions are in stocks that are not hard to borrow and hence the general collateral rate applies, which is small on average.

trading costs between selling a long position versus selling short. If short selling is indeed costlier, it is likely to be from opportunity cost (i.e., not being able to short). The execution costs of short selling versus selling long appear no different, even with uptick rules and other barriers to shorting. Results for international exchanges are similar. Figure A2 summarizes the results from Table II for the price impact measure.

Table A3 in the Appendix also reports pooled means of the trading cost measures, rather than Fama and MacBeth averages, over the full sample period and the post-2003 sample period. The estimates of trading costs are very similar.

Alternative cost measures. For comparison, Table A4 in the Appendix reports market impact costs relative to VWAP following equation (4), for our sample of U.S. trades. As Table A4 indicates, the trading costs when measured versus the VWAP are much lower than those in Table II that measured costs relative to the theoretical price. Instead of the 11.25 basis point average price impact in Table II, the average price impact measured relative to the VWAP is only 2.68 basis points, indicating that our manager paid 2.7 basis points more on average relative to the VWAP during the same trading day. This difference is roughly consistent with the temporary trading cost from Figure 1, where the demand for immediacy should show up as both a cost relative to the average price paid by traders as well as be reversed the following trading day.

However, for the same reasons we use the total cost (permanent + temporary price impact), we believe the theoretical price impact numbers are the most relevant for answering the questions we pose in this paper, which is what are the implementation costs and break even sizes of strategies designed to exploit the asset pricing anomalies we investigate? Our conclusions using the theoretical price impact numbers will therefore be conservative relative to those using alternative price impact measures from average traded prices like the VWAP, but represent the total cost facing an investor in these strategies.

Novy-Marx and Velikov (2015) examine the trading costs of a number of anomaly-based portfolios, including size, value, and momentum, using the effective bid-ask spread of Hasbrouck (2009). The effective spread, however, is a very poor estimate of the costs facing a large arbitrageur. For one, it does not take into account price impact from the size of the trade, which is the biggest cost facing a large trader. Second, as a result, it cannot be used to assess the capacity or break-even size of an anomaly. Finally, it also assumes market orders and immediate liquidity demand. As Novy-Marx and Velikov (2015) state, their costs "should thus be interpreted as the costs faced by a small liquidity demander," and therefore are not a good estimate of the costs facing a large arbitrageur. As a case in point, the effective spread of our manager is 1.53 basis points in U.S. stocks (essentially the manager's temporary price impact), which is orders of magnitude smaller than the costs Novy-Marx and Velikov (2015) estimate, even for the largest stocks in their sample. This is because a large institutional trader like our manager often trades within the spread, using limit orders and tries to supply rather than demand liquidity. Consequently, the trading costs estimated in Novy-Marx and Velikov (2015) do not represent those of a large arbitrageur and are, therefore, ill-equipped to answer market efficiency questions regarding these anomalies.

C. Exogenous trades from inflows

We claim and argue above that the trading costs we estimate are predominantly exogenous to the portfolio weights and set of securities being traded. To show some direct evidence consistent with that claim and to rule out any remaining concerns over the endogeneity of portfolio weights to transactions costs, we re-estimate our trading costs using only the first trade from new inflows coming from long-only mandates that adhere to a benchmark with a tight tracking error constraint. The initial trades from new inflows provide little scope for deviation or optimization and hence offer an exogenous estimate of trading costs.

Table III reports the mean, median, and value-weighted mean market impact measures for long-only trades coming from inflows only and for all other trades for comparison. The estimates across the two samples are very similar and show no systematic differences in market impact across the two types of trades. The last column performs a formal statistical test on the differences in trading costs that cannot reject that the two measures are identical. Cutting the trades by small and large cap stocks yields the same conclusions. Figure A3 summarizes the results, showing no difference in market impact costs from the exogenous initial inflow trades relative to all other trades in our sample. The evidence suggests that our trading cost estimates are fairly independent and exogenous to the portfolios being traded, and hence can be applied more broadly to other portfolios or settings. Put differently, if we were to use the trading costs estimated only from new inflows to tightly benchmarked portfolios, our results would be the same.

The exogeneity of our trading cost estimates allows us to extrapolate our cost estimates (under some assumptions) to other trade sizes, other trading strategies, and perhaps other stocks and time periods out of sample to get a reasonable measure of trading costs.

III. The Cross-Section of After-Trading Cost Returns

We examine the cross-section of after trading cost returns using some of the most common academic portfolios based on several well-known asset pricing anomalies. We begin with a very simple exercise. Using our database of live trades, we reconstruct the academic portfolios from those trades and apply the actual trading costs incurred by our manager from those trades to the portfolios. This exercise requires no estimate of trading costs. Unlike other papers in the literature that estimate trading costs from TAQ data, bid and ask quotes, or proprietary broker data, we simply use the actual costs paid by our manager on the actual trades and quantities traded by our manager in real time to build the anomaly/factor portfolios. As such, we examine the actual realized costs of these trades on the portfolios, requiring no modeling or estimation.

The advantages of using actual trading costs are clear – they provide a precise measure of real trading costs. The disadvantage is that we only have these costs for the sample period our manager traded and for the stocks it traded. Fortunately, as Table A4 shows, our manager traded 23,593 stocks in the U.S. and 28,756 stocks internationally over a fairly long time period from 1998 to 2013. Moreover, our manager traded portfolios very similar to and highly correlated with the academic portfolios used in the literature.

A. U.S. trade execution sample

We start by computing the actual trading costs and net returns to our portfolio strategies in the U.S. over the 1998 to 2013 period. We examine the strategies SMB, HML, and UMD, plus an equal-weighted combination (Combo) of all three strategies to assess the diversification benefits of combining value with momentum (and size) as shown in Asness, Moskowitz, and Pedersen (2013), Asness and Frazzini (2013), and Daniel and Moskowitz (2015) .

Specifically, we reconstruct SMB, HML, and UMD by first restricting the universe of stocks to those that appear in our trade execution database at portfolio formation, and then building the long/short portfolios following Ken French's data library. For each stock we estimate its market impact at each rebalance using the realized market impact of all trades executed in that stock in the prior six months. For example, if the UMD portfolio required buying Microsoft on June 30, 2003, we assume that the market impact of that purchase is equal to the average market impact of Microsoft trades between January 1, 2003, and June 30, 2003. We use six month windows of trading activity to maximize stock coverage and to reduce noise. This assumes that the trade size is equal to the average trade size of our manager over the past six months. In essence, we replicate SMB, HML, and UMD using only those stocks that were traded by the manager in the most recent six months at the actual sizes the manager traded and use the actual prices and actual trading costs borne for those trades.[13]

Since our manager was trading similar types of strategies during this time, many of the stocks in these portfolios are covered by the execution data. The portfolios from our data on average cover about 63% of the portfolio weights of these strategies. More importantly, the correlation of returns of our portfolios with those of Fama and French are extremely high (0.96 for HML, 0.97 for UMD, and 0.78 for SMB), indicating that the portfolios from our live trading database replicate the academic factor/anomaly portfolios well.

Panel A of Table IV reports results for the U.S. equity strategies over the sample of live trading data. The first row reports the average dollar volume of trades used to estimate trading costs at each rebalance. For example, when rebalancing SMB, our trading cost estimate for that month is based on an average of 9.69 billion dollars of trades. The second row reports the "implied fund size" from these trades. For example, given turnover (which we report near the bottom of the table as the sum of long and short trades per dollar) of SMB is 53% per month, our estimates are based on a fund size of 9.69/0.53 = 18.18 billion dollars. In other words, we are carving out a 18.2 billion dollar SMB portfolio from our live trading data that would generate the same dollar volume of trades in SMB stocks that our manager experienced. We do the same for HML and UMD. Effectively, our manager ran an 18.2 billion dollar SMB portfolio, a $9.42 billion HML portfolio, and a $5.21 billion UMD portfolio in U.S. equities.

The fourth row reports the realized trading costs of each strategy at the fund sizes above. For SMB, the realized annual trading cost is 1.47 percent. That cost comes from the turnover of the strategy (53% per month) times the average market impact per trade in the strategy (22.94 bps), both of which are reported in the last two rows of Table IV. HML has a slightly lower realized cost of 1.35 percent per year despite having higher turnover of 63%, because it has lower average market impact costs of 17.7 basis points. UMD has an annual trading cost drag of 3.03 percent per year due to its higher 119% turnover and an average market impact cost of 21.3 basis points. The combination of all three strategies has a realized cost of 1.46 percent per year with turnover of only 57% since

[13] This calculation produces actual trading costs on a stock-by-stock basis, but assumes that trading costs are the same no matter what portfolios the trades belong to. Given the independence between portfolios and trading costs we showed earlier, this assumption seems valid. If market makers knew the reason behind each trade and adjusted prices depending on the type of portfolio being traded (e.g., SMB versus UMD), then this assumption would be violated. However, given the anonymity of trading through the trading algorithm and the virtual impossibility of knowing which portfolio each trade belongs to, it seems quite reasonable to assume that these costs are independent of the portfolio.

value and momentum trades tend to offset each other, resulting in lower turnover which has real transactions costs benefits, with an average price impact of 21.2 basis points.

To compare the realized trading costs to the profitability of these strategies, the fifth row of Panel A of Table IV reports the break-even cost of each strategy, which we estimate using historical data over the largest sample of returns available. For example, for SMB we compute the average return in U.S. data from July 1926 to September 2013, which is 2.95 percent. We use the longest sample of returns available because estimating mean returns is notoriously difficult and the short sample from 1998 to 2013 may provide a distorted view of the true expected return to these strategies. For instance, the actual gross average return over the 1998 to 2013 sample period for SMB is 7.98 percent, which is about three times larger than its historical average. For completeness, the remaining rows report the average gross and net returns of the strategies (and their associated *t*-statistics) for both the full historical sample mean and the shorter trading sample mean. Given the 1.47 percent trading cost for SMB, this implies that a $18.2 billion SMB portfolio would deliver in expectation 1.48 percent per year net of trading costs return using the historical mean or 6.52 percent per year net return using the 1998 to 2013 mean. The former appears to be a much more reasonable and realistic expectation of the mean net return to SMB at this size. Hence, we focus on the longer and more reliable mean return estimate, which is also the estimate used by other studies (e.g., Korajczyk and Sadka (2004)) with which we compare.

For HML, using the historical average return from 1926 to 2013 of 4.95 percent, its net return (trading on average $6 billion) is 3.61 percent per year (*t*-statistic of 2.25). For UMD (trading $6.2 billion), its net expected return is 5.17 percent per year (*t*-statistic of 3.02). The combination portfolio generates a 3.93 percent net expected return (*t*-statistic of 6.66) at a fund size of nearly $17 billion. The net returns are similar even over the shorter 1998 to 2013 trading sample period, with a mean of 3.58 (*t*-statistic = 2.23). Hence, although the short sample period can lead to factor realizations quite different from their long-run means – being much greater for SMB and much weaker for UMD over this sample period – the combination portfolio diversifies these idiosyncratic realizations so that the shorter sample delivers similar net returns to the long-run historical mean estimate.

The significant positive net of trading cost returns to these strategies seem at odds with the conclusions reached in the literature. For instance, Lesmond, Schill, and Zhou (2003) and Korajczyk and Sadka (2004) find that the net returns to momentum strategies are zero at much lower fund sizes than our manager traded. Korajczyk and Sadka (2004) find that at $2-3 billion, momentum profits would be wiped out entirely by trading costs, whereas we show that our manager ran more than $6 billion in an identical momentum strategy that easily survives trading costs. As a point of comparison, Korajczyk and Sadka (2004) estimate that at a $2 billion fund size (the break-even size for NYSE stocks in a momentum fund) annual trading costs to a momentum strategy would be 57 bps per month or 6.8 percent per year, and at $3 billion (the break-even size for a momentum fund run across all NYSE/AMEX/Nasdaq stocks) the trading costs would be 9.6 percent per year. As Table IV shows, our manager realized actual trading costs of only 3.03 percent per year on the same momentum fund run on more than $5 billion in size – twice as large as what Korajczyk and Sadka (2004) claim is the break-even capacity.

Novy-Marx and Velikov (2015) also examine the net returns to momentum using the scale invariant effective spread of Hasbrouck (2009). They estimate an annual trading cost of 7.8 percent at any fund size, which is again much larger than our manager's actual cost, most of which is price impact. For the same momentum strategy studied by Novy-Marx and Velikov (2015), the effective spread paid by our manager was only 12.7 basis points, which is orders of magnitude lower. This evidence is consistent with the effective spread being a proxy for small liquidity demanders and a very bad proxy for the costs facing a large arbitrageur.

B. International trade execution sample

Panel B of Table IV reports the results for the international portfolios over the trading sample period. The results are remarkably similar, where the trading costs are very close to those for U.S. equities and the implied fund sizes are also similar (and often slightly higher) internationally.

The net returns in excess of trading costs are also positive and significant internationally, except for SMB, where historically the gross return to size is non-existent internationally.[14] However, SMB still exhibits the same trading costs as its U.S. counterpart. The fact that trading costs and net returns are similar in 20 other developed equity markets provides additional robustness for the trading costs associated with these strategies.

While the strategies we examine – size, value, and momentum – easily survive trading costs at the fund sizes our manager traded, not every strategy necessarily would. For example, Table A5 in the appendix calculates the trading costs of the one-month reversal strategy of Jegadeesh (1990) using the short-term reversal (STR) portfolio from Ken French's website. As the table shows, our manager effectively traded $5.16 ($18.24) billion in STR in the U.S. (internationally) for an implied fund size of $1.72 ($6.17) billion. At this size, the STR portfolio incurred transactions costs that

[14] See Asness, Frazzini, Israel, Moskowitz, and Pedersen (2015) on the size effect internationally. Alternatively, we could use the U.S. historical estimates for the expected returns to these strategies since they cover a much longer period and there is no a priori reason why the returns should be different internationally.

eliminated its gross positive returns. Using the full sample mean of STR of 5.65 percent (the largest mean estimate) the strategy realized costs of 6.73 percent, exceeding even the largest gross return estimate we could find. The costs on the international STR strategy are similar. Hence, for short-term reversals, which require a great deal more trading (turnover equals about 300% per month), the realized trading costs of our manager rendered the strategy unprofitable at the sizes our manager traded. Thus, not every strategy would survive these trading cost estimates.

IV. Estimated Trading Costs and Break-Even Capacity

In this section, we use our live trading cost database to estimate a model of trading costs that we then apply to different trade sizes to calculate break-even capacities across the anomaly portfolios. We can also use the model to estimate trading costs out of sample to different stocks and time periods by imposing additional structure. These estimates are compared to other models in the literature and will be used for trading cost optimization in Section V.

A. Trading cost model: What determines trading costs?

To build a trading cost model, we first try to understand what affects trading costs. We investigate what observable market, stock, and trade characteristics are related to the realized trading costs in our sample. We then estimate an econometric model for the price impact function based on these observable variables that we can then apply to various portfolio sizes and out of sample.

There are a number of different ways to model transactions costs. Proportional trading costs models, where the cost of trading does not vary with the size of the portfolio traded, such as the effective spread, do not match theory or data (or practice). Non-proportional trading cost models, where the cost of trading does vary with the size of the portfolio traded, are more realistic, but the shape of the price impact function is an important ingredient. Often these are modeled as concave or linear functions (Hasbrouck (1991), Hausman, Lo, and MacKinlay (1992), Keim and Madhavan (1996), Almgren (2001), and Almgren, Thum, Hauptmann, and Li (2005)).

Figure 2 plots the average and median price impact measures from our live trading data as a function of trade size, measured by fraction of daily trading volume traded. The plot shows a concave relation between market impact and trade size, with costs growing significantly at very large sizes (as a percentage of DTV). Fitting a curve to our sample of trade data (superimposed on the graph) statistically supports this function of price impact. Formal statistical tests clearly reject a linear model of price impact in favor of a non-linear function.

To capture this functional form, we use a square root process to model the relationship between price impact and trade size (fraction of DTV traded). Theory (Barra (1997), Kahn (1993), Grinold and Kahn (1999), and Loeb (1983)) argues for a square root function and Almgren (2001) and Almgren, Thum, Hauptmann, and Li (2005) find that a square root function does a good job describing their price impact data. We, too, find that a square root function describes our trade data well. Figure A5 in the appendix shows a log-log plot of market impact on trade size and finds that the coefficient is 0.31, which is close to the square root coefficient. Given theory and attempting to avoid overfitting, we use a square root function rather than a power law of 0.31 to model trading costs.

Table V presents our econometric model for trading costs. The explanatory variables used in the model are characterized into three groups. The first set of variables captures how trading costs vary with the size of a trade. These variables include the fraction of daily trading volume traded, which is the dollar size of the trade divided by that stock's average daily trading volume over the past six months, and the square root of the fraction of daily volume traded. The second group of variables include firm characteristics that may be related to the cost of trading, such as the size of the firm, which is the log of 1 plus the market value of equity [log(1+ME)], and the idiosyncratic volatility of firm's equity return, which is the standard deviation of the residuals from a regression of one-year daily stock returns of the firm on the corresponding value-weighted market index for that country (expressed as an annualized percentage). Finally, we include a set of market variables to capture variation in trading costs over time that result from different market conditions. We include the "VIX" to capture market volatility, which is the monthly variance of the country's value-weighted index computed using daily returns (expressed as an annualized percentage) and a linear time trend to capture the changes in aggregate trading costs over time. Table A6 in the appendix looks at other timing variables, including dummies for events like decimalization, and finds similar results. We also include the variable "Beta*IndexRet*buysell", which is the beta of the stock times the index return for that market times a variable "buysell" equal to 1 for buys and -1 for sells. "Beta" is the stock's predicted beta at time of order submission and "IndexRet" is the market index return over the life of the trade. This variable is used to account for contemporaneous (beta-adjusted) market returns over the trade period and to sign the trades correctly, since our data contains both buys and sells. This variable is only used to parameterize our model when looking at historical transactions. It plays no role in estimating future transactions costs or when applying the model out of sample, since we assume that daily market returns are unforecastable (i.e., E[IndexRet] = 0). Thus, when estimating transaction costs out of sample, we simply drop this term while holding the other parameters fixed.

The model is estimated separately for the full sample, U.S. trades only, and international trades only, using pooled regressions with country fixed effects. The dependent variable is the realized market impact from the live trading data, in basis points. The results in Table V are economically intuitive. Larger trades lead to higher transactions costs, where the square root term comes in significantly. Smaller and more volatile firms have higher transactions costs. An increase in market volatility is associated with higher transactions costs, which is consistent with the notion that market makers need to be compensated more in more volatile markets. These variables are statistically significant at the 5% level. The time trend is negative, as trading costs have declined. Table A6 shows that a big part of this decline is due to decimalization.

The R-squares of the regressions are about 7%, which is quite large for these types of models given the noisiness of trade data. Recall, the dependent variable, market impact, is essentially a daily return. So, a 7% R-squared is fairly impressive given the movement in daily returns.

As an additional test, we also add a variable which is the return on the stock being traded minus the return on a similar stock – in terms of size, BE/ME, and momentum/past 12-month return – not traded by our manager at that time (times the “buysell” indicator to get the right sign). The matched stock is based on the procedure of Daniel, Grinblatt, Titman, and Wermers (1997) that characteristically matches stocks along these three attributes. The idea here is to see if a stock being traded by our manager on a given day experiences different price movements than a similar stock (based on size, value, and momentum) that our manager did not trade that day. As Table V shows, on average our manager experiences an additional 4 bps of market impact on the stocks it trades relative to stocks of similar characteristics that it did not trade that day. This number may represent our manager’s demand for immediacy and the temporary price impact it incurs. The estimated coefficient is in line with our previous estimates of our manager’s temporary price impact. The coefficient estimates are all very similar across markets, with U.S. and international data giving similar parameterizations. The fact that the patterns and estimates are similar across these markets is reassuring and indicates our model is robust and can be applied reasonably well across markets.

We estimate our trading cost model over the entire sample period from 1998 to 2013 and account for trading costs declining through the sample period (from electronic trading, more venues, increased volumes, etc.) using the time trend and VIX. As a test of how reasonable the model is, we examine time variation in the market impact function by estimating it over three subperiods: 2004-2007, 2008-2009, and 2010-2013. Essentially, we estimate our model before, during, and after the financial crisis. We use Fama and MacBeth (1973) regressions and take an average of the year-by-year coefficients. Figure A6 plots the functions over the three subperiods. Costs appear to decline from 2004-2007 to 2010-2013, but jump up during the financial crisis in 2008-2009. The patterns in international markets mirror those of the U.S. This is consistent with anecdotal evidence that price impact, in addition to spreads, was greater during and immediately following the global financial crisis. We use the full sample coefficient estimates of our model to estimate trading costs for improved precision, which is a weighted average of the functions plotted in Figure A6, but our results are similar even if we allow all coefficients to vary through time using the three subperiod functions in Figure A6.

Finally, our price impact function is, of course, limited by the range of trading size our data provides. Hence, we cannot say what price impact might look like beyond these sizes. However, the range of trading volumes (as shown in Figure 2) used in our transactions cost model are mostly less than 5% of a stock's daily volume up to the max of 13.1%. This is a reasonable range of trading volumes for the types of strategies we investigate in this paper and for the type of trader, a large arbitrageur, we wish to model. Trades done at sizes well beyond these levels tend to reflect idiosyncratic trades, perhaps around an event, and are often informational trades. Thus, for most of our analysis, we will focus on using our model for trade sizes within this range. For sizes outside of this range we can extrapolate the model using the functional form, which is fine unless that function changes significantly at much higher %DTV than our manager traded. For trade sizes close to but outside of this range, the assumption that the functional form holds is likely reasonable. For sizes way outside of the range, we interpret the results with caution. Nevertheless, for comparison to other models in the literature, we share the same caveat that trading costs at extremely large sizes are rarely observed in the data and hence difficult to know the true cost.[15]

B. Break-even capacity

Using our estimated trading cost model from Table V, we calculate the break-even fund sizes/capacities of the anomaly/factor portfolios. The break-even size or capacity is the size of the fund that generates enough price impact that trading costs fully eliminate the average gross return to the strategy. Specifically, using the parameterized trading cost model from Table V we increase the trade size for each strategy to the point where price impact times the turnover of the strategy equals the expected gross return on the strategy. We then take the break-even trade size and back out the implied fund size that would generate those trades.

[15] One important point worth repeating is that our manager does not simply trade only when it expects costs to be low, or relatedly only when it can trade at a low %DTV. The trading algorithm does not make portfolio weight decisions, as evidenced by our costs being the same when only looking at exogenous trades from new inflows.

Table VI reports the results. Panel A examines the U.S. equity strategies and Panel B the international equity strategies. For break-even calculations we need an estimate of the gross return to each strategy. We start by using the most precise expected return estimate we have using the longest history of available data (1926 to 2013 in the U.S., 1986 to 2013 internationally). As the first column of Panel A of Table VI shows, for SMB, assuming an annual gross return of 2.95 percent, the break-even fund size is $275.52 billion. At that size, the average trade would incur 87.85 bps of price impact, which at 29% turnover per month would completely wipe out the 3.03 percent SMB premium. For HML, the break-even size is $214.28 billion assuming a gross return of 4.95, which is its historical average. For UMD, the break-even size is $56.16 billion. Although UMD has more than twice the turnover of HML and almost four times that of SMB, momentum also enjoys the largest return premium, which historically averages 8.20 percent per year. Finally, the combination portfolio of all three has a break-even capacity of just under $100 billion.

We also compute break-even capacities of these strategies using a more recent sample period from 1980 to 2013 to estimate their premia and to address concerns that the returns to these strategies have waned over time (in particular, SMB see Asness, Frazzini, Israel, Moskowitz, and Pedersen (2015)). In addition, for the portfolio optimization exercises in the next section, we use a "tradeable" sample of the largest 2,000 stocks from 1980 to 2013, where we have better estimates of their expected trading costs, and hence report their returns and non-optimized break-even sizes here for comparison. The returns to all three strategies are lower in the more recent sample period and hence break-even sizes are smaller, ranging from $26.6 billion for UMD to $153.8 billion for HML.

Panel B of Table VI reports the break-even sizes for the international sample, which range from just under $19 billion for UMD to $95 billion for HML, which are a bit smaller than the U.S. fund sizes, but still substantial.

C. Interpretation of break-even size

The break-even size analysis we conduct is similar to that employed in the literature (e.g., Korajczyk and Sadka (2004)) for ease of comparison, but what does it mean? A narrow and more precise interpretation is that these numbers represent the additional capacity our manager could run in each strategy before net returns are zero, ceteris paribus assuming the rest of the market remains the same. A broader interpretation is that these sizes represent the additional capacity available across all funds in the market assuming they face a similar price impact function as our manager. Thus, these numbers would represent the additional dollar amounts the market could tolerate before net returns ceased being positive (and assuming gross average returns remained the same as their historical averages). Of course, not knowing what the total dollars currently devoted to these strategies are, we cannot say anything about what the total market capacity is for these strategies.

Taken together, these results suggest that size, value, and momentum all survive trading costs at reasonably large fund sizes. Our numbers indicate that globally there is an additional $300 billion in size and value capacity and $80 billion in momentum capacity available. However, these numbers significantly understate capacity since they are based on standard academic portfolios that are not designed to minimize or manage trading costs. In the next section, we examine size, value, and momentum portfolios designed to minimize costs without incurring large tracking error or sacrificing gross returns and show that break-even capacities can be expanded by three to four times the numbers reported in Table VI.

C. Comparison with the literature

Before turning to portfolios optimized for trading costs, it is important to reconcile our results with those in the literature that estimate much higher transactions costs and far lower break-even sizes for similar strategies in U.S. stocks. Chen, Stanzl, and Watanabe (2002) estimate very small maximal fund sizes before costs eliminate profits on size, book-to-market, and momentum strategies. Lesmond, Schill, and Zhou (2003) find that trading costs eliminate the profits to momentum strategies at small fund sizes and Korajczyk and Sadka (2004) find that break-even fund sizes for long-only momentum portfolios are about $2 to a maximum of $5 billion. Our estimates of trading costs are an order of magnitude smaller than these studies and, consequently, our estimated break-even sizes are orders of magnitude larger. For example, taking the largest break-even fund size in the literature we could find – $5 billion from Korajczyk and Sadka (2004) for U.S. stock momentum – our break-even fund size for momentum is more than 10 times larger at $56 billion.

The main reason our results are so different is that previous studies estimate trading costs using aggregated daily or TAQ data that grossly overestimates the costs for a large institutional trader such as our manager. TAQ data approximates the average trade, which includes informed trades, retail trades, liquidity demanding trades, impatient trades, which all face costs much higher than those of a large institution patiently investing in these strategies. Hence, the previous literature's cost estimates are ill-suited to answer market efficiency questions with respect to these anomalies as they are not suitable cost estimates for the marginal trader/arbitrageur in these strategies.

For a more direct comparison, Panel A of Table VII compares the trading cost estimates from Korajczyk and Sadka (2004) (henceforth, KS) that uses TAQ data to our trading cost estimates on the common momentum portfolio UMD. KS estimate a linear cost function using TAQ data from 1993 to 2001. Our model uses more recent data from 1998 to 2013 and a square root function for trade size, which also lowers trading costs, but these differences are minor compared to the main difference which is use of TAQ data versus our live trading data.

We calibrate the KS model by taking their most generous break-even size of $5 billion on the momentum portfolio. As the first column of Table VII shows, the implied break-even average market impact per trade is 66.8 basis points, which would wipe out the entire momentum profits, using the historical mean return to momentum of 8.20 percent per year. The second column reports this exercise for the same $5 billion fund size using the more recent TAQ data that lines up with our sample period from 1998 to 2013. At $5 billion the annual trading costs to momentum using the more recent TAQ data are about 6% per year, leaving a net return to momentum of 2.2% per year. Hence, trading cost estimates using more recent TAQ data are lower, as expected, due to a declining trend in costs in markets. The third column reports the costs and net returns of the same $5 billion momentum portfolio under our trading cost model (henceforth FIM) using our proprietary live trading data.[16] The costs are less than half what we get from TAQ data over the same time period. For a $5 billion momentum portfolio, trading costs are 2.95% per year, leaving a 5.25 percent net average return to UMD. So, while the different time periods explain part of the reason why our cost estimates are lower than KS (2004), the main difference comes from using TAQ data versus our live institutional trade data.

The fourth column of Table VII reports the break-even fund size associated with the more recent 1998 to 2013 TAQ data, which is $10.6 billion. At $10.6 billion, the costs estimated from the most recent TAQ data would wipe out all of the 8.2 percent momentum premium. Using the same fund size, column five computes the costs under our model calibrated with live trade data and finds the annual trading costs to be only 3.8%, leaving a significant positive net return to momentum of 4.41% per year. The TAQ data simply do not provide a good approximation for the trading costs of a large institutional investor actually engaged in these strategies. Finally, to appreciate how different the trading cost estimates are from TAQ versus our real trading costs, the last column of Table VII reports the break-even fund size of UMD under our model, which is $56.16 billion or more than five times larger than what we get from recent TAQ data and more than 10 times larger than what KS (2004) claim (and that is for their most generous fund size estimate).

[16] Since our manager did not trade every stock in UMD nor trade the stocks necessarily at the exact time UMD required rebalancing, we use the trading cost model based off of actual live trades to estimate the costs of UMD using the market and stock's characteristics and our parameter estimates from Table V.

We argue that our trading cost estimates are much more reasonable for a large arbitrageur than those made from TAQ data. However, because we do not know the true capacity of the anomaly/factor portfolios we study, we cannot say for sure whether the TAQ trading cost estimates are way too high (our claim) or perhaps our estimates are too low. To address this issue, Panel B of Table VII makes another comparison between our estimates and the literature's. We conduct the same exercise on the CRSP value-weighted index. The first column of Panel B reports the break-even size of the index according to the KS model estimates. The CRSP index is a buy and hold cap-weighted index, but incurs 1.4% turnover per month due to new issues and events such as mergers, etc. According to the calibrated KS model, the capacity of the CRSP index would be $23.88 trillion, assuming a historical risk premium of 5.28 percent per year (its sample mean in excess of the Treasury bill rate from 1926 to 2013). The current capitalized value of the index is $25.7 trillion. So, at roughly $2 trillion less than the actual value of the index, the KS estimates imply that annual trading costs would entirely wipe out the equity premium (actually, they would make it slightly negative). This seems unreasonable. At the same fund size, our annual trading cost estimate for the CRSP index is 37 basis points, leaving a net equity premium of 4.91 percent.

Put differently, at the current size of the CRSP index – $25.7 trillion highlighted in the last two columns of Panel B – the KS estimates imply a -0.41% net return while our estimates imply a 4.89 percent equity premium after trading costs. While we can debate the magnitude of the gross equity premium, the trading costs that come out of the KS calibrated model seem exorbitant for essentially a buy and hold portfolio. At the market's current size, the KS estimates imply that the gross equity premium would have to be at least 5.70 percent per year to provide even a 0.01 positive net equity premium. To get a 1, 2, or even 3 percent equity premium, the gross excess return on the market would have to be 6.7, 7.7, or 8.7 percent per year, which are well beyond most model's theoretical and empirical predictions.[17]

Panel C of Table VII conducts a similar exercise for the S&P 500 index. The nice thing about the S&P 500 index is that tradeable funds are based off of it and so we can compare the actual trading costs of those funds, whereas the CRSP index is not tradeable. In addition, the S&P 500 is much closer to a buy and hold portfolio with only 0.4% turnover per month. We start in the first column of Panel C by calculating the trading cost and net return implied by the KS estimates at the current size of the S&P 500 index, which is $19,840 billion. At that size, KS estimates imply annual costs of 1.59% per year, which seems extraordinarily high for a portfolio that barely trades, meaning

[17] Most models have a tough time justifying the historical 5% equity premium (Hansen and Singlet (1982), Mehra and Prescott (1985)).

that the price impact estimates are enormous. Conversely, our trading cost estimates at that size for the S&P 500 are 0.11% or 11 basis points per year. Rather than use the entire S&P 500 capitalization as the fund size, the last two columns of Panel C report trading costs and net returns at the current dollar amount estimated to be benchmarked to the S&P 500, obtained from Standard & Poor's website – $7.8 trillion. According to the KS estimates, the $7.8 trillion of funds being benchmarked to the S&P 500 would experience an annual trading cost drag of 63 basis points per year, which seems to be an order of magnitude too high. In contrast, our estimate of trading costs on the S&P 500 at that size is 6 basis points (in the last column).

We can also compare these cost estimates to actual costs from live funds benchmarked to the S&P 500 index. Looking at the Vanguard S&P 500 fund and the iShares S&P 500 ETF, we obtain trading cost numbers from their prospectuses. Bearing in mind that these funds also may face additional turnover from inflows and outflows, which the index does not, the annual cost for the Vanguard S&P 500 fund is 12 basis points and for the iShares S&P 500 ETF it is 7 basis points per year. These numbers are exactly in line with our estimates and more than an order of magnitude smaller than the TAQ estimates. Hence, our costs seem very much in line with other institutional managers' live trading experience and are substantially less than estimates from the literature that grossly overestimate the cost of trading for a large institution.

Finally, our manager has also been running several index strategies based on one of the academic factors, momentum, since July 2009. Our manager runs a long-only momentum index in large and small cap U.S. stocks and internationally, and has mutual funds based off of these indexes. Table A7 in the appendix reports the live realized price impact costs of these funds, which are 6, 17.8, and 8 basis points in large cap, small cap, and international momentum, respectively. These numbers are in line with, and even lower than, our estimates from the historical trading data. The corresponding implied break-even fund sizes are therefore even larger than our estimates above.

Figure 3 highlights the difference between the KS price impact cost function estimated from TAQ data and our model estimated from proprietary live trades. The figure adds the estimated price impact from KS as a function of the fraction of DTV traded to Figure 2, which plots our trading cost model. As Figure 3 shows, the trading cost estimates for the same %DTV are vastly different, and the difference widens for larger trades. At 1% DTV, our model predicts 12.5 bps of price impact and the KS-TAQ model predicts 44.8 bps. At 10% DTV, our model implies 36 bps, whereas the KS model implies nearly 250 bps of price impact. Previous estimates of trading costs in the literature from aggregate transaction level data grossly overstate the costs for a large arbitrageur in these markets and, therefore, seem ill-equipped to address whether anomalies survive transactions costs. Our data from an actual large arbitrageur shows that actual trading costs are substantially lower and therefore the capacity of these strategies substantially higher than the literature portrays.

D. Out of sample net returns

Finally, our trading cost model can also be applied out of sample under a set of assumptions. For example, since returns to our strategies are available over a much longer sample period in both U.S. and international markets, we can assess the net expected returns of our strategies out of sample over these longer time periods. We use our trading cost model estimated in Table V and apply the coefficients to observable variables in the out of sample time periods. This assumes that the relation between the observable variables and trading costs from our model are the same over time. The model allows for time variation in costs as VIX or other observable variables change through time, but it assumes that the *relation* between those variables and costs remains the same through time. While this assumption may not be true, it provides a rough estimate of what trading costs might have looked like in earlier time periods if our model were accurate and predictive.

This exercise produces market impact estimates for the entire cross section of stocks from 1926 to 2013. We fix the average trade size (as a fraction of daily volume) equal to the median fraction of volume traded in our execution data. For stocks covered by our execution data, we take a conservative approach and take the maximum between the realized market impact (averaged over the past six months) and the predicted impact from the regression model. Table A7 in the Appendix reports the predicted price impact costs and net returns for U.S. stock portfolios from 1926 to 2013 and for international stocks from 1986 to 2013. The results are fairly consistent with those from the shorter and more recent sample period. All the strategies produce positive net expected returns after trading costs that are significant at the 5% level, except for SMB internationally.

Overall, our trading cost estimates suggest that size, value, and momentum all survive transactions costs at reasonably large sizes and that break-even capacities of these strategies are large and many times larger than previous estimates. However, the portfolios examined so far are those from the academic literature, which make no attempt to take expected trading costs into account in their design. Using our calibrated trading cost model, we now investigate whether these strategies can be successfully redesigned to minimize trading costs without incurring significant tracking error or “style drift” with respect to the desired factors, and what the tradeoffs are between tracking error and trading costs across the different anomalies.

V. Optimizing Portfolios for Trading Costs

To more fully address the trading cost efficiency of the anomaly/factor portfolios, we attempt to maximize the after-trading cost returns and capacity of each portfolio using trading cost optimization. As we have stated several times, our realized trading cost data comes from algorithmic trades that can alter the timing/patience of a trade, but do not make any portfolio weight or quantity decisions. In this section, we apply trading cost optimization using our model for expected price impact to alter quantities traded and allow the portfolio to change in response to expected trading costs. It is worth mentioning that if our manager was already altering these portfolios in response to expected trading costs, then this exercise would yield nothing.

We design "trading cost managed" versions of our portfolios that try to maximize after-trading cost returns of each strategy, subject to a tracking error constraint. Portfolios analyzed in the literature are not designed to optimize or pay attention to transactions costs in any way and hence may be trading-cost inefficient. In order to answer how cost-efficient various investment styles or factors are, it is crucial to evaluate how trading costs can be optimized for a given factor. Comparing the after-trading costs returns of the original academic versions of the factors/anomalies to those optimized for trading costs provides a sense of how large the improvements are in trading costs across the different anomalies.

A. Trading cost optimization

The objective is to maximize after-trading cost returns subject to maintaining the "style" or factor exposure of the original portfolio. We place a 0.75 percent (annualized) constraint on the amount of tracking error or style drift we allow the optimized portfolio to have. We want to optimize for trading costs but not at the expense of producing a portfolio that is too dissimilar from the factor itself.[18] The optimizer trades off the trading cost consequences of trading against the opportunity cost of not trading by minimizing expected trading costs, while imposing a cost for tracking error. To improve precision, we also constrain the trade size in any stock to less than 5% of a stock's average daily trading volume. Since the bulk of our trading cost realizations are less than this threshold (see Table I and Figure 2) our trading cost estimates are more precise in this range. While our break-even capacity analysis often extrapolates beyond this range, we recognize that those are imprecise numbers and treat them as rough approximations for capacity. However, when conducting portfolio

[18] For example, we could buy and hold a portfolio and never trade for the entire sample period in order to minimize trading costs, but this portfolio would not look anything like its intended style.

optimization, where stock-specific costs are required, we need to be more precise. Hence, we restrict the data and parameters to trades for which we have precise cost information.

The optimization problem is

$$\min_{\boldsymbol{w}} Total\ Trading\ Cost\ (\boldsymbol{w})$$

Subject to:

$$\text{(5)} \qquad \begin{aligned} &\textit{Tracking Error Constraint}: \sqrt{(\boldsymbol{w}-\boldsymbol{B})\boldsymbol{\Omega}(\boldsymbol{w}-\boldsymbol{B})} \leq 0.75\% \\ &\textit{\$1 long and \$1 short}: \quad \boldsymbol{w}'\boldsymbol{i} = 0 \text{ and } |\boldsymbol{w}|'\boldsymbol{i} = 2 \\ &\textit{Trading Constraint: Fraction of daily volume} <=5\% \end{aligned}$$

where $\boldsymbol{w}$ is the vector of chosen portfolio weights, $\boldsymbol{B}$ is the vector of original weights for the factor portfolio, $\boldsymbol{i}$ is a vector of 1's and $\boldsymbol{\Omega}$ is the CAPM-implied covariance matrix estimated using daily data over the prior 12 months. For simplicity, we use the CAPM as a risk model so the inputs of the covariance matrix are a stock's beta, its idiosyncratic volatility, and market volatility.[19] The trading cost function is based on column (4) of Table V.

Since we are interested in studying implementable portfolios, and to ease computational burden and because volume information and trading cost estimates are more noisy in the earlier time periods, we restrict our sample to the period 1980 to 2013 and use stocks with relatively active trading markets, mimicking the universe of liquid stocks traded by large institutional managers. We focus on the top 2,000 stocks in the U.S. based on their combined rank of size and daily volume, and do the same for the top 2,000 stocks outside of the U.S. across all other markets. We refer to this sample of 4,000 stocks as our "tradable" universe. These additional restrictions allow the optimization to run in reasonable time and are also consistent with the set of stocks in the live trading database that our trading cost model is based on.

In order to run the optimizations and estimate trading costs, we must input a portfolio size. Table VIII reports results where we start with a fund size of $200 million in net asset value (NAV) in 1980. Panel A reports results for the U.S. and Panel B internationally. Focusing on the U.S. equity portfolios, the second row reports the ending NAV for each strategy, which ranges from $1.7 billion for SMB to $4.9 billion for the combination of all factors and is a function of both the gross returns on each strategy as well as the trading costs of each strategy over the sample period. The next set of rows report the gross and net returns (and their *t*-statistics) of the non-optimized portfolios. At the

[19] Using the Fama and French (1993) three factor model augmented with a fourth momentum factor does not have a material impact on the results.

sizes traded, SMB and HML have their returns cut in half by trading costs, while UMD being a higher turnover strategy, has its return cut from 5.6 percent gross to 1.3 percent after trading costs.

The next set of rows report gross and net returns for the optimized portfolios we obtain from equation (5). First, the gross returns are very close to those for the non-optimized portfolios, suggesting that the tracking error constraint does a good job of capturing each factor and not allowing significant style drift. Or, put differently, there are little expected return consequences associated with these optimized factors/anomaly portfolios. Second, the optimized portfolios significantly reduce trading costs. So, while the gross returns are similar, the net returns of the optimized portfolios are much higher than the original portfolios. For SMB, the optimized portfolio produces a 2.02 percent return net of trading costs compared to only 0.87 percent non-optimized. For HML, the optimized strategy has a 3.02 net return versus 1.29 non-optimized.

These improvements come solely from the reduction in trading costs through portfolio design (save SMB, which has a slightly higher gross return than the original SMB portfolio). The next set of rows report the total trading costs before and after optimization, as well as the change in turnover and market impact that comprise those trading costs. Non-optimized SMB generates 88 basis points of total trading costs per year, but the optimized version of SMB generates only 13 basis points of trading costs. This reduction in trading costs comes from reducing the turnover of the strategy from 0.28 to 0.14 per month, but also from trading stocks with significantly lower expected market impact. While the original SMB portfolio trades stocks with an average 25.7 basis point market impact per trade, the optimized version trades stocks with only 8.2 basis points of average market impact. HML's trading cost reduction is even more dramatic, going from 1.89 percent per year to 55 basis points after optimization. The optimization reduces HML's turnover slightly, but cuts the average market impact per trade by more than 60%. Again, since the gross returns to the non-optimized versus optimized versions of HML are similar, this results in significant improvement in HML's net returns after trading costs. For UMD, the optimization has the most impact. The non-optimized version of UMD faces 4.31% trading cost drag per year, but an optimized version can reduce those costs to 1.99% per year, leaving significant after-trading-cost returns to momentum. The optimization reduces turnover by one third and market impact per trade by one third, resulting in significant trading cost savings.[20] The combination of all three factors also experiences a significant reduction in trading costs and reliably positive net expected returns.[21]

An arbitrageur facing trading costs who wishes to implement these strategies would both manage the turnover of the strategies and be strategic in which stocks he/she traded in order to minimize price impact costs. Our optimized portfolios achieve both and have a substantial impact on trading costs. Moreover, this reduction is achieved without incurring substantial tracking error or degradation in gross returns. The tracking errors are small, and the betas with respect to the original non-optimized portfolios (reported at the bottom of Table VIII) are very close to one.

Panel B of Table VIII reports results for the international strategies, and finds similar effects. SMB, HML, and UMD all experience significant improvements in trading costs through portfolio optimization that are similar in magnitude to those found in the U.S.

B. Tracking error frontier of trading costs and net Sharpe ratios

The results in Table VIII pertain to optimized portfolios with a tracking error constraint of 75 basis points. Table IX examines how the results change across factors as we vary the tracking error constraint, and compares the tradeoff between trading costs and tracking error across anomalies.

Table IX reports the total trading costs, net Sharpe ratios, and break-even fund sizes/capacities of the portfolios at ex ante tracking errors ranging from zero (the non-optimized/original portfolio) to 2% per year. Panel A reports the U.S. results and Panel B the international results. For each of the anomaly/factor portfolios, trading costs decrease monotonically as tracking error is allowed to increase, indicating that the optimization is doing what it is supposed to do. The reductions are substantial. For example, HML's 189 basis point annual trading cost in the U.S. can be reduced to 9 basis points through portfolio optimization that allows as much as 2% tracking error. UMD's 431 bps annual cost is reduced to 51 basis points at 2% tracking error. The combination of all three strategies can similarly be vastly improved from a trading cost perspective, declining from 381 bps to 18 bps at 2% tracking error. The numbers from the international sample are remarkably similar.

The next stanza of results reports the net Sharpe ratios of the strategies after trading costs for different tracking error constraints. As Table IX indicates, the net Sharpe ratio peaks at about 100-150 basis points of tracking error for all three strategies (and their combination). This indicates that at

[20] For the short-term reversal strategy, which faces a 10.82 trading cost per year, the optimization can reduce this to 6.17 percent, but the net returns are still not significant at a $200 million starting NAV.

[21] There are significant trading cost synergies from combining factors, especially value and momentum because of their negative correlation, that allow trading costs to be reduced even further, which is another virtue of combining value and momentum in the same portfolio (Asness (1997), Asness, Moskowitz, and Pedersen (2012), Daniel and Moskowitz (2012), and Israel and Moskowitz (2012)).

more tracking error than 1-1.5%, the decline in gross returns exceeds in magnitude the reduction in trading costs. At tracking error lower than 1%, trading costs decline faster than gross returns, leading to further net after-cost improvement in the portfolio. This evidence indicates that trading costs can be significantly reduced without incurring large return consequences or tracking error costs. The biggest improvements in net Sharpe ratio come from a combination of all three factors.

Figure 4 summarizes these results by plotting the total trading costs, gross, and net returns of each strategy across tracking error. As the graph illustrates, UMD seems to benefit the most from portfolio optimization as its trading costs decline the fastest as tracking error increases. However, the gross returns to UMD also decline the fastest as tracking error increases. Given that UMD is the highest turnover strategy, these results are intuitive. Around 1% tracking error is the tipping point for the tradeoff between trading costs and gross returns and hence maximizes net returns.[22]

C. Break-even capacity after optimization

How much do break-even sizes of these strategies improve from portfolio optimization? To answer this question, we first examine how trading costs change not just across tracking error, but also across fund size. Figure 5 plots the trading cost-tracking error frontiers of each anomaly at different starting fund sizes: $100 million, $200 million, $500 million, $1 billion, $2 billion, and $5 billion. The ending NAVs are determined according to each fund's net return path over the sample period. For example, a $1 billion fund in SMB would grow to $1.8 billion, but in HML would be $3.7 billion, and in UMD would be $6.1 billion. As the figure shows, trading costs are, of course, higher at larger NAVs, but the differences are muted at higher tracking error, where the portfolio optimizer can effectively reduce market impact. There is a steep decline in trading costs at larger sizes when the optimizer is given more latitude, enabling the portfolio to reduce turnover and avoid trading stocks with large expected market impact. At small fund sizes, the reduction in trading costs is much flatter, since price impact is less concerning at those sizes.

Applying this analysis to the trading cost optimizer, we calculate the break-even fund sizes of these strategies at different levels of allowable tracking error. The last two stanzas of Table IX report the break-even sizes for the optimized portfolios at different tracking error allowances, where we use both the in-sample mean return on each strategy over the more recent period as well as the longest historical period available to compute average returns for our break-even calculations. The zero

[22] It is worth repeating that the significant improvement in trading costs and net returns through portfolio optimization further indicates that the trading costs we estimate from the live trade data are largely independent from the portfolios being traded. If portfolio weights and trading costs had already been simultaneously optimized, and given the manager was running similar strategies, we would have found no further improvement from optimization.

tracking error break-even sizes are identical to those in Table V and represent the non-optimized capacities. As the table shows, there is a marked increase in fund size for all strategies at higher tracking error. SMB goes from $275.52 billion to $1.15 trillion in capacity after optimization at only 50 bps of tracking error. HML's capacity more than doubles (from $214 to $558 billion), and UMD's break-even capacity increases by slightly less than twice from $56 billion to just under $100 billion. At higher tracking error, the break-even sizes increase further, until the point where gross returns are reduced by more than trading costs are reduced. For example, for UMD that tradeoff is maximized at about 100 bps of tracking error, which is the maximum break-even fund size of $159.26 billion.

Panel B of Table IX repeats these calculations for the international portfolios and finds remarkably similar improvements. Figure 6 plots the break-even fund sizes for momentum in the U.S. and internationally using both the historical and in-sample risk premium estimates. In all cases, the largest break-even fund size occurs at 100 basis points of allowable tracking error.

Adding everything up, on a global scale, we estimate that through portfolio optimization with less than 1% tracking error, SMB, HML, and UMD have $1.9 trillion, $1.5 trillion, and $209 billion in capacity. A combination of all three strategies has almost $500 billion in capacity. These break-even sizes are orders of magnitude larger than those previously estimated in the literature. For example, Korajczyk and Sadka (2004) estimate at most a break-even fund size for momentum of $5 billion in the U.S., whereas we find that momentum capacity in the U.S. can reach $160 billion.

VI. Conclusion

We examine the trading costs, net of cost returns, and break-even fund sizes of equity strategies designed to capture several of the main asset pricing anomalies documented in the literature. Using a unique dataset of live trades from a large institutional investor, who engages in many similar strategies, we measure the trading costs of a large arbitrageur. Our trading costs are an order of magnitude smaller, and fund sizes orders of magnitude larger, than those claimed in the literature.

Two key factors contribute to our capacity estimates being several orders of magnitude higher. First, we use actual trading costs from a real-world arbitrageur, who seeks to avoid demanding liquidity in markets, to estimate price impact rather than aggregated trade and quote level data used in other studies. Second, we show that simple changes to portfolio design can significantly reduce trading costs without incurring much tracking error or factor/style drift, further increasing capacity. We use our unique data to build an empirical model of expected trading costs for use in portfolio optimization to design strategies that account for expected trading costs. The combination of these two features provides orders of magnitude larger break-even sizes.

Since these two innovations are what smart arbitrageurs would do in practice, we view our trading cost and break-even capacity estimates as much more relevant for answering questions regarding market efficiency with respect to these anomalies. Our results indicate that strategies based on size, value, and momentum can be deployed at very high asset size and still survive trading costs. Other strategies, such as short-term reversals, may not. Hence, the return premia associated with size, value, and momentum appear to be robust, sizeable, and implementable.

Our estimates and model provide a metric of real-world trading costs facing a large arbitrageur that we hope will be used to further study these and other investment strategies.

References

Almgren (2001)

Almgren, . . .

Asness, Cliff S. (1994), "Variables that Explain Stock Returns", Ph.D. Dissertation, University of Chicago.

Asness, Cliff S., and Andrea Frazzini (2012), "The Devil in HML's Details," Working Paper, AQR Capital Management.

Asness, Cliff S., Andrea Frazzini, Ronen Israel, Tobias J. Moskowitz, and Lasse Pedersen (2015), "Size Matters: If You Control Your Junk," Working Paper, AQR Capital Management.

Asness, Cliff S., Tobias J. Moskowitz, and Lasse H. Pedersen (Forthcoming), "Value and Momentum Everywhere," *The Journal of Finance*.

Engle, Robert, Robert Ferstenberg and Jeffrey Russell (2012), "Measuring and Modeling Execution Cost and Risk," *The Journal of Portfolio Management*.

Fama, E., French, K., 1992. The cross section of expected stock returns. Journal of Finance 47, 427–465.

Fama, Eugene F. and Kenneth R. French (1993), "Common Risk Factors in the Returns on Stocks and Bonds," *The Journal of Financial Economics*, vol. 33, pp. 3-56.

Fama, E., French, K., 1996. Multifactor explanations of asset pricing anomalies. Journal of Finance 51, 55–84.

Fama, E., French, K., 1998. Value versus growth: the international evidence. Journal of Finance 53, 1975–1999.

Fama, E., French, K., 2012. Size, value, and momentum in international stock returns. Forthcoming Journal of Financial Economics.

Fama, E., MacBeth, J., 1973. Risk, return, and equilibrium: empirical tests. Journal of Political Economy 607-636.

Garleanu, N. and Pedersen L.H., 2012, "Dynamic Trading with Predictable Returns and Transaction Costs", Journal of Finance, forthcoming.

Grinblatt, Mark and Tobias J. Moskowitz (2004), "Predicting Stock Price Movements from Past Returns: The Role of Consistency and Tax-Loss Selling", *The Journal of Financial Economics*, vol. 71, pp. 541-579.

Israel, Ronen and Tobias J. Moskowitz (Forthcoming), "The Role of Shorting, Firm Size, and Time on Market Anomalies." *The Journal of Financial Economics*.

Jegadeesh, Narasimhan and Sheridan Titman (1993), “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *The Journal of Finance*, vol. 48, pp. 65-91.

Keim, D.B. (1999), “An analysis of mutual fund design: the case of investing in small-cap stocks,” *The Journal of Financial Economics*, vol. 51, pp. 173-194.

Keim, D.B., Madhavan, A. (1997), “Execution costs and investment style: an interexchange analysis of institutional equity trades,”*Journal of Financial Economics*, vol. 46, pp. 265-292.

Korajczyk, Robert A., and Ronnie Sadka (2004), “Are momentum profits robust to trading costs?,” *Journal of Finance*, vol. 59, pp. 1039-1082.

Lehmann, Bruce (2003), “What We Measure in Execution Cost Measurement”, *Journal of Financial Markets*, vol. 6, pp. 227-231.

Lesmond, David A., Michael J. Schill, and Chunsheng Zhou (2003), “The illusory nature of momentum profits,” *Journal of Financial Economics*, vol. 71, pp. 349–380.

Perold, Andre F. (1988), “The implementation shortfall: paper vs. reality,” *Journal of Portfolio Management*, vol. 14, pp. 4–9.

Sensoy, Berk (2009), “Performance Evaluation and Self-Designated Benchmark Indexes in the Mutual Fund Industry,” *Journal of Financial Economics*, vol. 92, no. 1, pp. 25-39.

Werner, Ingrid (2003), “NYSE Order Flow, Spreads and Information”, *Journal of Financial Markets*, vol. 6, pp. 309-335.

Table I
Trade Execution Data Summary Statistics (1998 – 2013)

This table shows summary statistics of the trade execution database. Panel A reports the total amount traded by year. Amounts are in billion USD. "U.S." are trades executed in the United States. "International" are trades executed outside of the United States. "Large Cap" are trades in large capitalization securities, "Small Cap" are trades in small capitalization securities. The distinction between large and small cap is based on the portfolio's benchmark (typically large cap is the Russell 1000 and above and small cap is below the Russell 1000 in market cap, but typically within the Russell 2000 size universe). "Long Short" are trades executed in long-short accounts and "Long only" are trades executed in long-only accounts. Relaxed constraints portfolios (130-30 and 140-40) are classified as "Long only", though these portfolios are small relative to the trades made in other portfolios. Panel B reports times-series means of yearly cross-sectional summary statistics as of December of each year. Panel C reports Fama-MacBeth averages, where each calendar month we compute value-weighted means of each variable, with weights equal to the total amount traded, and the time series means of the cross sectional coefficients are reported in Panel C. "Fraction of daily average trading volume" is equal to the trade's dollar size divided by the stock's average prior one-year dollar volume. This table includes all available developed market equity transactions (cash equities and equity swaps) in our data between August 1998 and September 2013. * indicates partial year from August 31, 1998. ** indicates partial year up to September 30, 2013.

Panel A: Amount Traded (Billion USD)

		By region		By size		By portfolio type	
Year	Total	U.S.	International	Large Cap	Small Cap	Long short	Long only
1998*	2.96	1.29	1.67	2.96		2.96	
1999	5.29	1.99	3.30	5.29		5.29	
2000	1.99	0.76	1.23	1.99		1.86	0.13
2001	1.08	0.55	0.53	1.08		1.00	0.08
2002	4.21	0.71	3.50	4.21	0.00	1.40	2.81
2003	5.43	2.69	2.75	5.43	0.00	4.17	1.26
2004	10.00	2.95	7.05	9.99	0.01	6.38	3.62
2005	16.16	8.06	8.10	15.75	0.41	11.45	4.71
2006	67.01	34.79	32.22	64.23	2.78	44.69	22.31
2007	129.46	50.70	78.76	125.21	4.25	96.65	32.81
2008	108.29	25.06	83.24	104.27	4.02	69.30	38.99
2009	111.12	18.58	92.54	108.12	2.99	85.50	25.62
2010	117.17	29.15	88.02	113.78	3.38	91.94	25.23
2011	146.50	56.62	89.88	141.93	4.58	115.69	30.81
2012	179.09	121.39	57.70	173.41	5.68	141.97	37.13
2013**	141.18	92.87	48.31	136.04	5.14	95.21	45.98
Total	1,046.94	448.15	598.79	1,013.69	33.25	775.46	271.48

Panel B: Annual time series

	Mean	Median	Std	Min	Max
Number of stocks per year	3,256	3,732	1,592	386	5,105
Number of countries per year	17.5	19.0	4.2	8.0	21.0
Number of exchanges per year	24.6	26.0	7.1	11.0	34.0

Panel C: Fama MacBeth averages

	Mean	Median	Std	Min	Max
Average trade size (1,000$)	658	345	990	53	5,993
Fraction of average daily volume (%)	1.1	0.5	2.0	0.1	13.1
Trade horizon (days)	2.0	1.1	2.1	0.0	8.8

Table II
Trade Execution Data, Realized Trading Costs

This table shows average Market Impact (MI) and Implementation Shortfall (IS). Each calendar month, we compute average, median and weighted average cost ("vw mean") of all trades during the month. When computing weighted average cost, trades are weighted by their dollar amount. This table reports time-series averages of the cross sectional estimates. When computing time series averages, we weight each monthly observation by the number of stocks traded during the month. Statistics pertain to all available developed market equity transactions (cash equities and equity swaps) in our data between August 1998 and September 2013. The distinction between large and small cap is based on the portfolio's benchmark (typically large cap is the Russell 1000 and above and small cap is below the Russell 1000 in market cap, but typically within the Russell 2000 size universe). "Long Short" are trades executed in long-short accounts and "Long only" are trades executed in long-only accounts. Relaxed constraints portfolios (130-30 and 140-40) are classified as "Long only", though these portfolios are small relative to the trades made in other portfolios. Panel A reports statistics over the full sample period from August 1998 to September 2013 and Panel B from January 2003 to September 2013. Market Impact (MI) and Implementation Shortfall (IS) are in basis points and standard errors are reported at the bottom of each panel. Panel C reports MI (vw mean) for different types of trades: buy long, buy to cover, sell long, and sell short, by region and size.

Panel A: Full sample 1998 to 2013

	All	By region			By size		By portfolio type	
		US Nyse-Amex	US Nasdaq	Int.	Large cap	Small cap	Long short	Long only
MI mean	11.25	8.81	11.44	12.67	10.25	20.99	9.29	15.54
MI median	7.36	6.24	6.16	8.75	6.79	15.08	6.31	10.10
MI vw mean	17.06	14.39	16.53	17.69	16.52	25.50	16.30	17.23
IS mean	11.89	9.13	11.18	13.79	10.96	20.97	10.41	15.13
IS median	9.42	7.63	7.61	11.32	8.80	17.11	8.29	12.03
IS vw mean	18.34	16.04	18.93	18.82	17.77	28.03	17.58	18.20
Standard errors								
MI mean	0.83	0.90	1.81	0.80	0.86	1.40	0.99	0.94
MI median	0.45	0.66	0.85	0.58	0.48	1.22	0.54	0.79
MI vw mean	1.00	1.25	2.30	0.97	1.02	1.88	1.20	1.06
IS mean	1.08	1.13	2.16	1.09	1.13	1.67	1.29	1.17
IS median	0.66	0.89	1.27	0.85	0.69	1.62	0.77	0.97
IS vw mean	1.25	1.67	2.73	1.11	1.28	2.14	1.44	1.28

Panel B: Recent sample 2003 to 2013

	All	By Region			By Size		By Portfolio type	
		US Nyse-Amex	US Nasdaq	Int.	Large cap	Small cap	Long short	Long only
MI mean	11.21	8.92	11.84	12.41	10.18	21.21	9.23	15.56
MI median	7.22	6.09	6.12	8.55	6.62	15.15	6.12	10.08
MI vw mean	16.91	14.47	16.83	17.44	16.35	25.80	16.15	17.37
IS mean	11.88	9.22	11.73	13.54	10.92	21.17	10.39	15.14
IS median	9.29	7.49	7.63	11.15	8.64	17.19	8.11	12.06
IS vw mean	18.18	16.10	19.12	18.58	17.60	28.35	17.42	18.35
Standard errors								
MI mean	0.67	0.74	1.01	0.83	0.70	1.35	0.70	0.94
MI median	0.45	0.52	0.55	0.61	0.47	1.18	0.50	0.80
MI vw mean	1.00	1.21	1.67	1.03	1.00	1.88	1.17	1.02
IS mean	0.99	1.05	1.28	1.20	1.03	1.63	1.11	1.19
IS median	0.70	0.79	0.81	0.96	0.73	1.59	0.77	0.99
IS vw mean	1.31	1.74	2.16	1.22	1.32	2.14	1.45	1.28

Panel C: Market impact by trade type								
		Fraction of sample		All	By region		By size	
		Dollars	Trades		US	Int.	Large cap	Small cap
MI (VW-mean)	Buy Long	0.35	0.32	12.51	14.79	10.30	11.50	27.02
	Buy Cover	0.15	0.17	16.22	17.46	15.09	15.89	37.14
	Sell Long	0.32	0.32	18.18	11.06	22.46	17.59	27.43
	Sell Short	0.18	0.19	20.88	10.02	26.53	20.57	30.07
Differences	Buy Cover - Buy Long			3.72	2.66	4.79	4.39	10.12
	Sell Short - Sell Cover			2.70	-1.03	4.07	2.98	2.65
t-statistics	Buy Cover - Buy Long			1.14	0.57	1.36	1.28	0.71
	Sell Short - Sell Cover			0.99	-0.24	1.07	1.04	0.16

Table III
Trade Execution Data, Realized Trading Costs – Long-Only Inflows

This table reports average Market Impact (MI) from new inflows of funds into long-only funds. Each calendar month, we compute average, median and weighted average cost ("vw mean") of all trades during the month, where the weighted average costs are weighted by the dollar amount of trading. Time-series averages of the cross sectional estimates are reported, where each monthly observation is weighted by the number of stocks traded during the month. Statistics are reported for all available developed market equity transactions executed in long-only accounts (cash equities and equity swaps) in our data between August 1998 and September 2013 coming solely from "Inflows," which are defined as the first trade for a given long-only account. Market Impact is in basis points.

Long-only trades, 1998 - 2013	Trade type	Inflows only	All other trades	Difference	*t*-statistic
MI mean	All trades	16.95	15.54	1.40	0.20
MI median	All trades	12.51	10.23	2.28	0.35
MI vw mean	All trades	19.30	17.22	2.08	0.31
MI mean	Large cap	14.76	13.83	0.93	0.11
MI median	Large cap	9.88	9.13	0.75	0.09
MI vw mean	Large cap	11.59	16.71	-5.12	-0.66
MI mean	Small cap	20.94	20.56	0.37	0.07
MI median	Small cap	17.08	14.59	2.48	0.54
MI vw mean	Small cap	26.29	24.47	1.82	0.25

Table IV
Factor/Anomaly Portfolio Returns Net of Actual Trading Costs

Reported are portfolio returns gross and net of actual trading costs taken from the live trading database for size (SMB), value (HML), momentum (UMD), and an equally weighted composite portfolio (COMBO) of all three strategies. All the portfolios are the intersections of two portfolios formed on size and three portfolios formed on book to market or prior one-year returns. At the end of each calendar month stocks are assigned to two size-sorted portfolios based on their market capitalization. The size breakpoint for the U.S. sample is the median NYSE market equity. The size breakpoint for the international sample is the 80th percentile by country (which roughly matches the US 50th percentile breakpoint). We measure book to market and lagged book divided by current price and update monthly (Asness and Frazzini (2013)). We use one year return (in local currency) skipping the most recent month for momentum (UMD). Stocks are further ranked into three groups (low, neutral, high) based on NYSE-based breakpoints (or breakpoints based on the top 20% of market capitalization by country for the international sample). The size portfolio SMB is the average return on the three small portfolios minus the average return on the three big portfolios. The value portfolio, HML, is the average return on the two value portfolios minus the average return on the two growth portfolios. UMD is constructed in the same manner as the average of the two winner (high past year return) portfolios minus the average return of the loser (low past year return) portfolios. All portfolios are value-weighted, refreshed and rebalanced every calendar month to maintain value weights. The portfolios contain all available stocks in our trade execution data at portfolio formation. The sample period runs from August 1998 to September 2013. Country portfolios are aggregated into international portfolios using the country's total market capitalization as of the prior month. For each portfolio, "Dollar traded per month" is the monthly average dollar traded in each portfolio over the prior six months taken directly from the trade execution database. "Implied fund size" is equal to the monthly average dollar traded in each portfolio divided by its turnover. Market impact, MI, is the realized market impact in our data in basis points (annualized), *t*-statistics are reported below the coefficient estimates and 5% statistical significance is indicated in bold. The break-even cost is based on the factor's average return over the longest sample available (July 1926 to September 2013 in the US, January 1986 to September 2013 internationally).

	Panel A: U.S. trade execution sample, 1998 - 2013				**Panel B: International trade execution sample, 1998 - 2013**			
	SMB	HML	UMD	Combo	SMB	HML	UMD	Combo
Dollar traded per month (billion USD)	9.69	5.97	6.18	9.69	11.80	7.15	8.38	12.77
Implied fund size (billion USD)	18.18	9.42	5.21	16.91	17.88	10.09	6.85	19.74
Correlation to Fama-French	0.78	0.96	0.97	0.88	0.53	0.91	0.94	0.79
Realized cost	1.47	1.35	3.03	1.46	1.70	1.54	2.24	1.24
Break-even cost	2.95	4.95	8.20	5.39	-0.17	5.78	7.65	4.68
Realized minus breakeven	**-1.48**	**-3.61**	**-5.17**	**-3.93**	**1.87**	**-4.24**	**-5.40**	**-3.44**
Full sample historical mean:								
Return (Gross)	**2.95**	**4.95**	**8.20**	**5.39**	-0.17	**5.78**	**7.65**	**4.68**
	(2.72)	(3.10)	(4.79)	(9.13)	(-0.12)	(3.01)	(2.98)	(5.22)
Return (Net)	1.48	**3.61**	**5.17**	**3.93**	-1.87	**4.24**	**5.40**	**3.44**
	(1.40)	(2.25)	(3.02)	(6.66)	(-1.30)	(2.20)	(2.10)	(3.75)
Live trading sample mean:								
Return (Gross)	**7.98**	4.86	2.26	**5.04**	1.17	5.59	4.02	**3.59**
	(3.01)	(1.12)	(0.40)	(3.17)	(0.75)	(1.83)	(0.92)	(2.88)
Return (Net)	**6.52**	3.51	-0.77	**3.58**	-0.53	4.05	1.78	2.35
	(2.48)	(0.80)	-(0.14)	(2.23)	-(0.33)	(1.32)	(0.41)	(1.86)
Turnover (monthly)	0.53	0.63	1.19	0.57	0.66	0.71	1.22	0.65
MI (bps)	22.94	17.71	21.30	21.22	21.42	18.12	15.27	16.02

Table V
Trading Cost Regression Model Based Off of Live Trades

This table shows results from pooled regressions, where the left-hand side variable is a trade's Market Impact (MI), in basis points. The explanatory variables include the contemporaneous market returns, firm size, volatility and trade size (all measured at order submission). "Beta*IndexRet*buysell" is the contemporaneous (beta-adjusted) market return that controls for market movements occurring at the time of trade. Beta is the stock's predicted beta at the time of order submission. "indexRet" is the corresponding index return over the life of the trade (the movement of the market). "buysell" is a dummy equal to 1 for buy orders and -1 for sell orders. "Time trend" is a linear time trend. Size of Log of ME is equal to the log of 1 plus the market value of equity log(1+ME). ME is in billions of USD. "Fraction of daily volume" is equal to the trade's dollar size divided by the stock's average one-year dollar volume (in %). We include both a linear and square root function of fraction of daily volume. "Idiosyncratic Volatility" is the volatility of the residuals of a regression of one-year daily stock returns on the corresponding value-weighted benchmark (annualized %), "Market Volatility" is the monthly variance of the CRSP-value weighted index, computed using daily returns (annualized %). The DGTW-adjusted return is the return on the stock minus the return on a portfolio of similar stocks matched on size, book-to-market, and momentum (past one year return) from Daniel, et al. (1997), which is interacted with the buysell dummy. Country fixed effects are included when indicated, *t*-statistics are shown below the coefficient estimates and 5% statistical significance is indicated in bold. Standard errors are clustered by calendar month. Regression estimates are provided over the full sample, the US only, and outside of the US ("International").

	All sample					U.S.		International	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Beta*IndexRet*buysell	**0.21**	**0.21**	**0.21**	**0.21**	**0.19**	**0.21**	**0.20**	**0.20**	**0.17**
	(17.37)	(17.38)	(17.38)	(17.36)	(12.29)	(8.43)	(8.26)	(24.15)	(21.18)
Time trend	**-0.08**	-0.05	-0.07	-0.05	-0.05	0.00	0.00	**-0.09**	**-0.09**
	-(2.34)	-(1.61)	-(1.92)	-(1.33)	-(1.40)	-(0.02)	-(0.03)	-(3.80)	-(3.66)
Log of ME (billion USD)	**-4.10**	**-3.40**	**-2.72**	**-1.86**	**-1.74**	**-1.83**	**-1.65**	**-1.95**	**-1.90**
	-(11.19)	-(9.39)	-(7.84)	-(7.97)	-(7.53)	-(3.70)	-(3.81)	-(7.94)	-(7.53)
Fraction of daily volume		**1.54**	**0.81**	**0.80**	**1.00**	0.82	**1.34**	**0.72**	**0.69**
		(6.87)	(3.39)	(3.27)	(4.36)	(1.84)	(2.56)	(4.47)	(3.99)
Sqrt(fraction of daily volume)			**6.50**	**7.81**	**7.51**	**7.83**	**6.84**	**8.14**	**8.38**
			(3.50)	(4.14)	(4.08)	(2.10)	(1.97)	(7.27)	(7.32)
Idiosyncratic volatility				**0.10**	**0.12**	0.05	0.08	**0.22**	**0.22**
				(2.06)	(2.82)	(0.70)	(1.45)	(6.02)	(6.40)
Vix				**0.27**	**0.24**	**0.19**	**0.17**	**0.32**	**0.29**
				(4.47)	(4.56)	(2.62)	(2.69)	(3.53)	(3.63)
DGTW adjusted ret *buysell					**0.04**		**0.06**		**0.05**
					(2.78)		(2.45)		(4.18)
Observations (1,000s)	2,125	2,125	2,125	2,125	1,945	1,005	981	1,120	964
Adjusted R2	0.071	0.072	0.072	0.072	0.073	0.069	0.071	0.076	0.076
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes

Table VI
Break-Even Fund Sizes of Factor/Anomaly Portfolios

This table reports the break-even fund size capacities of the factor portfolios SMB, HML, UMD, and an equal-weighted combination (combo) of all three of them using the trading cost model from Table V. Panel A reports results for U.S. equity portfolios and Panel B for international equity portfolios. Break-even fund sizes are determined by the size of the fund that results in annual trading costs due to price impact that equal the estimated gross return on the portfolio and hence deliver a zero net of trading cost return. Annualized gross returns are estimated over the full sample period (which is 1926 to 2013 for the U.S. and 1986 to 2013 internationally) and over a more recent sample period (from 1980 to 2013 in the U.S. and from 1993 to 2013 internationally) and reported in the first row. The more recent sample periods offer more reliable data and estimates of trading costs and are used for portfolio optimization analysis, hence we report those results, too. The annualized turnover of each strategy is reported along with the break-even fund size and average fraction of daily volume traded, average market impact (in bps), and total cost of trading that are consistent with the break-even fund sizes.

Panel A: U.S. sample	Full Sample premium , 1926 - 2013				Recent sample premium, 1980 - 2013			
	SMB	HML	UMD	Combo	SMB	HML	UMD	Combo
Gross return (annualized %)	2.95	4.95	8.20	9.71	1.76	4.05	5.60	6.90
Turnover (monthly)	0.29	0.44	1.02	0.89	0.29	0.44	1.02	0.89
Break-even NAV (billion)	**275.52**	**214.28**	**56.16**	**98.69**	**102.21**	**153.78**	**26.60**	**54.64**
Average fraction of daily volume traded (%)	36.67	39.95	22.83	38.63	13.60	28.67	10.81	21.39
Average market impact (bps)	87.85	92.54	66.81	90.67	50.93	75.98	45.58	64.47
Total cost (annualized %)	3.03	4.93	8.20	9.71	1.76	4.05	5.60	6.90

Panel B: International sample	Full Sample premium , 1986 - 2013				Recent sample premium, 1993 - 2013			
	SMB	HML	UMD	Combo	SMB	HML	UMD	Combo
Gross return (annualized %)	-0.17	5.78	7.64	7.23	0.24	5.18	6.18	6.86
Turnover (monthly)	0.43	0.51	1.11	0.99	0.43	0.51	1.11	0.99
Break-even NAV (billion)	**0.00**	**95.48**	**18.87**	**23.40**	**0.00**	**79.66**	**12.34**	**21.17**
Average fraction of daily volume traded (%)	0.00	41.57	17.25	19.40	0.00	34.68	11.28	17.55
Average market impact (bps)	11.27	94.83	57.48	61.15	11.27	84.97	46.50	58.00
Total cost (annualized %)	0.59	5.78	7.64	7.23	0.59	5.18	6.18	6.86

Table VII
Comparison to Korajczyk and Sadka (2004) and TAQ Data

Panel A reports the total annualized trading costs and net returns to the momentum portfolio, UMD, at various fund sizes under various trading cost model estimates. The first column reports the total cost and net return to UMD at a fund size equal to the Korajczyk and Sadka (2004) (KS (2004)) estimated break-even fund size ($5 billion) under their tcost model, which uses TAQ data from 1993 to 2001 to estimate trading costs. We calibrate the KS (2004) model to a break-even fund size of $5 billion and report below the average fraction of daily volume traded and average market impact associated with that break-even calculation and a historical risk premium for UMD estimated over the longest possible sample period (1927 to 2013). The second column repeats this exercise for the same $5 billion fund size, but using more recent TAQ data from 1998 to 2013, matching our institutional live transaction data sample period, to estimate trading costs. The third column applies our tcost model using our proprietary institutional trading data to estimate trading costs (FIM (2015)) for the same $5 billion dollar fund size. The fourth column uses the implied break-even fund size associated with the more recent TAQ data from 1998 to 2013 and the fifth column reports results for that same fund size under our FIM (2015) tcost estimates. The sixth column reports the break-even fund size under our FIM (2015) tcost estimates. Panel B reports results for a similar exercise for the market CRSP value-weighted index, where the first two columns compare results using the KS (2004) estimates versus the FIM (2015) estimates for the break-even fund size that would wipe out the entire equity premium according to the calibrated KS (2004) model, and the third and fourth columns report results using the actual total market capitalization (as of June 2015) of the CRSP market index. Panel C reports results for the S&P 500 index, where the first two columns compare the KS (2004) estimates versus the FIM (2015) estimates for a fund size equal to the total market cap of the S&P 500, and the third and fourth columns report results for a fund size equal to the approximate value of assets benchmarked against the S&P 500 according to Dow Jones. Finally, we report at the bottom of the table the actual annualized trading costs of two S&P 500 index funds: Vanguard S&P 500 Index Fund and the iShares S&P 500 ETF.

Fund size =	KS break-even NAV	KS break-even NAV	KS break-even NAV	TAQ data break-even NAV	TAQ data break-even NAV	FIM break-even NAV
tcost estimate =	KS (2004)	TAQ data 1998-2013	FIM (2015)	TAQ data 1998-2013	FIM (2015)	FIM (2015)
	Panel A: UMD					
Gross return (annualized %)	8.20	8.20	8.20	8.20	8.20	8.20
NAV ($billion)	**5.00**	**5.00**	**5.00**	**10.60**	**10.60**	**56.16**
Average fraction of daily volume traded (%)	2.03	2.03	2.03	4.31	4.31	22.83
Average market impact (bps)	66.81	48.77	24.02	66.81	30.91	66.81
Total cost (annualized %)	**8.20**	**5.99**	**2.95**	**8.20**	**3.80**	**8.20**
Net return (annualized %)	0.00	2.21	5.25	0.00	4.41	0.00

Fund size = tcost estimate =	KS break-even NAV KS (2004)	KS break-even NAV FIM (2015)	Actual total market cap KS (2004)	Actual total market cap FIM (2015)
	Panel B: CRSP Value-Weighed Market index			
Gross return (annualized %)	5.28	5.28	5.28	5.28
Turnover (monthly)	1.4%	1.4%	1.4%	1.4%
NAV ($billion)	**23,877.50**	**23,877.50**	**25,724.49**	**25,724.49**
Average fraction of daily volume	144.93	144.93	156.14	156.14
Average market impact (bps)	3,180.74	221.09	3,425.03	233.62
Total cost (annualized %)	**5.28**	**0.37**	**5.69**	**0.39**
Net return (annualized %)	0.00	4.91	-0.41	4.89

Fund size = tcost estimate =	Actual total market cap KS (2004)	Actual total market cap FIM (2015)	Amount benchmarked to S&P 500 KS (2004)	Amount benchmarked to S&P 500 FIM (2015)
	Panel C: S&P 500 index*			
Gross return (annualized %)	5.28	5.28	5.28	5.28
Turnover (monthly)	0.4%	0.4%	0.4%	0.4%
NAV ($billion)	**19,840.00**	**19,840.00**	**7,800.00**	**7,800.00**
Average fraction of daily volume traded (%)	145.08	145.08	57.04	57.04
Average market impact (bps)	3,184.19	221.27	1,265.52	115.82
Total cost (annualized %)	**1.59**	**0.11**	**0.63**	**0.06**
Net return (annualized %)	3.69	5.17	4.65	5.22

*Vanguard S&P 500 Index Fund annual tcosts = 0.12% per year.

*iShares S&P 500 ETF annual tcosts = 0.07% per year.

Table VIII

Portfolios Optimized for Trading Costs

This table reports returns gross and net of trading costs for portfolios optimized to minimize trading costs. We report returns of Size (SMB), value (HML), momentum (UMD), and an equally weighted composite portfolio of all three (COMBO). Country portfolios are aggregated into International portfolios using the country's total market capitalization as of the prior month. This table includes all available tradable stocks in the combined CRSP Xpressfeed/Compustat global data. The sample period for U.S portfolios runs from January 1980 to September 2013. The sample period for International portfolios runs from January 1993 to September 2013. To select the tradable universe, at the end of each calendar month we rank all stocks based on their market capitalization and on their one-year average daily volume. We require stocks to have non-missing market capitalization, one-year daily volume, and one year of prior return history to compute market beta and idiosyncratic volatility. We sum the market cap and volume rank and select the top 2,000 U.S. securities and the top 2,000 international securities based on the combined rank. In the international sample we select the top *cw**2,000 securities in each country where *cw* is the country's market capitalization weight. Portfolio are optimized for trading costs by minimizing the expected total dollar trading costs subject to a 0.75% tracking error constraint relative to the unconstrained portfolio, a maximum trade constraint equal to 5% of the stock's average daily volume and a total dollar long/short notional size of 1$ per side per 1$ dollar of NAV. For each security, the predicted market impact MI is computed using coefficients from column (4) Table V. Returns and costs are annualized, MI is in basis points, *t*-statistics are reported below the coefficient estimates and 5% statistical significance is indicated in bold.

Starting NAV: 200M	Panel A: U.S. tradable sample				Panel B: International tradable sample			
	SMB	HML	UMD	Combo	SMB	HML	UMD	Combo
Starting Nav (Million, USD)	200.0	200.0	200.0	200.0	200.0	200.0	200.0	200.0
Ending Nav (Million, USD)	1,714.8	2,336.0	1,813.7	4,910.1	324.1	843.9	592.5	959.0
Non-optimized excess return (gross)	1.76	4.05	**5.60**	**6.90**	0.24	**5.18**	6.18	**6.86**
	(1.13)	(1.80)	(2.06)	(4.77)	(0.18)	(2.08)	(1.86)	(4.11)
Non-optimized excess return (net)	0.87	2.16	1.29	**3.10**	-0.82	3.45	2.51	**3.73**
	(0.56)	(0.96)	(0.47)	(2.14)	(-0.60)	(1.39)	(0.75)	(2.24)
Optimized exess return (gross)	2.16	3.92	5.01	**6.48**	-0.04	**5.36**	5.30	**6.25**
	(1.43)	(1.81)	(1.92)	(4.46)	(-0.03)	(2.23)	(1.64)	(3.76)
Optimized excess return (net)	2.02	3.37	3.02	**5.17**	-0.28	**4.78**	3.56	**5.09**
	(1.34)	(1.56)	(1.16)	(3.57)	(-0.20)	(1.99)	(1.10)	(3.06)
Total trading costs (non-optimized)	0.88	1.89	4.31	3.81	1.06	1.73	3.68	3.13
Total trading costs	0.13	0.55	1.99	1.31	0.24	0.58	1.74	1.16
Turnover (non-optimized)	0.29	0.44	1.02	0.89	0.43	0.51	1.11	0.99
Turnover	0.14	0.29	0.71	0.51	0.28	0.37	0.88	0.72
MI (non-optimized, bps)	25.66	35.41	35.07	35.53	20.31	28.36	27.64	26.45
MI (bps)	8.17	15.88	23.35	21.54	7.13	12.86	16.45	13.49
Sharpe ratio (gross, non-optimized)	0.19	0.31	0.36	0.82	0.04	0.46	0.41	0.91
Sharpe ratio (net, non-optimized)	0.10	0.17	0.08	0.37	-0.13	0.31	0.17	0.49
Sharpe ratio (gross)	0.25	0.31	0.33	0.77	-0.01	0.49	0.36	0.83
Sharpe ratio (net)	0.23	0.27	0.20	0.62	-0.04	0.44	0.24	0.67
Beta to non-optimized	0.96	0.96	0.96	0.99	1.03	0.96	0.97	0.99

Table IX
Tracking Error Frontier of Trading Costs, Net Sharpe Ratios, and Capacity

Reported are the total trading costs, net Sharpe ratios, and break-even fund sizes or capacities of size (SMB), value (HML), momentum (UMD), and an equally weighted composite portfolio of all three (COMBO). Panel A reports results for U.S. equities and Panel B for international equities, where country portfolios are aggregated using the country's total market capitalization as of the prior month using all available "tradable" stocks in the combined CRSP Xpressfeed/Compustat global data. The sample period for U.S portfolios is January 1980 to September 2013 and for International portfolios is January 1993 to September 2013. The "tradable" universe are stocks are defined as follows: at the end of each calendar month we rank all stocks based on their market capitalization and on their one-year average daily volume. We require stocks to have non-missing market capitalization, one-year daily volume, market beta, and idiosyncratic volatility. We sum the market cap and volume rank and select the top 2,000 U.S. securities and the top 2,000 International securities based on the combined rank. In the international sample we select the top *cw**2,000 securities in each country where *cw* is the country's market capitalization weight. Portfolios are optimized for trading costs by minimizing the expected total dollar trading costs subject to the various tracking error constraints relative to the unconstrained portfolio, a maximum trade constraint equal to 5% of the stock's average daily volume and a total dollar long/short notional size of $1 per side per $1 dollar of NAV. For each security, the predicted market impact MI is computed using coefficients from column (4) Table V. The starting NAV for every portfolio is 200 Million USD, as in Table IX. Trading costs are annualized, in percent. Sharpe ratios are annualized. Break-even fund sizes are computed as the dollar amount needed to make trading costs from price impact equal to the average excess return on the portfolio. Reported are break-even fund sizes (in $billions) using the sample estimate of the average excess returns, as well as the historical estimate using the longest history possible (1926 to 2013 in the U.S. and 1986 to 2013 internationally). Results are reported for optimized portfolios that allow ex ante tracking error of 0, 50, 75, 100, 150, and 200 basis points annually.

	Panel A: U.S. tradable sample, 1980 - 2013				**Panel B: International tradable sample, 1993 - 2013**			
	SMB	HML	UMD	Combo	SMB	HML	UMD	Combo
Tracking error				Total trading costs				
0	0.88	1.89	4.31	3.81	1.06	1.73	3.68	3.13
50	0.23	0.80	2.54	1.85	0.39	0.78	2.15	1.60
75	0.13	0.55	1.99	1.31	0.24	0.58	1.74	1.16
100	0.08	0.38	1.53	0.92	0.16	0.43	1.45	0.83
150	0.04	0.18	0.92	0.42	0.07	0.23	0.89	0.38
200	0.03	0.09	0.51	0.18	0.03	0.12	0.49	0.18
				Sharpe ratio (net)				
0	0.10	0.17	0.08	0.37	-0.13	0.31	0.17	0.49
50	0.19	0.24	0.18	0.58	-0.05	0.39	0.24	0.71
75	0.23	0.27	0.20	0.62	-0.04	0.44	0.24	0.67
100	0.25	0.29	0.21	0.63	-0.02	0.49	0.25	0.69
150	0.27	0.31	0.19	0.56	0.21	0.59	0.23	0.65
200	0.27	0.28	0.11	0.47	0.30	0.62	0.16	0.62
				Break-even fund size (using sample risk premium)				
Excess return =	1.76	4.05	5.60	6.90	0.24	5.18	6.18	6.86
0	102.21	153.78	26.60	54.64	0.00	79.66	12.34	21.17
50	534.70	420.17	51.66	129.85	0.00	216.54	24.61	45.77
75	635.66	658.30	71.68	261.81	0.31	336.19	30.48	81.77
100	920.89	827.52	86.92	291.73	0.31	399.11	35.55	104.53
150	1,206.12	1,026.48	75.02	332.89	4.00	493.92	34.51	135.73
200	1,364.25	1,367.30	66.91	504.29	5.54	655.73	30.81	207.08
				Break-even fund size (using historical risk premium)				
Excess return =	3.03	4.93	8.20	9.71	-0.17	5.78	7.64	7.23
0	275.52	214.28	56.16	98.69	0.00	95.48	18.87	23.40
50	1,152.37	558.16	99.91	217.54	0.00	252.96	35.68	49.88
75	1,348.19	860.83	133.82	419.06	0.00	389.54	43.59	88.44
100	1,891.60	1,074.73	159.26	464.18	0.00	461.14	50.37	112.76
150	2,435.00	1,326.00	139.60	526.00	0.00	569.06	48.99	146.13
200	2,735.00	1,755.00	126.00	780.00	0.00	752.58	44.03	222.14

Figure 1
Event-Time Average Market Impact, 1998 – 2013, U.S. Trades

This figure plots the average Market Impact (MI) in our data for all trades made within one day duration, where averages are reported for 30 min. intervals during the day and two-hour intervals overnight and for the next trading day. The data includes all available developed market equity transactions (cash equities and equity swaps) in our database between August 1998 and September 2013. Market Impact is in basis points. The shaded areas indicate 95% confidence bands.

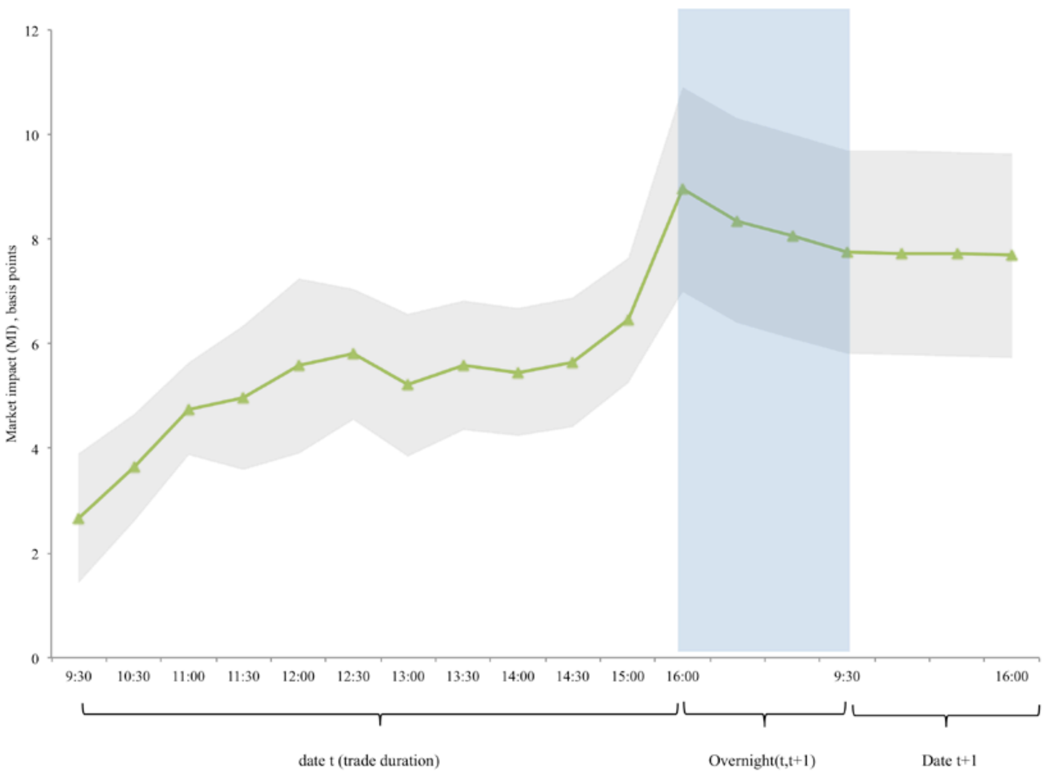

Figure 2
Market Impact by Fraction of Trading Volume

Plotted is the average market impact (MI) for actual live trades from our proprietary database. We sort all trades in our dataset into 30 bins based on their fraction of daily volume and compute average market impact for each bucket. This table includes all available developed market equity transactions (cash equities and equity swaps) in our data between August 1998 and September 2013. Market Impact is in basis points (annualized).

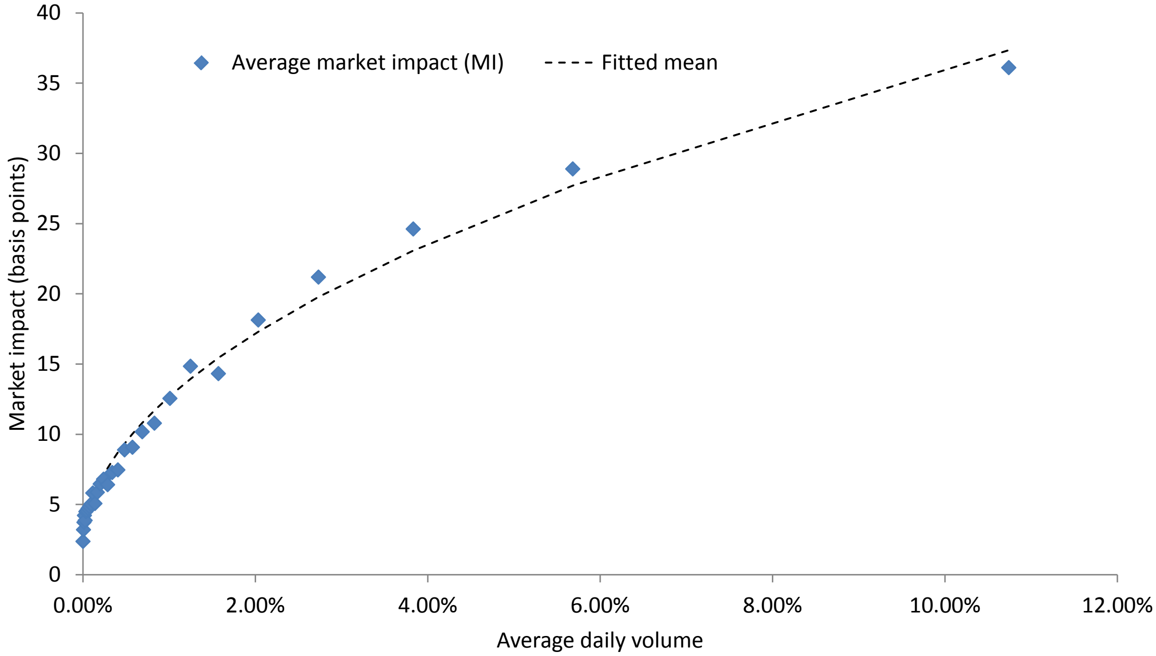

Figure 3
Comparing Market Impact Functions

Plotted is the actual market impact of live trades from our database as a function of the fraction of daily volume traded from Figure 2 versus the implied market impact from the Korajczyk and Sadka (2004) price impact function, which is calibrated to TAQ data, at the same fraction of daily volume values. The fitted mean and linear fit between the data points are also plotted. Market impact is in basis points (annualized).

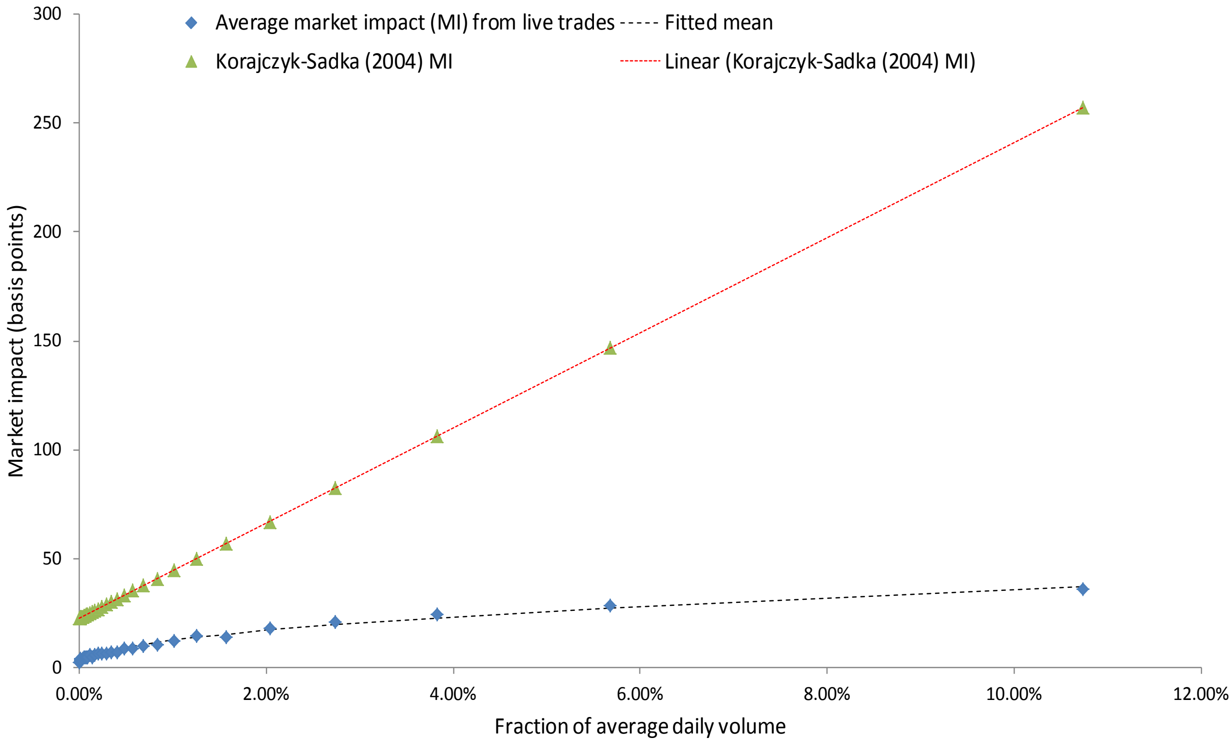

Figure 4
Tracking Error Frontier for Trading Costs, Gross Returns, and Net Returns

Plotted are the total trading costs, gross returns, and net returns to optimized portfolios based on size (SMB), value (HML), momentum (UMD), and an equally weighted composite portfolio of all three (COMBO) across various levels of tracking error. Portfolios are constructed from all available tradable stocks in CRSP for the sample period January 1980 to September 2013. The tradable universe is selected by sorting stocks at the end of each calendar month based on their market capitalization and one-year average daily volume and based on the sum of the ranks on market cap and dollar volume, we select the top 2,000 U.S. securities. Portfolios are optimized for trading costs by minimizing the expected total dollar trading costs subject to the various tracking error constraints relative to the unconstrained portfolio, a maximum trade constraint equal to 5% of the stock's average daily volume and a total dollar long/short notional size of $1 per side per $1 dollar of NAV. For each security, the predicted market impact MI is computed using coefficients from column (4) Table V. The starting NAV for every portfolio is 200 Million USD. Trading costs are annualized, in percent. Results are reported for optimized portfolios that allow ex ante tracking error of 0, 50, 100, 150, and 200 basis points annually.

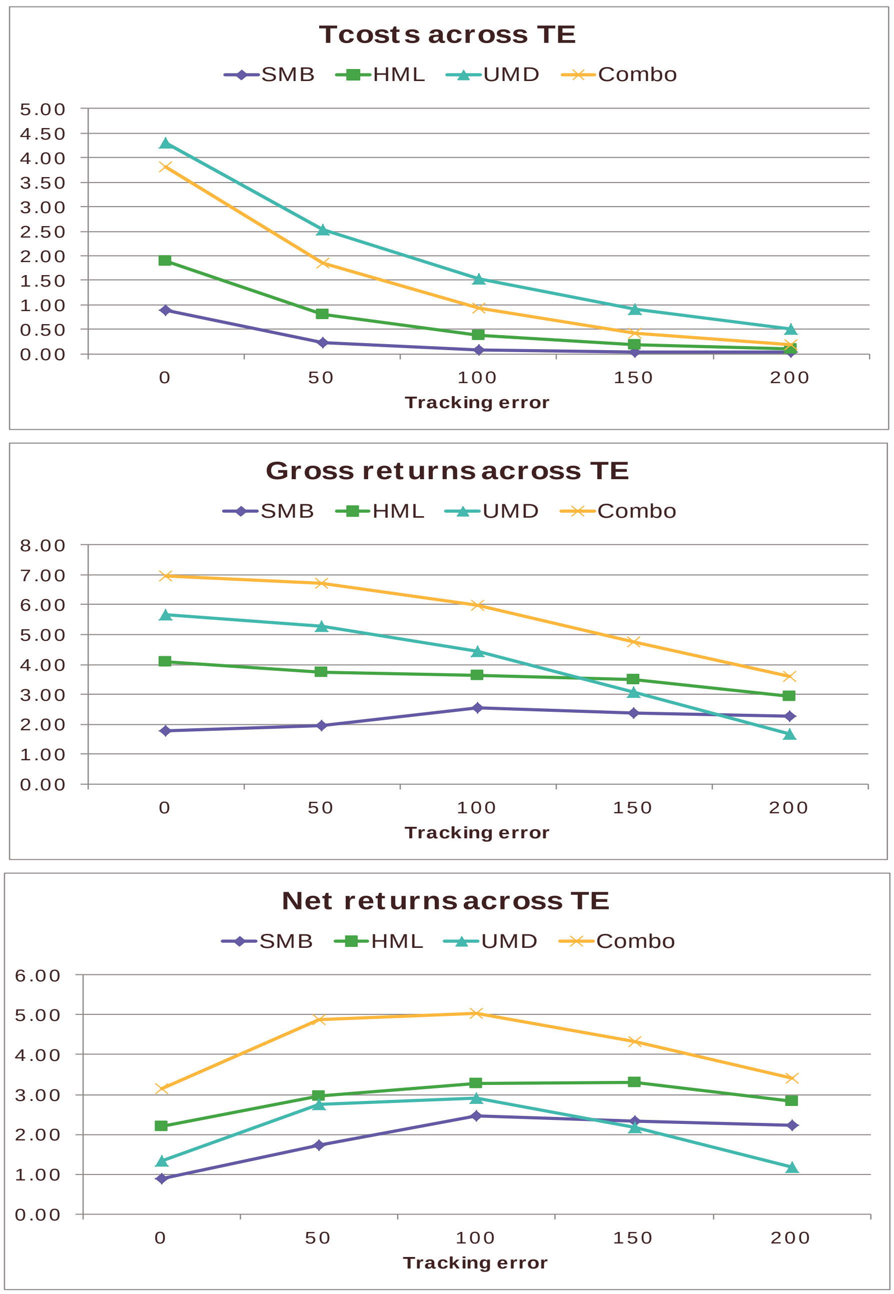

Figure 5
Tracking Error Frontier for Trading Costs at Different Fund Sizes

Plotted are the total trading costs to optimized portfolios based on size (SMB), value (HML), momentum (UMD), and an equally weighted composite portfolio of all three (COMBO) across various levels of tracking error and for various starting fund sizes. Portfolios are constructed from all available tradable stocks in CRSP for the sample period January 1980 to September 2013. The tradable universe is selected by sorting stocks at the end of each calendar month based on their market capitalization and one-year average daily volume and based on the sum of the ranks on market cap and dollar volume, we select the top 2,000 U.S. securities. Portfolios are optimized for trading costs by minimizing the expected total dollar trading costs subject to the various tracking error constraints relative to the unconstrained portfolio, a maximum trade constraint equal to 5% of the stock's average daily volume and a total dollar long/short notional size of $1 per side per $1 dollar of NAV. For each security, the predicted market impact MI is computed using coefficients from column (4) Table V. Trading costs are annualized, in percent. Results are reported for optimized portfolios that allow ex ante tracking error of 0, 50, 100, 150, and 200 basis points annually across various starting NAV sizes of $100mm, $200mm, $500mm, $1 billion, $2 billion, and $5 billion. The starting NAV values are for 1980 and grow according to each portfolio's return path until 2013. The figures report the ending NAV for each strategy at each size at the end of the sample period in 2013.

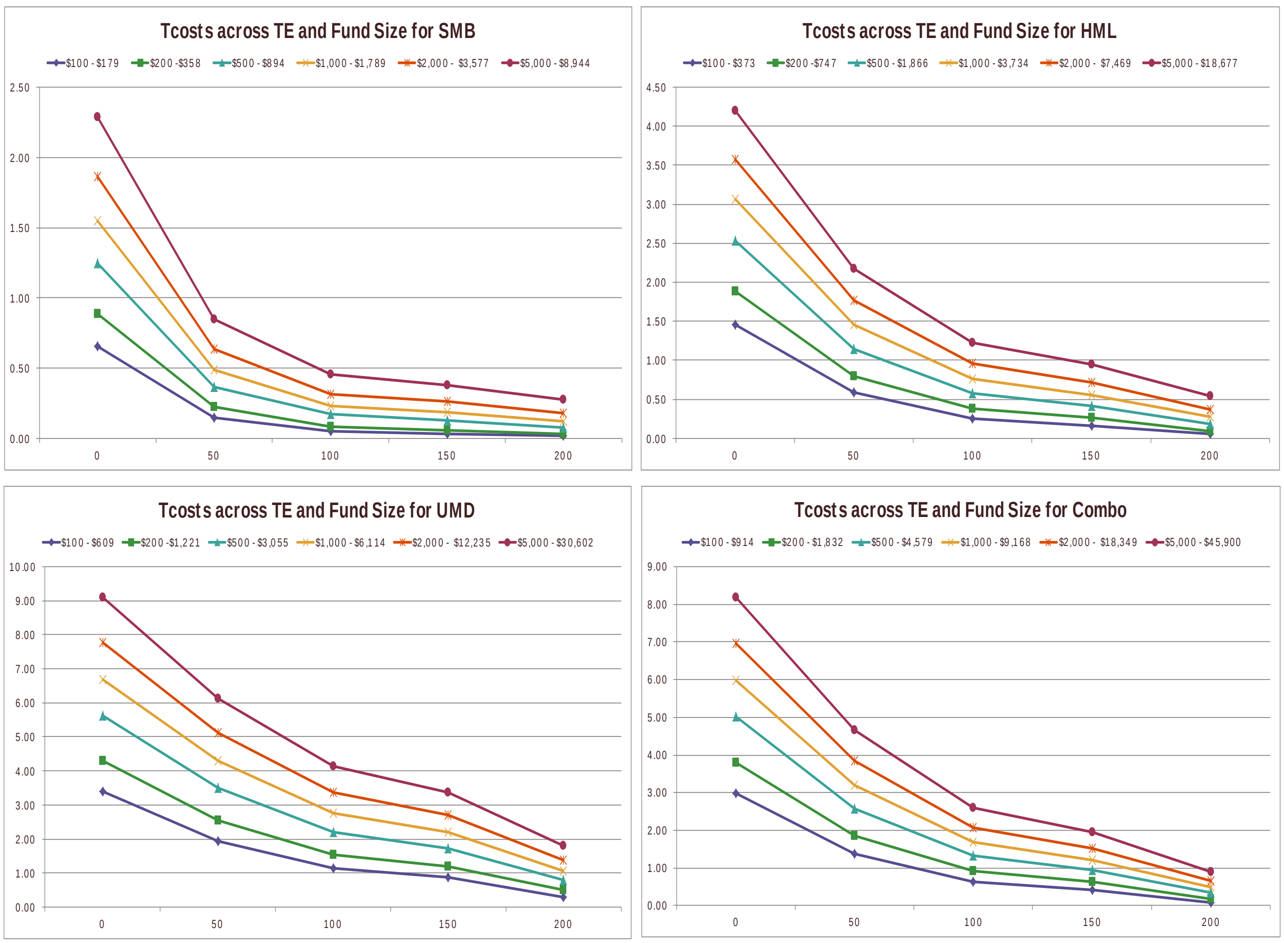

Figure 6
Tracking Error Frontier of Momentum Break-Even Capacity

The top figure plots the break-even fund sizes of momentum portfolios (UMD) among U.S. equities, and the bottom figure plots for international equities, optimized for trading costs for different levels of tracking error ranging from 0, 50, 75, 100, 150, and 200 basis points annually. The break-even fund sizes are computed using the price impact function estimated in column (4) of Table V and portfolios are optimized for trading costs by minimizing the expected total dollar trading costs subject to the various tracking error constraints relative to the unconstrained portfolio, a maximum trade constraint equal to 5% of the stock's average daily volume and a total dollar long/short notional size of $1 per side per $1 dollar of NAV. Break-even fund sizes (in $billions) are plotted using the sample estimate of the average excess returns (1980 to 2013 for the U.S. and 1993 to 2013 internationally), as well as the historical estimate using the longest history possible (1926 to 2013 in the U.S. and 1986 to 2013 internationally).