Flow-Based Asset Pricing: Two Trading Desk Separation Theorem*

Yu An[†] Yinan Su[‡] Chen Wang[§]

This version: May 12, 2022

First version: May 2, 2022

Abstract

In any market with uninformative flows, the maximum Sharpe-ratio portfolio can be separated into two. The first portfolio uses only fundamental information to maximize Sharpe ratio. The second portfolio provides liquidity to uninformative flows and maximizes price impact ratio, which is defined as a portfolio's price impact over its fundamental risk. We develop the factor model of price impacts to empirically investigate the maximum-price-impact-ratio (MPIR) portfolio. For U.S. equity mutual fund flows, we find that the MPIR portfolio constructed using flows into Fama and French (1993) factors is a good choice.

Keywords: asset pricing, cross section, flow, price impact, risk

JEL Codes: G12

*We thank Federico Bandi, Hank Bessembinder, Robin Greenwood, Paolo Pasquariello, Nagpurnanand Prabhala, Dongho Song, Zhaogang Song, Semih Üslü, Adrien Verdelhan, and Zeyu Zheng for helpful comments. All errors are our own.

[†]Carey Business School, Johns Hopkins University; yua@jhu.edu; https://sites.google.com/view/yu-an.

[‡]Carey Business School, Johns Hopkins University; ys@jhu.edu; https://www.suyinan.com/.

[§]Mendoza College of Business, University of Notre Dame; chen.wang@nd.edu; https://chenwang.one/.

1 Introduction

We price assets in a market with uninformative flows. Price impacts arise because liquidity providers (e.g., banks, non-bank dealers, and hedge funds) are averse to absorbing fundamental risks induced by flows.[1]

Two trading desk separation theorem. We prove that in any market with uninformative flows, the maximum Sharpe-ratio portfolio can be separated into two. The first portfolio uses only fundamental information to maximize the Sharpe ratio. A fundamental-investing desk that conducts Fama-French-style analysis can obtain this portfolio. The second portfolio maximizes the price impact ratio, which is defined as a portfolio's price impact over its fundamental risk. We term this portfolio the *maximum-price-impact-ratio (MPIR)* portfolio. A liquidity-provision desk that trades against uninformative flows can obtain this portfolio. The two desks operate *separately* to maximize their own objective. Importantly, the correct performance measure for liquidity provision is *price impact ratio, not Sharpe ratio*. By maximizing price impact ratio, liquidity provision complements fundamental investing and maximizes the overall Sharpe ratio. The Sharpe-ratio-maximizing investor allocates risk between the two desks, proportional to their respective Sharpe ratio and price impact ratio. This theorem is a mathematical identity under extremely general and simple setting, irrespective of any distribution assumptions or factor-model and equilibrium structures.

The theorem implies that any Sharpe-ratio-maximizing investor should simultaneously be a fundamental investor and a liquidity provider. Excessive specialization on one role may reflect Sharpe ratio left on the table. For example, Koijen and Yogo (2019) find that institutions vary drastically in demand elasticity, suggesting their tendency to specialize.

Relatedly, Duffie (2012) argues that it is hard in practice to eliminate the proprietary-

[1] Liquidity providers' risk aversion can arise from rational or institutional reasons. See Shleifer and Vishny (1997), Duffie (2012), He and Krishnamurthy (2013), Adrian, Etula, and Muir (2014), Brunnermeier and Sannikov (2014), He, Kelly, and Manela (2017), Bessembinder, Jacobsen, Maxwell, and Venkataraman (2018), Du, Tepper, and Verdelhan (2018), Andersen, Duffie, and Song (2019), Anderson, Du, and Schlusche (2021), Du, Hébert, and Huber (2021), He, Khorrami, and Song (2021), among others.

trading incentive of market-making desks at bank dealers, which is an objective of the Volcker rule. The theorem offers one potential solution by changing the performance evaluation measure of the market-making desks from the Sharpe ratio to price impact ratio.

Estimating MPIR portfolio. The theorem has strong empirical implications. Instead of directly searching for the maximum Sharpe-ratio portfolio in a market with uninformative flows, researchers can *separately* solve two easier problems. The first problem searches for the maximum Sharpe-ratio portfolio using only fundamental information. Since the seminal work of Fama and MacBeth (1973) and Gibbons, Ross, and Shanken (1989), numerous papers have made progress on this problem using the factor model of expected returns. The second problem searches for the MPIR portfolio. We are the first. The MPIR portfolio is economically different from the maximum Sharpe-ratio portfolio. We therefore develop the new factor model of price impacts to empirically identify the MPIR portfolio.[2] We develop the flow-based Fama-MacBeth regression to estimate the MPIR portfolio and the flow-based Gibbons-Ross-Shanken test to the diagnose MPIR portfolio. For U.S. equity mutual fund flows, we find that the MPIR portfolio constructed using flows into Fama and French (1993) factors is a good choice.

We do not claim that we find the uniquely true MPIR portfolio. Being the first to empirically investigate the MPIR portfolio means that we likely missed the true one. We do not claim that our factor-model approach is the only way of getting the MPIR. Factor model is also not the only way of getting the maximum Sharpe ratio, but is arguably the most commonly used and impactful approach.

Estimating MPIR portfolio solves the problem faced by the liquidity-provision desk of the rational investor. To simplify terminology, we call this desk *liquidity provider* hereafter.

For empirical measures of uninformative flows, we use mutual fund flow-induced trading.[3] The idea is that mutual fund investors are mostly uninformed retail investors. Their decision

[2] An (2022) derives factor model of price impacts by generalizing Ross (1976) arbitrage pricing theory.

[3] Mutual fund flow-induced trading has been used by various studies including Coval and Stafford (2007), Greenwood and Thesmar (2011), Lou (2012), Ben-David, Li, Rossi, and Song (2021), and Li (2021).

to buy or sell mutual fund shares reveals little information about the future fundamentals of underlying stocks. Experiencing retail investor flows, mutual funds tend to buy or sell stocks in proportion to their existing holdings, generating stock-level uninformative flows. We use monthly U.S. equity mutual fund flows data from 2000 through 2020. Our test assets are Fama and French (1993) 5×5 portfolios sorted on size and book-to-market equity.

In the first stage (1), we run a time-series regression of each test asset's flow $f_{n,t}$ on market (MKT), small-minus-big (SMB), and high-minus-low (HML) factor flows $q_{k,t}$,

$$\underbrace{f_{n,t}}_{\text{asset flow}} = \sum_{k\in\{\text{MKT,SMB,HML}\}} \underbrace{b_{n,k}}_{\text{flow beta}} \underbrace{q_{k,t}}_{\text{factor flow}} + \underbrace{e_{n,t}}_{\text{idiosyncratic flow}} . \tag{1}$$

Flow betas are *mimicking portfolio weights*. The 25 flow betas of MKT factor $b_{n,\text{MKT}}$ represent how \$1 flow into MKT factor drives the 25 asset flows. Flow betas $b_{n,\text{MKT}}$ are the portfolio that liquidity providers in equilibrium trade against for every dollar of MKT flow.

Importantly, our flow betas or mimicking portfolio weights are not necessarily the same as the Fama-French portfolio weights. Therefore, our regression (1) concerning flow betas neither replicates nor can be derived from the usual Fama-French regression. Whether our flow betas $b_{n,k}$ correctly recover the Fama-French portfolio weights is a meaningful economic test. We show that our recovery works well only if the proposed factor flows $q_{k,t}$ match the principal components of cross-sectional asset flows $f_{n,t}$.

Turning to results of (1), the regression R^2 ranges from 50% for small companies to 80% for large companies, showing that flows into MKT, SMB, and HML factors explain large portions of common variation in asset flows. The 25 MKT-factor flow betas $b_{n,\text{MKT}}$ are remarkably close to the fraction of each asset's market capitalization to total stock market capitalization. That is, MKT flow betas almost perfectly recover Fama-French portfolio weights, implying that MKT flow does explain cross-sectional asset flows. The SMB flow betas $b_{n,\text{SMB}}$ are positive for small companies and negative for large companies. The sum of all positive $b_{n,\text{SMB}}$ is close to 1 and the sum of all negative $b_{n,\text{SMB}}$ is close to -1, consistent with Fama-French SMB portfolio weights. SMB flow also explains cross-sectional asset flows.

For HML flow betas $b_{n,\text{HML}}$, the recovery works less well. Only the BL and B2 test assets (Biggest companies with Lowest and 2nd lowest book-to-market equity) have negative flow betas, and all other assets have positive flow betas. This result is reasonable because BL and B2 assets combined have over 50% of the total stock market capitalization. Nevertheless, our flow betas differ from Fama-French HML portfolio weights, suggesting that HML flow may not perfectly explain cross-sectional asset flows. We stress that this is not a failure of Fama and French (1993)—they never intended to explain cross-sectional flows.

In the second stage (2), we run a panel regression of asset returns $r_{n,t}$ on the product of factor flows $q_{l,t}$ and quantity of risk $\text{cov}(\mathbf{R}_0(n), \mathbf{b}_k^\top \mathbf{R}_0)$,

$$\underbrace{r_{n,t}}_{\text{asset return}} = \sum_{k,l \in \{\text{MKT,SMB,HML}\}} \underbrace{\lambda_{k,l}}_{\text{price of risk}} \underbrace{q_{l,t}}_{\text{factor flow}} \underbrace{\text{cov}(\mathbf{R}_0(n), \mathbf{b}_k^\top \mathbf{R}_0)}_{\text{quantity of risk}} + \xi_{n,t}. \tag{2}$$

Importantly, we use flow betas to form mimicking portfolio $\mathbf{b}_k = (b_{1,k}, b_{2,k}, \ldots, b_{25,k})^\top$ for each factor k. In equilibrium, liquidity providers trade against this portfolio $\mathbf{b}_k$ for every dollar of factor flow $q_{k,t}$. We use one-year return[4] to proxy the fundamental return $\mathbf{R}_0$. For each asset n, the quantity of risk induced by every dollar of factor-k flow is the covariance between the asset's fundamental return $\mathbf{R}_0(n)$ and the factor's $\mathbf{b}_k^\top \mathbf{R}_0$.

The quantity of risk differentiates our regression (2) from the usual price-impact regressions, which directly regress returns on flows. Economically, liquidity providers are averse to the fundamental *risks* induced by flows, not flows per se. Our regression obtains the price of risk $\lambda_{k,l}$, measuring whether factor-l flow impacts factor-k price. We impose a theory-founded structural restriction when estimating (2), such that the 3×3 prices of risk $\lambda_{k,l}$ have 3 degrees of freedom. Economically, 3 factors have 3 prices of risk, not 9 prices of risk.

Turning to results of (2), we find that the fundamental risk induced by SMB-factor flows has the greatest explanatory power for the cross section of price impacts, while the MKT and HML factors have roughly the same explanatory power. This fact may be surprising

[4] One-year returns are the most commonly used long-horizon returns in the asset pricing literature. See Koijen and van Nieuwerburgh (2011) for a review.

because MKT flow is about six times as volatile as SMB and HML flows. This is because the SMB and HML factors have a much higher price of risk than the MKT factor, and the quantity of risk induced by every dollar of flow differs. Not all flows are created equal—the price of risk and the quantity of risk matter for the cross section of price impacts.

Diagnosing MPIR portfolio. The second-stage regression (2) implies a MPIR portfolio at the end of each month t, based on the estimated prices of risk $\lambda_{k,l}$ and flow betas $\mathbf{b}_k$ and the observed factor flows $q_{k,t}$ in month t. Importantly, the second stage (2) and the constructed MPIR portfolio do not depend on idiosyncratic flows $e_{n,t}$, the residual from first-stage regression (1). Idiosyncratic flows are the portion of asset flows $f_{n,t}$ that cannot be explained by factor flows $q_{k,t}$. What is the relationship between the second stage (2) and the corresponding MPIR portfolio?

We theoretically answer this question by generalizing Gibbons, Ross, and Shanken (1989). Under the true second-stage model (2), any asset's idiosyncratic flow should not have *anomaly* price impacts on any asset. We prove that the joint χ^2 test statistic for anomaly price impacts equals the difference between factor-plus-idiosyncratic-flow-implied and factor-flow-implied squared MPIR. That is, diagnosing second-stage model (2) is the same thing as diagnosing the model-implied MPIR portfolio.

Turning to empirical results, the price dislocation of our factor-model-implied MPIR portfolio on average reverts to zero in six months. This reversion suggests that the MPIR portfolio constructed using Fama and French (1993) factor flows likely reflects liquidity providers' aversion to absorbing fundamental risks. We then add the anomaly impacts caused by idiosyncratic flows to our factor model and construct the new MPIR portfolio. The price dislocation of the anomaly MPIR portfolio barely reverts even after one year. This evidence suggests that the anomaly price impacts caused by idiosyncratic flows are not due to liquidity providers' risk aversion. The sharp contrast in the price reversion between factor and idiosyncratic flows supports our factor-model approach to identifying the MPIR portfolio.

Related literature. From a theoretical perspective, the two trading desk separation theorem connects the Fama-French-style asset pricing that studies expected returns with demand-system asset pricing (Koijen and Yogo (2019)) and microstructure literature that studies price impacts.[5] Although we focus on price impacts, our objective differs significantly from the existing literature. We estimate and diagnose the MPIR portfolio using factor model, rather than estimating any particular price impacts. For the purpose of getting the maximum Sharpe ratio in markets with uninformative flows, it suffices to estimate *only* the MPIR portfolio and then conduct Fama-French fundemental analysis.

From a methodological perspective, we show that the classic factor models can be applied to study not only expected returns but also price impacts. Figure 1 presents our paper's roadmap and analogy to standard asset pricing (Fama and MacBeth (1973); Gibbons, Ross, and Shanken (1989); Fama and French (1993)). Most importantly, flows (changes in quantity) are not just one type of non-tradable returns (changes in price). *Flows reflect liquidity providers' equilibrium portfolio changes.* This fundamental economic distinction motivates us to develop the new factor-model framework to study price impacts and MPIR portfolio. Modern approaches to factor models, especially those applying machine-learning techniques, could potentially be applied to our factor model of price impacts. See Cochrane (2009), Giglio, Kelly, and Xiu (2021), and Nagel (2021) for overviews.[6] Our work relates to Kozak, Nagel, and Santosh (2018). Our factor model allows for considerable empirical flexibility beyond the CARA-Gaussian price impact model, and nests it as a special case.[7]

[5]For related demand-system asset pricing and microstructure literature, see, for example, Hasbrouck (1988), Hasbrouck (1991), Lo and Wang (2000), Hasbrouck and Seppi (2001), Greenwood (2005), Coval and Stafford (2007), Cremers and Mei (2007), Andrade, Chang, and Seasholes (2008), Greenwood and Thesmar (2011), Lou (2012), Boulatov, Hendershott, and Livdan (2013), Chang, Hong, and Liskovich (2015), Pasquariello and Vega (2015), Bouchaud, Bonart, Donier, and Gould (2018), Da, Larrain, Sialm, and Tessada (2018), Frazzini, Israel, and Moskowitz (2018), Li, Pearson, and Zhang (2020), Schmickler (2020), Ben-David, Li, Rossi, and Song (2021), Huang, Song, and Xiang (2021), Li (2021), Koijen, Richmond, and Yogo (2021), Peng and Wang (2021), Gabaix and Koijen (2022), Haddad, Huebner, and Loualiche (2022), Li and Lin (2022), and Pavlova and Sikorskaya (2022).

[6]Notable papers include Daniel and Titman (1997), Ang, Hodrick, Xing, and Zhang (2006), Lewellen and Nagel (2006), Bali and Cakici (2008), Rapach, Strauss, and Zhou (2010), Asness, Moskowitz, and Pedersen (2013), Frazzini and Pedersen (2014), Harvey, Liu, and Zhu (2016), Stambaugh and Yuan (2017), Linnainmaa and Roberts (2018), Kozak, Nagel, and Santosh (2018), Kelly, Pruitt, and Su (2019), Hou, Xue, and Zhang (2020), Lettau and Pelger (2020), Bryzgalova, Pelger, and Zhu (2021), Chen (2021), Dou, Kogan, and Wu (2021), Holcblat, Lioui, and Weber (2022), and Novy-Marx and Velikov (2022).

[7]Compared with our model, the CARA-Gaussian model additionally requires that all factors must have

Figure 1. Flow-based asset pricing and analogy with standard asset pricing

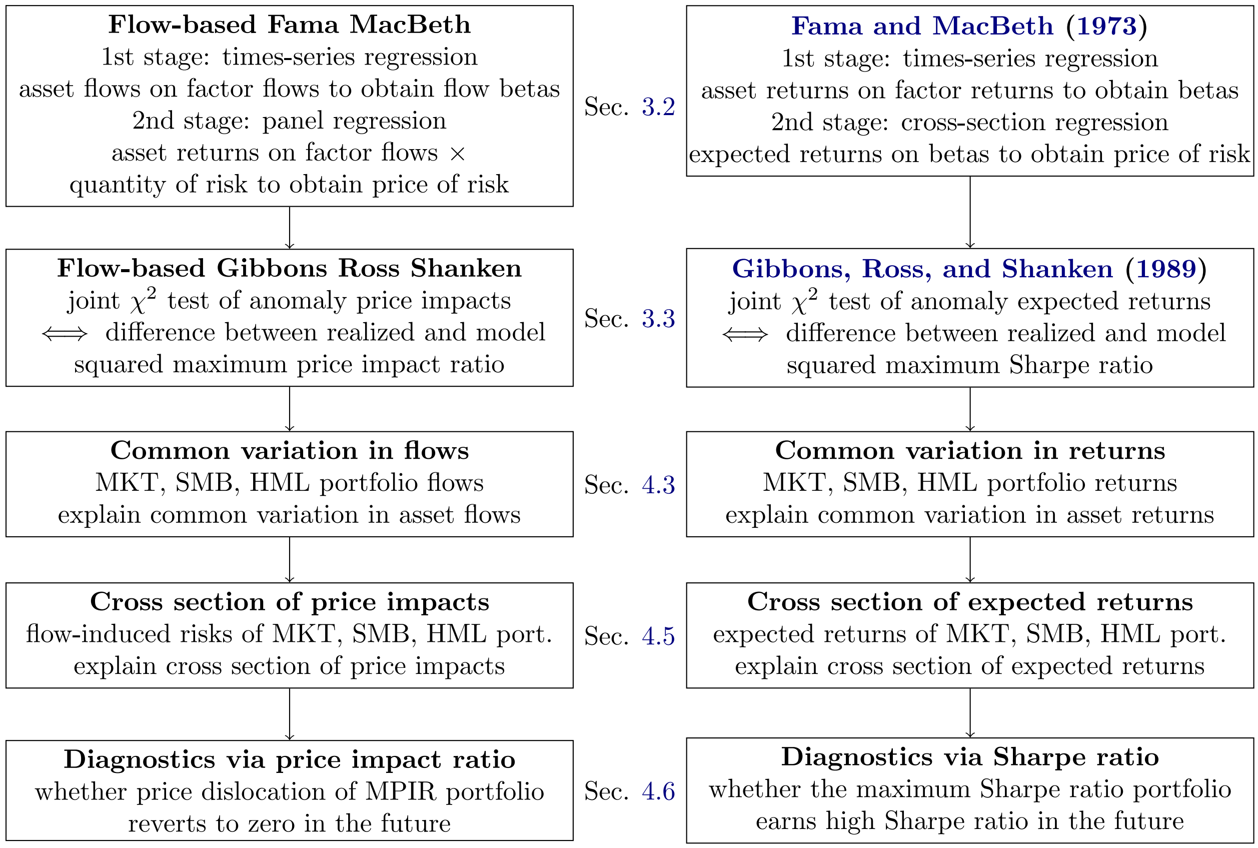

From an empirical perspective, we show that for U.S. equity mutual fund flows, the MPIR portfolio constructed using flows into Fama and French (1993) factors is a good choice. Our empirical exercise potentially suffers from the joint-hypothesis-testing issue of both uninformativeness of flows and risk aversion of liquidity providers. Thus, having a good instrument for uninformative flows is particularly important, and we borrow from the widely used mutual fund flow-induced trading instrument of Coval and Stafford (2007) and Lou (2012). More systematic ways of identifying uninformative flows by exploring detailed institutional and household holdings[8] could potentially be applied to our setting, including the investment-universe instrument of Koijen and Yogo (2019) and the granular instrument variables of Gabaix and Koijen (2021).

the same price of risk. We reject this null hypothesis in the data.

[8]Papers that exploit household holdings include Balasubramaniam, Campbell, Ramadorai, and Ranish (2021), Betermier, Calvet, Knüpfer, and Kvaerner (2021), and Egan, MacKay, and Yang (2021).

2 Two Trading Desk Separation Theorem

As a convention, we use bold font notation for matrices and vectors. We use $\mathbf{A}^\top$ to denote the transpose of matrix $\mathbf{A}$. We use $\mathbb{E}[\mathbf{b}]$ and $\mathrm{var}(\mathbf{b})$ to denote the expectation and variance-covariance matrix of a vector of random variables $\mathbf{b}$.

Asset $n = 1, 2, \ldots, N$ has a flow f_n at time 0, where $\mathbf{f} = (f_1, f_2, \ldots, f_N)^\top$. Asset n has a payoff X_n at time 1, where $\mathbf{X} = (X_1, X_2, \ldots, X_N)^\top$. Without loss of generality, we assume that $\mathrm{var}(\mathbf{X})$ has full rank.[9] Flows are uninformative about payoffs. The risk-free rate between time 0 and 1 is a constant R_F. Under flow $\mathbf{f}$, we denote time-0 price of asset n as $P_n(\mathbf{f})$. Under flow $\mathbf{f}$, asset returns are defined as

$$\mathbf{R}(\mathbf{f}) = \left(\frac{X_1}{P_1(\mathbf{f})}, \frac{X_2}{P_2(\mathbf{f})}, \ldots, \frac{X_N}{P_N(\mathbf{f})} \right)^\top. \tag{3}$$

The fundamental returns $\mathbf{R}(\mathbf{0})$ are asset returns when there are no flows. Price impacts are defined as the percentage price change with and without flow $\mathbf{f}$,

$$\Delta\mathbf{p}(\mathbf{f}) = \left(\frac{P_1(\mathbf{f}) - P_1(\mathbf{0})}{P_1(\mathbf{0})}, \frac{P_2(\mathbf{f}) - P_2(\mathbf{0})}{P_2(\mathbf{0})}, \ldots, \frac{P_N(\mathbf{f}) - P_N(\mathbf{0})}{P_N(\mathbf{0})} \right)^\top. \tag{4}$$

The maximum-Sharpe-ratio portfolio under flow $\mathbf{f}$ is defined as

$$\mathbf{c}^*(\mathbf{f}) \equiv \mathrm{var}(\mathbf{R}(\mathbf{f}))^{-1}\mathbb{E}[\mathbf{R}(\mathbf{f}) - R_F\mathbf{1}] = \arg\max_{\mathbf{c}\in\mathbb{R}^N} \frac{\mathbb{E}[\mathbf{c}^\top(\mathbf{R}(\mathbf{f}) - R_F\mathbf{1})]}{\sigma(\mathbf{c}^\top\mathbf{R}(\mathbf{f}))}. \tag{5}$$

The *maximum-price-impact-ratio (MPIR)* portfolio under flow $\mathbf{f}$ is defined as

$$\tilde{\mathbf{c}}^*(\mathbf{f}) \equiv \mathrm{var}(\mathbf{R}(\mathbf{0}))^{-1}\Delta\mathbf{p}(\mathbf{f}) = \arg\max_{\mathbf{c}\in\mathbb{R}^N} \frac{\mathbf{c}^\top\Delta\mathbf{p}(\mathbf{f})}{\sigma(\mathbf{c}^\top\mathbf{R}(\mathbf{0}))}, \tag{6}$$

which has the largest price impact $\mathbf{c}^\top\Delta\mathbf{p}(\mathbf{f})$ relative to its fundamental risk $\sigma(\mathbf{c}^\top\mathbf{R}(\mathbf{0}))$. We

[9] Otherwise, one can choose lower-dimensional linearly independent portfolios, and rotate the payoff $\mathbf{X}$ and flow $\mathbf{f}$ to those portfolios. See Section 2.2 of An (2022) for details on constructing portfolio flows.

translate dollar-amount portfolio weights $\mathbf{c}^*(\mathbf{f})$ and $\tilde{\mathbf{c}}^*(\mathbf{f})$ to the unit of shares,[10]

$$\mathbf{h}^*(\mathbf{f}) = \left(\frac{c_1^*(\mathbf{f})}{P_1(\mathbf{f})}, \frac{c_2^*(\mathbf{f})}{P_2(\mathbf{f})}, \ldots, \frac{c_N^*(\mathbf{f})}{P_N(\mathbf{f})}\right)^\top, \tag{7}$$

$$\tilde{\mathbf{h}}^*(\mathbf{f}) = \left(\frac{\tilde{c}_1^*(\mathbf{f})}{P_1(\mathbf{0})}, \frac{\tilde{c}_2^*(\mathbf{f})}{P_2(\mathbf{0})}, \ldots, \frac{\tilde{c}_N^*(\mathbf{f})}{P_N(\mathbf{0})}\right)^\top. \tag{8}$$

We stress that the model does not depend on any distribution assumptions or factor-model and equilibrium structures. Appendix A.1 proves the following theorem.

THEOREM 1 (Two trading desk separation). *We have*

$$\underbrace{\mathbf{h}^*(\mathbf{f})}_{\text{max. Sharpe ratio port. with flow}} = \overbrace{\mathbf{h}^*(\mathbf{0})}^{\text{max. Sharpe ratio port. without flow}} - R_F \underbrace{\tilde{\mathbf{h}}^*(\mathbf{f})}_{\text{max. price impact ratio port.}}. \tag{9}$$

The return volatility of portfolio $\mathbf{h}^*(\mathbf{0})$ *equals the maximum Sharpe ratio without flow. The return volatility of portfolio* $\tilde{\mathbf{h}}^*(\mathbf{f})$ *equals the maximum price impact ratio.*

Equation (9) shows that the maximum Sharpe ratio portfolio $\mathbf{h}^*(\mathbf{f})$ under flow $\mathbf{f}$ can be separated into two. The first portfolio $\mathbf{h}^*(\mathbf{0})$ maximizes the Sharpe ratio in the same economy but without flow or, equivalently, maximizes the fundamental-return-driven Sharpe ratio. A fundamental-investing desk that conducts Fama-French-style analysis can obtain this portfolio. The second portfolio maximizes the price impact ratio under flow $\mathbf{f}$. A liquidity-provision desk that trades against uninformative flows can obtain this portfolio, also explaining the negative sign in (9). The two desks operate *separately* to maximize their own objective. Importantly, the correct performance measure for liquidity provision is *price impact ratio, not Sharpe ratio.* By maximizing the price impact ratio, liquidity provision complements fundamental investing and maximizes the overall Sharpe ratio.

The Sharpe-ratio-maximizing investor allocates risk between the two separate desks. Because of diversification benefits, the amount of allocated risk is roughly proportional to the

[10]In equation (8), the denominator is asset prices $P_n(\mathbf{0})$ without flows. This is because the definition of MPIR in (6) is also based on fundamental returns $\mathbf{R}(\mathbf{0})$.

respective Sharpe ratio and price impact ratio. This is an approximation because of the multiplication by the risk-free rate R_F in (9). The precise risk allocation is proportional to the volatility of the SDF $M(\mathbf{0})$ in the economy without flow and the volatility of the *change* of the SDF $M(\mathbf{f}) - M(\mathbf{0})$ with and without flow.[11] In any reasonable model, price impact ratio is increasing in flow amounts.[12] In periods with larger uninformative flows, the investor optimally allocates greater risk exposure to liquidity provision; in periods with smaller flows, the investor allocates greater risk exposure to fundamental investing.

How is it possible that we obtain Theorem 1 with seemingly no assumptions? The trick is the setup itself. The economies with and without flows correspond to two states of world at time 0, and the fundamental payoffs realize at time 1. We do so because the short-term price impacts created by uninformative flows take a long time to revert until fundamental payoffs realize. Modeling flows as two states of world at time 0, instead of adding an extra time period, reveals the important short-term versus long-term economics.

3 Empirical Methods

In this section, we develop the flow-based Fama and MacBeth (1973) regression to estimate the linear factor model of price impacts and the flow-based Gibbons, Ross, and Shanken (1989) test to diagnose the model.

3.1 An Overview of Linear Factor Model of Price Impacts

We provide a brief overview of the linear factor model of price impacts. An (2022) presents the full theoretical foundation.

The setup is exactly the same as that in Section 2, but with additional structures. We study only unexpected flows, so requires $\mathbb{E}(\mathbf{f}) = \mathbf{0}$. The model has K factors. Each factor k is a portfolio of N assets, with weight $b_{n,k}$ in asset n. We write an $N \times 1$ vector $\mathbf{b}_k =$

[11] Hansen and Jagannathan (1991) prove that the volatility of the SDF $M(\mathbf{0})$ equals the maximum Sharpe ratio over R_F. An (2022) proves that the volatility of the change of the SDF $M(\mathbf{f}) - M(\mathbf{0})$ equals the MPIR.

[12] For example, in any linear price impact model, price impact ratio is proportional to flow amounts.

$(b_{1,k}, b_{2,k}, \dots, b_{N,k})^\top$ and an $N \times K$ matrix $\mathbf{B} = (b_1, b_2, \dots, b_K)$. The flow into factor k is q_k, and we write a $K \times 1$ vector $\mathbf{q} = (q_1, q_2, \dots, q_K)^\top$.

The N-asset K-factor price impact model is

$$\Delta \mathbf{p}(\mathbf{f}) = \sum_{k=1}^{K} \sum_{l=1}^{K} \lambda_{k,l} q_l \text{cov}(\mathbf{R}_0, \mathbf{b}_k^\top \mathbf{R}_0) = \text{var}(\mathbf{R}_0) \mathbf{B} \mathbf{\Lambda} \mathbf{q}. \tag{10}$$

The $K \times K$ matrix $\mathbf{\Lambda} = \{\lambda_{k,l}\}$ is the price of risk. The term $\lambda_{k,l}$ measures whether flow q_l into portfolio l impacts the price of portfolio k. The quantity of risk is $\text{cov}(\mathbf{R}_0, \mathbf{b}_k^\top \mathbf{R}_0)$.

The $K \times K$ prices of risk $\mathbf{\Lambda} = \{\lambda_{k,l}\}$ are not K^2 free parameters, but instead they have only K degrees of freedom that satisfy the following orthogonalization. Specifically, there exists some $K \times K$ invertible matrix $\mathbf{O}$ such that

$$\mathbf{O}^{-1} \mathbf{\Lambda} \mathbf{O} \text{ is a diagonal matrix } \tilde{\mathbf{\Lambda}} = \text{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_K), \tag{11}$$

$$\text{var}(\mathbf{O}^\top \mathbf{B}^\top \mathbf{R}_0) \text{ is an identity matrix } \mathbf{I}_K, \tag{12}$$

$$\text{var}(\mathbf{O}^{-1} \mathbf{q}) \text{ is a diagonal matrix } \mathbf{\Pi} = \text{diag}(\pi_1, \pi_2, \dots, \pi_K). \tag{13}$$

This orthogonalization requirement is called Irrelevance of Uncorrelated Flows (IUF).[13] Intuitively, K factors should have K prices of risk, not K^2 prices of risk. The K original factors may have correlated fundamental returns (i.e., $\text{var}(\mathbf{B}^\top \mathbf{R}_0)$ is not diagonal) and correlated flows (i.e., $\text{var}(\mathbf{q})$ is not diagonal). To understand cross-asset impacts, we orthogonalize the K factors by finding some rotation matrix $\mathbf{O}$ and define

$$\tilde{\mathbf{B}} \equiv \mathbf{B}\mathbf{O} \quad \text{and} \quad \tilde{\mathbf{q}} \equiv \mathbf{O}^{-1} \mathbf{q}. \tag{14}$$

The K orthogonalized factors have portfolio weights $\tilde{\mathbf{B}} = (\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2, \dots, \tilde{\mathbf{b}}_K)$ and flows $\tilde{\mathbf{q}} = (\tilde{q}_1, \tilde{q}_2, \dots, \tilde{q}_K)^\top$. These orthogonalized factors have uncorrelated and unit fundamental risk

[13] For readers who are familiar with the notations of An (2022), we have switched to use notations with tilde to represent orthogonalized model and notations without tilde to represent the original model. We do this because our paper refers to the original model much more frequently than the orthogonalized model.

Table 1. Flow-based Fama-MacBeth regression

Data pre-processing	Panel data for N asset returns and flows $\{r_{n,t}\}$ and $\{f_{n,t}\}$, panel data for K factor flows $\{q_{k,t}\}$, and estimated fundamental risk $\mathrm{var}(\mathbf{R}_0)$
First-stage regression	Time-series regression of asset flows on factor flows to obtain $b_{n,k}$
Orthogonalization	Find rotation $\mathbf{O}$ to obtain $\tilde{q}_{k,t}$ and $\tilde{b}_{n,k}$ for the orthogonalized model
Second-stage regression	Panel regression of orthogonalized $\Delta\mathbf{p}(\mathbf{f}) = \sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_k \mathrm{cov}(\mathbf{R}_0, \tilde{\mathbf{b}}_k^\top \mathbf{R}_0)$ to obtain the price of risk $\tilde{\mathbf{\Lambda}} = \mathrm{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \ldots, \tilde{\lambda}_K)$
Rotation to original model	Obtain the price of risk $\mathbf{\Lambda} = \{\lambda_{k,l}\}$ for original model $\Delta\mathbf{p}(\mathbf{f}) = \sum_{k=1}^{K} \sum_{l=1}^{K} \lambda_{k,l} q_l \mathrm{cov}(\mathbf{R}_0, \mathbf{b}_k^\top \mathbf{R}_0)$

(i.e., $\mathrm{var}(\tilde{\mathbf{B}}^\top \mathbf{R}_0) = \mathbf{I}_K$ by (12)) and uncorrelated flows (i.e., $\mathrm{var}(\tilde{\mathbf{q}})$ is diagonal by (13)). The rotation matrix $\mathbf{O}$ is generally unique. Intuitively, the K orthogonalized factors are truly "uncorrelated" with each other, so should not have any cross impacts. The orthogonalized price impact model is

$$\Delta\mathbf{p}(\mathbf{f}) = \sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_k \mathrm{cov}(\mathbf{R}_0, \tilde{\mathbf{b}}_k^\top \mathbf{R}_0) = \mathrm{var}(\mathbf{R}_0)\tilde{\mathbf{B}}\tilde{\mathbf{\Lambda}}\tilde{\mathbf{q}}. \tag{15}$$

where the $K \times K$ price-of-risk matrix $\tilde{\mathbf{\Lambda}} = \mathrm{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \ldots, \tilde{\lambda}_K)$ is diagonal. The IUF eliminates the off-diagonal terms from the price-of-risk matrix $\tilde{\mathbf{\Lambda}}$.

3.2 Flow-Based Fama-MacBeth Regression

Table 1 illustrates our estimation procedure. We rotate original model (10) to obtain the orthogonalized model (15). Under the orthogonalized model, the estimation is a flow-based Fama-MacBeth regression. In what follows, we explain each step in detail.

Data pre-processing. Our data inputs are $\{r_{n,t}\}$ and $\{f_{n,t}\}$, the return and flow of asset n at time t for $n = 1, 2, \ldots, N$ and $t = 1, 2, \ldots, T$. Empiricists need to choose asset flows that are plausibly uninformative about assets' fundamental returns. Because the model assumes $\mathbb{E}[\mathbf{f}] = \mathbf{0}$, we remove the time-series mean of flow $f_{n,t}$ for each asset n. Accordingly, price impact under our model (10) also has a zero mean, so we remove the time-series mean of return $r_{n,t}$ for each asset n.

The time series of K factor flows $\{q_{k,t}\}$ for $k = 1, 2, \ldots, K$ and $t = 1, 2, \ldots, T$ is the primitive input, which empiricists need to specify. Empiricists have the freedom to find the most economically reasonable $q_{k,t}$ for the cross section of price impacts. As with asset flow $f_{n,t}$, we remove the time-series mean of each factor flow $q_{k,t}$.

We estimate fundamental-risk matrix $\text{var}(\mathbf{R}_0)$ and take it as a known model input. To estimate $\text{var}(\mathbf{R}_0)$ that is mostly free from the confounding short-term price impact, we empirically use long-horizon returns. We normalize the estimated $\text{var}(\mathbf{R}_0)$ to annual frequency. We use long-horizon returns to approximate fundamental risk because this approximation is transparent and easy-to-interpret. After clarifying the price impact model estimation, we explain in footnote 16 how to jointly estimate fundamental risk $\text{var}(\mathbf{R}_0)$ and the price impact model by imposing an extra restriction to the data.

First-stage regression. For each asset $n = 1, 2, \ldots, N$, we run a time-series regression of asset flow $f_{n,t}$ on factor flows $q_{1,t}, q_{2,t}, \ldots, q_{K,t}$,

$$f_{n,t} = \sum_{k=1}^{K} b_{n,k} q_{k,t} + e_{n,t}. \tag{16}$$

This time-series regression recovers the beta $b_{n,k}$ of asset-n flow to factor-k flow.

Flow betas $b_{n,k}$ $(n = 1, 2, \ldots, N)$ are mimicking portfolio weights for factor flow $q_{k,t}$. Regression (16) implies that, holding all else equal, a one-unit flow $q_{k,t}$ into factor k leads to a $b_{n,k}$-unit flow into asset n. By definition, portfolio weights are the rule for allocating a one-dollar flow into the factor portfolio to different assets. If asset n has weight $b_{n,k}$ in factor k, then, for every one-dollar factor flow, the factor portfolio allocates $\$b_{n,k}$ to asset n.

Therefore, as long as we choose the correct unit of flow, flow betas are portfolio weights.[14]

The residual $e_{n,t}$ is the idiosyncratic component of flow into asset n at time t. Under the true model, idiosyncratic flows have no price impacts. This idea implies the flow-based Gibbons, Ross, and Shanken (1989) test presented in Section 3.3.

Orthogonalization. We find some $K \times K$ invertible matrix $\mathbf{O}$ that satisfies orthogonalization conditions (12) and (13). Appendix A.2 provides details for solving for $\mathbf{O}$. We write the first-stage flow beta $b_{n,k}$ as an $N \times K$ matrix $\mathbf{B}$. We construct the orthogonalized factor portfolios and factor flows

$$\mathbf{B}\mathbf{O} = \tilde{\mathbf{B}} = (\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2, \ldots, \tilde{\mathbf{b}}_K) = \{\tilde{b}_{n,k}\}, \tag{17}$$

$$\mathbf{O}^{-1} \begin{pmatrix} q_{1,t} \\ q_{2,t} \\ \cdots \\ q_{K,t} \end{pmatrix} = \begin{pmatrix} \tilde{q}_{1,t} \\ \tilde{q}_{2,t} \\ \cdots \\ \tilde{q}_{K,t} \end{pmatrix}, \tag{18}$$

for all times $t = 1, 2, \ldots, T$. The orthogonalized factor portfolios $\tilde{\mathbf{B}}$ form an $N \times K$ matrix, with the k-th factor being $\tilde{\mathbf{b}}_k$ and the (n,k)-th entry being $\tilde{b}_{n,k}$.

Second-stage regression. We estimate price impact under the orthogonalized model. Specifically, we denote the (n,m)-th element of fundamental-risk matrix $\text{var}(\mathbf{R}_0)$ as $v_{n,m}$. The orthogonalized model (15) implies the following panel regression,

$$r_{n,t} = \sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) + \xi_{n,t}. \tag{19}$$

The term $\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k}$ is the covariance $\text{cov}(\mathbf{R}_0(n), \tilde{\mathbf{b}}_k^\top \mathbf{R}_0)$ between asset n's fundamental return and orthogonalized factor k's fundamental return, so is the quantity of risk[15] per unit

[14] An (2022) demonstrates this idea formally by generalizing the standard asset-pricing framework.

[15] Instead of pre-estimating fundamental-risk matrix $\text{var}(\mathbf{R}_0)$, we can equivalently form portfolios $\tilde{\mathbf{b}}_k^\top \mathbf{R}_0$ using the first-stage regression and orthogonalization. We then estimate each asset's fundamental-return beta to orthogonalized factors. The fundamental-return beta equals $\text{cov}(\mathbf{R}_0(n), \tilde{\mathbf{b}}_k^\top \mathbf{R}_0)$ because $\text{var}(\tilde{\mathbf{b}}_k^\top \mathbf{R}_0) = 1$.

of flow $\tilde{q}_{k,t}$. The regressor $\tilde{q}_{k,t}\left(\sum_{m=1}^{N} v_{n,m}\tilde{b}_{m,k}\right)$ is the risk induced by factor-k flow in period t. This regression identifies the price of risk $\tilde{\mathbf{\Lambda}} = \text{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \ldots, \tilde{\lambda}_K)$.

Rotation to the original model. We recover the original price impact model (10) under factor flow $q_{1,t}, q_{2,t}, \ldots, q_{K,t}$. Specifically, equation (11) implies $\mathbf{\Lambda} = \mathbf{O}\tilde{\mathbf{\Lambda}}\mathbf{O}^{-1}$.

Standard Errors. Like the standard Fama-MacBeth regression, our estimation also faces a generated-regressor issue. To compute standard errors, we use bootstrapping. Specifically, we sample with replacement the panel $\{r_{n,t}, f_{n,t}, q_{k,t}\}$ across times $t = 1, 2, \ldots, T$. The bootstrapping is theoretically justified because our model is static and imposes no intertemporal restrictions. For each bootstrapping sample, we compute the corresponding parameter estimates. We compute standard errors based on estimates from all bootstrapping samples.

3.3 Flow-Based Gibbons-Ross-Shanken Test

In this subsection, we provide the flow-based Gibbons, Ross, and Shanken (1989) test, a χ^2 test of anomaly price impact. We prove that the test statistic is the difference between realized and model-implied squared maximum price impact ratios.

We estimate an unrestricted version of model (19) as

$$r_{n,t} = \sum_{m=1}^{N} a_{n,m} e_{m,t} + \sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) + \xi_{n,t}. \tag{20}$$

The idiosyncratic flow $e_{m,t}$ of asset m at time t is the residual from first-stage regression (16). The coefficient $a_{n,m}$ measures the anomaly price impact on asset n by the idiosyncratic flow of asset m. The true model implies the null hypothesis

$$H_0 : a_{n,m} = 0 \text{ for all } n \text{ and } m. \tag{21}$$

That is, the idiosyncratic flow of any asset m should not generate price impacts for any asset n, including itself. Let $\hat{\mathbf{a}}$ be the $N^2 \times 1$ vector of parameter estimates $\hat{a}_{n,m}$. As period T

tends to infinity, the asymptotic χ^2 test statistic for the null hypothesis is

$$T\hat{\mathbf{a}}^\top \frac{\text{var}(\hat{\mathbf{a}})^{-1}}{T}\hat{\mathbf{a}} \sim \chi^2_{N^2}. \tag{22}$$

The cross section of price impacts. To understand the χ^2 test statistic, we study the cross section of price impacts. Price impact model (19) not only explains how flows impact returns but also implies the following return decomposition,

$$r_{n,t} = \underbrace{\sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right)}_{\text{price impact}} + \underbrace{\xi_{n,t}}_{\text{fundamental return}}. \tag{23}$$

Only the portion of $r_{n,t}$ that covariates with the flow can be interpreted as price impacts. If the price moves without any flow, then this variation does not stem from the price impact, but instead reflects the fundamental return. That is, the regression residual $\xi_{n,t}$ is the fundamental return on asset n in period t.

Therefore, we define the price impact component $\bar{r}_{n,t}$ of asset n at time t as

$$\bar{r}_{n,t} = \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right), \tag{24}$$

where we use $\hat{\lambda}_k$ to denote the estimates of $\tilde{\lambda}_k$. The cross section of price impacts at time t is the $N \times 1$ vector $\bar{\mathbf{r}}_t = (\bar{r}_{1,t}, \bar{r}_{2,t}, \ldots, \bar{r}_{N,t})^\top$. The price impact ratio of any portfolio $\mathbf{c} \in \mathbb{R}^N$ at time t is defined as

$$\frac{\mathbf{c}^\top \bar{\mathbf{r}}_t}{\sigma(\mathbf{c}^\top \mathbf{R}_0)}. \tag{25}$$

The numerator is the price impact of portfolio $\mathbf{c}$ at time t and the denominator is the fundamental risk of portfolio $\mathbf{c}$. An (2022) proves the flow-based Hansen-Jagannathan bound and shows that the maximum price impact ratio across all portfolios is

$$\max_{\mathbf{c}\in\mathbb{R}^N} \frac{\mathbf{c}^\top \bar{\mathbf{r}}_t}{\sigma(\mathbf{c}^\top \mathbf{R}_0)} = \sqrt{\bar{\mathbf{r}}_t^\top \text{var}(\mathbf{R}_0)^{-1}\bar{\mathbf{r}}_t}. \tag{26}$$

The MPIR portfolio has the largest price dislocation relative to its fundamental risk. We define the time-series average of the model-implied squared maximum price impact ratio as

$$\hat{\theta}_q^2 \equiv \frac{1}{T}\sum_{t=1}^{T} \bar{\mathbf{r}}_t^\top \text{var}(\mathbf{R}_0)^{-1}\bar{\mathbf{r}}_t. \tag{27}$$

Similarly, the price impact component under the unrestricted model (20) is

$$\check{r}_{n,t} = \sum_{m=1}^{N} \hat{a}_{n,m} e_{m,t} + \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right). \tag{28}$$

The cross section of price impacts at time t under model (20) is $\check{\mathbf{r}}_t = (\check{r}_{1,t}, \check{r}_{2,t}, \ldots, \check{r}_{N,t})^\top$. We define the time-series average of the realized squared maximum price impact ratio as

$$\hat{\theta}^{*2} \equiv \frac{1}{T}\sum_{t=1}^{T} \check{\mathbf{r}}_t^\top \text{var}(\mathbf{R}_0)^{-1}\check{\mathbf{r}}_t. \tag{29}$$

Recall that we have pre-estimated the fundamental risk matrix $\text{var}(\mathbf{R}_0) = \{v_{n,m}\}$ to estimate (23), and the error term $\xi_{n,t}$ is the fundamental return under the true model. Therefore, an implicit restriction of the true model is

$$\text{var}(\mathbf{R}_0) = H\mathbf{\Sigma}_\xi, \tag{30}$$

where $\mathbf{\Sigma}_\xi$ is the variance matrix of the cross-sectional error term $\boldsymbol{\xi}_t = (\xi_{1,t}, \xi_{2,t}, \ldots, \xi_{N,t})^\top$. Because we normalize $\text{var}(\mathbf{R}_0)$ to the annual frequency, H is the ratio of the annual frequency to the higher frequency at which we estimate the price impact model.[16] Theoretically, this

[16]In this paper, we pre-estimate $\text{var}(\mathbf{R}_0)$ from long-horizon returns and do not impose this restriction (30). In the ongoing project, we impose this restriction in the data to jointly estimate the fundamental risk $\text{var}(\mathbf{R}_0)$ and the price impact model.

restriction (30) connects price impact ratios (27) and (29) to the χ^2 test statistic (22).

Appendix A.3 proves the following flow-based Gibbons, Ross, and Shanken (1989) result.

THEOREM 2. *Assuming that $e_{n,t}$, $\tilde{b}_{n,k}$, and $\tilde{q}_{k,t}$ are observed and the cross-sectional error term $\boldsymbol{\xi}_t$ is i.i.d. over time, we have, almost surely in the limit of T tending to infinity,*

$$\hat{\mathbf{a}}^\top \frac{\text{var}(\hat{\mathbf{a}})^{-1}}{T} \hat{\mathbf{a}} = H(\hat{\theta}^{*2} - \hat{\theta}_q^2). \tag{31}$$

Equation (31) shows that, as period T tends to infinity, the χ^2 test statistic for the anomaly price impact in (22) equals the difference between the realized and model-implied average squared maximum price impact ratios. The multiplication by H normalizes the price impact ratio to annual frequency. This result is parallel to the original Gibbons, Ross, and Shanken (1989), who show that the χ^2 test statistic for the anomaly expected return equals the difference between the realized and model-implied squared maximum Sharpe ratios. Note that $e_{n,t}$, $\tilde{b}_{n,k}$, and $\tilde{q}_{k,t}$ are estimated from the first-stage regression and subsequent orthogonalization, so the correct χ^2 test statistic should account for the generated-regressor issue. A generalization of the Shanken (1992) correction is likely doable, and we leave this for future research.

The empirical implication of Theorem 2 is that, to *jointly* test the N^2 anomaly price impacts $a_{n,m}$ of all assets, it suffices to study the MPIR portfolio. Section 4.6 applies this idea to develop the price-impact-ratio diagnostic.

Simplifying the regression for anomaly price impacts. If we care only about the point estimate of $a_{n,m}$, we simply run the following asset-by-asset time-series regression,

$$r_{n,t} = \sum_{m=1}^{N} a_{n,m} e_{m,t} + \zeta_{n,t}. \tag{32}$$

Idiosyncratic flow $e_{m,t}$ is the residual from the first-stage regression (16), so by construction $e_{m,t}$ is orthogonal to any factor flow $q_{k,t}$. Appendix A.4 uses this fact to show that panel regression (20) reduces to time-series regression (32).

4 Empirical Results: A Three-Factor Model of Price Impacts

In this section, we estimate and diagnose the MPIR portfolio using a three-factor model of price impacts.

4.1 Data and Empirical Measures

In this subsection, we describe our data and construct empirical measures.

First, we construct monthly stock-level returns and flows. To obtain uninformative flows at the stock level, we use mutual fund flow-induced trading from Coval and Stafford (2007) and Lou (2012). The idea is that mutual fund investors are mostly uninformed retail investors. Their decision to buy or sell mutual fund shares reveals little information about the future fundamentals of underlying stocks. Experiencing retail investor flows, mutual funds tend to buy or sell stocks in proportion to their existing holdings, generating stock-level uninformative flows.

We obtain monthly mutual fund returns and characteristics from the CRSP Survivorship-Bias-Free Mutual Fund database and quarterly holdings from the Thomson/Refinitiv Mutual Fund Holdings Data (S12). Our mutual fund sample contains both active and passive mutual funds. The sample period runs from January 2000 through September 2020. Our sample starts in 2000 because mutual funds constituted a smaller proportion of the U.S. equity market before then.[17]

Mutual fund returns and total net assets (TNA) are essential for constructing mutual fund flows. To obtain high accuracy for these inputs, we cross-check mutual funds' monthly returns and TNA from CRSP with those from Morningstar (return and TNA) and Thomson/Refinitiv (TNA). We also manually correct a number of data input errors. We present the details in Appendix B. Correcting measurement errors is especially important because we measure flow in dollar amounts—a small error for a big fund causes significant bias.

[17]For example, figure 7.1 in Investment Company Institute (2021) shows that only 5.7% of U.S. households held mutual funds in 1980. This ratio grew steadily to 45% in 2000 and has remained stable since then.

We denote the TNA of mutual fund $m = 1, 2, \ldots, M$ at the end of month $t = 1, 2, \ldots, T$ as $\text{TNA}_{m,t}$ and the mutual fund's net-of-fee return in month t as $\check{r}_{m,t}$. We define the dollar flow into mutual fund m in month t as

$$\check{f}_{m,t} = \text{TNA}_{m,t} - \text{TNA}_{m,t-1}(1 + \check{r}_{m,t}). \tag{33}$$

We denote the weights of stock $j = 1, 2, \ldots, J$ in mutual fund m's holding at the end of month $t-1$ as $\check{w}_{j,m,t-1}$. To obtain $\check{w}_{j,m,t-1}$, we multiply mutual fund m's holding of stock j (in stock share units) by the stock price at the end of month $t-1$ and then divide by $\text{TNA}_{m,t-1}$. Because mutual fund holdings are disclosed at quarterly frequency, we compute the weights using the the last disclosed holding and corresponding price and TNA information before the end of month $t-1$. The mutual fund flow-induced trading $\hat{f}_{j,t}$ in stock j in month t is defined as

$$\hat{f}_{j,t} = \sum_{m=1}^{M} \check{w}_{j,m,t-1} \check{f}_{m,t}. \tag{34}$$

The return $\hat{r}_{j,t}$ on stock j in month t is obtained from CRSP.

Second, we aggregate stock-level flow to the Fama and French (1993) factor level. We construct the weight $w_{j,k,t}$ of stock j in factor portfolio k in month t, where $k = \text{MKT}, \text{SMB}, \text{HML}$. We write the portfolio weights $w_{j,k,t}$ in month t as a $J \times 3$ matrix $\mathbf{W}_t$. We define factor flows in month t as

$$\begin{pmatrix} q_{\text{MKT},t} \\ q_{\text{SMB},t} \\ q_{\text{HML},t} \end{pmatrix} = (\mathbf{W}_t^\top \mathbf{W}_t)^{-1} \mathbf{W}_t^\top \begin{pmatrix} \hat{f}_{1,t} \\ \hat{f}_{2,t} \\ \cdots \\ \hat{f}_{J,t} \end{pmatrix}. \tag{35}$$

To obtain factor flows, we multiply stock flows by the *pseudoinverse* of portfolio weights $(\mathbf{W}_t^\top \mathbf{W}_t)^{-1} \mathbf{W}_t^\top$. This procedure differs from constructing factor returns, where one multiplies stock returns by the transpose $\mathbf{W}_t^\top$. An (2022) provides the theoretical foundation for this construction, and here we provide intuitions. By definition of portfolio weights, if portfolio-k flow increases by \$1, stock-j flow increases by $\$w_{j,k,t}$. Flow first goes into the portfolio, and then is allocated according to portfolio weights. We observe the cross section of stock flow $\hat{f}_{1,t}, \hat{f}_{2,t}, \ldots, \hat{f}_{J,t}$ in month t and we back out the factor flows. This is a cross-sectional regression for each month t,

$$\hat{f}_{j,t} = w_{j,\text{MKT},t} q_{\text{MKT},t} + w_{j,\text{SMB},t} q_{\text{SMB},t} + w_{j,\text{HML},t} q_{\text{HML},t} + \epsilon_{j,t}, \tag{36}$$

where regressors are portfolio weights $w_{j,k,t}$. This is why we use pseudoinverse $(\mathbf{W}_t^\top \mathbf{W}_t)^{-1}\mathbf{W}_t^\top$ in (35). Factor flows $q_{\text{MKT},t}$, $q_{\text{SMB},t}$, $q_{\text{HML},t}$ provide the best approximation in the OLS sense of the cross-sectional stock flows $\hat{f}_{j,t}$.

In Appendix C.1, we use a three-stock example to illustrate why the pseudoinverse of portfolio weights $(\mathbf{W}_t^\top \mathbf{W}_t)^{-1}\mathbf{W}_t^\top$, rather than the transpose $\mathbf{W}_t^\top$, leads to the correct measure of factor flows. The resulting measures differ as long as factor portfolio weights are not orthogonal to each other. The MKT, SMB, and HML portfolio weights are not orthogonal.[18]

We choose flows into Fama-French portfolios to estimate factor model of price impacts and construct MPIR portfolio, because Fama and French (1993) show that these portfolios are important risk factors. So if we can additionally show that these factor flows explain common variation in asset flows, then these factors explain both flows and fundamental risks.[19] We stress that we use MKT, SMB, and HML factor flows. Nowhere do we use MKT, SMB, and HML factor returns.

Third, we aggregate stock-level return and flow to Fama and French (1993) 5×5 test assets. The 25 test asset weights are constructed in the same way as in Fama and French (1993), based on sorted size and book-to-market equity. The weight of stock j in test asset

[18]We compute the correlation between factor portfolio weights for each month and then compute the average correlation over time. The average correlations between the portfolio weights of (MKT, SMB), (MKT, HML), and (HML, SMB) are -0.82, -0.23, and -0.20, respectively.

[19]In theory, doing a principal component analysis (PCA) on flow-induced risk achieves the maximum in-sample explanatory power. The PCA on flow-induced risk is not a simple PCA on flow. See section 2.8 of An (2022) for details. We follow Fama and French (1993) for transparency and ease of economic interpretation.

$n = 1, 2, \ldots, 25$ at month t is $u_{j,n,t}$. We write the portfolio weights $u_{j,n,t}$ in month t as a $J \times 25$ matrix $\mathbf{U}_t$. We define the flow $f_{n,t}$ of and return $r_{n,t}$ on test asset n in month t as

$$\begin{pmatrix} f_{1,t} \\ f_{2,t} \\ \cdots \\ f_{25,t} \end{pmatrix} = (\mathbf{U}_t^\top \mathbf{U}_t)^{-1} \mathbf{U}_t^\top \begin{pmatrix} \hat{f}_{1,t} \\ \hat{f}_{2,t} \\ \cdots \\ \hat{f}_{J,t} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} r_{1,t} \\ r_{2,t} \\ \cdots \\ r_{25,t} \end{pmatrix} = \mathbf{U}_t^\top \begin{pmatrix} \hat{r}_{1,t} \\ \hat{r}_{2,t} \\ \cdots \\ \hat{r}_{J,t} \end{pmatrix}. \tag{37}$$

Fourth, to estimate the fundamental risk matrix $\text{var}(\mathbf{R}_0)$ that is mostly free from the confounding short-term price impact, we calculate the variance-covariance matrix between the one-year returns of 25 test assets.[20]

4.2 The Playing Field

In this subsection, we summarize test asset and factor flows and estimate raw price impacts.

In the top two panels of Figure 2, we plot the time series of test asset and factor flows at monthly frequency. In the two middle panels, we plot the cumulative sum of test assets and factor flows. In terms of the 25 asset flows, big companies naturally exhibit the widest variation because of their large market capitalization. Asset flows exhibit strong commonality. The MKT factor exhibits the widest variation in flows, but SMB and HML flows also exhibit sizable variation.[21]

The BL asset (the Biggest companies with the Lowest book-to-market equity) and the MKT factor experience a total outflow from mutual funds of about \$200 billion over the past twenty years, and the SMB and HML factors each experience inflows of about \$50 billion. Other asset flows roughly average out to zero. Asset and factor flows also display a strong

[20] We use long-horizon returns to approximate fundamental risk because this approximation is transparent and easy-to-interpret. See footnote 16 for how to jointly estimate fundamental risk $\text{var}(\mathbf{R}_0)$ and the price impact model by imposing an extra restriction to the data.

[21] Flow into the MKT factor induced by mutual fund trading is significantly less volatile than aggregate mutual fund flow, which we show in Figure B.1 of Appendix B.3. These are distinct objects. The variation in aggregate mutual fund flow consists of many factors, and flow into the MKT factor is just one of them.

Figure 2. Time series of 25 test asset flows and 3 factor flows

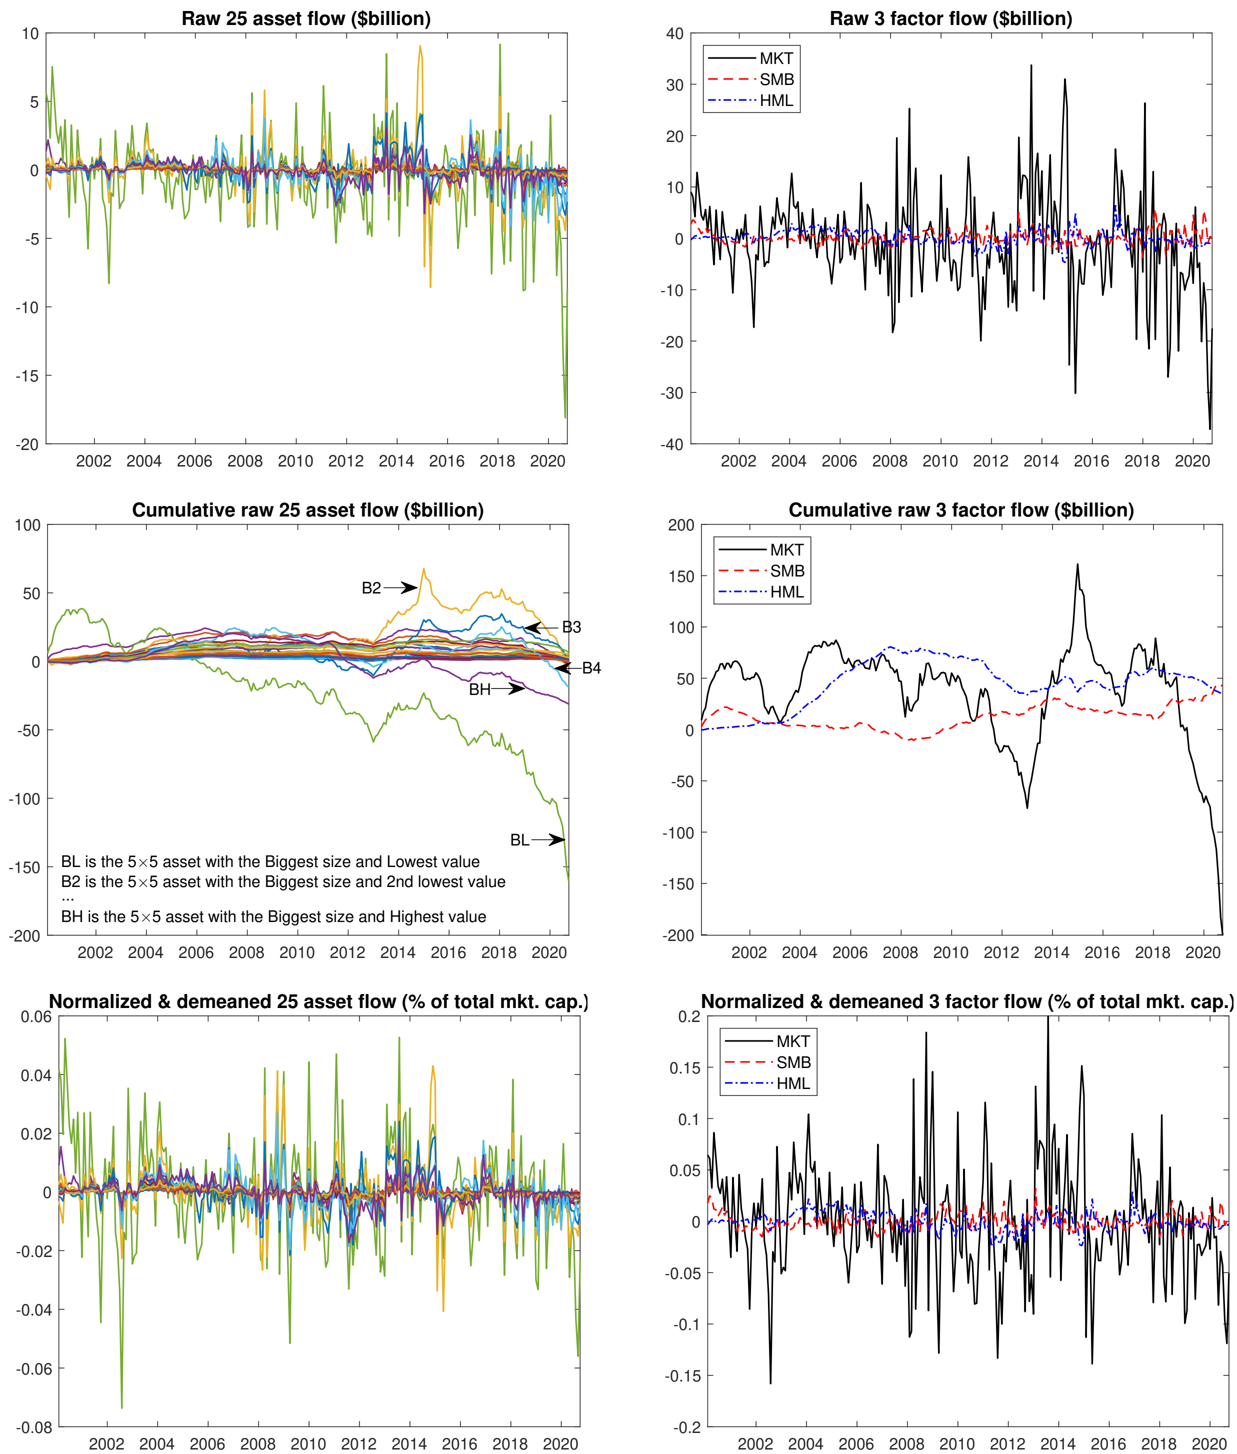

Notes: In the top two figures, we plot the time series of test asset and factor flows at monthly frequency. In the two middle figures, we plot the cumulative sum of test asset and factor flows. In the bottom two figures, we divide test asset and factor flows in month t by the total stock market capitalization of month $t-1$ and then subtract the time-series mean. The sample period runs from January 2000 through September 2020.

serial correlation.[22] Our static model is silent regarding the impact of expected flows and the serial correlation of flows, but these facts motivate future research.

Asset and factor flows tend to exhibit greater volatility in the latter portion of the sample. To address this effect, we normalize all asset and factor flows in month t by total stock market capitalization in month $t-1$ (defined as the sum of all 25 test assets). We do not normalize asset or factor flows by their respective market capitalizations. Our normalization preserves the relative magnitude between flows in the cross section, which is crucial for obtaining the correct flow beta. Our model concerns only unexpected flows, so we remove the unconditional time-series mean of asset and factor flows.

The bottom two panels of Figure 2 show the time series of normalized and demeaned asset and factor flows, which are used in later estimations. The MKT factor can experience inflow and outflow as large as 0.1% of total stock market capitalization. SMB and HML exhibit smaller variations in flow, but we later show that these two factors also contribute greatly to the cross section of price impacts. The pairwise correlation between the flows of (MKT, SMB), (MKT, HML), and (HML, SMB) are only 0.13, 0.20, and -0.01, respectively.

We remove the unconditional time-series mean of return $r_{n,t}$ for each asset n. We define the raw price impact multiplier η_n from the asset-by-asset time-series regression,

$$r_{n,t} = \eta_n f_{n,t}/w_n + \epsilon_{n,t}, \tag{38}$$

where w_n is the market capitalization weight of test asset n. To obtain w_n, we calculate the ratio of the market capitalization of asset n to the total stock market for each month, and take the simple average over time. Recall that flow $f_{n,t}$ is normalized by total stock market capitalization, so the ratio $f_{n,t}/w_n$ measures the flow into asset n in month t as a fraction of the asset's market capitalization. Using $f_{n,t}/w_n$ on the right-hand side of regression (38)

[22]The AR(1) coefficients for the 25 asset flows range from 0.19 to 0.65. The AR(1) coefficients for the MKT, SMB, and HML flows are 0.24, 0.36, and 0.57, respectively. These coefficients are all statistically significant.

Figure 3. Illustration of cross impact for 5×5 assets

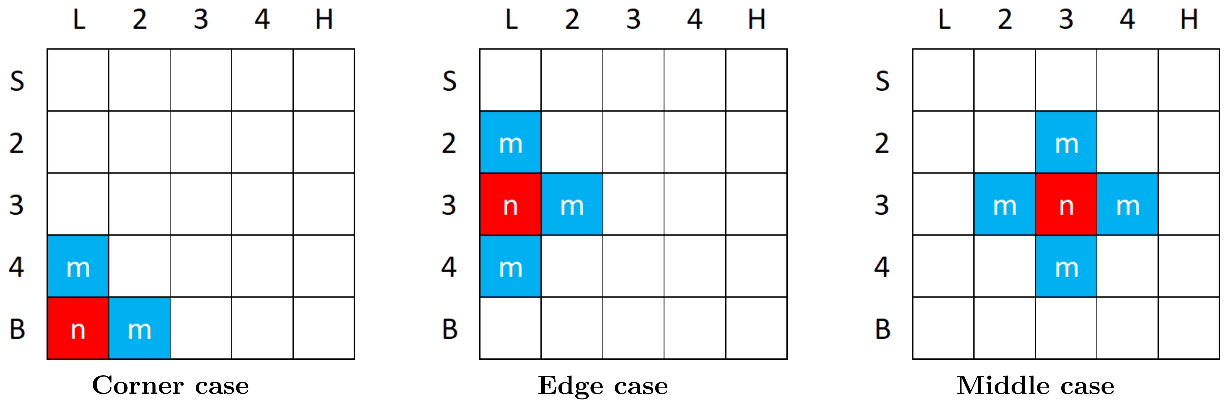

Notes: This figure illustrates the raw cross impact multiplier ϕ_n for asset n. The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). We find all assets m that are neighboring to asset n. We calculate the average flow into these neighboring assets for each month t. To obtain multiplier ϕ_n, we regress the demeaned return $r_{n,t}$ of asset n on the average flow.

ensures that our η_n is comparable to the standard price impact multiplier measure.[23]

The return $r_{n,t}$ on asset n is affected not only by its own flow $f_{n,t}/w_n$ but also by flow $f_{m,t}/w_m$ into other assets m. To account for all potential cross impacts, one would estimate

$$r_{n,t} = \sum_{m=1}^{N} a_{n,m} f_{m,t}/w_m + \epsilon_{n,t}. \tag{39}$$

However, it is impractical to estimate and interpret $25 \times 25 = 625$ parameters $a_{n,m}$ given the monthly frequency flow and return data. For all 25 test assets, the regression R^2 is about 20%, but many of the estimates $a_{n,m}$ are negative or statistically insignificant.[24] Regression (39) suffers from over-fitting problems and lacks statistical power.

We propose a remedy to this estimation problem by taking advantage of the sorted feature of 5×5 assets. Specifically, we define the raw cross impact multiplier ϕ_n of asset n from the

[23]Our theory implies that one should not apply any cross-sectional normalization to flow $f_{n,t}$. We divide by w_n in (38), which has only one effect of scaling the multiplier η_n.

[24]We report the regression R^2 in Table C.1 of Appendix C.3.

asset-by-asset time-series regression,

$$r_{n,t} = \phi_n \left(\sum_{m \text{ neighboring to } n} f_{m,t}/w_m \right) \Big/ (\text{number of } m \text{ neighboring to } n) + \epsilon_{n,t}. \tag{40}$$

Figure 3 illustrates the idea. We find all assets m that are neighboring to asset n. Because 5×5 assets are sorted, these assets m are similar to asset n in terms of size and book-to-market equity. When asset n is in a corner, there are two neighboring assets m; when n is at an edge, there are three m; when n is in the middle, there are four m. We calculate the average flow (normalized by the respective market capitalization weight w_m) into these neighboring assets for each month t. The raw cross-impact multiplier ϕ_n measures whether the neighboring asset flow impacts the return $r_{n,t}$ on asset n. Focusing on the neighboring assets' cross impacts significantly increases the statistical power of regression (40). Our goal is to show that, after imposing our price impact model, the residual cross impacts are small. The more easily we can find raw cross impacts, the harder is the task our model faces when explaining these impacts.

Table 2 presents the R^2, point estimates of multipliers, and t-statistics for raw price impact regression (38) and raw cross-impact regression (40). For an average 5×5 asset, the price impact from its own flow and from neighboring assets' flows both explain about 10% of return variations. We subsequently show that this component of return variations has a large effect on future return predictability. For many assets, the R^2 obtained from cross-impact regression (40) is actually larger than the R^2 obtained from price impact regression (38). This evidence supports our approach of jointly studying the cross section of price impacts.

The point estimates of multipliers η and ϕ range between 10 and 20. Our multiplier estimates are on the larger side of estimates reported in the literature.[25] For example, Gabaix and Koijen (2022) estimate a market-level multiplier of 5. The difference occurs because 5×5 assets are sorted by size and book-to-market equity, and we subsequently show

[25]Note that the price impact multiplier is the inverse of demand elasticity. See Table 1 of Gabaix and Koijen (2022) for a summary of estimates in the literature.

Table 2. Raw price impact and raw cross impact

	raw price impact					raw cross impact				
	Low	2	3	4	High	Low	2	3	4	High
	regression R^2					regression R^2				
Small	6.59%	7.67%	10.28%	10.56%	9.64%	8.85%	8.93%	11.44%	10.96%	12.04%
2	9.57%	9.49%	10.98%	10.41%	8.68%	7.97%	9.81%	12.05%	10.47%	9.34%
3	8.58%	10.17%	10.77%	9.65%	8.30%	9.48%	11.60%	10.60%	11.79%	10.18%
4	12.26%	10.61%	9.26%	10.63%	7.72%	9.00%	10.42%	10.79%	8.87%	9.33%
Big	3.25%	3.07%	2.72%	1.71%	5.99%	6.12%	8.98%	4.59%	5.43%	5.71%
	multiplier η					multiplier ϕ				
Small	15.48	12.57	11.21	10.75	11.26	16.60	13.30	12.31	11.08	13.05
2	13.82	10.23	10.53	10.45	11.37	14.29	12.34	12.28	11.67	13.68
3	13.33	11.44	12.06	11.59	12.24	15.63	13.83	12.70	14.17	14.11
4	19.49	16.13	16.21	17.81	13.69	19.38	18.11	18.62	16.60	17.54
Big	15.60	11.77	11.05	9.26	14.37	15.64	19.32	14.12	14.67	15.51
	$t(\eta)$					$t(\phi)$				
Small	4.35	4.39	4.89	5.07	4.35	5.00	4.89	4.98	5.29	4.88
2	4.90	4.58	4.58	5.05	4.27	4.61	4.62	5.29	4.96	4.71
3	4.30	4.93	5.23	4.95	4.78	4.45	5.38	4.87	5.41	5.09
4	4.92	4.51	4.09	4.58	5.20	4.36	4.39	4.59	4.28	5.11
Big	2.16	2.47	2.20	1.63	3.90	3.30	4.10	2.75	3.70	3.49

Notes: This table presents the R^2, point estimates of multipliers, and t-statistics for raw price impact regression (38) and raw cross impact regression (40). The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). To obtain the raw price impact, we run asset-by-asset time-series regressions of each asset's demeaned return on its own flow. To obtain the raw cross impact, we run asset-by-asset time-series regressions of each asset's demeaned return on its neighboring assets' average flow. The t-statistics are calculated from heteroskedasticity-robust standard errors.

that the SMB and HML factors have a higher price of risk than the MKT factor.

Table C.2 of Appendix C.3 runs the joint price impact and cross impact regression,

$$r_{n,t} = \eta_n f_{n,t}/w_n + \phi_n \left(\sum_{m \text{ neighboring to } n} f_{m,t}/w_m \right) \Big/ (\text{number of } m \text{ neighboring to } n) + \epsilon_{n,t}. \tag{41}$$

However, because flow into an asset is highly correlated with flow into neighboring assets, regression (41) suffers from multicollinearity. Specifically, the regression R^2 of (41) improves very little compared to the maximum of the R^2 from raw price impact regression (38) and the R^2 from raw cross impact regression (40). For each asset n, the point estimates of η_n and ϕ_n from regression (41) are usually one positive and one negative, and both lack statistical significance. The multicollinearity concern leads us to run separate regressions (38) and (40).

Table 2 is our playing field—our goal is to explain these raw impacts. In what follows, we estimate our price impact model with MKT, SMB, and HML flows.

4.3 Common Variation in Flows

In this subsection, we show that MKT, SMB, and HML flows explain the common variation in 5×5 asset flows. This is the first stage of our flow-based Fama-MacBeth regression.

We run the following asset-by-asset time-series regression,

$$f_{n,t} = b_{n,\text{MKT}} q_{\text{MKT},t} + b_{n,\text{SMB}} q_{\text{SMB},t} + b_{n,\text{HML}} q_{\text{HML},t} + e_{n,t}. \tag{42}$$

Table 3 reports the regression R^2, point estimates and t-statistics for regression coefficients $b_{n,\text{MKT}}$, $b_{n,\text{SMB}}$, and $b_{n,\text{HML}}$ as well as the market capitalization weight w_n of the 5×5 assets. The regression R^2 ranges from 50% for small companies to 80% for large companies. Factor flows explain large portions of the common variation in asset flows. In Figure C.1 of Appendix C.3, we report idiosyncratic flows $e_{n,t}$ and find that idiosyncratic flows exhibit significantly lower commonality than raw flows in Figure 2.

Flow betas $b_{n,\text{MKT}}$, $b_{n,\text{SMB}}$, and $b_{n,\text{HML}}$ are *mimicking portfolio weights*. Regression (42) implies that, holding all else equal, a one-dollar increase[26] in MKT flow increases asset-n flow by \$$b_{n,\text{MKT}}$. The 25 flow betas $b_{n,\text{MKT}}$ are the portfolio that liquidity providers in equilibrium trade against for every dollar of MKT flow.

Importantly, our flow betas or mimicking portfolio weights are not necessarily the same as the Fama-French portfolio weights. Therefore, our regression (42) concerning flow betas neither replicates nor can be derived from the usual Fama-French regression. Whether our

[26]Recall that we normalize all asset and factor flows by the same total stock market capitalization.

Table 3. First-stage regression: asset flows on factor flows

	Low	2	3	4	High	Low	2	3	4	High	Low	2	3	4	High
	market cap. weight w_n					regression R^2									
Small	0.0042	0.0034	0.0040	0.0045	0.0044	52.00%	56.67%	59.78%	60.60%	54.68%					
2	0.0081	0.0072	0.0071	0.0065	0.0042	60.13%	66.87%	62.54%	66.44%	69.29%					
3	0.0142	0.0130	0.0107	0.0091	0.0060	60.75%	59.89%	66.10%	62.44%	64.81%					
4	0.0371	0.0272	0.0190	0.0148	0.0112	54.64%	58.42%	57.43%	67.53%	67.70%					
Big	0.3651	0.1790	0.1079	0.0844	0.0478	84.32%	86.20%	86.27%	87.37%	73.05%					
	Low	2	3	4	High	Low	2	3	4	High	Low	2	3	4	High
	MKT flow beta $b_{n,\text{MKT}}$					**SMB flow beta** $b_{n,\text{SMB}}$					**HML flow beta** $b_{n,\text{HML}}$				
Small	0.0022	0.0019	0.0026	0.0040	0.0033	0.0366	0.0308	0.0371	0.0360	0.0343	0.0279	0.0289	0.0415	0.0493	0.0502
2	0.0076	0.0076	0.0072	0.0074	0.0051	0.0869	0.0755	0.0693	0.0582	0.0376	0.0582	0.0745	0.0738	0.0717	0.0486
3	0.0145	0.0133	0.0108	0.0089	0.0052	0.1180	0.0882	0.0592	0.0465	0.0278	0.0909	0.1098	0.0970	0.0871	0.0606
4	0.0407	0.0267	0.0172	0.0149	0.0132	0.1256	0.0316	-0.0038	-0.0001	-0.0030	0.0984	0.1374	0.1017	0.0865	0.0816
Big	0.3168	0.1846	0.1140	0.0838	0.0590	0.0186	-0.2975	-0.2277	-0.2495	-0.1090	-0.2496	-0.1967	0.0711	0.2212	0.1997
	sum of $b_{n,\text{MKT}}$ coefficients					sum of $b_{n,\text{SMB}}$ coefficients					sum of $b_{n,\text{HML}}$ coefficients				
		all	positive	negative			all	positive	negative			all	positive	negative	
		0.97	0.97	0.00			0.13	1.02	-0.89			1.52	1.97	-0.45	
	$t(b_{n,\text{MKT}})$					$t(b_{n,\text{SMB}})$					$t(b_{n,\text{HML}})$				
Small	4.35	4.85	4.19	6.99	4.38	9.27	9.47	8.33	6.48	5.05	8.32	10.51	10.30	12.70	10.61
2	7.74	9.27	8.47	10.10	8.87	10.91	9.22	8.20	7.62	6.82	11.35	16.93	13.80	15.14	15.26
3	8.62	9.25	10.50	8.60	7.06	10.08	7.32	6.20	5.69	5.13	9.60	11.19	16.05	14.34	14.35
4	10.14	10.17	9.24	11.13	10.35	4.73	2.08	-0.36	-0.01	-0.43	4.49	8.35	10.32	11.44	11.31
Big	19.51	24.17	29.02	26.30	11.09	0.28	-8.62	-9.42	-15.20	-5.53	-2.92	-4.73	3.50	13.96	7.71

Notes: To obtain the results reported in this table, we run asset-by-asset time-series regressions (42) of asset flows on factor flows. The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). We report the regression R^2, point estimates and t-statistics for flow betas $b_{n,\text{MKT}}$, $b_{n,\text{SMB}}$, and $b_{n,\text{HML}}$ as well as the market capitalization weight w_n of the 5×5 assets. To obtain w_n, we calculate the ratio of the market capitalization of asset n to the total stock market for each month, and take the simple average over time. The t-statistics are calculated from heteroskedasticity-robust standard errors.

flow betas $b_{n,k}$ correctly recover the Fama-French portfolio weights is a meaningful economic test. Our asset-by-asset time-series regression (42) imposes no cross-sectional restrictions. Appendix C.2 shows through a simple example that our recovery works well only if the proposed factor flows $q_{k,t}$ match the principal components of cross-sectional asset flows $f_{n,t}$.

First, we study the flow beta of MKT factor $b_{n,\text{MKT}}$. The coefficient of 0.3168 for the BL test asset (the Biggest companies with the Lowest book-to-market equity) means that if MKT flow increases by \$1, the BL test asset flow increases by \$0.3168. For all assets n, flow beta $b_{n,\text{MKT}}$ is remarkably close to the market capitalization weight w_n. All $b_{n,\text{MKT}}$ are positive, and their sum is very close to one. This evidence implies that MKT flow does explain cross-sectional asset flows.

Second, we study the flow beta of SMB factor $b_{n,\text{SMB}}$. For small companies, $b_{n,\text{SMB}}$ is positive; for big companies, $b_{n,\text{SMB}}$ is negative. This is consistent with the small-minus-big construction. The sum of all positive $b_{n,\text{SMB}}$ is 1.02 and the sum of all negative $b_{n,\text{SMB}}$ is -0.89. The absolute value of both sums are very close to one, which is consistent with the Fama-French SMB portfolio construction. This evidence implies that SMB flow also explains cross-sectional asset flows.

Third, we study the flow beta of HML factor $b_{n,\text{HML}}$. Here, the recovery works less well. Only BL and B2 assets (Biggest companies with Lowest and 2nd lowest book-to-market equity) have negative flow betas, and all other assets have positive flow betas. The sum of positive $b_{n,\text{HML}}$ is almost 2 and the sum of negative $b_{n,\text{HML}}$ is about -0.5. Our flow betas differ from Fama-French HML portfolio weights, suggesting that HML flow may not perfectly explain cross-sectional asset flows.

We stress that this is not a failure of Fama and French (1993)—they never intended to explain cross-sectional flows. Rather, this evidence suggests the potential for future research to improve the construction of the HML factor to better explain both cross-sectional returns and flows. In this paper, we stick with the original Fama and French (1993) construction. HML flow contains extremely useful information about cross-sectional flows. Our orthogonalization is a powerful tool that extracts this information.

4.4 Model Orthogonalization

In this subsection, we orthogonalize the MKT, SMB, and HML factors. The economic motivation is that the original MKT, SMB, and HML factors could have cross impacts between each other because of correlated flows and correlated fundamental risks. If we directly run a price impact regression, we would have $3 \times 3 = 9$ prices of risk for 3 factors. The orthogonalized MKT, SMB, and HML factors have uncorrelated flows and uncorrelated fundamental risks. The structural restriction, which is called Irrelevance of Uncorrelated Flows (IUF), implies that orthogonalized factors have no cross impacts and have in total 3 prices of risk. See Section 3.1 and An (2022) for theoretical foundations.

In the top-left panel of Table 4, we report the correlations and standard derivations of factor flows $q_{\text{MKT},t}$, $q_{\text{SMB},t}$, and $q_{\text{HML},t}$. We use flow betas $\mathbf{b}_{\text{MKT}} = \{b_{n,\text{MKT}}\}$, $\mathbf{b}_{\text{SMB}} = \{b_{n,\text{SMB}}\}$, and $\mathbf{b}_{\text{HML}} = \{b_{n,\text{HML}}\}$ to form factor mimicking portfolios using the 25 test assets. In equilibrium, liquidity providers trade against the portfolio $\mathbf{b}_{\text{MKT}}$ for every dollar of MKT flow. The quantity of risk induced by every dollar of MKT flow is precisely the fundamental-return risk of this portfolio $\mathbf{b}_{\text{MKT}}$. We report the correlations and standard derivations of the fundamental returns of mimicking portfolios $\mathbf{b}_{\text{MKT}}^{\top}\mathbf{R}_0$, $\mathbf{b}_{\text{SMB}}^{\top}\mathbf{R}_0$, and $\mathbf{b}_{\text{HML}}^{\top}\mathbf{R}_0$ in the top-right panel.[27] We stress that the returns of mimicking portfolios are not MKT, SMB, and HML returns constructed by Fama and French (1993).

MKT flow is about six times as volatile as SMB and HML flows, and their pairwise correlation is low. Mimicking portfolios have highly correlated return risk, especially between MKT and HML mimicking portfolios. Because the original MKT, SMB, and HML returns have very low correlation, this evidence shows that the mimicking portfolio weights $\mathbf{b}_{\text{SMB}}$ and $\mathbf{b}_{\text{HML}}$ do not recover the original SMB and HML weights.[28] As discussed before, the

[27] Recall that we have estimated the fundamental risk matrix $\text{var}(\mathbf{R}_0)$ using one-year returns.

[28] During our sample period, the pairwise correlation between the one-year returns on (MKT, SMB), (MKT, HML), and (HML, SMB) are 0.11, -0.24, and -0.17, respectively.

Table 4. Model orthogonalization

original factors

	flow				**fundamental risk of mimicking portfolios**			
	MKT	SMB	HML		MKT	SMB	HML	
	flow correlation			flow std.	return correlation			return std.
MKT	1	0.13	0.20	5.58×10^{-4}	1	0.63	0.89	0.17
SMB	0.13	1	-0.01	0.82×10^{-4}	0.63	1	0.70	0.11
HML	0.20	-0.01	1	0.90×10^{-4}	0.89	0.70	1	0.36

orthogonalized factors

	flow				**fundamental risk of mimicking portfolios**			
	$\tilde{\mathbf{b}}_1$	$\tilde{\mathbf{b}}_2$	$\tilde{\mathbf{b}}_3$		$\tilde{\mathbf{b}}_1$	$\tilde{\mathbf{b}}_2$	$\tilde{\mathbf{b}}_3$	
	flow correlation			flow std.	return covariance			return std.
$\tilde{\mathbf{b}}_1$	1	0	0	10.32×10^{-5}	1	0	0	1
$\tilde{\mathbf{b}}_2$	0	1	0	1.32×10^{-5}	0	1	0	1
$\tilde{\mathbf{b}}_3$	0	0	1	0.63×10^{-5}	0	0	1	1

orthogonalization matrix

$$\tilde{\mathbf{b}}_1 = 6 \times (0.86 \times \mathbf{b}_{\text{MKT}} + 0.02 \times \mathbf{b}_{\text{SMB}} + 0.07 \times \mathbf{b}_{\text{HML}})$$
$$\tilde{\mathbf{b}}_2 = 6 \times (-2.02 \times \mathbf{b}_{\text{MKT}} + 0.15 \times \mathbf{b}_{\text{SMB}} + 0.98 \times \mathbf{b}_{\text{HML}})$$
$$\tilde{\mathbf{b}}_3 = 6 \times (0.08 \times \mathbf{b}_{\text{MKT}} + 2.12 \times \mathbf{b}_{\text{SMB}} - 0.51 \times \mathbf{b}_{\text{HML}})$$

	Low	2	3	4	High
Small	0.03	0.03	0.04	0.05	0.04
2	0.07	0.08	0.08	0.08	0.05
3	0.13	0.12	0.10	0.09	0.05
4	0.27	0.20	0.13	0.11	0.10
Big	1.54	0.84	0.59	0.49	0.37

orthogonalized factor $\tilde{\mathbf{b}}_1$

	Low	2	3	4	High
Small	0.17	0.17	0.24	0.27	0.28
2	0.32	0.41	0.41	0.38	0.26
3	0.46	0.56	0.49	0.44	0.32
4	0.19	0.51	0.39	0.33	0.32
Big	-5.30	-3.66	-1.17	0.06	0.36

orthogonalized factor $\tilde{\mathbf{b}}_2$

	Low	2	3	4	High
Small	0.38	0.30	0.35	0.31	0.28
2	0.93	0.74	0.66	0.52	0.33
3	1.23	0.79	0.46	0.33	0.17
4	1.32	-0.00	-0.35	-0.26	-0.28
Big	1.16	-3.09	-3.05	-3.80	-1.97

orthogonalized factor $\tilde{\mathbf{b}}_3$

Notes: In this table, we report the correlations and standard derivations of flows and returns of mimicking portfolios for the MKT, SMB, and HML factors and for orthogonalized factors $\tilde{\mathbf{b}}_1$, $\tilde{\mathbf{b}}_2$, and $\tilde{\mathbf{b}}_3$. We also report the orthogonalization matrix that constructs $\tilde{\mathbf{b}}_1$, $\tilde{\mathbf{b}}_2$, and $\tilde{\mathbf{b}}_3$ from flow betas $\mathbf{b}_{\text{MKT}}$, $\mathbf{b}_{\text{SMB}}$, and $\mathbf{b}_{\text{HML}}$. The three bottom figures report the portfolio weights of orthogonalized factors $\tilde{\mathbf{b}}_1$, $\tilde{\mathbf{b}}_2$, and $\tilde{\mathbf{b}}_3$.

Fama and French (1993) construction of the HML portfolio (and to some extent the SMB portfolio) may not perfectly explain cross-sectional stock flows—this is not their objective.

Nevertheless, the Fama and French (1993) construction of SMB and HML contains extremely useful information about cross-sectional asset flows. Our orthogonalization extracts this information. Specifically, in the bottom half of Table 4, we construct orthogonalized factors $\tilde{\mathbf{b}}_1$, $\tilde{\mathbf{b}}_2$, and $\tilde{\mathbf{b}}_3$ from flow betas $\mathbf{b}_{\text{MKT}}$, $\mathbf{b}_{\text{SMB}}$, and $\mathbf{b}_{\text{HML}}$. The orthogonalized factors have uncorrelated flows and uncorrelated and unit return risks. The orthogonalized factor $\tilde{\mathbf{b}}_1$ is effectively the MKT factor (multiplication by 6 normalizes $\tilde{\mathbf{b}}_1$ to unit risk). The orthogonalized factor $\tilde{\mathbf{b}}_2$ is the HML factor minus twice the MKT factor. This $\tilde{\mathbf{b}}_2$ mostly picks out the cross-sectional variation in flows in the HML direction. This interpretation is also evident from the 5×5 portfolio weights of $\tilde{\mathbf{b}}_2$. Almost all the actions in the HML direction revolve around the BL and B2 test assets,[29] which combined have over 50% of total stock market capitalization (see the top-left panel of Table 3). The orthogonalized factor $\tilde{\mathbf{b}}_3$ is twice the SMB factor minus half the HML factor. This $\tilde{\mathbf{b}}_3$ mostly picks out the cross-sectional variation in flows in the SMB direction. What matters to the price impact is not flow per se, but flow-induced risk, which we show next.

4.5 The Cross Section of Price Impacts

In this subsection, we show that the fundamental risk induced by MKT, SMB, and HML flows explains about 70% of the cross-sectional variation of price impacts. This is the second stage of the flow-based Fama-MacBeth regression.

Using the orthogonalized factors $\tilde{\mathbf{b}}_1 = \{\tilde{b}_{n,1}\}$, $\tilde{\mathbf{b}}_2 = \{\tilde{b}_{n,2}\}$, and $\tilde{\mathbf{b}}_3 = \{\tilde{b}_{n,3}\}$ and corresponding factor flows $\tilde{q}_{1,t}$, $\tilde{q}_{2,t}$, and $\tilde{q}_{3,t}$, we run the price impact regression

$$r_{n,t} = \sum_{k=1}^{3} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) + \xi_{n,t}. \tag{43}$$

The term $\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k}$ is the covariance $\mathrm{cov}(\mathbf{R}_0(n), \tilde{\mathbf{b}}_k^\top \mathbf{R}_0)$ between asset n's fundamental return and orthogonalized factor k's fundamental return,[30] so is the quantity of risk per unit of flow $\tilde{q}_{k,t}$. The regressor $\tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right)$ is the risk induced by factor-k flow in month t. This regression obtains the prices of risk $\tilde{\lambda}_k$.

In the left panel of Table 5 we report the estimated prices of risk $\tilde{\mathbf{\Lambda}} = \mathrm{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$.

[29] The BL and B2 test assets are the Biggest companies with the Lowest and 2nd lowest book-to-market equity, including companies such as Facebook and Tesla.

[30] Recall that $v_{n,m}$ is the (n, m)-th entry of the fundamental risk matrix $\mathrm{var}(\mathbf{R}_0)$.

Table 5. Second-stage regression: returns on factor flows × quantity of risk

price-of-risk coefficients

orthogonalized factors				**original factors**				**original factors without IUF**		
	$\tilde{\mathbf{b}}_1$	$\tilde{\mathbf{b}}_2$	$\tilde{\mathbf{b}}_3$		MKT	SMB	HML		MKT	SMB	HML
$\tilde{\mathbf{b}}_1$	680			MKT	1,128	-760	-5,519	MKT	540	2,569	-3,360
$\tilde{\mathbf{b}}_2$		3,852		SMB	-208	14,846	-2,066	SMB	65	3,166	-4,360
$\tilde{\mathbf{b}}_3$			15,335	HML	-178	-2,746	3,893	HML	-8	-2,24	3,393

price-of-risk t-statistics

orthogonalized factors				**original factors**				**original factors without IUF**		
	$\tilde{\mathbf{b}}_1$	$\tilde{\mathbf{b}}_2$	$\tilde{\mathbf{b}}_3$		MKT	SMB	HML		MKT	SMB	HML
$\tilde{\mathbf{b}}_1$	3.50			MKT	3.16	-0.31	-1.63	MKT	1.38	1.08	-1.76
$\tilde{\mathbf{b}}_2$		1.97		SMB	-0.95	2.18	-1.34	SMB	0.16	0.61	-2.41
$\tilde{\mathbf{b}}_3$			2.21	HML	-1.26	-1.65	2.20	HML	-0.06	-0.17	4.57
regression $R^2 = 7.25\%$				rotation from orthogonal factors				regression $R^2 = 9.11\%$		

Notes: In this table, we report the point estimates and t-statistics of prices of risk. We run the second-stage regression of demeaned returns on factor flows times quantity of risk. In the left panel we report the diagonal price-of-risk matrix $\tilde{\mathbf{\Lambda}}$ for orthogonalized factors. In the middle panel, we report $\mathbf{\Lambda}$ under original factors, through rotation from $\tilde{\mathbf{\Lambda}}$ under the orthogonalized model. For the right panel, we estimate the 3×3 matrix $\mathbf{\Lambda}$ directly without imposing the IUF restriction. We bootstrap $10,000$ times to obtain the t-statistics.

All prices of risk are statistically significant. The factor $\tilde{\mathbf{b}}_3$, which corresponds mostly to the flow-induced risk in the SMB direction, has the highest price of risk. The factor $\tilde{\mathbf{b}}_2$, which corresponds mostly in the HML direction, also has a decently high price of risk. The factor $\tilde{\mathbf{b}}_1$, which corresponds mostly in the MKT direction, has the lowest price of risk.

The CARA-Gaussian price impact model is a special case of our factor model of price impacts by further requiring all prices of risk to be the same. The significantly different prices of risk in the left panel of Table 5 strongly reject the CARA-Gaussian model. The rejection of the CARA-Gaussian model supports the "commonality-in-flow paradox" and the associated flow-induced risk channel of cross-asset price impacts in An (2022).

The middle panel of Table 5 reports the prices of risk $\mathbf{\Lambda} = \{\lambda_{k,l}\}$ under original factors

$$r_{n,t} = \sum_{k,l \in \{\text{MKT,SMB,HML}\}} \lambda_{k,l} q_{l,t} \left(\sum_{m=1}^{N} v_{n,m} b_{m,k} \right) + \xi_{n,t}. \tag{44}$$

We obtain $\mathbf{\Lambda}$ by rotating from orthogonalized $\tilde{\mathbf{\Lambda}}$ via the orthogonalization matrix in the bottom panel of Table 4, not by running a regression.[31]

We first study the diagonal terms of $\mathbf{\Lambda}$. Consistent with the evidence obtained from orthogonalized factors, SMB has the highest price of risk, followed by HML and MKT. All diagonal prices of risk are statistically significant.

We then study the off-diagonal terms of $\mathbf{\Lambda}$. The -5,519 coefficient at the top-right corner of the matrix implies that flow into the HML factor negatively affects the MKT return in the current month. Equivalently, flow into growth relative to value stocks increases the MKT return. Other entries can be interpreted similarly. Many of these off-diagonal terms are marginally statistically significant. This evidence shows that, to understand the MKT, SMB, and HML factors' explanatory power for the cross section of price impacts, one cannot just study these factors in isolation. Yet inspecting the off-diagonal prices of risk directly is not very informative. In the next subsection, we use the maximum price impact ratio to decompose how MKT, SMB, and HML contribute to the cross section of price impacts.

To illustrate the power of our IUF orthogonalization, we conduct the following placebo test and report the results in the right panel of Table 5. We run the unrestricted model (44) directly, such that the 3×3 matrix $\mathbf{\Lambda}$ is 9 free parameters. With 9 free parameters, the regression R^2 is just 9%. When we impose the IUF restriction and have only three free parameters in the left panel, we already obtain an R^2 of over 7%. We also compare the estimated prices of risk $\mathbf{\Lambda}$ with and without IUF restriction in the middle and right panels. The point estimates are broadly similar, but the IUF version provides much stronger statistical power. For example, the diagonal prices of risk for MKT and SMB are statistically significant with the IUF restriction but insignificant without IUF. The evidence shows that IUF orthogonalization significantly increases the statistical power and loses little economics.

Anomaly price impact. We study anomaly price impact under the model. The residual $e_{n,t}$ from the first-stage regression (42) is the idiosyncratic flow of asset n in month t.

[31] See equation (11) for technical details.

To see whether an asset's idiosyncratic flow has any anomaly impact on its own price, we run an unrestricted version of the second-stage regression (43),

$$r_{n,t} = \check{\eta}_n e_{n,t}/w_n + \sum_{k=1}^{3} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) + \epsilon_{n,t}. \tag{45}$$

As shown in Appendix A.4, to obtain the point estimates of multiplier $\check{\eta}_n$, it suffices to run the asset-by-asset time-series regression,

$$r_{n,t} = \check{\eta}_n e_{n,t}/w_n + \epsilon_{n,t}. \tag{46}$$

Compared with the raw impact regression (38), the anomaly price impact regression (46) substitutes raw flow $f_{n,t}$ with idiosyncratic flow $e_{n,t}$. Similarly, to obtain the anomaly cross impact, we substitute raw flow $f_{n,t}$ with idiosyncratic flow $e_{n,t}$ in regression (40),

$$r_{n,t} = \check{\phi}_n \left(\sum_{m \text{ neighboring to } n} e_{m,t}/w_m \right) \Bigg/ (\text{number of } m \text{ neighboring to } n) + \epsilon_{n,t}. \tag{47}$$

Table 6 reports the R^2, point estimates, and t-statistics for anomaly price impact regression (46) and anomaly cross impact regression (47). For an average 5×5 asset, idiosyncratic flows explain about 3% of return variations. Recall that Table 2 shows that raw flows explain about 10% of return variations. The multipliers of idiosyncratic flows mostly drop in magnitude compared with those of raw flows. Some multipliers in Table 6 even have insignificantly negative point estimates. Consistently, our second-stage panel regression has an R^2 of about 7%. Our three-price-of-risk model explains about 70% of cross-sectional price impacts.

4.6 Maximum Price Impact Ratio Portfolio

In this subsection, we study the maximum price impact ratio portfolio.

Table 6. Anomaly price impact and anomaly cross impact

	anomaly price impact					**anomaly cross impact**				
	Low	2	3	4	High	Low	2	3	4	High
	regression R^2					regression R^2				
Small	1.08%	1.17%	2.19%	2.22%	1.51%	2.23%	1.58%	2.63%	2.37%	2.34%
2	3.23%	2.43%	1.73%	1.43%	0.66%	2.10%	2.69%	2.39%	1.48%	1.22%
3	3.02%	2.09%	2.27%	1.05%	0.78%	4.04%	2.89%	2.07%	2.06%	1.33%
4	6.45%	4.82%	3.72%	4.00%	0.55%	3.47%	3.88%	4.31%	1.89%	0.73%
Big	0.21%	0.30%	0.05%	0.34%	0.81%	3.13%	2.50%	0.61%	1.31%	0.19%
	multiplier $\check{\eta}$					multiplier $\check{\phi}$				
Small	9.04	7.45	8.15	7.85	6.63	13.74	9.47	9.86	8.64	10.16
2	12.71	9.00	6.82	6.68	5.66	12.98	11.23	10.02	7.96	8.96
3	12.62	8.18	9.51	6.23	6.31	17.30	12.40	9.66	11.21	10.10
4	21.00	16.87	15.75	19.16	6.40	21.11	19.46	23.29	15.54	10.08
Big	10.01	-9.98	3.84	-11.58	10.21	23.44	27.43	13.98	19.48	6.75
	$t(\check{\eta})$					$t(\check{\phi})$				
Small	1.57	1.59	2.61	2.62	2.01	2.10	1.84	2.62	2.75	2.55
2	2.61	2.30	2.09	2.20	1.32	2.07	2.40	2.81	2.20	2.14
3	2.53	2.48	2.77	1.90	1.53	2.68	2.88	2.57	2.80	2.04
4	3.71	3.08	2.61	2.78	1.32	2.77	2.86	3.18	2.23	1.33
Big	0.59	-0.72	0.31	-0.88	1.00	2.62	2.39	1.17	1.70	0.73

Notes: In this table we report the R^2, point estimates, and t-statistics for anomaly price impact regression (46) and anomaly cross impact regression (47). The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). The idiosyncratic flow is the residual from the first-stage flow regression. To obtain the anomaly price impact, we run asset-by-asset time-series regressions of each asset's demeaned return on its idiosyncratic flow. To obtain the anomaly cross impact, we run asset-by-asset time-series regressions of each asset's demeaned return on its neighboring assets' average idiosyncratic flow. We bootstrap 10,000 times to obtain the t-statistics.

Under our factor model (43), the price impact component of asset n in month t is

$$\bar{r}_{n,t} = \sum_{k=1}^{3} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right), \tag{48}$$

which is the regression fitted value without the error term. We form the cross section of price impacts in month t as an $N \times 1$ vector $\bar{\mathbf{r}}_t = (\bar{r}_{1,t}, \bar{r}_{2,t}, \ldots, \bar{r}_{N,t})^\top$. The portfolio that obtains the maximum price impact ratio, or equivalently, experiences the most severe model-implied flow-induced price dislocation $\mathbf{c}^\top \bar{\mathbf{r}}_t$ relative to its fundamental risk $\sigma(\mathbf{c}^\top \mathbf{R}_0)$

$$\max_{\mathbf{c}\in\mathbb{R}^N} \frac{\mathbf{c}^\top \bar{\mathbf{r}}_t}{\sigma(\mathbf{c}^\top \mathbf{R}_0)} \tag{49}$$

in month t is

$$\tilde{\mathbf{c}}_t^* = \sum_{k=1}^{3} \tilde{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k. \tag{50}$$

Equation (50) shows that the MPIR portfolio $\tilde{\mathbf{c}}_t^*$ is a linear combination of the orthogonalized factor portfolio $\tilde{\mathbf{b}}_k$, with weights equal to the product of the price of risk $\tilde{\lambda}_k$ and orthogonalized factor flow $\tilde{q}_{k,t}$ in month t.

To empirically study the price dislocation of the MPIR portfolio, we calculate the return on each asset n from month t to month $t+i$, which we denote as $r_{n,t\to t+i}$. We remove the time-series mean of $r_{n,t\to t+i}$ for each asset n and holding horizon i, so that our calculation is not confounded by Sharpe ratio. We write the demeaned $r_{n,t\to t+i}$ as an $N\times 1$ vector $\mathbf{r}_{t\to t+i}$. The price dislocation of the MPIR portfolio $\tilde{\mathbf{c}}_t^*$ from month $t-1$ to $t+i$ (normalized by its fundamental risk) is defined as

$$\kappa_{t-1\to t+i} = \frac{(\tilde{\mathbf{c}}_t^*)^\top (\bar{\mathbf{r}}_t + \mathbf{r}_{t\to t+i})}{\sigma((\tilde{\mathbf{c}}_t^*)^\top \mathbf{R}_0)}. \tag{51}$$

The solid line in Figure 4 shows the means of $\kappa_{t-1\to t+i}$ over all months t, and the shaded area represents the 95% confidence intervals of the means. Our model has an average MPIR of about 0.1 in the current month from $t-1$ to t . The price dislocation completely reverts in about six months. If liquidity providers trade against the MPIR portfolio $\tilde{\mathbf{c}}_t^*$ and hold the position for six months, they earn, on average, an annualized[32] price impact ratio of about 0.2 (we later study the time-series variation in MPIR). This reversion of price dislocation is consistent with the interpretation that our factor-model-implied MPIR portfolio reflects liquidity providers' aversion to absorbing fundamental risks.

[32]In definition (51), the denominator is always the annualized return risk of portfolio $\tilde{\mathbf{c}}_t^*$ for varying horizons i. To annualize the ratio $\kappa_{t-1\to t+i}$, we annualize the numerator.

Figure 4. Average price dislocation of the MPIR portfolio

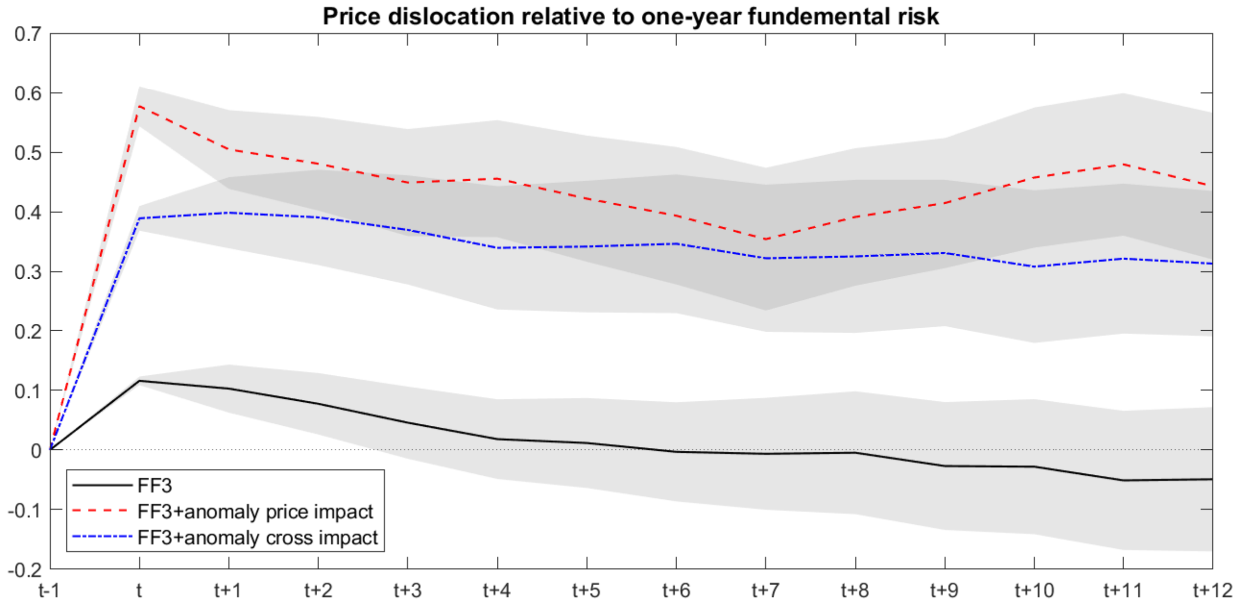

Notes: This figure shows the average price dislocation from month $t-1$ to $t+i$ of the MPIR portfolio $\tilde{\mathbf{c}}_t^*$ (normalized by its fundamental risk) for our price impact model with MKT, SMB, and HML factor flows (FF3), FF3 plus anomaly price impact $\check{\eta}$, and FF3 plus anomaly cross impact $\check{\phi}$. The shaded area represents the 95% confidence intervals of the means. We measure fundamental risk using one-year return. For each model, we compute the MPIR portfolio $\tilde{\mathbf{c}}_t^*$ using equation (50). We then compute the normalized price dislocation of $\tilde{\mathbf{c}}_t^*$ from month $t-1$ to $t+i$ using equation (51). Lastly, we compute the time-series means of the normalized price dislocation and the standard deviation of the mean for each horizon i.

We then diagnose the anomaly price impact $\check{\eta}$ in Table 6. The flow-based Gibbons-Ross-Shanken test in Section 3.3 implies that instead of studying $\check{\eta}$ one-by-one, it suffices to study the MPIR portfolio constructed using our factor plus anomaly impacts. The cross section of price impacts in month t under the model with anomaly price impact (45) is an $N \times 1$ vector $\check{\mathbf{r}}_t = (\check{r}_{1,t}, \check{r}_{2,t}, \ldots, \check{r}_{N,t})^\top$, where

$$\check{r}_{n,t} = \check{\eta}_n e_{n,t}/w_n + \sum_{k=1}^{3} \tilde{\lambda}_k \tilde{q}_{k,t} \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right). \tag{52}$$

We compute the new MPIR portfolio $\tilde{\mathbf{c}}_t^* = \text{var}(\mathbf{R}_0)^{-1} \check{\mathbf{r}}_t$ and then follow the same procedure as before. The computation of our factor model plus anomaly cross impact $\check{\phi}$ is similar.

Figure 4 shows the time-series average price dislocation of MPIR portfolio $\tilde{\mathbf{c}}_t^*$ under these two types of anomaly impacts. The MPIR from month $t-1$ to t increases significantly relative to our factor model. Our factor model has only 3 free parameters, and the anomaly model has $3 + 25$ parameters. Under the anomaly model, one can find some MPIR portfolio with a large price dislocation from month $t-1$ to t but small fundamental risk. The price dislocation of the anomaly MPIR portfolio barely reverts, however, even after one year. The lack of price reversion suggests that the anomaly price impacts caused by idiosyncratic flows are not caused by liquidity providers' risk aversion.[33] The sharp contrast in the price reversion between factor and idiosyncratic flows supports our factor-model approach to identify the MPIR portfolio.

Decomposition of the maximum price impact ratio. We use MPIR to decompose how the MKT, SMB, and HML factors explain the cross section of price impacts. Under our factor model, the cross section of price impacts is $\bar{\mathbf{r}}_t = (\bar{r}_{1,t}, \bar{r}_{2,t}, \ldots, \bar{r}_{N,t})^\top$, where

$$\bar{r}_{n,t} = \sum_{k,l \in \{\text{MKT,SMB,HML}\}} \lambda_{k,l} q_{l,t} \left(\sum_{m=1}^{N} v_{n,m} b_{m,k} \right). \tag{53}$$

This equation is obtained from (48) through rotation to the MKT, SMB, and HML factors.

For any subset of factors $\mathcal{L} \subseteq \{\text{MKT,SMB,HML}\}$, we compute the price impact

$$\bar{r}_{n,t} = \sum_{k \in \{\text{MKT,SMB,HML}\}} \sum_{l \in \mathcal{L}} \lambda_{k,l} q_{l,t} \left(\sum_{m=1}^{N} v_{n,m} b_{m,k} \right). \tag{54}$$

Equation (54) shuts down the flow of factors other than $\mathcal{L}$. Nevertheless, the $\mathcal{L}$ flows could still impact the prices of all of the MKT, SMB, and HML factors. Using (54), we compute the MPIR for each month t.

Figure 5 shows the time series of MPIR for MKT+SMB+HML, MKT+SMB, and MKT. The full model (MKT+SMB+HML) exhibits a pretty volatile MPIR. In periods with extreme flows, the MPIR can be as high as 0.4. That is, the portfolio's price dislocation is about

[33] We cannot give a positive answer on the cause of non-reverting anomaly impacts. Conjectures include the permanent demand effects of Koijen and Yogo (2019) and Gabaix and Koijen (2022). It is also possible that although mutual-fund flow-induced trading is largely uninformed at factor levels, this trading might contain some information at specific 5×5 asset level, causing the permanent impacts.

Figure 5. Decomposition of the maximum price impact ratio

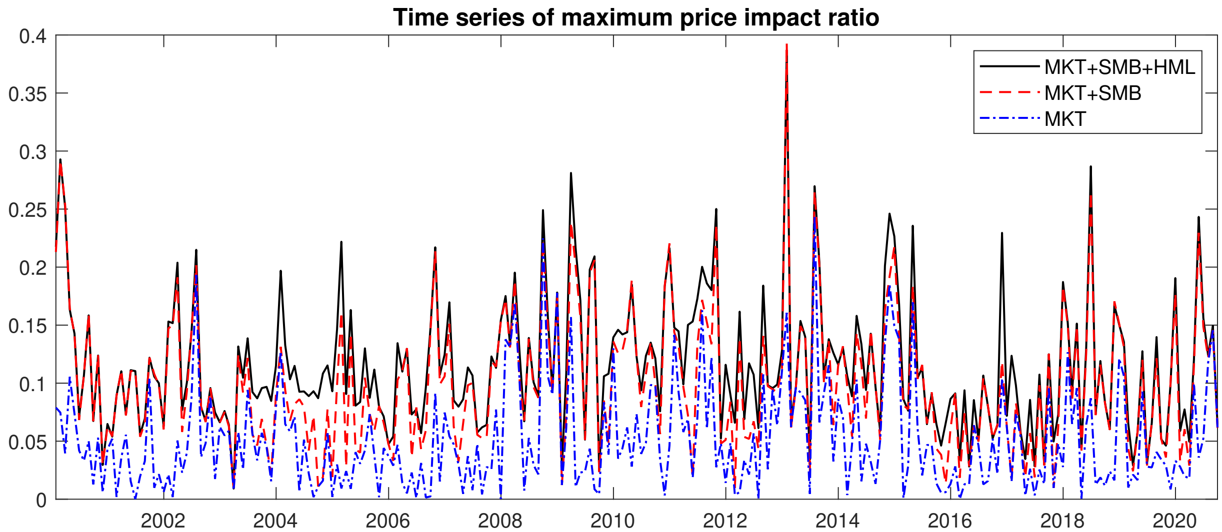

time-series mean of the maximum price impact ratio							
	MKT	SMB	HML	MKT+SMB	MKT+HML	SMB+HML	MKT+SMB+HML
mean	0.051	0.076	0.042	0.101	0.074	0.096	0.116
t-stats	(17.94)	(20.58)	(18.31)	(26.97)	(23.63)	(27.81)	(31.91)

Notes: We compute the MPIR of a given subset of factors $\mathcal{L} \subseteq \{\text{MKT,SMB,HML}\}$. We shut down the flow of factors other than $\mathcal{L}$ and compute the corresponding cross section of price impacts. We then compute the corresponding MPIR for each month t. The top figure plots the time series of MPIR and the bottom table provides the time-series mean of MPIR and the t-statistics of the mean.

40% of its annual risk. This is a very large price dislocation for liquidity providers to trade against. Not all explanatory power for the cross section of price impacts arises from the MKT factor. The SMB factor also contributes greatly.

The bottom panel in Figure 5 shows the time-series means of MPIR for all possible model combinations. The SMB factor provides the highest explanatory power for the cross section of price impacts, and the MKT and HML factors have roughly the same explanatory power. This fact may be somewhat surprising because MKT flow is about six times as volatile as the SMB and HML flows (Table 4). The reason is that the SMB and HML factors have higher prices of risk than the MKT factor (Table 5), and the quantity of risk induced by every dollar of flow also differs. Not all flows are created equal—the price of risk and the quantity of risk matter for the cross section of price impacts.

5 Conclusion

We show that in any market with uninformative flows, the maximum Sharpe-ratio portfolio can be separated into two portfolios. The first portfolio uses only fundamental information to maximize Sharpe ratio. The second portfolio provides liquidity to uninformative flows and maximizes price impact ratio, which is defined as a portfolio's price impact over its fundamental risk. The rational investor allocates risk exposure between the two portfolios, proportional to their respective Sharpe ratio and price impact ratio.

We develop the factor model of price impacts to empirically investigate the maximum-price-impact-ratio (MPIR) portfolio. We develop the flow-based Fama-MacBeth regression to estimate the MPIR portfolio and the flow-based Gibbons-Ross-Shanken test to the diagnose MPIR portfolio. For U.S. equity mutual fund flows, we find that the MPIR portfolio constructed using flows into Fama and French (1993) factors is a good choice.

References

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.

An, Yu, 2022, Flow-based arbitrage pricing theory, Working paper, Johns Hopkins University.

Andersen, Leif, Darrell Duffie, and Yang Song, 2019, Funding value adjustments, *Journal of Finance* 74, 145–192.

Anderson, Alyssa G, Wenxin Du, and Bernd Schlusche, 2021, Arbitrage capital of global banks, Working paper, University of Chicago.

Andrade, Sandro C, Charles Chang, and Mark S Seasholes, 2008, Trading imbalances, predictable reversals, and cross-stock price pressure, *Journal of Financial Economics* 88, 406–423.

Ang, Andrew, Robert J Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.

Asness, Clifford S, Tobias J Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929–985.

Balasubramaniam, Vimal, John Y Campbell, Tarun Ramadorai, and Benjamin Ranish, 2021, Who owns what? A factor model for direct stockholding, *Journal of Finance* Forthcoming.

Bali, Turan G, and Nusret Cakici, 2008, Idiosyncratic volatility and the cross section of expected returns, *Journal of Financial and Quantitative Analysis* 43, 29–58.

Ben-David, Itzhak, Jiacui Li, Andrea Rossi, and Yang Song, 2021, Ratings-driven demand and systematic price fluctuations, *Review of Financial Studies* Forthcoming.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.

Betermier, Sebastien, Laurent E Calvet, Samuli Knüpfer, and Jens Kvaerner, 2021, What do the portfolios of individual investors reveal about the cross-section of equity returns? Working paper, McGill University.

Bouchaud, Jean-Philippe, Julius Bonart, Jonathan Donier, and Martin Gould, 2018, *Trades, quotes and prices: financial markets under the microscope* (Cambridge University Press).

Boulatov, Alex, Terrence Hendershott, and Dmitry Livdan, 2013, Informed trading and portfolio returns, *Review of Economic Studies* 80, 35–72.

Brunnermeier, Markus K, and Yuliy Sannikov, 2014, A macroeconomic model with a financial sector, *American Economic Review* 104, 379–421.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2021, Forest through the trees: Building cross-sections of stock returns, Working paper, London Business School.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, 2015, Regression discontinuity and the price effects of stock market indexing, *Review of Financial Studies* 28, 212–246.

Chen, Andrew Y, 2021, The limits of p-hacking: Some thought experiments, *Journal of Finance* 76, 2447–2480.

Cochrane, John H, 2009, *Asset pricing: Revised edition* (Princeton university press).

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

Cremers, K. J. Martijn, and Jianping Mei, 2007, Turning over Turnover, *The Review of Financial Studies* 20, 1749–1782.

Da, Zhi, Borja Larrain, Clemens Sialm, and José Tessada, 2018, Destabilizing financial advice: Evidence from pension fund reallocations, *Review of Financial Studies* 31, 3720–3755.

Daniel, Kent, and Sheridan Titman, 1997, Evidence on the characteristics of cross sectional variation in stock returns, *Journal of Finance* 52, 1–33.

Dou, Winston, Leonid Kogan, and Wei Wu, 2021, Common fund flows: Flow hedging and factor pricing, Working paper, University of Pennsylvania.

Du, Wenxin, Benjamin M Hébert, and Amy Wang Huber, 2021, Are intermediary constraints priced? Working paper, University of Chicago.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2018, Deviations from covered interest rate parity, *Journal of Finance* 73, 915–957.

Duffie, Darrell, 2012, Market making under the proposed Volcker rule, Working paper, Stanford University.

Egan, Mark L, Alexander MacKay, and Hanbin Yang, 2021, What drives variation in investor portfolios? evidence from retirement plans, Working paper, Harvard University.

Fama, Eugene F, and Kenneth R French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F, and James D MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Frazzini, Andrea, Ronen Israel, and Tobias J Moskowitz, 2018, Trading costs, Working paper, Yale University.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gabaix, Xavier, and Ralph SJ Koijen, 2021, Granular instrumental variables, Working paper, Harvard University.

Gabaix, Xavier, and Ralph SJ Koijen, 2022, In search of the origins of financial fluctuations: The inelastic markets hypothesis, Working paper, Harvard University.

Gibbons, Michael R, Stephen A Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 1121–1152.

Giglio, Stefano, Bryan Kelly, and Dacheng Xiu, 2021, Factor models, machine learning, and asset pricing, *Annual Review of Financial Economics* Forthcoming.

Greenwood, Robin, 2005, Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage, *Journal of Financial Economics* 75, 607–649.

Greenwood, Robin, and David Thesmar, 2011, Stock price fragility, *Journal of Financial Economics* 102, 471–490.

Haddad, Valentin, Paul Huebner, and Erik Loualiche, 2022, How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing, Working paper, UCLA.

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Harvey, Campbell R, Yan Liu, and Heqing Zhu, 2016, . . . and the cross-section of expected returns, *Review of Financial Studies* 29, 5–68.

Hasbrouck, Joel, 1988, Trades, quotes, inventories, and information, *Journal of Financial Economics* 22, 229–252.

Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.

Hasbrouck, Joel, and Duane J Seppi, 2001, Common factors in prices, order flows, and liquidity, *Journal of Financial Economics* 59, 383–411.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

He, Zhiguo, Paymon Khorrami, and Zhaogang Song, 2021, Commonality in credit spread changes: Dealer inventory and intermediary distress, *Review of Financial Studies* Forthcoming.

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–70.

Holcblat, Benjamin, Abraham Lioui, and Michael Weber, 2022, Anomaly or possible risk factor? simple-to-use tests, Working paper, University of Chicago.

Hou, Kewei, Chen Xue, and Lu Zhang, 2020, Replicating anomalies, *Review of Financial Studies* 33, 2019–2133.

Huang, Shiyang, Yang Song, and Hong Xiang, 2021, Noise trading and asset pricing factors, Working paper, The University of Hong Kong.

Investment Company Institute, 2021, Investment company fact book, https://www.ici.org/system/files/2021-05/2021_factbook.pdf.

Kelly, Bryan T, Seth Pruitt, and Yinan Su, 2019, Characteristics are covariances: A unified model of risk and return, *Journal of Financial Economics* 134, 501–524.

Koijen, Ralph S. J., and Stijn van Nieuwerburgh, 2011, Predictability of returns and cash flows, *Annual Review of Financial Economics* 3, 467–491.

Koijen, Ralph SJ, Robert J Richmond, and Motohiro Yogo, 2021, Which investors matter for equity valuations and expected returns? Working paper, University of Chicago.

Koijen, Ralph SJ, and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.

Lettau, Martin, and Markus Pelger, 2020, Factors that fit the time series and cross-section of stock returns, *Review of Financial Studies* 33, 2274–2325.

Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Li, Jennifer Jie, Neil D Pearson, and Qi Zhang, 2020, Impact of demand shocks on the stock market: Evidence from Chinese IPOs, Working paper, INSEAD.

Li, Jiacui, 2021, What drives the size and value factors? Working paper, University of Utah.

Li, Jiacui, and Zihan Lin, 2022, Prices are less elastic at more aggregate levels, Working paper, University of Utah.

Linnainmaa, Juhani T, and Michael R Roberts, 2018, The history of the cross-section of stock returns, *Review of Financial Studies* 31, 2606–2649.

Lo, Andrew W, and Jiang Wang, 2000, Trading volume: definitions, data analysis, and implications of portfolio theory, *Review of Financial Studies* 13, 257–300.

Lou, Dong, 2012, A flow-based explanation for return predictability, *Review of Financial Studies* 25, 3457–3489.

Nagel, Stefan, 2021, *Machine learning in asset pricing* (Princeton University Press).

Novy-Marx, Robert, and Mihail Velikov, 2022, Betting against betting against beta, *Journal of Financial Economics* 143, 80–106.

Pasquariello, Paolo, and Clara Vega, 2015, Strategic cross-trading in the US stock market, *Review of Finance* 19, 229–282.

Pavlova, Anna, and Taisiya Sikorskaya, 2022, Benchmarking intensity, Working paper, London Business School.

Peng, Cameron, and Chen Wang, 2021, Factor demand and factor returns, Working paper, LSE.

Rapach, David E, Jack K Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.

Ross, Stephen A, 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–60.

Schmickler, Simon, 2020, Identifying the price impact of fire sales using high-frequency surprise mutual fund flows, Working paper, Princeton University.

Shanken, Jay, 1992, On the estimation of beta-pricing models, *Review of Financial Studies* 5, 1–33.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Stambaugh, Robert F, and Yu Yuan, 2017, Mispricing factors, *Review of Financial Studies* 30, 1270–1315.

Warther, Vincent A., 1995, Aggregate mutual fund flows and security returns, *Journal of Financial Economics* 39, 209–235.

Appendices

The appendices provide additional theoretical and empirical details.

A Proofs

In this appendix, we provide proofs omitted in the main text.

A.1 Proof of Theorem 1

First, we have

$$\mathrm{var}(\mathbf{R}(\mathbf{0})) = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))\mathrm{var}(\mathbf{R}(\mathbf{f}))\mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f})). \tag{A.1}$$

This equation simplifies to

$$\mathrm{var}(\mathbf{R}(\mathbf{f}))^{-1} = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))\mathrm{var}(\mathbf{R}(\mathbf{0}))^{-1}\mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f})) \tag{A.2}$$

Second, we have

$$\mathbf{R}(\mathbf{0}) = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))\mathbf{R}(\mathbf{f}). \tag{A.3}$$

This equation implies that

$$\mathbf{R}(\mathbf{0}) - R_F\mathbf{1} + R_F\mathbf{1} = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))(\mathbf{R}(\mathbf{f}) - R_F\mathbf{1} + R_F\mathbf{1}), \tag{A.4}$$

which further simplifies to

$$\mathbf{R}(\mathbf{f}) - R_F\mathbf{1} = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))^{-1}(\mathbf{R}(\mathbf{0}) - R_F\mathbf{1} - R_F\Delta\mathbf{p}(\mathbf{f})). \tag{A.5}$$

Taking expectation on both sides, we have

$$\mathbb{E}[\mathbf{R}(\mathbf{f})] - R_F\mathbf{1} = \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))^{-1}(\mathbb{E}[\mathbf{R}(\mathbf{0})] - R_F\mathbf{1} - R_F\Delta\mathbf{p}(\mathbf{f})). \tag{A.6}$$

Therefore, we have

$$\begin{aligned}\mathbf{c}^*(\mathbf{f}) &= \mathrm{var}(\mathbf{R}(\mathbf{f}))^{-1}(\mathbb{E}[\mathbf{R}(\mathbf{f})] - R_F\mathbf{1}) \\ &= \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))\mathrm{var}(\mathbf{R}(\mathbf{0}))^{-1}(\mathbb{E}[\mathbf{R}(\mathbf{0})] - R_F\mathbf{1} - R_F\Delta\mathbf{p}(\mathbf{f})) \\ &= \mathrm{diag}(\mathbf{1} + \Delta\mathbf{p}(\mathbf{f}))(\mathbf{c}^*(\mathbf{0}) - R_F\tilde{\mathbf{c}}^*(\mathbf{f})),\end{aligned} \tag{A.7}$$

where the first equality uses definition (5), the second equality uses (A.2) and (A.6), and the third equality uses definitions (5) and (6). Using definitions (7) and (8), we see that (A.7) is equivalent to (9).

The return volatility of portfolio $\mathbf{h}^*(\mathbf{0})$ is

$$\begin{aligned}\sigma(\mathbf{h}^*(\mathbf{0})^\top \mathbf{X}) &= \sigma(\mathbf{c}^*(\mathbf{0})^\top \mathbf{R}(\mathbf{0})) = \sqrt{\mathbf{c}^*(\mathbf{0})^\top \mathrm{var}(\mathbf{R}(\mathbf{0}))\mathbf{c}^*(\mathbf{0})} \\ &= \sqrt{\mathbb{E}[\mathbf{R}(\mathbf{0}) - R_F\mathbf{1}]^\top \mathrm{var}(\mathbf{R}(\mathbf{0}))^{-1}\mathbb{E}[\mathbf{R}(\mathbf{0}) - R_F\mathbf{1}]},\end{aligned} \tag{A.8}$$

which equals the maximum Sharpe ratio without flow by definition (5).

Similarly, the return volatility of portfolio $\tilde{\mathbf{h}}^*(\mathbf{f})$ is

$$\sigma(\tilde{\mathbf{h}}^*(\mathbf{f})^\top \mathbf{X}) = \sigma(\tilde{\mathbf{c}}^*(\mathbf{f})^\top \mathbf{R}(\mathbf{0})) = \sqrt{\tilde{\mathbf{c}}^*(\mathbf{f})^\top \mathrm{var}(\mathbf{R}(\mathbf{0}))\tilde{\mathbf{c}}^*(\mathbf{f})} = \sqrt{\Delta\mathbf{p}(\mathbf{f})^\top \mathrm{var}(\mathbf{R}(\mathbf{0}))^{-1}\Delta\mathbf{p}(\mathbf{f})}, \tag{A.9}$$

which equals the maximum price impact ratio by definition (6).

A.2 Details for Solving the Rotation Matrix

Our goal is to find some $K \times K$ invertible matrix $\mathbf{O}$, such that

$$\mathbf{O}^\top \mathbf{B}^\top \mathrm{var}(\mathbf{R}_0)\mathbf{B}\mathbf{O} = \mathbf{I}_K, \tag{A.10}$$

$$\mathbf{O}\mathbf{\Pi}\mathbf{O}^\top = \mathrm{var}(\mathbf{q}) = \frac{1}{T}\begin{pmatrix} \sum_{t=1}^T q_{1,t}^2 & \sum_{t=1}^T q_{1,t}q_{2,t} & \cdots & \sum_{t=1}^T q_{1,t}q_{K,t} \\ \sum_{t=1}^T q_{2,t}q_{1,t} & \sum_{t=1}^T q_{2,t}^2 & \cdots & \sum_{t=1}^T q_{2,t}q_{K,t} \\ \cdots & \cdots & \cdots & \cdots \\ \sum_{t=1}^T q_{K,t}q_{1,t} & \sum_{t=1}^T q_{K,t}q_{2,t} & \cdots & \sum_{t=1}^T q_{K,t}^2 \end{pmatrix}, \tag{A.11}$$

where $\mathbf{\Pi} = \mathrm{diag}(\pi_1, \pi_2, \ldots, \pi_K)$ is some $K \times K$ diagonal matrix.

First, we carry out Cholesky decomposition and obtain

$$\mathbf{B}^\top \mathrm{var}(\mathbf{R}_0)\mathbf{B} = \mathbf{U}^\top \mathbf{U}, \tag{A.12}$$

where $\mathbf{U}$ is an $K \times K$ upper triangular matrix. Second, we carry out eigenvalue decomposition

$$\left(\mathbf{U}\mathrm{var}(\mathbf{q})\mathbf{U}^\top\right)\mathbf{G} = \mathbf{G}\mathbf{\Pi}, \tag{A.13}$$

where $\mathbf{\Pi} = \mathrm{diag}(\pi_1, \pi_2, \ldots, \pi_K)$, and $\mathbf{G}$ is an orthogonal $K \times K$ matrix satisfying $\mathbf{G}^\top \mathbf{G} = \mathbf{I}_K$.

We claim that $\mathbf{O} = \mathbf{U}^{-1}\mathbf{G}$ satisfies the orthogonalization conditions (A.10) and (A.11).

First, we see that

$$\mathbf{O}^\top \mathbf{B}^\top \text{var}(\mathbf{R}_0)\mathbf{B}\mathbf{O} = \mathbf{G}^\top (\mathbf{U}^\top)^{-1}\mathbf{U}^\top \mathbf{U}\mathbf{U}^{-1}\mathbf{G} = \mathbf{I}_K. \tag{A.14}$$

Second, because $\mathbf{q} = \mathbf{O}\tilde{\mathbf{q}}$, we have

$$\text{var}(\mathbf{q}) = \mathbf{O}\text{var}(\tilde{\mathbf{q}})\mathbf{O}^\top. \tag{A.15}$$

From (A.13), we have

$$\mathbf{U}\text{var}(\mathbf{q})\mathbf{U}^\top \mathbf{U}\mathbf{O} = \mathbf{U}\mathbf{O}\boldsymbol{\Pi}. \tag{A.16}$$

Eliminating the invertible matrix $\mathbf{U}$ on both sides, we obtain

$$\text{var}(\mathbf{q})\mathbf{U}^\top \mathbf{U}\mathbf{O} = \mathbf{O}\boldsymbol{\Pi}. \tag{A.17}$$

Plugging (A.15) into (A.17), we obtain

$$\mathbf{O}\boldsymbol{\Pi} = \mathbf{O}\text{var}(\tilde{\mathbf{q}})\mathbf{O}^\top \mathbf{U}^\top \mathbf{U}\mathbf{O} = \mathbf{O}\text{var}(\tilde{\mathbf{q}}), \tag{A.18}$$

where we have used $\mathbf{O}^\top \mathbf{U}^\top \mathbf{U}\mathbf{O} = \mathbf{G}^\top \mathbf{G} = \mathbf{I}_K$. We therefore obtain

$$\text{var}(\tilde{\mathbf{q}}) = \boldsymbol{\Pi}, \tag{A.19}$$

which proves (A.11).

A.3 Proof of Theorem 2

First, we simplify the χ^2 test statistics in (31). We write a $1\times N$ vector $\mathbf{e}_t = (e_{1,t}, e_{2,t}, \ldots, e_{N,t})$ and an $N \times 1$ vector $\mathbf{a}_n = (a_{n,1}, a_{n,2}, \ldots, a_{n,N})^\top$, and we write regression (20) as

$$r_{n,t} = \mathbf{e}_t \mathbf{a}_n + \sum_{k=1}^{K} \tilde{\lambda}_k \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) \tilde{q}_{k,t} + \xi_{n,t}. \tag{A.20}$$

We define the $T\times N$ matrix $\mathbf{e} = (\mathbf{e}_1; \mathbf{e}_2; \ldots; \mathbf{e}_T)$ and the $N\times 1$ vector $\boldsymbol{\xi}_n = (\xi_{n,1}, \xi_{n,2}, \ldots, \xi_{n,T})^\top$. We show in Appendix A.4 that we can run an asset-by-asset time-series regression to obtain the point estimator of $\mathbf{a}_n$ as

$$\hat{\mathbf{a}}_n = (\mathbf{e}^\top \mathbf{e})^{-1}\mathbf{e}^\top \begin{pmatrix} r_{n,1} \\ r_{n,2} \\ \cdots \\ r_{n,T} \end{pmatrix} = \mathbf{a}_n + (\mathbf{e}^\top \mathbf{e})^{-1}\mathbf{e}^\top \boldsymbol{\xi}_n, \tag{A.21}$$

where we have used the fact that, for any $n = 1, 2, \ldots, N$ and $k = 1, 2, \ldots, K$,

$$\sum_{t=1}^{T} e_{n,t}\tilde{q}_{k,t} = 0, \tag{A.22}$$

because $e_{n,t}$ is the residual from the first-stage regression.

Therefore, we have, for any m and n,

$$\text{cov}(\hat{\mathbf{a}}_n, \hat{\mathbf{a}}_m) = (\mathbf{e}^\top \mathbf{e})^{-1}\mathbf{e}^\top \text{cov}(\boldsymbol{\xi}_n, \boldsymbol{\xi}_m)\mathbf{e}(\mathbf{e}^\top \mathbf{e})^{-1} = (\mathbf{e}^\top \mathbf{e})^{-1}\boldsymbol{\Sigma}_\xi(n, m), \tag{A.23}$$

where we have used the assumption that $\boldsymbol{\xi}_t$ is i.i.d. over time. The term $\boldsymbol{\Sigma}_\xi(n, m)$ is the (n, m)-th element of the cross-sectional variance-covariance matrix of $\boldsymbol{\xi}_t$. When constructing the χ^2 test statistic, we use the asymptotically consistent sample variance-covariance matrix $\hat{\boldsymbol{\Sigma}}_\xi$ for the true $\boldsymbol{\Sigma}_\xi$. We denote $\hat{\mathbf{a}} = (\hat{\mathbf{a}}_1; \hat{\mathbf{a}}_2; \ldots; \hat{\mathbf{a}}_N)$ as the $N^2 \times 1$ vector of parameter estimates. By equation (A.23), we have

$$\text{var}(\hat{\mathbf{a}}) = \hat{\boldsymbol{\Sigma}}_\xi \otimes (\mathbf{e}^\top \mathbf{e})^{-1}, \tag{A.24}$$

where $\otimes$ represents the Kronecker product. Therefore, we have

$$\hat{\mathbf{a}}^\top \frac{\text{var}(\hat{\mathbf{a}})^{-1}}{T} \hat{\mathbf{a}} = \hat{\mathbf{a}}^\top \left(\hat{\boldsymbol{\Sigma}}_\xi^{-1} \otimes (\mathbf{e}^\top \mathbf{e}/T) \right) \hat{\mathbf{a}}. \tag{A.25}$$

Under the null hypothesis of $\mathbf{a} = \mathbf{0}$, we have

$$
\begin{aligned}
&\hat{\mathbf{a}}^\top \left(\hat{\boldsymbol{\Sigma}}_\xi^{-1} \otimes (\mathbf{e}^\top \mathbf{e}/T) \right) \hat{\mathbf{a}} \\
&= \sum_{n=1}^{N} \sum_{m=1}^{N} \left((\mathbf{e}^\top \mathbf{e})^{-1} \mathbf{e}^\top \boldsymbol{\xi}_n \right)^\top \hat{\boldsymbol{\Sigma}}_\xi^{-1}(n,m) (\mathbf{e}^\top \mathbf{e}/T) (\mathbf{e}^\top \mathbf{e})^{-1} \mathbf{e}^\top \boldsymbol{\xi}_m && \text{(A.26)} \\
&= \frac{1}{T} \sum_{n=1}^{N} \sum_{m=1}^{N} \hat{\boldsymbol{\Sigma}}_\xi^{-1}(n,m) \boldsymbol{\xi}_n^\top \mathbf{e} (\mathbf{e}^\top \mathbf{e})^{-1} \mathbf{e}^\top \boldsymbol{\xi}_m \\
&= \frac{1}{T} \sum_{n=1}^{N} \sum_{m=1}^{N} \hat{\boldsymbol{\Sigma}}_\xi^{-1}(n,m) \boldsymbol{\psi}_n^\top \boldsymbol{\psi}_m && \text{(A.27)} \\
&= \frac{1}{T} \begin{pmatrix} \boldsymbol{\psi}_1 \\ \boldsymbol{\psi}_2 \\ \dots \\ \boldsymbol{\psi}_N \end{pmatrix}^\top \left(\hat{\boldsymbol{\Sigma}}_\xi^{-1} \otimes \mathbf{I}_T \right) \begin{pmatrix} \boldsymbol{\psi}_1 \\ \boldsymbol{\psi}_2 \\ \dots \\ \boldsymbol{\psi}_N \end{pmatrix}, && \text{(A.28)}
\end{aligned}
$$

where we define

$$\boldsymbol{\psi}_n = \mathbf{e}(\mathbf{e}^\top \mathbf{e})^{-1} \mathbf{e}^\top \boldsymbol{\xi}_n, \tag{A.29}$$

as the projection of $\boldsymbol{\xi}_n$ onto the idiosyncratic flow space. In step (A.26), we use block matrix multiplication for every N elements and $\hat{\boldsymbol{\Sigma}}_\xi^{-1}(n,m)$ is the (n,m)-th element of $\hat{\boldsymbol{\Sigma}}_\xi^{-1}$. In step (A.27), we use the fact that the projection matrix $\mathbf{e}(\mathbf{e}^\top \mathbf{e})^{-1}\mathbf{e}^\top$ is idempotent. In step (A.28), we use the block-matrix multiplication in reverse direction, with each $\boldsymbol{\psi}_n$ as a $T \times 1$ vector.

We define $\boldsymbol{\psi}_t = (\psi_{1,t}, \psi_{2,t}, \dots, \psi_{N,t})^\top$. In this way, $\boldsymbol{\psi}_n$ is the time-series variation in $\psi_{n,t}$ for a given asset n, and $\boldsymbol{\psi}_t$ is the cross-sectional variation in $\psi_{n,t}$ for a given time t. By rearranging $\boldsymbol{\psi}_n$ into $\boldsymbol{\psi}_t$, we have

$$\frac{1}{T} \begin{pmatrix} \boldsymbol{\psi}_1 \\ \boldsymbol{\psi}_2 \\ \dots \\ \boldsymbol{\psi}_N \end{pmatrix}^\top \left(\hat{\boldsymbol{\Sigma}}_\xi^{-1} \otimes \mathbf{I}_T \right) \begin{pmatrix} \boldsymbol{\psi}_1 \\ \boldsymbol{\psi}_2 \\ \dots \\ \boldsymbol{\psi}_N \end{pmatrix} = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \hat{\boldsymbol{\Sigma}}_\xi^{-1} \boldsymbol{\psi}_t. \tag{A.30}$$

This step may not be immediately obvious, but writing out the elements of the $NT \times NT$ matrix $\hat{\boldsymbol{\Sigma}}_\xi^{-1} \otimes \mathbf{I}_T$ helps understanding.

Because $\boldsymbol{\xi}_t$ is i.i.d. over time, the strong law of large numbers implies that the sample variance-covariance matrix $\hat{\boldsymbol{\Sigma}}_\xi$ converges to the true $\boldsymbol{\Sigma}_\xi$ almost surely as T tends to infinity. Restriction (30) implies that $\boldsymbol{\Sigma}_\xi = \text{var}(\mathbf{R}_0)/H$. Therefore, in the limit of T tending to infinity, we have almost surely,

$$\hat{\mathbf{a}}^\top \frac{\mathrm{var}(\hat{\mathbf{a}})^{-1}}{T} \hat{\mathbf{a}} = \frac{H}{T} \sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1} \boldsymbol{\psi}_t. \tag{A.31}$$

Next, we transform the maximum squared price impact ratio in (31). We define $\check{\mathbf{r}}_n = (\check{r}_{n,1}, \check{r}_{n,2}, \ldots, \check{r}_{n,T})^\top$. Using equations (28), (A.21) and (A.29), we have

$$\check{\mathbf{r}}_n = \mathbf{e}(\mathbf{e}^\top \mathbf{e})^{-1} \mathbf{e}^\top \boldsymbol{\xi}_n + \sum_{k=1}^{K} \hat{\lambda}_k \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) \tilde{\mathbf{q}}_k = \boldsymbol{\psi}_n + \sum_{k=1}^{K} \hat{\lambda}_k \left(\sum_{m=1}^{N} v_{n,m} \tilde{b}_{m,k} \right) \tilde{\mathbf{q}}_k, \tag{A.32}$$

where $\tilde{\mathbf{q}}_k = (\tilde{q}_{k,1}, \tilde{q}_{k,2}, \ldots, \tilde{q}_{k,T})^\top$. We define $\check{\mathbf{r}}_t = (\check{r}_{1,t}, \check{r}_{2,t}, \ldots, \check{r}_{N,t})^\top$. In this way, $\check{\mathbf{r}}_n$ is the time-series variation in $\check{r}_{n,t}$ for a given asset n, and $\check{\mathbf{r}}_t$ is the cross-sectional variation in $\check{r}_{n,t}$ for a given time t. We have

$$\check{\mathbf{r}}_t = \boldsymbol{\psi}_t + \mathrm{var}(\mathbf{R}_0) \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k. \tag{A.33}$$

The realized maximum squared price impact ratio at time t is

$$\begin{aligned}
&\check{\mathbf{r}}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1} \check{\mathbf{r}}_t \\
=&\left(\boldsymbol{\psi}_t^\top + \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k^\top \mathrm{var}(\mathbf{R}_0) \right) \mathrm{var}(\mathbf{R}_0)^{-1} \left(\boldsymbol{\psi}_t + \mathrm{var}(\mathbf{R}_0) \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k \right) \\
=&\boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1} \boldsymbol{\psi}_t + 2\boldsymbol{\psi}_t^\top \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k + \left(\sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k^\top \right) \mathrm{var}(\mathbf{R}_0) \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k \\
=&\boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1} \boldsymbol{\psi}_t + 2\boldsymbol{\psi}_t^\top \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k + \sum_{k=1}^{K} \hat{\lambda}_k^2 \tilde{q}_{k,t}^2,
\end{aligned} \tag{A.34}$$

where, in the last step, we use the orthogonalization $\tilde{\mathbf{B}}^\top \mathrm{var}(\mathbf{R}_0) \tilde{\mathbf{B}} = \mathbf{I}_K$. The time-series average is

$$\begin{aligned}
\hat{\theta}^{*2} =& \frac{1}{T} \sum_{t=1}^{T} \check{\mathbf{r}}_n^\top \mathrm{var}(\mathbf{R}_0)^{-1} \check{\mathbf{r}}_n \\
=& \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1} \boldsymbol{\psi}_t + \frac{2}{T} \sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \sum_{k=1}^{K} \hat{\lambda}_k \tilde{q}_{k,t} \tilde{\mathbf{b}}_k + \frac{1}{T} \sum_{t=1}^{T} \sum_{k=1}^{K} \hat{\lambda}_k^2 \tilde{q}_{k,t}^2.
\end{aligned} \tag{A.35}$$

Note that for any $n = 1, 2, \dots, N$ and $k = 1, 2, \dots, K$, we have

$$\begin{aligned}\sum_{t=1}^{T} \tilde{q}_{k,t}\psi_{n,t} &= \sum_{t=1}^{T} \tilde{q}_{k,t}\mathbf{e}_t(\mathbf{e}^\top\mathbf{e})^{-1}\mathbf{e}^\top\boldsymbol{\xi}_n \\ &= \left(\sum_{t=1}^{T} \tilde{q}_{k,t}e_{1,t}, \sum_{t=1}^{T} \tilde{q}_{k,t}e_{2,t}, \dots, \sum_{t=1}^{T} \tilde{q}_{k,t}e_{N,t}\right)(\mathbf{e}^\top\mathbf{e})^{-1}\mathbf{e}^\top\boldsymbol{\xi}_n = 0. \end{aligned} \tag{A.36}$$

Therefore, we know that

$$\hat{\theta}^{*2} = \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1}\boldsymbol{\psi}_t + \frac{1}{T}\sum_{t=1}^{T}\sum_{k=1}^{K} \hat{\lambda}_k^2 \tilde{q}_{k,t}^2. \tag{A.37}$$

A similar calculation gives

$$\hat{\theta}_q^2 = \frac{1}{T}\sum_{t=1}^{T}\sum_{k=1}^{K} \hat{\lambda}_k^2 \tilde{q}_{k,t}^2. \tag{A.38}$$

Therefore, we have

$$\hat{\theta}^{*2} - \hat{\theta}_q^2 = \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1}\boldsymbol{\psi}_t. \tag{A.39}$$

Using (A.31), we have almost surely in the limit of T tending to infinity,

$$\hat{\mathbf{a}}^\top \frac{\mathrm{var}(\hat{\mathbf{a}})^{-1}}{T}\hat{\mathbf{a}} = \frac{H}{T}\sum_{t=1}^{T} \boldsymbol{\psi}_t^\top \mathrm{var}(\mathbf{R}_0)^{-1}\boldsymbol{\psi}_t = H(\hat{\theta}^{*2} - \hat{\theta}_q^2). \tag{A.40}$$

A.4 Proof of the Simplifying Regression

We show that regression (20),

$$r_{n,t} = \sum_{m=1}^{N} a_{n,m}e_{m,t} + \sum_{k=1}^{K} \tilde{\lambda}_k \tilde{q}_{k,t}\left(\sum_{m=1}^{N} v_{n,m}\tilde{b}_{m,k}\right) + \xi_{n,t}, \tag{A.41}$$

reduces to regression (32) for the purpose of estimating $a_{n,m}$.

First, because the idiosyncratic flow $e_{m,t}$ is the residual of the first-stage regression (16), we know by construction that $\sum_{t=1}^{T} q_{k,t}e_{m,t} = 0$. Because each $\tilde{q}_{k,t}$ is a linear combination of $q_{1,t}, q_{2,t}, \dots, q_{K,t}$, we know that $\sum_{t=1}^{T} \tilde{q}_{k,t}e_{m,t} = 0$.

Second, we stack the panel regression (A.41) in vector form and define the $NT \times 1$ vector

$$\mathbf{r} = (r_{1,1}, r_{1,2}, \dots, r_{1,T}, r_{2,1}, r_{2,2}, \dots, r_{2,T}, \dots, r_{N,1}, r_{N,2}, \dots, r_{N,T})^\top. \tag{A.42}$$

We can rewrite the panel regression (A.41) in vector form as

$$\mathbf{r} = \sum_{n=1}^{N} \sum_{m=1}^{N} a_{n,m} \mathbf{e}_{n,m} + \sum_{k=1}^{K} \lambda_k \mathbf{y}_k + \boldsymbol{\xi}. \tag{A.43}$$

Each vector $\mathbf{e}_{n,m}$ is an $NT \times 1$ vector with only the $(n-1)T+1$-th to nT-th entry ranging from $e_{m,1}$ to $e_{m,T}$ and all other entries equaling zero. Each $\mathbf{y}_k$ is an $NT \times 1$ vector in the form

$$\mathbf{y}_k = \Bigg(\underbrace{\tilde{q}_{k,1} \sum_{m=1}^{N} v_{1,m} \tilde{b}_{m,k}, \ldots, \tilde{q}_{k,T} \sum_{m=1}^{N} v_{1,m} \tilde{b}_{m,k}}_{T \text{ terms}}, \ldots, \underbrace{\tilde{q}_{k,1} \sum_{m=1}^{N} v_{N,m} \tilde{b}_{m,k}, \ldots, \tilde{q}_{k,T} \sum_{m=1}^{N} v_{N,m} \tilde{b}_{m,k}}_{T \text{ terms}} \Bigg)^{\top}. \tag{A.44}$$

The vector $\boldsymbol{\xi}$ simply stacks all error terms $\xi_{n,t}$.

Note that, for any n, m, and k, we have

$$\mathbf{e}_{n,m}^{\top} \mathbf{y}_k = \sum_{l=1}^{N} v_{n,l} \tilde{b}_{l,k} \sum_{t=1}^{T} \tilde{q}_{k,t} e_{m,t} = 0, \tag{A.45}$$

where the last equality uses the first step in the proof. As a result, to estimate coefficient $a_{n,m}$, it suffices to run regression

$$\mathbf{r} = \sum_{n=1}^{N} \sum_{m=1}^{N} a_{n,m} \mathbf{e}_{n,m} + \boldsymbol{\zeta}. \tag{A.46}$$

This regression further reduces to asset-by-asset time-series regression (32).

B Construction and Cleaning of Mutual Fund Flows

In this Appendix, we present the details involved in constructing and cleaning mutual fund flows.

Our primary data source is the CRSP Survivorship-Bias-Free Mutual Fund database. We start with return and total net assets (TNA) data at the share-class level from all funds. A mutual fund may include multiple share classes. We first drop observations with no valid CRSP identifier, `crsp_fundno`. A small number of fund share classes report multiple TNA in a given month. These are likely data duplication, so we keep only the first observation of the month. We end up with 8,591,018 share-class×month observations. In what follows, we discuss the cleaning steps for returns and TNA separately at the share-class level. After cleaning, we aggregate the share-class level data to the fund level.

B.1 Return Cleaning

We first correct data errors in monthly net returns, `mret`.

First, we address extreme positive returns. We study the case in which a certain fund share has returns greater than 100% and has existed for more than one year.[34] We manually check the entire time series of each share class in this subsample. The majority of these extreme returns reflect misplaced decimal points, which confound returns in decimal and percentage formats. For these cases, we divide the erroneous returns by 100.

Second, we address extreme negative returns. Similarly, we obtain a subset of monthly fund observations, in which a certain fund share has existed for more than one year and has returns that are lower than −50%. With extreme negative returns, we need to distinguish data errors from significantly negative returns before a fund's closure. Thus, we manually check only the subsample of negative returns that occur at least one year prior to the last observation of a closed fund. For this subsample, we manually check whether these extreme returns reflect data-input errors. For the cases with misplaced decimal points, we divide the erroneous returns by 100.

B.2 TNA Cleaning

Unlike many prior studies that construct percentage mutual fund flows, we study dollar-amount flows in order to preserve the relative magnitudes in the cross section. The mutual fund size distribution features a very long right tail. Winsorizing the extreme dollar-amount TNA likely removes valid large values together with input errors, generating significant bias. We devise an algorithm to identify and correct erroneous observations of TNA:

1. Using the sample with corrected returns, we calculate dollar flows as in (33) at the share-class level.

2. We study the top and bottom 0.5% of all dollar flows.[35] We manually check the TNA time series of all share classes in this subsample. We identify several common error types:

[34] We require the one-year threshold because mutual fund returns and TNAs during the first year are sometimes inaccurate in the CRSP database. For example, returns and TNAs can be stale or reported using a placeholder number such as 0.1. We address these issues in later steps by cross-checking alternative database.

[35] The choice of the top and bottom 0.5% is motivated by the distribution of dollar flows, where most extreme values tend to occur at these tails.

- Misplaced decimal points (usually by hundredths or thousandths).
- Stale TNA observations from CRSP, typically when a fund reorganizes its share class offering (e.g., adding a new share class and moving assets into a single share class).
- CRSP sometimes sets TNA = 0.1 for the first few months of a new fund or a new share class.

We correct the misplaced decimal issue immediately. For funds suffering from the latter two problems, we obtain their TNA from Morningstar.[36] Morningstar's TNA data (`Net_Assets_ShareClass_Monthly`) suffer to a lesser extent from these issues than CRSP's TNA data. We arrive at this conclusion by further cross-checking other third-party vendors (e.g., Yahoo Finance and Bloomberg Terminal). Hence, when a fund's CRSP TNA deviates more than 50% from its Morningstar TNA, we use the Morningstar TNA.

3. We repeat the previous steps one more time to make sure that we have accounted for most, if not all, extreme errors.

4. We compare our cleaned TNA with total assets (`assets`) from Thomson/Refinitiv Holdings data.[37] Following Coval and Stafford (2007) and Lou (2012), we drop observations whenever our cleaned TNA deviate more than 50% from `assets` from Thomson/Refinitiv.

Using cleaned return and TNA data, we calculate dollar flows at the share-class level using (33). We obtain a fund's flows by adding up the flows of all share classes in the fund. The final sample contains 1,613,579 fund×month observations.

B.3 Cross-Validating the Data-Cleaning Procedure

We cross-validate our data-cleaning procedure. For each month, we compute the aggregate mutual fund flows in dollar amounts. We compare our aggregate flow measures with those from alternative sources, including the Investment Company Institute (ICI) and the Flow of Funds (FoF).

The ICI provides aggregate monthly mutual fund flows. We obtain a version of ICI aggregate flows data from 2007 to 2020. We use the ICI's Total Equity mutual fund flows, which feature a close coverage scope to mutual funds in our sample.

[36] We merge the CRSP and Morningstar databases using a fund's ticker.
[37] We merge the two databases via the linking table MFLINKS, which is provided by WRDS.

Figure B.1. Time series of aggregate mutual fund flows from various sources

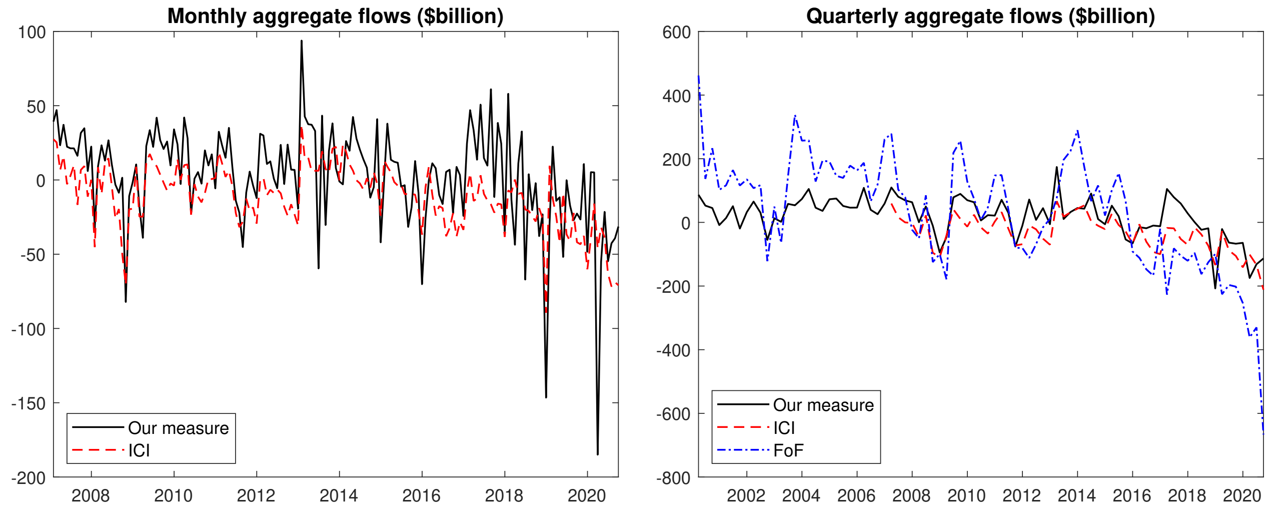

Notes: The left panel plots the monthly time series of our measure of aggregate mutual fund flows and Investment Company Institute (ICI) flows. The right panel plots the quarterly time series of our measure, ICI flows, and Flow of Funds (FoF) flows.

The FoF data (now known as "Financial Accounts of the United States - Z.1") are published quarterly by the Federal Reserve Board. We use mutual fund flow (Line 28) of Corporate Equities (Table 223) unadjusted flows (FU). We use the December 2021 vintage of the data, because the Federal Reserve revises historical FoF data every quarter.

Figure B.1 plots the time series of aggregate mutual fund flows from various sources. The left panel shows the monthly time series of our measure and ICI flow. The right panel shows the quarterly time series of all three sources. Our measure of aggregate mutual fund flows is broadly consistent with the other two sources. The correlation between our aggregate flow measure and ICI flow is 0.68 at the monthly level and 0.80 at the quarterly level. The correlation between our measure and FoF flow is 0.55 at the quarterly level.

The differences in Figure B.1 between the three measures of aggregate flows likely reflect differences in mutual fund coverage. Specifically, the ICI flow tracks virtually all U.S. equity mutual funds that invest in both domestic and world equity markets.[38] The FoF flow, sourced from unpublished ICI data, aggregates unadjusted flows into and out of all U.S. mutual funds (including variable annuity long-term mutual funds). The FoF flow is calculated based on mutual fund assets in common stock, preferred stock, and rights and warrants.[39] In comparison, our mutual fund sample contains U.S. mutual funds that are covered by CRSP.

[38] The ICI is a trade association for the mutual fund industry, and virtually all U.S. mutual funds are ICI members (Warther, 1995).

[39] Information source: https://www.federalreserve.gov/apps/fof/SeriesAnalyzer.aspx?s=FA653064100&t=F.223&suf=Q.

CRSP collects historical data from various sources.[40] Due to the nature of data collection process, CRSP's coverage is less complete than ICI's coverage.

C Additional Empirical Details

In this appendix, we provide additional empirical details.

C.1 An Illustration of Factor-Flow Construction

We illustrate factor-flow construction. This example shows why multiplying stock-level flow by the pseudoinverse of portfolio weights instead of the transpose is correct.

We consider three stocks, with flows equal to \$2, \$4, and \$2, respectively. There are two factor portfolios. The first portfolio equally weights stocks one and two, and the second portfolio equally weights stocks two and three. The portfolio weight matrix is therefore

$$\mathbf{W} = \begin{pmatrix} \frac{1}{2} & 0 \\ \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} \end{pmatrix}. \tag{C.1}$$

The pseudoinverse recovers the correct magnitude of flow into factor portfolios,

$$(\mathbf{W}^\top \mathbf{W})^{-1}\mathbf{W}^\top \begin{pmatrix} 2 \\ 4 \\ 2 \end{pmatrix} = \begin{pmatrix} 4 \\ 4 \end{pmatrix}, \tag{C.2}$$

because \$4 flows into portfolios one and two exactly lead to the stock-level flow \$2, \$4, and \$2. In contrast, the transpose leads to an incorrect measure of flow into factor portfolios,

$$\mathbf{W}^\top \begin{pmatrix} 2 \\ 4 \\ 2 \end{pmatrix} = \begin{pmatrix} 3 \\ 3 \end{pmatrix}. \tag{C.3}$$

C.2 An Illustration of Recovering Portfolio Weights

We use an example to illustrate when flow betas recover the original portfolio weights.

[40] The sources include the Fund Scope Monthly Investment Company Magazine, the Investment Dealers Digest Mutual Fund Guide, Investor's Mutual Fund Guide, the United and Babson Mutual Fund Selector, and the Wiesenberger Investment Companies Annual Volumes.

There are two assets. Their flows $f_{1,t}$ and $f_{2,t}$ are driven by a single factor q_t,

$$\begin{pmatrix} f_{1,t} \\ f_{2,t} \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \end{pmatrix} q_t. \tag{C.4}$$

In the data, we observe only asset flows $f_{1,t}$ and $f_{2,t}$, and do not observe the true q_t. Suppose that we choose an incorrect factor portfolio weight $\mathbf{W} = (1, 1)^\top$.

We use the pseudoinverse in equation (35) to construct the factor flow

$$q_t = (\mathbf{W}^\top \mathbf{W})^{-1} \mathbf{W}^\top \begin{pmatrix} f_{1,t} \\ f_{2,t} \end{pmatrix} = \frac{f_{1,t} + f_{2,t}}{2} = \frac{3}{2} q_t. \tag{C.5}$$

We then run the first-stage time-series regression (42) of asset flow $f_{i,t}$ on factor flow q_t for $i = 1, 2$, and obtain the flow betas

$$b_i = \frac{\sum_t q_t f_{i,t}}{\sum_t q_t^2} = \frac{2i}{3}. \tag{C.6}$$

The mimicking portfolio weights are $\mathbf{b} = (2/3, 4/3)^\top$. We see that $\mathbf{b}$ does not equal the original factor portfolio weight $\mathbf{W} = (1, 1)^\top$.

In contrast, if we choose any factor portfolio weight $\mathbf{W} = (a, 2a)^\top$ that is proportional to the correct principal-component weight $(1, 2)^\top$, then our procedure indeed recovers $\mathbf{b} = \mathbf{W}$. This example shows that our procedure recovers the original portfolio weights $\mathbf{W}$ only if $\mathbf{W}$ matches the principal components of cross-sectional flows.

C.3 Additional Empirical Results

In Table C.1, we report the R^2 of total price impact regression (39),

$$r_{n,t} = \sum_{m=1}^{N} a_{n,m} f_{m,t} / w_m + \epsilon_{n,t}. \tag{C.7}$$

The regression R^2 is about 20% for all 25 test assets. In unreported results, we also find that many of the points estimates of $a_{n,m}$ are either statistically insignificant or negative, suggesting that this regression likely suffers from over-fitting problems.

In Table C.2, we report the R^2, point estimates of multipliers, and t-statistics for the

Table C.1. R^2 of total price impact regression

	Low	2	3	4	High
Small	23.58%	24.52%	23.86%	24.86%	25.88%
2	21.91%	20.72%	24.09%	21.98%	21.03%
3	21.06%	22.10%	20.49%	22.68%	20.38%
4	21.05%	20.11%	22.02%	22.17%	22.90%
Big	19.89%	19.00%	18.32%	16.05%	20.28%

Notes: This table reports the R^2 of total price impact regression (47). The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). We run asset-by-asset time-series regressions of each asset's demeaned return on all 25 assets' flows.

joint price impact and cross impact regression,

$$r_{n,t} = \eta_n f_{n,t}/w_n + \phi_n \left(\sum_{m \text{ neighboring to } n} f_{m,t}/w_m \right) \Big/ (\text{number of } m \text{ neighboring to } n) + \epsilon_{n,t}. \tag{C.8}$$

The results show that this regression suffers severe multicollinearity concern. Specifically, the regression R^2 of (41) improves very little compared to the maximum of the R^2 from raw price impact regression (38) and the R^2 from raw cross impact regression (40) in Table 2. For each asset n, the point estimates of η_n and ϕ_n from regression (41) are usually one positive and one negative, and both lack statistical significance.

In Figure C.1, we report the time series of 25 test asset idiosyncratic flows. The idiosyncratic flow is the residual from the first-stage regression (42). The idiosyncratic flows exhibit significantly lower commonality than the raw flows in Figure 2.

Table C.2. Joint price impact and cross impact regression

	regression R^2									
	Low	2	3	4	High					
Small	9.25%	9.05%	11.45%	10.97%	12.23%					
2	9.82%	9.82%	12.07%	10.57%	9.36%					
3	9.56%	11.83%	10.84%	12.12%	10.20%					
4	12.78%	10.81%	10.79%	10.66%	9.34%					
Big	6.24%	9.40%	4.68%	6.75%	6.40%					
	multiplier η					multiplier ϕ				
	Low	2	3	4	High	Low	2	3	4	High
Small	-11.01	-5.36	-1.22	1.66	-4.53	26.16	18.34	13.53	9.44	17.45
2	21.42	1.49	-1.94	4.64	1.59	-8.99	10.61	14.37	6.63	11.93
3	-5.17	-7.19	7.79	-7.11	-1.59	21.23	21.73	4.65	21.67	15.64
4	29.83	9.56	-0.30	20.34	-0.71	-12.89	7.87	18.92	-2.78	18.30
Big	-5.14	-6.43	-3.76	-14.69	8.77	18.68	23.83	17.24	25.53	7.44
	$t(\eta)$					$t(\phi)$				
Small	-1.16	-0.55	-0.20	0.19	-0.83	2.73	1.95	1.88	1.10	2.83
2	2.09	0.17	-0.23	0.51	0.18	-0.82	0.96	1.69	0.65	1.19
3	-0.44	-0.82	0.73	-0.81	-0.22	1.56	2.22	0.38	2.24	1.99
4	2.92	0.77	-0.03	1.64	-0.11	-1.16	0.55	1.54	-0.22	2.07
Big	-0.49	-1.04	-0.47	-1.80	1.12	2.73	3.60	2.09	4.20	0.82

Notes: This table presents the R^2, point estimates of multipliers, and t-statistics for joint price impact and cross impact regression (C.8). The 5×5 assets are sorted based on size (small to big) and book-to-market equity (low to high). We run asset-by-asset time-series regressions of each asset's demeaned return on its own flow and its neighboring assets' average flow. The t-statistics are calculated from heteroskedasticity-robust standard errors.

Figure C.1. Time series of 25 test asset idiosyncratic flows

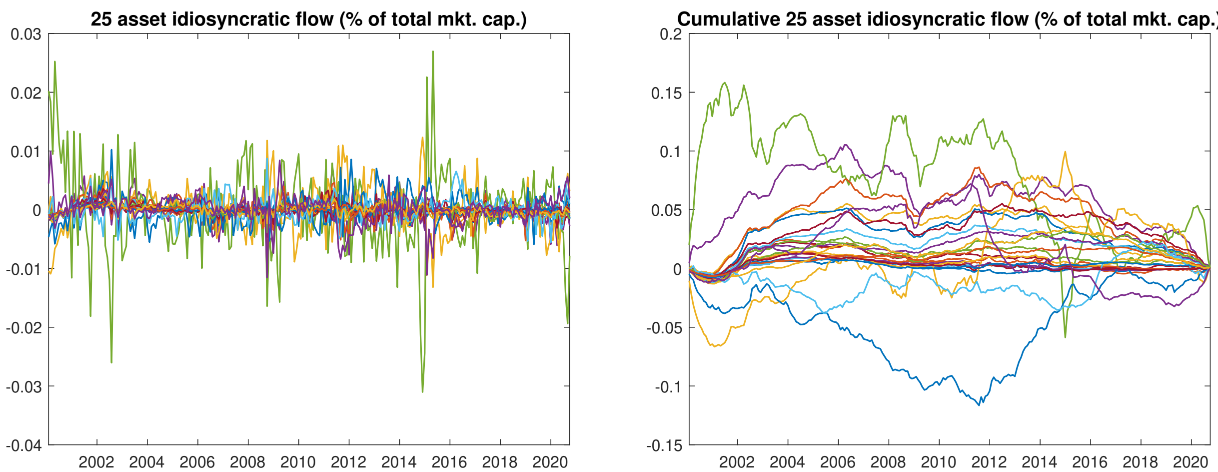

Notes: The idiosyncratic flow is the residual from the first-stage flow regression. In the left panel, we plot the time series of 25 asset idiosyncratic flows at monthly frequency. In the right panel, we plot the cumulative sum of 25 asset idiosyncratic flows. The sample period runs from January 2000 through September 2020.