The volatility puzzle of the beta anomaly[⋆]

Pedro Barroso[a], Andrew Detzel[b,*], Paulo Maio[c,d]

[a] *Universidade Católica Portuguesa, Católica-Lisbon School of Business and Economics, Palma de Cima, 1649-023 Lisboa, Portugal*
[b] *Baylor University, Waco, TX 76712, USA*
[c] *Hanken School of Economics, Department of Finance and Economics, Helsinki, Finland*
[d] *Brazilian School of Public and Business Administration, Getulio Vargas Foundation, Brazil*

Abstract

This paper shows that leading theories of the beta anomaly fail to explain the anomaly's conditional performance. Abnormal returns and Sharpe ratios of betting-against-beta (BAB) factors rise following months with below-median realized volatility, even controlling for mispricing, limits to arbitrage, lottery preferences, analyst disagreement, and sentiment. Moreover, the leverage constraints theory counterfactually predicts that market and BAB Sharpe ratios increase with volatility. We further show that institutional investors shift their demand from high- to low-beta stocks as volatility increases, and the resulting price impact is sufficient to explain the difference in abnormal BAB returns between high- and low-volatility states.

JEL classification: G11, G12, G17

Keywords: Betting-against-beta, time-varying risk, realized volatility, risk factors, scaled factors, anomalies, lottery preferences, leverage constraints

---

⋆Nikolai Roussanov was the editor for this article. The present work has benefited from comments and suggestions from Vikas Agrawal, Eser Arisoy, Marie Brière, Doina Chichernea, Sara Easterwood, Paul Karehnke, Ralph Koijen, Scott Murray, Konark Saxena, Li Yang, and numerous seminar participants in Católica-Lisbon, Georgia State University, IESE, ESADE, Paris-Dauphine, Queensland University of Technology, the 2017 Australasian Banking and Finance Conference, World Finance Conference, European Financial Management Association, 11th Portuguese Finance Network Conference, the Spanish Finance Association annual meeting, the Frontiers of Factor Investing Conference at Lancaster University, the 14th annual Risk Management Conference organized by National University of Singapore, and the 2022 China International Conference in Finance. We thank in particular the discussants Tomohiro Ando, George Chalamandaris, Yvonne Kreis, Tim Kroencke, Yang Liu, and Nuno Soares. We are grateful to Kenneth French, Amit Goyal, Lasse Pedersen, and Robert Stambaugh for providing stock market data. This study subsumes several previous working papers including "The volatility puzzle of the low-risk anomaly," and "Managing the risk of the beta anomaly." Pedro Barroso acknowledges the support from FCT – Portuguese Foundation of Science and Technology for the project "UID/GES/00407/2020."

*Corresponding author

Email addresses: pedro.barroso@ucp.pt (Pedro Barroso), andrew_detzel@baylor.edu (Andrew Detzel), paulo.maio@hanken.fi (Paulo Maio)

"We keep regurgitating the data to find yet one more variation of the size, value, or momentum anomaly, when the mother of all inefficiencies may be standing right in front of us—the risk anomaly."

—Robin Greenwood quoted in Ang (2014), page 332.

1. Introduction

Understanding the relationship between risk and return is perhaps the most central pursuit of asset pricing. One of the oldest and most well-known facts about this relationship is that stocks with low market betas earn positive abnormal returns relative to the CAPM of Sharpe (1964) and Lintner (1965), with the opposite result holding for high-beta stocks. This anomaly dates to at least Friend and Blume (1970) and Black, Jensen and Scholes (1972) but it remains robust in the fifty years of data after these seminal studies and its cause is still actively debated, lacking a single universally accepted explanation (see, e.g., Frazzini and Pedersen, 2014; Liu, Stambaugh and Yuan, 2018; Asness, Frazzini, Gormsen and Pedersen, 2020). In addition to academic interest, "defensive equity" strategies have experienced "massive capital inflows" in the words of Novy-Marx and Velikov (2022) and provide a foundation for a multitude of exchange-traded funds and other investment products.

In this paper, we take a novel approach to evaluating leading theories of the beta anomaly based on their ability to explain the conditional returns of "betting-against-beta" (BAB) factors that buy low-beta stocks and short high-beta stocks. Overwhelming evidence exists that risk and risk premia vary over time, and the true cause of an anomaly must match this time variation (see, e.g., Cochrane, 2011). The literature also posits several theories of the beta anomaly that, suspiciously, are all capable of explaining the unconditional returns of BAB factors even though they are based on very different economic mechanisms. As noted by Nagel and Singleton (2011), exploiting conditional variation in anomaly performance provides more statistical power to reject false explanations that "pass" unconditional asset pricing tests. We use lagged realized volatility of the BAB factor as conditioning information following Moreira and Muir (2017) and Cederburg, O'Doherty, Wang and Yan (2020), who show that volatility timing beta factors results in significant performance improvements because their Sharpe ratios rise following low realizations of volatility. We also show that the loadings of BAB returns on several asset pricing factors fall with volatility, so that the returns of BAB are most anomalous when risk is low. The main results of this paper show that none of the leading theories of the beta anomaly are up to the task of explaining why low-risk stocks perform so well in low-risk times.

We consider six leading explanations for the beta anomaly: leverage constraints, risk factors missing from the CAPM, limits to arbitrage, lottery preferences, sentiment, and analyst disagreement. The leverage constraints theory, proposed by Black (1972) and expanded by Frazzini and Pedersen (2014), posits that leverage-constrained investors with relatively low risk aversion bid up the prices of high-beta stocks because of their high expected returns, rendering their CAPM alphas negative. Consistent with this theory, Adrian, Etula and Muir (2014), Frazzini and Pedersen (2014), Boguth and Simutin (2018), Jylhä (2018), and Lu and Qin (2021) show that proxies for leverage constraints forecast returns on BAB factors. However, previous studies ignore the predictions this theory makes about the relationship between volatility and subsequent BAB performance. We compare several calibrations of Frazzini and Pedersen's equilibrium model that generate different levels of BAB volatility. They show that the Sharpe ratios of both BAB and the market portfolio should increase with these factors' respective volatilities, the opposite of the results found in the data. These counterfactual predictions arise from a key assumption in the model that all investors are risk averse and therefore demand a positive risk-return tradeoff (see, e.g., Merton, 1973).

Turning to multifactor explanations, Novy-Marx and Velikov (2022) show that the Fama and French (2018) six-factor model (FF6) prices BAB because low-beta stocks move like those with robust profitability, low asset growth, and high momentum. We confirm this result, but also show that the FF6 fails to explain the conditional performance of BAB. When lagged volatility is "low" (below its median value), the BAB Sharpe ratio more than doubles compared to its entire sample estimate. At the same time, the factor loadings that explain BAB returns unconditionally shrink significantly, or even flip signs, resulting in significantly positive alpha that is one percentage point higher per month than in high-volatility states. Thus, BAB returns appear most anomalous precisely when they require the least risk to earn. We further show that this pattern extends to the CAPM and the Fama and French (1993) three-factor model and to models with economically motivated factors based on the "unfiltered" consumption of Kroencke (2017) and the intermediary leverage measure of Adrian et al. (2014), which unconditionally prices BAB. In particular, no factor model we consider can explain the abnormal returns on BAB conditional on volatility.

We next investigate whether the remaining theories of the beta anomaly rectify the failure of the FF6 in low-volatility states. Liu et al. (2018) posit that the beta anomaly could arise from the strong cross-sectional correlation between idiosyncratic volatility (IVOL) and beta. This argument expands on that of Stambaugh, Yu and Yuan (2015), who argue IVOL is a proxy for arbitrage risk and that it interacts with limitations on short selling to magnify overpricing. Using the mispricing measure of Stambaugh, Yu and Yuan (2012), Liu et al. (2018) show that the cross-sectional relationship between beta and abnormal returns disappears when excluding "overpriced" high-IVOL stocks or when controlling for IVOL or mispricing.

Motivated by these results, we construct three new BAB factors: one that excludes overpriced high-IVOL stocks, one based on betas that are cross-sectionally orthogonalized to IVOL, and one based on betas that are cross-sectionally orthogonalized to mispricing. Consistent with Liu et al. (2018), all three factors earn an insignificant unconditional alpha; however, this alpha becomes positive when volatility is low and negative when volatility is high, with a significant difference between the two states of 0.8 to 1.0 percentage points per month. The Sharpe ratios of these factors also increase by 0.5 to 0.7 going from high- to low-volatility months. Thus, limits to arbitrage do not explain the conditional performance of the beta anomaly.

Bali, Cakici and Whitelaw (2011) and Bali, Brown, Murray and Tang (2017) argue that investor demand for stocks with lottery-like payoffs could bid up the prices of high-risk stocks at the expense of subsequently earning negative abnormal returns. They measure a stock's lottery demand in a given month using the average of the highest five daily returns in that month, and the latter study shows it subsumes the cross-sectional relationship between beta and returns. Hypothetically, lottery demand could explain the time-varying performance of BAB factors as long as this demand is high when BAB volatility is low. However, similar to our limits-to-arbitrage tests, we form BAB factors based on betas orthogonalized to lottery demand, and show they earn significantly higher Sharpe ratios and alphas in low-volatility states than high-volatility ones. Thus, lottery demand cannot explain the conditional performance of the beta anomaly.

Next we consider two explanations of the beta anomaly that predict time-series variation in BAB returns. Antoniou, Doukas and Subrahmanyam (2016) posits that high-beta stocks are especially exposed to sentiment and find that the anomaly is significant only in high-sentiment months. Similarly, Hong and Sraer (2016) argue that beta amplifies disagreement about the economy, which, coupled with short-sale constraints, renders high-beta stocks overpriced when disagreement is high. Consistent with this argument, they show that the beta anomaly is significant only following relatively high levels of analyst earnings forecast dispersion. However, we find that the negative relationship between BAB alphas and lagged volatility holds regardless of the level of analyst disagreement or sentiment.

For the abnormal returns of BAB to be negatively related to lagged volatility, it must be the case that high-beta stocks become less "overvalued" relative to low-beta stocks as the volatility of BAB increases. Next, we investigate whether institutional investors' trading can explain this dynamic relationship by estimating a demand system following Koijen and Yogo (2019). On average, we find that institutional investors have a significant preference for high-beta stocks. However, going from low- to high-volatility months, this preference falls to essentially zero and becomes statistically insignificant, while so-called "household" investors consistently prefer low-beta stocks regardless of volatility. To determine whether this time-varying demand can explain the conditional performance of the beta anomaly, we perform a counterfactual experiment in which we drop institutional investors representing 10% of aggregate assets under management who have the greatest tendency to reduce the betas of their equity portfolios when volatility increases, and then reallocate their assets to remaining investors. Doing so causes low-beta stocks to appreciate relative to high-beta stocks by 4.2 percentage points on average, by 8.4 percentage points when BAB volatility is low, and 1.5 percentage points when volatility is high. The 6.9 percentage point difference between low- and high-volatility states is sufficient to fully eliminate the conditional alpha of BAB based on an estimate derived from the Campbell and Shiller (1988) decomposition following Han, Roussanov and Ruan (2022).

The results from the counterfactual experiment are consistent with a nascent literature arguing that widely used performance evaluation contracts incentivize institutional managers to overweight high-beta stocks, even if their alphas are negative, because, on average, their high returns exceed those of unit-beta benchmarks like the S&P 500 (e.g., Baker, Bradley and Wurgler, 2011 and Christoffersen and Simutin, 2017). The same incentives also penalize tracking error volatility and so, as BAB volatility increases, managers have the incentive to retreat from their positions in high-beta stocks toward their benchmarks, leading them to sell high-beta stocks and buy low-beta ones.[1] Some institutions do not face explicit-benchmark-based incentives per se, but still have pressure to beat an equity index while reducing risk when volatility is high. For example, Ellul, Jotikasthira, Kartasheva, Lundblad and Wagner (2022) shows that life insurers "reach for yield" (on average) to beat equity indexes because their largest liabilities are variable annuities that make payouts based on the performance of the stock market. However, these insurers also face internal risk limits and risk-based capital requirements, so it stands to reason that, similar to mutual- and pension-fund managers, they face pressure to reduce betas of their equity portfolios when volatility increases substantially. Regardless of the precise mechanism, the counterfactual experiment clearly shows that the demand dynamics of institutional managers whose portfolio betas are the most sensitive to volatility impact stock prices in a manner that can fully account for the time-variation in the beta anomaly relative to volatility.

Overall, the results in this paper challenge current explanations from the extensive literature on the beta anomaly. Our study is most closely related to Asness et al. (2020), who examine the compatibility of this anomaly with essentially the same leading explanations as we do. Their key innovation is creating betting-against-risk factors that separate the cross-sectional effects of the systematic risk in beta from those of idiosyncratic volatility. In contrast, our key innovation focuses on the time series, which leads us to

[1] Several investments professionals we discussed these results with had observed similar behavior, namely managers retreating towards benchmarks when risk of positions increases. One common stated reason is that managers are especially concerned about outflows or termination from poor performance when risk increases and investors are paying relatively close attention.

very different conclusions. They find that multiple theories contribute to explaining the unconditional beta anomaly while we find that none of these theories are consistent with the conditional performance of BAB. More broadly, our results show that five decades worth of attempts to explain a single anomaly fail to produce a single theory that can explain this anomaly's returns conditional on volatility. This fact highlights the importance of not ignoring conditioning information in asset pricing, which is still widely done in practice.

This paper proceeds as follows. Section 2 specifies our data sources and the construction of our BAB portfolios. Section 3 presents motivational evidence on the performance of volatility-managed BAB factors. Section 4 evaluates leading theories of the beta anomaly. Section 5 reports institutional demand results. Section 6 concludes.

2. Data and variable construction

This section describes our data sources and the construction of our betting-against-beta (BAB) factor.

2.1. Data sources

We obtain returns on U.S. common stocks from CRSP, accounting data from COMPUSTAT, institutional common stock holdings from the Thomson Reuters Institutional Holdings Database, and analyst forecasts from I/B/E/S. We obtain return data for the risk-free rate and the six-factor asset pricing model of Fama and French (2018), FF6, from the website of Kenneth French. FF6 includes the market (MKT) and value (HML) factors of Fama and French (1993), as well as factors based on size (SMB), robust (high)-minus-weak (low) profitability (RMW), conservative (low)-minus-aggressive (high) real investment (CMA), and momentum (MOM) following Carhart (1997). We use the stock-level mispricing measure of Stambaugh et al. (2012), Stambaugh et al. (2015), and Stambaugh and Yuan (2017) obtained from the website of Robert Stambaugh. They construct their measure each month as the average of the stock's rankings with respect to 11 variables associated with prominent anomalies relative to the Fama and French (1993) three-factor model (FF3) such that a higher level of mispricing negatively predicts subsequent abnormal returns. Lastly, we obtain the monthly sentiment index of Baker and Wurgler (2007) orthogonalized to economic conditions from the website of Jeffrey Wurgler, the "unfiltered" consumption growth time series of Kroencke (2017) from Tim Kroencke's website, and yields on BAA- and AAA-rated bonds from the St. Louis Fed's website. Due to data availability, our sample period is from July 1965 to December 2016.[2]

[2] Links to the websites cited in this section are: Kenneth French: `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`; Robert Stambaugh: `http://finance.wharton.upenn.edu/~stambaug/`; Tim Kroencke: `https://sites.google.com/site/kroencketim/`; and the St. Louis Fed: `https://fred.stlouisfed.org/series/BAA`. The 11 anomalies

2.2. Betting-against-beta factor

We estimate betas following Liu, Stambaugh and Yuan (2018). Each stock-month, (n,t), we estimate CAPM regressions using the prior 60 monthly returns (36 month minimum), applying a Dimson (1979) correction for nonsynchronous trading:

$$r_{n,\tau} - r_{f,\tau} = a_{n,t} + \beta_{n,t,0} MKT_\tau + \beta_{n,t,1} MKT_{\tau-1} + \epsilon_\tau, \qquad \tau = t-59, ..., t, \tag{1}$$

where MKT_τ is the return on the market in excess of the risk-free rate, $r_{f,\tau}$. We then shrink the resulting estimates, $\hat{\beta}^{ts}_{n,t} := \hat{\beta}_{n,t,0} + \hat{\beta}_{n,t,1}$, toward one using the procedure of Vasicek (1973) to mitigate estimation error:

$$\hat{\beta}_{n,t} := w_n \hat{\beta}^{ts}_{n,t} + (1 - w_n) \times 1, \tag{2}$$

where:

$$w_n = \frac{1/\hat{\sigma}^2(\hat{\beta}^{ts}_{n,t})}{1/\hat{\sigma}^2(\hat{\beta}^{ts}_{n,t}) + 1/\hat{\sigma}^2(\beta)}. \tag{3}$$

The $\hat{\sigma}(\hat{\beta}^{ts}_{n,t})$ is the standard error of $\hat{\beta}^{ts}_{n,t}$ and $\hat{\sigma}^2(\beta) := \hat{\sigma}^2_{cs}(\hat{\beta}^{ts}_{n,t}) - \overline{\hat{\sigma}^2(\hat{\beta}^{ts})}$ is a measure of the cross-sectional variance of true betas, where $\hat{\sigma}^2_{cs}(\hat{\beta}^{ts}_{n,t})$ is the cross-sectional variance of $\hat{\beta}^{ts}_{n,t}$ and $\overline{\hat{\sigma}^2(\hat{\beta}^{ts})}$ is the cross-sectional mean of $\hat{\sigma}^2(\hat{\beta}^{ts}_{n,t})$.[3]

At the beginning of each month, we sort all common stocks in CRSP into value-weighted decile portfolios based on their $\hat{\beta}_{n,t}$ (estimated through month $t-1$), denoting the return on the lowest-beta (highest-beta) portfolio as $r_{L,t}$ ($r_{H,t}$). Following Novy-Marx and Velikov (2022), we define our BAB factor to target beta neutrality as follows:

$$BAB_t = r_{L,t} - r_{H,t} - (\beta_{L,t-1} - \beta_{H,t-1}) MKT_t, \tag{4}$$

where $\beta_{L,t-1}$ ($\beta_{H,t-1}$) denotes the value-weighted average of shrunk $\hat{\beta}_{n,t-1}$ across stocks in the low-beta (high-beta) portfolio. Without hedging the market exposure with $(\beta_{L,t-1} - \beta_{H,t-1})MKT_t$, BAB would

used in the mispricing measure are: failure probability (e.g., Campbell, Hilscher and Szilagyi, 2008), O-Score (Ohlson, 1980), net stock issuance (Loughran and Ritter, 1995), composite equity issuance (Daniel and Titman, 2006), accruals (Sloan, 1996), net operating assets (Hirshleifer, Hou, Teoh and Zhang, 2004), momentum (Levy, 1967; Jegadeesh and Titman, 1993; and Fama and French, 1996), gross profitability (Novy-Marx, 2013), asset growth (Titman, Wei and Xie, 2004b; Lyandres, Sun and Zhang, 2007; Xing, 2007; and Cooper, Gulen and Schill, 2008), return on assets (Fama and French, 2006), and abnormal capital investment (Titman, Wei and Xie, 2004a).

[3] In extremely rare cases, $\hat{\sigma}^2(\beta)$ can be negative, which causes w_n to exceed one and so we set $w_n = 1$ (no shrinkage) in these cases.

have insignificant mean returns but significantly positive alpha. The beta neutrality, transforms this alpha into positive mean returns, which is necessary for Sharpe ratio comparisons that are central to our study.[4] Frazzini and Pedersen (2014) propose a widely cited alternative construction of a BAB factor that targets beta neutrality by scaling each leg of the factor by the inverse of their betas. However, this scaling results in factors that are "net-long," placing more weight on r_L than r_H, which Liu et al. (2018) and Han (2022) show results in abnormal returns that are unrelated to the beta anomaly. Novy-Marx and Velikov (2022) further argue that the Frazzini and Pedersen (2014) net-long scaling puts extreme weight on small-cap stocks because it uses the long and short legs of BAB to hedge each other's beta, whereas the "direct hedging" that we use hedges with the value-weighted market factor. Using value-weighted portfolios also mitigates another concern raised by Novy-Marx and Velikov (2022), specifically that the rank-weighting scheme of Frazzini and Pedersen (2014) puts extreme weights in micro-cap stocks.

3. The performance of volatility-managed BAB

Several recent studies find that volatility-managed versions of numerous asset pricing factors, which increase leverage following low realizations of volatility, produce significant Sharpe ratio gains relative to their unmanaged counterparts.[5] To set the stage for our main results, this section replicates these findings for our specification of BAB and discusses their implications for this factor's time-varying performance.

Following Barroso and Santa-Clara (2015), we define the volatility-managed betting-against-beta strategy as:

$$BAB^{\sigma}_{t+1} = \left(\frac{c}{\hat{\sigma}_t}\right) BAB_{t+1}, \tag{5}$$

where BAB_{t+1} denotes the month-$t+1$ buy-and-hold ("unmanaged") return and $\hat{\sigma}_t$ is the the realized volatility of the last 21 daily returns of the prior month, $BAB_{d,t},\ d = 1, ..., 21$:

$$\hat{\sigma}_t = \sqrt{\sum_{d=1}^{21} BAB^2_{d,t}}. \tag{6}$$

[4]Table 3 below shows that the hedge reduces the ex post market beta of BAB to a trivial -0.02 ($t = -0.31$) and untabulated results show that the strategy $(\beta_{L,t-1} - \beta_{H,t-1})MKT_t$ has a CAPM alpha that is essentially zero (-0.05%, $t = -0.87$) and its abnormal return never drives our results.

[5]Barroso and Santa-Clara (2015), Daniel and Moskowitz (2016), Eisdorfer and Misirli (2020), and Barroso and Detzel (2021) demonstrate these results for the market portfolio, momentum, the beta anomaly, and the distress anomaly. Moreira and Muir (2017) and Cederburg et al. (2020) expand them to numerous factors. See also Kirby and Ostdiek (2012), who consider volatility timing strategies for portfolio allocations across stock portfolios. Much of this literature follows from Fleming, Kirby and Ostdiek (2001, 2003) and Marquering and Verbeek (2004), who demonstrate large utility gains from volatility timing allocations across several asset classes.

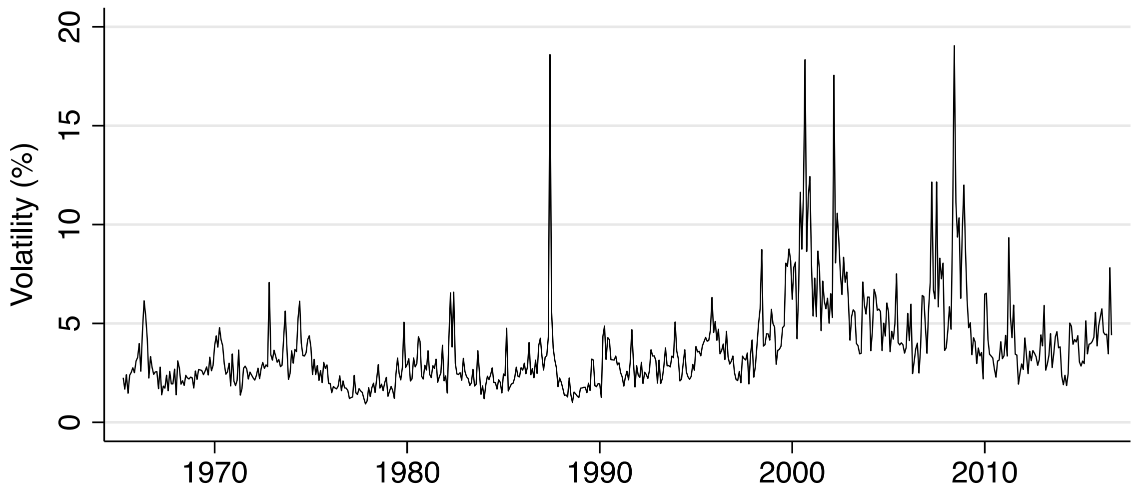

Fig. 1. Monthly time series of realized volatility of the betting-against-beta factor
This figure depicts the monthly time-series of realized volatility, defined in Eq. (6), of the betting-against-beta factor. The sample period is July 1968 to December 2016.

This timing approach is simple to construct, avoids hindsight biases documented by Liu, Tang and Zhou (2019), and does not produce as extreme weights as the common alternative practice of scaling by realized variance, (e.g., Fleming et al., 2001, 2003; Kirby and Ostdiek, 2012; and Moreira and Muir, 2017). The scalar c is arbitrary and has no impact on statistical inference, but for expositional purposes, we choose it to equate the standard deviation of the returns on the managed and unmanaged factors. Figure 1 presents a monthly time-series of the realized volatility of BAB defined by Eq. (6).

Table 1 reports descriptive statistics for the performance of BAB and BAB$^{\sigma}$. By construction, BAB$^{\sigma}$ has the same standard deviation, but about 50% higher mean return than the unmanaged BAB (0.67% vs. 0.45%).[6] As a result, BAB$^{\sigma}$ has a signifcantly higher Sharpe ratio than BAB (0.50 vs. 0.34). The economic and statistical significance of this Sharpe ratio improvement is noteworthy since Cederburg et al. (2020) find that only eight out of 103 managed factors have significantly higher Sharpe ratios than their unmanaged counterparts, and investors would find it difficult to realize benefits implied by other measures of performance, such as alphas, in real time. Volatility management also reduces higher order risk of BAB$^{\sigma}$, flipping the skew from negative to positive and decreasing the excess kurtosis.

[6] Untabulated results show similar increases in mean returns from volatility management for BAB factors made from betas estimated following Frazzini and Pedersen (2014), Novy-Marx and Velikov (2022), and Welch (2022), all ranging from 0.15 to 0.25 percentage points per month.

Table 1
Performance of BAB and its volatility-managed counterpart

At the beginning of each month, we sort all stocks into value-weighted decile portfolios based on beta. From these sorts, we form a betting-against-beta factor, BAB, that goes long the lowest-beta portfolio, short the highest-beta portfolio, and hedges market risk based on the ex ante betas of each portfolio (see, Eq. (4) for details). This table presents descriptive statistics of the returns on BAB and BAB$^{\sigma}$, which leverages BAB each month proportionally to its inverse realized volatility estimated over the previous month. The statistics shown are: mean excess return (Mean), standard deviation (SD), skewness (Skew), excess kurtosis (Kurt), and Sharpe ratio (SR). The sample period is July 1965 to December 2016. The parenthesis contains a z-statistic from a Jobson and Korkie (1981) test of the difference in Sharpe ratio between BAB and BAB$^{\sigma}$.

	Mean (%)	SD (%)	Skew	Kurt	SR
BAB	0.45	4.64	−0.91	8.06	0.34
BAB$^{\sigma}$	0.67	4.64	0.01	3.68	0.50
					(2.23)

The finding of Table 1 that BAB$^{\sigma}$ earns a higher Sharpe ratio than BAB indicates that the Sharpe ratio of BAB increases when volatility falls in the time-series; and this fact raises the bar for explanations of the beta anomaly. To see why, suppose that arbitrage opportunities do not exist so that assets are priced by a stochastic discount factor, M_t. Standard arguments (see, e.g., Cochrane, 2005, Ch. 1) show that the maximum Sharpe ratio in the market is $\sigma_t(M_{t+1})/E_t(M_{t+1})$, and the conditional Sharpe ratio of BAB is:

$$\mathrm{SR}_t(BAB) := \frac{E_t(BAB_{t+1})}{\sigma_t(BAB_{t+1})} = -\rho_t(M_{t+1}, BAB_{t+1})\frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})}, \tag{7}$$

where ρ_t denotes the conditional correlation coefficient. For Eq. (7) to hold, the negative relation between the volatility and Sharpe ratio of BAB necessarily requires that when its volatility *falls*, BAB loads more heavily on priced risk factors or the maximum Sharpe ratio in the economy rises. Especially given the commonality of volatility across factors, these facts are counterintuitive since investors should demand higher compensation for risk when volatility rises (see, e.g., Moreira and Muir, 2017). Alternatively, BAB may earn higher abnormal returns when volatility falls due to mispricing, but this possibility is also counterintuitive since one would expect arbitrage risk to be positively related to volatility.

4. Evaluating explanations of the beta anomaly

In this section, we investigate whether leading theories of the beta anomaly are consistent with the strong performance of BAB conditional on low volatility.

4.1. The leverage constraints theory

Black (1972) and Frazzini and Pedersen (2014) posit that high-beta stocks have embedded leverage that has economic value for investors facing borrowing constraints. Since they cannot borrow, less risk averse investors will instead invest heavily in high-beta stocks in equilibrium, bidding up their prices and lowering their expected returns relative to the CAPM. Sophisticated financial intermediaries able to leverage and assume the opposite side of the trade (e.g. hedge funds "betting against beta") are exposed to funding risk and margin requirements, which limits their capacity to eliminate the anomaly completely. Consistent with this theory, Adrian et al. (2014), Frazzini and Pedersen (2014), Boguth and Simutin (2018), Jylhä (2018), and Lu and Qin (2021) all show that various proxies for the tightness of leverage constraints forecast the returns on BAB factors. However, prior studies generally ignore the implications that the leverage constraints theory has for the predictive relationship between BAB's volatility and its subsequent returns.[7] In this section, we identify this predictive relation in their model by comparing several calibrations with different levels of volatility.

In the Frazzini and Pedersen (2014) model, each agent i has initial wealth $W_{i,t}$ and forms a portfolio of N securities. The number of shares of each security she wants to own is given by the vector $\boldsymbol{x}_i = (x_{i1}, ...x_{iN})$. The agent has a quadratic utility function and maximizes the respective objective function:

$$\max_{\boldsymbol{x}_i} \boldsymbol{x}_i' \left[E_t(\boldsymbol{P}_{t+1} + \boldsymbol{\delta}_{t+1}) - (1 + r_f)\boldsymbol{P}_t\right] - \frac{\gamma_i}{2}\boldsymbol{x}_i'\boldsymbol{\Omega}_t\boldsymbol{x}_i, \tag{8}$$

where $\boldsymbol{P}_{t+1}$ and $\boldsymbol{\delta}_{t+1}$ are the price and dividend at time $t+1$ (both are N-by-1 vectors), γ_i is a risk aversion parameter, and $\boldsymbol{\Omega}$ is the covariance matrix of $\boldsymbol{P}_{t+1} + \boldsymbol{\delta}_{t+1}$. Agents face the restriction that invested wealth cannot exceed the amount required to cover margin requirements:

$$m_{i,t} \sum_{n=1}^{N} x_{i,n} P_{n,t} \leq \ W_{i,t}, \tag{9}$$

where $m_{i,t}$ is the margin constraint. An investor that cannot use leverage, for example, has $m_{i,t} = 1$, meaning that the value of her risky portfolio of assets cannot exceed her wealth. The higher the $m_{i,t}$, the tighter becomes the margin requirement. This feature plays a crucial role in the model as investors have different levels of risk aversion and some, the less risk-averse ones, can be constrained by margin requirements, creating a shadow cost of funding constraints.

[7] This is arguably a similar omission as the common practice of motivating cross-sectional asset pricing studies with the ICAPM of Merton (1973) while ignoring the time-series predictions of the model like the positive risk-return tradeoff (see, e.g., Barroso and Maio, 2024). Moreira and Muir (2017) also note that this positive tradeoff is a feature of all leading macro-finance asset pricing models.

In equilibrium, for any security, the expected return should then be:

$$E\left(r_{n,t+1}\right) = r_f + \psi_t + \gamma \mathrm{cov}_t(r_{n,t+1}, r_{M,t+1})\boldsymbol{P}_t'\boldsymbol{x}^*, \tag{10}$$

where $r_{n,t+1}$, $r_{M,t+1}$ are, respectively, the security and market returns, ψ_t is an endogenous variable reflecting the aggregate shadow cost of leverage (a function of the average Lagrange multiplier of the constraint in Eq. (9)), and $\boldsymbol{x}^*$ is the N-by-1 vector with the total number of shares outstanding in the economy that must all be owned in equilibrium. Applying Eq. (10) to the market itself yields:

$$E\left(r_{M,t+1}\right) = r_f + \psi_t + \gamma \sigma_t^2(r_{M,t+1})\boldsymbol{P}_t'\boldsymbol{x}^*. \tag{11}$$

Solving Eq. (11) for $\boldsymbol{P}_t'\boldsymbol{x}^*$, plugging into Eq. (10), and rearranging terms yields the main equilibrium relationship of the model:

$$E\left(r_{n,t+1}\right) - r_f = \psi_t(1 - \beta_{n,t}) + \beta_{n,t}\left[E\left(r_{M,t+1}\right) - r_f\right]. \tag{12}$$

In particular, $\psi_t(1-\beta_{n,t})$ is the alpha of stock n. As long as some agents are constrained ($\psi_t > 0$), this alpha is negative for high-beta stocks ($\beta_{n,t} > 1$) and positive for low-beta stocks.

Eq. (11) highlights a previously neglected feature of the leverage constraints model: a positive market risk-return tradeoff. Because investors are risk averse, the representative risk aversion γ will be positive and the expected return and Sharpe ratio of the market should rise with volatility, counterfactual to the empirical results that the predictive relationship between market volatility and expected return is weak, or even negative, in the data, and the market Sharpe ratio falls with volatility (see, e.g., Moreira and Muir, 2017). We next show the model is also inconsistent with the conditional performance of BAB.

Frazzini and Pedersen (2014) provide the calibration exercise of their model in the Internet Appendix (page B18) for values "chosen to roughly match the empirical volatilities and correlations of the asset returns." They consider an economy with two agents and two assets, which have the same expected payoff of $\frac{1}{2} \times 100/(1+r_f)$. The payoffs have variances of 40 and 205 and a covariance of 84. The risk-free rate is 3.6% and total wealth split by the two agents is 100. The share of wealth of each agent, their risk aversion parameters, and margin requirements are the only exogenous variables needed to fully characterize equilibrium.

The first three columns of Table 2 replicate baseline calibrations of Frazzini and Pedersen (2014) and the remaining columns show results for different levels of volatility. In the standard CAPM scenario, investors are not constrained and so no BAB effect emerges in equilibrium. But a sizable BAB effect appears when the

Table 2
Calibrations of the Frazzini and Pedersen (2014) leverage constraints model

In the model, two agents with different relative risk aversion, shares of wealth, and leverage constraints (m1 and m2) trade a pair of risky assets that together comprise the market. The first column shows a baseline calibration in which leverage constraints do not bind and the CAPM holds. The following two columns replicate scenarios with leverage constraints from Frazzini and Pedersen (2014). Columns (4) to (7) present calibrations that each multiply the baseline covariance matrix of asset payoffs by a "volatility scalar" ranging from 0.25 to 2 (see Section 4.1 for details). The last column shows the combined effects of high volatility with high margin requirements. The first six rows specify the exogenous variables. Remaining rows specify the endogenous outcome variables, which include the annual volatility, expected excess return, Sharpe ratio, and beta of, the low-risk (L) asset, the high-risk asset (H), the market portfolio (MKT), and the betting-against-beta (BAB) factor. The last row contains the shadow cost of leverage constraints for the less risk averse investor (ψ).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Exogenous variables								
Volatility scalar	1	1	1	0.25	0.5	1.5	2	2
Risk aversion 1	1	1	1	1	1	1	1	1
Risk aversion 2	1	10	10	10	10	10	10	10
Wealth share 1	0.5	0.5	0.8	0.8	0.8	0.8	0.8	0.8
m1	1	1	1.2	1.2	1.2	1.2	1.2	1.5
m2	1	1	0	0	0	0	0	0
Endogenous variables								
Vol L	13%	14%	14%	7%	9%	18%	21%	22%
Vol H	30%	33%	33%	15%	22%	42%	51%	55%
Vol MKT	21%	23%	23%	11%	15%	29%	35%	37%
Vol BAB	7%	7%	7%	3%	5%	9%	11%	12%
Excess return L	3%	9%	10%	3%	6%	14%	17%	26%
Excess return H	6%	16%	15%	4%	8%	22%	28%	37%
Excess return MKT	4%	12%	13%	3%	7%	18%	22%	31%
Excess return BAB	0%	4%	6%	2%	3%	8%	9%	16%
SR MKT	0.20	0.55	0.55	0.33	0.44	0.61	0.64	0.84
SR BAB	0.00	0.47	0.79	0.51	0.66	0.81	0.78	1.32
β_L	0.6	0.6	0.6	0.6	0.6	0.6	0.6	0.6
β_H	1.4	1.4	1.4	1.4	1.4	1.4	1.5	1.5
ψ	0%	4%	7%	2%	4%	9%	10%	19%

less risk-averse investor is constrained, which occurs in columns (2) and (3). Frazzini and Pedersen (2014) argue that the results in column (3) are the most suitable to explain the observed BAB effect in the data and adopt it as their baseline calibration, so we focus on that scenario.

In columns (4) through (7), we multiply the covariance matrix of payoffs in the original calibration by a volatility scalar ranging from 0.25 to 2. The expected return of the BAB portfolio monotonically increases from 2% in the low volatility scenario to 9% in the high volatility scenario. The increase in expected excess returns is so large that, in spite of higher volatility, the Sharpe ratio of the strategy increases in equilibrium. Both results are the opposite of the empirical results documented in Table 1, where expected returns are either constant or decreasing with volatility and Sharpe ratios fall with volatility.

A second concern is that margin requirements typically increase with volatility. Brokers and central clearing counterparties routinely raise margin requirements on their clients when volatility is high to guard against the elevated likelihood of loss. In column (8), we combine the high volatility scenario in column (7) with higher margin requirements. In this scenario, the BAB Sharpe ratio is even higher than in column (7) since the less risk averse investor becomes even more constrained.

Overall, we conclude that the leverage constraints model cannot explain the variation in BAB returns and Sharpe ratios driven by volatility. Intuitively, it is hard to reconcile the economic rationale of the model with the absence of risk-return tradeoffs for both BAB and the market necessary to generate these factors' volatility-timing benefits. In the model, the constrained investors, who are more risk tolerant and therefore cause the beta anomaly, are not unsophisticated, optimistic, or driven by lottery preferences. Rather, they are rational utility-maximizing investors that choose, in equilibrium, to accept low CAPM alphas in high-beta stocks to benefit from their embedded leverage. As risk averse investors, however, they still require larger returns to hold risky investments when volatility increases.

4.2. Multifactor explanations of the beta anomaly

Novy-Marx and Velikov (2022) show that the FF6 prices BAB due to loadings on profitability, investment, and momentum factors. In this subsection, we examine whether this result continues to hold conditional on the level of BAB volatility and more thoroughly document the time-variation in BAB abnormal returns relative to volatility and other predictors.

As a baseline for the FF6 results, the first two columns of Panel A of Table 3 present whole-sample spanning regressions of BAB on the factors from the CAPM and Fama and French (1993) three-factor model, respectively. These columns essentially replicate the beta anomaly from prior studies, specifically that the CAPM and FF3 fail to price BAB even *unconditionally*, with alphas of 0.44% to 0.46% per month (with significance at the 5% or 1% level). The loading on the market factor in the CAPM regression is statistically indistinguishable from zero (-0.02; $t = -0.31$), indicating that the market hedge in BAB's construction is effective. In the FF3 regression, the slope on the market factor is nonzero due to the inclusion of correlated right-hand-side factors: SMB and HML. The second column, "FF3," shows that BAB loads heavily and negatively on SMB, and positively on HML, consistent with low (high) beta stocks moving like large caps (small-cap) value (growth) stocks.

While the CAPM and FF3 cannot clear the bar of pricing BAB unconditionally, the third column of Panel A of Table 3 shows a different story for FF6: Consistent with Novy-Marx and Velikov (2022), the FF6 alpha of BAB is economically and statistically insignificant (-0.08% per month; $t = -0.45$,) so the FF6 evidently

Table 3
Performance of BAB conditional on volatility

Panel A reports regressions of the returns on BAB (defined in Table 1) on those of the factors from the CAPM, Fama and French (1993), and Fama and French (2018) six-factor model (FF6). Rows labeled "Alpha" present the intercepts from these regressions. Panel B reports: a regression using the FF6 model over the subsample period of months with below-median lagged volatility ("Low"), the subsample period with above-median volatility ("High"), and the corresponding low-minus-high differences ("L−H"). Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. Panel C reports the Sharpe ratios (SR) of BAB over the subsamples specified by the column headings in Panel B. The bottom row of Panel C reports a p-value for the low-minus-high difference in Sharpe ratios based on a GMM estimation with heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

	Panel A: Whole-sample regressions			Panel B: FF6 by volatility state		
	CAPM	FF3	FF6	Low	High	L−H
Alpha (%)	0.46	0.44	−0.08	0.56	−0.41	0.97
	(2.49)	(2.71)	(−0.45)	(2.92)	(−1.61)	(3.03)
MKT	−0.02	0.16	0.31	0.18	0.44	−0.25
	(−0.31)	(3.12)	(6.62)	(3.53)	(6.35)	(−2.93)
SMB		−0.66	−0.59	−0.59	−0.60	0.01
		(−9.01)	(−8.52)	(−7.57)	(−5.77)	(0.11)
HML		0.33	0.15	0.17	0.06	0.11
		(3.57)	(1.53)	(1.45)	(0.43)	(0.60)
RMW			0.32	−0.16	0.49	−0.65
			(3.33)	(−1.05)	(3.94)	(−3.34)
CMA			0.62	0.23	0.78	−0.55
			(4.82)	(1.37)	(4.48)	(−2.25)
MOM			0.30	0.12	0.37	−0.25
			(3.67)	(2.08)	(3.63)	(−2.09)
R^2 (%)	−0.14	21.40	34.24	25.33	40.62	
			Panel C: Sharpe ratio			
SR				0.72	0.13	0.59
p-value(SR)				-	-	0.040

prices the beta anomaly unconditionally due largely to positive (and significant) loadings on RMW, CMA, and MOM. Next, we investigate whether the FF6 can clear the higher bar of pricing BAB conditionally.

Panel B of Table 3 shows the results from spanning regressions of the returns on BAB on those of the FF6 factors over different sample periods. The first column, "Low," reports a regression of BAB returns on the FF6 factors following months where realized volatility is below its median value. It shows that the alpha increases significantly compared to the "unconditional" case from Panel A and flips signs to a positive, and economically significant, 0.56% per month (with significance at the 1% level). In contrast, the second column of Panel B, "High," reports a similar regression following months where realized variance is above its median value and shows that BAB earns a negative alpha of −0.41% (albeit statistically insignificant). The fourth column, "L−H," presents low-minus-high differences between the second and third columns of Panel B and shows that the alpha is a full 0.97 percentage points higher in months with low volatility compared to those with high volatility and this difference is strongly significant (1% level).

The large increase in abnormal returns of BAB going from high- to low-volatility states reflects the confluence of two patterns. First, Panel B shows that BAB's Sharpe ratio rises more than five fold (0.13 to 0.72), going from high- to low-volatility states, with the difference significant at the 5% level. Second, BAB's exposure to the profitability, investment, and momentum factors that price BAB unconditionally all fall as volatility decreases (actually, turning negative in the case of RMW). As discussed above, Eq. (7) shows that for a risk-based explanation to hold, it must be the case that BAB loads more heavily on risk factors when volatility is low, or for the maximum Sharpe ratio in the economy to increase proportionally with that of BAB. However, since the FF6 factor loadings fall with volatility, for the former possibility to hold, it must be the case that BAB loads very heavily on risk factors that are both missing from the FF6 and economically significant enough to justify Sharpe ratios that exceed 0.70. Even if this were the case, and the correlation between BAB and the stochastic discount factor remained the same, the maximum Sharpe ratio in the economy would also have to (more than) double from its unconditional level to that in the half of months with below-median BAB volatility given corresponding increases in BAB Sharpe ratio of 0.34 (Table 1) to 0.72 ("Low"). But, if the unconditional market Sharpe ratio of 0.5 is deemed a "puzzle," as in Mehra and Prescott (1985), its more than doubling must be even more so.[8] Overall, the evidence in Table 3 supports explanations of the low-risk anomaly that do not require a risk-factor structure in returns because BAB returns seem highest precisely when they require bearing the least observable risk. In interpreting these results, it is also worth noting that "risk-based" explanations need not be rational. Kozak, Nagel and Santosh (2018) show that, in the absence of near arbitrage, asset returns can conform to a factor model even if the fundamental source of return predictability is entirely driven by sentiment. Indeed, Stambaugh et al. (2015) and Stambaugh and Yuan (2017) even classify several characteristics (e.g., momentum, profitability, investment) as mispricing that are the basis for three risk factors in the FF6. In any case, even if the premiums on MOM, RMW, and CMA are entirely driven by sentiment, it is still the case that these factors do not explain the conditional returns on BAB since their correlations with BAB are low precisely when the Sharpe ratio of BAB is highest.

The results in Panel B of Table 3 can be thought of as a special case of a more general asset-pricing model of the form:

$$BAB_t = \alpha_0 + \boldsymbol{\alpha}_1'\mathbf{z}_{t-1} + \sum_{k=1}^{K}(\gamma_{k,0} + \boldsymbol{\gamma}_{k,1}'\mathbf{z}_{t-1})f_{k,t} + \epsilon_t, \tag{13}$$

where $\mathbf{z}_{t-1}$ is a vector of demeaned lagged instruments and $f_{k,t}$ is the return on factor $k = 1, ..., K$. Such

[8] Sharpe ratios of approximately twice the market are often considered "good deals" that are unlikely to survive arbitrage. See, e.g., Ross (1976), Shanken (1992), and Cochrane and Saá-Requejo (2000). MacKinlay (1995) argues that a Sharpe ratio of 0.6 borders on implausible for anomaly factors.

models are widely used in conditional asset-pricing tests following the seminal study of Shanken (1990). If a factor model prices BAB unconditionally, α_0 will be zero, and if it prices BAB conditionally, then every element of the vector $\boldsymbol{\alpha_1}$ will be zero as well. The γ parameters govern how factor betas comove with the corresponding instruments. In the case of Table 3, the only instruments are discrete dummy variables for whether a month t has "high" or "low" volatility. One could instead use the continuous realized BAB volatility and control for other instruments, which is important given that multiple variables are known to predict the future returns on BAB and there is no reason to believe that BAB's volatility *subsumes* every other predictor of BAB's performance. Thus, to more thoroughly understand the conditional mispricing of BAB, we present estimates of the α parameters that characterize the abnormal returns in Eq. (13) and Table IA.1 in the Internet Appendix presents estimated γ parameters that characterize the slopes.

Following Cederburg and O'Doherty (2016), $\mathbf{z}_{t-1}$ includes realized BAB volatility in addition to the default spread, the market dividend-price ratio, and the ex ante beta of BAB ($\beta_{L,t-1} - \beta_{H,t-1}$ in Eq. (4)). Section 4.5 and Internet Appendix Section IA.2 expands the set of instruments further. The instruments are demeaned and standardized such that α_0 is the average abnormal return from the model and elements of $\boldsymbol{\alpha}_1$ represent the impact of a one-standard deviation change in the instrument on the conditional abnormal return. Each pair of columns corresponds to a model, with odd-numbered columns using only the log realized volatility of BAB as an instrument, and even-numbered columns presenting estimates from the unrestricted model given by Eq. (13) (all four instruments).[9]

The first six columns of Table 4 present results for the CAPM, Fama and French (1993) three-factor model (FF3), and FF6. The results show that even conditional versions of either the CAPM and FF3 fail to price BAB unconditionally, as the estimates of α_0 are economically large (above 0.40%) and statistically significant in all four cases. Only for the two conditional specifications of FF6 do we achieve statistically insignificant estimated (unconditional) abnormal returns. We can see that volatility negatively predicts abnormal returns in the single-instrument specifications, albeit this effect is only statistically significant for FF6—a one-standard deviation decrease in BAB volatility predicts an increase in abnormal return of 0.43 percentage points (t-ratio of -2.54). Importantly, Columns (2), (4), and (6) show that including the other three instruments makes the incremental effect of volatility stronger: a one-standard deviation decrease in

[9] Cederburg and O'Doherty (2016) show that a conditional CAPM with similar instruments prices the beta anomaly, although Liu et al. (2018) show this finding is not robust to the technique used to estimate beta. Asness et al. (2020) also argue that the tests from Cederburg and O'Doherty (2016) suffer from low power to reject the null that the conditional CAPM holds and are ruled out by the simple fact that the BAB factors have a *constant* beta of zero by construction so that time-varying betas cannot explain their returns. Some conditional CAPM studies restrict the $\boldsymbol{\alpha}_1 = 0$. We do not do this because time-variation in abnormal returns associated with volatility is central to our study, but this has no impact on α_0 since the elements of $\mathbf{z}_{t-1}$ are demeaned.

Table 4
Performance of BAB relative to conditional factor models

This table reports α_0 and $\boldsymbol{\alpha}_1$ estimates from regressions of the form:

$$BAB_t = \alpha_0 + \boldsymbol{\alpha}_1' \mathbf{z}_{t-1} + \sum_{k=1}^{K} (\gamma_{k,0} + \boldsymbol{\gamma}_{k,1}' \mathbf{z}_{t-1}) f_{k,t} + \epsilon_t,$$

where f_k is the return on factor $k = 1, .., K$. In odd-numbered columns, the vector of instruments, $\mathbf{z}_{t-1}$ only includes the log realized volatility on BAB (VOL). In even-numbered columns, $\mathbf{z}_{t-1}$ also includes the difference between BAA and AAA bond yields (DEF), the log dividend-price ratio on the CRSP value-weighted index (DP), and the difference in market betas between the long and short legs of BAB ($Beta$) before hedging the market exposure. The instruments are demeaned and standardized so that the $\boldsymbol{\alpha}_1$ coefficients are the impact, in percentage points, of a one standard deviation shock to the instrument. In Columns (1) and (2), "CAPM," the only factor is MKT; in Columns (3) and (4), "FF3," the factors are MKT, SMB, and HML; in Columns (5) and (6), "FF6," the factors are all six in the FF6 model; in Columns (7) and (8), the only factor is the mimicking portfolio for growth in the unfiltered consumption series of Kroencke (2017); and in Columns (9) and (10), the only factor is the mimicking portfolio for intermediary leverage of Adrian et al. (2014) (see Section 4.2 for details). Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

	CAPM		FF3		FF6		Consumption		Leverage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
α_0	0.43	0.41	0.43	0.45	0.18	0.30	0.47	0.51	0.25	0.25
	(2.32)	(2.23)	(2.60)	(2.93)	(1.14)	(1.91)	(2.55)	(2.83)	(1.35)	(1.33)
α_1^{VOL}	−0.33	−0.67	−0.29	−0.69	−0.43	−0.59	−0.27	−0.52	−0.41	−0.60
	(−1.35)	(−2.23)	(−1.51)	(−3.07)	(−2.54)	(−2.92)	(−1.19)	(−1.70)	(−1.84)	(−1.96)
α_1^{DEF}		−0.06		0.26		0.40		0.20		0.09
		(−0.20)		(1.16)		(2.08)		(0.76)		(0.33)
α_1^{DP}		−0.31		−0.31		−0.16		−0.44		−0.27
		(−0.98)		(−1.33)		(−0.63)		(−1.36)		(−0.81)
α_1^{Beta}		−0.21		−0.17		0.02		−0.17		−0.20
		(−1.02)		(−0.99)		(0.11)		(−0.82)		(−0.86)

BAB volatility predicts an increase in abnormal returns of around 0.7 percentage points for both CAPM and FF3, with strong significance (1% or 5% level). In the case of FF6, such effect is only marginally smaller with an estimate of -0.59% (t-ratio of -2.92) for α_1^{VOL}.

The last four columns of Table 4 present estimates from models with factors based on macroeconomic conditions. Columns (7) and (8) use a single-factor model similar to the consumption CAPM of Breeden, Gibbons and Litzenberger (1989), but based on a mimicking portfolio for the "unfiltered" annual consumption growth series of Kroencke (2017). This portfolio is formed by projecting the consumption growth series onto the annual excess returns of the base assets used in the FF6 model and then sampling the resulting portfolio at the monthly frequency (see the Internet Appendix for details). Columns (9) and (10) use a single-factor model based on the mimicking portfolio of intermediary leverage from Adrian et al. (2014) based on Federal Reserve Flow of Funds data. They form this factor by projecting the quarterly time series of intermediary leverage growth onto the returns of Fama and French (1993) and Carhart (1997) size, value, and momentum portfolios.

The results show that the consumption growth portfolio fails to price BAB with a significant α_0 of 0.47%–

0.51% per month (Columns (7) and (8)). The leverage model renders α_0 insignificant for BAB (columns (9) and (10)). This evidence is partially consistent with Adrian et al. (2014) and the leverage constraints theory of the beta anomaly. However, BAB volatility negatively predicts subsequent BAB abnormal returns relative to the leverage model, with essentially the same economic magnitude as the empirically motivated FF6 model, although the statistical signifcance is marginal. Similarly to the three equity models, including the other instruments strengthens the negative effect of lagged volatility on BAB's abnormal returns.

In summary, Tables 3 and 4 show that the CAPM and FF3 can't price BAB unconditionally, and while the FF6 can price BAB unconditionally, the returns on BAB become more anomalous when its volatility falls, regardless of model considered. Importantly, the economic significance of such effect is similar across all five models (when all instruments are considered) with a one-standard deviation *decrease* in volatility predicting a 0.5-0.7 percentage point *increase* in monthly abnormal returns.

4.3. Limits to arbitrage

Shleifer and Vishny (1997) and Pontiff (2006) argue that bearing IVOL risk is an important holding cost of arbitrage, and, consistent with this argument, many studies show that anomaly returns increase in the cross section with high levels of IVOL (see, e.g., Ali, Hwang and Trombley, 2003; Mashruwala, Rajgopal and Shevlin, 2006; Doukas, Kim and Pantzalis, 2010; Li and Zhang, 2010; and McLean, 2010). Stambaugh et al. (2012) observe that a disproportionate amount of capital exploits underpricing relative to overpricing and show, presumably as a result of this "arbitrage asymmetry," that overpricing is generally more prevalent than underpricing, and even more so in times of optimistic sentiment. Expanding on this work, Stambaugh et al. (2015) posit that arbitrage risk and asymmetry combine to explain the IVOL puzzle. IVOL deters arbitrage in both underpriced and overpriced stocks, but the effect in overpriced stocks is larger due to the asymmetry. Hence, the effect in overpriced stocks dominates when sorting stocks based on IVOL alone. Consistent with this argument, they show that IVOL is positively related to FF3 alpha in underpriced stocks, but the average relation across all stocks is negative. Liu et al. (2018) extend these results and show that the unconditional beta anomaly results from the strong positive cross-sectional correlation between beta and IVOL. Specifically, they show that (unconditional) BAB abnormal returns are absent in stocks that are not "high-IVOL" (in the highest IVOL quintile) or "overpriced" (highest mispricing quintile). In this subsection, we investigate whether the conditional performance of the beta anomaly is consistent with the anomaly's limits-to-arbitrage explanation.

We form five BAB factors: one that drops high-IVOL stocks, one that drops "overpriced" stocks, one that drops high-IVOL and overpriced stocks, one based on beta orthogonalized cross-sectionally with respect to IVOL, and one based on betas orthogonalized with respect to mispricing. If the limits-to-arbitrage theory explains the conditional performance of BAB, these factors should lose the alphas seen in Table 3 in the low- and high-volatility states.

To form the factor that uses betas orthogonalized with respect to IVOL, each month, t, we estimate a cross-sectional regression of the beta estimates of each stock on the corresponding IVOL for that month:

$$\hat{\beta}_{n,t} = a_t + b_t IVOL_{n,t} + \nu_{n,t} \tag{14}$$

We then form BAB factors in month $t+1$ as in Eq. (4) based on the residuals from these regressions, $\{\nu_{n,t}\}_n$, in lieu of the raw betas. By construction, the $\nu_{n,t}$ isolate variation in beta that is orthogonal to idiosyncratic volatility. The BAB factor based on betas orthogonalized with respect to mispricing is formed similarly, but using mispricing on the right-hand-side of Eq. (14).

Panel A of Table 5 presents FF6 alphas, similar to those in Table 3, for the BAB factor that drops stocks in top quintile of IVOL. Consistent with Table 3, the column labeled "Unc." shows that over the full sample, the BAB factor earns a largely insignificant alpha. However, when volatility is low, these alphas become positive, and negative when volatility is high, with a spread of 0.93 percentage points between the two states that is significant at the 1% level. The corresponding Sharpe ratio of BAB controlling for IVOL is almost sixfold in high- compared to low-volatility states (0.13 to 0.77), and this gap is significant at the 5% level.

In Panel B of Table 5, the overpriced stocks, those that are in the top quintile of mispricing, are excluded when constructing the BAB factor. Panel C presents results for a BAB factor that drops overpriced high-IVOL stocks. The results from both panels are at least as strong as those in Panel A of Table 5. The spread in BAB alphas between low- and high-volatility states is around one percentage point per month and significant at the 1% level in both Panels B and C. Similarly, the Sharpe ratio increases ninefold from high- to low-volatility states (from 0.08 to 0.74) when overpriced high-IVOL stocks are dropped. When excluding only overpriced stocks (Panel B), the Sharpe ratio becomes 0.57 in low-volatility states (from zero in high-volatility states).

Panel D of Table 5 presents results based on the factor that is based on betas orthogonalized to IVOL. The results are qualitatively similar to the results in Panel A that drop high-IVOL stocks. The low-volatility-minus-high-volatility spread in BAB's alphas has about the same magnitude (0.90 percentage points) and statistical significance (t =2.82).

In Panel E of Table 5, we report the results for yet another version of the BAB factor: the one based on betas orthogonalized to mispricing. Despite showing a lower magnitude than in Panel B, the low-volatility-

Table 5

Conditional performance of BAB controlling for idiosyncratic volatility and mispricing

In Panel A (B) we make a BAB factor that goes long the quintile's low-beta portfolio, short the high-beta portfolio, and hedges market risk as in Eq. (4) after dropping stocks in the highest IVOL (mispricing) quintile. In Panel C, we make a similarly constructed BAB factor after dropping stocks that are in both the top mispricing quintile and the top IVOL quintile. In Panels D and E, respectively, we make BAB factors after orthogonalizing beta with respect to IVOL and mispricing using month-by-month cross-sectional regressions of beta on these characteristics. The first row of each panel presents FF6 alphas from regressions of the returns on the corresponding BAB factor on those of the FF6 factors. Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The four columns report: an unconditional regression over the full sample ("Unc."), one using the subsample of months with below-median lagged volatility ("Low"), one using the subsample period with above-median volatility ("High"), and the corresponding low-minus-high differences ("L−H"). The rows labeled SR report the Sharpe ratios (SR) of the BAB factors over the subsample periods specified by the column headings. The bottom row of each panel reports a p-value for the low-minus-high differences in Sharpe ratios based on GMM estimation with heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

	Unc.	Low	High	L−H
		Panel A: Dropping high-IVOL stocks		
Alpha (%)	−0.00	0.54	−0.39	0.93
	(−0.02)	(2.83)	(−1.48)	(2.85)
SR	0.37	0.77	0.13	0.64
p-value(SR)	-	-	-	0.024
		Panel B: Dropping overpriced stocks		
Alpha (%)	−0.26	0.39	−0.62	1.01
	(−1.54)	(2.00)	(−2.33)	(3.07)
SR	0.20	0.57	0.00	0.57
p-value(SR)	-	-	-	0.045
		Panel C: Dropping overpriced high-IVOL stocks		
Alpha (%)	−0.10	0.54	−0.48	1.02
	(−0.59)	(2.78)	(−1.86)	(3.15)
SR	0.31	0.74	0.08	0.66
p-value(SR)	-	-	-	0.022
		Panel D: Sorting on beta orthogonalized to IVOL		
Alpha (%)	−0.14	0.45	−0.45	0.90
	(−0.83)	(2.31)	(−1.78)	(2.82)
SR	0.30	0.61	0.13	0.48
p-value(SR)	-	-	-	0.094
		Panel E: Sorting on beta orthogonalized to mispricing		
Alpha (%)	−0.12	0.39	−0.37	0.76
	(−0.69)	(1.81)	(−1.38)	(2.20)
SR	0.30	0.71	0.07	0.64
p-value(SR)	-	-	-	0.024

minus-high-volatility spread in BAB's alphas is still both economically (0.76 percentage points) and statistically significant (5% level). Moreover, the gain in Sharpe ratio (across volatility states) is marginally higher than in Panel B (with significance also at the 5% level).

Overall, the results in Table 5 show that the limits to arbitrage captured by IVOL and mispricing do not explain the strength of the beta anomaly conditional on low volatility. In other words, our evidence is not consistent with the arbitrage asymmetry explanation of the beta anomaly.

4.4. Lottery preferences

The lottery preferences theory of the beta anomaly posits that some investors demand high-risk stocks, even if they have poor average performance, because they offer a possibility of very high returns. Consistent with this explanation, several studies find evidence that lottery demand drives down expected returns of risky stocks and plays at least a partial role explaining both the beta and IVOL anomalies (e.g., Bali, Cakici and Whitelaw, 2011; Conrad, Kapadia and Xing, 2014; and Bali et al., 2017). Bali et al. (2011) and Bali et al. (2017) measure lottery demand with a variable called MAX, which is defined for each stock and month as the average of the stock's highest five daily returns over that month. They find that MAX subsumes the low-risk anomaly. Asness et al. (2020) argue that stocks can have a high MAX either because their returns are volatile or because they are right-skewed. They measure lottery demand with a variable called SMAX, which scales a given stock's MAX by the stock's standard deviation of daily returns in the month MAX is measured. As a result, SMAX captures skewness on a stock while neutralizing the mechanical effect of volatility.

For the conditional performance of BAB to be consistent with the lottery theory, lottery demand must be lower in high-volatility periods when risky stocks appear less "overvalued" relative to low-risk stocks. But lower demand of lotteries in high-vol states is not consistent with extant facts on lottery preferences. Kumar (2009), and cites therein, all find that lottery demand for a given investor increases during economic downturns when volatility is relatively high. Kumar further argues: "When volatility is high, investors might believe that the extreme return events observed in the past are more likely to be realized again." However, an opposite effect can be at work as high-volatility periods are also "bad times" when investors commonly classified as "sentiment traders" tend to exit the market and these traders may be hard to distinguish empirically from traders characterized by lottery preferences (e.g., Antoniou, Doukas and Subrahmanyam, 2013; Antoniou et al., 2016). Thus, increases in volatility can increase demand of risky stocks by lottery-motivated traders that stay in the market, while simultaneously decreasing the number of such traders in the market, rendering the net effect an empirical question.

We test whether lottery demand can explain the conditional performance of BAB using a similar method as in Table 5, specifically forming BAB factors based on betas that are orthogonalized to lottery demand via Eq. (14). Panel A (B) of Table 6 reports performance statistics for the BAB factors that control for SMAX (MAX) conditional on volatility. Similarly to the results in Table 5, Column "Unc." shows that the FF6 yields an insignificant alpha on BAB, while the alpha rises to a positive and significant 0.49% when volatility is low, and falls to a negative (albeit insignificant) -0.37% when volatility is high. This produces a spread between low- and high-volatility states of 0.86 percentage points that is strongly significant ($t = 2.67$).

Table 6
Conditional performance of BAB controlling for lottery demand

In Panel A (B) we orthogonalize beta with respect to lottery demand proxied by SMAX (MAX) using month-by-month cross-sectional regressions of beta on lottery demand. We then make a BAB factor that goes long the low-beta portfolio, short the high-beta portfolio, and hedges market risk as in Eq. (4) using these orthogonalized betas. The first row of each panel presents FF6 alphas from regressions of the returns on the corresponding BAB factor on those of the FF6 factors. Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The four columns report: an unconditional regression over the full sample ("Unc."), one using the subsample of months with below-median lagged volatility ("Low"), one using the subsample period with above-median volatility ("High"), and the corresponding low-minus-high differences ("L−H"). The rows labeled SR report the Sharpe ratios (SR) of the BAB factors over the subsample periods specified by the column headings. The bottom row of each panel reports a p-value for the low-minus-high differences in Sharpe ratios based on GMM estimation with heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

	Unc.	Low	High	L−H
Panel A: Sorting on beta orthogonalized to SMAX				
Alpha (%)	−0.10	0.49	−0.37	0.86
	(−0.61)	(2.39)	(−1.49)	(2.67)
SR	0.33	0.69	0.12	0.58
	-	-	-	0.043
Panel B: Sorting on beta orthogonalized to MAX				
Alpha (%)	−0.21	0.41	−0.59	1.00
	(−1.28)	(2.13)	(−2.29)	(3.11)
SR	0.25	0.61	0.04	0.57
	-	-	-	0.046

Similarly, the Sharpe ratio of BAB rises significantly (5% level) from high- to low-volatility months by 0.58. The results are at least as strong when using MAX in lieu of SMAX, as indicated in Panel B of Table 6: The spread in alphas between low- and high-volatility states increases to one percentage point (also with significance at the 1% level), while the gap (across volatility states) in Sharpe ratio is basically the same (0.57). Overall, Table 6 shows that lottery demand can not explain BAB returns in different volatility states.

4.5. Sentiment, analyst disagreement, and the beta anomaly

The results above consider explanations of the beta anomaly based on beta's cross-sectional correlation with stock-level characteristics such as mispricing or IVOL. Next, we consider explanations whose evidence is based on time-series predictability.

Antoniou et al. (2016) and Liu et al. (2018) show that the beta anomaly is nearly absent in pessimistic periods, consistent with mispricing of beta being caused by unsophisticated investors drawn to high-beta stocks when they feel optimistic. Barroso and Detzel (2021) document that volatility timing the market portfolio is only profitable when sentiment is high, and significantly diminishes performance during pessimistic periods, reflecting the results of Yu and Yuan (2011) that the time-series risk-return tradeoff of the market is negative in optimistic periods, but positive otherwise. Motivated by these findings, we investigate whether sentiment explains the abnormal returns on BAB when volatility is low.

To do so, following Liu et al. (2018), Panel A of Table 7 presents FF6 alphas of BAB in months divided into four regimes according to whether prior-month Baker and Wurgler (2007) sentiment and BAB realized volatility are "high" or "low" relative to their median value. This panel shows that BAB earns alphas that are 0.68 to 1.28 percentage points higher in low-volatility states than in high-volatility states depending on sentiment. The average alpha spread between low- and high-volatility states is 0.98 percentage points across the two sentiment regimes and is significant at the 1% level. If anything, the negative relationship between BAB's volatility and its subsequent alpha is relatively strong when sentiment is low, as indicated by the strongly significant (1% level) alpha of 1.28%. However, the corresponding high-minus-low-sentiment difference (in the low-minus-high-volatility alpha spreads) is -0.60 percentage points, yet, this estimated difference is statistically insignificant ($t = -0.92$). Another way to see the economic significance of volatility-timing BAB is by looking at its "unconditional" performance (with respect to sentiment), that is comparing the two "extreme" states: BAB's alpha varies from -0.70% (with significance at the 10% level) in the low sentiment-high volatility state to 0.56% (with significance at the 5% level) in the high sentiment-low volatility state, which represents a gap as large as 1.26 percentage points per month. Overall, we conclude from Panel A that sentiment does not subsume volatility in predicting abnormal returns on BAB.

Hong and Sraer (2016) argue that, compared to low-beta stocks, high-beta stocks are more sensitive to aggregate disagreement about the growth prospects of firms, experience greater divergence of opinion about their payoffs, and are thus more overpriced due to short-sales constraints. They propose a measure of aggregate disagreement, which is the beta-weighted average of the cross-sectional standard deviation of long-term earnings growth forecasts from IBES. Panel B of Table 7 investigates whether disagreement subsumes the volatility-timing benefits of BAB by estimating alphas over four regimes similar to Panel A, but using disagreement in lieu of volatility. The spread between low- and high-volatility regimes ranges from 0.92 to 0.58 percentage points across disagreement regimes, averaging 0.75 percentage points (albeit only with marginal significance, $t = 1.70$). This spread is larger within the low-disagreement regime, though with no statistical significance. By looking at the two "extreme" states, we see a similar pattern to that registered in Panel A: BAB's alpha varies from -0.73% in the low disagreement-high volatility state to 0.45% in the high disagreement-low volatility state, which represents a similar gap to that observed for sentiment (1.18 percentage points). Yet, in contrast to the case of sentiment, those alpha estimates are not individually statistically significant. Overall, Panel B shows that disagreement does not appear to subsume volatility, however, the statistical significance of BAB's performance is more marginal than in Panel A.

To provide a robustness check for Table 7, Table IA.2 in the Internet Appendix presents continuous scaled

Table 7

Abnormal returns of BAB in periods of high and low volatility, sentiment, and analyst disagreement

We assign each month to one of four regimes ($j = 1, ..., 4$) according to whether prior-month investor sentiment and realized volatility of BAB (defined in Table 1) are above or below their median values. Panel A reports alphas estimated in each of these regimes from a regression of the form:

$$BAB_t = \sum_{j=1}^{4} D_{j,t} \left(\alpha_j + \beta_j MKT_t + s_j SMB_t + h_j HML_t + m_j MOM_t + r_j RMW_t + c_j CMA_t\right) + \epsilon_t, \tag{15}$$

where $D_{j,t}$ equals one if month t is in regime j, and zero otherwise. Panel B reports similar estimates as Panel A, but using the time series of beta-weighted analyst disagreement about long-term earnings growth forecasts in lieu of sentiment. Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

Panel A: Sentiment

	BAB volatility			
Sentiment	Low	High	Low−High	Average
Low	0.57	−0.70	1.28	−0.07
	(2.14)	(−1.84)	(2.73)	(−0.28)
High	0.56	−0.12	0.68	0.22
	(1.98)	(−0.33)	(1.51)	(0.99)
High−Low	−0.01	0.59	−0.60	0.29
	(−0.03)	(1.14)	(−0.92)	(0.89)
Average	0.57	−0.41	0.98	
	(2.91)	(−1.59)	(3.02)	

Panel B: Disagreement

	BAB volatility			
Disagreement	Low	High	Low−High	Average
Low	0.19	−0.73	0.92	−0.27
	(0.62)	(−1.22)	(1.37)	(−0.80)
High	0.45	−0.13	0.58	0.16
	(1.11)	(−0.32)	(1.01)	(0.55)
High-Low	0.26	0.60	−0.35	0.43
	(0.50)	(0.83)	(−0.39)	(0.97)
Average	0.32	−0.43	0.75	
	(1.26)	(−1.20)	(1.70)	

variable regressions similar to those in Table 4, but including disagreement and sentiment as instruments. This table shows the inferences from Table 7 are robust, the negative effect of volatility on BAB alphas is always significant at the 5% level irrespective of controlling for sentiment or disagreement (in addition to the other instruments). It also shows that disagreement remains at least a marginally significant predictor of BAB abnormal returns after controlling for volatility.

Overall, the results from Table 7, bolstered by Internet Appendix Table IA.2, challenge the view that sentiment or disagreement might subsume the effect of volatility in predicting abnormal returns on BAB.

4.6. Market volatility vs. BAB volatility

Realized variance of returns is highly correlated across factors (Moreira and Muir, 2017; Eisdorfer and Misirli, 2020; and DeMiguel, Martín-Utrera and Uppal (2022)). Table 8 investigates whether the conditional

Table 8
Abnormal returns of BAB in periods of high- and low- own-factor and market volatility

This table assigns months to one of four regimes based according to whether prior-month realized market volatility and realized BAB volatility are above or below their median value and presents FF6 alphas of BAB (defined in Table 1) in each of the four regimes. Regression slopes and intercepts are estimated separately in each regime as in Eq. (15). Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The sample period is July 1965 to December 2016.

Market	BAB volatility			
volatility	Low	High	Low−High	Average
Low	0.56	0.14	0.42	0.35
	(2.61)	(0.38)	(0.99)	(1.64)
High	0.40	−0.43	0.83	−0.02
	(0.96)	(−1.32)	(1.57)	(−0.06)
High−Low	−0.16	−0.57	0.41	−0.37
	(−0.35)	(−1.16)	(0.60)	(−1.08)
Average	0.48	−0.15	0.63	
	(2.05)	(−0.59)	(1.84)	

performance of BAB documented in this section is driven by its own volatility or that of the market. Similar to the sentiment results from the previous table, market volatility clearly does not subsume BAB volatility. The average (across market volatility regimes) alpha spread between low- and high-BAB volatility states is 0.63 percentage points and marginally significant, while the average (across BAB volatility regimes) spread between high- and low-market volatility states is an insignificant (even at the 10% level) −0.37 percentage points. Thus, it appears to be the specific risk of BAB that drives its performance in low-volatility states, not the risk of the market.[10]

5. Institutional ownership and betting against beta

The results above show that the conditional performance of BAB is a puzzle relative to leading explanations of the beta anomaly. For this performance to exist, it must be the case that high-beta stocks become less "overvalued" relative to low-beta stocks when the volatility of these factors increases. In this section, we use the demand system framework of Koijen and Yogo (2019) to investigate whether changes in institutional demand for high- and low-beta stocks can explain these changes in valuations sufficiently to produce the spread in BAB abnormal returns we observe between high- and low-volatility states. Even if they do not, documenting who trades on the beta anomaly relative to volatility provides important facts for the next generation of beta anomaly theories to explain.

Numerous studies document incentives that many institutional investors have that plausibly lead them

[10] In the Internet Appendix, we also consider BAB performance conditional on volatility and lagged market return in lieu of market volatility in Table 8.

to have a relatively high demand for high-beta stocks when volatility is low. Baker et al. (2011) show theoretically that fund managers judged relative to a unit-beta benchmark like the S&P 500, which is standard in mutual funds, pension funds, and many hedge funds, will overweight high-beta stocks, even if they have negative alpha, because their returns exceed those of the benchmark on average (and they face leverage constraints). Consistent with this theory, Christoffersen and Simutin (2017) show that, in the cross-section, mutual fund managers facing the greatest pressure to beat benchmarks tilt their portfolios toward high-beta stocks and Boguth and Simutin (2018) show that active mutual funds have betas that are greater than one on average. While standard performance evaluation incentives reward beating the benchmark on average, they also penalize excessive tracking-error volatility. This tradeoff gives managers an incentive to overweight high-beta stocks unconditionally; but, when these stocks' contribution to tracking-error volatility increases, managers have the incentive to shift their portfolio weights towards those of the benchmark. When BAB volatility increases, the potential contribution of high- and low-beta stocks to tracking error volatility relative to unit-beta indexes should mechanically increase as well because both quantities are driven by the market variance and the residual variance of extreme-beta stocks.[11] Taken together, while managers prefer to unconditionally overweight high-beta stocks, an increase in BAB volatility encourages them to shift their portfolio weights towards those of the benchmark, which increases demand for low-beta stocks and decreases demand for high-beta stocks.

Equity fund managers at insurance companies and banks do not necessarily face the same explicit benchmark-based contracts as their counterparts working for pension funds and mutual funds, but they also plausibly have incentives to outperform an equity index on average while retreating from risky positions in highly volatile states. Christoffersen and Simutin (2017) show that pension funds put particular emphasis on benchmarks when evaluating performance of mutual funds they invest in (and this contributes to the beta anomaly, at least unconditionally). Del Guercio (1996) argues that equity portfolios of banks largely exist to serve pension fund and private trust clients and therefore these portfolios' plausibly face similar pressures as mutual funds working for pension clients. Banks also have internal and regulatory risk limits that could force reducing risk in sufficiently volatile times. Ellul et al. (2022) shows that life insurers "reach for yield" to beat equity indexes because their largest liabilities are variable annuities that make payouts based on the performance of the stock market. At the same time, these insurers also face internal risk limits and risk-based capital requirements, so it stands to reason that they face pressure to reduce betas of their equity portfolios

[11]To see this, note that the tracking error volatility of a fund return, r_n, relative to the market benchmark, r_{MKT}, is given by: $\sqrt{\sigma^2(r_n - r_{MKT})} = \sqrt{(1-\beta_n)^2\sigma^2(r_{MKT}) + \sigma^2(\epsilon_n)}$.

when volatility increases substantially. Said differently, these insurers earn revenue by having higher returns than those of the equity market, but must reduce their positions in risky assets in high-risk times.

5.1. Demand system

The market has N stocks, $n = 1, ..., N$, and an outside asset, $n = 0$, traded by I investors consisting of $I - 1$ institutions required to file form 13f (those with at least \$100 million of assets under management), denoted $i = 1, ..., I - 1$, and "households," denoted $i = I$, who hold all stocks not held by 13f institutions. Lower case variables denote natural logarithms of uppercase variables, e.g., $me_t(n) = \log(ME_t(n))$; and bold font denotes vectors. Asset n has a vector of observable characteristics $\boldsymbol{x}_t(n) \in \mathbb{R}^K$ besides market value of equity, $ME_t(n)$, and the beta defined in Eq. (2), which we denote in this section as $\beta_t(n)$. Investor $i's$ assets under management are denoted $A_{i,t}$ and the "investment universe" of assets they are allowed to hold is denoted as $\mathcal{N}_{i,t} \subset \{1, ..., N\}$.

The weight of asset n in investor i's portfolio is denoted $w_{i,t}(n)$ and is modeled by:

$$w_{i,t}(n) = \frac{\delta_{i,t}(n)}{1 + \sum_{m=1}^{N} \delta_{i,t}(m)}, \tag{16}$$

where $\delta_{i,t}(n)$ is the weight normalized by investor i's weight in the outside asset:

$$\delta_{i,t}(n) := \frac{w_{i,t}(n)}{w_{i,t}(0)} = \exp\left(\theta_{i,t}^{me} me_t(n) + \theta_{i,t}^{\beta}\beta_t(n) + \boldsymbol{\theta}_t^{x\prime}\boldsymbol{x}_t(n) + \theta_{0,i,t}\right)\epsilon_{i,t}(n). \tag{17}$$

That weights sum to one across positions requires:

$$w_{i,t}(0) = \frac{1}{1 + \sum_{m=1}^{N} \delta_{i,t}(m)}. \tag{18}$$

Eq. (17) maps the demand for a given stock to its characteristics and latent demand, $\epsilon_{i,t}(n)$, in a way that varies across institutions and time. Weights and characteristics are observable, so we can estimate Eq. (17) via nonlinear GMM for each institution-quarter using the moment condition:

$$E\left(\epsilon_{i,t}(n) | \widehat{me}_t(n), \boldsymbol{x}_t(n), \beta_t(n)\right) = 1. \tag{19}$$

Since $me_t(n)$ is endogenous, we use the instrumental variable of Koijen and Yogo (2019), defined as:

$$\widehat{me}_{i,t}(n) := \log\left(\sum_{j \neq i} A_{j,t} \frac{1_{j,t}(n)}{1 + \sum_{m \in \mathcal{N}_{i,t}} 1_{j,t}(m)}\right), \tag{20}$$

where $1_{j,t}(n)$ equals one if $n \in \mathcal{N}_{j,t}$ and zero otherwise. This instrument depends only on the investment universe of other investors and the wealth distribution, which are exogenous (to investor i) by assumption. It can be interpreted as the counterfactual market equity, at the market clearing price, if other investors were to hold an equal-weighted portfolio within their investment universe.[12] The characteristics in $\boldsymbol{x}_t(n)$ are: log book equity, profitability, investment, and dividends to book equity, all defined the same as Koijen and Yogo (2019).[13] The choice of book equity, profitability, and investment is explicitly motivated by the Fama-French five-factor model, which is helpful given that we are trying to explain abnormal returns relative to factors based on these characteristics.

5.2. Institutional demand for high- and low-beta stocks

For each institution-quarter since 1980q1 (when 13f data become available), we estimate the demand system, Eqs. (17) and (19), via nonlinear GMM following Koijen and Yogo (2019). The "Unc." row of Table 9 presents time-series means of quarterly assets-under-management-weighted averages, $\bar{\theta}_t^{\beta}$, of the parameter $\theta_{i,t}^{\beta}$, by institution type. This row shows that, on average, all institutional investors have a positive preference for beta. In contrast, households have a negative preference for beta on average and the "Low vol" and "High vol" rows show that household demand remains negative in both volatility states.[14] In contrast, all institutional investor types significantly drop their demand for high-beta stocks as volatility goes from low to high. For "all" institutions, this drop is essentially to zero (0.01, $t = 0.55$), and in the case of investment advisors (which includes hedge funds), it even becomes negative (-0.06, $t = -2.57$). Taken together, this evidence is consistent with institutional demand (controlling for size, value, profitability, and investment) driving an "overvaluation" of high-beta stocks in states with low BAB volatility compared to those with high volatility.

5.3. Market clearing and counterfactual experiments

Assuming that shares outstanding are exogenous in the short run, we can relate the demand system to prices via market clearing:

$$ME_t(n) = \sum_{i=1}^{I} A_{i,t} w_{i,t}(n). \tag{23}$$

---

[12]Koijen and Yogo (2019) also show the demand system results are robust to using other natural modifications to this instrument.

[13]To estimate the demand system, we use the replication code from Ralph Koijen's website, `https://koijen.net/index.html`, modified only to use our definition of beta from Section 2.

[14]The household coefficients should be interpreted with a caveat: An increase in volatility can coincide with the reduction in value of many institutions' portfolio values, and therefore some small institutions can be reduced to "households" following high-volatility quarters.

Table 9
Demand system coefficients by institution type

For each institution-quarter, (i,t), we estimate the following demand system for portfolio weights via GMM:

$$E\left(\epsilon_{i,t}(n)|\widehat{me}_t(n), \beta_{i,t}(n), \boldsymbol{x}_t(n)\right) = 1, \tag{21}$$

$$\frac{w_{i,t}(n)}{w_{i,t}(0)} = \exp\left(\theta_{i,t}^{me} me_t(n) + \theta_{i,t}^{\beta}\beta_t(n) + \boldsymbol{\theta}_t^{x\prime}\boldsymbol{x}_t(n) + \theta_{0,i,t}\right)\epsilon_{i,t}(n), \tag{22}$$

where $n = 1, ..., N_t$ indexes stocks that exist at time t, and $w_{i,t}(n)/w_{i,t}(0)$ is the portfolio weight in stock n normalized by the weight in the outside asset. This table reports time-series means of quarterly assets under management-weighted averages, $\bar{\theta}_t^{\beta}$, of the parameter $\theta_{i,t}^{\beta}$, over: the full sample (row "Unc."), the subsample with below-median realized volatility ("Low vol"), and the subsample with above-median volatility ("High vol"), as well as low-minus-high difference ("Low−High"). Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The sample period is 1980q1 to 2016q4.

	Households	Banks	Insurance Companies	Investment Advisors	Mutual Funds	Pension Funds	Other	All Institutions
Unc.	−0.10	0.06	0.08	0.06	0.10	0.16	0.19	0.09
	(−6.42)	(4.19)	(3.51)	(3.01)	(7.23)	(7.20)	(4.75)	(5.93)
Low vol	−0.10	0.12	0.14	0.19	0.19	0.30	0.33	0.18
	(−4.30)	(4.76)	(3.33)	(5.61)	(6.99)	(7.21)	(4.45)	(6.43)
High vol	−0.11	−0.00	0.03	−0.06	0.03	0.04	0.08	0.01
	(−4.76)	(−0.03)	(1.20)	(−2.57)	(2.11)	(1.77)	(1.81)	(0.55)
Low−High	0.00	0.13	0.11	0.25	0.16	0.26	0.25	0.17
	(0.10)	(4.28)	(2.42)	(6.11)	(5.21)	(5.46)	(2.92)	(5.41)

Market clearing and Eq. (17) implicitly defines market price as a function of observable characteristics and latent demand. Our primary strategy to quantify the impact of change in institutional demand on BAB abnormal returns is to perform a counterfactual experiment and then compare the resulting valuation impact implied by market clearing across high- and low-volatility states. In this experiment, we seek to identify institutions with the incentives described above, namely those who aggressively lower their betas when BAB volatility increases, reallocate these institutions' assets to other investors, and calculate counterfactual stock prices via Eq. (23). If these institutions drive the time-varying beta anomaly documented in Section 4, then dropping these institutions from the sample and reallocating their assets to other institutions should cause the beta anomaly to attenuate when volatility is low, which will manifest in counterfactual valuations as a reduction in the price of high-beta stocks and an increase in valuation of low-beta stocks. Following Han et al. (2022), we then use the Campbell and Shiller (1988) decomposition to map how these changes in valuations across low- and high-volatility states compare to those required to produce the BAB alphas we see in the data.

Armed with counterfactual stock prices, we compute the price impact measure of a portfolio p (e.g., high-or low-beta stocks) as:

$$PI_t(p) := \frac{ME_t^{CF}(p) - ME_t(p)}{ME_t(p)} = \frac{\sum_{n\in p} ME_t^{CF}(n) - \sum_{n\in p} ME_t(n)}{\sum_{n\in p} ME_t(n)}, \tag{24}$$

where $ME_t^{CF}(\cdot)$ denotes the market capitalization of a given asset in the counterfactual. To evaluate these price impact measures, we next derive a benchmark estimate of the price impact that would result from eliminating the alpha of the BAB factor. Applying the Campbell and Shiller (1988) decomposition to an arbitrary portfolio p with (log) return $r_{p,t}$:

$$r_{p,t} \approx \kappa_0 + \kappa_1 pd_{p,t} - pd_{p,t-1} + dg_{p,t}, \tag{25}$$

where $pd_{p,t}$ denotes the log price-dividend ratio, $dg_{p,t}$ denotes log dividend growth, and $\kappa_1 = 0.996$.[15] Solving forward for $pd_{p,t}$ and taking expectations yields:

$$pd_{p,t} = \frac{\kappa_0}{1-\kappa_1} + E_t\left(\sum_{h=0}^{\infty} \kappa_1^h (dg_{p,t+1+h} - r_{p,t+1+h})\right). \tag{26}$$

Next, we assume that $r_{p,t}$ follows the factor structure:

$$r_{p,t} = \alpha_{p,t-1} + \boldsymbol{\beta}'_{p,t-1}\boldsymbol{f}_t + \epsilon_{p,t}, \tag{27}$$

and the abnormal return decays according to an autoregressive process:

$$\alpha_{p,t} = \phi_p \alpha_{p,t-1} + \eta_{p,t}, \tag{28}$$

with $0 < \phi_p < 1$. If we further assume that dividends and the factor structure of the asset are the same in the counterfactual as they are in reality, the percentage price impact of a portfolio p in this experiment relative to the baseline case is given by (see Appendix Appendix A for details):

$$PI_t(p) = \exp\left(\frac{\alpha_{p,t} - \alpha_{p,t}^{CF}}{1-\phi\kappa_1}\right) - 1. \tag{29}$$

[15] Our alpha estimates are based on monthly returns and this choice of κ_1 comes from Campbell and Ammer (1993) for the market portfolio. This parameter is less than, but close to one (it is approximately one minus the monthly dividend yield, which is assumed to be 4% per year). For simplicity, we follow Han et al. (2022) and assume that κ_1 is constant across portfolios but our inferences are robust to any plausible choices of κ_1 for each portfolio.

If an asset is "overvalued" in the real-world ($\alpha_p < 0$), but "fairly priced" in the counterfactual ($\alpha_p^{CF} = 0$), then $PI_t(p) < 0$, and vice versa.

Using L and H to denote the low- and high-beta portfolios, respectively, we can define the price impact of the BAB portfolio as $PI_t(L-H) = PI_t(L) - PI_t(H)$. Using subscripts LV and HV to denote a representative month in a low- and high-volatility state, respectively, we then compare the price impact of the BAB portfolio across low- and high-volatility states:

$$\begin{aligned} PI_{LV-HV}(L-H) &:= PI_{LV}(L-H) - PI_{HV}(L-H) \\ &= \exp\left(\frac{\alpha_{L,LV} - \alpha_{L,LV}^{CF}}{1-\phi\kappa_1}\right) - \exp\left(\frac{\alpha_{H,LV} - \alpha_{H,LV}^{CF}}{1-\phi\kappa_1}\right) \\ &\quad - \left[\exp\left(\frac{\alpha_{L,HV} - \alpha_{L,HV}^{CF}}{1-\phi\kappa_1}\right) - \exp\left(\frac{\alpha_{H,HV} - \alpha_{H,HV}^{CF}}{1-\phi\kappa_1}\right)\right]. \end{aligned} \tag{30}$$

What price impact would be required to eliminate the spread in FF6 BAB alphas of 0.97 percentage points between low- and high-volatility states reported in Table 3? Consider a counterfactual universe in which $\alpha^{CF} = 0$. In the real world, untabulated regressions show the low-beta portfolio in BAB earns an FF6 alpha of 0.17% in the low-volatility state, and −0.18% in the high-volatility state. The high-beta portfolio earns alpha of −0.37% and 0.27% in these states, respectively. Departing from Han et al. (2022), based on our results in Tables 3 and 4, we further assume that $\alpha_{p,t} = a_p + b_p\hat{\sigma}_t$, for some constant $b_p < 0$, where $\hat{\sigma}_t$ is the log realized volatility of BAB. This assumption implies that the AR(1) coefficient, ϕ, in Eq. (28), is equal to the AR(1) coefficient of $\hat{\sigma}_t$, which we estimate to be 0.77. Under these assumptions and Eq. (30), the price impact required to render $\alpha^{CF} = 0$ is:

$$\begin{aligned} PI_{LV-HV}(L-H) &= \exp\left(\frac{0.17\%}{1-0.77\cdot 0.996}\right) - \exp\left(\frac{-0.37\%}{1-0.77\cdot 0.996}\right) \\ &\quad - \left[\exp\left(\frac{-0.18\%}{1-0.77\cdot 0.996}\right) - \exp\left(\frac{0.27\%}{1-0.77\cdot 0.996}\right)\right] \\ &= 4.24\%. \end{aligned} \tag{31}$$

Using the FF3 model, ignoring volatility, and only considering the high-beta portfolio, Han et al. (2022) estimate $\phi = 0.5$, which lowers the figure in Eq. (31) to 2.0%.

Table 10 reports time-series means of the repricing statistics estimated each period from our counterfactual experiment. Following Han et al. (2022), this experiment seeks to identify institutions that most intensely face the incentives described above, drop them from the sample, transfer their assets to the remaining institutions, and then determine the resulting price impact on high- and low-beta stocks. Each quarter, we rank institutions based on the sensitivity of the betas of their portfolios to volatility over a rolling 40-quarter window and "shut down" the institutions that most aggressively reduce their betas when volatility rises, and reallocate their assets under management to all remaining investors. To do so, for each institution-quarter, (i, t), we estimate the institution-level beta, $\beta_{i,t}$, as the weighted average of the stock-level betas in that institution's portfolio. We then estimate the sensitivity of this beta to BAB volatility by estimating the following rolling regression (20 observations minimum):

$$\begin{aligned} \beta_{i,\tau} &= a_{i,t} + b_{i,t}\hat{\sigma}_{\tau-1} + \nu_{i,\tau}, \\ \tau &= t-40, ..., t-1 \end{aligned} \tag{32}$$

We then rank all institutions in quarter t based on their estimated $b_{i,t}$ coefficients and then drop the institutions with the most negative $b_{i,t}$ whose cumulative assets under management is 10% of that of all investors combined.[16] Intuitively, this procedure identifies the institutional investors that face the greatest pressure to reduce their positions in high-beta stocks when volatility increases.

Table 10 shows that, on average (row "Unc"), dropping institutions with high beta sensitivity causes low-beta stocks to appreciate by 4.23% more than high-beta stocks. In low-volatility states (row "Low vol"), where the beta anomaly is the strongest, dropping the "volatility-sensitive" institutions leads to a greater impact on relative valuations than in high-volatility states, consistent with higher relative "overvaluation" of high-beta stocks when volatility is low. The average spread in BAB valuation changes between the two states is 6.91%, more than the 4.24% price impact from Eq. (31) required to produce the difference in BAB alphas between low- and high-volatility states (and several times more than the 2.0% price impact implied by the Han et al., 2022 estimate of ϕ).

Overall, Table 10 clearly shows that institutions whose betas drop the most with increasing volatility impact prices in a manner consistent with contributing to the time-varying beta anomaly and the demand from those representing just 10% of aggregate assets under management can account for essentially all of the conditional performance of BAB.

[16]With the 20-observation minimum, the sample period begins in 1985q1, with the twenty quarters from 1980q1 to 1984q4 used to identify institutions' beta sensitivity to volatility used in 1985q1.

Table 10
Counterfactual repricing statistics

Based on the estimated demand system defined in Table 9, each quarter, we compute the price impact statistics, defined in Eq. (24), of the low- and high-beta portfolios ("$PI(L)$" and "$PI(H)$," respectively) in BAB from a counterfactual experiment. In this experiment, we drop institutional investors with the greatest sensitivity of portfolio betas with respect to volatility, estimated via Eq. (32), and then reallocate their assets to all remaining institutions. Each panel reports time-series means of these price impact statistics (in units of %) over the whole sample ("Unc"), months with below-median BAB volatility ("Low vol"), and months with above-median volatility ("High vol"). Parentheses below point estimates contain t-statistics based on White (1980) heteroskedasticity-robust standard errors. The sample period is 1985q1 to 2016q4.

	$PI(L)$	$PI(H)$	Low−High
Unc	2.21	−2.02	4.23
	(4.41)	(−3.27)	(3.98)
Low vol	3.97	−4.47	8.44
	(5.45)	(−6.30)	(6.20)
High vol	1.07	−0.45	1.53
	(1.66)	(−0.52)	(1.06)
Low−High	2.90	−4.01	6.91
	(2.98)	(−3.59)	(3.49)

6. Conclusion

For fifty years, the finance literature has struggled to explain the puzzling weak cross-sectional relationship between beta and return, which challenges foundational principles of asset pricing. This literature posits a wide range of explanations for the anomaly. Perhaps the earliest of them, the leverage constraints theory argues that high-beta stocks appear overvalued because they give relatively risk-tolerant investors the chance to earn high returns without borrowing. More recently, some argue that investors value the lottery-like payoffs of high-risk stocks or are naively optimistic about their prospects. Others argue that limits to arbitrage allow this anomaly to persist or that it can be explained by missing risk factors.

These explanations all have empirical backing, however, prior studies do not rigorously test them in a conditional setting. The abnormal returns and Sharpe ratios of betting-against-beta factors based on the anomaly rise dramatically when the ex ante volatility of these factors is low. This leads us to inquire which explanations can simultaneously match the cross-sectional risk-return patterns and this element of time-series predictability. The resulting exercises show that this dual achievement is more demanding than what the current state of theory can offer.

We show that the leverage constraints theory counterfactually predicts a positive relationship between volatility and subsequent Sharpe ratios of beta factors, intuitively because it depends critically on the assumption of risk averse investors who naturally demand a positive risk-return tradeoff in both the cross-section and time-series. Empirically, we find that multifactor asset pricing models fail to explain the conditional returns of betting-against-beta factors. The loadings that explain their returns unconditionally generally shrink when realized variance falls such that they earn most anomalous returns precisely when they have the lowest risk. This fact remains true for beta factors that control for the cross-sectional effects of lottery preferences and limits to arbitrage. Overall, we bring a new time-series dimension to test leading theories of the beta anomaly and find these theories generally fail to explain this anomaly's conditional performance.

The results in this paper highlight serious problems associated with ignoring conditioning information in asset pricing research. Five decades worth of explanations for the beta anomaly uniformly fail to explain its performance conditional on volatility. While many studies estimate conditional asset pricing models, these studies are generally the exception, and not the norm, with unconditional tests still dominating the literature.

Towards an explanation of our findings, we show that institutional investors have strong demand for high-beta stocks when volatility is low, but this demand diminishes as risk increases. We further show quantitatively that the value impact of this shifting demand is sufficient to account for the entire difference in abnormal returns on the betting-against-beta strategy between high- and low-volatility regimes. Incentives from widely used benchmark-based performance-evaluation contracts and various institutional liabilities and risk constraints can explain these trading patterns. They motivate institutions to robustly demand high-beta stocks when volatility is low to beat unit-beta benchmarks like the S&P500 on average; but they also punish tracking error, which encourages a retreat from high- to low-beta stocks when volatility increases. Regardless of the precise sources of incentives, the conditional performance of the betting-against-beta strategy and our institutional demand results substantially raise the bar for the next generation of theories of the low-risk anomaly.

Appendix A. Price impact benchmark

Here we derive the price impact that would obtain if alpha of the BAB factor was zero in all states. This quantity serves as a benchmark to determine how large repricing statistics from counterfactual experiments would have to be in order to explain the difference in BAB abnormal returns between low- and high-volatility states. We follow Han et al. (2022), who make a similar statistic for a high-beta portfolio defined similarly to ours. Consider the Campbell and Shiller (1988) decomposition, applied to an arbitrary portfolio p:

$$r_{p,t} = \kappa_0 + k_1 pd_{p,t} - pd_{p,t-1} + dg_{p,t}, \tag{A.1}$$

with $k_1 \approx 0.996$. It follows from Eq. (A.1) that:

$$pd_{p,t} = \frac{\kappa_0}{1-\kappa_1} + E_t\left(\sum_{h=0}^{\infty} \kappa_1^h (dg_{p,t+1+h} - r_{p,t+1+h})\right) \tag{A.2}$$

Next, assume that $r_{p,t}$ follows the factor structure:

$$r_{p,t} = \alpha_{p,t-1} + \boldsymbol{\beta}'_{p,t-1}\boldsymbol{f}_t + \epsilon_{p,t}, \tag{A.3}$$

and α has the following dynamics:

$$\alpha_{p,t} = \phi_p \alpha_{p,t-1} + \eta_{p,t}, \tag{A.4}$$

with $0 < \phi < 1$. Combining the above and the law of iterated expectations yields:

$$pd_{p,t} = \frac{k_0}{1-k_1} + \overline{dg - \beta f}^{\infty}_{p,t} - \frac{\alpha_{p,t}}{1-\phi\kappa_1}, \tag{A.5}$$

where $\overline{dg - \beta f}^{\infty}_{p,t} := E_t\left(\sum_{h=0}^{\infty} \kappa_1^h (dg_{p,t+1+h} - \boldsymbol{\beta}'_{p,t+h}\boldsymbol{f}_{t+1+h})\right)$. Let the percentage price impact of an asset p resulting from setting $\alpha_t \equiv 0$, while keeping all else equal to the "real world," be denoted $PI_t(p)$. In particular, assuming the path of factor loadings and level of dividends do not change, the term $\overline{dg - \beta f}^{\infty}_{p,t}$ coincides with its real-word value and:

$$\begin{aligned} PI_t(p) &= \frac{ME_t^{CF}(p) - ME_t(p)}{ME_t(p)} \\ &= \left(\frac{PD_{p,t}^{CF} - PD_{p,t}}{PD_{p,t}}\right) \\ &= \exp\left(pd_{p,t}^{CF} - pd_{p,t}\right) - 1 \\ &= \exp\left(\frac{\alpha_{p,t} - \alpha_{p,t}^{CF}}{1-\phi\kappa_1}\right) - 1, \end{aligned} \tag{A.6}$$

where we use the fact that the level of dividends is the same in both worlds. In our counterfactual experiment, we consider the impact of setting institutional preferences for beta to zero, that is, $\theta_{it}^{\beta} = 0$. We estimate the relative valuation impact, $PI_t(L-H) := PI_t(L) - PI_t(H)$, across low- and high-beta portfolios. We then compare this difference across low- and high-volatility states:

$$
\begin{aligned}
PI_{LV-HV}(L-H) :=& \\
& \exp\left(\frac{\alpha_{L,LV}-\alpha^{CF}_{L,LV}}{1-\phi\kappa_1}\right)-\exp\left(\frac{\alpha_{H,LV}-\alpha^{CF}_{H,LV}}{1-\phi\kappa_1}\right) \\
& -\left[\exp\left(\frac{\alpha_{L,HV}-\alpha^{CF}_{L,HV}}{1-\phi\kappa_1}\right)-\exp\left(\frac{\alpha_{H,HV}-\alpha^{CF}_{H,HV}}{1-\phi\kappa_1}\right)\right].
\end{aligned}
\tag{A.7}
$$

The α_t are assumed to be continuously time varying, while the estimates in Table 3 are based on discrete high/low-volatility regimes. Thus, Eq. (A.7) should be thought of as the price impact required to eliminate the alphas of BAB in "typical" months where the $\alpha_t = \alpha_{HV}$ or $\alpha_t = \alpha_{LV}$ depending on whether t is in the high- or low-volatility regime, respectively.

References

Adrian, T., Etula, E., Muir, T., 2014. Financial intermediaries and the cross-section of asset returns. Journal of Finance 69, 2557–2596.

Ali, A., Hwang, L.S., Trombley, M.A., 2003. Arbitrage risk and the book-to-market anomaly. Journal of Financial Economics 69, 355–373.

Ang, A., 2014. Asset management: A systematic approach to factor investing. Oxford University Press.

Antoniou, C., Doukas, J.A., Subrahmanyam, A., 2013. Cognitive dissonance, sentiment, and momentum. Journal of Financial and Quantitative Analysis 48, 245–275.

Antoniou, C., Doukas, J.A., Subrahmanyam, A., 2016. Investor sentiment, beta, and the cost of equity capital. Management Science 62, 347–367.

Asness, C., Frazzini, A., Gormsen, N.J., Pedersen, L.H., 2020. Betting against correlation: Testing theories of the low-risk effect. Journal of Financial Economics 135, 629–652.

Baker, M., Bradley, B., Wurgler, J., 2011. Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. Financial Analysts Journal 67, 40–54.

Baker, M., Wurgler, J., 2007. Investor sentiment in the stock market. Journal of Economic Perspectives 21, 129–152.

Bali, T.G., Brown, S.J., Murray, S., Tang, Y., 2017. A lottery-demand-based explanation of the beta anomaly. Journal of Financial and Quantitative Analysis 52, 2369–2397.

Bali, T.G., Cakici, N., Whitelaw, R.F., 2011. Maxing out: Stocks as lotteries and the cross-section of expected returns. Journal of Financial Economics 99, 427–446.

Barroso, P., Detzel, A., 2021. Do limits to arbitrage explain the benefits of volatility-managed portfolios? Journal of Financial Economics 140, 744–767.

Barroso, P., Maio, P.F., 2024. The risk-return tradeoff among equity factors. Journal of Empirical Finance (forthcoming).

Barroso, P., Santa-Clara, P., 2015. Momentum has its moments. Journal of Financial Economics 116, 111–120.

Black, F., 1972. Capital market equilibrium with restricted borrowing. Journal of Business 45, 444–455.

Black, F., Jensen, M.C., Scholes, M.S., 1972. The capital asset pricing model: Some empirical tests, in: Jensen, M. (Ed.), Studies in the Theory of Capital Markets. Praeger Publishers Inc.

Boguth, O., Simutin, M., 2018. Leverage constraints and asset prices: Insights from mutual fund risk taking. Journal of Financial Economics 127, 325–341.

Breeden, D.T., Gibbons, M.R., Litzenberger, R.H., 1989. Empirical tests of the consumption-oriented capm. Journal of Finance 44, 231–262.

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

Campbell, J.Y., Ammer, J., 1993. What moves the stock and bond markets? a variance decomposition for long-term asset returns. Journal of Finance 48, 3–37.

Campbell, J.Y., Hilscher, J., Szilagyi, J., 2008. In search of distress risk. Journal of Finance 63, 2899–2939.

Carhart, M.M., 1997. On persistence in mutual fund performance. Journal of Finance 52, 57–82.

Cederburg, S., O'Doherty, M.S., 2016. Does it pay to bet against beta? On the conditional performance of the beta anomaly. Journal of Finance 71, 737–774.

Cederburg, S., O'Doherty, M.S., Wang, F., Yan, X.S., 2020. On the performance of volatility-managed portfolios. Journal of Financial Economics 138, 95–117.

Christoffersen, S.E.K., Simutin, M., 2017. On the demand for high-beta stocks: Evidence from mutual funds. Review of Financial Studies 30, 2596–2620.

Cochrane, J., Saá-Requejo, J., 2000. Beyond arbitrage: Good-deal asset price bounds in incomplete markets. Journal of Political Economy 108, 79–119.

Cochrane, J.H., 2005. Asset Pricing. Princeton University Press, Princeton, N.J.

Cochrane, J.H., 2011. Presidential address: Discount rates. Journal of Finance 66, 1047–1108.

Conrad, J., Kapadia, N., Xing, Y., 2014. Death and jackpot: Why do individual investors hold overpriced stocks? Journal of Financial Economics 113, 455–475.

Cooper, M.J., Gulen, H., Schill, M.J., 2008. Asset growth and the cross-section of stock returns. Journal of Finance 63, 1609–1651.

Daniel, K., Titman, S., 2006. Market reactions to tangible and intangible information. Journal of Finance 61, 1605–1643.

Daniel, K.D., Moskowitz, T.J., 2016. Momentum crashes. Journal of Financial Economics 122, 221–247.

Del Guercio, D., 1996. The distorting effect of the prudent-man laws on institutional equity investments. Journal of Financial Economics 40, 31–62.

DeMiguel, V., Martín-Utrera, A., Uppal, R., 2022. A multifactor perspective on volatility-managed portfolios, London Business School working paper.

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

Doukas, J.A., Kim, C.F., Pantzalis, C., 2010. Arbitrage risk and stock mispricing. Journal of Financial and Quantitative Analysis , 907–934.

Eisdorfer, A., Misirli, E.U., 2020. Distressed stocks in distressed times. Management Science 66, 2452–2473.

Ellul, A., Jotikasthira, C., Kartasheva, A., Lundblad, C.T., Wagner, W., 2022. Insurers as Asset Managers and Systemic Risk. The Review of Financial Studies 35, 5483–5534.

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

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

Fama, E.F., French, K.R., 2006. Profitability, investments and average returns. Journal of Financial Economics 82, 491–518.

Fama, E.F., French, K.R., 2018. Choosing factors. Journal of Financial Economics 128, 234–252.

Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. Journal of Finance 56, 329–352.

Fleming, J., Kirby, C., Ostdiek, B., 2003. The economic value of volatility timing using "realized" volatility. Journal of Financial Economics 67, 473–509.

Frazzini, A., Pedersen, L.H., 2014. Betting against beta. Journal of Financial Economics 111, 1–25.

Friend, I., Blume, M., 1970. Measurement of portfolio performance under uncertainty. American Economic Review 60, 561–575.

Han, X., 2022. Understanding the performance of components in betting against beta. Critical Finance Review 11, 1–36.

Han, X., Roussanov, N., Ruan, H., 2022. Mutual fund risk shifting and risk anomalies.

Hirshleifer, D., Hou, K., Teoh, S.H., Zhang, Y., 2004. Do investors overvalue firms with bloated balance sheets? Journal of Accounting and Economics 38, 297–331.

Hong, H., Sraer, D.A., 2016. Speculative betas. Journal of Finance 71, 2095–2144.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance 48, 65–91.

Jobson, J.D., Korkie, B.M., 1981. Performance hypothesis testing with the Sharpe and Treynor measures. Journal of Finance , 889–908.

Jylhä, P., 2018. Margin requirements and the security market line. Journal of Finance 73, 1281–1321.

Kirby, C., Ostdiek, B., 2012. It's all in the timing: Simple active portfolio strategies that outperform naive diversification. Journal of Financial and Quantitative Analysis 47, 437–467.

Koijen, R.S.J., Yogo, M., 2019. A demand system approach to asset pricing. Journal of Political Economy 127, 1475–1515.

Kozak, S., Nagel, S., Santosh, S., 2018. Interpreting factor models. Journal of Finance 73, 1183–1223.

Kroencke, T.A., 2017. Asset pricing without garbage. The Journal of Finance 72, 47–98.

Kumar, A., 2009. Who gambles in the stock market? Journal of Finance 64, 1889–1933.

Levy, R.A., 1967. Relative strength as a criterion for investment selection. Journal of Finance 22, 595–610.

Li, D., Zhang, L., 2010. Does q-theory with investment frictions explain anomalies in the cross section of returns? Journal of Financial Economics 98, 297–314.

Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. Review of Economics and Statistics , 13–37.

Liu, F., Tang, X., Zhou, G., 2019. Volatility-managed portfolio: Does it really work? Journal of Portfolio Management 46, 38–51.

Liu, J., Stambaugh, R.F., Yuan, Y., 2018. Absolving beta of volatility's effects. Journal of Financial Economics 128, 1–15.

Loughran, T., Ritter, J.R., 1995. The new issues puzzle. Journal of Finance 50, 23–51.

Lu, Z., Qin, Z., 2021. Leveraged funds and the shadow cost of leverage constraints. Journal of Finance 76, 1295–1338.

Lyandres, E., Sun, L., Zhang, L., 2007. The new issues puzzle: Testing the investment-based explanation. Review of Financial Studies 21, 2825–2855.

MacKinlay, A., 1995. Multifactor models do not explain deviations from the capm. Journal of Financial Economics 38, 3–28.

Marquering, W., Verbeek, M., 2004. The economic value of predicting stock index returns and volatility. Journal of Financial and Quantitative Analysis 39, 407–429.

Mashruwala, C., Rajgopal, S., Shevlin, T., 2006. Why is the accrual anomaly not arbitraged away? The role of idiosyncratic risk and transaction costs. Journal of Accounting and Economics 42, 3–33.

McLean, R.D., 2010. Idiosyncratic risk, long-term reversal, and momentum. Journal of Financial and Quantitative Analysis 45, 883–906.

Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. Journal of Monetary Economics 15, 145–161.

Merton, R.C., 1973. An intertemporal capital asset pricing model. Econometrica 41, 867–887.

Moreira, A., Muir, T., 2017. Volatility-managed portfolios. Journal of Finance 72, 1611–1644.

Nagel, S., Singleton, K.J., 2011. Estimation and evaluation of conditional asset pricing models. Journal of Finance 66, 873–909.

Novy-Marx, R., 2013. The other side of value: The gross profitability premium. Journal of Financial Economics 108, 1–28.

Novy-Marx, R., Velikov, M., 2022. Betting against betting against beta. Journal of Financial Economics 143, 80–106.

Ohlson, J.A., 1980. Financial ratios and the probabilistic prediction of bankruptcy. Journal of Accounting Research 18, 109–131.

Pontiff, J., 2006. Costly arbitrage and the myth of idiosyncratic risk. Journal of Accounting and Economics 42, 35–52.

Ross, S.A., 1976. The arbitrage theory of capital asset pricing. Journal of Economic Theory 13, 341–360.

Shanken, J., 1990. Intertemporal asset pricing: An empirical investigation. Journal of Econometrics 45, 99–120.

Shanken, J., 1992. The current state of arbitrage pricing theory. Journal of Finance 47, 1569–1574.

Sharpe, W.F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance 19, 425–442.

Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.

Sloan, R.G., 1996. Do stock prices fully reflect information in accruals and cash flows about future earnings? Accounting Review 71, 289–315.

Stambaugh, R.F., Yu, J., Yuan, Y., 2012. The short of it: Investor sentiment and anomalies. Journal of Financial Economics 104, 288–302.

Stambaugh, R.F., Yu, J., Yuan, Y., 2015. Arbitrage asymmetry and the idiosyncratic volatility puzzle. Journal of Finance 70, 1903–1948.

Stambaugh, R.F., Yuan, Y., 2017. Mispricing factors. Review of Financial Studies 30, 1270–1315.

Titman, S., Wei, K.C.J., Xie, F., 2004a. Capital investments and stock returns. The Journal of Financial and Quantitative Analysis 39, 677–700.

Titman, S., Wei, K.J., Xie, F., 2004b. Capital investments and stock returns. Journal of Financial and Quantitative Analysis 39, 677–700.

Vasicek, O.A., 1973. A note on using cross-sectional information in bayesian estimation of security betas. Journal of Finance 28, 1233–1239.

Welch, I., 2022. Simply better market betas. Critical Finance Review 11, 37–64.

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.

Xing, Y., 2007. Interpreting the value effect through the q-theory: An empirical investigation. Review of Financial Studies 21, 1767–1795.

Yu, J., Yuan, Y., 2011. Investor sentiment and the mean–variance relation. Journal of Financial Economics 100, 367–381.