

What Drives Booms and Busts in Value?

John Y. Campbell, Stefano Giglio, and Christopher Polk¹

First draft: March 2023

This version: May 2025

Abstract

Value investing delivers volatile returns, with large drawdowns during both market booms and busts. This paper interprets these returns through an intertemporal CAPM, which predicts that aggregate cash flow, discount rate, and volatility news all move value returns. We document that indeed these shocks explain a large fraction of quarterly value returns over the last 60 years. We also distinguish between the intra-industry and inter-industry components of value, showing that the ICAPM explains the former better. Finally, we develop a novel methodology to perform this decomposition at the daily frequency, using it to interpret value returns during the Covid-19 pandemic.

¹Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, and NBER. Email john_campbell@harvard.edu. Phone 617-496-6448. Giglio: Yale School of Management, 165 Whitney Ave, New Haven, CT 06520, CEPR, and NBER. Email stefano.giglio@yale.edu. Polk: Department of Finance, London School of Economics, London WC2A 2AE, UK, and CEPR. Email c.polk@lse.ac.uk. We thank Hongyu Wu for excellent research assistance.

1 Introduction

Value investors seek to buy stocks that are cheap in the sense that their market prices are low relative to accounting measures of their value. This investment approach has a long history predating modern finance theory (see, for example, Graham, Dodd, and Cottle 1934). A value strategy can be implemented using various alternative valuation ratios such as the dividend yield, earnings-price ratio, or book-market ratio, but since the influential work of Fama and French (1992, 1993) academic research on value has primarily emphasized the book-market ratio.²

If a value strategy uses an accounting measure of a stock’s value that captures its rationally expected future cash flows, discounted at a constant rate, then by the logic of Campbell and Shiller (1988a), the corresponding valuation ratio should be a good proxy for the rationally expected long-run return on the stock. This insight implies that a value portfolio should contain stocks with high average returns, either because they are risky stocks with high equilibrium returns, or because they are stocks that have been underpriced by irrational investors.

Whether value returns are high on average or not, they are very far from riskless. Figure 1 illustrates this by plotting the return to the standard Fama-French long-short value strategy, HML, over the period 1963Q1–2022Q4. The raw quarterly log return to HML is smoothed for readability using an exponentially weighted moving average with a half-life of two years. The solid black line in the figure shows large swings in HML returns, with value booms in the late 1970s, mid 1980s, and mid 2000s, and short but sharp value busts in the early 1970s, the early 1980s, the early 1990s, the turn of the millennium, the global financial crisis of 2008-09, and—most striking of all—the Covid-19 pandemic of 2020-21.

In this paper we ask what factors drive such booms and busts in value, and more generally the volatility of value returns. We argue that three factors relevant to value investors, derived from data on the aggregate stock market using an intertemporal CAPM (Merton 1973, ICAPM), go a long way to explain the volatile history of value returns, both to HML

²Early academic papers on value investing using the dividend yield as a valuation ratio include Litzenberger and Ramaswamy (1982) and Miller and Scholes (1982). Basu (1977, 1983) uses the earnings-price ratio, and Stattman (1980) and Rosenberg, Reid, and Lanstein (1985) use the book-market ratio. Goncalves and Leonard (2023) propose an alternative way to estimate the fundamental value of a company that may obviate a potential deterioration of book equity as a measure of fundamental value. See also earlier work by Frankel and Lee (1998) as well as Cong, George, and Wang (2023).

itself and particularly to its intra-industry component.

We make two preliminary observations. First, value booms and busts do not line up with returns on the aggregate stock market. Figure 1 illustrates this fact by plotting the log real return on the CRSP value-weighted index, smoothed in the same way as the HML return, as a dashed red line. The earlier value busts coincide with high aggregate stock returns, but value busts in the global financial crisis and the Covid-19 pandemic coincide with sharp declines in the aggregate market. Overall, there is little correlation between the two lines in Figure 1, consistent with the fact that HML has a negative beta with the aggregate market that is close to zero in absolute value.

Second, value booms and busts primarily result from returns to value within industries. We create an intra-industry value strategy that holds stocks whose book-market ratios are high relative to other stocks in the same industry.³ We regress HML onto the return of this strategy and call the fitted value intra-industry HML, or HML_{Intra} , and the residual inter-industry HML, or HML_{Inter} . With this procedure, HML is the sum of the two components, and inter-industry HML picks up industry shocks that are orthogonal to intra-industry value returns. Figure 2 plots the log returns to HML and its two components over the same sample period and using the same smoothing procedure as Figure 1. The importance of intra-industry HML, the dashed green line in the figure, is immediately visible. However, many value busts do coincide with poor returns to inter-industry HML, the dash-dotted red line, despite the zero unconditional correlation (by construction) between intra-industry and inter-industry HML.

If booms and busts in value are not explained by aggregate stock returns, are there other aggregate shocks that do explain them? We argue that the answer lies in shocks that are relevant to value investors according to the ICAPM. Specifically, we follow Campbell and Vuolteenaho (CV 2004) in distinguishing market cash-flow news (upward revisions in rationally expected market dividends, discounted to the infinite future at a constant rate)

³We are not the first to decompose value into intra- and inter-industry components. Cohen, Polk, and Vuolteenaho (2003) provide a variance decomposition of firm-level BE/ME and find “that the book-to-market effect in returns is mostly an intra-industry effect.” Other papers include Cohen and Polk (1998) (the first to study intra- and inter-industry value strategies), Lewellen (1999), and Asness, Porter, and Stevens (2000). Daniel, Grinblatt, Titman, and Wermers (1997) use Cohen and Polk’s (1998) adjustment in their popular characteristic-based approach to performance evaluation. Industry adjustment is widely-used in practice; see, for example, Israel, Jiang, and Ross (2018), who argue that careful industry adjustment of value characteristics is an example of “craftsmanship alpha” that skilled fund managers can provide.

and negative discount-rate news (downward revisions in rationally expected market returns, again discounted to the infinite future at a constant rate). The total return on the market equals cash-flow news plus negative discount-rate news, and so a stock’s covariance with the market can be expressed as its covariance with cash-flow news plus its covariance with negative discount-rate news. Dividing these covariances by the variance of the return on the market, a stock’s beta with the market can be written as the sum of its cash-flow and discount-rate betas.

In a homoskedastic setting, CV show that conservative long-term investors charge a different risk price for exposure to these two types of aggregate shocks, with cash-flow beta (“bad beta”) commanding a higher price than discount-rate beta (“good beta”). CV further show that value stocks have higher cash-flow betas but lower discount-rate betas and overall market betas than growth stocks. This pattern may reflect the lower duration of value stocks, whose cash flows are concentrated in the near term, relative to growth stocks.⁴ According to the CV model, value stocks are likely to underperform in a period when the aggregate market is driven down by bad cash flow news, but also in a period when the market is driven up by declining discount rates. There need be no strong unconditional relationship between value returns and aggregate stock returns since both types of shocks affect the aggregate market.

Campbell, Giglio, Polk, and Turley (CGPT 2018) extend the CV model to include shocks to the variance of stock returns. Using a model in which the market return and all predictor variables share a common process for volatility, CGPT show that long-run market volatility shocks (upward revisions in the rationally expected value of future market return variances, discounted to the infinite future at a constant rate) are relevant to long-term investors because they reduce the quality of investment opportunities. An asset that performs well when a volatility shock occurs hedges against this deterioration in investment opportunities. Conservative long-term investors charge a negative risk price for exposure to volatility shocks that depends on the volatility process and risk aversion, but not on any other preference parameters. CGPT find that value stocks have lower volatility betas than growth stocks, perhaps reflecting the optionality of growth-company investment opportunities that increases the prices of growth stocks in an uncertain environment. According to the CGPT

⁴A recent literature stresses this duration effect. See Dechow, Sloan, and Soliman (2004), Lettau and Wachter (2007), Weber (2018), and Gormsen and Lazarus (2021).

model, a period with unusual upward shocks to volatility is likely to be a period in which value underperforms.

In this paper, we estimate the VAR model specified by CGPT, using quarterly data from 1926Q3 through 2022Q4. The model generates quarterly news shocks similar to those estimated in CV and CGPT.⁵ In the period since 1963Q3 where the return to value is anomalously high relative to its market beta (the “modern period” studied by CV and CGPT), we estimate the betas of HML and its intra-industry and inter-industry components with the three shocks to cash flows, discount rates, and volatility. We find similar betas for HML to those reported in prior research. A novel result is that intra-industry HML has a considerably higher cash-flow beta and considerably lower volatility beta than inter-industry HML. Since cash-flow and volatility betas are highly compensated in the ICAPM, with a positive risk price for cash-flow beta and a negative risk price for volatility beta, this pattern helps to explain the higher average return to intra-industry HML than to inter-industry HML.

Over our full sample period since 1963Q3, the three ICAPM shocks explain almost half (48%) of the variance of HML, 37% of the variance of intra-industry HML, but only 14% of the variance of inter-industry HML. In a more recent sample period starting in 1990 (a period in which we can also perform the analysis at a higher frequency, as described below), the explanatory power of the ICAPM is almost unchanged at 47% for HML, 37% for intra-industry HML, and 13% for inter-industry HML. These statistics suggest that the pattern of ICAPM shocks has the potential to explain recent value returns and particularly the returns to intra-industry HML.

Our ICAPM shocks provide more than just a statistical fit to HML’s fluctuations: they also help us understand the main economic drivers of HML returns. Our analysis identifies what news was most important at different times, highlighting that different periods see different mixes of the news shocks moving HML. There are periods when different news shocks have offsetting effects on HML (for example the 1970s, with cash-flow news offsetting

⁵Our approach uses the CGPT unrestricted VAR to measure the three shocks driving value returns. In contrast, the structured approach of Campbell, Giglio, and Polk (2013) imposes the cross-sectional restrictions of the two-beta CV ICAPM when estimating the time-series VAR to sharpen equity premium forecasts and the resulting decomposition of unexpected market returns into discount-rate and cash-flow news. Given our focus on the time-series decomposition of returns (as opposed to the cross-sectional fit of the model), in this paper we estimate the VAR from the time-series information alone, as in CGPT.

discount rate news and variance news), and other periods when all news terms move HML in the same direction (for example the technology boom in the late 1990s).

Whereas the majority of the literature studying value returns (including CGPT) has focused on *cross-sectional* patterns, the focus of this paper is understanding the economic forces driving its *time-series* fluctuations. The importance of this question – specifically as it concerns the fluctuations of value investing – was recently emphasized by Rob Vishny at the “Chicago School of Finance at 125” conference in October 2023: “What I find lacking is the ability [...] to explain big shifts in investor sentiment or valuation dispersion after the fact. Forget about predicting, we can’t predict...But after the fact, can’t we do better explaining it?” Of course, Vishny’s challenge to the field is also closely related to the R^2 question posed by Roll (1988). The time series patterns we document matter not only because we would like to understand the proximate causes of different episodes: They also represent additional moments that can be used to guide the calibration of theoretical models that aim to explain the behavior of value and growth stocks.⁶

The Covid-19 pandemic presents a particularly interesting case study. During the pandemic, all three types of news arrived at a rapid pace and market returns were unusually volatile. These facts meant that major shocks occurred at a higher frequency than quarterly. To understand the returns to value in 2020, we adapt the CGPT framework and estimate a daily VAR. We start our sample period at the beginning of 1990 and end at the end of 2022Q4. Focusing on this sub-period allows us to include the VIX as an additional state variable relevant for predicting future volatility. We show that the daily version of the ICAPM gives a good account of HML and intra-industry HML returns throughout this sub-period and specifically in 2020. The high-frequency variation of returns in 2020 is primarily driven by cash-flow news, because volatility and discount-rate news had largely offsetting effects on HML and intra-industry HML.

While the ICAPM helps us to understand the 2020 returns to intra-industry HML, it has much lower explanatory power for inter-industry HML. To gain some insight into value returns across industries, we measure earnings growth shocks to value industries relative to growth industries. We show that the Covid-19 pandemic particularly hurt the earnings of

⁶For example, economic explanations for the value effect include sentiment (e.g., Baker and Wurgler 2006), production-based asset pricing models (e.g., the technology shocks of Kogan et al. 2013), and justifications linked to equity duration (e.g., Lettau and Wachter 2007). Our findings present a new set of time-series facts linked to aggregate shocks that these different types of models of the value premium should explain.

value industries, which tend to be physical businesses, relative to growth industries. This sectoral shock, rather than an aggregate shock, particularly affected inter-industry HML during 2020.

The main focus of our paper is the time-series variability of value returns, rather than gradual shifts in the average return to value. However, many commentators have noted that the average return to HML has been lower in recent decades, and particularly since the global financial crisis in 2008, than it was in the earlier part of our sample period. This fact is clearly visible in Figure 1, and Figure 2 shows that it primarily reflects a lower average return to intra-industry value. It is not clear whether this phenomenon reflects a structural shift or simply a sequence of unfavorable shocks to value stocks.⁷ To shed some light on this issue, we study whether gradual changes in the loadings of HML on ICAPM news shocks might have caused a repricing of value stocks that could help explain the low recent returns to value. We find that indeed, the data are consistent with this explanation.

The organization of the paper is as follows. Section 2 sets up a quarterly VAR model and explains how we estimate the three aggregate ICAPM shocks to market cash flows, market discount rates, and market variance. Section 3 estimates the betas of HML and its two components with these three shocks, and studies the time-series evolution of the realized quarterly returns to HML, intra-industry HML, and inter-industry HML that can be explained by the ICAPM.

Section 4 extends this quarterly analysis in several ways. It explores the ability of the ICAPM to explain the returns to other factor and anomaly portfolios, documents the robustness of the results to other ways to decompose HML into intra-industry and inter-industry components, explains the link between the explanatory power of the ICAPM for HML and the ICAPM-implied Sharpe ratio of HML, and discusses the role of the value spread – one of the explanatory variables CGPT include in the VAR – in driving our results.

Section 5 sets up a daily model and reestimates the ICAPM shocks at a daily frequency. We then use this daily ICAPM to understand the returns to HML and its two components during the Covid-19 pandemic. Section 6 studies the cross-sectional pattern of industry earnings shocks in 2020 and argues that this pattern contributed to the poor performance

⁷For example, Fama and French (2020) are unable to reject the hypothesis that the value premium is the same in both halves of the July 1963 - June 2019 sample. However, if they forecast the value premium using Cohen, Polk, and Vuolteenaho's (2003) value spread, they do find some evidence of a second-half decline in the expected value premium.

of inter-industry HML at that time.

Section 7 uses rolling regressions to track gradual shifts over time in the loadings of HML and its components on ICAPM news shocks. We present some suggestive evidence that these loadings have changed in a way that increases the expected return to HML going forward, but may have lowered realized HML returns as value stocks were repriced to reflect their altered risks.

Section 8 concludes. An online appendix (Campbell, Giglio, and Polk 2025) provides additional supporting evidence.

2 A Quarterly VAR Model for Aggregate Stock Returns and Volatility

2.1 News terms and the dynamics of the economy

In this paper, we build on CGPT (2018) to describe the dynamics of the economy. Like them, we assume that the economy is well described by the following first-order vector autoregression (VAR):

$$\mathbf{x}_{t+1} = \bar{\mathbf{x}} + \mathbf{\Gamma} (\mathbf{x}_t - \bar{\mathbf{x}}) + \sigma_t \mathbf{u}_{t+1}, \quad (1)$$

where $\mathbf{x}_t$ is an $n \times 1$ vector of state variables which includes the market return r_t as its first element, $\bar{\mathbf{x}}$ and $\mathbf{\Gamma}$ are an $n \times 1$ vector and an $n \times n$ matrix of constant parameters, and $\mathbf{u}_{t+1}$ is a vector of shocks to the state variables normalized so that its first element has unit variance. As in CGPT, we assume that $\mathbf{u}_{t+1}$ has a constant variance-covariance matrix $\mathbf{\Sigma}$, with element $\Sigma_{11} = 1$, and include σ_t^2 , the variance of returns, in the state vector as the second element. We define $n \times 1$ vectors $\mathbf{e}_1$ and $\mathbf{e}_2$, all of whose elements are zero except for a unit first element in $\mathbf{e}_1$ and a unit second element in $\mathbf{e}_2$. Finally, whereas $\mathbf{u}_{t+1}$ refers to the scaled VAR innovations (i.e., normalized by σ_t), we refer to η_{t+1} as the raw, or unscaled, VAR innovations: $\eta_{t+1} = \sigma_t \mathbf{u}_{t+1}$.

CGPT derive the first-order conditions of a long-term investor with Epstein-Zin utility that is holding the market in equilibrium. The paper shows that equilibrium pricing from these first order conditions (i.e., the stochastic discount factor associated with this investor's Euler equation) involves three priced shocks: to (aggregate) cash flow news, discount rates,

and variance. These news terms can be expressed as a function of the VAR parameters:

$$N_{DR,t+1} = \underbrace{\mathbf{e}'_1 \rho \mathbf{\Gamma} (\mathbf{I} - \rho \mathbf{\Gamma})^{-1}}_{\lambda_{DR}} \eta_{t+1}, \quad (2)$$

$$N_{CF,t+1} = \underbrace{(\mathbf{e}'_1 + \mathbf{e}'_1 \rho \mathbf{\Gamma} (\mathbf{I} - \rho \mathbf{\Gamma})^{-1})}_{\lambda_{CF}} \eta_{t+1}, \quad (3)$$

$$N_{V,t+1} = \underbrace{\rho \mathbf{e}'_2 (\mathbf{I} - \rho \mathbf{\Gamma})^{-1}}_{\lambda_V} \eta_{t+1}, \quad (4)$$

which we refer to as discount rate news, cash flow news, and variance news respectively. The parameter ρ is the loglinearization parameter of the Campbell-Shiller approximation, and should be thought of as close to one (we follow CGPT and use 0.95 in our empirical work). Each news term loads linearly on the innovations η_{t+1} , and we refer to the loadings as λ_{DR} , λ_{CF} , and λ_V respectively.

In CGPT's model, the risk premium that an asset yields depends on its exposure to these three news terms.⁸ While that paper focuses on the cross-sectional implications of the model, in this paper we investigate the *time-series* relation between the returns of the value (HML) portfolio and the realizations of these news shocks. That is, we investigate to what extent the fluctuations in HML can be explained by its exposure to these three aggregate shocks.

We begin in the next section by estimating the VAR parameters and studying the time series of the news terms. After that, we relate these estimated news terms to the returns of HML.

2.2 VAR specification

To implement the VAR estimation, we need to choose the variables that are included in the state vector. We use the same six state variables as in CGPT (2018), with one exception: we replace their interest rate variable with the term yield spread.⁹ The data are quarterly, from

⁸In their model, the stochastic discount factor is: $m_{t+1} - \mathbb{E}_t m_{t+1} = -\gamma N_{CF,t+1} - [-N_{DR,t+1}] + \frac{1}{2} \omega N_{V,t+1}$, where γ represents risk aversion and ω is a parameter that is a deterministic function of γ .

⁹The published version of CGPT uses the 3-month Treasury bill yield, but we follow CV, Campbell, Giglio, and Polk (2013), and the NBER working paper version of CGPT and use the term yield spread instead. CGPT replaced the term yield spread with the 3-month Treasury bill yield following the suggestion

1926Q3 to 2022Q4. Given that the VAR is used to extract news about cash flows, discount rates, and volatility for long horizons, and contains very persistent variables, using this long sample is important to obtain precise estimates of the dynamics. In fact, CGPT show that this VAR specification is a good forecaster of stock return variance at very long horizons.

The first variable in the VAR is the log real return on the market, r_M , the difference between the log return on the Center for Research in Securities Prices (CRSP) value-weighted stock index and the log return on the Consumer Price Index. This portfolio is a standard proxy for the aggregate wealth portfolio, but CGPT confirm the robustness of their findings to alternative proxies that delever the market return by combining it in various proportions with Treasury bills. The second variable is expected market variance ($EVAR$). This variable is meant to capture the variance of market returns, σ_t^2 , conditional on information available at time t , so that innovations to this variable can be mapped to the N_V term described above.

To construct $EVAR_t$, we proceed as follows. We first construct a series of within-quarter *realized* variance of daily returns for each time t , $RVAR_t$. We then run a regression of $RVAR_{t+1}$ on lagged realized variance ($RVAR_t$) as well as the other five state variables at time t . This regression then generates a series of predicted values for $RVAR$ at each time $t + 1$, that depend on information available at time t : $\widehat{RVAR}_{t+1}$. Finally, we define our expected variance at time t to be exactly this predicted value at $t + 1$:

$$EVAR_t \equiv \widehat{RVAR}_{t+1}. \quad (5)$$

Note that though we describe our methodology in a two-step fashion where we first estimate $EVAR$ and then use $EVAR$ in a VAR, this choice is only for interpretability. Indeed, this approach to modeling $EVAR$ can be considered a simple renormalization of equivalent results we would find from a VAR that included $RVAR$ directly.¹⁰

of a reviewer in the publication process who argued that pricing the returns on test assets in excess of that rate, as in CGPT's NBER working paper, required that the T-bill rate was included in the VAR. CGPT ultimately derived a more general GMM pricing equation that allowed for any reference asset and used the market return as a benchmark, a natural choice given their theory. Though their switch to these more general pricing tests eliminated the reviewer's original concern, for transparency's sake, CGPT did not revert back to using the term yield spread in their VAR. Nevertheless, all of our conclusions (as well as those of CGPT) are robust to using either the 3-month bill yield or the term yield spread in the VAR.

¹⁰As discussed in CGPT, observations are weighted based on $RVAR$ in the first stage and then reweighted using $EVAR$ in the second stage, so our two-stage approach in practice is not exactly the same as a one-stage

The third variable is the log of the S&P 500 price-smoothed earnings ratio (PE) adapted from Campbell and Shiller (1988b), where earnings are smoothed over ten years.¹¹ The fourth is the term yield spread ($TERM$) measured as the difference between the log yield on 10-year Treasuries and the log yield on 3-month Treasuries, obtained from the Federal Reserve Bank of St. Louis. The fifth is the small-stock value spread (VS); this variable is the difference between the log average book-to-market ratio of the small value portfolio and the log book-to-market ratio of the small growth portfolio, where the two portfolios are from the six elementary portfolios available on Ken French’s website. The details of the construction are reported in Campbell, Giglio, and Polk (2013).

The sixth and final variable is the default spread (DEF), defined as the difference between the log yield on Moody’s BAA and AAA bonds, obtained from the Federal Reserve Bank of St. Louis. We include the default spread in part because that variable is known to track time-series variation in expected real returns on the market portfolio (Fama and French, 1989), but also because shocks to the default spread should to some degree reflect news about aggregate default probabilities, which in turn should reflect news about the market’s future cash flows and volatility.¹²

We follow CGPT in estimating our VAR, and refer the reader to that paper for details. Appendix Table 1 reports the estimates of the regression of $RVAR$ onto the lagged variables (used to build EVAR) and the VAR parameter estimates. All the estimates are similar to those in CGPT.

2.3 Estimated news terms

Table 1 presents the standard deviation/correlation matrix of the news terms, estimated by combining the estimated VAR parameters and innovations. The standard deviations are broadly comparable with those reported by CGPT, and the correlations are qualitatively similar. Cash-flow news is positively correlated with negative discount-rate news, implying

approach. CGPT explore the robustness of their results to several different ways of estimating their VAR, including using a $RVAR$ -weighted, single-step estimation approach.

¹¹The PE ratio reported on Robert Shiller’s website compares a monthly average of daily closing prices to smoothed earnings. To facilitate our use of these data at the daily frequency, we replace Shiller’s monthly average price with the relevant closing price at month-end.

¹²Again, to facilitate our use of these data at the daily frequency, for both the term spread and the default spread, we use the version of the series in question that reflects daily closing yields rather than average yields.

that positive cash-flow shocks are associated with declines in discount rates that amplify the response of stock prices. Cash-flow news is negatively correlated with variance news, implying that positive cash-flow shocks are associated with reductions in equity market risk. Quantitatively, the correlations we estimate are somewhat larger in absolute value than those reported by CGPT.

Figure 3 plots the three news terms, smoothed for readability using an exponentially-weighted moving average (EWMA) with a half-life of two years. While the news terms are produced for the entire sample starting in 1926 (the sample we use to estimate the VAR), we focus here on the “modern sample” starting in 1963, during which, as shown by Campbell and Vuolteenaho (2004) and CGPT, the cross-section of risk premia is significantly better described by the ICAPM than by the CAPM.

The vertical line in the graph marks the year 1990. As discussed in the introduction, in this paper we pay special attention to the post-1990 period, a period for which we observe additional variables at the daily frequency, that allow us to increase the frequency of the estimation of the news terms.

The patterns in the figure are consistent with previous research. Most of the variation in aggregate returns comes from discount-rate news as shown in Campbell (1991). Campbell and Vuolteenaho (2004) emphasize how patterns in cash-flow news are distinct from those in discount-rate news, which the figure broadly confirms. Cash-flow news is positive in the 1960s and then is mostly negative during the 1970s and 1980s. The 1990s bring positive cash-flow news which continues throughout the early 2000s until the Global Financial Crisis, a market downturn driven by the arrival of negative cash-flow news. This large negative shock is consistent with Campbell, Giglio, and Polk (2013) who document that in the history of booms and busts in the U.S., the 2007-2009 downturn was one of a few that were driven mainly by cash-flow news. Consistent with the results in CGPT, volatility news spikes during the tech boom and the Global Financial Crisis.

Relative to this previous work, our updated sample includes the Covid-19 episode of 2020, marked by a combination of negative cash flow news, positive volatility news, and positive discount rate news. The figure also shows that the news arriving in early 2020 was quickly reversed over the following quarters. The fast reversal suggests that higher frequency data (which we analyze later in the paper) will be useful to obtain a more complete understanding of the Covid-19 episode.

3 A Quarterly ICAPM Perspective on HML and its Components

In this section we study the time series of returns of HML and explore to what extent they can be explained by exposure to the three aggregate news shocks. To facilitate our interpretation of the exposure of HML to these three news shocks, we decompose Fama and French’s (1993) value strategy into an intra-industry and an inter-industry component.

3.1 Decomposing HML

The construction of HML, following Fama and French (1993), involves double-sorting stocks by their market capitalization and their book-to-market ratio. Because different industries include firms of different sizes with different book-to-market ratios, HML generally overweights certain industries relative to others, that is, it is not industry-neutral. For example, banks tend to be value companies, and thus tend to be overweighted in HML. To better understand the drivers of HML returns, we decompose HML into an intra-industry component that has no industry characteristic exposure and an inter-industry component that captures the effects of industry tilts on HML. We refer to the first as HML_{Intra} , and to the second as HML_{Inter} . The focus on this decomposition is motivated by previous evidence (e.g., Cohen, Polk and Vuolteenaho 2003) that has shown that value is in large part an intra-industry phenomenon. We therefore explore the possibility that the ICAPM news terms might explain the returns to the two components differently.

There are multiple ways to perform this decomposition, some of which have been explored in previous literature. The approach we use builds an HML portfolio separately for each industry (following the same procedure as Fama and French 1993, and using the 48 industries of Fama and French 1997), and then aggregates the returns of the industry-specific HML portfolios by weighting the industries by their market capitalization. This procedure yields a long-short intra-industry HML portfolio which, by construction, has no industry characteristic exposure. We next construct the inter-industry component as the residual of a regression of HML onto the intra-industry component. We define HML_{Inter} as the residual of this regression. Finally, we construct HML_{Intra} as the rescaled intra-industry component, where the scaling is such that a regression of HML onto HML_{Intra} has a coefficient of exactly

1. Thanks to this rescaling, we can then write:

$$HML = HML_{Intra} + HML_{Inter}, \quad (6)$$

where the two components are orthogonal.¹³

We begin by describing some of the main features of HML, HML_{Intra} , and HML_{Inter} in the modern sample (1963Q3-2022Q4). The risk premia of the three portfolios are 0.99%, 1.43%, and -0.44% per quarter, respectively. The HML risk premium is positive (and statistically significant with a t -statistic of 2.50) over the modern sample, despite the lower premium earned by value over the last 15 years. The HML premium is more than fully attributed to the intra-industry component (t -statistic of 4.16), as the inter-industry risk premium is estimated to be significantly negative (t -statistic of -2.26).

Table 2 shows how HML, HML_{Intra} , and HML_{Inter} (built using our benchmark specification) relate to each other. The table has three panels, with the first one focusing on the entire modern sample (1963Q3-2022Q4), and the other two panels on the pre-1990Q2 and post-1990Q1 samples, respectively.¹⁴

Starting from the modern sample (Panel A), the first row reports the results of a regression of HML on the intra-industry component. Given the way intra-industry HML is scaled, this regression has by construction a coefficient of 1. The R^2 of the regression is high at 76%, indicating that time-series variation in the intra-industry component of HML is closely related to variation in HML. The second row reports the regression of HML_{Intra} onto HML (which mechanically has the same R^2). Both regressions show a statistically significant alpha (negative in the first regression and positive in the second regression), which indicates that the intra-industry component of HML commands a risk premium that goes beyond that of HML.

The third row shows that HML_{Inter} has a large negative alpha with respect to HML and a much lower correlation with HML than does HML_{Intra} . Together, these results suggest that the intra-industry component is highly priced by investors, whereas the inter-industry

¹³A similar intra-industry HML portfolio is available on Robert Novy-Marx's website at <https://rnm.simon.rochester.edu/>. We only consider industries for which we have at least one stock in each of the four portfolios used to build HML (large growth, small growth, large value, small value), so that an HML portfolio can be built within that industry.

¹⁴We choose this specific split because we have a complete quarter of daily news only starting in 1990Q2.

component, i.e., the component of HML that is uncorrelated with HML_{Intra} , carries a much smaller risk premium.

The last three rows of the panel study the market exposure of the three portfolios. Row (4) shows that HML is weakly negatively related to the market – its beta is -0.17, and the R^2 of the projection onto the market is only 6%. This regression indicates that exposure to the market alone will not be sufficient to understand HML’s fluctuations. The same is true for HML_{Intra} and HML_{Inter} ; we will need to go beyond the market to understand variation in their returns. Finally, the table also shows that while the intra-industry portfolio is responsible for all of HML’s CAPM alpha, the inter-industry portfolio is responsible for most of HML’s market beta.

Panels B and C of the table repeat the analysis splitting the sample in 1990. The results highlighted above continue to hold in the two samples, with few exceptions. HML_{Intra} explains all of the alpha *and* slightly more of the market beta of HML before 1990. In addition, the overall alpha of HML is, as well known, much lower in the recent period. This decrease in HML alpha is mirrored by a decrease in the alpha of HML_{Intra} .

As we have already discussed, Figure 2 shows the time series of (EWMA-smoothed) log returns on HML and its two components in the modern period. Several of the patterns of Table 2 are visible in the figure. The fluctuations in HML closely match those of HML_{Intra} , during booms and busts. HML_{Inter} instead seems to be less correlated with HML and HML_{Intra} , though it does tend to experience crashes at the same time (e.g., during the Covid-19 pandemic). In the remainder of this paper, we seek to explain the booms and busts in HML and its two components using the three ICAPM shocks that we have estimated.

3.2 HML and the ICAPM news terms

We next study how HML loads on the aggregate news terms estimated from the VAR. We also explore how the two components, HML_{Intra} and HML_{Inter} , load on them: to the extent that the two components load differentially on these news terms, this analysis can provide an interesting new test of CGPT’s neoclassical explanation for the value effect and can also facilitate a deeper understanding of the sources of HML’s fluctuations.

Table 3 shows the three portfolios’ market betas, as well as their betas with respect to

the ICAPM news terms plotted in Figure 2.¹⁵ The table reports univariate betas on the left side, and multivariate betas on the right side.¹⁶ We begin with the top panel, that uses the entire modern sample. As highlighted in Table 2, we find that HML has a negative market beta, confirming the findings of CV and CGPT. HML’s negative CAPM beta comes from the strategy’s negative discount-rate beta, which masks a positive cash-flow beta. In addition, HML has a negative variance beta. These findings continue to hold in a multiple regression context.¹⁷

The last column of the table reports the R^2 of the multiple regression. The ICAPM news terms explain 48%, nearly half, of the time-series variation in HML, a striking increase from the mere 6% explained by the market factor (see Table 2). This result is important because it tells us that the ICAPM news terms are not only useful in explaining the cross-section of risk premia (as demonstrated by CGPT), they are also helpful in explaining period-by-period fluctuations in value returns, and therefore they can be useful in interpreting the (proximate) sources of HML’s booms and busts (that is, its volatility).

The next two rows of the table focus on the intra- and inter-industry components of HML. They reveal several novel facts about value and the intertemporal CAPM. First, HML_{Inter} is responsible for a significant proportion of HML’s discount-rate beta (and, as noted above, of HML’s market beta). However, almost all of HML’s cash-flow beta comes from the HML_{Intra} component, which is also responsible for almost all of the variance beta of HML. HML_{Inter} , instead, has essentially a zero cash-flow and an economically small variance beta. These facts are broadly consistent with the differences in average returns on these two components of HML. The component of value contributing to HML’s CAPM alpha is also the component

¹⁵Note that Table 2 estimates market betas using simple excess returns on the market while Table 3 estimates betas using the unexpected component of log real market returns. Thus, the estimates should not be exactly equal to each other.

¹⁶When reporting the results of the simple and multiple regressions, we follow CV and CGPT and measure ICAPM betas by scaling the covariance of value returns with each ICAPM shock by the total variance of unexpected market returns, not by the variance of the particular shock in question.

¹⁷Our focus is on the R^2 in these regressions rather than using the resulting betas to price the cross-section. Nevertheless, the standard errors associated with our estimates of beta in this table, as well as in all subsequent tables that use estimated news terms, require a caveat. For simplicity, we report OLS standard errors which do not take into account VAR sampling error. Thus, while the t -statistics in those tables are often large in absolute value, the true statistical precision of these estimates is likely to be lower. However, we note that our conclusions regarding the statistical significance of HML’s ICAPM betas are consistent with those in Campbell and Vuolteenaho (2004) and Campbell, Giglio and Polk (2013), both of whom explicitly account for VAR sampling error when estimating ICAPM beta standard errors.

that provides the lion’s share of exposure to the two ICAPM betas that command relatively large risk prices. The component of value that does not contribute to HML’s CAPM alpha is also the component that is mostly sensitive to the temporary component of aggregate returns, which carries the smallest risk price in the ICAPM.

The ability of the ICAPM to explain the fluctuations of these portfolios is also markedly different between HML_{Intra} and HML_{Inter} . While the ICAPM has the highest (time-series) explanatory power for HML, it has similar power for HML_{Intra} , but performs significantly worse in explaining the fluctuations of HML_{Inter} . This difference in explanatory power indicates that (unpriced) risk factors that are orthogonal to the priced factors of either the CAPM or the ICAPM drive most of the time-series variation in HML_{Inter} . Finally, all these patterns continue to hold when we examine separately the pre-1990 and the post-1990 samples, as the next two panels of the table show.

Figure 4 illustrates visually the ability of the market and the ICAPM news terms to explain the time-series fluctuations of HML. Panel A shows the EWMA-smoothed log returns on HML, together with the fitted value from the multivariate regressions of Table 3 (the “ICAPM fit”, red dashed), and the fitted value from the regression on the market alone (the “CAPM fit”, green dotted).¹⁸

The figure shows the remarkable success of the ICAPM in explaining the variability of HML since the 1960s: the ICAPM fit closely tracks HML during its various booms and busts. In particular, the ICAPM explains a large fraction of HML’s performance during the major crashes that occurred during this period. These value busts include several crashes in the 70s and 80s, the tech cycle in the 90s, the financial crisis, and the Covid-19 pandemic (though during the Covid-19 pandemic, as we study in greater detail below, there is a nontrivial part of unexplained variation in HML, reflecting the uniqueness of the Covid-19 episode). In contrast, the market factor explains almost none of the variation of HML during the last 60 years. That said, there are some features of the returns of HML that are not fully captured by the ICAPM. For example, the ICAPM cannot explain the protracted underperformance of HML in the last 15 years: while the actual return of HML has been persistently negative over that period, the fitted value fluctuates around zero, before turning clearly negative

¹⁸As we are smoothing log returns, we generate the corresponding fitted values from regressions of log returns on news terms. The resulting betas are qualitatively similar, with extremely minor differences. We report those regressions in the Appendix (Table 5).

starting in 2015.

In Panel B, we decompose the returns of the ICAPM fit for HML into the part that is attributed to discount rate, cash flow, and variance news. HML’s various cycles in HML and in HML_{Intra} have different proximate causes. Some periods see offsetting effects of the different news terms (for example, in the 70s, when discount rate and variance news contributed positively to the HML and HML_{Intra} returns, while cash flow news contributed negatively). Other times, all three news terms push HML’s return in the same direction – for example, positively in the 2000s before the financial crisis, and negatively during the Covid-19 crash.

Appendix Figures 5 and 6 repeat the exercise for HML_{Intra} and HML_{Inter} , respectively. They show that the ICAPM has a similar (though slightly lower) explanatory power for the fluctuations in HML_{Intra} , whereas the shocks to HML_{Inter} are difficult to explain for either the CAPM or the ICAPM (and this inability is evident in crashes, including the financial crisis and the Covid-19 episode, as well). Moreover, whereas the variation in HML_{Inter} over time is mostly attributed to discount rate news, the various cycles in HML_{Intra} are very similar to those of HML. Given the close relation between HML and HML_{Intra} , it is not surprising that the decomposition looks highly similar across the two.

To summarize, in this subsection we have shown that the ICAPM news terms help explain a large fraction of the variability of HML’s return (and the intra-industry component) over time, and can be used to understand the drivers of the booms and busts of value investing. Importantly, the analysis shows that the proximate causes of HML’s fluctuations are not always the same over time, highlighting the usefulness of the approach developed in this paper to understand the different causes underlying different episodes.

4 Further Quarterly Analysis

4.1 Other factors and anomalies

In this subsection, we document how exposure to ICAPM news terms varies more broadly for other factors and anomalies in the empirical literature. Given our focus on value investing, we first decompose HML into intra- and inter-industry components within the large-cap and small-cap portions of the stock market. Panel A of Table 4 shows that the improvement in R^2

is higher for both the small-cap and large-cap intra-industry components of HML. We also find that the small-cap inter-industry component is much better explained by the ICAPM, primarily because that component is more sensitive to aggregate discount-rate news. Within both the large-cap and small-cap component of HML, we continue to find that the intra-industry component is riskier in terms of cash-flow and variance betas. Since both these small-cap and large-cap value strategies derive most of their abnormal return from their intra-industry component, this more granular analysis is consistent with CGPT’s ICAPM.

Panel B of the table reports the extent to which the other non-market Fama-French-Carhart factors load on the ICAPM. We find that the ICAPM produces relatively high R^2 s for the size and investment factors. The ICAPM R^2 is 20% for SMB, Fama and French’s size factor, which is higher than its CAPM R^2 of 15%. The ICAPM R^2 for Fama and French’s investment factor, CMA, is 36%, more than double the CAPM R^2 of 14%. For both of these factors, the ICAPM loadings are of the correct sign to help explain the premia on those two factors.

However, the ICAPM R^2 is rather poor for the Fama-French profitability factor, RMW. That factor’s ICAPM R^2 is only 3%, and actually lower than its CAPM R^2 because of the adjustment for the loss of two degrees of freedom. Though the ICAPM R^2 for the Carhart momentum factor, MOM, is 12%, which is higher than its CAPM R^2 of 6%, all three of momentum’s ICAPM betas are of the wrong sign to explain the positive average return to momentum trading. Of course, the inability of the ICAPM to explain momentum may not be that surprising, as most researchers view a risk-based explanation of the momentum effect as unlikely.

In the fifth row of Panel B, we measure the CAPM and ICAPM betas of the fundamental-to-market (F-M) factor of Goncalves and Leonard (2023). Consistent with the ICAPM, F-M has a positive cash-flow beta that is economically and statistically significant when measured in a multiple regression. However, in contrast to HML and HML_{Intra} , F-M has a positive, though statistically insignificant, multiple variance beta.¹⁹ The ICAPM R^2 for F-M is higher than the CAPM R^2 of 7% but remains fairly modest at 13%.

Finally, the last two rows of the Panel B report the results for two different long-short portfolios sorted on duration, measured, respectively, as in Golcalves (2021) and Weber

¹⁹In unreported pricing tests, we find that F-M does not subsume HML_{Intra} , though HML_{Intra} does not subsume F-M either.

(2018). The duration and value patterns have been linked by the literature (e.g., Lettau and Wachter 2007), and the table shows that, indeed, for these two portfolios the ICAPM explains approximately double the time-series variation explained by the CAPM.

Of course, there are hundreds of anomalous trading strategies documented in the empirical asset-pricing literature. To underscore the improvement in explanatory power that the ICAPM provides, we turn to the 193 anomalies replicated in Hou, Xue, and Zhang (2020). One benefit of those data, relative to otherwise similar datasets, such as the open-source data of Chen and Zimmermann (2022), is that Hou, Xue, and Zhang (2020) categorize their 193 anomalies into six subsets (strategies based on either value-versus-growth, intangibles, investment, profitability, momentum, or frictions). We rely on their categorizations to guide and constrain our analysis, as Table 4 documents that the ICAPM’s improvement is far from uniform across the five Fama-French-Carhart non-market factors.

In the six panels of Figure 5, we report the cdfs of the resulting adjusted R^2 s when explaining the quarterly returns on the anomalies within each of these six subsets. Consistent with our findings in Table 4, the improvement in explanatory power for the ICAPM relative to the CAPM is clearly evident for the value, intangibles, and investment classes of anomalies. These findings confirm that the ICAPM does substantially better than the CAPM in explaining the realized returns on an economically important portion of the cross-section of average returns. However, for the remaining three classes of anomalies – profitability, momentum, and frictions – the ICAPM does not offer that much improvement relative to the CAPM.

4.2 Robustness of the HML decomposition

For robustness, we also consider several alternative ways to decompose HML into intra-industry and inter-industry components, reporting detailed results for these alternatives in the appendix. In a first alternative decomposition, we retain the structure of our basic approach, but use 17 aggregated Fama-French industries rather than 48.

In a second alternative decomposition, we construct the intra-industry HML portfolio by following the same procedure Fama and French (1993) use to build HML, using stocks from all industries but cross-sectionally demeaning the book-to-market ratio within each industry. This procedure is the one used in Novy-Marx (2012). As also discussed in that

paper, this intra-industry component has some residual industry characteristic exposure, which can be offset with corresponding positions in the underlying industries, obtaining an intra-industry HML component that is fully industry-neutral, that is, it has zero exposure to industry characteristics. As before, we construct HML_{Inter} as the residual of a projection of HML onto the intra-industry component and build HML_{Intra} by rescaling the intra-industry component so the regression coefficient is 1.

In a third alternative decomposition, we compute an inter-industry HML portfolio, $HML_{\text{InterDirect}}$, defined as the portfolio of industries held in the same proportion as HML's industry characteristic exposures. Subtracting $HML_{\text{InterDirect}}$ from HML gives a portfolio that is, by construction, industry-neutral (i.e., it has zero industry characteristic exposure), which is our second alternative version of HML_{Intra} . We again construct HML_{Inter} as the residual of a projection of HML onto the intra-industry component, and build HML_{Intra} by rescaling the intra-industry component so the regression coefficient is 1.

The appendix shows that all three of these alternative HML decompositions deliver results that are broadly comparable to those we report in the paper for our basic decomposition.²⁰ To confirm that our finding is robust, we not only vary the way in which we decompose HML, we also vary a key aspect of the VAR. In particular, we replace the smoothed PE variable with forward PE , the log ratio of the price of the S&P 500 to expected one-year earnings from IBES.²¹ Appendix Table 9 reports the resulting ICAPM betas for HML and its intra- and inter-industry components. We find that the resulting R^2 s are very similar to those from our baseline specification in Table 3.

4.3 The link between time-series R^2 and Sharpe ratio

We have emphasized the explanatory power of the three ICAPM shocks for the time-series behavior of HML and its intra-industry component. But how does this explanatory power

²⁰Specifically, the results for the 17 industries are in Appendix Tables 2, 6, 10, and 17, and Appendix Figures 1, 7, 10, 16, 19, and 22; the results for the alternative construction for HML_{Intra} are in Appendix Tables 3, 7, 11, and 18, and Appendix Figures 2, 8, 11, 17, 20, and 23; the results for the last alternative approach are in Appendix Tables 4, 8, 12, and 19, and Appendix Figures 3, 9, 12, 18, 21, and 24. Appendix Figure 4 plots the series together.

²¹To measure forward-PE, we use the quarterly series kindly shared by Ricardo De la O on his website, which starts from 1976Q1. Given that forward-PE and our baseline measure of PE (CAPE) are very close in 1976Q1, we join the two series together, using CAPE before 1976Q1 and the forward-PE after that. We then use this combined series to estimate the VAR and perform the empirical analysis.

relate to the Sharpe ratio that an investor should expect to earn on these portfolios if the ICAPM prices assets? The answer is straightforward. If we combine the ICAPM shocks using the correct weights to create a proxy for the log stochastic discount factor implied by the ICAPM and regress a portfolio log return onto this proxy, then the model-implied Sharpe ratio of the portfolio is proportional to the square root of the R^2 statistic from this regression.²² The coefficient of proportionality is simply the standard deviation of the SDF, multiplied by the negative of the sign of the regression coefficient.

We conduct this exercise in Table 5. The table combines the three ICAPM shocks using a risk aversion coefficient of 7.2 estimated by CGPT, along with the variance risk price of 24.9 implied by this estimate. In both the full modern sample and its two subsamples, we find R^2 statistics slightly above one quarter for HML and intra-industry HML, and R^2 statistics close to 1% for inter-industry HML. The full-sample estimates imply quarterly Sharpe ratios of 0.22, 0.22, and 0.05 respectively. We confirm in the Appendix (Appendix Tables 10, 11, and 12) that these findings are robust to using our alternative HML decompositions.

4.4 Is the high HML R^2 due to having the value spread in the VAR?

Of course, the explanatory power we find is in part due to the inclusion of the value spread in the VAR, and shocks to the value spread should be negatively correlated with the return on HML. Indeed, a regression of HML on the value spread shock has an R^2 of 52%. However, it is important to recognize that the value spread shock is combined with all of the other five shocks driving the VAR when computing the ICAPM news terms (via the λ vectors). Moreover, the λ vectors are not optimized to describe the cross-section of average returns or to explain the realized returns on HML. Instead, they are chosen based on the ability of the variables in the VAR, such as the value spread, to forecast future aggregate returns and risks, as well as on the dynamics of the state variables, as the economic logic of the ICAPM

²²We perform this regression using log returns because CGPT's ICAPM delivers a log SDF. The formula follows from equation (4.43) on p. 94 of Campbell (2018). Squaring both sides of the equation, dividing by the variance of the excess log portfolio return, multiplying and dividing the right-hand side by the variance of the log SDF, and rearranging terms shows that the squared Sharpe ratio is the R^2 statistic times the variance of the log SDF. Taking the signed square root delivers the desired result. Note that the Sharpe ratio from this equation uses the approximation (discussed also in CGPT) that $E_t[r_{p,t+1}] - r_{f,t+1} + \frac{\sigma_{p,t}^2}{2} \simeq E_t[R_{p,t+1} - R_{f,t+1}]$, and that the variance of log and level returns are approximately the same.

explicitly requires.

Another way to see that our results do not follow mechanically from the inclusion of the value spread in the VAR is to study the explanatory power of ICAPM models for other anomalies, where for each anomaly we include in the VAR a valuation spread based on sorting stocks on the anomaly characteristic rather than on the book-market ratio.

We first explore this in the context of HML_{Intra} and HML_{Inter} , in Appendix Table 13. For that table, we have built two separate versions of the value spread: VS_{Intra} and VS_{Inter} . The first row of the table replicates our baseline result – the high explanatory power of our VAR with the small-stock value spread for the return of HML . The second row of the table includes both VS_{Intra} and VS_{Inter} in the VAR, and uses the resulting news terms to explain the HML return. This version results in a higher R^2 than in the CAPM, though lower than in the baseline case. The lower R^2 is possibly due to the increase in noise from adding a less useful state variable to the VAR (which adds several more parameters to estimate). The third row uses only the VS_{Intra} in the VAR and reports the R^2 of HML_{Intra} , and the last row uses only the VS_{Inter} in the VAR and reports the R^2 of HML_{Inter} .

The table shows that the inclusion of the value spread corresponding to the long-short portfolio whose return we are trying to explain does not *guarantee* a high R^2 . This finding is directly visible in the last row: it's not because we include the inter-industry value spread in the VAR that the news terms can explain fluctuations in HML_{Inter} . In fact, the R^2 for HML_{Inter} is lower when using its own value spread (10%, from this table) than in the baseline specification using the standard value spread (14%, from Table 3). These results help exclude mechanical effects as an explanation for our main analysis.

We also conduct a similar exercise in Appendix Table 14 and Appendix Figure 13, but looking at a large number of anomalies, and find that the improvement in fit from the CAPM to the ICAPM is lower in these placebo exercises compared to the case when the value spread is included in the VAR.

5 A Daily ICAPM Decomposition of Inter- and Intra-industry HML Returns

So far, we have conducted all the analysis at the quarterly frequency, following Campbell and Vuolteenaho (2004), Campbell, Giglio and Polk (2013), and CGPT. As these papers discuss, a quarterly first-order VAR allows us to capture the long-term dynamics of the state variables very well and can be estimated based on nearly a century of data.

A shortcoming of the quarterly VAR specification is that it produces news terms at that same (quarterly) frequency. This aspect is sufficient for studying fluctuations in HML at relatively low frequencies, but is of limited help to understand episodes that unfold at higher frequencies. An interesting example is the Covid-19 market crash and subsequent recovery, which occurred rapidly over the span of a few quarters in 2020.

In this section we extend the VAR methodology to estimate the ICAPM news terms at the daily frequency. We do so by combining information from the quarterly VAR (specifically, the mapping between innovations in the state variables and aggregate news terms) and from a daily VAR that captures the evolution of the state variables at high frequency. We then use the estimated daily news terms to zoom into the Covid-19 episode.

5.1 Methodology

Our methodology combines the quarterly VAR estimated over the full sample (1926Q3 to 2022Q4), with a daily VAR estimated over a shorter sample (1990-01-04 to 2022-12-31), to obtain *daily* series for N_{DR} , N_{CF} , and N_V . The daily VAR is used to obtain within-quarter (daily) innovations in the state variables. The quarterly VAR is used to map these innovations in the state variables to long-run news about future returns, cash flows, and variance.

As discussed in section 2.1, under the assumptions maintained by CV, Campbell, Giglio and Polk (2013), and CGPT, namely, that the quarterly VAR (with state vector $\mathbf{x}_t$) captures the long-run dynamics of those variables well, the news terms in equations (2), (3), and (4) can be written as a linear function of the *innovations* in $\mathbf{x}_t$, $\eta_{t+1} \equiv \mathbf{x}_{t+1} - E_t \mathbf{x}_{t+1}$. The mapping from VAR innovations to the news terms is captured by the vectors $\lambda_{DR}, \lambda_{CF}$ and λ_V . For example, for the case of volatility, we have $N_{V,t+1} = \lambda_V \eta_{t+1}$, as described

in section 2.1. The quarterly news shocks are therefore constructed by first computing quarterly innovations in $\mathbf{x}_t$ (i.e., the innovations from the quarterly VAR, η_{t+1}), and then mapping those innovations into long-term news (via the λ s, estimated from the same VAR).

Using a quarterly VAR, one can therefore only compute news shocks at the quarterly frequency. The procedure we present in this section augments this setup with a higher-frequency VAR that allows us to track the innovations in $\mathbf{x}_t$ *within* the quarter. To do so, we focus on a (shorter) sample for which we observe, at the daily frequency, a set of state variables that can be used to construct *daily* innovations in $\mathbf{x}_t$, that can be then mapped into daily news terms via the vectors λ .

We proceed as follows. Denote days by the subscript s , and call $\tilde{\mathbf{x}}_s$ the vector of daily variables (which, in general, is different from the vector of variables that enter the quarterly VAR, $\mathbf{x}_s$). We begin by estimating a first-order daily VAR, which gives us daily innovations in $\tilde{\mathbf{x}}_s$ as well as its daily dynamics (e.g., the daily transition matrix $\tilde{\Gamma}$).

Next, we use the estimated daily VAR to compute daily innovations in the state variables that enter the quarterly VAR, $\mathbf{x}_s$. To do so, we proceed as follows. First, recall that the first element of $\mathbf{x}_s$ is the trailing 60-trading-day log real return of the market. The daily innovation in this variable is:

$$\mathbf{e}'_1(\mathbf{x}_s - E_{s-1}\mathbf{x}_s) = \sum_{j=0}^{59} r_{s-j} - E_{s-1} \left[\sum_{j=0}^{59} r_{s-j} \right] = r_s - E_{s-1}[r_s] = \mathbf{e}'_1(\tilde{\mathbf{x}}_s - E_{s-1}\tilde{\mathbf{x}}_s)$$

where the log real market return is also the first element of the daily state vector $\tilde{\mathbf{x}}_s$; this result follows from the fact that as of time $s-1$, the returns for the 59 days up to $s-1$ are known.

The second element in $\mathbf{x}_s$ is $EVAR_s$, that is, $E_s[\sum_{j=1}^{60} R_{s+j}^2]$. Using the estimated intercept and transition matrix from the daily VAR, we can compute $E_s[\sum_{j=1}^{60} R_{s+j}^2]$, and therefore its daily innovations, as long as R_s^2 , the squared market return, is included in the daily state vector. The other elements of the quarterly VAR (PE ratio, term spread, default spread, and value spread) are observed at the daily frequency, so the daily innovations can be read directly from the daily VAR as long as these variables are included in $\tilde{\mathbf{x}}_s$.

Motivated by the discussion above, our daily state vector $\tilde{\mathbf{x}}_s$ includes the daily log market return r_s , the squared market return, R_s^2 , the PE ratio, the term spread, the default spread, and the value spread. In addition, we also include in $\tilde{\mathbf{x}}_s$ variables that are useful to predict

returns and variances at high frequencies, and that are available at the daily frequency in the sample we use for estimation (1990-01-04 to 2022-12-31): the 60-day rolling sum of the log market return, the 60-day rolling sum of the squared return, and the VIX (which is particularly useful to predict realized variance, see, for example, Berger et al. 2020).

Once we have obtained daily innovations in the quarterly state variables, ($\tilde{\eta}_s \equiv \mathbf{x}_s - E_{s-1}\mathbf{x}_s$), we combine them with the λ vectors that map the innovations in the quarterly state variables into the three news terms, obtaining daily news series. For example, for the case of volatility, we will have:

$$N_{V,s+1} = \lambda_V \tilde{\eta}_{s+1}$$

Before moving to the empirical results, it is useful to discuss how we can interpret the combination of VARs estimated at different frequencies and using different state variables. Ideally, we would want to use one specification for the entire analysis – the “true” time-series model incorporating the investors’ full information set, with all the correct predictors (e.g., including VIX), at the highest possible frequency. In practice, however, some series are not available at the daily frequency for a very long period, and correctly capturing the long-term dynamics that play an essential role when computing these news terms would likely require a much richer specification than simply a first-order daily VAR.

Our procedure combining the two VARs at different frequencies can be interpreted as an approximation of the way an investor facing data limitations similar to the ones we face would compute the news shocks. To learn the mapping between the current state variables $\mathbf{x}_t$ and long-term expectations of returns, cash flows, and volatility (i.e., to compute the λ s for the three news terms), one may want to sacrifice some potentially informative variables like the VIX that are only available for a short period, in exchange for having a long time series for the remaining variables. Therefore, an investor might want to use the quarterly VAR with a long time series to learn about the long-run responses of returns and volatility and estimate the λ s.

Given that the additional variables we include in the daily VAR (especially the VIX) have been shown to have predictive power mostly for the short-run evolution of volatility (again, see Berger et al. 2020), we do not believe that their omission from the quarterly VAR would induce a substantial bias for estimating λ . However, to identify the daily shocks to $\mathbf{x}_t$, the investor may want to instead use the full daily VAR, with all predictors available at higher frequency, even if for a shorter sample, since the forecasts from the daily VAR only

need to cover a 60 trading day horizon as an input for the news terms.²³

5.2 Results

We report the results of the daily VAR estimation in Appendix Table 20. Many of the patterns are similar to the ones we find in the quarterly VAR: for example, the PE ratio and the default spread help predict returns. Consistent with previous research, we find that the VIX plays an important role in predicting future realized variance (squared returns) even after controlling for lagged realized variance.

After building the daily news terms as described in the previous section, we next compare the estimated daily news with the news from the quarterly VAR. Note that, in theory, the daily news do not exactly sum up to the quarterly news within each quarter. The reason is that the innovations in $\mathbf{x}_s$ from the daily VAR (from which the news terms are constructed) do *not* add up to the quarterly innovations $\mathbf{x}_t$, since in the quarterly VAR we use time- t information to predict $\mathbf{x}_t$ over the entire following quarter, whereas in the daily analysis we try to predict $\mathbf{x}_s$ day by day.²⁴

That said, we do find that empirically the two are quite close: the sums of the daily news terms N_{DR} , N_{CF} , and N_V within each quarter are highly correlated with the quarterly news series obtained directly from the quarterly VAR. The correlations of the news terms are 0.95, 0.97 and 0.71 respectively for N_{DR} , N_{CF} , and N_V . Appendix Figure 14 shows the three news terms (one per panel) estimated using the quarterly VAR (solid line) and the daily VAR (dashed line). The figure smooths the series using the same EWMA as in Figure 3.²⁵

²³An alternative approach would be to estimate a mixed-frequency model to jointly estimate both the daily news and the long-run dynamics of the variables in $\mathbf{x}_t$. Work by Ghysels, Santa-Clara, and Valkanov (2005, 2006) and Ghysels, Sinko, and Valkanov (2007) has documented the usefulness of such a technique. However, this would be computationally more complex, and we leave it for future research. In contrast, our approach is simple to implement and easily interpretable.

²⁴More formally, the innovations from the quarterly VAR can be written, using the daily notation, as $\eta_{s+60} = \mathbf{x}_{s+60} - E_s \mathbf{x}_{s+60}$. Since the daily innovations are $\tilde{\eta}_{s+1} = \mathbf{x}_{s+1} - E_s \mathbf{x}_{s+1}$, then

$$\sum_{j=1}^{60} \tilde{\eta}_{s+j} = \sum_{j=1}^{60} (\mathbf{x}_{s+j} - E_{s+j-1} \mathbf{x}_{s+j}) \neq \mathbf{x}_{s+60} - E_s \mathbf{x}_{s+60} = \eta_{s+60}$$

²⁵We initialize the EWMA filter to zero.

We next re-examine the ICAPM explanation for the fluctuations in HML using the daily news. We begin by reporting in Appendix Table 14 the exposure of HML to the three news terms over the 1990Q2-2022Q4 sample. Each row in the panel corresponds to a different regression of the three portfolios (HML, HML_{Intra} , and HML_{Inter}), onto the three news terms, using a different construction of the news series and data frequency. Specifically, the first row of each panel reports the regression using the quarterly news series (at the quarterly frequency): that is, the same as in Panel C of Table 3, i.e., for the post-1990 sample. As before, we find that HML has a negative discount rate beta, positive cash flow beta, and negative variance beta; and again, the regression has a high time-series R^2 for HML and HML_{Intra} , and much lower R^2 for HML_{Inter} .

The second row of each panel uses the daily news series aggregated to a quarterly frequency (by summing the daily news within the quarter), and then estimates the betas of the three portfolios with respect to the news terms using quarterly data. The results are very similar to the ones in the first row, showing that at the quarterly frequency, there is little difference between using the daily or quarterly VAR to estimate the news terms.

The last rows of each panel report daily betas, obtained using the daily news terms (and the daily returns on the three portfolios). For comparability with the quarterly analysis, we include 59 lags of each news term (i.e., roughly a quarter of daily news), though results are qualitatively similar if we only estimate betas using the contemporaneous daily shock. Appendix Table 15 reports the sum of the betas on each news term from that regression.²⁶ The R^2 is lower (as we are now trying to explain daily fluctuations in HML, HML_{Intra} and HML_{Inter}), but is still quite high (27%, 26%, and 15% respectively). The signs of the daily betas are almost always the same as in the quarterly case, at least for those quarterly estimates that are statistically significant.

Overall, Appendix Table 15 and Appendix Figure 14 show that at lower frequencies (quarterly), estimating the news daily makes little difference to the results. The main advantage of the daily news estimation is the ability to zoom into specific episodes, studying them at high frequency.

Appendix Figure 15 uses the beta estimates from this table to show the time series of the fitted HML (Panel A), HML_{Intra} (Panel B) and HML_{Inter} (Panel C) against the original series.

²⁶Appendix Figures 25, 26 and 27 plot how the estimated coefficients evolve as a function of lag; Appendix Table 16 repeats the analysis using log returns.

Each figure reports the ICAPM fit obtained using quarterly data (dashed red), using the daily news series aggregated to the quarterly frequency (solid red), and using the daily regression (green). There are two interesting patterns in this figure. First, fitting the portfolios using quarterly as opposed to daily regressions yields a better overall fit – more so for HML_{Intra} and HML_{Inter} than for HML . This worse fit may arise from either additional measurement error in the daily series or from asset returns that do not immediately reflect information in the daily news.

Second, and despite the worse fit, for all three series, the daily ICAPM does a better job in capturing the most severe booms and busts of HML , including, and most interestingly, the Covid-19 pandemic. The ICAPM fit for HML_{Inter} is poor regardless of the estimation frequency.

We conclude this section by zooming into the Covid-19 episode of 2020. Figures 6, 7, and 8 show the news terms, the ICAPM fit, and the components of that ICAPM fit, respectively, during 2020. In Figure 6, we plot both the quarterly and daily news terms, and we do *not* smooth them using an EWMA as in the previous figures so that the daily variation appears clearly (we instead report cumulative news terms, demeaned within this year to focus on the high-frequency variation). Clearly, the quarterly news can only give us a very coarse picture of the events of 2020, since it includes just four data points for the year. The daily news give us a much higher frequency account of the events, including the very rapid developments in February 2020, and the arrival of large news terms at other points in the year. Perhaps most interesting are the dynamics of volatility, in Panel C: the daily news series captures an extreme, but transitory, peak that is in large part missed by the quarterly series, and is reverted substantially before the end of the second quarter.

Figure 7 shows how HML and its components fared during this period. The figure shows that the daily ICAPM fit tracks the high-frequency movements of HML during the year, in a way that is almost entirely missed by the quarterly news. During the Covid-19 pandemic, even HML_{Inter} 's fluctuations are explained by the ICAPM to some degree – though the explanatory power for HML and HML_{Intra} , is, as usual, much higher.

As discussed above, the Covid-19 pandemic was characterized by the arrival of large cash-flow, discount rate, and volatility shocks. Exposure to these three shocks goes a long way to explain not only why HML dropped substantially at the beginning of 2020, but also why it recovered towards the end of the year.

We again decompose the ICAPM fit into the part that is attributed to discount rate, cash flow, and variance news focusing on 2020. Figure 8 Panel A reports the decomposition for HML, whereas the other two panels show the decomposition for HML_{Intra} and HML_{Inter} .

The figure shows that though HML was buffeted by all three shocks in 2020, the volatility and discount-rate news shocks largely offset each other. In contrast, cash-flow news is the primary driver of the component of value returns attributable to the ICAPM. Consistent with our previous results, this decomposition of HML’s ICAPM fit is mainly attributable to the intra-industry component of HML. In contrast, variation in the ICAPM fit of the inter-industry component of HML comes primarily from discount-rate news.

To conclude, in this section we showed that the ICAPM news terms help explain a large fraction of the variability of HML’s return (and the intra-industry component) over time, and can be used to understand the drivers of the booms and busts of value investing.

We summarize our findings with a variance decomposition of value returns. Specifically, Appendix Table 20 reports the percentage of return variation explained by each of the three ICAPM shocks, at both the quarterly and daily frequency, for HML, HML_{Intra} , and HML_{Inter} . Panel A shows that the ICAPM explains roughly half of the variation in quarterly HML returns over the 1963Q3-2022Q4 sample. The majority of that explanatory power comes from the intra-industry component of HML. That component is primarily explained by cash-flow (16%) and variance news (14%). In contrast, most of the variation in inter-industry HML is unexplained (85%), and any explanatory power comes entirely from aggregate discount-rate news. When we move to the daily frequency, we continue to find that the ICAPM explains relatively more of the intra-industry component of HML (27% compared to only 17% for inter-industry HML). Interestingly, during 2020, there was a significant increase in the contribution of cash-flow news to the variation in both components of HML (21% in both instances).

6 Earnings Shocks and Industry Shocks During the Covid-19 Pandemic

In the previous sections, we showed that the ICAPM model is able to explain a non-trivial portion of the movements of HML during the pandemic. However, the explanatory power of

the ICAPM comes relatively more from the intra-industry component of HML. The inter-industry component of HML (that is, the return of the portfolio of industries that HML is exposed to) also suffered during the pandemic, in a way that the ICAPM does not explain.

In this section, we look directly at the earnings of the industries that HML is exposed to during 2020, a period that saw large industry-specific shocks related to the Covid-19 pandemic. We show that HML underperformed during the pandemic beyond its exposure to aggregate news, because it happened to be long industries whose earnings were poor during the pandemic, and short industries whose earnings did relatively better during the pandemic. These earnings patterns do not simply reflect exposure to aggregate cash flow news, which we believe is the reason why the ICAPM does not explain well the returns of the inter-industry component of HML during this period.

To study the earnings of the industries that HML is exposed to, we proceed as follows. We first compute the weight of HML in each industry (that is, the weights used to build $HML_{InterDirect}$). Then, we separately consider a portfolio that buys all industries in which HML has long exposure (L), and the portfolio that buys all industries in which HML has short exposure (S). We refer to the former as “value industries” and to the latter as “growth industries”.

For each of the two portfolios, $p = \{L, S\}$, call w_{it}^p the weight of stock i in HML as of the end of quarter t . For each stock i , MV_{it} is the market value at the end of quarter t , E_{it} the total earnings during quarter t , and BE_{it} the total book equity as of end of quarter t . For each portfolio $p = \{L, S\}$, we construct an aggregated measure of earnings during quarter t scaled by the previous quarter’s book equity:

$$ScaledEarn_t^p = \frac{\sum_i w_{it-1}^p E_{i,t} / MV_{it-1}}{\sum_i w_{it-1}^p BE_{it-1} / MV_{it-1}} \quad (7)$$

Note that for a value-weight portfolio, weights are proportional to market cap, so that this formula reduces to the total earnings over the previous-quarter total book equity of the portfolio. We construct these scaled earnings for the four quarters of 2020, and then, starting from $t = 2020q1$, we build a cumulated series (“cumulated scaled earnings”) during the year. We residualize the earnings series of the long and short HML portfolios relative to analogous measures for the aggregate market, to account for different exposure of the two portfolios to aggregate earning shocks.

The top panel of Figure 9 shows these (residualized) cumulated scaled earnings for the portfolio of value industries and growth industries, with black and red lines respectively. The figure shows that during 2020, value industries had relatively flat earnings, whereas growth industries tended to have relatively good earnings, contributing to a particularly negative cash-flow shock for the HML_{Inter} portfolio.

Next, we decompose these results by industry, in the bottom panel of the figure. For each of the 48 industries, we compute the contribution of that industry to the cumulative scaled earnings of the value and growth portfolios during 2020 (described in equation 7, and plotted in the top panel). The bottom panel plots this contribution (again residualized to the aggregate market) against the weight of that industry in HML. So the industries on the left side of the figure are “growth industries” (portfolio S), and industries on the right side are “value industries” (portfolio L). The figure shows that HML has little exposure to many majority of industries. Some industries, however, are significantly weighted by HML over this period, positively (e.g., oil and banks), or negatively (business services, drugs).

The earnings of the value industries in 2020 were mixed, as some (banks) had good earnings and others (oil) had negative earnings during the pandemic. The performance of the growth industries (that enter negatively into HML) was instead much more positive, for reasons that are clearly idiosyncratic to the pandemic. This fact explains the patterns we observe at the aggregate level in the top panel of the figure.

Overall, these results show that the Covid-19 pandemic featured important industry-specific shocks to earnings, that went beyond what can be explained by the different industries’ exposure to cash flow news. Our findings help explain why the ICAPM can account for only a small fraction of the fluctuations in HML_{Inter} over this episode.

7 Time-variation in HML’s ICAPM Betas

So far, the focus of the paper has been on studying the unconditional time-series relation between the returns of HML and the ICAPM news shocks. In that analysis we have assumed that the exposures of HML to the shocks are constant, and we have estimated them using the entire sample period or relatively long subsamples. In this section, we document interesting variation in those exposures over time. While a full analysis of the dynamics of the conditional exposures is beyond the scope of this paper, the evidence suggests that these

dynamics may have contributed to the well-known underperformance of HML over the last 15 years: an increase in HML’s exposures to priced news shocks led to a higher expected return and lower realized returns on HML as this risk was repriced by investors.

We explore these time patterns in HML’s beta with respect to the news shocks in Table 6 and Figure 10. The figure shows five-year rolling window estimates of the discount-rate news (Panel A), cash-flow news (Panel B), and variance news (Panel C) betas of HML, HML_{Intra} , and HML_{Inter} . We plot a vertical bar for the window ending in 2006Q4, as the subsequent years are known to have very low realized returns on value-minus-growth strategies (for example, the cumulative return of HML over this period has been -50%). The figure shows important changes in the betas of HML and its components during this period, with a marked upward trend in cash-flow betas and discount-rate betas, and a downward trend in volatility betas (which became more negative over time).²⁷ Panel A of Table 8 summarizes these differences by reporting the values as of 2006 and as of 2022, and shows that, once again, HML and HML_{Intra} behave similarly (and differently from HML_{Inter} , which displays much less of a pattern in betas over time).

These changes in betas imply an increase in the expected return on HML implied by the ICAPM model (taking the prices of risk as constant). In the ICAPM, HML commands a high risk premium because it is exposed to cash-flow news (positively) and volatility news (negatively), both of which carry large risk compensations; the exposure to discount rate news carries a much lower risk premium. Panel A of Table 6 and Figure 10 show that over the last 15 years, HML became riskier because of increased exposure to both cash-flow risk and variance risk. Intuitively, we would then expect the risk premium to HML to have increased, leading to relatively lower prices for value stocks and lower returns to HML as these changes unfolded.

Evidence for this repricing is presented in Panel B of Table 6. The left two columns show the average realized return over the sample as well as a dummy variable for the change in the average return after 2006, confirming an economically large decline in average returns over the last 15 years. Though the change in HML’s average return is not statistically significant (-1.74%, with a t -statistic of -1.78), the corresponding estimate for HML_{Intra} is (-2.05%, with a t -statistic of -2.63). The right two columns show, respectively, the change in risk premium implied by the model over this period (using the prices of risk estimated from CGPT) as

²⁷Appendix Figure 28 repeats the figure for the alternative constructions of HML_{Intra} and HML_{Inter}

well as the implied “repricing return,” computed from the model under the assumption that the changes in betas were unanticipated by investors. Specifically, we multiply the change in risk premia implied by the ICAPM by -1 times the duration of the typical stock, which we assume to be 20 to be consistent with our assumption for ρ , and then divide by the number of years in the post-2006 sample. Since the risk premia estimates from CGPT are already in quarterly units, this calculation gives us the average quarterly repricing return. The fourth column of the table shows that the repricing return was slightly more negative than the post-2006 decline in the average realized return for HML, and was about 80% as negative as the post-2006 decline for intra-industry HML. Finally, consistent with the repricing hypothesis, we note that the value spread has widened since 2006, from 1.30 at the end of 2006 to 1.63 at the end of 2022, a 25% increase. The intra-industry value spread also increased by 36% over this period from 1.16 at the end of 2006 to 1.58 at the end of 2022.

While this is a simple back-of-the-envelope calibration that ignores important forces (including potential changes in risk prices during this period as well as the extent to which the change in betas was anticipated), the fact that the implied return due to the repricing of the exposure is relatively close to the observed change in the data suggests that repricing of risk might have been at least one contributor to the dismal performance of HML since 2006.²⁸

8 Conclusion

We show that the ICAPM, in stark contrast to the static CAPM, explains a significant portion of the realized returns on value-minus-growth strategies (like Fama and French’s (1993) HML) and thus can help us understand the large swings in the fortunes of value investors over the last 60 years. A key implication of the ICAPM is that investors price exposure to three main sources of risk, captured by cash-flow news, discount-rate news, and

²⁸One issue that does not play a large role in our analysis is the fact that HML is a managed portfolio. Cohen, Polk, and Vuolteenaho (2003) were the first to highlight the difficulty in applying a VAR-based decomposition to a managed portfolio (see their Appendix A). Indeed, explicitly because of this issue, Campbell, Polk, and Vuolteenaho (2010) first estimate a firm-level VAR and then aggregate the resulting news terms to compute discount-rate and cash-flow news for portfolios of value and growth stocks. See Lochstoer and Tetlock (2020) who echo these points and follow Campbell, Polk, and Vuolteenaho’s methodology (2010). However, this issue does not affect our repricing exercise, as the change in beta we measure is associated with the portfolio, not the stock.

volatility news. Previous literature has shown that HML is significantly exposed to all news terms, in a way that can explain its high expected return. In this paper, we show that the exposure to the three shocks actually accounts for a large part of the time-series variation in HML’s returns: the differential exposures of value and growth stocks to these three key drivers of returns, measured using a quarterly vector autoregression, explain roughly half of its realized quarterly returns.

The classic value strategy bets on the convergence of valuation ratios both across and within industries. However, the intra-industry component of such a bet typically contributes much more to value’s average outperformance. We exploit this distinction to discipline our ICAPM interpretation of the drivers of the returns to value investing. If the ICAPM is a useful model of expected returns, the intra-industry component of value should load more on ICAPM shocks, especially those that command a higher price of risk.

Our results fare well in the face of this refinement. We find that the ICAPM does a much better job of explaining the realized returns on the intra-industry component of a value-minus-growth bet. Moreover, intra-industry value is especially sensitive to cash-flow and volatility news, the two ICAPM factors that command a much higher risk premium. In contrast, the ICAPM does not explain inter-industry value as well, and any explanatory power comes primarily from discount-rate news, the ICAPM factor that commands a much lower risk premium.

We also study the explanatory power of ICAPM shocks for the returns to other investment strategies that have been explored in the enormous literature on equity market factors and anomalies. We find that the ICAPM has good explanatory power for the Fama-French size and investment factors but not for the profitability factor or the Carhart momentum factor. We obtain comparable results when we look at 193 equity anomalies replicated and categorized by Hou, Xue, and Zhang (2020): anomalies related to value, intangibles, and investment are better explained by the ICAPM than anomalies related to profitability, momentum, or frictions.

Previous research typically measures ICAPM news terms at the monthly or quarterly frequency, so that the crucial low-frequency dynamics of aggregate risks and returns can be measured precisely, over as long a sample as possible. However, aggregate discount-rate, cash-flow, and variance news can instantaneously move market prices in a substantial fashion. As a consequence, the potential benefits of an ICAPM perspective to explain *realized* returns,

especially the sorts of sharp booms and busts that one observes in value returns, may be limited by this aspect of existing methods.

We fill this gap by developing a novel methodology that marries low-frequency and high-frequency VARs to generate daily ICAPM news terms. Our technique retains the valuable information about the low-frequency component of aggregate risks and returns that a long sample based on infrequent data (e.g., accounting-based valuation ratios) can bring while also adding the incremental information in novel high-frequency data, often only available in recent subsamples (e.g., VIX), concerning the high-frequency dynamics of those key low-frequency state variables driving aggregate risks and returns.

With our new method in hand, we then confirm that the aforementioned conclusions about the relative importance of the CAPM and the ICAPM across the intra- and inter-industry components of *quarterly* value returns are also evident in the *daily* movements of value-minus-growth returns, especially during the Covid-19 pandemic. Cash-flow news drove a significant proportion of the returns on intra-industry value in 2020, while inter-industry value was not well-explained by the ICAPM. Instead, shocks to the relative earnings of value and growth industries were key drivers of the returns to inter-industry value during the pandemic.

Finally, we also provide some evidence that the riskiness of value and growth stock (as captured by their exposures to ICAPM news shocks) has changed significantly in the last 15 years. Specifically, exposure to cash-flow and discount-rate shocks increased, and exposure to volatility shocks became stronger (more negative), making the HML portfolio riskier for an ICAPM investor. This is consistent with the observation that HML has underperformed over this period, as the higher risk exposures would have implied higher discount rates, lower prices and simultaneously lower realized returns as HML's risks were repriced. While this is not the only possible explanation for the underperformance, our analysis suggests it may have been a potential contributor.

To conclude, the large literature on the value premium and other asset pricing anomalies has mostly focused on explaining cross-sectional patterns in average returns. The main reason for this focus is that theory has clear predictions about the relation between risk exposures and risk premia, but existing models are often silent on the sources of time-series variation in returns (especially the component driven by unpriced risks). Studying the time-series fluctuations of returns for diversified portfolios like HML, including attributing them to

priced and unpriced risks – as we do in our paper – is a step towards a deeper understanding of anomaly returns, and can guide the development of theoretical models that aim to explain their fundamental sources.

References

- Asness, Clifford, R. Burt Porter, and Ross Stevens, 2000, “Predicting Stock Returns Using Industry Relative Firm Characteristics”, Working paper, University of Florida.
- Baker, Malcolm, and Jeffrey Wurgler, 2006, “Investor sentiment and the cross-section of stock returns”, *Journal of Finance* 61, 1645–1680.
- Basu, Sanjoy, 1977, “Investment Performance of Common Stocks in Relation to their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis”, *Journal of Finance* 32, 663–682.
- Basu, Sanjoy, 1983, “The Relation Between Earnings Yield, Market Value, and Return for NYSE Common Stocks: Further Evidence”, *Journal of Financial Economics* 12, 129–156.
- Berger, David, Ian Dew-Becker, and Stefano Giglio, 2020, “Uncertainty Shocks as Second-Moment News Shocks”, *Review of Economic Studies* 87, 40–76.
- Campbell, John Y, Stefano Giglio, and Christopher Polk, 2013, “Hard Times”, *Review of Asset Pricing Studies* 3, 95–132.
- Campbell, John Y, Stefano Giglio, and Christopher Polk, 2022, “Appendix to What Drives Booms and Busts in Value?”, appendix available online.
- Campbell, John Y., Stefano Giglio, Christopher Polk, and Robert Turley, 2018, “An Intertemporal CAPM with Stochastic Volatility”, *Journal of Financial Economics* 128, 207–233.
- Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2010, “Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns”, *Review of Financial Studies* 23, 305–344.
- Campbell, John Y. and Robert J. Shiller, 1988a, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors”, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y. and Robert J. Shiller, 1988b, “Stock Prices, Earnings, and Expected Dividends”, *Journal of Finance* 43, 661–676.
- Campbell, John Y. and Tuomo Vuolteenaho, 2004, “Bad Beta, Good Beta”, *American Economic Review* 94, 1249–1275.

- Chen, Andrew and Tom Zimmerman, 2022, “Open-source Cross-sectional Asset Pricing ”, *Critical Finance Review* 12, 207-264.
- Cohen, Randolph B., and Christopher Polk, 1998, “The Impact of Industry Factors in Stock Returns”, Working paper, London School of Economics.
- Cohen, Randolph, Christopher Polk, and Tuomo Vuolteenaho, 2003, “The Value Spread”, *Journal of Finance* 58, 609–641.
- Cong, Lin, Nathan George, and Guojun Wang, 2023, “RIM-based Value Premium and Factor Pricing Using Value-Price Divergence”, *Journal of Banking and Finance* 149, 1-22.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, “Measuring Mutual Fund Performance with Characteristics-Based Benchmarks”, *Journal of Finance* 52, 1035–1058.
- Dechow, Patricia M., Richard G. Sloan, and Mark T. Soliman, 2004, “Implied Equity Duration: A New Measure of Equity Risk”, *Review of Accounting Studies* 9, 197–228.
- Eisfeldt, Andrea L. , Edward Kim, and Dimitris Papanikolaou, 2020, “Intangible Value”, NBER Working Paper No. 28056.
- Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2014, “The Value and Ownership of Intangible Capital”, *American Economic Review* 104, 189–194.
- Fama, Eugene F. and Kenneth R. French, 1989, “Business Conditions and Expected Returns on Stocks and Bonds”, *Journal of Financial Economics* 25, 23–50.
- Fama, Eugene F. and Kenneth R. French, 1992, “The Cross-Section of Expected Stock Returns”, *Journal of Finance* 47, 427–465.
- Fama, Eugene F. and Kenneth R. French, 1993, “Common Risk Factors in the Returns on Stocks and Bonds”, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F. and Kenneth R. French, 1997, “Industry Costs of Equity”, *Journal of Financial Economics* 43, 153–193.
- Fama, Eugene F. and Kenneth R. French, 2020, “The Value Premium”, University of Chicago working paper.
- Frankel, Richard and Charles M.C. Lee, 1998, “Accounting Valuation, Market Expectation, and Cross-Sectional Stock Returns”, *Journal of Accounting and Economics* 25, 283–319.

- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, “There Is a Risk-Return Trade-Off After All”, *Journal of Financial Economics* 76 509–548.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2006, “Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies”, *Journal of Econometrics* 131 59–95.
- Ghysels, Eric, Arthur Sinko, and Rossen Valkanov, 2007, “MIDAS Regressions: Further Results and New Directions”, *Econometric Reviews* 26 53–90.
- Goncalves, Andrei and Gregory Leonard, 2023, “The Fundamental-to-Market Ratio and the Value Premium Decline”, *Journal of Financial Economics* 47, 382–405.
- Gormsen, Niels J. and Eben Lazarus, 2021, “Duration-Driven Returns”, forthcoming *Journal of Finance*.
- Graham, Benjamin, David Le Fevre Dodd, and Sidney Cottle, 1934, *Security Analysis*, McGraw-Hill, New York.
- Hall, Robert E., 2000, “E-Capital: The Link between the Stock Market and the Labor Market in the 1990s ”, *Brookings Papers on Economic Activity* 2000(2), 73–118.
- Hou, Kewei, Chen Xue, and Lu Zhang, 2020, “Replicating anomalies ”, *Review of Financial Studies* 33, 2019–2133.
- Israel, Ronen, Sarah Jiang, and Adrienne Ross, 2018, “Craftsmanship Alpha: An Application to Style Investing”, *Journal of Portfolio Management* 44, 23–39.
- Kogan, Leonid, and Dimitris Papanikolaou, 2013, “Firm characteristics and stock returns: The role of investment-specific shocks ”, *Review of Financial Studies*, 26, 2718–2759.
- Lettau, Martin, and Jessica A. Wachter, 2007, “Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium”, *Journal of Finance* 62, 55–92.
- Lev, Baruch, and Anup Srivastava, 2022, “Explaining the Recent Failure of Value Investing”, *Critical Finance Review*, forthcoming
- Lewellen, Jonathan, 1999, “The Time-Series Relations Among Expected Return, Risk, and Book-to-Market”, *Journal of Financial Economics* 54, 5–44.
- Litzenberger, Robert and Krishna Ramaswamy, 1982, “The Effects of Dividends on Common Stock Prices: Tax Effects or Information Effects?”, *Journal of Finance* 37, 429–443.

- Lochstoer, Lars and Paul Tetlock, 2020, “What drives anomaly returns?”, *Journal of Finance* 75, 1417–1455.
- Merton, Robert C., 1973, “An Intertemporal Capital Asset Pricing Model”, *Econometrica* 41, 867–887.
- Miller, Merton and Myron Scholes, 1982, “Dividends and Taxes: Some Empirical Evidence”, *Journal of Political Economy* 90, 1118–1141.
- Roll, Richard, 1988, “Presidential address: R2. ”, *Journal of Finance* 43: 51-566.
- Rosenberg, Barr, Kenneth Reid, and Ronald Lanstein, 1985, “Persuasive Evidence of Market Inefficiency”, *Journal of Portfolio Management* 11, 9–17.
- Stattman, Dennis, 1980, “Book Values and Stock Returns”, *The Chicago MBA: A Journal of Selected Papers* 4, 25–45.
- Weber, Michael, 2018, “Cash Flow Duration and the Term Structure of Equity Returns”, *Journal of Financial Economics* 128, 486–503.

Table 1: **Cash-flow, Discount-rate, and Variance News for the Market Portfolio**

The table shows the correlations (off-diagonals) and standard deviations (diagonals) of the return shock ($N_{CF} - N_{DR}$), cash-flow news (N_{CF}), negative discount-rate news ($-N_{DR}$), and volatility news (N_V) implied by the VAR model of Appendix Table 1.

News corr/std	$N_{CF} - N_{DR}$	N_{CF}	$-N_{DR}$	N_V
$N_{CF} - N_{DR}$	0.104	0.643	0.952	-0.068
N_{CF}	0.643	0.034	0.378	-0.444
$-N_{DR}$	0.952	0.378	0.086	0.095
N_V	-0.068	-0.444	0.095	0.026

Table 2: **HML: Intra- and Inter-industry Decomposition and CAPM pricing**

The table reports a decomposition of Fama and French's (1993) HML into its intra- and inter-industry components. We first create an HML portfolio within each industry in the same way as HML. We then combine these industry-specific HML portfolios using industry market weights. We define industries using Fama and French's (1997) mapping of SICC into 48 industries, based on historical classifications from Ken French's website. We scale the resulting composite portfolio so that a full-sample regression of HML on that scaled portfolio, the first regression in the table, has a unit loading. We define the scaled portfolio as HML_{Intra} and the constant and residual from the regression as HML_{Inter} . The sample is 1963Q3-2022Q4, which is then split into two sub-samples in Panels B (1963Q3-1990Q1) and C (1990Q2-2022Q4). We report t -statistics in parentheses.

Panel A: Full-sample estimates					
		constant	HML_{Intra}	HML	$\widehat{R}^2$
(1)	HML	-0.44% (-2.17)	1.00 (27.00)		75.4%
(2)	HML_{Intra}	0.68% (3.96)		0.76 (27.00)	75.4%
(3)	HML_{Inter}	-0.68% (-3.96)		0.24 (8.74)	24.1%
(4)	HML	1.29% (3.28)		-0.17 (-3.86)	5.5%
(5)	HML_{Intra}	1.55% (4.43)		-0.07 (-1.70)	0.8%
(6)	HML_{Inter}	-0.26% (-1.36)		-0.11 (-4.82)	8.6%

Panel B: Pre-1990 estimates					
		constant	HML _{Intra}	HML	$\widehat{R}^2$
(1)	HML	-0.49% (-1.73)	1.00 (17.31)		73.8%
(2)	HML _{Intra}	0.87% (3.75)		0.74 (17.31)	73.8%
(3)	HML _{Inter}	-0.87% (-3.75)		0.26 (5.99)	24.8%
(4)	HML	1.76% (3.77)		-0.25 (-4.65)	16.3%
(5)	HML _{Intra}	2.12% (4.95)		-0.13 (-2.72)	5.7%
(6)	HML _{Inter}	-0.36% (-1.47)		-0.12 (-4.17)	13.4%

Panel C: Post-1990 estimates					
		constant	HML _{Intra}	HML	$\widehat{R}^2$
(1)	HML	-0.40% (-1.37)	1.00 (20.27)		75.9%
(2)	HML _{Intra}	0.54% (2.16)		0.76 (20.27)	75.9%
(3)	HML _{Inter}	-0.54% (-2.16)		0.24 (6.44)	23.7%
(4)	HML	0.82% (1.37)		-0.10 (-1.52)	1.0%
(5)	HML _{Intra}	1.01% (1.92)		-0.01 (-0.12)	-0.8%
(6)	HML _{Inter}	-0.19% (-0.66)		-0.10 (-2.97)	5.7%

Table 3: **Cash-flow, Discount-rate, and Variance Betas**

The table shows the estimated market ($\hat{\beta}$), cash-flow ($\hat{\beta}_{CF}$), discount-rate ($\hat{\beta}_{DR}$), and variance ($\hat{\beta}_V$) betas for HML and its intra- and inter-industry components. The left side of the table runs simple regressions while the right side of the table estimates a multiple regression with all three ICAPM News terms as regressors. The resulting point estimates in both the simple and multiple ICAPM regressions are scaled as in Campbell, Giglio, Polk, and Turley (2018). The sample is 1963Q3-2022Q4, which is then split into two sub-samples in Panels B (1963Q3-1990Q1) and C (1990Q2-2022Q4). We report t -statistics in parentheses.

		simple regressions				multiple regression			
		$\hat{\beta}$	$\hat{\beta}_{DR}$	$\hat{\beta}_{CF}$	$\hat{\beta}_V$	$\hat{\beta}_{DR}$	$\hat{\beta}_{CF}$	$\hat{\beta}_V$	$\widehat{R}^2$
Panel A: Full Sample									
(1)	HML	-0.16 (-3.55)	-0.24 (-6.77)	0.09 (6.42)	-0.11 (-10.12)	-0.27 (-7.95)	0.09 (7.50)	-0.05 (-4.66)	47.8%
(2)	HML _{Intra}	-0.07 (-1.67)	-0.14 (-4.35)	0.08 (6.70)	-0.09 (-9.24)	-0.15 (-4.78)	0.07 (6.21)	-0.05 (-4.68)	37.3%
(3)	HML _{Inter}	-0.09 (-4.22)	-0.10 (-5.49)	0.01 (1.20)	-0.02 (-3.24)	-0.11 (-5.34)	0.02 (2.48)	0.00 (-0.32)	13.6%
Panel B: Pre-1990									
(1)	HML	-0.25 (-4.74)	-0.30 (-6.99)	0.04 (2.49)	-0.10 (-8.14)	-0.27 (-4.75)	0.07 (4.48)	-0.04 (-2.16)	50.4%
(2)	HML _{Intra}	-0.14 (-2.95)	-0.19 (-4.74)	0.05 (3.20)	-0.07 (-6.47)	-0.17 (-3.06)	0.06 (4.19)	-0.03 (-1.99)	37.7%
(3)	HML _{Inter}	-0.11 (-3.90)	-0.11 (-4.41)	0.00 (-0.41)	-0.03 (-3.41)	-0.10 (-2.65)	0.01 (0.62)	0.00 (-0.36)	13.9%
Panel C: Post-1990									
(1)	HML	-0.07 (-1.03)	-0.20 (-3.54)	0.13 (6.47)	-0.13 (-6.72)	-0.28 (-6.01)	0.13 (6.09)	-0.06 (-3.16)	47.0%
(2)	HML _{Intra}	0.01 (0.19)	-0.10 (-1.95)	0.11 (6.33)	-0.11 (-6.60)	-0.15 (-3.53)	0.09 (4.83)	-0.06 (-3.32)	37.1%
(3)	HML _{Inter}	-0.08 (-2.50)	-0.10 (-3.67)	0.02 (1.68)	-0.02 (-1.69)	-0.12 (-4.25)	0.03 (2.40)	0.00 (-0.02)	13.0%

Table 4: **Other Empirical Factors: ICAPM Betas**

The table shows the estimated market ($\hat{\beta}$), cash-flow ($\hat{\beta}_{CF}$), discount-rate ($\hat{\beta}_{DR}$), and variance ($\hat{\beta}_V$) betas for other well-known empirical factors. Panel A reports the results for the intra- and inter-industry components of the large-cap and small-cap components of HML while Panel B reports the results for the other Fama-French-Carhart factors (SMB, RMW, CMA, and MOM), for the FM factor of Goncalves and Leonard (2023), and for two short duration trading strategies. The sample is 1963Q3-2022Q4 except for the two short duration trading strategies. The strategy labeled $DUR_{Goncalves}$ is from Goncalves (2021) and covers the sample period 1973Q3-2022Q4. The strategy labeled DUR_{Weber} is from Weber (2018) and covers the sample period 1963Q3-2014Q2. The left side of the table runs simple regressions while the right side of the table estimates a multiple regression with all three ICAPM News terms as regressors. The resulting point estimates in both the simple and multiple ICAPM regressions are scaled as in Campbell, Giglio, Polk, and Turley (2018). We report t -statistics in parentheses.

Panel A: HML Decomposition: Large-cap and Small-cap Components

		simple regressions					multiple regression			
		$\hat{\beta}$	$\widehat{R^2}$	$\hat{\beta}_{DR}$	$\hat{\beta}_{CF}$	$\hat{\beta}_V$	$\hat{\beta}_{DR}$	$\hat{\beta}_{CF}$	$\hat{\beta}_V$	$\widehat{R^2}$
(1)	HML_{Intra}^{Large}	-0.05 (-1.32)	0.3%	-0.13 (-3.73)	0.07 (6.14)	-0.08 (-7.27)	-0.15 (-4.31)	0.07 (5.78)	-0.04 (-3.12)	28.6%
(2)	HML_{Inter}^{Large}	0.01 (0.41)	-0.4%	-0.02 (-0.68)	0.03 (3.20)	-0.03 (-3.49)	-0.02 (-0.59)	0.02 (2.18)	-0.02 (-1.97)	5.7%
(3)	HML_{Intra}^{Small}	-0.07 (-1.67)	0.7%	-0.15 (-4.10)	0.07 (5.91)	-0.09 (-8.90)	-0.15 (-4.03)	0.06 (5.10)	-0.06 (-4.79)	32.7%
(4)	HML_{Inter}^{Small}	-0.20 (-6.99)	16.8%	-0.22 (-8.55)	0.00 (0.23)	-0.02 (-2.17)	-0.27 (-9.03)	0.03 (3.46)	0.02 (1.92)	26.4%

Panel B: Other Empirical Factors

		simple regressions					multiple regression			
		$\widehat{\beta}$	$\widehat{R^2}$	$\widehat{\beta}_{DR}$	$\widehat{\beta}_{CF}$	$\widehat{\beta}_V$	$\widehat{\beta}_{DR}$	$\widehat{\beta}_{CF}$	$\widehat{\beta}_V$	$\widehat{R^2}$
(1)	SMB	0.25 (6.49)	14.8%	0.18 (4.76)	0.08 (7.18)	-0.02 (-1.34)	0.12 (2.95)	0.07 (5.15)	-0.01 (-0.38)	20.3%
(2)	RMW	-0.09 (-2.96)	3.2%	-0.09 (-2.94)	-0.01 (-1.34)	0.00 (-0.29)	-0.09 (-2.57)	0.00 (-0.24)	0.01 (0.47)	2.5%
(3)	CMA	-0.18 (-6.15)	13.5%	-0.22 (-8.82)	0.03 (2.80)	-0.05 (-6.12)	-0.25 (-9.13)	0.04 (5.00)	-0.01 (-1.33)	35.7%
(4)	MOM	-0.21 (-3.88)	5.6%	-0.14 (-2.80)	-0.07 (-4.58)	0.05 (3.36)	-0.17 (-2.93)	-0.04 (-1.99)	0.06 (3.13)	11.7%
(5)	F-M	-0.24 (-3.92)	7.3%	-0.26 (-4.83)	0.02 (0.93)	-0.02 (-1.26)	-0.33 (-5.32)	0.05 (2.52)	0.03 (1.56)	13.0%
(6)	DUR _{Goncalves}	-0.35 (-5.21)	11.9%	-0.39 (-6.71)	0.03 (1.31)	-0.04 (-2.09)	-0.48 (-7.38)	0.08 (3.56)	0.03 (1.39)	22.8%
(7)	DUR _{Weber}	-0.43 (-5.73)	13.6%	-0.49 (-7.52)	0.04 (1.63)	-0.12 (-5.28)	-0.50 (-6.70)	0.07 (3.09)	-0.03 (-1.29)	27.2%

Table 5: **Stochastic Discount Factor Regressions**

The table reports the estimated SDF beta ($\hat{\beta}_{SDF}$) and the associated unadjusted $\hat{R}^2$ for the log returns of HML and its intra- and inter-industry components. The ICAPM SDF is $m_{t+1} - E_t m_{t+1} = -\gamma N_{CF,t+1} - [-N_{DR,t+1}] + \frac{1}{2}\omega N_{V,t+1}$ where $\gamma=7.2$ and $\omega=24.9$, as estimated in Table 4 column (3) of Campbell, Giglio, Polk, and Turley (2018). The ICAPM-implied Sharpe ratio is proportional to the square root of the $\hat{R}^2$. The coefficient of proportionality is the standard deviation of the SDF, multiplied by the negative of the sign of the regression coefficient. The volatility of the SDF is 0.4238, 0.3357, and 0.4849 in Panel A (1963Q3-2022Q4), Panel B (1963Q3-1990Q1), and Panel C (1990Q2-2022Q4), respectively. We report t -statistics in parentheses.

		$\hat{\beta}_{SDF}$	$\hat{R}^2$	Implied Sharpe Ratio
Panel A: Full Sample				
(1)	HML	-0.07 (-9.43)	27.4%	0.22
(2)	HML _{Intra}	-0.06 (-9.48)	27.6%	0.22
(3)	HML _{Inter}	-0.01 (-1.96)	1.6%	0.05
Panel B: Pre-1990				
(1)	HML	-0.08 (-6.37)	27.9%	0.18
(2)	HML _{Intra}	-0.07 (-6.22)	26.9%	0.17
(3)	HML _{Inter}	-0.01 (-1.56)	2.3%	0.05
Panel C: Post-1990				
(1)	HML	-0.07 (-7.09)	28.0%	0.26
(2)	Intra HML	-0.06 (-7.26)	29.0%	0.26
(3)	Inter HML	-0.01 (-1.31)	1.3%	0.06

Table 6: **Repricing from Cash-flow, Discount-rate, and Variance Beta Changes**

The table measures the change in ICAPM betas and the resulting change in ICAPM expected return from 2006Q4 to 2022Q4. Panel A reports discount-rate ($\hat{\beta}_{DR}$), cash-flow ($\hat{\beta}_{CF}$), and variance ($\hat{\beta}_V$) betas for HML, HML_{Intra}, and HML_{Inter} as of 2006Q4 and 2022Q4. We estimate these betas using monthly observations over a five-year window (see Figure 10 for the evolution of these betas over the 199002-202212 sample). Panel B reports the change in the average return in the 2007Q1-2022Q4 subsample relative to the 199002-202212 full sample as well as the implied change in the ICAPM expected return based on the ICAPM beta changes in Panel A. We compute the average quarterly repricing return over the 2006Q4-2022Q4 sample that is implied by that change, assuming the entire change was unanticipated. We map ICAPM betas to ICAPM risk premium using the estimates from Table 4 column (3) of Campbell, Giglio, Polk, and Turley (2018). We report t -statistics in parentheses.

Panel A: ICAPM Beta Changes

		$\widehat{\beta}_{DR}$		$\widehat{\beta}_{CF}$		$\widehat{\beta}_V$	
		2006	2022	2006	2022	2006	2022
(1)	HML	-0.07 (-1.29)	-0.03 (-0.43)	0.04 (3.40)	0.14 (5.85)	0.00 (0.16)	-0.10 (-5.39)
(2)	HML _{Intra}	-0.01 (-0.19)	-0.03 (-0.40)	0.04 (3.53)	0.12 (5.26)	0.01 (0.89)	-0.08 (-4.87)
(3)	HML _{Inter}	-0.06 (-1.42)	0.00 (-0.11)	0.00 (-0.11)	0.02 (1.26)	-0.01 (-0.93)	-0.01 (-1.20)

Panel B: Repricing Implications

		Average quarterly returns		Pricing the Δ in ICAPM Betas	
		Full	Post 2006	$\Delta \mu_{ICAPM}^{quarterly}$	Average quarterly repricing return
(1)	HML	1.37% (2.00)	-1.74% (-1.78)	1.51%	-1.89%
(2)	HML _{Intra}	1.79% (3.28)	-2.05% (-2.63)	1.32%	-1.61%
(3)	HML _{Inter}	-0.42% (-1.01)	0.31% (0.52)	0.20%	-0.28%




Figure 1: The figure plots the log return on HML (solid black line) and r_M , the log real return on the market (dashed red line), for the Compustat period from 1963Q3-2022Q4. The series are smoothed with a trailing exponentially weighted moving average in which the decay parameter is set to 0.08 per quarter, and the smoothed series is generated, for example, as $MA_t(\text{HML}) = 0.08\text{HML}_t + (1-0.08)MA_{t-1}(\text{HML})$. This decay parameter implies a half-life of two years. The vertical line indicates the start of the subperiod 1990Q2-2022Q4.




Figure 2: We plot the results from the Table 2 decomposition of HML into its intra- and inter-industry components for the Compustat period from 1963Q3-2022Q4. The solid black line shows the smoothed log return to Fama and French's (1993) HML; the dashed green line shows the smoothed log return to the intra-industry component of HML; and the dashed-dotted red line shows the smoothed log return to the inter-industry component. The series are smoothed with a trailing exponentially weighted moving average in which the decay parameter is set to 0.08 per quarter, and the smoothed series is generated, for example, as $MA_t(\text{HML}) = 0.08\text{HML}_t + (1-0.08)MA_{t-1}(\text{HML})$. This decay parameter implies a half-life of two years. The vertical line indicates the start of the subperiod 1990Q2-2022Q4.




Figure 3: We plot smoothed ICAPM News for the Compustat subsample from 1963Q3-2022Q4 as generated by the VAR in Appendix Table 1. The solid black line shows the negative of smoothed aggregate discount-rate news; the dashed green line shows smoothed aggregate cash-flow news; and the dashed-dotted red line shows smoothed aggregate variance news. The series are smoothed with a trailing exponentially weighted moving average in which the decay parameter is set to 0.08 per quarter, and the smoothed series is generated, for example, as $MA_t(N^{CF}) = 0.08N_t^{CF} + (1-0.08)MA_{t-1}(N^{CF})$. This decay parameter implies a half-life of two years. The vertical line indicates the start of the subperiod 1990Q2-2022Q4.




Figure 4: In Panel A, we explain time-series variation in demeaned HML. Smoothed log returns are plotted with a solid black line while smoothed ICAPM (CAPM) fitted values are plotted with a dashed red (dotted green) line. In Panel B, we plot the components of the ICAPM fit for HML. The solid black line is the smoothed ICAPM fit from Panel A; the dashed green line shows the smoothed contribution of N_{DR} ; the dashed blue line shows the smoothed contribution of N_{CF} ; and the dashed red line shows the smoothed contribution of N_V .




Figure 5: We explain the returns on the six subsets of anomalies replicated by Hou, Xue, and Zhang (2020) for the 1967Q1-2022Q4 period, plotting the resulting cdfs describing the R^2 's distributions for the CAPM (blue line) and the ICAPM (red line).




Figure 6: We plot cumulative ICAPM News for the 2020-01-02 to 2020-12-31 subsample as generated by the quarterly VAR in Appendix Table 1 and daily VAR in Appendix Table X. The solid (dashed) green line shows the negative of accumulated quarterly (daily) N_{DR} ; the solid (dashed) blue line shows accumulated quarterly (daily) N_{CF} ; and the solid (dashed) red line shows accumulated quarterly (daily) N_V .




Figure 7: We plot various cumulative ICAPM fits for HML , HML_{Intra} , and HML_{Inter} from the quarterly and daily VARs for the 2020-01-02 to 2020-12-31 subsample.




Figure 8: We plot the components of the ICAPM fit for HML, HML_{Intra} , and HML_{Inter} for the 2020-01-02 to 2020-12-31 subsample. The solid black line shows the smoothed ICAPM fit; the dashed green line shows the smoothed contribution of N_{DR} ; the dashed blue line shows the smoothed contribution of N_{CF} ; and the dashed red line shows the smoothed contribution of N_V .




Figure 9: The top panel shows the cumulative earnings during 2020 (scaled by total book equity in 2019q4, $ScaledCumulEarn_{t,t+h}^p$) for the long side (black line) and short side (red line) of the inter-industry HML portfolio. Both are residualized to the market earnings. The bottom panel reports HML's weights in each industry at the end of 2019 on the X axis and that industry's contribution to HML's cumulative earnings in 2020 on the Y axis, again residualized to the aggregate earnings.




Figure 10: We plot discount-rate ($\hat{\beta}_{DR}$) (Panel A), cash-flow ($\hat{\beta}_{CF}$) (Panel B), and variance ($\hat{\beta}_V$) betas (Panel C) for HML (solid black line), HML_{Intra} (dashed red line), and HML_{Inter} (dotted green line) that we estimate using monthly observations within five-year rolling windows. The vertical line indicates the end of 2006Q4, since then, the cumulative return to HML has been -50%.