

A Tale of Two Market Returns: The Broad Market Factor and The Idiosyncratic Financial Factor*

Sung Je Byun[†] Johnathan A. Loudis[‡] Lawrence D.W. Schmidt[§]

May 2026

Abstract

We construct a Broad Market Factor (BMF), which is a proxy for the value-weighted equity return on *all* firms in the US economy (public and private). The BMF differs from the standard Value-weighted Market Factor (VMF), which reflects the value-weighted equity return on *public* firms. We define the difference between the VMF and the BMF to be the Idiosyncratic Financial Factor (IFF). The IFF carries no risk premium and is uncorrelated with all macroeconomic proxies for investor marginal utility we consider. CAPM betas and, consequently, discount rates are underestimated when measured with respect to the VMF compared to the BMF for most portfolios. Size factors become redundant and the size anomaly is resolved when the VMF is replaced by the BMF in standard factor models. The intertemporal risk-return relation is substantially stronger when one replaces the VMF with the BMF. The unifying explanation for these results is that the IFF adds unpriced risk to the VMF, distorting both cross-sectional and time-series estimates of exposure to priced market risk.

* *JEL classification*: C15; C58; G12; G17

* *Keywords*: Idiosyncratic risk; Market factor; Size effect; Size factors; Intertemporal risk-return relation

*This is a substantially revised version of an earlier manuscript which was circulated under the title, “Idiosyncratic financial risk and a reevaluation of the market risk-return tradeoff,” which itself was a substantially revised version of an even earlier manuscript circulated under the title, “Real risk or paper risk? Mis-measured factors, granular measurement errors, and empirical asset pricing tests.” We are grateful to John Campbell, John Cochrane, Zhi Da, Xavier Gabaix, Stefano Giglio, Daniel Greenwald, Evan Jo (discussant), Leonid Kogan, Ralph Koijen (discussant), Arvind Krishnamurthy, Eben Lazarus, Gang Li (discussant), Lira Mota, Jonathan Parker, Seth Pruitt, Cesare Robotti (discussant), Andrea Tamoni, Jun Tu (discussant), Yaxuan Wen (discussant), Nancy Xu (discussant) and seminar participants at the American Finance Association, the Northern Finance Association, the Financial Intermediation Research Society, the BI/SHOF asset pricing conference, the China International Conference in Finance, the Eastern Finance Association, the NBER Asset Pricing meeting, the City University of Hong Kong, McGill, MIT, Rutgers, Stanford, the University of Rochester, the University of Southern California, and the University of Wisconsin-Milwaukee for helpful discussions related to the paper. We thank Brice Green, Luxi Han, Maziar Kazemi, Paymon Khorrami, Kyle Kost, and Huben Liu for outstanding research assistance. The views expressed are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas, the Federal Reserve System, or the National Bureau of Economic Research.

[†]Supervisory Policy, Federal Reserve Bank of Dallas, SungJe.Byun@dal.frb.org.

[‡]Mendoza College of Business, University of Notre Dame, jloudis@nd.edu.

[§]MIT Sloan School of Management and NBER, ldws@mit.edu.

1 Introduction

There are more than six million firms in the United States, but only about four thousand are publicly traded.¹ Moreover, firms that go public are not selected at random from the corporate sector: the IPO process favors large firms from industries with high valuations. Even within public firms, the distribution of firm size is highly skewed towards large firms. Thus, the aggregate stock market return is heavily influenced by a small subset of selected firms. If these overrepresented firms load on a factor that does not affect the aggregate return on all firms or, more generally, the return on investor wealth, this component of the aggregate stock market return is likely not priced (Mayers, 1973; Roll, 1977).

In this paper, we construct an empirical proxy for this factor, which we call the Idiosyncratic Financial Factor (IFF). The IFF is the difference between the standard Value-weighted Market Factor (VMF)—the value-weighted return on *public* firms—and a “Broad Market Factor” (BMF)—an alternative market factor that we construct to reflect the value-weighted return on *all* firms in the US economy, both public and private. The IFF arises because the aggregate return on the selected sample of public firms loads on a systematic factor that diversifies away in the BMF. Hence, the VMF loads on the IFF whereas the BMF does not. Empirically, the IFF’s time series average is indistinguishable from zero, and it is uncorrelated with empirical proxies for investor marginal utility from multiple classes of macro-finance models. Thus, the IFF is a factor that affects returns in financial markets, but is “idiosyncratic” with respect to the broader economy. Our evidence suggests that the IFF adds unpriced risk to the VMF, which distorts standard estimates of priced market risk and creates unintended consequences for asset pricing tests and estimated discount rates.

Our results can be organized into four main contributions. First, we build and calibrate a model that features selection into public markets. Firms that are more or less likely to be selected load differently on an unpriced factor, which generates an IFF. The IFF can be expressed as the difference between the model’s VMF (the value-weighted return on public firms) and its BMF (the value-weighted return on all firms), and it adds unpriced risk to the VMF. The model makes different predictions depending on whether the VMF or BMF

¹The estimate of all firms is based on IRS data aggregated by Kwon et al. (2024) and that of public firms is based on CRSP data, both as of 2018.

is used as a market factor in standard asset pricing tests. Second, we develop a method that uses newly digitized administrative data from Kwon et al. (2024) to construct empirical proxies for the BMF and IFF, and provide evidence consistent with the IFF not being priced. Third, we provide empirical evidence consistent with several model predictions. Most notably, CAPM betas and discount rates are underestimated when measured with respect to the VMF compared to the BMF for most portfolios. Furthermore, size factors become redundant and the size anomaly is resolved when the VMF is replaced by the BMF in standard factor models. Fourth, the intertemporal risk-return relation is substantially stronger when one replaces the VMF with the BMF. Thus, we provide a joint explanation for three long-standing asset pricing “puzzles”: 1) the VMF-CAPM underestimates discount rates for most portfolios, 2) the size anomaly and fact that size factors lead to pricing improvements in VMF-based factor models (even when the size premium is small), and 3) the weak intertemporal risk-return relation implied by the VMF. Our unifying explanation for these puzzles is that the IFF adds unpriced risk to the VMF which distorts estimates of priced market risk in the cross section and time series.

We begin with two motivating facts about firm selection into public markets. First, a small number of large firms accounts for a disproportionate share of public equity by market value. For instance, in 2020, the largest five firms in the CRSP universe comprised about 20% of the total market value. Second, using sales as an observable size metric for all firms, we show that the public equity market overrepresents big firms and underrepresents small firms relative to their sales shares among all firms (public and private). Thus, small firms collectively generate a much higher fraction of aggregate economic output than implied by their share in the public equity market. Size is the most salient dimension along which public firms differ from all firms, but we also show that there are persistent and dynamic differences in industry composition. This compositional disconnect between public and all firms could cause a wedge between the value-weighted return on public versus all firms.

Motivated by this evidence, we build and calibrate a static representative agent model that features selection into public markets where the probability of going public increases with firm size (though this could be thought of as a proxy for any dimension along which public and private firms may differ). Thus, as in the data, big firms are overrepresented in

the public market. We also assume—and later provide supporting empirical evidence—that there is an unpriced factor on which big and small firms differentially load. The combination of selection and differential loadings on the unpriced factor contaminates the VMF with unpriced risk. This unpriced factor – the IFF – captures a wedge between the VMF and the BMF (the true market factor in the model). The model generates a number of predictions which we test empirically. The predictions also hold in an extension with heterogeneous agents and incomplete markets under conditions which we fail to reject in the data.

The main challenge in building an empirical proxy for the BMF is that private firms’ market values and returns are unobserved. To address this, we develop a methodology for reweighting public firm returns using sales and industry information so that their aggregate approximates the value-weighted return on all firms. This methodology generates a BMF with lower weights on large firms and higher weights on small firms relative to their value weights in the VMF. As in our model, we compute the empirical IFF as the difference between the VMF and BMF. We show our results are robust to using other BMF aggregation methods that downweight large firms and upweight small firms relative to their value weights, implying that our results are not particularly sensitive to various implementation choices.

We next test a key modeling assumption: the IFF has no risk premium. Most directly, its time series average is an economically small and statistically insignificant -0.51% per year. Furthermore, and consistent with our model, in most of the multi-factor models we consider the IFF generates a negative alpha when the VMF is used as the market factor, and no alpha when the BMF is used. An equivalent interpretation of these results is that the VMF is spanned by models that use the BMF as the market factor, but the BMF is not spanned by models that use the VMF.

We provide additional empirical evidence that the IFF is not priced by estimating local projections of a number of variables linked to investor marginal utility in different classes of asset pricing models onto both the BMF and the IFF. Responses to the BMF are large, statistically significant, and have signs consistent with what we would expect from theory. By contrast, variables show no systematic response to the IFF, and coefficients are generally insignificant or have the opposite sign compared to those from the BMF. These results are consistent with our claim that the IFF is not priced and our notion that it is “idiosyncratic”

with respect to the broader economy.

Consistent with our model, VMF betas are typically lower than BMF betas as are the associated CAPM-implied discount rates, especially for portfolios with smaller market capitalizations and higher betas. These differences can be large: The cross-sectional standard deviation of VMF–BMF beta differences is 0.13, which is about two-thirds of the cross-sectional standard deviation of VMF betas themselves (0.19). These differences matter for pricing errors: average alphas under the BMF are effectively zero, whereas they are positive under the VMF across a broad set of test portfolios. Equivalently, while the VMF implies discount rates that are too low on average, the BMF implies discount rates that are correct on average. Using price wedges (van Binsbergen et al., 2023) as an alternative metric, we again find that the BMF implies significantly lower mispricing than the VMF.

Our model provides a simple explanation for the size anomaly and the role of size factors in standard factor models (e.g., Banz, 1981; Fama and French, 1992, 1993). When the VMF contains unpriced risk mostly related to the returns on big firms (the IFF), it underestimates most other stocks’ exposure to priced market risk (the BMF), causing small stocks to earn positive VMF alphas. Size factors, which reflect returns on strategies that buy small stocks and sell large ones, act as control variables that separate priced BMF risk from unpriced IFF risk and align VMF-implied risk premia with the correct BMF-implied risk premia.

We investigate these model-implied connections between the IFF and size in two ways. First, consistent with prior literature, we show that standard VMF-based factor models yield positive alphas on all size-sorted portfolios below the top decile and that these alphas increase as portfolio size decreases. However, when the VMF is replaced by the BMF these alphas and their relationship with size are effectively eliminated. Second, and again consistent with the literature, size factors are not spanned by VMF-based factor models, which is part of the rationale for including size factors in these models in the first place. However, when the VMF is replaced by the BMF, size factors are spanned in all the factor models we study. In other words, size factors improve a model’s ability to price assets when the VMF is used as the market factor, but not when the BMF is used (Barillas and Shanken, 2017).

Finally, we conduct a test in the time series. If the IFF is not priced, it contributes to the VMF’s conditional variance but not its risk premium, weakening and destabilizing the risk-

return relation. Consistent with this intuition, we find that the BMF generates a stronger risk-return relation than the VMF. Replicating prior results, the coefficient relating VMF returns to its conditional variance is positive but insignificant (e.g., Nelson, 1991; Glosten et al., 1993; Koopman and Hol Uspensky, 2002). When using the BMF, the coefficient is two to four times larger and is statistically significant across all specifications we investigate.

Contribution to the literature: Our paper contributes to the long literature on mis-measured market risk factors dating back to Roll’s (1977) critique. Researchers have sought to address this critique by augmenting the standard VMF with observable proxies for omitted components of wealth.² Our approach, closer in spirit to Daniel et al. (2020), instead removes an unpriced component from the VMF.³ Building on Mayers (1973), we highlight a form of mismeasurement that arises when investors hold substantial wealth outside public equity markets, particularly in private firms, and motivate how this can contaminate the VMF with an unpriced component (the IFF), which we remove to form the BMF.

Regardless of the alternative market factor constructions noted above, the VMF-CAPM remains a popular model for estimating discount rates in practice.⁴ This is despite the fact that it produces a security market line that is “too flat” (e.g., Black et al., 1972; Fama and MacBeth, 1973; Fama and French, 1992). A related but distinct issue is that for most long-only portfolios, the CAPM implies discount rates that are too low compared to average returns. Equivalently, most portfolios generate positive CAPM alphas. For example, this is true for industry-sorted portfolios, size-sorted portfolios (see Gibbons et al., 1989, for both), book-to-market sorted portfolios (Fama and French, 2004), and size by book-to-

²Examples include other asset classes (Stambaugh, 1982), human capital (e.g., Mayers, 1973; Campbell, 1996; Jagannathan and Wang, 1996; Heaton and Lucas, 2000; Campbell et al., 2016), and aggregate financial and non-financial holdings of households (Cederburg and O’Doherty, 2019).

³Unlike Daniel et al. (2020), who study disconnects between characteristics-sorted portfolios and asset covariances, we focus on a specific composition issue in the VMF. Our adjustment is both conceptually and empirically distinct: The BMF (IFF) has a correlation of only 39% (47%) with their corrected market factor (i.e., their corrected market factor is far from perfectly correlated with the BMF and it loads on the IFF).

⁴For example, Graham and Harvey (2001) report that 73.5% of CFOs use the CAPM to estimate the cost of capital. More recently, Graham (2022) shows that the CAPM is still the most common model that firms use to estimate discount rates, although its use among small firms has diminished compared to findings in Graham and Harvey (2001). Our results provide one potential explanation for this finding: small-cap firms are the ones whose CAPM-implied discount rates are most distorted by the IFF. Relatedly, Gormsen and Huber (2024) show that CAPM beta and size are two of the strongest predictors of firms’ perceived costs of capital. Additionally, Decaire and Graham (2024) report that among sell-side analysts who report their method for computing discount rates, 97% use the CAPM to estimate when valuing companies they follow.

market sorted portfolios (Fama and French, 1993). We provide similar evidence across a broad set of portfolios and offer a simple explanation: the average underpricing is the result of mismeasured market risk, which the BMF resolves. To be clear, we do not claim to “resurrect” the CAPM. For example, although the BMF-CAPM generates smaller alphas than the VMF-CAPM, it does not eliminate them in the broad set of portfolios we study. However, it does imply that they are zero on average, compared to VMF-CAPM alphas which are positive on average. Thus, the BMF-CAPM helps resolve the VMF-CAPM’s tendency to underestimate discount rates *on average*.

Starting with Fama and French (1993), another tack to address issues with the VMF-CAPM is to incorporate additional characteristics-based factors into pricing models. Most factor models include at least two: a value-weighted market factor and a size factor that buys small stocks and sells big ones. There are still debates about what economic mechanism generates the size effect and whether size factors earn positive risk premia.⁵ Regardless, factor models include size factors because they improve pricing power, even in the post-1993 sample in which the size effect is relatively small. Our novel explanation for this practice operates regardless of the magnitude of the size premium: because size factors also load on the IFF, they serve as control variables for unpriced IFF risk that contaminates the VMF. Most other explanations of size revolve around small firms being exposed to extra priced risks, but our explanation is that small firm market risk is underestimated when measured with respect to the IFF-contaminated VMF. In line with this mechanism, we provide evidence that size-sorted portfolios are correctly priced and that size factors are spanned—and therefore redundant—in factor models when one replaces the VMF with the BMF.

We also contribute to the literature that studies the intertemporal risk-return relation. Many studies have focused on various approaches to better estimate conditional variance, and it has been found that results are sensitive to modeling choices.⁶ Irrespective of variance

⁵See van Dijk (2011) and Alquist et al. (2018) for summaries of this debate. More recently, Asness et al. (2018) and Hou and van Dijk (2019) provide evidence that the size effect can be resurrected if one controls for firm quality or profitability shocks, respectively. See also Berk (1995) for a general argument for why a size anomaly is likely to emerge in the presence of omitted sources of risk. Furthermore, even if the size anomaly is small, it can still have a large effect on allocational efficiency in the real economy if it is persistent (van Binsbergen and Opp, 2019).

⁶Using GARCH-type volatility forecasting models, for example, Bollerslev et al. (1988), Chou (1988), Baillie and DeGennaro (1990), Campbell and Hentschel (1992), and Chou et al. (1992) find a positive risk-

measurement issues, our critique relates to misspecification in the VMF itself and applies even when a researcher has access to a “perfect” measure of conditional VMF variance. Accordingly, we find that the BMF, which purges the VMF unpriced IFF risk, yields a stronger risk-return relation regardless of how we measure conditional variance.

2 Data

All Firms: Our analyses makes use of novel data recently collected and digitized from historical IRS publications by Kwon et al. (2024, hereafter KMZ). The KMZ data covers all firms in the US economy (both public and private), and provides information about the number of firms and total sales by groups of firms segmented into seven SIC-based industries and firm sales bins from 1959–2013.⁷ From 2013–2018, only a breakdown by sales is available. These data help us assess the representativeness of public firms for the set of all firms, and are a key input for constructing our BMF.

Publicly Traded Firms: For listed firms, we use standard data from the CRSP/Compustat merged database. We obtain monthly return and market capitalization measures for public firms (common shares on the NYSE, AMEX, or NASDAQ) from CRSP, along with annual sales and other financial information from Compustat. We use CRSP SIC industry codes to match with KMZ’s SIC-based industry designations. Our data begins July 1963—when Compustat financials become available—and ends in December 2021. Our test assets are 382 value-weighted characteristics-sorted portfolios from Ken French’s Data Library. These frequently-studied portfolios are listed in Online Appendix Table OA.3.

3 Motivating evidence

In this section, we provide evidence of a disconnect between the composition of public firms and that of all firms (public and private) in the US.

return relation, whereas Nelson (1991) and Glosten et al. (1993) find a negative relation when using extended GARCH models with so-called leverage effects. Using alternative approaches to GARCH forecasting models, French et al. (1987), Harvey (1989), Ludvigson and Ng (2007) and Ghysels et al. (2005) find a positive relation, while Pagan and Hong (1991) and Whitelaw (1994) find a negative relation.

⁷The industries are mining, construction, manufacturing, utilities, trade, finance, and services. These correspond to the following two-digit SIC codes, respectively: 10–14, 15–17, 20–39, 40–49, 50–59, 60–67, and 70–89. We do not include an eighth industry studied by KMZ (agriculture, SIC codes 01–09) due to the scarcity of publicly-listed agriculture firms (only 0.65% of observations).

3.1 Public equities are dominated by a few large firms

Popular news media is rife with stories highlighting concerns that recent stock market performance is dominated by a few mega-cap companies such as the “Magnificent Seven”. As of December 2023, these seven tech stocks made up more than 30% of the S&P 500 Index and had a larger weight in the MSCI All Country World Index than *all of its constituents from Japan, France, China, and the UK*.⁸ This concern is not unique to the present. For instance, the “Nifty Fifty” dominated the market in the late 1960s and early 1970s. Often, as is the case for the Magnificent Seven, these firms cluster in related industries.

Figure 1 plots the value share (by market capitalization) of the largest 5, 10, 25, and 50 stocks in the CRSP universe of public firms. From 1926 to 2021, the largest five stocks are associated with between 10 and 25 percent of overall market value. The combined weight of the top 25 stocks ranges from a low of approximately 20% during the 1980s to a high of almost 50% in the 1930s. These shares decline somewhat over the sample in part because the CRSP universe expands, as is clear from the number of constituent stocks plotted on the right axis. Nevertheless, consistent with the anecdotes above, the market is highly concentrated throughout the sample, and it is striking that one can barely detect any changes in these shares even when the number of stocks included in the index jumps dramatically upward.

3.2 Public firm composition differs from that of all firms

The market value concentration among public firms highlighted above likely overstates concentration in the overall economy since it neglects the value of private firms. To investigate this issue, we compare the composition of public and all firms using observable sales data from CRSP/Compustat and KMZ, respectively. These data imply that, on average, public firms comprise about 40% of total U.S. firm sales during our sample period, leaving scope for a compositional disconnect between the two sets of firms.

Figure 2 plots time-series average sales and value shares for sales-sorted decile portfolios

⁸Source: “It’s the Magnificent Seven’s Market. The Other Stocks Are Just Living in It” from December 17, 2023. The article also highlights how the aggregate price-earnings ratio of the S&P 500 is heavily influenced by the above-average multiples of these seven stocks. This concentration is also concerning to practitioners. An October 2025 article from the Wall Street Journal titled “The Warning Signs Lurking Below the Surface of a Record Market” stated “‘The strength in megacap tech has obfuscated what appears to be increasing signs of weakness through much of the real economy,’ said Bob Elliott, chief executive at Unlimited Funds. ‘And the pricing of that weakness has accelerated in recent weeks.’”

of public firms using NYSE breakpoints, and compares them with matched portfolios representing all firms. Red and blue bars show market value and sales shares, respectively, for *public* firms as fractions of total public firm market value or sales. Both shares are similar within each portfolio and rise with portfolio number. Thus, concentration among public firms by size is similar whether we measure it using value or sales shares, and these shares are mostly concentrated among large firms. For example, public firms in the top sales decile account for 56% of total public firm sales, while firms in the bottom decile only account for 2.0% of sales. Thus, mirroring the market value concentration in Figure 1, total public firm sales is concentrated in the subset of largest firms.

Yellow bars show the sales shares for *all* firms that fall in each bin as fractions of total sales among all firms. In contrast to the public firm shares, these follow a U-shaped pattern across size deciles. The sales share of the bottom decile is 30% and that of the top decile is 29%, indicating small firms carry larger, and large firms smaller, sales shares than those implied by public firms alone. Thus, small firms represent a larger share of the overall economy, and large firms a smaller one, than one would infer from public firms' sales shares. We also find differences in industry-level sales shares between public and all firms, though these are less pronounced than those based on size. As an example, manufacturing firms are particularly overrepresented in the public markets in the early part of the sample. For brevity, we relegate these industry-based comparisons to Online Appendix OA.2.

Relation to mechanisms from the literature on the firm size distribution: The above compositional differences are not surprising in light of evidence from the literature on firm dynamics. It is well documented that the distribution of firm sizes has a fat right tail (Axtell, 2001), and multiple theoretical mechanisms have been proposed to explain it (see, e.g., Luttmer, 2010; Gabaix, 2016, for surveys). Both Autor et al. (2020) and Kwon et al. (2024) emphasize how an increasing share of output is generated by a small set of large, “superstar firms”. Given the large scale of these firms' operations, most superstars—often the most productive producers—will likely value access to public capital markets more than their smaller competitors. In addition, the strength of these forces is heterogeneous across industries, so some industries are also likely to be overrepresented in the public markets.

The evidence we document above shows that the public equity market overrepresents a

selected subset of all firms, motivating a mechanism that could generate the IFF. Specifically, the VMF may include a component relevant for firms that select into public markets that diversifies away in broader measures of wealth (such as the BMF). In the remainder of the paper, we explore evidence related to this mechanism both theoretically and empirically.

4 Theoretical model

We next construct and calibrate a model to highlight our proposed mechanism that generates the IFF, and to build testable predictions in a well-defined setting. Motivated by results in the previous section, the model features two firm types with one overrepresented in the public market, and with heterogeneity in firms’ exposures to systematic risk factors. We provide an outline of the model here and relegate details to Online Appendix OA.3.

4.1 Cash flow risks

In our economy, there are three sources of cash flow risk defined as

$$f \sim \mathcal{N}(0, \sigma_f^2), \quad g \sim \mathcal{N}(0, \sigma_g^2), \quad \text{and} \quad z \sim \mathcal{N}(0, \sigma_z^2) \quad \text{with} \quad f \perp g \perp z. \quad (1)$$

Given assumptions we outline below, f is the only risk factor that drives uncertainty about the aggregate dividend, whereas g and z are additional factors that redistribute cash flows between different firm types.

Firms are indexed by i and partitioned into two types, $\tau_i \in \{S, B\}$. These types could represent any dimension over which public and all firms differ, but we refer to them as “small” and “big” firms since this is the most salient dimension over which they differ in the data.⁹ There are a total of N_S small firms, N_B big firms, and $N = N_S + N_B$ total firms with $N_S > N_B$, as in the data. Let $\theta = N_S/N$ be the fraction of S firms. Consistent with our motivating evidence, we assume that the probability a small firm goes public, p_S^{pub} , is less than the probability a large firm goes public, p_B^{pub} . Each firm has an additional characteristic, $X_i \in \{-1, 1\}$, which serves as a simple stand-in for loadings on other systematic factors that may cause the CAPM to fail. We assume $X_i \perp \tau_i$ for simplicity, although this condition can

⁹Although our main focus is on the size dimension, we explore the relationship between the BMF, IFF, and other common characteristics such as value at the end of Section 8.3.

be relaxed with similar implications. Let each firm i 's dividend be

$$d_i = \begin{cases} \frac{1}{N} [1 + \lambda_S \bar{d} + \gamma_S f + \lambda_S g + \lambda_S X_i z] & \text{if } \tau_i = S \\ \frac{1}{N} [1 + \lambda_B \bar{d} + \gamma_B f + \lambda_B g + \lambda_B X_i z] & \text{if } \tau_i = B \end{cases} \quad (2)$$

where we define

$$\lambda_S = -\frac{1 - \theta}{\theta} \quad \text{and} \quad \lambda_B = 1 \quad (3)$$

so that neither \bar{d} , g , nor z affect the aggregate dividend, which is given by

$$D_1 = 1 + \bar{\gamma} f, \quad \text{where} \quad \bar{\gamma} \equiv \theta \gamma_S + (1 - \theta) \gamma_B. \quad (4)$$

Furthermore, $\bar{d} > 0$ so expected dividends are larger for big firms than for small firms. We abstract away from adding additional idiosyncratic firm-level shocks for simplicity.¹⁰

Interpretation of the redistributive shock g : Without loss of generality, we can always define the factors in a multivariate model such that aggregate output only loads on the first factor (f) and not on the other factors. In our model, g stands in for any force that redistributes cash flows between firms that are underrepresented and those that are overrepresented in the public market without affecting f . Empirically, firms that are overrepresented are likely to be those for which the benefits of being listed exceed the costs, a feature of the stylized model we develop to explicitly microfound our empirical BMF construction methodology (see Section 5 and Online Appendix OA.5). Such firms could include, for example, large firms that benefit from economies of scale (e.g., “superstar firms,” consistent with our motivation in Section 3), firms with higher valuation multiples that have more valuable growth options (see, e.g., Kim and Weisbach, 2008; Chemmanur et al., 2010), or firms from industries with greater listing benefits (e.g., capital-intensive manufacturing firms may benefit more than service firms). See Online Appendix OA.5.2 for additional motivation regarding the existence of redistributive shocks among firms along these dimensions. Thus, the *IFF* (a scaled version of g in the model, which we define below) also reflects a redistributive shock between under- and overrepresented firms.

¹⁰Implicitly, we are assuming that either N_S and N_B are sufficiently large for the law of large numbers to approximately hold for an average of idiosyncratic shocks, or, equivalently, we can interpret N_S and N_B as measures over a continuum of atomistic firms.

4.2 Preferences

There is a unit mass of ex-ante identical agents, each of whom is endowed with one share of every firm, with preferences given by

$$U_0 = -\exp(-\alpha C_0) - \rho \mathbb{E}_0 [\exp(-\alpha C_1 - \omega z)] \quad \text{with} \quad C_1 = W_1 = W_0 R_w \quad (5)$$

where C_0 is initial consumption, α is absolute risk aversion, $W_0 > 0$ is initial wealth, and R_w is the return on wealth (i.e., the value-weighted return on all firms, public and private). Agents optimally consume their entire endowments (i.e., there is no savings technology), so $C_t = D_t$.¹¹ The additional z shifter has a price of risk given by ω and acts as a stand-in for any other force that causes the consumption CAPM to fail. This feature tractably captures a property which frequently emerges in macrofinance models: each investor's intertemporal marginal rate of substitution depends on consumption growth (which depends only on f , as shown in equation 4) as well as the shifter z which captures other priced factors (see, e.g., Cochrane, 2017). In our empirical implementations below, we use various characteristics-based factors from the literature to capture the economic forces represented by z .

4.3 Risk premia, factor returns, and betas

In equilibrium, optimality implies that the expected excess return on any firm i is given by

$$\mathbb{E}[R_i] - R_f = \beta_i^w \mathbb{E}[R_w - R_f] + \beta_i^z \mathbb{E}[R_z - R_f], \quad (6)$$

where R_f is the risk-free rate, R_z is the return on a tradable factor-mimicking portfolio for z , and β_i^w and β_i^z are betas with respect to R_w and R_z .¹² R_w and R_z satisfy

$$R_w - R_f = \frac{R_f (\alpha \bar{\gamma}^2 \sigma_f^2 + \bar{\gamma} f)}{1 - \alpha \bar{\gamma}^2 \sigma_f^2} \quad \text{and} \quad R_z - R_f = \omega \sigma_z^2 + z, \quad (7)$$

so they are affine functions of f and z , respectively. Since the two factors R_w and R_z price all assets and both factors are uncorrelated with g , g is not priced in this economy.

¹¹We focus attention on a symmetric equilibrium in which no trade takes place and all agents choose to consume their endowments, so $C_t = D_t$. Since agents will equate their intertemporal marginal rates of substitution state by state, the same optimality conditions will hold for all firms (whether listed or unlisted) regardless of whether markets are complete or unlisted firms are tradable.

¹²See Online Appendix OA.3 for analytical expressions for these betas and all subsequent betas defined in this section. The notational convention we use for betas is that the subscript represents the asset of interest and the superscript represents the factor or asset with respect to which the beta is being measured.

The aggregate value-weighted return on *public* firms, the *VMF* in the model, is

$$VMF = \underbrace{\beta_{VMF}^w (R_w - R_f)}_{\equiv BMF} + \underbrace{\beta_{VMF}^g g}_{\equiv IFF}, \quad (8)$$

where $\beta_{VMF}^w > 0$ is the model-implied VMF's loading on R_w and $\beta_{VMF}^g > 0$ is its loading on g . Whereas the BMF only loads on f (through its dependence on R_w), the VMF also loads on g since big firms are overrepresented in the public market. The IFF is orthogonal to the BMF and has no risk premium (i.e., a mean of zero) because it is proportional to g .

We assume throughout that some firms of each type are public, claims to cash flows of all public firms are tradable, and S and B firms' returns are linearly independent. Thus, we can define our model's Small Minus Big factor as

$$SMB \equiv \tilde{R}_S - \tilde{R}_B = \beta_{SMB}^{BMF} BMF + \beta_{SMB}^{IFF} IFF, \quad (9)$$

where \tilde{R}_S reflects the return on a portfolio comprised primarily of S-type firms and \tilde{R}_B reflects the return on a portfolio comprised primarily of B-type firms.¹³ As the second equality indicates, we can express this factor as a linear combination of the BMF and the IFF. Because SMB only loads on priced BMF and unpriced IFF risk, it is priced by the BMF but not by the VMF (due to the VMF's exposure to the IFF). This feature is important for explaining why including size factors in VMF-based factor models improves model pricing, both in our theoretical model and in our empirical results, which we discuss below.

Due to the VMF's differential loadings on f and g compared to the BMF, betas with respect to each of these factors will differ. In particular, for any asset or portfolio p

$$\beta_p^{BMF} = \frac{\beta_p^w}{\beta_{VMF}^w} \quad \text{and} \quad \beta_p^{VMF} = \beta_p^{BMF} \underbrace{\frac{\sigma_{BMF}^2}{\sigma_{VMF}^2}}_{\substack{\text{attenuation} \\ \text{bias}}} + \beta_p^{IFF} \underbrace{\frac{\sigma_{IFF}^2}{\sigma_{VMF}^2}}_{\substack{\text{additive} \\ \text{bias}}}, \quad (10)$$

and the difference between these betas can be expressed as

$$\beta_p^{VMF} - \beta_p^{BMF} = (\beta_p^{IFF} - \beta_p^{BMF}) \frac{\sigma_{IFF}^2}{\sigma_{VMF}^2}. \quad (11)$$

¹³We choose this construction rather than one that goes long exclusively S-type firms and short exclusively B-type firms to recognize the fact that, when calibrating to the SMB factor in the data, the constituents of the related empirical long-short portfolios may not perfectly conform to the strict definitions of S- and B-type firms in the model. The particular compositions of \tilde{R}_S and \tilde{R}_B are governed by calibrated parameters.

If $\beta_p^{BMF} > \beta_p^{IFF}$ the difference will be negative and if $\beta_p^{BMF} < \beta_p^{IFF}$ the difference will be positive. There two main forces which determine this beta gap. The first is an attenuation bias proportional each portfolio’s BMF beta, analogous to how an OLS regression coefficient is biased downwards in the presence of classical measurement error. The second is an additive bias proportional to each portfolio’s exposure to the IFF. Since small portfolios load negatively on the IFF and tend to have larger betas, we expect the two biases to reinforce each other in this case yielding small-cap portfolio beta gaps that are typically negative. Since large portfolios load positively on the IFF and tend to have smaller betas, we expect these two biases to counteract each other in this case yielding large-cap portfolio beta gaps that are closer to zero. Since there are more small stocks than big stocks (both in the model and in practice), we expect this beta gaps for most portfolios to be negative. In other words, we expect VMF betas to be biased downward relative to BMF betas in most cases.

4.4 Model extension with heterogeneity and non-traded assets

In our baseline model, all consumption is assumed to be financed by tradable financial assets, which we interpret as claims to profits of both listed and unlisted firms. It also imposes that all workers have symmetric preferences and endowments; hence, there is an equilibrium in which all agents consume their endowments and the above asset pricing predictions for listed firms hold whether or not claims to unlisted firms are tradable.

In reality, investors’ consumption is often financed by many other claims to nonfinancial assets that are difficult if not impossible to trade and are often subject to significant, un-diversifiable risk: most notably labor/entrepreneurial income and real estate. Exposures to these risks are heterogeneous across investors, both ex-ante (since nontradable endowments differ across investors) and ex-post (due to uninsurable idiosyncratic risk). The literature on heterogeneous household asset exposures (e.g., Mankiw and Zeldes, 1991; Heaton and Lucas, 2000; Malloy et al., 2009) therefore motivates a potential concern about whether incomplete risk sharing or market segmentation could affect the pricing of g and, by extension, the IFF.

To address this question, we extend our baseline model in Online Appendix OA.4 by introducing non-traded human capital and workers of different “types” (S- and B-types). In the extended model, markets for human capital are fully segmented: agents can neither trade

nor insure against labor income risk except by forming portfolios in financial markets. Each type’s systematic labor income mirrors the factor structure of the corresponding firm type, and additionally features uninsurable idiosyncratic shocks that are drawn from a distribution which may depend on the aggregate state. We refer to this new asset class as “human capital” for concreteness and because it aligns with some of the evidence we present in our local projection impulse response analysis in Subsection 6.3. However, it can be interpreted more generally as any source of additional cash flows available to investors, including ownership stakes in private firms that are non-tradable.¹⁴

In this richer setting, the price of g -risk is determined by two channels. The first is a systematic labor income channel: g is priced if the population-weighted labor income exposure to g does not cancel. The second is an idiosyncratic risk channel: g is priced if the population-weighted average of certainty equivalent utility over state-dependent idiosyncratic labor income risk loads on g . Recall that we interpret labor income broadly as capturing additional sources of aggregate resources (and background risk) not captured by a claim to firm ownership. See Proposition 1 in Online Appendix OA.4 for additional details.

When aggregate risk-adjusted labor income does not load on g , the extended model is isomorphic to the baseline model albeit with different parameters: g is unpriced, the IFF has no risk premium, and the baseline asset-pricing predictions continue to hold (see Online Appendix OA.4.8 for details). However, it is possible that additional resources available to consume (unrelated to the BMF) are lower or risk sharing is worse when small firms do poorly relative to large firms. These empirically plausible departures would imply a negative price of g -risk and hence a negative IFF premium (see Online Appendix OA.4.9 for details). This sign is consistent with our empirical point estimate of the IFF factor premium (see Subsection 6.1 below), which is itself statistically and economically insignificant, suggesting that the departures from our baseline model explored in this extension are small in practice. Thus, for simplicity, we focus on our baseline model when developing the IFF’s empirical implications.

¹⁴Since all investors in the model are assumed to be on their Euler equations, the human capital process stands in for additional resources available to investors who participate in financial markets.

4.5 Model calibration and implications

Our baseline model includes 15 parameters, which we estimate by minimizing the distance relative to an overidentified system of moments. Calibration details appear in Online Appendix OA.3, with all model parameters summarized in Table OA.1. Using the calibrated parameters, we simulate 10,000 draws of the fundamental shocks (f , g , and z), the main factors (R_w , BMF, and VMF), and returns on 382 portfolios. Each portfolio is a weighted combination of small and large firm portfolios, as defined in equation (OA.7), with weights chosen to match BMF betas on the set of 382 test portfolios we use in our empirical analysis. Portfolio-specific z loadings are randomly drawn from the empirical distribution of implied z premia based on our empirical data. Additional details on the simulated asset returns, which we next use to illustrate model implications, are provided in Online Appendix OA.3.3.

Figure 3 illustrates key model implications. Panel A shows that VMF betas are 0.14 lower than BMF betas on average, with beta differences having a standard deviation of 0.18. Given the mean BMF and VMF betas are 1.21 and 1.08, respectively, this gap is substantial. There is not just a level shift in the VMF betas relative to the BMF betas—the slope of the VMF–BMF beta relationship is below one, illustrating the stronger attenuation for high-beta portfolios that contain more small firms. Panel B shows that VMF alphas are on average higher and larger in magnitude than BMF alphas, reflecting the downward bias in VMF betas noted in Panel A. The alpha dispersion arises from heterogeneity in portfolios’ loadings on z , which is not controlled for in these CAPM alpha estimates to be analogous to empirical CAPM pricing results we present below in Section 7.

Panel C plots VMF and BMF alphas for 10 size-sorted portfolios constructed by grouping the 382 portfolios into size deciles. We control for z here to capture predictions related to a multifactor setting, which will allow for a clean comparison with empirical results in Section 8 which control for empirical factors unrelated to size when investigating implications for the size anomaly. The VMF alphas are positive for small portfolios and negative for big portfolios, while BMF alphas are effectively zero—any deviations reflect simulation noise. The VMF misprices these portfolios because it is exposed to the unpriced g shock. As a result, market risk is underestimated for small portfolios and overestimated for big ones, consistent with these portfolios’ negative and positive exposures to g , respectively. The

model also predicts that the IFF has a negative VMF alpha because the IFF, overweighted in large stocks, has a positive VMF beta but earns no risk premium.

Panel D compares factor model-implied risk premia from two VMF-based models and one BMF-based model. All models control for z . One VMF model also includes the SMB factor, while the other does not. Including the SMB is isomorphic to adding a g factor into the model due to small and big firms' differential loadings on g . The BMF model correctly prices assets, producing risk premia consistent with the model's true values. In contrast, the VMF model without SMB implies risk premia (equivalently, discount rates) that are too low on average because its betas are lower than BMF betas on average. Once SMB is added, the VMF model's risk premia align perfectly with those from the BMF model, as SMB corrects for the g -driven contamination in the VMF.

5 BMF and IFF construction methodology

In this section, we summarize a methodology for constructing an empirical proxy for our model's BMF with details and a microfoundation—related to the idea that firms with higher valuation multiples benefit more from going public—relegated to Online Appendix OA.5.

5.1 Methodology

Our goal is to construct an index whose return reflects the value-weighted return on *all* firms, public and private. R_{all} is an empirical proxy for R_w in our model, defined as

$$R_{all} = \frac{\sum_{i=1}^N V_i R_i}{\sum_{i=1}^N V_i} \quad (12)$$

where N is the total number of all firms, V_i is the market value of firm i , and R_i is its return. Unfortunately, we do not observe private firm market values or returns. However, we can define an alternative set of values that, when applied to public firm returns, delivers a consistent estimator for R_{all} if we adjust for: 1. the probability that a firm is listed (i.e., selection), and 2. the economic scale that we want to represent (i.e., market value).

We begin by double-sorting all firms into portfolios p based on the seven KMZ-defined industries, then by within-industry sales deciles using NYSE breakpoints. Crucially, we can compute sales aggregates for these groups of public and all firms using Compustat and KMZ

data, respectively. Motivated by evidence from the literature that growth-like firms (which tend to have higher valuation multiples) are more likely to go public (see, e.g., Kim and Weisbach, 2008; Chemmanur et al., 2010), we assume that, *among firms in group p*, a firm's probability of going public is independent of its sales and linearly related to its valuation multiple, as measured by its market value-to-sales ratio.

Under these assumptions, we can obtain an adjusted weight w_i^* that, when used to aggregate public firm returns, yields a proxy for the value-weighted return on all firms:

$$\hat{R}_{all} = \sum_{i \in \mathcal{L}} w_i^* R_i, \quad \text{where} \quad V_i^* = \underbrace{\frac{V_{p(i)}^{pub}}{S_{p(i)}^{pub}}}_{\text{valuation mult.}} \underbrace{\frac{S_{p(i)}^{all}}{S_{p(i)}^{pub}}}_{\text{selection adj.}} \underbrace{S_i}_{\text{size}} \quad \text{and} \quad w_i^* = \frac{V_i^*}{\sum_{i \in \mathcal{L}} V_i^*}, \quad (13)$$

where \mathcal{L} denotes the set of public firms in N , $p(i)$ reflects the portfolio set to which the public firm i belongs (for our main construction, industry \times sales decile), $V_{p(i)}^{pub}$ and $S_{p(i)}^{pub}$ denote, respectively, the total market value and sales of public firms of type $p(i)$, and $S_{p(i)}^{all}$ is the total sales of all firms (public and private) associated with the industry-sales group $p(i)$. The ratio $S_{p(i)}^{all}/S_{p(i)}^{pub}$ acts as an adjustment for the selection of public firms to scale the observed firm's sales S_i so that it better reflects all firms. The valuation multiple $V_{p(i)}^{pub}/S_{p(i)}^{pub}$ translates this quantity into an implied market value. Under reasonable assumptions, \hat{R}_{all} provides a consistent estimate of R_{all} when the number of firms is large (see Online Appendix OA.5).

Consistent with our model's BMF definition in equation (8), we identify the empirical BMF by running a full-sample regression of the VMF on \hat{R}_{all} as

$$VMF_{t+1} = \underbrace{\beta_{VMF}^{all} \left(\hat{R}_{all,t+1} - R_{f,t} \right)}_{\equiv BMF_{t+1}} + \underbrace{a + \varepsilon_{t+1}}_{\equiv IFF_{t+1}}. \quad (14)$$

By construction, the VMF has a loading of one on the BMF and the IFF is orthogonal to the BMF. One key assumption in our model is that the risk premium on g is zero. This implies $a = 0$ (i.e., the empirical IFF has a mean of zero), which we test below.

Before proceeding, we pause to acknowledge potential limitations this approach. We condition on observable variables in order to reweight public firms in order to capture the unobservable return on private firms. Other unobservable characteristics of private firms could correlate with risk exposures which are not shared by public firms with similar ob-

servables. While the assumptions underlying our estimator are unlikely to hold exactly, our primary objective is to capture the sources of systematic risk captured by private firms rather than to obtain an exact point estimate for the level of wealth. Hence, while the Roll (1977) critique related to a mismeasured market proxy still applies to \hat{R}_{all} , we conjecture that it more accurately captures the contribution from private firms relative to ignoring their existence completely—as is the case for the *VMF*—and therefore \hat{R}_{all} is likely to work better in asset pricing tests. The remainder of our paper investigates this empirical question.

5.2 Empirical implementation

To implement the methodology from Section 5.1, we sort stocks into seven KMZ industries then into within-industry decile portfolios by sales using NYSE breakpoints (except for the construction industry, which we sort into tercile portfolios by sales due to limited data), yielding a total of 63 portfolios. These are the portfolios for which we must measure $V_{p(i)}^{pub}$, $S_{p(i)}^{pub}$, and $S_{p(i)}^{all}$ to implement equation (13).

We compute $V_{p(i)}^{pub}$, $S_{p(i)}^{pub}$, and S_i directly from CRSP and Compustat data. We estimate $S_{p(i)}^{all}$ using KMZ data on total sales broken into industry and sales bins available from 1963–2013. Since this data ends before our sample end date, we apply sales totals from 2013 to the subsequent periods in our sample when necessary.¹⁵ See Online Appendix OA.5.3 for additional details on the empirical implementation of our methodology.

Figure 4 Panels A–B plot value weights over time of sales-sorted portfolios 1 and 10 in the *VMF* compared to those in the *BMF* implied by equation (13). The lowest sales decile receives larger value weights while the top decile receives smaller value weights in the *BMF* compared to in the *VMF*. Although Figure 4 Panels A–B show that the *BMF* places relatively more weight on small firms and relatively less weight on large firms than does the *VMF*, the *BMF* is distinct from an equal-weighted portfolio.¹⁶ Panels C–D plot value weights for industry portfolios in the *VMF* compared to those implied by our methodology

¹⁵Sales shares vary little from year to year, so using data from 2013 (or 2018, in our robustness check that uses sales-sorted portfolios discussed below) to apply in the last few years of our sample is unlikely to have a material impact on our estimated *BMF*. While the IRS has discontinued reports that aggregate by both industry and size, analogous tabulations could be recomputed with microdata from other statistical agencies.

¹⁶The weights in our *BMF* computed according to equation (13) are more evenly distributed than value weights in the *VMF*, but are distinct from a simple equal-weighted index (see Online Appendix Figure OA.2).

for constructing the BMF. The panels show that industry value weights in the BMF are more evenly distributed and less time-varying than those in the VMF. Thus, our BMF uses weights that are distinct from those in the VMF both in the size and industry dimensions, although the most salient distinction is in the size dimension.

5.3 Alternative BMF constructions for robustness

Because the largest compositional differences between public and all firms arise from size rather than industry, we first construct an alternative BMF using a base set of 10 sales-sorted portfolios and the aggregated sales-bin-only data from KMZ to implement equation (13). This construction yields similar weighting implications across size and industry portfolios as in our main specification.¹⁷

Our main BMF construction requires a significant amount of data from the IRS (via KMZ), as well as several assumptions. As an alternative, we next consider using statistical factors to identify alternative BMFs which use our main BMF only for the purpose of characterizing its loadings on the statistical factors. We restrict attention to statistical factors which respect a salient feature of our BMF’s construction: small firms receive more weight and large firms receive less weight in the BMF relative to the VMF. In particular, we estimate statistical factors using the following approaches:

- **PCA:**¹⁸ We use principal components analysis (PCA) applied to equal-weighted or capped value-weighted versions of our 382 portfolios of test assets. The capped value weighted portfolios place moderate weights on large stocks compared to their value weights but significantly less weight on the smallest of stocks than equal-weighted portfolios, thus improving tradability while limiting excessive weights on large stocks.¹⁹
- **IPCA:** We use the latent instrumented-principal-component-analysis (IPCA) factors from Kelly et al. (2019).²⁰

¹⁷See Online Appendix Figures OA.3–OA.6 compared to Figures 2, OA.1, OA.2, and 4, respectively.

¹⁸There is precedent for using PCA to identify pricing factors as in Chamberlain and Rothschild (1983) and Connor and Korajczyk (1986, 1988). More recent examples include Kozak et al. (2018), Pukthuanthong et al. (2019), Kelly et al. (2019), and Giglio and Xiu (2021).

¹⁹We obtain the equal-weighted portfolio returns directly from Ken French’s Data Library. We construct the analogous capped-value-weighted portfolios ourselves using the same sorting variables (obtained from CRSP and Compustat) but applying the capped-value-weighting methodology from Jensen et al. (2023).

²⁰We use the factors identified without assuming an intercept because they are priced correctly under the

While we could project the BMF onto multiple factors, we project it onto the first factor from each approach for two main reasons. First, that there is a high correlation (around 0.98–0.99) between our main BMF and each of the first factors (see Table 1, discussed below). Second, we cannot reject the null hypothesis that our main BMF is priced by each of these first factors, with little gain in explanatory power from including other factors (see Online Appendix Table OA.4). Thus, we use the first PC from each method above in place of $\hat{R}_{all,t+1}$ to implement the decomposition from (14), yielding three alternative BMFs and IFFs.²¹

All four alternative BMFs closely match our main BMF and generate similar quantitative results. The near-perfect correlation between our main BMF—which relies on IRS data and several methodological assumptions—and statistical factors represents an important finding: it implies that the compositional effect we identify corresponds to the dominant dimension of return variation, and is not an artifact of the particular data or assumptions that underlie our main construction. Further, the BMF’s close correspondence with first principal components—single factors that, by construction, explain the most return variation across assets—gives it a natural interpretation within the Arbitrage Pricing Theory framework of Ross (1976) and connects our analysis to the statistical factor identification literature (e.g., Chamberlain and Rothschild, 1983; Connor and Korajczyk, 1986, 1988; Kozak et al., 2018). This connection provides a new economic interpretation of what these model-free approaches identify: a broad “level” pricing factor that corrects for the compositional bias inherent in the value-weighted market portfolio. For brevity, all results in the main text are based on our main BMF and associated IFF unless otherwise noted. We reference appendix results that use these alternative BMFs (and corresponding IFFs) in footnotes where relevant.

null of the model and the ones that received the most emphasis in Kelly et al. (2019). Since their factors are only available from July 1964 through May 2014, all analyses with IPCA factors use that date range.

²¹Results are robust to projecting our main BMF onto additional PCs and using these projections in place of $\hat{R}_{all,t+1}$, but we omit these results for brevity.

6 The BMF and IFF: An initial characterization

6.1 BMF and IFF summary statistics

Table 1 reports VMF, BMF, and IFF summary statistics. We include results for our main BMF as well as the alternative BMFs described in Section 5.3. Panel A presents statistics related to the BMF and Panel B presents statistics related to the IFF. Average returns on the VMF and the various BMFs are similar at approximately 7% per year. The BMF has a lower volatility than the VMF, is less negatively skewed, and has higher kurtosis.

If the IFF captures an important source of priced risk, we would expect it to earn positive excess returns on average. However, the average return on our main IFF is economically small at -0.51% per year and we cannot reject the null hypothesis that it is equal to zero, with similar results for other IFF specifications. This result is consistent with our model's assumption that g carries no risk premium. Furthermore, the average IFF return is equivalent to the alpha from a regression of the VMF on our main BMF (since it is orthogonal to the BMF by construction), and implies that the BMF prices the VMF with the BMF explaining about 83% of its variation. Equivalently, the ratio of IFF-to-VMF variance is 17%, so the fraction of total variability in the VMF induced by the IFF is nontrivial. These results are consistent with our hypothesis that the main source of risk in the VMF for which investors demand a positive risk premium is captured by the BMF and provides prima facie evidence against the inclusion of the IFF in a proxy for systematic, priced risk.²²

We provide correlation matrices for the BMF and IFF variants in Panels A and B, respectively. All BMF variants are highly correlated (0.98 or higher) with our baseline BMF and also with each other. This, paired with their similar average returns, explains their similar performance to our main BMF in robustness checks.

6.2 IFF pricing

We next run IFF factor spanning tests to test the model implication that the IFF has a negative alpha with respect to the VMF and zero alpha with respect to the BMF, as

²²We plot cumulative return time series for our BMF and IFF variants in Online Appendix Figure OA.7. Cumulative BMF returns are large and perform similarly across variants. By contrast, cumulative IFF returns are essentially flat (but slightly negative).

highlighted in Figure 3 Panel C. To control other factors beyond the market factor that may price assets (i.e., the role played by z in the model), we run the tests using nine standard factor models including 1) the CAPM, 2) the Fama and French (1993) three-factor model (“FF3”), 3) the Fama and French (2015) five-factor model (“FF5”), 4) the FF5 model augmented with the Carhart (1997) momentum factor (“FF5C”), 5) the Hou et al. (2015) four-factor model (“q4”), 6) the Stambaugh and Yuan (2017) four-factor model (“SY”), 7) the Daniel et al. (2020) three-factor model (“DHS”), 8) the Hou et al. (2021) five-factor model (“q5”), and 9) the Chabi-Yo et al. (2025) three factor ICAPM-inspired model (“ICAPM”).

Although the only pricing factors in our model are the BMF and z , our model predicts that the IFF alpha will be zero if a size factor is included in the pricing model regardless of whether the BMF or the VMF is used as the market factor. This is because, even though the size factor is not itself an independent pricing factor, it allows both model types to span the IFF due to its exposure to g . So, the above predictions regarding IFF alphas only hold when models do not include size factors. Thus, in our main tests, we remove size factors from factor models that have them. However, we do include these size factors in auxiliary IFF spanning tests which we discuss below.

Table 2 presents results from regressing the IFF on factors from each model using either the VMF (Panel A) or the BMF (Panel B) as the market factor with size factors excluded from the spanning regressions. When we use the VMF, the IFF alphas are always negative and typically statistically significant, consistent with our model’s implications. The intuition is that if the IFF is not priced but loads positively on the VMF (due to the IFF and VMF both being overweight in large stocks), it will generate a negative alpha. By contrast, when we use the BMF, the alphas are much lower in magnitude and we cannot reject the null that the IFF alphas are equal to zero (except for the q5 model), again consistent with our model.

Table 2 also presents results from regressing the IFF on factors from each model using either the VMF (Panel C) or the BMF (Panel D) as the market factor only for models that include size factors with size factors included in the spanning regressions. Although results that use the VMF are still all negative and typically statistically significant, their magnitudes decrease dramatically compared to those in Panel A. This result is directionally consistent with the model prediction that IFF alphas shrink to zero when size factors are included in

the tests. However, the alphas are not exactly zero, which is likely due to a combination of the fact that the IFF loads on other empirical factors in these factor models (whereas it does not load on z by assumption in our theoretical model) and the fact that empirical SMB factors are not perfectly explained by the BMF and the IFF as in our theoretical model. For example, regressing the FF3 SMB factor on the BMF and the IFF yields an R^2 of 65%. Similar to results in Panel B, results in Panel D show that we typically cannot reject the null that IFF alphas are zero when the BMF is used as the market factor even when size factors are included in the spanning regressions, again consistent with our model.²³

Market factor spanning test interpretation: Because the VMF is the sum of the BMF and the IFF, there is an alternative interpretation of the spanning tests presented in Table 2. Namely, the alphas from VMF-based models are exactly the same in magnitude but opposite in sign as the alphas that would obtain from a BMF spanning test. Similarly, the alphas from BMF-based models are exactly the same in magnitude but opposite in sign as the alphas that would obtain from a VMF spanning test. In other words, results in Panel A imply that the BMF is not spanned by models that use the VMF as the market factor (excluding the FF3 model) and produces positive alphas. By contrast, results in Panel B imply that the VMF is spanned by models that use the BMF as the market factor (excluding the q5 model). Similar conclusions follow from results in Panels C and D where size factors are included in these models. Thus, using the BMF in place of the VMF should improve the pricing ability of these models (Barillas and Shanken, 2017).

6.3 The IFF and the macro-financial disconnect

If the IFF captures unpriced risk, it should not correlate with macroeconomic aggregates that proxy for investor marginal utility. In the context of our model, this follows from the assumption that g does not affect aggregate consumption and therefore does not enter the stochastic discount factor (SDF). To test this, we estimate local projection impulse responses (LPIRs; Jordà, 2005) with respect to the BMF and the IFF for fifteen empirical proxies of investor marginal utility (macroeconomic aggregates and returns on other asset classes)

²³We report robustness checks on the results in Table 2 in Online Appendix Table OA.5 using the alternative BMF constructions described in Section 5.3.

drawn from several different classes of asset pricing models.²⁴ Mapping this exercise to the model, these variables can act as empirical proxies for priced risk factors which enter the SDF either directly as consumption growth proxies or indirectly via the exogenous component z . Details on these variables and their theoretical connections to marginal utility are provided in Online Appendix OA.7.1, and details on the LPIR estimation, which follows Schmidt (2025), are provided in Online Appendix OA.7.2.

Figure 5 plots the LPIRs. The blue lines correspond with responses to the BMF, whereas the red lines correspond with responses to the IFF.²⁵ Variables are signed such that increases in their values correspond with decreases in marginal utility (i.e., we would expect these variables to have positive LPIR coefficients with respect to the BMF).²⁶

The primary take-away from this analysis is that the BMF robustly forecasts contemporaneous and future increases in macroeconomic aggregates associated with decreases in marginal utility, whereas the IFF does not. The statistically insignificant responses of consumption growth, GDP growth, and total income growth with respect to the IFF (Panels (A)–(C)) imply that the potential redistributive effects of non-traded labor income that we explore in our extended model (Online Appendix OA.4) which would cause the IFF to be priced are likely weak in practice. We find similar results for income growth of the top 0.1% of earners and proprietary income growth (Panels (F)–(G)), which reflect the incomes of agents who are more likely to be marginal investors in public financial markets. We find mixed results for state-dependent income risk proxies: point estimates for the Herskovic et al. (2016) CIV factor are negative and statistically significant (consistent with a negative

²⁴Our choice of macroeconomic aggregates and other time series are log per capita consumption growth (FRED), log GDP growth (FRED), total US income growth (World Inequality Database, <https://wid.world/data/>), real uncertainty (Jurado et al., 2015), top 0.1% US income growth (World Inequality Database, <https://wid.world/data/>), non-farm proprietary income growth (FRED) (Heaton and Lucas, 2000), the income risk measure based on cross-sectional skewness in household income growth (Schmidt, 2025), the unemployment rate (FRED), the common idiosyncratic volatility (CIV) factor (Herskovic et al., 2016), the intermediary capital risk factor (He et al., 2017), the common mutual fund flow factor (Dou et al., 2024), the credit spread (Moody’s seasoned Baa corporate bond yield relative to the yield on 10-year treasury), the Bloomberg Barclays US Aggregate Bond Index, the Hedge Fund Research fund-of-funds diversified hedge fund return index (HFRIFOFD), and the FTSE NAREIT US Real Estate Index.

²⁵We compare LPIRs with respect to the BMF and the VMF in Online Appendix OA.7.3. The LPIRs with respect to the VMF are typically slightly lower than those with respect to the BMF, which, as we explain there, are as expected.

²⁶Note that we expect permanent response increases for non-stationary variables such as Consumption Growth (Panel A), but mean-reverting responses for stationary time series such as Income Risk (Panel G).

risk price), but they are slightly positive but insignificant for the Schmidt (2025) skewness index. Hence, we do not find much support for the assumptions required for g to be priced in our baseline or extended model. Stepping outside of our theoretical framework, regardless of one’s preferred microfoundation for risk premia, a stark contrast emerges between the BMF and the IFF. While the BMF is tightly linked with variables that proxy for investor marginal utility, such a link is tenuous at best (and sometimes incorrectly signed) for the IFF. This is the sense in which we deem the IFF “idiosyncratic” with respect to the broader economy.

7 CAPM Pricing with VMF vs BMF market factors

In this section, we explore the pricing implications of our BMF versus the VMF in the canonical CAPM. Since the CAPM is the de facto model for estimating discount rates among many practitioners, it is important to understand whether the VMF leads to predictable biases. After presenting our findings in Section 7.1, Section 7.2 inspects the mechanism by studying the relationship between VMF and BMF betas.

7.1 Alphas and price wedges under the VMF versus BMF

Here, we explore the empirical pricing implications of VMF- and BMF-based CAPM models via two pricing metrics: CAPM alphas and CAPM price wedges (van Binsbergen et al., 2023). Figure 6 Panel A plots histograms of VMF and BMF alphas for our 382 test portfolios. It replicates a well-known issue with the VMF: it places most assets above the security market line (SML, i.e., the relation between expected excess returns and betas predicted by the CAPM), implying positive alphas and an underestimation of discount rates on average. By contrast, the BMF produces alphas that are centered around zero with similar dispersion. Consistent with the model, the VMF-CAPM yields discount rates that are too low on average, but the BMF implies discount rates that are “correct” on average. However, pricing errors remain, as happens in the model due to omitted z exposures.

Figure 6 Panel C plots VMF alphas against BMF alphas and shows that VMF alphas are generally larger than the corresponding BMF alphas—consistent with the model implication in Figure 3 Panel B.²⁷ Figure 6 Panel E plots alpha differences against the “beta gap,” which

²⁷We report robustness checks on the results in Figure 6 Panel C in Online Appendix Figure OA.10 using the alternative BMF constructions described in Section 5.3.

we define as the difference between VMF and BMF betas (i.e., $\beta^{VMF} - \beta^{BMF}$). The strong relationship between alpha differences and beta gaps is expected under two alternative single-factor CAPMs when the two factors have similar means, as is the case for the VMF and the BMF.²⁸ Hence, the superior pricing under the BMF is the result of these beta differences, which we argue is due to the BMF beta’s ability to better capture market risk.

Figure 6 Panel B plots histograms of VMF- and BMF-based (negative) price wedges for the 382 value-weighted portfolios. van Binsbergen et al. (2023) define price wedges to be the log difference between an asset’s current price and its price implied by a particular SDF model (i.e., $PW_t \equiv \log\left(P_t/\tilde{P}_t\right)$, where P_t is the current price and \tilde{P}_t is the model-implied price using the model’s SDF to discount and cumulate expected future dividends and a terminal value).²⁹ We report negative price wedges so that the sign convention with respect to mispricing is consistent with that implied by alphas (i.e., positive signs reflect underpricing and vice versa). Under the VMF-CAPM, the price wedge distribution is centered far away from zero implying most portfolios are underpriced. By contrast, under the BMF-CAPM, the price wedge distribution shifts left towards zero (indicating less implied mispricing) though it remains centered around a positive mean. Panel D plots VMF against BMF price wedges and reveals the same pattern observed for alphas—VMF wedges are generally larger. Panel F plots price wedge differences against beta gaps, again revealing that most valuation differences can be explained by differences in betas, as we further discuss in section 7.2.

Table 3 Panel A presents summary statistics for the VMF- and BMF-based alphas, along with formal tests for differences in these summary statistics across the two models. Specifically, we report averages, medians, 25th quantiles, and 75th quantiles for raw alphas along with their absolute values. We also test for differences in these statistics applied to the absolute alphas. For the set of alphas from all portfolios (left portion of Panel A), we reject the null that absolute alphas under the VMF are smaller than those under the

²⁸To be more precise, it is straightforward to show that the differences in alphas can be expressed as $\alpha_i^{VMF} - \alpha_i^{BMF} = \beta_i^{BMF} \cdot \mathbb{E}[BMF] - \beta_i^{VMF} \cdot \mathbb{E}[VMF]$. Thus, when $\mathbb{E}[BMF] \approx \mathbb{E}[VMF]$, alpha differences are mostly explained by beta gaps.

²⁹We use the same estimation and simulation procedures as in van Binsbergen et al. (2023) to estimate CAPM-implied price wedges that they denote using PW^* . These are the price wedges that result from an estimation procedure that corrects for one of two estimation biases highlighted by van Binsbergen et al. (2023) and which they find yields results similar to those that correct for both biases (albeit with a much simpler estimation procedure).

BMF at the 10% level for the 25th quantiles, median, and 75th quantiles, and are close to rejecting for the averages. Restricting the sample to the 50% of portfolios with the largest beta gaps strengthens these results: we reject the null that absolute alphas under the VMF are smaller than those under the BMF at the 5% level for all statistics. In other words, there is statistical support for the model’s prediction that the BMF yields lower average, median, and 25th/75th quantiles for absolute alphas across our test portfolios. The evidence is particularly strong for portfolios with the largest beta gaps, which, according to our model, are those with the most mispricing under the VMF and the highest statistical power.

Table 3 Panel B reports analogous statistics for price wedges. Although average, median, and 25th/75th quantiles of absolute price wedges are lower under the BMF than under the VMF (left portion of Panel B), the differences are only close to significant at the 10% level. However, among the portfolios with the largest beta gaps, all differences become significant at the 10% level. Overall, the results suggest that the BMF produces smaller absolute price wedges than the VMF, with the strongest evidence for the largest beta gap portfolios.

7.2 Mechanism: Differences between VMF and BMF betas

We next examine the sources of the VMF–BMF beta differences that underpin the BMF’s superior pricing performance. Equations (10)–(11) from the model show that the IFF introduces two sources of bias in VMF betas relative to BMF betas: 1) an attenuation bias proportional to the BMF beta, and 2) an additive bias proportional to the IFF beta. By the related logic in Section 4.3, we expect most BMF betas to be larger than VMF betas, but with differences that vary in the cross section due to BMF and IFF beta variation.

Figure 7 Panel A plots histograms of VMF and BMF betas for the 382 test portfolios. Consistent with our interpretation of equations (10)–(11), the average VMF beta (1.09) is below the average BMF beta (1.21). Panel B plots VMF betas against their corresponding BMF betas and shows that, portfolio by portfolio, VMF betas are typically lower but that the relationship varies in the cross section. Differences are economically significant: the cross-sectional standard deviation of the VMF–BMF beta gap is 0.13—roughly two-thirds of the cross-sectional standard deviation of VMF betas (0.19). These findings support the model’s prediction that VMF betas are generally biased downward relative to BMF betas

but with differences that vary in the cross section similar to those in Figure 3 Panel A.³⁰

The VMF–BMF beta gaps highlighted above are the primary driver of the pricing improvements under the BMF illustrated in Figure 6. A key implication is that CAPM discount rates estimated using the VMF are systematically too low, with the gap increasing in market beta. Furthermore, our results offer a potential rationale for the finding from Decaire and Graham (2024) that sell-side analysts use higher market betas than typical VMF estimates would imply (with the gap between VMF and analyst betas also increasing in beta) to price assets (see their Figure 6 Panel C).

In addition to the tendency of small-cap portfolios to have higher market betas than large-cap portfolios, a key driver of the beta gaps according to equations (10)–(11) is our model’s assumption that small portfolios load negatively on the IFF and large portfolios load positively on it. Figure 8 plots IFF betas against two measures of portfolio size—average log market capitalization (Panel A) and the negative SMB beta (Panel B). Consistent with our model’s assumption, small portfolio IFF betas are negative, large portfolio betas are positive, and the betas increase with size. To the extent that IFF betas increase in portfolio size and BMF betas decrease in portfolio size, equation (11) implies that beta gaps should be negative for small-cap portfolios but increasing in portfolio size. Figure 8 Panels C and D confirm this relationship: beta gaps increase with both portfolio size measures, and the largest gaps are negative beta gaps associated with small-cap portfolios. Consistent with our interpretation of equations (10)–(11), the additive bias term mostly offsets the attenuation term for the largest portfolios where the beta gaps are close to zero. Dotted blue lines plot model-implied relationships between the variables in each panel and demonstrate similar relationships as in the data.

Given the link between beta gaps and size, and the fact that the BMF overweights small stocks relative to their VMF weights, one might be concerned that the BMF only improves pricing for small-cap portfolios and, potentially, at the expense of pricing large-cap portfolios. To assess it, we divide the 382 value-weighted portfolios into three groups by average market capitalization and compare resulting VMF and BMF alpha distributions in

³⁰We report robustness checks on the results in Figure 7 Panel B in Online Appendix Figure OA.9 using the alternative BMF constructions described in Section 5.3.

Figure 9. Although both CAPM models generate sizable alpha dispersion, the BMF delivers better pricing for both small- and mid-cap portfolios without sacrificing performance for large-caps. The large-cap result aligns with Figure 8 Panels C and D: BMF betas for large portfolios are similar to VMF betas, implying similar pricing accuracy across the BMF and VMF for that portfolio group.

We formally test for differences in VMF and BMF alphas by portfolio size in Table 4, which reports statistics on raw and absolute alpha distributions by size group. Among the smallest portfolios, we reject the null that absolute VMF alphas are lower than absolute BMF alphas at the 5% significance level for their averages, medians, and 25th/75th quantiles. For mid-cap portfolios, we reject the null at the 10% level for the median and 25th quantile, and are close to rejecting for the average and 75th quantile. For the large-cap portfolios, we cannot reject the null, although the average, median, and 25th/75th quantiles of absolute VMF alphas remain higher than those of absolute BMF alphas. Thus, the BMF prices small- and mid-cap portfolios better than the VMF without sacrificing performance for large-caps.

8 The IFF, the size anomaly, and the role of size factors

Next, we test our model's predictions related to the size anomaly and the role of size factors. In Section 8.1, we test model predictions about the ability to price size-sorted portfolios with and without size factors. In Section 8.2, we test model predictions about whether or not size factors are spanned and the extent to which size factors help to reduce pricing errors using our full set of test portfolios. Section 8.3 provides additional evidence on the mechanism.

8.1 The BMF prices size-sorted portfolios without size factors

In Figure 8 above, we show that beta gaps are highly correlated with portfolio size (in our model with two firm types, this correlation is 100%). Thus, size sorted portfolios provide a natural laboratory for testing a key prediction of our model: alphas are larger for portfolios which have large beta gaps when size factors are not included. This is the force that causes the multi-factor alphas reported in Figure 3 Panel C to be monotonically decreasing in portfolio size when the VMF is used as a pricing factor and size factors are not included. By contrast, these alphas are zero when the BMF is used instead.

Figure 10 reports a direct empirical analogue of this test, showing alphas for 10 value-weighted size-sorted portfolios estimated using nine standard factor models described in Section 6.2. To test the model predictions, this analysis excludes size factors from the six models that have them (FF3, FF5, FF5C, q4, SY, and q5). Panel A reports results using the VMF as the market factor in each model. In this case, the small-cap portfolio (decile 1) yields positive and statistically significant alphas across all models at the 10% level, except under the CAPM and FF3 models, where alphas are positive but insignificant. Within each model, alphas generally decline monotonically with increasing portfolio size. Panel B reports results when the VMF is replaced by the BMF in each model. In this case, small-cap portfolio alphas are statistically indistinguishable from zero in all but two cases (SY and FF5C), where they are slightly negative. Within each model, BMF alphas show no consistent pattern with portfolio size, and their magnitudes are much smaller compared to those under the VMF. Overall, these findings are consistent with their model analogues in Figure 3 Panel C.³¹

Consistent with results from Figure 6 Panel A that show VMF-CAPM alphas are typically positive across a broad set of portfolios, VMF multi-factor alphas are positive in all size-sorted portfolios except for the large-cap (decile 10). Stated differently, *multi-factor models that use the VMF without size factors (including the CAPM) imply discount rates that are too low for the average stock in each size category outside of the top decile*. By contrast, as shown in Figure 6 Panel A for a broad set of portfolios and as we show here for size-sorted portfolios, using the BMF prices all these portfolios correctly on average. Thus, replacing the BMF with the VMF improves the “calibration” of models that exclude size factors across most size categories, including fairly large and mid-cap stocks, without degrading its performance in the top decile, echoing our findings in Figure 9 about CAPM alphas.

In our model, adding a size factor to VMF-based models improves pricing even when the size premium is small, since the size factor still differentiates between the assets’ exposures to priced BMF and unpriced IFF risk. Consistent with this logic, we find similar results to those in Figure 10 when we restrict our sample to 1993–2021 (see Online Appendix Figure OA.12). During this sample, the SMB factor earned an average annualized return of 1.6%,

³¹We report robustness checks on the results in Figure 10 Panel B in Online Appendix Figure OA.11 using the alternative BMF constructions described in Section 5.3.

roughly half its 3.2% average from 1963–2021. Despite this decline, size-sorted portfolio alphas remain quantitatively similar (albeit less precisely estimated in the shorter sample). Thus, irrespective of the ongoing debate over the declining size premium, standard models continue to misprice stocks outside the top size decile when size factors are excluded.

8.2 The BMF spans and obviates the need for size factors

A key prediction of our model is that size factors are not spanned (i.e., priced) by the VMF due to their loading on the IFF, but are spanned by the BMF. Thus, they should improve pricing power in models that use the VMF as the market factor, but not in models that use the BMF. We test these implications by running size factor spanning tests, and testing the ability of factor models to price a broad set of portfolio when we either include or exclude size factors in those models. Furthermore, to the extent that empirical factors do not span the IFF (like z in the model), including other factors in these tests should not allow VMF-based models to span size factors.

Spanning tests: Table 5 reports alphas from size factor spanning tests, where size factors from each of the six models that include size factors are regressed on other model factors and either the VMF (Panel A) or the BMF (Panel B) is used as the market factor. In Panel A, size factor alphas are statistically and economically significant in all models except FF3, consistent with the motivation for adding size factors—adding unspanned factors improves model pricing (Barillas and Shanken, 2017). In contrast, when the BMF is used, size factors are spanned in all models: their alphas are both economically and statistically insignificant, indicating size factors do not improve pricing in this case. Reported VMF and BMF betas show that size factors load more on the BMF than on the VMF, allowing BMF models to span them via increased market risk. Thus, substituting the BMF for the VMF effectively eliminates the need for size factors without reducing pricing performance.³²

Pricing error reductions from size factors: Another way to show that replacing the VMF with the BMF eliminates the need for size factors is to test each model’s ability to price a broad set of test assets. Thus, we estimate alphas for each of our 382 value-weighted

³²Consistent with the notion that size factors generate pricing improvements in factor models even when the size premium is small, spanning results are similar if we limit the sample to the 1993–2021 subsample. We report these results in Online Appendix Table OA.6 for brevity.

portfolios using four variants of each of the six models that include size factors: models that use 1) the VMF with the size factor included, 2) the VMF with the size factor excluded, 3) the BMF with the size factor included, and 4) the BMF with the size factor excluded. Figure 11 reports the resulting mean absolute alphas across all 382 portfolios for each model. When the VMF is used (red bars), mean absolute alphas are consistently higher when size factors are excluded. By contrast, when the BMF is used (blue bars), mean absolute alphas are comparable between the models that include and those that exclude the size factors.

We formally test for differences in the mean absolute alphas between models in Table 6. Panel A reports mean absolute alphas across all six factor models that use the VMF, comparing cases where the size factors are included versus excluded. It also reports p-values testing the null that mean absolute alphas are smaller when size factors are omitted. In every model, mean absolute alphas are larger when size factors are excluded, and we strongly reject the null for all models except FF3, where we are close to rejecting the null at the 10% level. Overall, including size factors in VMF-based models yields statistically and economically significant pricing improvements. Panel B reports analogous results when the BMF is used in each model. Here, we fail to reject the null that mean absolute alphas are lower in models that exclude size factors across all models. Moreover, the magnitudes of mean absolute alphas are nearly identical across models whether size factors are included or not. Finally, Panel C compares mean absolute alphas across models that use either the VMF or BMF as a market factor within models that either include or exclude size factors. For models that include size factors, we cannot reject the null that VMF-based alphas are less than BMF-based alphas (first two rows). However, for models that exclude size factors, we reject this null in four out of the six models (last two rows), and are close to rejecting at the 10% level in one additional case (FF5C).

In summary, adding size factors to VMF-based models improves pricing performance. In contrast, adding size factors to BMF-based models yields no meaningful improvement. Moreover, BMF-based models achieve similar pricing performance as VMF-based models that include size factors whether size factors are included in the BMF-based models or not.

8.3 Mechanism: How size factors modulate estimated risk premia

Although SMB does not act as an independent pricing factor in our model, if the VMF is used as a market factor in place of the BMF, then SMB does play a role: it allows the model to separate exposures to the priced (BMF) versus the unpriced (IFF) components of the VMF, thus aligning VMF-based risk premia with the true risk premia as we show in Figure 3 Panel D. This explanation is quite distinct from typical justifications for including size factors in factor models, which often revolve around the argument that small stocks are exposed to *additional risks* to which large firms are not exposed. Here, we provide empirical evidence consistent with our model’s interpretation of size factors as control variables that help align VMF-based model-implied risk premia with those from BMF-based models.³³

Figure 12 Panel A plots VMF- versus BMF-based model-implied risk premia for our test portfolios across the six factor models that have size factors, but exclude them in the estimation. In this case, the risk premia implied by VMF-based models are typically too low compared to those implied by BMF-based models, a pattern which is mimicked in our model (represented by the dotted blue line in the figure). However, including size factors in VMF-based models sharply increases the alignment of their implied risk premia with those from BMF-based models, as we show in Panel B.

Panel C plots model-implied risk premia from BMF-based factor models that include size factors against those that exclude them. In this case, model-implied risk premia are nearly identical regardless of whether size factors are included, consistent with the alpha results in Figure 11 and Table 6. This pattern is also consistent with results from our model, represented by the dotted blue line in this figure. Thus, including size factors improves VMF-based model pricing by bringing the model-implied risk premia more in line with those from BMF-based models, regardless of whether the latter include size factors.³⁴

Our model has only one dimension over which the composition of public and all firms differ, which we refer to as “size.” Although size is the most salient, there are many dimen-

³³By “model-implied risk premia,” we mean the following. Suppose F reflects a set of factors for a particular model and β the vector of factor loadings. The expected excess returns for a portfolio p can be expressed as $\mathbb{E}[R_p - R_f] = \alpha_p + \beta'_p \mathbb{E}[F]$. The model-implied risk premium is then given by $\beta'_p \mathbb{E}[F]$ where we measure $\mathbb{E}[R_p - R_f]$ and $\mathbb{E}[F]$ using their sample counterparts in the data.

³⁴We report robustness checks on the results in Figure 12 in Online Appendix Figure OA.13 using the alternative BMF constructions described in Section 5.3.

sions over which public and all firms could differ. These differences could have important implications for factors beyond size if those factors also load strongly on the IFF. We investigate the quantitative importance of such potential differences in a manner similar to our analysis of the size anomaly. Namely, we run factor spanning tests for each factor in each factor model that we investigate when using either the VMF or the BMF as the market factor. The main takeaway from this analysis is that, unlike for the size factor spanning results reported in Table 5, alphas from spanning tests on other factors are similar whether a VMF- or BMF-based model is used. The main driver of this result is that, unlike size factors, other factors are not strongly correlated with the IFF and thus are priced similarly whether the VMF or the BMF is used as a market factor. For brevity, we relegate the details of this analysis to Online Appendix OA.8.

9 The IFF and the intertemporal risk-return relation

Thus far, our model and empirical analyses have focused on the IFF’s cross-sectional implications. In this section, we study its implications for the intertemporal risk-return relation.

9.1 Econometric intuition

Consider a standard regression that tests for the intertemporal risk-return relation as

$$R_{M,t+1} - R_{f,t} = a + \gamma\sigma_t^2 + \varepsilon_{t+1} \quad (15)$$

where $R_{M,t+1} - R_{f,t}$ represents the realized excess market return in period $t+1$, σ_t^2 represents its conditional volatility, and γ reflects the intertemporal risk-return relation. This test can be motivated by a number of studies (see, e.g., French et al., 1987; Glosten et al., 1993).

Assume the true market factor is the BMF and its risk premium obeys the relationship in equation (15). If one had a perfect measure of the BMF’s variance, $\sigma_{BMF,t}^2$, one could consistently estimate γ by running a regression of BMF returns on $\sigma_{BMF,t}^2$. If instead we tried to estimate γ using the VMF and a perfect σ measure of its conditional variance, $\sigma_{VMF,t}^2$,

under mild assumptions our estimate would converge to³⁵

$$\text{plim}_{T \rightarrow \infty} \hat{\gamma} = \frac{\text{COV} [\sigma_{VMF,t}^2, VMF_{t+1}]}{\text{VAR} [\sigma_{VMF,t}^2]} = \gamma \cdot \frac{\text{VAR} [\sigma_{BMF,t}^2] + \text{COV} [\sigma_{BMF,t}^2, \sigma_{IFF,t}^2]}{\text{VAR} [\sigma_{BMF,t}^2] + \text{VAR} [\sigma_{IFF,t}^2] + 2 \cdot \text{COV} [\sigma_{BMF,t}^2, \sigma_{IFF,t}^2]}, \quad (16)$$

where VAR is the unconditional variance operator, COV is the unconditional covariance operator, and $\sigma_{IFF,t}^2$ is the conditional IFF variance. The multiplicative term on γ induces a downward bias if $\text{VAR} [\sigma_{IFF,t}^2] \geq -\text{COV} [\sigma_{BMF,t}^2, \sigma_{IFF,t}^2]$. This condition holds if $\sigma_{BMF,t}^2$ and $\sigma_{IFF,t}^2$ are positively correlated, which is likely true empirically. This logic implies that, in the presence of the IFF, VMF-based estimates of γ are likely biased downward.

9.2 The BMF yields a stronger intertemporal risk-return relation

We next investigate whether empirical evidence supports the econometric intuition from the previous subsection. To this end, we estimate conditional VMF or BMF variance, σ_t^2 , using a GARCH(1,1) specification as

$$r_{t+1} = \phi_0 + \phi_1 \cdot r_t + \gamma \cdot \sigma_t^2 + u_{t+1}, \quad (17)$$

$$u_{t+1} = \sigma_t v_{t+1}, \quad (18)$$

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 u_t^2 + \alpha_2 \sigma_t^2, \quad (19)$$

where r_{t+1} is either VMF_{t+1} or BMF_{t+1} , u_t is a forecasting error from the previous period, v_{t+1} has unit variance, and γ is an estimate of the intertemporal risk-return relation. We find evidence of meaningful differences in the resulting BMF and VMF volatilities, thus giving scope for differences in the related γ estimates.³⁶

While we adopt the GARCH(1,1) as our baseline variance forecasting model, there are many GARCH extensions aimed at better capturing pervasive features in the data such as volatility persistence, asymmetry, and the leptokurtic distribution of financial asset returns.

³⁵This is true if we assume that the conditional covariance between the BMF and the IFF is zero and $\mathbb{E}_t[IFF_{t+1}] = 0$. Table 1 shows that the unconditional expectation of the IFF is not different than zero, and equation (14) implies that the unconditional covariance between the IFF and the BMF is zero. In unreported results we run rolling regressions of the VMF on the BMF and find that the resulting slope coefficient does not vary much implying that our unconditional BMF scaling in equation (14) approximately holds conditionally, implying that both of these assumed conditions hold approximately in the data.

³⁶We provide conditional volatility estimates from the baseline GARCH specification in Online Appendix Figure OA.14. As expected from the unconditional analysis in Table 1, the conditional BMF volatility is typically lower than the conditional VMF volatility. Additionally, the difference exhibits substantial variation at both high and low frequencies.

Thus, we complement our baseline GARCH(1,1) with estimates from two additional nonlinear GARCH models: the threshold GARCH (GJR-GARCH(1,1)) from Glosten et al. (1993) and the exponential GARCH (E-GARCH(1,1)) from Nelson (1991). These two models are widely used in the empirical literature to capture “leverage effects.”³⁷ In addition to using conditional variance as a proxy for time-varying risk as in equation (17), we also consider using conditional volatility as in Baillie and DeGennaro (1990) and French et al. (1987) by replacing σ_t^2 with σ_t . Additional details related to the models and estimation are provided in Online Appendix OA.9.

Table 7 reports γ estimates under the baseline GARCH (Panel A), the GJR-GARCH (Panel B) and the E-GARCH (Panel C). Each column reports results using a different market return proxy. The first column uses the VMF, the second column uses our main BMF specification, and the remainder use our alternative BMF specifications described in Section 5.3. The reported γ differences are differences between γ s estimated using the VMF and γ s estimated using different BMF specifications. The first set of results in each panel defines risk according to the variance specification in equation (17), and the second set of results defines risk according to the volatility specification when σ_t^2 is replaced by σ_t . The BMF generates positive and statistically significant γ estimates across all specifications. By contrast, the VMF generates positive estimates that are always statistically insignificant and lower than our BMF-based estimates. These results are consistent with the econometric intuition from the previous section. Namely, the IFF introduces unpriced risk into the VMF, leading to a weaker measured intertemporal risk-return relation.³⁸

10 Conclusion

We provide evidence that the standard value-weighted market factor contains an “idiosyncratic” component—the IFF—that has a mean which is statistically indistinguishable from zero and is unrelated to any of the proxies for investor marginal utility we consider from a

³⁷For reviews of this literature, see French et al. (1987), Schwert (1990), Franses and Van Dijk (1996), Poon and Granger (2003), and Brownlees et al. (2012).

³⁸Although we focus on the conditional mean-variance relationship of the stock market for brevity, the IFF’s effect on predictive regression instability is more general. Consistent with this, Schmidt (2025) shows that involuntary job loss and several standard predictors, including realized variance, forecast the BMF far more effectively than the VMF.

broad class of models. The IFF arises because the aggregate return on the selected sample of *public* firms loads on a systematic factor that diversifies away in the aggregate return on *all* firms. We show, both theoretically and empirically, how the existence of the IFF can be used to explain three long-standing asset pricing “puzzles” related to a general underestimation of discount rates under the CAPM, the size anomaly, and the weak intertemporal risk-return relation. Our unifying explanation is that the IFF adds unpriced risk to the VMF, which generates distortions in measured market risk relative to that under the BMF.

While value weighting has many advantages (e.g., lower rebalancing costs, theoretical justification from the CAPM), we uncover a potential disadvantage related to its tendency to overweight firms that are overrepresented in the public market. Although we focus on implications with respect to the VMF due to its historical importance and continued widespread use as a pricing factor, we do not fully explore implications for other pricing factors that traditionally use value weights. For instance, our results could provide a microfoundation for the need to control for size when constructing characteristic-efficient factors from the original characteristics-based value-weighted factors in Daniel et al. (2020). Furthermore, our mechanism could help explain the “puzzling” result that the CAPM works better on announcement days (Savor and Wilson, 2014) and that the CAPM works well for jumps in market returns but not for diffusive shocks (Bollerslev et al., 2016).³⁹ We leave these as potentially interesting paths for future research.

³⁹Our framework could explain these results if the conditional variance of the BMF relative to that of the IFF is larger on announcement/jump days, causing standard measures to more accurately reflect the quantity of risk on these days (both in the time series and cross section).

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used generative AI tools from OpenAI, Claude, and Cursor.ai in order to assist with copy editing, quality control, and model derivations. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

- ALQUIST, R., R. ISRAEL, AND T. MOSKOWITZ (2018): “Fact, Fiction, and the Size Effect,” *The Journal of Portfolio Management*, 45, 3–30.
- ASNESS, C., A. FRAZZINI, R. ISRAEL, T. J. MOSKOWITZ, AND L. H. PEDERSEN (2018): “Size matters, if you control your junk,” *Journal of Financial Economics*, 129, 479–509.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. V. REENEN (2020): “The Fall of the Labor Share and the Rise of Superstar Firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- AXTELL, R. L. (2001): “Zipf Distribution of U.S. Firm Sizes,” *Science*, 293, 1818–1820.
- BAI, J. AND S. NG (2006): “Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions,” *Econometrica*, 74, 1133–1150.
- BAILLIE, R. T. AND R. P. DEGENNARO (1990): “Stock returns and volatility,” *Journal of Financial and Quantitative Analysis*, 25, 203–214.
- BANSAL, R., D. KIKU, I. SHALIASTOVICH, AND A. YARON (2014): “Volatility, the macroeconomy, and asset prices,” *The Journal of Finance*, 69, 2471–2511.
- BANZ, R. W. (1981): “The relationship between return and market value of common stocks,” *Journal of financial economics*, 9, 3–18.
- BARILLAS, F. AND J. SHANKEN (2017): “Which Alpha?” *Review of Financial Studies*, 30, 1316–1338.
- BASAK, S. AND D. CUOCO (1998): “An equilibrium model with restricted stock market participation,” *The Review of Financial Studies*, 11, 309–341.
- BERK, J. B. (1995): “A Critique of Size-Related Anomalies,” *Review of Financial Studies*, 8, 275–286.
- BLACK, F. (1976): “Studies in stock price volatility changes, Proceedings of the 1976 Business Meeting of the Business and Economic Statistics Section,” *American Statistical Association*, 177–181.
- BLACK, F., M. C. JENSEN, AND M. S. SCHOLES (1972): “The capital asset pricing model: Some empirical tests,” *Studies in the Theory of Capital Markets*.
- BOLLERSLEV, T., R. F. ENGLE, AND J. M. WOOLDRIDGE (1988): “A capital asset pricing model with time-varying covariances,” *Journal of political Economy*, 96, 116–131.

- BOLLERSLEV, T., S. Z. LI, AND V. TODOROV (2016): “Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns,” *Journal of Financial Economics*, 120, 464–490.
- BROWNLEES, C., R. ENGLE, AND B. KELLY (2012): “A practical guide to volatility forecasting through calm and storm,” *Journal of Risk*, 14, 3.
- BYUN, S. (2016): “The usefulness of cross-sectional dispersion for forecasting aggregate stock price volatility,” *Journal of Empirical Finance*, 36, 162–180.
- CAMPBELL, J. Y. (1996): “Understanding risk and return,” *Journal of Political economy*, 104, 298–345.
- CAMPBELL, J. Y. AND L. HENTSCHEL (1992): “No news is good news: An asymmetric model of changing volatility in stock returns,” *Journal of Financial Economics*, 31, 281–318.
- CAMPBELL, S. D., S. DELIKOURAS, D. JIANG, AND G. M. KORNIOTIS (2016): “The human capital that matters: Expected returns and high-income households,” *The Review of Financial Studies*, 29, 2523–2563.
- CARHART, M. M. (1997): “On Persistence in Mutual Fund Performance,” *Journal of Finance*, 52, 57–82.
- CATHERINE, S. (2022): “Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle,” *Review of Financial Studies*, 35, 4016–4054.
- CEDERBURG, S. AND M. S. O’DOHERTY (2019): “Understanding the Risk-Return Relation: The Aggregate Wealth Proxy Actually Matters,” *Journal of Business and Economic Statistics*, 37, 721–735.
- CHABI-YO, F., A. GONCALVES, AND J. A. LOUDIS (2025): “An Intertemporal Risk Factor Model,” *Management Science*, 71, 6319–7222.
- CHAMBERLAIN, G. AND M. ROTHSCHILD (1983): “Factor Structure, and Mean-Variance Analysis on Large Asset Markets,” *Econometrica*, 51, 1281–1304.
- CHEMMANUR, T. J., S. HE, AND D. K. NANDY (2010): “The Going-Public Decision and the Product Market,” *Review of Financial Studies*, 23, 1855–1908.
- CHOD, J. AND E. LYANDRES (2011): “Strategic IPOs and Product Market Competition,” *Journal of Financial Economics*, 100, 45–67.
- CHOU, R. Y. (1988): “Volatility persistence and stock valuations: Some empirical evidence using GARCH,” *Journal of Applied Econometrics*, 3, 279–294.
- CHOU, R. Y., R. F. ENGLE, AND A. KANE (1992): “Measuring risk aversion from excess returns on a stock index,” *Journal of Econometrics*, 52, 201–224.
- CHRISTIE, A. A. (1982): “The stochastic behavior of common stock variances: Value, leverage and interest rate effects,” *Journal of Financial Economics*, 10, 407–432.
- COCHRANE, J. H. (2009): *Asset pricing: Revised edition*, Princeton university press.
- (2017): “Macro-Finance,” *Review of Finance*, 21, 945–985.
- CONNOR, G. AND R. A. KORAJCZYK (1986): “Performance measurement with the arbitrage pricing theory: A new framework for analysis,” *Journal of Financial Economics*, 15, 373–394.
- (1988): “Risk and return in an equilibrium APT: Application of a new test methodology,” *Journal of Financial Economics*, 21, 255–289.
- CONSTANTINIDES, G. M. AND D. DUFFIE (1996): “Asset pricing with heterogeneous consumers,” *Journal of Political economy*, 104, 219–240.


- CROUZET, N. AND N. R. MEHROTRA (2020): “Small and Large Firms over the Business Cycle,” *American Economic Review*, 110, 3549–3601.
- DANIEL, K., L. MOTA, S. ROTTKE, AND T. SANTOS (2020): “The Cross-Section of Risk and Return,” *Review of Financial Studies*, 33, 1927–1979.
- DAVIS, S. J. AND J. HALTIWANGER (1992): “Gross Job Creation, Gross Job Destruction, and Employment Reallocation,” *The Quarterly Journal of Economics*, 107, 819–863.
- DECAIRE, P. H. AND J. R. GRAHAM (2024): “Valuation Fundamentals,” Working paper.
- DOIDGE, C., G. A. KAROLYI, AND R. M. STULZ (2017): “The U.S. listing gap,” *Journal of Financial Economics*, 123, 464–487.
- DONANGELO, A., F. GOURIO, M. KEHRIG, AND M. PALACIOS (2019): “The cross-section of labor leverage and equity returns,” *Journal of Financial Economics*, 132, 497–518.
- DOU, W. W., L. KOGAN, AND W. WU (2024): “Common Fund Flows: Flow Hedging and Factor Pricing,” *Journal of Finance, Forthcoming*.
- FAMA, E. F. AND K. R. FRENCH (1992): “The cross-section of expected stock returns,” *Journal of Finance*, 47, 427–465.
- (1993): “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, 33, 3–56.
- (2004): “The Capital Asset Pricing Model: Theory and Evidence,” *Journal of Economic Perspectives*, 18, 25–46.
- (2015): “A Five-Factor Asset Pricing Model,” *Journal of Financial Economics*, 116, 1–22.
- FAMA, E. F. AND J. D. MACBETH (1973): “Risk, return, and equilibrium: Empirical tests,” *Journal of Political Economy*, 81, 607–636.
- FRANSES, P. H. AND D. VAN DIJK (1996): “Forecasting stock market volatility using (nonlinear) GARCH models,” *Journal of Forecasting*, 229–235.
- FRENCH, K. R., G. W. SCHWERT, AND R. F. STAMBAUGH (1987): “Expected stock returns and volatility,” *Journal of Financial Economics*, 19, 3–29.
- GABAIX, X. (2016): “Power Laws in Economics: An Introduction,” *Journal of Economic Perspectives*, 30, 185–206.
- GERTLER, M. AND S. GILCHRIST (1994): “Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms,” *The Quarterly Journal of Economics*, 109, 309–340.
- GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2005): “There is a risk-return trade-off after all,” *Journal of Financial Economics*, 76, 509–548.
- GIBBONS, M. R., S. A. ROSS, AND J. SHANKEN (1989): “A test of the efficiency of a given portfolio,” *Econometrica*, 57, 1121–1152.
- GIGLIO, S. AND D. XIU (2021): “Asset Pricing with Omitted Factors,” *Journal of Political Economy*, 129, 1947–1990.
- GLOSTEN, L. R., R. JAGANNATHAN, AND D. E. RUNKLE (1993): “On the relation between the expected value and the volatility of the nominal excess return on stocks,” *Journal of Finance*, 48, 1779–1801.
- GORMSEN, N. J. AND K. HUBER (2024): “Firms’ perceived cost of capital,” Tech. rep., National Bureau of Economic Research.
- GRAHAM, J. R. (2022): “Presidential Address: Corporate Finance and Reality,” *Journal of Finance*, 77, 1975–2049.

- GRAHAM, J. R. AND C. R. HARVEY (2001): “The theory and practice of corporate finance: evidence from the field,” *Journal of Financial Economics*, 60, 187–243.
- GUIO, L., L. PISTAFERRI, AND F. SCHIVARDI (2005): “Insurance within the Firm,” *Journal of Political Economy*, 113, 1054–1087.
- GUVENEN, F., S. OZKAN, AND J. SONG (2014): “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122, 621–660.
- HAMILTON, J. D. (1994): *Time Series Analysis*, Princeton, New Jersey: Princeton University Press.
- HARTMAN-GLASER, B., H. LUSTIG, AND M. Z. XIAOLAN (2019): “Capital Share Dynamics When Firms Insure Workers,” *The Journal of Finance*, 74, 1707–1751.
- HARVEY, C. R. (1989): “Time-varying conditional covariances in tests of asset pricing models,” *Journal of Financial Economics*, 24, 289–317.
- HE, Z., B. KELLY, AND A. MANELA (2017): “Intermediary asset pricing: New evidence from many asset classes,” *Journal of Financial Economics*, 126, 1–35.
- HEATON, J. AND D. LUCAS (2000): “Portfolio choice and asset prices: The importance of entrepreneurial risk,” *Journal of Finance*, 55, 1163–1198.
- HERSKOVIC, B., B. KELLY, H. LUSTIG, AND S. VAN NIEUWERBURGH (2016): “The common factor in idiosyncratic volatility: Quantitative asset pricing implications,” *Journal of Financial Economics*, 119, 249–283.
- HOU, K. AND M. A. VAN DIJK (2019): “Resurrecting the Size Effect: Firm Size, Profitability Shocks, and Expected Stock Returns,” *Review of Financial Studies*, 32, 2850–2889.
- HOU, K., C. XUE, AND L. ZHANG (2015): “Digesting Anomalies: An Investment Approach,” *Review of Financial Studies*, 28, 650–705.
- (2021): “An Augmented q-Factor Model with Expected Growth,” *Review of Finance*, 25, 1–41.
- JAGANNATHAN, R. AND Z. WANG (1996): “The Conditional CAPM and the Cross-Section of Expected Returns,” *Journal of Finance*, 51, 3–53.
- JENSEN, T. I., B. KELLY, AND L. H. PEDERSEN (2023): “Is There a Replication Crisis in Finance?” *Journal of Finance*, 78, 2465–2518.
- JORDÀ, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American Economic Review*, 95, 161–182.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KELLY, B. T., S. PRUITT, AND Y. SU (2019): “Characteristics are covariances: A unified model of risk and return,” *Journal of Financial Economics*, 134, 501–524.
- KIM, W. AND M. S. WEISBACH (2008): “Motivations for public equity offers: An international perspective,” *Journal of Financial Economics*, 87, 281–307.
- KOGAN, L., D. PAPANIKOLAOU, A. SERU, AND N. STOFFMAN (2017): “Technological Innovation, Resource Allocation, and Growth,” *The Quarterly Journal of Economics*, 132, 665–712.
- KOOPMAN, S. J. AND E. HOL USPENSKY (2002): “The stochastic volatility in mean model: empirical evidence from international stock markets,” *Journal of Applied Econometrics*, 17, 667–689.
- KOZAK, S., S. NAGEL, AND S. SANTOSH (2018): “Interpreting factor models,” *Journal of Finance*, 73, 1183–1223.

- KWON, S. Y., Y. MA, AND K. ZIMMERMANN (2024): “100 Years of Rising Corporate Concentration,” *American Economic Review*, 114, 2111–2140.
- LETTAU, M. AND M. PELGER (2020): “Factors That Fit the Time Series and Cross-Section of Stock Returns,” *Review of Financial Studies*, 33, 2274–2325.
- LUCAS, ROBERT E., J. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 631–653.
- LUDVIGSON, S. C. AND S. NG (2007): “The empirical risk–return relation: A factor analysis approach,” *Journal of Financial Economics*, 83, 171–222.
- LUSTIG, H., S. VAN NIEUWERBURGH, AND A. VERDELHAN (2013): “The wealth–consumption ratio,” *The Review of Asset Pricing Studies*, 3, 38–94.
- LUTTMER, E. G. J. (2010): “Models of Growth and Firm Heterogeneity,” *Annual Review of Economics*, 2, 547–576.
- MAKSIMOVIC, V. AND P. PICHLER (2001): “Technological Innovation and Initial Public Offerings,” *The Review of Financial Studies*, 14, 459–494.
- MALLOY, C. J., T. J. MOSKOWITZ, AND A. VISSING-JØRGENSEN (2009): “Long-run stockholder consumption risk and asset returns,” *Journal of Finance*, 64, 2427–2479.
- MANKIW, N. G. (1986): “The Equity Premium and the Concentration of Aggregate Shocks,” *Journal of Financial Economics*, 17, 211–219.
- MANKIW, N. G. AND S. P. ZELDES (1991): “The consumption of stockholders and non-stockholders,” *Journal of Financial Economics*, 29, 97–112.
- MAYERS, D. (1973): “Nonmarketable assets and the determination of capital asset prices in the absence of a riskless asset,” *The Journal of Business*, 46, 258–267.
- NELSON, D. B. (1991): “Conditional heteroskedasticity in asset returns: A new approach,” *Econometrica*, 347–370.
- NEWBY, W. K. AND K. D. WEST (1987): “A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix,” *Econometrica*, 703–708.
- (1994): “Automatic Lag Selection in Covariance Matrix Estimation,” *Review of Economic Studies*, 61, 631–653.
- PAGAN, A. R. AND Y. HONG (1991): “Nonparametric estimation and the risk premium,” *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, W. Barnett, J. Powell and GE Tauchen (eds), 51–75.
- PANAGEAS, S. ET AL. (2020): “The implications of heterogeneity and inequality for asset pricing,” *Foundations and Trends® in Finance*, 12, 199–275.
- PAPANIKOLAOU, D. (2011): “Investment Shocks and Asset Prices,” *Journal of Political Economy*, 119, 639–685.
- PARKER, J. A. AND A. VISSING-JØRGENSEN (2009): “Who bears aggregate fluctuations and how?” *American Economic Review*, 99, 399–405.
- POON, S. H. AND C. W. GRANGER (2003): “Forecasting volatility in financial markets: A review,” *Journal of Economic Literature*, 41, 478–539.
- PUKTHUANThONG, K., R. ROLL, AND A. SUBRAHMANYAM (2019): “A protocol for factor identification,” *Review of Financial Studies*, 32, 1573–1607.
- ROLL, R. (1977): “A critique of the asset pricing theory’s tests Part I: On past and potential testability of the theory,” *Journal of financial economics*, 4, 129–176.
- ROSS, S. A. (1976): “The arbitrage theory of capital asset pricing,” *Journal of Economic Theory*, 13, 341–360.

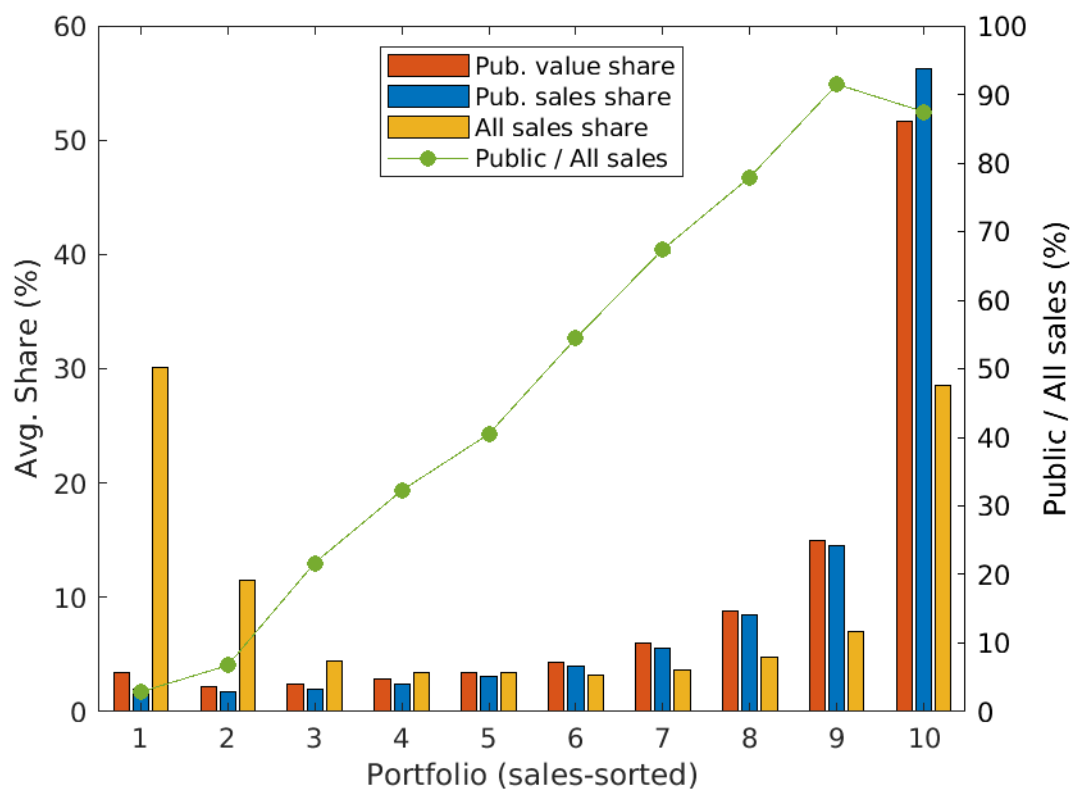
- (1977): “The capital asset pricing model (capm), short-sale restrictions and related issues,” *Journal of Finance*, 32, 177–183.
- SAVOR, P. AND M. WILSON (2014): “Asset pricing: A tale of two days,” *Journal of Financial Economics*, 113, 171–201.
- SCHMIDT, L. D. (2025): “Climbing and falling off the ladder: Asset pricing implications of labor market event risk,” *Journal of Financial Economics*, 172, 104131.
- SCHWERT, G. W. (1990): “Stock returns and real activity: A century of evidence,” *Journal of Finance*, 45, 1237–1257.
- STAMBAUGH, R. F. (1982): “On the Exclusion of Assets from Tests of the Two-Parameter Model: A Sensitivity Analysis,” *Journal of Financial Economics*, 10, 237–268.
- STAMBAUGH, R. F. AND Y. YUAN (2017): “Mispricing Factors,” *Review of Financial Studies*, 30, 1270–1315.
- STOCK, J. H. AND M. W. WATSON (2002): “Forecasting using principal components from a large number of predictors,” *Journal of the American Statistical Association*, 97, 1167–1179.
- VAN BINSBERGEN, J. H., M. BOONS, C. C. OPP, AND A. TAMONI (2023): “Dynamic asset (mis)pricing: Build-up versus resolution anomalies,” *Journal of Financial Economics*, 147, 406–431.
- VAN BINSBERGEN, J. H. AND C. C. OPP (2019): “Real Anomalies,” *Journal of Finance*, 74, 1659–1706.
- VAN DIJK, M. A. (2011): “Is size dead? A review of the size effect in equity returns,” *Journal of Banking and Finance*, 35, 3263–3274.
- WHITELAW, R. F. (1994): “Time variations and covariations in the expectation and volatility of stock market returns,” *Journal of Finance*, 49, 515–541.

Figure 1: Value share of largest stocks in the CRSP value-weighted index



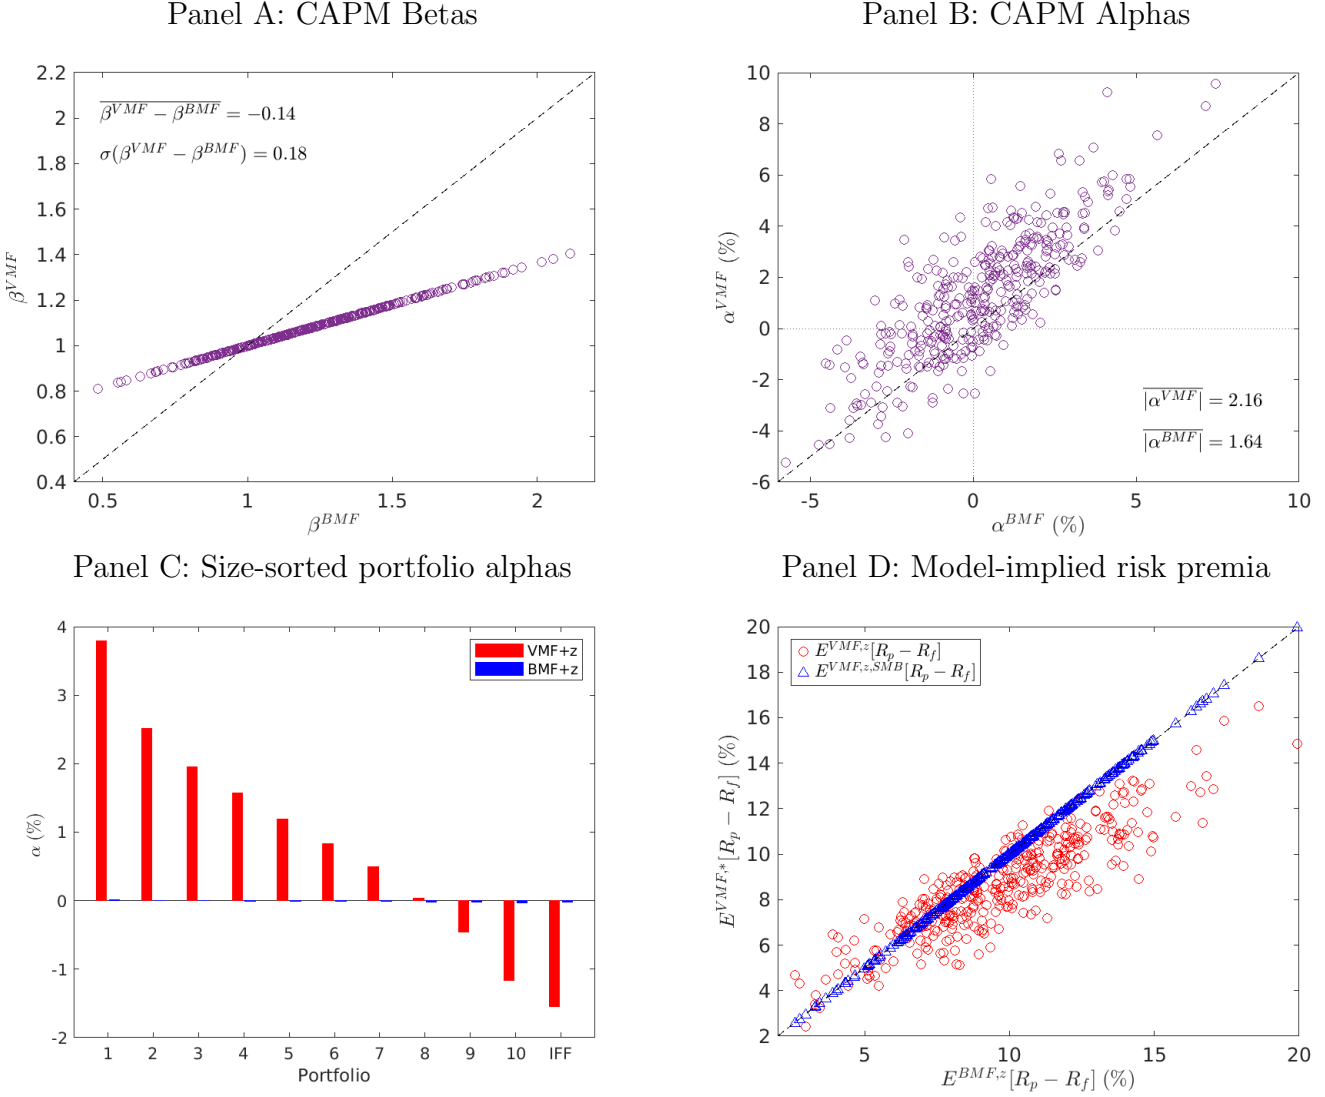
Note. This figure plots the cumulative value shares of individual stocks for the 5, 10, 25, and 50 largest stocks in the CRSP value-weighted index (left axis). The top 2.5% line, also on the left axis, plots the weight associated with stocks whose market capitalization puts them among the largest stocks with the rank cutoff based on 2.5% of the total firms in the NYSE. We use the NYSE breakpoint to avoid mechanical decreases in the series as many (mostly small) AMEX and NASDAQ firms are added to the database over time. The right axis plots the total number of stocks included in the CRSP index as well as in the NYSE only, where the former includes two large jumps as the CRSP sample coverage expands. See Section 3.1 for additional details.

Figure 2: Sales-sorted portfolio shares in the aggregate market



Note. This figure plots average sales and value shares for 10 sales-sorted portfolios of either public or all firms that use NYSE breakpoints. The red bars plot value shares among portfolios of public firms in the aggregate market for public firms. The blue bars plot sales shares among portfolios of public firms in the aggregate market for public firms. The yellow bars plot sales shares among portfolios of all firms in the aggregate market for all firms. The green line plots the fraction of sales accounted for by public firms in each bin relative to the total sales among all firms in the bin. The all-firm sales shares (in yellow) are calculated using data from Kwon et al. (2024) according to the industry-sales double-sorting methodology described in Section 5.2 with more details provided in Online Appendix OA.5.3. Data are from 1963–2021. See Section 3.2 for additional commentary.


Figure 3: Model implications




Note: This figure highlights four main implications from our theoretical model based on simulation results. Panel A plots VMF betas against BMF betas. Panel B plots VMF alphas against BMF alphas and reports their mean absolute values. Panel C plots size-sorted portfolio alphas along with IFF alphas estimated with respect to either the VMF or BMF (controlling for z in both cases). Panel D plots model implied risk premia from two VMF-based factor models against those from a BMF-based model, with all models controlling for z in this case. One of the VMF-based models additionally controls for the model's size factor, SMB, defined in equation (9). Note that the BMF-based model correctly prices assets in our theoretical model. See Section 4 for additional details.

Figure 4: Value shares in aggregate market over time


Panel A: Sales-sorted portfolios 1 & 10 (public)




Panel B: Sales-sorted portfolios 1 & 10 (all)



Panel C: Industry portfolios (public)

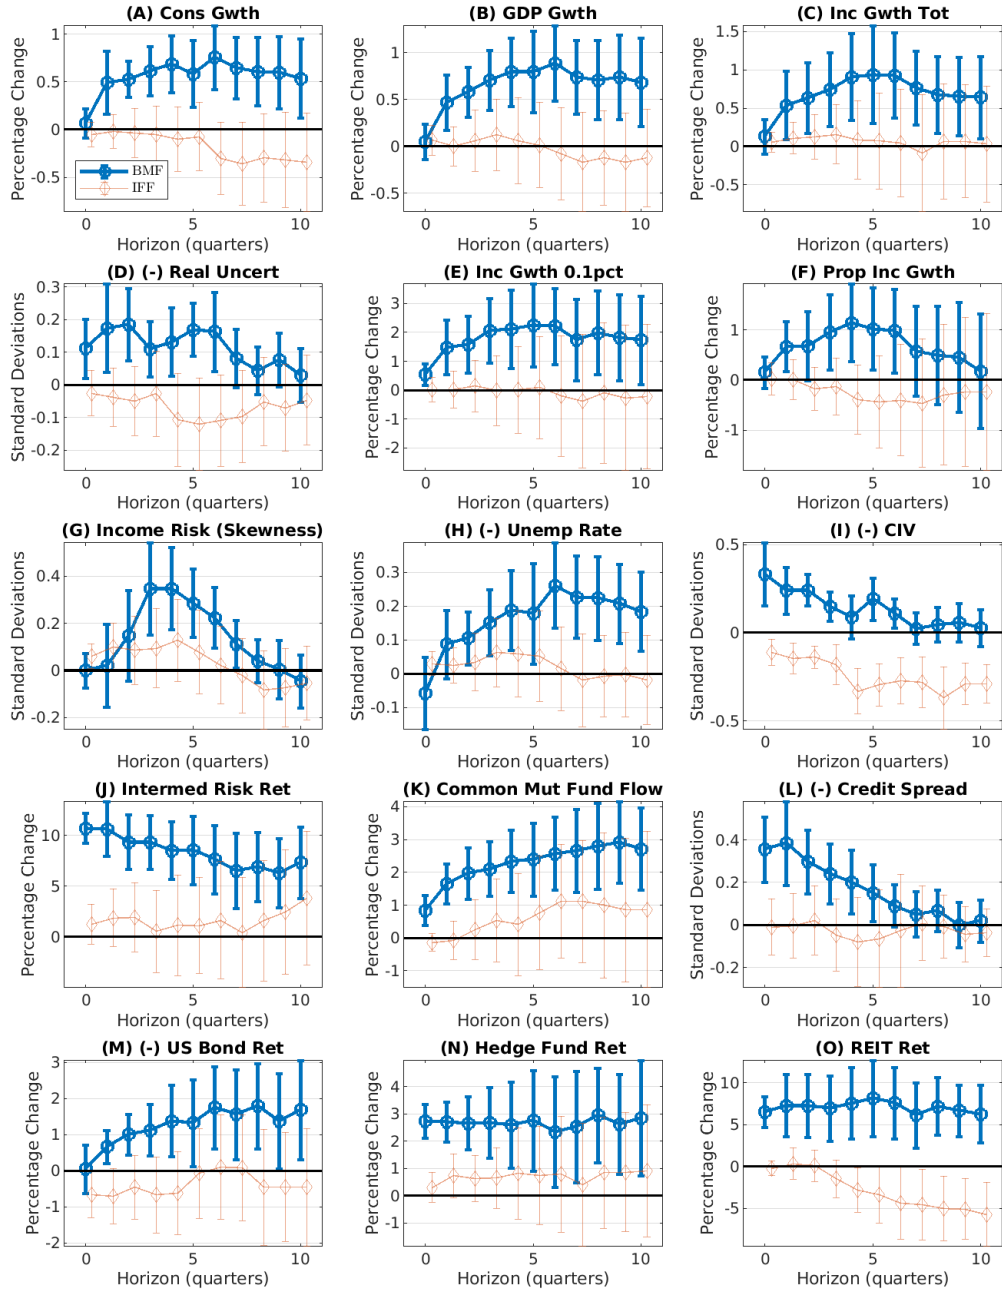


Panel D: Industry portfolios (all)



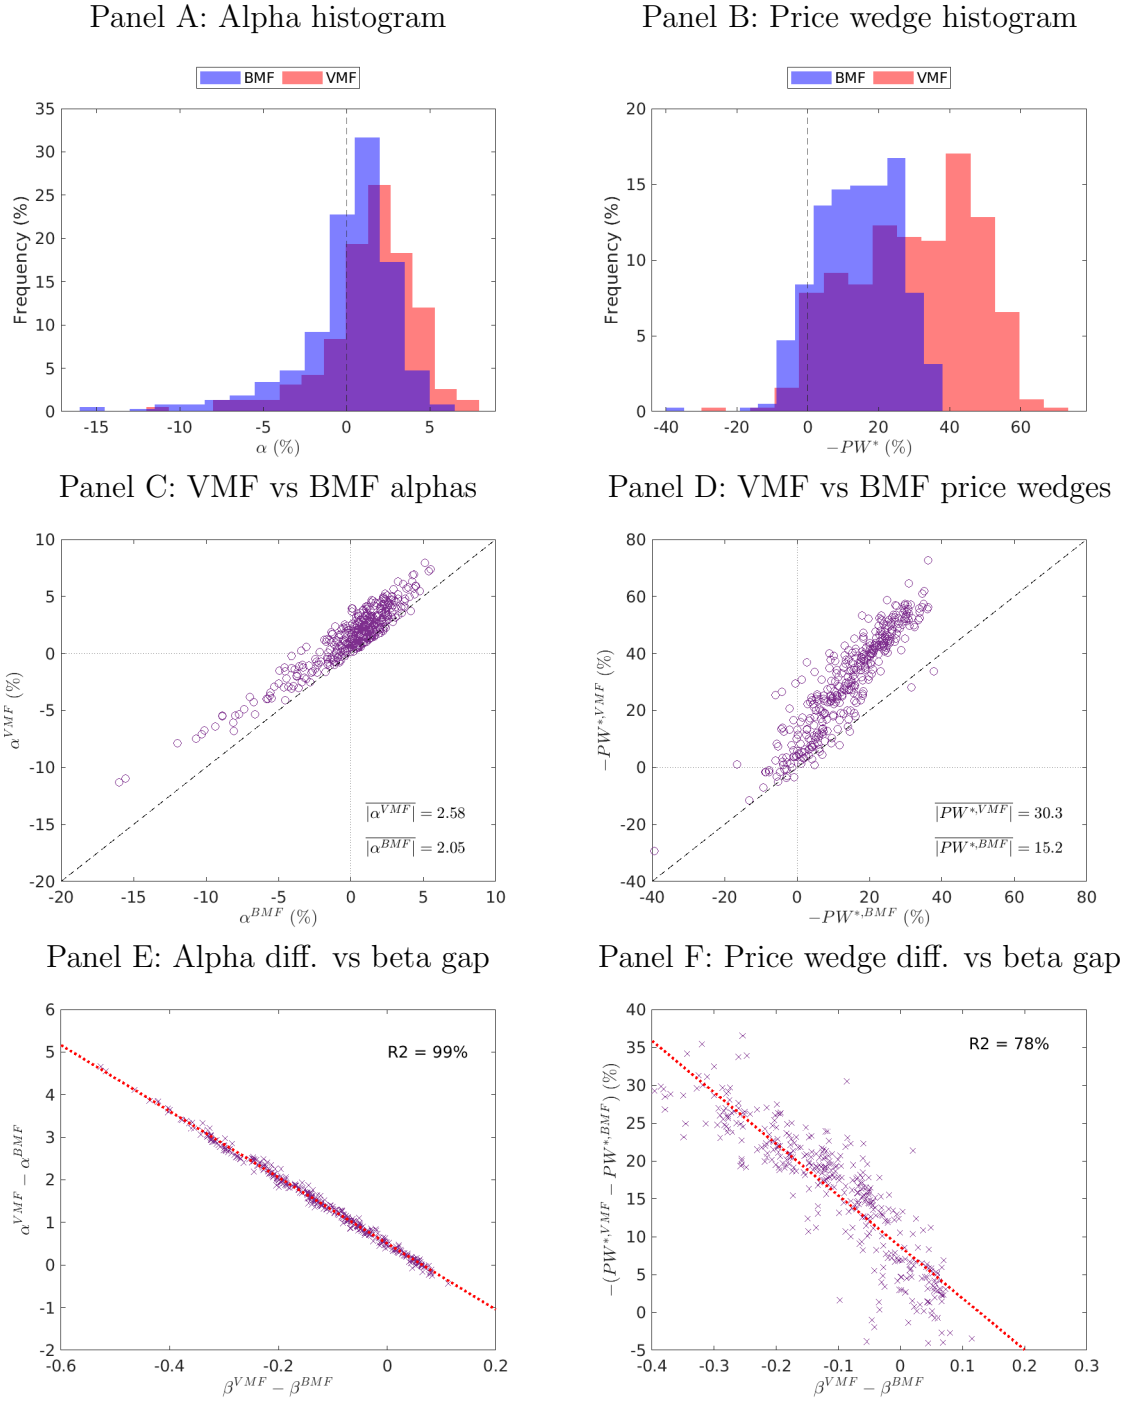
Note: This figure plots the value shares in the aggregate market for public firms and all firms (public and private) over time. Panels A and B plot the value shares of sales-sorted portfolios 1 and 10 (using NYSE breakpoints) among public and all firms, respectively. Panels C and D plot the value shares by industry among public and all firms, respectively. Value shares for all firms are calculated according to the methodology described in Section 5 and by double-sorting firms into portfolios based on industry then within-industry sales using NYSE breakpoints (i.e., according to our main BMF specification). In particular, stock-level implied value weights are calculated according to equation (13) and then aggregated by sales (Panel B) or industry (Panel D). Data are from 1963–2021. See Section 5 for additional details on the construction of these measures and Section 5.2 for specific commentary.

Figure 5: Local projection impulse responses to the BMF and IFF



Note. This figure plots local projection impulse responses of selected macroeconomic aggregates and financial return time series to the BMF and IFF. The indexes are defined in Section 6.3 and are signed so that increases in the variable value correspond to decreases in marginal utility in standard models. A “(-)” in the panel title indicates that the negative value of a time series is used to conform to this logic. Monthly returns are compounded to quarterly frequency to match that of the macroeconomic aggregates. The point estimates represent the impact of a standard deviation increase in the predictor variable on each response variable. Response variables are in units of percent in the case of return or growth time series (e.g., consumption growth, US bond returns, etc.) or units of standard deviation in the case of stationary time series (e.g., unemployment rate, real uncertainty, etc.). The initial point estimates (i.e., when the horizon is equal to zero) correspond to the contemporaneous response of each macroeconomic aggregate to each predictor variable. The error bars indicate 95% confidence intervals computed using Newey-West standard errors with 10 lags. Data are from 1963–2021 (conditional on variable availability). See Section 6.3 for additional details.

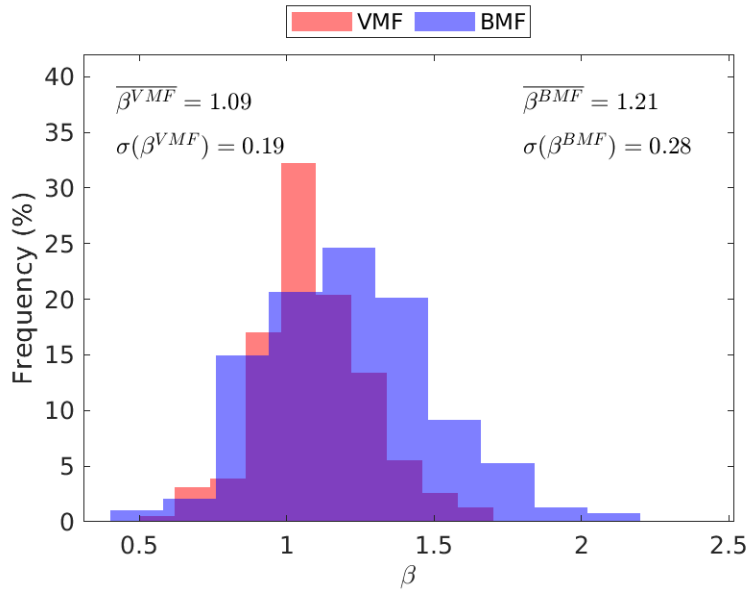
Figure 6: VMF versus BMF alphas and price wedges




Note. This figure plots VMF versus BMF alphas and (negative) price wedges (van Binsbergen et al., 2023) estimated for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3. We plot negative price wedges so that positive values correspond with model-implied underpricing, similar to the standard sign convention for alphas. Panel A plots alpha histograms. Panel B plots price wedge histograms. Panel C plots VMF alphas against BMF alphas and reports their mean absolute values. Panel D plots VMF price wedges against BMF price wedges and reports their mean absolute values. Panel E plots alpha differences against beta differences. Panel F plots price wedge differences against beta differences. The dotted red lines in Panels E and F reflect OLS best-fit lines for the data. Data are from 1963–2021. See Section 7.1 for additional details.

Figure 7: Betas under the VMF versus BMF

Panel A: Beta histograms (VMF vs BMF)




Panel B: VMF versus BMF betas




Note. This figure plots univariate VMF and BMF betas estimated for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3. Panel A plots histograms for the VMF and BMF betas. Panel B plots the VMF betas against the BMF betas with the dotted 45-degree line to help with the comparison. The dotted red line is the OLS best-fit for the data. The dotted blue line is the OLS best-fit line for our theoretical model's implied VMF and BMF betas (based on results reported in Figure 3 Panel A). Data are from 1963–2021. See Section 7.2 for additional details.

Figure 8: IFF betas and size


Panel A: β^{IFF} vs avg. log market equity




Panel B: β^{IFF} vs $-\beta^{SMB}$



Panel C: $\beta^{VMF} - \beta^{BMF}$ vs avg. log market equity




Panel D: $\beta^{VMF} - \beta^{BMF}$ vs $-\beta^{SMB}$



Note. This figure plots univariate IFF betas and beta gaps (the difference between univariate VMF and BMF betas) against two portfolio size metrics for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3. Panel A plots IFF betas against each portfolio's log size (computed as the average monthly log market capitalization over time). Panel B plots IFF betas against each portfolio's negative univariate size factor loading where the size factor we use is SMB from the Fama-French three-factor model (Fama and French, 1993). Panel C plots beta gaps against each portfolio's log size. Panel D plots beta gaps against each portfolio's negative univariate size factor loading ($-\beta^{SMB}$) is increasing as a portfolio's returns behave more like big stocks in the short leg of the SMB factor. The dotted red lines are the OLS best-fit lines, and the reported r-squared values are associated with these. The dotted blue lines are OLS best-fit lines for our theoretical model's implied IFF betas and beta gaps plotted against either average log market capitalization (Panels A and C) or negative model-implied SMB betas (Panels B and D). In the latter case, model-implied SMB betas are estimated with respect to a model-implied SMB factor that we define in equation (9). Data are from 1963–2021. See Section 7.2 for additional details.

Figure 9: Distribution of VMF versus BMF alphas by portfolio size





Note. This figure plots VMF and BMF alphas estimated for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 binned into three groups based on each portfolio's log size (computed as the average monthly log market capitalization over time). The horizontal red and blue lines represent the median values within each bin. The lower and upper whiskers represent the 10th and 90th percentiles within each bin, respectively. Data are from 1963–2021. See Section 7.2 for additional details.

Figure 10: Size-sorted decile portfolio alphas

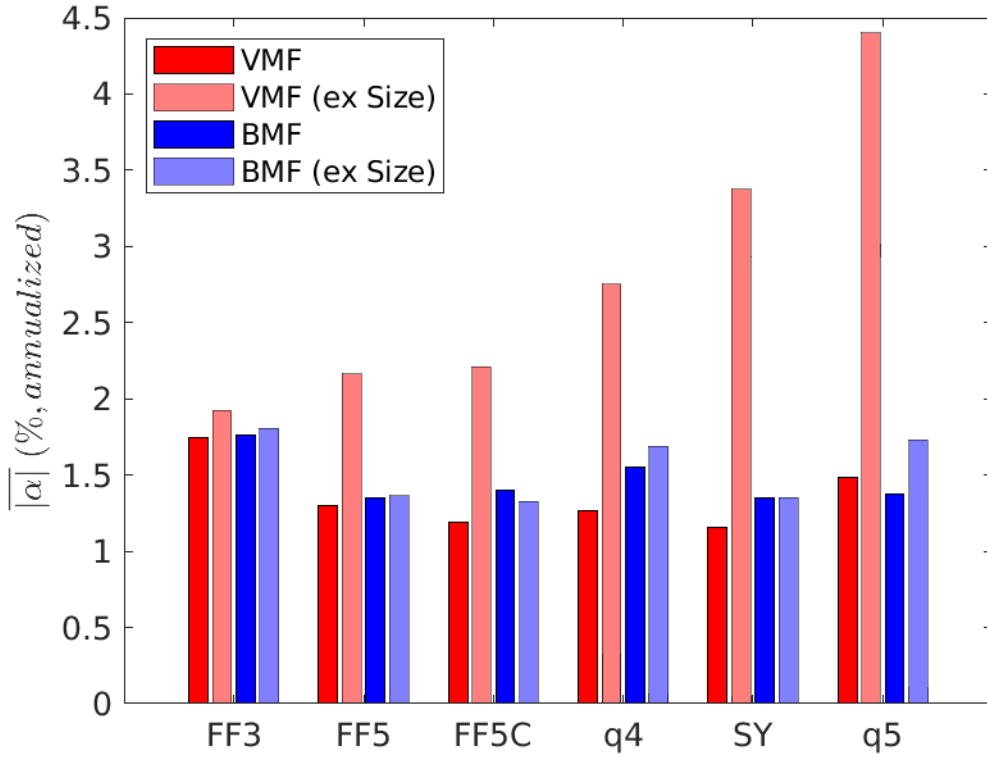
Panel A: VMF-based models

Panel B: BMF-based models



Note. This figure plots alphas for value-weighted size-sorted decile portfolios (a subset of our set of 382 characteristics-sorted portfolios described in Online Appendix Table OA.3) for each of the nine factor models described in Section 6.2. We drop size factors from the six models with size factors (FF3, FF5, FF5C, q4, SY, and q5) in this analysis because, consistent with our theoretical model and our empirical results, size factors help correctly price size-sorted portfolios when the VMF is used as the market factor (Panel A), but add no pricing power when the BMF is used (Panel B). The error bars represent 90% confidence intervals on each alpha estimate based on Newey and West (1987, 1994) standard errors. Data are from 1963–2021 (conditional on factor model availability). See Section 8.1 for additional details.

Figure 11: Mean absolute alphas for a large number of test assets




Note. This figure plots the mean absolute alphas for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 as test assets across the six standard factor models described in Section 6.2 that include size factors. These models are plotted in chronological order based on publication date. We study four versions of each model: 1) “VMF”: Models that use the VMF as the market factor and include size factors, 2) “VMF (ex Size)”: Models that use the VMF as the market factor and exclude size factors, 3) “BMF”: Models that use the BMF as the market factor and include size factors, and 4) “BMF (ex Size)”: Models that use the BMF as the market factor and exclude size factors. Data are from 1963–2021 (conditional on factor availability). See Section 8.2 for additional details.

Figure 12: Size factors and VMF- versus BMF-based model-implied risk premia

Panel A: VMF vs BMF models (ex size) Panel B: VMF vs BMF models (w/ size) Panel C: BMF (w/ size) vs BMF (ex size)

56



Note. This figure plots factor model-implied risk premia for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 for the six standard factor models described in Section 6.2 that include size factors. Comparisons are made between models that use either the VMF or BMF as the market factor, and models that either include or exclude the size factors. Panel A plots results for VMF- versus BMF-based models when size factors are excluded from the models. Panel B plots results for VMF- versus BMF-based models with size factors included in the VMF-based models but excluded from the BMF-based models. Panel C plots results for BMF-based models that either include size factors (y axis) or exclude size factors (x axis). The dotted red lines are OLS best-fit lines across all models and portfolios. The dotted blue lines are relationships implied by our theoretical model. For Panel A, we estimate this relationship as an OLS best-fit line for the theoretical model's corresponding risk premia documented in Figure 3 Panel D (i.e., the red circular markers). For Panels B and C, the model implies that the VMF and BMF-based risk premia fall on the 45-degree line, which is what we plot in these cases. Data are from 1963–2021 (conditional on factor availability). See Section 8.3 for additional details.

Table 1: Summary statistics for the VMF, BMF, and IFF

	Statistic						Correlations					
	Mean (%)	t(Mean)	Stdev. (%)	Sharpe ratio	Skew.	Kurt.	<i>VMF</i>	Ind- sales	Sales- only	EW PCA 1	Cap VW PCA 1	IPCA 1
Panel A: Market portfolio proxies												
<i>VMF</i>	6.79***	[3.23]	15.36	0.44	-0.56	5.07	1.00	0.91	0.90	0.92	0.95	0.84
<i>BMF</i> (ind-sales)	7.30***	[3.54]	13.99	0.52	-0.35	5.76	0.91	1.00	1.00	0.99	0.98	0.98
<i>BMF</i> (sales-only)	7.34***	[3.69]	13.77	0.53	-0.31	5.90	0.90	1.00	1.00	0.99	0.98	0.98
<i>BMF</i> (EW PCA 1)	7.31***	[3.68]	14.09	0.52	-0.34	5.84	0.92	0.99	0.99	1.00	0.99	0.97
<i>BMF</i> (cap VW PCA 1)	7.13***	[3.62]	14.52	0.49	-0.50	5.57	0.95	0.98	0.98	0.99	1.00	0.94
<i>BMF</i> (IPCA 1)	7.04***	[3.45]	13.16	0.53	0.09	6.00	0.84	0.98	0.98	0.97	0.94	1.00
Panel B: IFF proxies												
<i>IFF</i> (ind-sales)	-0.51	[-0.54]	6.34	-0.08	-0.50	6.07	0.41	1.00	0.98	0.96	0.88	0.95
<i>IFF</i> (sales-only)	-0.55	[-0.55]	6.82	-0.08	-0.39	5.80	0.44	0.98	1.00	0.97	0.90	0.96
<i>IFF</i> (EW PCA 1)	-0.52	[-0.55]	6.13	-0.09	-0.37	5.24	0.40	0.96	0.97	1.00	0.96	0.95
<i>IFF</i> (cap VW PCA 1)	-0.35	[-0.43]	5.02	-0.07	-0.25	4.53	0.33	0.88	0.90	0.96	1.00	0.84
<i>IFF</i> (IPCA 1)	-1.31	[-1.03]	8.42	-0.16	-0.49	6.27	0.54	0.95	0.96	0.95	0.84	1.00

Note: This table reports summary statistics for monthly excess returns on the VMF, BMF variants, and respective IFFs. Panel A reports results for the VMF and BMF variants. Panel B reports results for the IFF variants. We consider five different variants of the BMF: 1) “BMF (ind-sales)” is our main BMF measure constructed using the industry-sales-based double-sorting procedure described in Section 5.2, 2) “BMF (sales-only)” is an alternative BMF specification constructed using the sales-based single-sorting procedure described in Section 5.3 (note that the first two BMF specifications are based on our public-to-all-firm reweighting methodology described in Section 5.1, but using different portfolio sorts), 3) “BMF (EW PCA 1)” is an alternative that projects our main BMF onto the first PC from a PCA using our set of 382 equal-weighted portfolios described in Online Appendix Table OA.3, 4) “BMF (cap VW PCA 1)” is an alternative that projects our main BMF onto the first PC from a PCA using our set of 382 capped value-weighted portfolios described in Online Appendix Table OA.3, and 5) “BMF (IPCA 1)” is an alternative that projects our main BMF onto the first instrumented principal component from Kelly et al. (2019). Additional details on our alternative BMFs can be found in Section 5.3. The IFFs are computed by subtracting the respective BMF variants from the VMF. We report the mean, standard deviation, skewness, and kurtosis for each variable. We also provide correlation matrices among the VMF/BMF and the VMF/IFF. All statistics are annualized and in percent. We annualize the mean and standard deviation statistics by multiplying monthly frequency values by 12 and $\sqrt{12}$, respectively. T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance of the mean values at 90%, 95%, and 99%, respectively. Data are from 1963–2021, except for the IPCA-based BMF/IFF, for which we are limited to the 1964–2014 from Kelly et al. (2019) when the average VMF value was 5.73% (this is what yields a more negative average IFF in this case compared to the others despite the BMF having a similar average as other BMFs).

Table 2: IFF spanning tests

	CAPM	FF3	FF5	FF5C	q4	SY	DHS	q5	ICAPM
Panel A: VMF as the market factor (ex size factors) [Alternate interpretation: BMF spanning test]									
$\alpha^{VMF, exSize}$	-1.67**	-1.10	-2.18***	-3.00***	-3.53***	-4.24***	-3.51***	-5.19***	-2.20***
	[-2.04]	[-1.38]	[-2.73]	[-3.52]	[-3.86]	[-4.51]	[-3.20]	[-5.68]	[-2.76]
Panel B: BMF as the market factor (ex size factors) [Alternate interpretation: VMF spanning test]									
$\alpha^{BMF, exSize}$	-0.51	0.19	-0.19	-0.86	-1.46	-0.64	-0.44	-2.00**	-0.15
	[-0.65]	[0.24]	[-0.24]	[-1.02]	[-1.62]	[-0.69]	[-0.44]	[-2.13]	[-0.19]
Panel C: VMF as the market factor (with size factors) [Alternate interpretation: BMF spanning test]									
$\alpha^{VMF, wSize}$	-	-0.42	-0.66	-1.55***	-1.61**	-1.24**	-	-2.15***	-
	-	[-1.06]	[-1.32]	[-3.28]	[-2.21]	[-2.13]	-	[-3.12]	-
Panel D: BMF as the market factor (with size factors) [Alternate interpretation: VMF spanning test]									
$\alpha^{BMF, wSize}$	-	-0.03	-0.06	-1.26**	-1.05	0.00	-	-1.11	-
	-	[-0.07]	[-0.11]	[-2.27]	[-1.23]	[0.00]	-	[-1.38]	-

Note: This table reports IFF factor spanning test results (alphas) for each of our nine standard factor models described in Section 6.2. Panel A reports results when the VMF is used as the market factor in each model with size factors excluded from the spanning regressions (for models that have size factors). Panel B reports results when the BMF is used with size factors excluded from the spanning regressions (for models that have size factors). Panel C reports results when the VMF is used as the market factor in each model only for models that include size factors with size factors included in the spanning regressions. Panel D reports results when the BMF is used as the market factor in each model only for models that include size factors with size factors included in the spanning regressions. An alternative interpretation of the results in Panels A and C are that they reflect VMF spanning tests, whereas those from Panels B and D reflect BMF spanning tests (with alphas from those tests having the opposite sign but same t-statistics as those reported). Alphas are annualized and in percent. T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. Data are from 1963–2021 (conditional on factor availability). See Section 6.2 for additional details.

Table 3: VMF versus BMF alpha and price wedge summary statistics

Operator (\mathcal{O}):	All portfolios				Above median beta gap ports.			
	Avg.	q25	Med.	q75	Avg.	q25	Med.	q75
Panel A: Alphas								
$\mathcal{O}(\alpha^{VMF})$	1.50	0.41	1.86	3.24	1.87	1.00	2.65	4.07
$\mathcal{O}(\alpha^{BMF})$	0.05	-0.82	0.60	1.78	-0.40	-1.40	0.60	1.97
$\mathcal{O} \alpha^{VMF} $	2.58	1.22	2.26	3.66	3.41	2.08	3.25	4.43
$\mathcal{O} \alpha^{BMF} $	2.05	0.68	1.41	2.59	2.58	0.85	1.71	3.33
$\mathcal{O} \alpha^{VMF} - \mathcal{O} \alpha^{BMF} $	0.53	0.53*	0.84*	1.07*	0.83**	1.23**	1.54**	1.09**
$p(\mathcal{O} \alpha^{VMF} < \mathcal{O} \alpha^{BMF})$	0.124	0.099	0.086	0.092	0.046	0.027	0.022	0.035
Panel B: Price wedges								
$\mathcal{O}(-PW^{*,VMF})$	29.90	16.52	30.90	43.54	42.57	37.39	42.97	50.95
$\mathcal{O}(-PW^{*,BMF})$	14.14	5.72	14.82	23.04	20.34	15.85	22.26	26.36
$\mathcal{O} PW^{*,VMF} $	30.29	16.59	30.90	43.54	42.69	37.39	42.97	50.95
$\mathcal{O} PW^{*,BMF} $	15.23	6.34	15.02	23.11	20.79	15.95	22.26	26.36
$\mathcal{O} PW^{*,VMF} - \mathcal{O} PW^{*,BMF} $	15.06	10.25	15.87	20.43	21.90*	21.45*	20.71*	24.59*
$p(\mathcal{O} PW^{*,VMF} < \mathcal{O} PW^{*,BMF})$	0.148	0.226	0.152	0.117	0.081	0.074	0.076	0.086

Note: This table reports pricing results under the VMF- or BMF-based CAPM using our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 as test assets. Panel A reports results for CAPM alphas and their absolute values. Panel B reports results for (negative) CAPM price wedges and their absolute values (van Binsbergen et al., 2023). We report the average, 25th quantile, median, and 75th quantile for the distributions of alphas and price wedges. The “Operator (\mathcal{O})” notation denotes which of these operators is being applied in each column. The argument in $\mathcal{O}(\cdot)$ denotes to what variable the operator is being applied (i.e., the “Avg.” column associated with the “ $p(\mathcal{O}|\alpha^{VMF}| < \mathcal{O}|\alpha^{BMF}|)$ ” reports the p-value associated with the test that the average absolute VMF alpha is less than the average absolute BMF alpha among our set of 382 portfolios). The leftmost columns report statistics for all portfolios. The rightmost columns report statistics only for portfolios whose beta gaps are above the median beta gap. P-values for the alpha statistics are based on our bootstrapping methodology described in Online Appendix Section OA.10. P-values for the price wedge statistics are based on the bootstrapping methodology from van Binsbergen et al. (2023). In both cases, we compute the p-values as the fraction of bootstrapped simulations that satisfy the test criterion in the respective rows. Data are from 1963–2021. Additional details can be found in Section 7.1.

Table 4: VMF versus BMF alpha summary statistics by portfolio size

Operator (\mathcal{O}):	Small				Mid.				Large			
	Avg.	q25	Med.	q75	Avg.	q25	Med.	q75	Avg.	q25	Med.	q75
$\mathcal{O}(\alpha^{VMF})$	2.35	1.66	3.24	4.34	1.39	0.43	2.03	3.07	0.77	0.01	0.89	1.81
$\mathcal{O}(\alpha^{BMF})$	-0.12	-0.75	0.80	2.11	-0.02	-1.03	0.67	1.79	0.29	-0.67	0.54	1.55
$\mathcal{O} \alpha^{VMF} $	3.61	2.23	3.35	4.65	2.62	1.65	2.43	3.41	1.52	0.62	1.29	2.20
$\mathcal{O} \alpha^{BMF} $	2.48	0.80	1.63	3.06	2.19	0.83	1.59	2.89	1.50	0.57	1.25	2.06
$\mathcal{O} \alpha^{VMF} - \mathcal{O} \alpha^{BMF} $	1.14**	1.43**	1.72**	1.59**	0.43	0.82*	0.85*	0.53	0.02	0.05	0.04	0.14
$p(\mathcal{O} \alpha^{VMF} < \mathcal{O} \alpha^{BMF})$	0.030	0.021	0.013	0.031	0.130	0.097	0.086	0.114	0.675	0.534	0.575	0.618

Note: This table reports VMF- or BMF-based CAPM alpha statistics by portfolio size using our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 as test assets. We report the average, 25th quantile, median, and 75th quantile for the alphas and their absolute values. The “Operator (\mathcal{O})” notation denotes which of these operators is being applied in each column. The argument in $\mathcal{O}(\cdot)$ denotes to what variable the operator is being applied (i.e., the “Avg.” column associated with the “ $p(\mathcal{O}|\alpha^{VMF}| < \mathcal{O}|\alpha^{BMF}|)$ ” reports the p-value associated with the test that the average absolute VMF alpha is less than the average absolute BMF alpha among our set of 382 portfolios). The leftmost, middle, and rightmost columns report statistics for small-cap, mid-cap, and large-cap portfolios, respectively. Size is measured as the average monthly log market capitalization over time and portfolios are split by terciles. P-values are based on our bootstrapping methodology described in Online Appendix Section OA.10 and are computed as the fraction of bootstrapped simulations that satisfy the test criterion in the last row. Data are from 1963–2021. Additional details can be found in Section 7.2.

Table 5: Size factor spanning tests

	Factor model					
	FF3	FF5	FF5C	q4	SY	q5
Panel A: VMF as the market factor						
α^{VMF}	1.44	3.26**	3.10**	4.56**	6.47***	7.39***
	[1.00]	[2.16]	[2.21]	[2.55]	[4.44]	[3.98]
β^{VMF}	0.19***	0.15***	0.15***	0.14***	0.09**	0.09**
	[6.39]	[4.01]	[4.37]	[3.50]	[2.22]	[2.18]
$adjR^2$	0.10	0.18	0.18	0.15	0.10	0.19
Panel B: BMF as the market factor						
α^{BMF}	-0.40	0.23	-0.70	0.84	1.16	1.80
	[-0.34]	[0.17]	[-0.60]	[0.50]	[1.01]	[1.02]
β^{BMF}	0.43***	0.42***	0.44***	0.40***	0.44***	0.38***
	[14.37]	[11.96]	[13.82]	[10.51]	[12.35]	[9.39]
$adjR^2$	0.35	0.38	0.40	0.34	0.32	0.34

Note: This table reports size factor spanning test results using the six standard factor models described in Section 6.2 that include size factors (i.e., results from regressing size factors on other model factors). Panel A reports results when the VMF is used as the market factor in each model. Panel B reports results when the BMF is used. α^{VMF} (α^{BMF}) is the alpha when the VMF (BMF) is used (annualized and in percent). β^{VMF} (β^{BMF}) is the size factor beta with respect to the VMF (BMF). T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. Data are from 1963–2021 (conditional on factor availability). See Section 8.2 for additional details.

Table 6: Mean absolute alphas and size factors

	FF3	FF5	FF5C	q4	SY	q5
Panel A: VMF as the market factor						
$ \alpha_{wSize}^{VMF} $	1.74	1.30	1.19	1.27	1.16	1.48
$ \alpha_{exSize}^{VMF} $	1.92	2.17	2.21	2.75	3.37	4.40
$ \alpha_{exSize}^{VMF} - \alpha_{wSize}^{VMF} $	0.18	0.87***	1.02***	1.49***	2.22***	2.92***
$p\left(\alpha_{exSize}^{VMF} - \alpha_{wSize}^{VMF} \leq 0\right)$	0.101	0.004	0.007	0.002	0.002	0.000
Panel B: BMF as the market factor						
$ \alpha_{wSize}^{BMF} $	1.76	1.35	1.40	1.56	1.35	1.37
$ \alpha_{exSize}^{BMF} $	1.81	1.36	1.33	1.69	1.35	1.73
$ \alpha_{exSize}^{BMF} - \alpha_{wSize}^{BMF} $	0.04	0.01	-0.08	0.13	0.01	0.36
$p\left(\alpha_{exSize}^{BMF} - \alpha_{wSize}^{BMF} \leq 0\right)$	0.255	0.262	0.585	0.299	0.423	0.131
Panel C: Comparing VMF- and BMF-based models						
$ \alpha_{wSize}^{VMF} - \alpha_{wSize}^{BMF} $	-0.02	-0.05	-0.21	-0.29	-0.19	0.11
$p\left(\alpha_{wSize}^{VMF} - \alpha_{wSize}^{BMF} \leq 0\right)$	0.70	0.96	0.83	0.88	0.97	0.46
$ \alpha_{exSize}^{VMF} - \alpha_{exSize}^{BMF} $	0.12	0.80*	0.88	1.07**	2.02***	2.67***
$p\left(\alpha_{exSize}^{VMF} - \alpha_{exSize}^{BMF} \leq 0\right)$	0.444	0.074	0.101	0.037	0.009	0.000

Note: This table reports mean absolute alphas (and their differences) using our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 as test assets. Alphas are computed with respect to the six standard factor models described in Section 6.2 that include size factors. The mean absolute alphas reflect those plotted in Figure 11. Panel A reports results when the VMF is used as the market factor in each model. Panel B reports results when the BMF is used. Panel C reports results that compare VMF- to BMF-based model results. Data are from 1963–2021 (conditional on factor availability). See Section 8.2 for additional details.

Table 7: Risk-return relation coefficient (γ) across GARCH models

Risk defn.	Parameters	VMF	BMF (ind-sales)	BMF (sales-only)	BMF EW PCA 1	BMF cap VW PCA 1	BMF IPCA 1
Panel A: GARCH							
Variance (σ_t^2)	γ	2.47 [1.23]	6.58*** [2.72]	7.96*** [2.86]	8.01*** [3.05]	7.32*** [2.71]	8.56** [2.53]
	$\gamma^{VMF} - \gamma^{BMF}$		-4.11	-5.48	-5.54	-4.85	-5.71
Volatility (σ_t)	γ	0.19 [1.03]	0.54*** [2.65]	0.66*** [2.85]	0.69*** [3.04]	0.65*** [2.69]	0.73*** [2.76]
	$\gamma^{VMF} - \gamma^{BMF}$		-0.35	-0.47	-0.50	-0.46	-0.52
Panel B: GJR-GARCH							
Variance (σ_t^2)	γ	1.01 [0.54]	4.05** [2.29]	4.70** [2.54]	4.22** [2.52]	3.28** [1.97]	9.18*** [3.26]
	$\gamma^{VMF} - \gamma^{BMF}$		-3.04	-3.69	-3.21	-2.27	-7.41
Volatility (σ_t)	γ	0.15 [0.79]	0.41** [2.53]	0.48*** [2.87]	0.45*** [2.81]	0.41** [2.41]	0.80*** [3.54]
	$\gamma^{VMF} - \gamma^{BMF}$		-0.26	-0.33	-0.30	-0.26	-0.57
Panel C: E-GARCH							
Variance (σ_t^2)	γ	1.61 [0.86]	5.17** [2.56]	5.92*** [2.72]	6.19*** [2.95]	6.04*** [2.79]	9.78*** [3.19]
	$\gamma^{VMF} - \gamma^{BMF}$		-3.57	-4.31	-4.59	-4.43	-7.60
Volatility (σ_t)	γ	0.18 [1.01]	0.36** [2.31]	0.41*** [2.69]	0.46*** [3.09]	0.49*** [2.69]	0.72*** [3.11]
	$\gamma^{VMF} - \gamma^{BMF}$		-0.19	-0.24	-0.29	-0.31	-0.51

Note: This table compares γ estimates based on the VMF with those based on five different BMF variants: 1) “BMF (ind-sales)” is our main BMF measure constructed using the industry-sales-based double-sorting procedure described in Section 5.2, 2) “BMF (sales-only)” is an alternative BMF specification constructed using the sales-based single-sorting procedure described in Section 5.3 (note that the first two BMF specifications are based on our public-to-all-firm reweighting methodology described in Section 5.1, but using different portfolio sorts), 3) “BMF (EW PCA 1)” is an alternative that projects our main BMF onto the first PC from a PCA using our set of 382 equal-weighted portfolios described in Online Appendix Table OA.3, 4) “BMF (cap VW PCA 1)” is an alternative that projects our main BMF onto the first PC from a PCA using our set of 382 capped value-weighted portfolios described in Online Appendix Table OA.3, and 5) “BMF (IPCA 1)” is an alternative that projects our main BMF onto the first instrumented principal component from Kelly et al. (2019). Additional details on our alternative BMFs can be found in Section 5.3. We consider three popular GARCH volatility forecasting models including a GARCH (1,1) (Panel A), a threshold GARCH model of Glosten et al. (1993) (Panel B) and an exponential GARCH model of Nelson (1991) (Panel C). Within each panel, the first (second) set of results is based on the conditional variance (volatility) as a proxy for risk. Rows labeled $\gamma^{VMF} - \gamma^{BMF}$ report the difference between γ s estimated using the VMF and the respective BMF variants. We report t-statistics based on outer product standard errors (Hamilton, Time Series Analysis, 1994, pp. 133–148) in brackets. */**/** represent statistical significance at 90%, 95%, and 99%, respectively. Data are from 1963–2021. See Section 9 for additional details.

Online Appendix

“A Tale of Two Market Returns: The Broad Market Factor and The Idiosyncratic Financial Factor”

Sung Je Byun, Johnathan A. Loudis, and Lawrence D.W. Schmidt

OA.1 Summary of additional tables & figures

In this section, we summarize the figures and tables in the Online Appendix related to the additional results and robustness checks referenced in the main text. Please see references to these figures and tables in the main text for additional details.

OA.1.1 Figures

- Figure OA.1 plots the sales shares in the aggregate market by industry for public firms and for all firms based on our methodology described in Section 5 (based on our main industry-sales-based portfolio sorting approach). See Section 3.2 and Online Appendix OA.2 for additional details.
- Figure OA.2 plots information about the distribution of firm-level weights in our main BMF aggregation using the industry-sales-based portfolio sorting approach described in Section 5. See Section 5.2 and footnote 16 for additional details.
- Figures OA.3, OA.4, OA.5, and OA.6 provide information about sales and value weights of portfolios and firms implied by our methodology in Section 5 in our alternative BMF based on the sales-only portfolio sorting approach to constructing the BMF described in Section 5.3 (i.e., as opposed to our main industry-sales-based portfolio sorting approach described in Section 5.2). These figures are analogous to those based on the industry-sales-based portfolio sorting approach from Figures 2, OA.1, OA.2, and 4 , respectively.
- Figure OA.7 plots BMF and IFF cumulative return time series. See Section 6.1 (footnote 22) for more details.
- Figure OA.8 compares local projection impulse responses (LPIRs) with respect to the BMF against those with respect to the VMF. See Section 6.3 (footnote 25) and Online

Appendix OA.7.3 for additional details.

- Figure OA.9 plots robustness checks on the VMF versus BMF betas plotted in Figure 7 Panel B using alternative BMF specifications described in Section 5.3. These are meant to test the model implication highlighted in Figure 3 Panel B. See Section 7.2 for additional details.
- Figure OA.10 plots robustness checks on the VMF versus BMF alphas plotted in Figure 6 Panel C using alternative BMF specifications described in Section 5.3. These are meant to test the model implication highlighted in Figure 3 Panel A. See Section 7.1 for additional details.
- Figure OA.11 plots robustness checks on alphas for size-sorted portfolios plotted in Figure 6 Panel B using alternative BMF specifications described in Section 5.3. These are meant to test the model implication highlighted in Figure 3 Panel C. See Section 8.1 for additional details.
- Figure OA.12 plots alphas for value-weighted size-sorted decile portfolios (a subset of our 382 characteristics-sorted portfolios described in Online Appendix Table OA.3) for each of the nine factor models described in Section 6.2 when either the VMF or the BMF is used as the market factor in each model for the 1993–2021 subsample. See Section 8.1 for additional details.
- Figure OA.13 plots robustness checks on model-implied return alignment plotted in Figure 12 using alternative BMF specifications described in Section 5.3. These are meant to test the model implication highlighted in Figure 3 Panel D. See Section 8.3 for additional details..
- Figure OA.14 plots conditional VMF and BMF volatility estimates from the baseline GARCH specification in equations (17)–(19). See Section 9.2 for additional details.
- Figure OA.15 plots size-by-book-to-market-sorted portfolio beta gaps. See Online Appendix OA.8 for additional details.

- Figure OA.16 plots results from factor spanning tests run for each factor and each model we study. See Online Appendix OA.8 for additional details.

OA.1.2 Tables

- Table OA.1 summarizes our model parameters and their calibrated values. See Section 4 and Online Appendix OA.3 for additional details.
- Table OA.2 summarizes our model-implied moments based on calibrated parameter values reported in Table OA.1. See Section 4 and Online Appendix OA.3 for additional details.
- Table OA.3 summarizes the set of 382 portfolios we use in our PCA analysis and as test assets. These portfolios were obtained from Kenneth French’s website. In some cases (as in our PCA exercise described in Section 5.3), we use equal-weighted versions of the portfolios. In others (as in our asset pricing tests), we use value-weighted versions of the portfolios. In yet others (as in our alternative capped value-weighted PCA exercise described in Section 5.3), we construct our own versions of these portfolios since they are not publicly available.
- Table OA.4 reports results from regressions of our main BMF on various PCs and IPCs described in Section 5.3. See Section 5.3 for additional details.
- Table OA.5 reports robustness checks on the IFF spanning tests reported in Table 2 using alternative BMF specifications described in Section 5.3. See Section 6.2 for additional details.
- Table OA.6 reports size factor spanning test results for the 1993–2021 subsample. See Section 8.2 for additional details.
- Table OA.7 reports HML factor spanning test results. See Online Appendix OA.8 for additional details.

OA.2 Motivating evidence: Industry-based sales shares

Figure OA.1 plots the sales shares across the seven KMZ industries over time among public firms (Panel A), all firms (Panel B), the top 2.5% of public firms (Panel C), and the top 2.5% of public and private firms (the last two are both according to NYSE market equity breakpoints). These shares are computed using on our main industry-by-sales public-to-all-firm matching procedure described in Section 5, and in more detail in Online Appendix OA.5.

Comparing Panels A and B reveals striking differences in industry composition between public firms and all firms, some of which are dynamic and some of which are persistent. For instance, there is a steady decline in the sales share of the manufacturing industry during this sample among both sets of firms, but it is less drastic among all firms. There is also a notable increase in the sales share of the services industry (which includes the technology sector) among public firms around 2015, which is not present in all firms. However, the sales share of the services industry was consistently higher among all firms until around 2015. Some industries demonstrate consistently higher sales shares in one universe of firms versus the other. For instance, the trade industry has a consistently larger sales share among all firms compared to public firms. The key takeaway is that there are some stark differences in terms of the industry composition of public firms versus all firms in the US.

Panels C and D show that these compositional differences can be even more dramatic among the largest firms in the economy, which, owing to their size, can have outsized influence on value-weighted indices such as the VMF. For instance, consistent with recent concerns in the popular news media, there has been rapid growth in the sales share of services firms among public firms (Panel C, yellow region) since 2015, which is much more muted among all firms represented in Panel D (note that the service industry contains the high-tech sector).

OA.3 Theoretical model: Additional details

In this section, we provide additional details related to our theoretical model and its calibration. For reference, we summarize all 15 of our model parameters in Table OA.1. Note that we do not consider the number of small and big firms (N_S and N_B , respectively) as parameters since they are not needed for our calibration or simulations (we only need know

the share of small firms, θ).

OA.3.1 Additional model-implied return and beta expressions

Here we provide additional details related to model-implied returns and betas. The risk premium-beta relationship in equation (6) can also be expressed in covariance form as

$$\mathbb{E}[R_i] - R_f = (\alpha W_0) \text{COV}[R_i, R_w] + \omega \text{COV}[R_i, R_z], \quad (\text{OA.1})$$

Applying this pricing equation to the wealth portfolio implies

$$W_0 = R_f^{-1} [1 - \alpha \bar{\gamma}^2 \sigma_f^2] \quad \mathbb{E}[R_w] - R_f = \frac{R_f \alpha \bar{\gamma}^2 \sigma_f^2}{1 - \alpha \bar{\gamma}^2 \sigma_f^2}. \quad (\text{OA.2})$$

The risk-free rate is determined by equating the marginal utility of current consumption with the risk-free rate times the expected marginal utility of future wealth as

$$R_f = \frac{1}{\rho} \exp \left[\alpha (1 - D_0) - \frac{1}{2} (\alpha^2 \bar{\gamma}^2 \sigma_f^2 + \omega^2 \sigma_z^2) \right]. \quad (\text{OA.3})$$

Thus, R_f is determined by the usual time discounting, expected consumption growth, and precautionary savings terms.

Note that equation (OA.1) implies that the SDF is linear in the factors according to

$$M = \frac{1}{R_f} [1 - \alpha \bar{\gamma} f - \omega z]. \quad (\text{OA.4})$$

Using the SDF representation and dividend expressions from equation (2), it is straightforward to show that prices are given by

$$P_{i(S)} = \frac{1}{N} \frac{1}{R_f} [1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2 - \omega \lambda_S X_i \sigma_z^2] \quad P_{i(B)} = \frac{1}{N} \frac{1}{R_f} [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2 - \omega X_i \sigma_z^2] \quad (\text{OA.5})$$

Thus, the realized return on each stock is given by its realized dividend at date 1 (from

equation (2)) divided by its price at date 0, which is given by

$$R_{i(\tau)} = \frac{R_f (1 + \lambda_\tau \bar{d} + \gamma_\tau f + \lambda_\tau g + \lambda_\tau X_i z)}{1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2 - \omega \lambda_\tau X_i \sigma_z^2}. \quad (\text{OA.6})$$

We can also aggregate the dividends of small and big firms from equation (2) and price them using the SDF in equation (OA.4) to show that the return on the portfolios of all small and all big firms can be expressed as

$$R_S = \frac{R_f (1 + \lambda_S \bar{d} + \gamma_S f + \lambda_S g)}{[1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2]} \quad R_B = \frac{R_f (1 + \bar{d} + \gamma_B f + g)}{[1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]}, \quad (\text{OA.7})$$

where there is no dependence on z since X has a mean of zero and is orthogonal to firm size, τ . Note that these return expressions actually apply to portfolio of all (small or big) firms, public-only (small or big) firms, and private (small or big) firms since within a size type there is no differential probability of going public.

We can also express the returns on individual stocks from equation (OA.6) in terms of their betas with respect to the BMF and IFF as

$$R_{i(\tau)} = R_f + \beta_{i(\tau)}^{BMF} \cdot BMF + \beta_{i(\tau)}^{IFF} \cdot IFF + \beta_{i(\tau)}^z (R_z - R_f) \quad (\text{OA.8})$$

where

$$\beta_{i(\tau)}^{BMF} = \frac{\beta_{i(\tau)}^w}{\beta_{VMF}^w} = \frac{1}{\beta_{VMF}^w} \frac{\gamma_\tau}{\bar{\gamma}} \frac{(1 - \alpha \bar{\gamma}^2 \sigma_f^2)}{(1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2 - \omega \lambda_\tau X_i \sigma_z^2)} \quad (\text{OA.9})$$

$$\beta_{i(\tau)}^{IFF} = \frac{\beta_{i(\tau)}^g}{\beta_{VMF}^g} = \frac{1}{\beta_{VMF}^g} \frac{R_f \lambda_\tau}{1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2 - \omega \lambda_\tau X_i \sigma_z^2}, \text{ and} \quad (\text{OA.10})$$

$$\beta_{i(\tau)}^z = \frac{R_f \lambda_\tau X_i}{1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2 - \omega \lambda_\tau X_i \sigma_z^2}. \quad (\text{OA.11})$$

We can aggregate these returns to form the aggregate small and big portfolio returns in

equation (OA.7) and alternatively express them in beta form as

$$R_\tau = R_f + \beta_\tau^{BMF} \cdot BMF + \beta_\tau^{IFF} \cdot IFF \quad (\text{OA.12})$$

where $\tau \in \{S, B\}$ and where

$$\beta_\tau^{BMF} = \frac{\beta_\tau^w}{\beta_{VMF}^w} = \frac{1}{\beta_{VMF}^w} \frac{\gamma_\tau}{\bar{\gamma}} \frac{1 - \alpha \bar{\gamma}^2 \sigma_f^2}{1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2}, \text{ and} \quad (\text{OA.13})$$

$$\beta_\tau^{IFF} = \frac{\beta_\tau^g}{\beta_{VMF}^g} = \frac{1}{\beta_{VMF}^g} \frac{R_f \lambda_\tau}{1 + \lambda_\tau \bar{d} - \alpha \gamma_\tau \bar{\gamma} \sigma_f^2}. \quad (\text{OA.14})$$

Note that the aggregated small and big portfolios do not load on z because it is diversified in those portfolios, but they do load on g (i.e., the IFF).

We can construct the VMF by aggregating small and big portfolio returns accounting for the probabilities of small and big firms going public, which yields equation (8) repeated here for convenience

$$VMF^{model} = \underbrace{\beta_{VMF}^w (R_w - R_f)}_{\equiv BMF^{model}} + \underbrace{\beta_{VMF}^g g}_{\equiv IFF^{model}}, \quad (\text{OA.15})$$

where

$$\beta_{VMF}^w = \frac{(1 - \alpha \bar{\gamma}^2 \sigma_f^2) \left[1 + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \frac{\gamma_B}{\bar{\gamma}} \right]}{1 - \alpha \bar{\gamma}^2 \sigma_f^2 + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]} \text{ and} \quad (\text{OA.16})$$

$$\beta_{VMF}^g = \frac{R_f \left[\left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \right]}{\{1 - \alpha \bar{\gamma}^2 \sigma_f^2\} + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]}. \quad (\text{OA.17})$$

OA.3.2 Model calibration details

Our model has 15 parameters that need to be calibrated (see Table OA.1 for a summary of the parameters). We set $\rho = 0.999$ (note this is a monthly discount factor as we calibrate our model using monthly data). We calibrate three of the parameters (θ , ρ_S^{pub} , and ρ_B^{pub}) using direct analogs in the data. We pin down the remaining parameters using an overidentified system of moments.

For the purposes of this calibration, we match some model moments associated with small and big firm portfolios. We use the long and short legs of the Fama and French (1993) SMB factor (obtained from Ken French’s website) as proxies for these portfolios in the data, respectively. To recognize the fact that the constituents of these portfolios in the data may not adhere to the strict definitions of S- and B-type firms in the model (for instance, in the observed publicly-traded small-cap portfolio, it may contain some firms the model would consider as B-type, either based strictly on size or some other omitted characteristic such as value), we add two additional calibration parameters that modulate the fraction of S-type firms in the long leg of the model’s SMB factor and the fraction of B-type firms in the short leg of the model’s SMB factor. These are the parameters listed as w_S and w_B in Table OA.1, respectively. These parameters help us match model-implied small and big portfolio moments with those in the data (particularly, the loading of the empirical SMB factor on our empirical BMF and IFF). Given the returns on the S- and B-type portfolios defined in equation (OA.7), the returns on these mixed portfolios can be expressed as:

$$\tilde{R}_S = w_S R_S + (1 - w_S) R_B \quad \tilde{R}_B = w_B R_B + (1 - w_B) R_S, \quad (\text{OA.18})$$

We use model-implied moments from these “mixed” portfolios to match data-implied moments from the SMB small- and large-cap portfolios as well as the loadings of SMB on our BMF and IFF (see below for additional details on these model-implied moments).

We use the relative fraction of public versus all firm sales (based on KMZ data) among firms in different size categories to estimate the probability of firms in each category going public (see Section 5 for a description of how we estimate sales aggregates for the set of public versus all firms). Thus, to calibrate the model, we use our data to estimate the sales

of firms in the set of all (public) small firms defined as by the terciles discussed above and call this S_S^{all} (S_S^{pub}). We do the same for a set of big firms and call this S_B^{all} (S_B^{pub}). As in our BMF construction (see Section 5.1), we assume that the probability of going public is related to the ratio of public to private sales for each firm type as

$$\rho_S^{pub} = \frac{S_S^{pub}}{S_S^{all}}, \text{ and} \quad (\text{OA.19})$$

$$\rho_B^{pub} = \frac{S_B^{pub}}{S_B^{all}}, \quad (\text{OA.20})$$

which gives us our calibrated ρ_S^{pub} and ρ_B^{pub} values. Next, we observe the number of public small and big firms and we call these N_S^{pub} and N_B^{pub} . We can then apply inverse probability weights to get the implied total fraction of small firms among all firms as

$$\theta = \frac{N_S^{pub} * \frac{S_S^{all}}{S_S^{pub}}}{N_S^{pub} * \frac{S_S^{all}}{S_S^{pub}} + N_B^{pub} * \frac{S_B^{all}}{S_B^{pub}}}, \quad (\text{OA.21})$$

which give us our calibrated θ value.

We calibrate the remaining eleven parameters as follows. First, note that σ_z is not identified independently of ω , so we simply set its value to the calibrated value of σ_f . This leaves ten remaining parameters. We pin down the next eight (α , σ_f , σ_g , γ_S , γ_B , \bar{d} , w_S , and w_B) using an overidentified system of 18 moment conditions, which are functions of the parameters. Given the returns and beta expressions from the previous section along with expressions from Section 4, it is straightforward to derive the following 18 moment conditions.

$$\mathbb{E}[R_w] = \frac{R_f}{1 - \alpha \bar{\gamma}^2 \sigma_f^2}$$

$$\text{VAR} [R_w]^{1/2} = \left[\frac{R_f^2 \bar{\gamma}^2 \sigma_f^2}{(1 - \alpha \bar{\gamma}^2 \sigma_f^2)^2} \right]^{1/2}$$

$$\mathbb{E} [VMF] = \frac{R_f \left[1 + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) (1 + \bar{d}) \right]}{\{1 - \alpha \bar{\gamma}^2 \sigma_f^2\} + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]} - R_f$$

$$\begin{aligned} \text{VAR} [VMF]^{1/2} = & \left[\left\{ \frac{R_f \left[\bar{\gamma} + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \gamma_B \right]}{\{1 - \alpha \bar{\gamma}^2 \sigma_f^2\} + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]} \right\}^2 \sigma_f^2 \right. \\ & \left. + \left\{ \frac{R_f \left[\left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \right]}{\{1 - \alpha \bar{\gamma}^2 \sigma_f^2\} + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) [1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2]} \right\}^2 \sigma_g^2 \right]^{1/2} \end{aligned}$$

$$\beta_{VMF}^w = \frac{\left[1 + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \frac{\gamma_B}{\bar{\gamma}} \right]}{\left[1 + \left(\frac{\rho_B^{pub} - \rho_S^{pub}}{\rho_S^{pub}} \right) (1 - \theta) \left[\frac{1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2}{1 - \alpha \bar{\gamma}^2 \sigma_f^2} \right] \right]}$$

$$\text{VAR} [IFF]^{1/2} = \text{VAR} [\beta_{VMF}^g]^{1/2} = [\beta_{VMF}^{2g} \sigma_g^2]^{1/2}$$

$$\mathbb{E} [\tilde{R}_S] = w_S \mathbb{E} [R_S] + (1 - w_S) \mathbb{E} [R_B]$$

$$\begin{aligned}\text{VAR} \left[\tilde{R}_S \right]^{1/2} &= \left(w_S^2 \text{VAR} [R_S] + (1 - w_S)^2 \text{VAR} [R_B] \right. \\ &\quad \left. + 2w_S(1 - w_S) \text{VAR} [R_S]^{1/2} \text{VAR} [R_B]^{1/2} \rho_{SB} \right)^{1/2}\end{aligned}$$

$$\mathbb{E} \left[\tilde{R}_S \right] = w_S \mathbb{E} [R_S] + (1 - w_S) \mathbb{E} [R_B]$$

$$\begin{aligned}\text{VAR} \left[\tilde{R}_B \right]^{1/2} &= \left(w_B^2 \text{VAR} [R_B] + (1 - w_B)^2 \text{VAR} [R_S] \right. \\ &\quad \left. + 2w_B(1 - w_B) \text{VAR} [R_B]^{1/2} \text{VAR} [R_S]^{1/2} \rho_{SB} \right)^{1/2}\end{aligned}$$

$$\tilde{\beta}_S^w = w_S \beta_S^w + (1 - w_S) \beta_B^w$$

$$\tilde{\beta}_S^{IFF} = w_S \beta_S^{IFF} + (1 - w_S) \beta_B^{IFF}$$

$$\tilde{\beta}_B^w = w_B \beta_B^w + (1 - w_B) \beta_S^w$$

$$\tilde{\beta}_B^{IFF} = w_B \beta_B^{IFF} + (1 - w_B) \beta_S^{IFF}$$

$$\tilde{\beta}_S^{VMF} = w_S \beta_S^{VMF} + (1 - w_S) \beta_B^{VMF}$$

$$\tilde{\beta}_B^{VMF} = w_B \beta_B^{VMF} + (1 - w_B) \beta_S^{VMF}$$

$$\tilde{\beta}_{SMB}^{BMF} = \tilde{\beta}_S^{BMF} - \tilde{\beta}_B^{BMF}$$

$$\tilde{\beta}_{SMB}^{IFF} = \tilde{\beta}_S^{IFF} - \tilde{\beta}_B^{IFF}$$

where “ $\tilde{\beta}_S$ ” and “ $\tilde{\beta}_B$ ” reflect betas of the \tilde{R}_S and \tilde{R}_B portfolios with respect to various factors. Note that these $\tilde{\beta}$ s and other moments associated with \tilde{R}_S and \tilde{R}_B are expressed above in terms of betas and other moments associated with the R_S and R_B portfolios, which are given as functions of model parameters as:

$$\mathbb{E}[R_S] = \frac{R_f (1 + \lambda_S \bar{d})}{1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2}$$

$$\text{VAR}[R_S]^{1/2} = \left[\frac{R_f^2 (\gamma_S^2 \sigma_f^2 + \lambda_S^2 \sigma_g^2)}{(1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2)^2} \right]^{1/2}$$

$$\mathbb{E}[R_B] = \frac{R_f (1 + \bar{d})}{1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2}$$

$$\text{VAR} [R_B]^{1/2} = \left[\frac{R_f^2 (\gamma_B^2 \sigma_f^2 + \sigma_g^2)}{(1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2)^2} \right]^{1/2}$$

$$\beta_S^w = \frac{\gamma_S}{\bar{\gamma}} \frac{1 - \alpha \bar{\gamma}^2 \sigma_f^2}{1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2}$$

$$\beta_S^{IFF} = \frac{R_f \lambda_S}{\beta_{VMF,g} (1 + \lambda_S \bar{d} - \alpha \gamma_S \bar{\gamma} \sigma_f^2)}$$

$$\beta_B^w = \frac{\gamma_B}{\bar{\gamma}} \frac{1 - \alpha \bar{\gamma}^2 \sigma_f^2}{1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2}$$

$$\beta_B^{IFF} = \frac{R_f}{\beta_{VMF,g} (1 + \bar{d} - \alpha \gamma_B \bar{\gamma} \sigma_f^2)}$$

$$\beta_S^{VMF} = \frac{\beta_S^w \beta_{VMF}^w \text{VAR} [R_w] + \beta_S^g \beta_{VMF}^g \sigma_g^2}{(\beta_{VMF}^w)^2 \text{VAR} [R_w] + (\beta_{VMF}^g)^2 \sigma_g^2}$$

$$\beta_B^{VMF} = \frac{\beta_B^w \beta_{VMF}^w \text{VAR} [R_w] + \beta_B^g \beta_{VMF}^g \sigma_g^2}{(\beta_{VMF}^w)^2 \text{VAR} [R_w] + (\beta_{VMF}^g)^2 \sigma_g^2}$$

$$\rho_{SB} = \frac{\beta_S^w \beta_B^w \text{VAR} [R_w] + \beta_S^{IFF} \beta_B^{IFF} \text{VAR} [IFF]}{\text{VAR} [R_S]^{1/2} \text{VAR} [R_B]^{1/2}}$$

We use our empirical proxy for R_w , \hat{R}_{all} , described in Section 5 for the purposes of estimating data moments when implementing the above equations.

This leaves two more parameters to identify (C_0 and ω). Given the parameters identified above, we pin down C_0 by solving for the value that equates the expression for R_f in equation (OA.3) equal to the average R_f value in our sample, \hat{R}_f , as

$$\hat{R}_f = \frac{1}{\rho} \exp \left[\alpha (1 - C_0) - \frac{1}{2} (\alpha^2 \bar{\gamma}^2 \sigma_f^2 + \omega^2 \sigma_z^2) \right].$$

Thus, the model-implied risk-free rate will exactly match its average in our data. Finally, we pin down ω by setting the risk premium on z to be equal to the mean absolute alpha across our 382 portfolios measured with respect to our empirical BMF as

$$|\bar{\alpha}| = \omega \sigma_z^2.$$

A summary of all calibrated parameter values can be found in Table OA.1. Finally, we provide a summary of all the above moments measured in the data along with their exact model-implied values (given the parameters in Table OA.1) and values implied by our 10,000 simulations in Table OA.2.

OA.3.3 Simulating portfolio returns

We simulate returns on 382 portfolios constructed to match our empirical BMF betas by forming portfolios of small and big stock portfolios whose returns are described in equation (OA.7) where we select the small portfolio weight, w_S , such that the combined portfolio has a BMF beta for each of our 382 portfolios equal to our estimated BMF beta from our empirical exercise (using the model-implied BMF beta expressions in equation (OA.13)).

Furthermore, because these portfolios have average excess returns that cannot be fully explained by the BMF, we induce random loading, X_p , for each portfolio on R_z as follows.⁴⁰ First, we take the difference between model-implied excess returns under the Fama-French five-factor model and the CAPM (both using our empirical BMF as the market factor) to pin down a distribution of implied z risk premia across these portfolios (i.e., z reflects factors outside the CAPM that are needed to price assets). This distribution has an average value

⁴⁰Recall that z is meant to be a shifter that captures any source of risk premia in the data that is not explained by R_w (equivalently, not explained by the BMF).

close to zero, consistent with our model’s assumptions about the average z risk premium. To be consistent with our model’s null that z loadings (X) are independent across asset types, we randomly select z risk premia from the empirical distribution (described above) for each of our simulated portfolios and set the portfolio’s z loading (X_p) to be the sampled portfolio’s z premium divided by the total z premium in the model ($\omega\sigma_z^2$). Note that the alternative model (here the Fama-French five-factor model) does not affect the qualitative model-implied results. This is because the average difference between BMF- and other model-implied expected returns is close to zero (on average) and we just use z exposure in the model to generate a wedge between BMF-implied and true expected returns similar to that which we observe in the data.

Thus, the DGP for our simulated portfolio returns can be expressed as

$$R_p - R_f = w_S R_S + (1 - w_S) R_B + X_p (R_z - R_f). \quad (\text{OA.22})$$

OA.4 Model extension: Heterogeneous agents with non-traded human capital

This section extends our baseline model from Section 4 to incorporate an additional source of resources available to finance consumption: human capital, represented by heterogeneous agents (S-workers and B-workers) with non-traded labor income subject to both systematic and idiosyncratic risk. We refer to this component of wealth as human capital for concreteness and because it aligns with some of the evidence we present in our local projection impulse response analysis in Subsection 6.3. However, as we explain in Remark 3, it can be interpreted more generally as any non-traded asset (with a similar payoff structure), such as ownership stakes in private firms that are non-tradable.

The key question we address in the extended model is whether the redistributive shock g enters the SDF once we account for nontraded human capital. We show below that the answer depends on the loading of aggregate *risk-adjusted* labor income on g . If this loading is zero, then the extended economy preserves the baseline implication that g is unpriced and, since the IFF is a scaled version of g , the IFF has no risk premium. If this loading is

nonzero, then g can be priced. We present the general model first, state restrictions under which the loading of idiosyncratic risk-adjusted labor income on g is zero, and then explore plausible departures from these restrictions. These departures point toward a negative IFF premium, but our empirical work suggests that any such premium is economically small and statistically indistinguishable from zero.

OA.4.1 Aggregate risk factors

We retain the aggregate risk structure from the baseline model in its entirety. There are N_S small firms and N_B big firms with $\theta = N_S/N$ ($N = N_S + N_B$), three orthogonal shocks

$$f \sim \mathcal{N}(0, \sigma_f^2), \quad g \sim \mathcal{N}(0, \sigma_g^2), \quad z \sim \mathcal{N}(0, \sigma_z^2),$$

and dividends

$$d_i = \begin{cases} \frac{1}{N} [1 + \lambda_S \bar{d} + \gamma_S f + \lambda_S g + \lambda_S X_i z] & \text{if } \tau_i = S, \\ \frac{1}{N} [1 + \lambda_B \bar{d} + \gamma_B f + \lambda_B g + \lambda_B X_i z] & \text{if } \tau_i = B, \end{cases}$$

where $\lambda_S = -(1 - \theta)/\theta$, $\lambda_B = 1$, $X_i \in \{-1, 1\}$ with $X_i \perp \tau_i$, and $\bar{d} > 0$. The aggregate dividend is $D_1 = 1 + \bar{\gamma}f$ with $\bar{\gamma} \equiv \theta\gamma_S + (1 - \theta)\gamma_B$, and the aggregate S and B dividends are

$$D_S \equiv \sum_{i:\tau_i=S} d_i = \theta(1 + \lambda_S \bar{d} + \gamma_S f + \lambda_S g), \quad (\text{OA.23})$$

$$D_B \equiv \sum_{i:\tau_i=B} d_i = (1 - \theta)(1 + \lambda_B \bar{d} + \gamma_B f + \lambda_B g), \quad (\text{OA.24})$$

where the $X_i z$ terms cancel within each type by the assumption that X_i is independent of type τ_i .

OA.4.2 Agents and endowments

We introduce **two types of agents**. A fraction $\mu \in (0, 1)$ are S -workers and a fraction $1 - \mu$ are B -workers. All agents have the same CARA preferences as in the baseline:

$$U_0^j = -\exp(-\alpha C_0^j) - \rho \mathbb{E}_0[\exp(-\alpha C_1^j - \omega z)], \quad (\text{OA.25})$$

where $\alpha > 0$ is absolute risk aversion, ω is the price of z -risk, and ρ is the discount factor.

Each agent receives two endowments:

1. **Financial endowment (tradable).** Every agent holds one share of every firm, regardless of type.
2. **Human capital endowment (nontradable).** Agent j of type $\tau(j) \in \{S, B\}$ receives labor income (for compactness of notation we suppress the dependence of τ on j)

$$Y_\tau^j = \underbrace{\bar{Y}_0 + \kappa_\tau(\gamma_\tau f + \lambda_\tau g)}_{\equiv Y_\tau \text{ (type-}\tau \text{ aggregate labor income)}} + \varepsilon_j, \quad (\text{OA.26})$$

where $\bar{Y}_0 \geq 0$ is common and nonstochastic, $\kappa_\tau > 0$ governs the type-specific scale of systematic labor income risk, and ε_j is an idiosyncratic shock with $\mathbb{E}[\varepsilon_j | f, g] = 0$.

The risky systematic component of each worker's labor income mirrors the per-firm dividend factor structure of the corresponding firm type, scaled by κ_τ . We assume the z -factor does not appear in Y_τ^j for simplicity, and including it would not affect our intuition developed herein regarding g pricing. Both worker types receive the same expected labor income $\mathbb{E}[Y_\tau^j] = \bar{Y}_0$ since $\mathbb{E}[f] = \mathbb{E}[g] = \mathbb{E}[\varepsilon_j] = 0$. While we refer to the idiosyncratic shock ε_j as an uninsurable labor income shock for simplicity, an alternative interpretation is that ε_j captures nondiversifiable exposures to idiosyncratic firm-level shocks which must be absorbed by business owners (e.g., due to moral hazard frictions).⁴¹

Remark 1 (Tradable spanning of f and g) *Given our maintained assumption that returns of S and B firms are linearly independent and span f and g and both types of firms are represented in public markets, agents can freely choose any linear exposure to aggregate risk by forming portfolios of publicly traded firms.*

Remark 2 (Separation of level and risk.) *The formulation in equation (OA.26) separates labor income into a nonstochastic level \bar{Y}_0 and a risk component $\kappa_\tau(\gamma_\tau f + \lambda_\tau g) + \varepsilon_j$.*

⁴¹Recall that, as discussed further below, extending the model to allow for idiosyncratic firm-level cash flow shocks is straightforward and has no implications for our asset pricing results provided that the number of firms is large. To keep the interpretation that ε_j is uninsurable, the key additional restriction to preserve implications from this section is that investors cannot short sell their own firms, a standard restriction in the literature.

With CARA preferences, the portfolio choice problem depends on the covariance structure of labor income—governed by κ_τ and the distribution of ε_j —but not on the level \bar{Y}_0 , which affects only the risk-free rate. This separation is what allows us to vary κ_S and κ_B independently while keeping the expected endowment common across types.

OA.4.3 Idiosyncratic risk structure

Similar to Constantinides and Duffie (1996), Herskovic et al. (2016), and Schmidt (2025), we allow the distribution of ε_j to depend on the aggregate state. We assume the conditional moment-generating function is exponentially affine with type-specific coefficients:

$$\mathbb{E}[e^{t\varepsilon_j} \mid f, g] = \exp(\gamma_{0,\tau}(t) + \gamma_{f,\tau}(t)f + \gamma_{g,\tau}(t)g), \quad \tau \in \{S, B\}, \quad (\text{OA.27})$$

where the coefficients may depend on $t \in \mathbb{R}$ and on the worker type. The zero-conditional-mean restriction requires $\gamma'_{0,\tau}(0) = \gamma'_{f,\tau}(0) = \gamma'_{g,\tau}(0) = 0$. Define

$$\Gamma_{0,\tau} \equiv -\frac{1}{\alpha}\gamma_{0,\tau}(-\alpha), \quad \Gamma_{f,\tau} \equiv -\frac{1}{\alpha}\gamma_{f,\tau}(-\alpha), \quad \Gamma_{g,\tau} \equiv -\frac{1}{\alpha}\gamma_{g,\tau}(-\alpha), \quad (\text{OA.28})$$

the certainty-equivalent coefficients implied by the conditional MGF at $t = -\alpha$ (the argument relevant for CARA expected utility). Equivalently, conditional on (f, g) , the idiosyncratic shock is equivalent to a change in expected consumption of $\Gamma_{0,\tau} + \Gamma_{f,\tau}f + \Gamma_{g,\tau}g$. Schmidt (2025) shows how linear combinations of normals and compound Poisson with state-dependent coefficients can fit into this class. This specification nests several important special cases. Setting $\gamma_{f,\tau}(\cdot) = \gamma_{g,\tau}(\cdot) = 0$ recovers state-independent idiosyncratic risk (the distribution of ε_j does not depend on the aggregate state). If in addition $\gamma_{0,\tau}(t) = 1/2\sigma^2t^2$, $\gamma_{f,\tau}(t) = 1/2bt^2$, and $\gamma_{g,\tau}(t) = 1/2ct^2$, then $\varepsilon_j \stackrel{iid}{\sim} N(0, \sigma^2 + bf + cg)$ for both types and, for example, $\Gamma_{g,\tau} = -(\alpha/2)c$.⁴² Setting all three coefficient functions to zero eliminates idiosyncratic risk entirely.

⁴²As is common in this literature, our way of specifying a moment generating function implicitly assumes that some technical restrictions on parameters are satisfied so that ε_j has a valid distribution in each state, such as nonnegative variances in Gaussian specifications ($\sigma^2 + bf + cg \geq 0$ in our example) or nonnegative arrival intensities in compound-Poisson/event-risk specifications. Researchers typically impose these restrictions formally through the choice of state space or handle remote edge cases by truncating or smoothly extrapolating the affine specification. We take the latter interpretation when such boundary cases arise.

OA.4.4 Terminal consumption and budget constraint

There is no production or storage technology; output is determined by the exogenous dividend and labor income processes. Agents trade risky assets and a risk-free bond (in zero net supply) at date 0 and consume all wealth at date 1. Agent j of type τ holds a financial portfolio with gross payoff W_1^j . Terminal consumption is

$$C_1^j = Y_\tau^j + W_1^j = Y_\tau + \varepsilon_j + W_0 R_p^j, \quad (\text{OA.29})$$

where W_0 is initial financial wealth (the equilibrium value of the stock market) and R_p^j is the gross return on the portfolio of financial wealth.

OA.4.5 Factoring out idiosyncratic risk

A key step in our derivations exploits the exponentially affine structure of the conditional MGF: given (f, g) , the idiosyncratic shock ε_j is independent of z and of the portfolio choice, so conditioning on (f, g) allows us to absorb ε_j into modified factor loadings. Applying the law of iterated expectations to agent j 's expected utility yields

$$\mathbb{E} \left[e^{-\alpha C_1^j - \omega z} \right] = \mathbb{E} \left[e^{-\alpha(Y_\tau + \varepsilon_j + W_0 R_p^j) - \omega z} \right] = e^{-\alpha \Gamma_{0,\tau}} \cdot \mathbb{E} \left[e^{-\alpha(\tilde{Y}_\tau + W_0 R_p^j) - \omega z} \right], \quad (\text{OA.30})$$

where *idiosyncratic-risk-adjusted labor income* \tilde{Y}_τ is defined as

$$\tilde{Y}_\tau = \bar{Y}_0 + \phi_{f,\tau} f + \phi_{g,\tau} g, \quad (\text{OA.31})$$

with the effective loadings coming from the conditional MGF equation (OA.27)

$$\phi_{f,\tau} \equiv \kappa_\tau \gamma_\tau + \Gamma_{f,\tau}, \quad \phi_{g,\tau} \equiv \kappa_\tau \lambda_\tau + \Gamma_{g,\tau}. \quad (\text{OA.32})$$

The terms $\Gamma_{f,\tau}$ and $\Gamma_{g,\tau}$ capture how the certainty equivalent cost of the idiosyncratic shock ε_j varies with the aggregate shocks f and g . When idiosyncratic risk is state-independent ($\Gamma_{f,\tau} = \Gamma_{g,\tau} = 0$), the effective loadings reduce to the systematic loadings $\kappa_\tau \gamma_\tau$ and $\kappa_\tau \lambda_\tau$ and $\tilde{Y}_\tau = Y_\tau$. When there is no idiosyncratic risk, additionally $\Gamma_{0,\tau} = 0$ and the factoring equation (OA.30) is trivial.

The multiplicative constant $e^{-\alpha\Gamma_{0,\tau}}$ does not depend on the portfolio choice \mathbf{x}^j , so it cancels from the first-order conditions. The portfolio choice problem only depends on the risk-adjusted labor income \tilde{Y}_τ . We restrict attention to a symmetric equilibrium in which all agents of the same type select the same portfolio.

OA.4.6 Portfolio choice

Let \mathbf{x}^j denote the vector of portfolio weights in risky assets, so that $R_p^j = R_f + \mathbf{x}^{j\prime}(\mathbf{R} - R_f\mathbf{1})$, with the residual fraction $1 - \mathbf{1}'\mathbf{x}^j$ invested in the risk-free asset. With the factoring in equation (OA.30), the portfolio choice problem reduces to

$$\max_{\mathbf{x}^j} \left\{ \mathbb{E}[R_p^j] - \frac{\alpha W_0}{2} \text{Var}(R_p^j) - \alpha \text{Cov}(\tilde{Y}_\tau, R_p^j) - \omega \text{Cov}(z, R_p^j) \right\}. \quad (\text{OA.33})$$

The first-order condition for asset i is

$$\mathbb{E}[R_i] - R_f = \alpha W_0 \text{Cov}(R_i, R_p^j) + \alpha \text{Cov}(R_i, \tilde{Y}_\tau) + \omega \text{Cov}(R_i, z). \quad (\text{OA.34})$$

In vector form, let Σ denote the return covariance matrix and let Σ^\dagger denote its Moore-Penrose pseudoinverse. We use the pseudoinverse because Σ need not be invertible: small- and big-firm returns have common f and g loadings within type, generating collinearity in the return covariance matrix.

$$\mathbf{x}^\tau = \underbrace{\frac{1}{\alpha W_0} \Sigma^\dagger [\mathbb{E}[\mathbf{R}] - R_f \mathbf{1} - \omega \text{Cov}(\mathbf{R}, z)]}_{\mathbf{x}^* \text{ (common component)}} - \underbrace{\frac{1}{W_0} \Sigma^\dagger \text{Cov}(\mathbf{R}, \tilde{Y}_\tau)}_{\boldsymbol{\delta}^\tau \text{ (hedging demand, type-specific)}}. \quad (\text{OA.35})$$

Each agent's optimal portfolio has two key components:

1. A common component \mathbf{x}^* that is identical across agent types, which captures the exposure of aggregate financial wealth to f as well as the impact of the exogenous utility shifters z .
2. The hedging demand $\boldsymbol{\delta}^\tau$, which differs across types and is proportional to $\text{Cov}(\mathbf{R}, \tilde{Y}_\tau)$. Each type tilts its financial portfolio away from assets that covary with its risk-adjusted labor income.

OA.4.7 Aggregation and pricing

Define the population-weighted risk-adjusted labor-income loadings:

$$\bar{\phi}_f \equiv \mu \phi_{f,S} + (1 - \mu) \phi_{f,B}, \quad \bar{\phi}_g \equiv \mu \phi_{g,S} + (1 - \mu) \phi_{g,B}. \quad (\text{OA.36})$$

Symmetrically with Y , define risk-adjusted aggregate labor income \tilde{Y} as

$$\tilde{Y} = \mu \tilde{Y}_S + (1 - \mu) \tilde{Y}_B = \bar{Y}_0 + \bar{\phi}_f f + \bar{\phi}_g g. \quad (\text{OA.37})$$

Then, our main pricing result is given by Proposition 1.

Proposition 1 (General equilibrium pricing) *In the general model with type-specific κ_τ , population share μ , and state-dependent idiosyncratic risk:*

(a) *The covariance pricing equation for any tradable asset i is*

$$\mathbb{E}[R_i] - R_f = \alpha(\bar{\gamma} + \bar{\phi}_f) \text{Cov}[R_i, f] + \alpha \bar{\phi}_g \text{Cov}[R_i, g] + \omega \text{Cov}[R_i, z]. \quad (\text{OA.38})$$

(b) *The following linear SDF prices all tradable assets:*

$$M = \frac{1}{R_f} [1 - \alpha(\bar{\gamma} + \bar{\phi}_f) f - \alpha \bar{\phi}_g g - \omega z], \quad (\text{OA.39})$$

where the risk-free rate is

$$R_f = \frac{1}{\rho} \exp \left[\alpha(1 + \bar{Y}_0 - C_0) - \frac{1}{2} (\alpha^2 (\bar{\gamma} + \bar{\phi}_f)^2 \sigma_f^2 + \alpha^2 (\bar{\phi}_g)^2 \sigma_g^2 + \omega^2 \sigma_z^2) + \alpha \bar{\Gamma}_0 \right], \quad (\text{OA.40})$$

with $\bar{\Gamma}_0 \equiv \mu \Gamma_{0,S} + (1 - \mu) \Gamma_{0,B}$ capturing the state-independent certainty-equivalent component of idiosyncratic risk.

(c) *The price of g -risk, $\alpha \bar{\phi}_g$, decomposes into a systematic channel and an idiosyncratic-risk channel according to*

$$\alpha \bar{\phi}_g = \alpha \frac{\text{Cov}(g, \tilde{Y})}{\sigma_g^2} = \underbrace{\alpha [\mu \kappa_S \lambda_S + (1 - \mu) \kappa_B]}_{(\text{systematic channel})} + \underbrace{\alpha [\mu \Gamma_{g,S} + (1 - \mu) \Gamma_{g,B}]}_{(\text{idiosyncratic-risk channel})}. \quad (\text{OA.41})$$

Proof. From equation (OA.34) and $\tilde{Y}_\tau = \bar{Y}_0 + \phi_{f,\tau}f + \phi_{g,\tau}g$, the type- τ first-order condition can be written as

$$\mathbb{E}[R_i] - R_f = \alpha W_0 \text{Cov}(R_i, R_p^\tau) + \alpha \phi_{f,\tau} \text{Cov}(R_i, f) + \alpha \phi_{g,\tau} \text{Cov}(R_i, g) + \omega \text{Cov}(R_i, z).$$

Since $\tilde{Y}_\tau = \bar{Y}_0 + \phi_{f,\tau}f + \phi_{g,\tau}g$ and the shocks are orthogonal, the hedging demand in equation (OA.35) is

$$\boldsymbol{\delta}^\tau = \frac{\phi_{f,\tau}}{W_0} \boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, f) + \frac{\phi_{g,\tau}}{W_0} \boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, g). \quad (\text{OA.42})$$

The population-weighted aggregate hedging demand is

$$\begin{aligned} \bar{\boldsymbol{\delta}} &\equiv \mu \boldsymbol{\delta}^S + (1 - \mu) \boldsymbol{\delta}^B \\ &= \frac{\bar{\phi}_f}{W_0} \boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, f) + \frac{\bar{\phi}_g}{W_0} \boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, g). \end{aligned} \quad (\text{OA.43})$$

Market clearing requires $\mathbf{x}^* - \bar{\boldsymbol{\delta}} = \mathbf{x}^{\text{mkt}}$. This market-clearing condition is unchanged if initial tradable equity ownership is allocated differently across worker types. Such a change alters the trades each type must make to reach its desired final portfolio, but with unconstrained trading and CARA preferences it does not alter final dollar risky positions or their population-weighted aggregate. Combining market clearing with the type-level first-order condition gives

$$\mathbb{E}[R_i] - R_f = \alpha W_0 \text{Cov}(R_i, R_p^{\text{mkt}}) + \alpha \bar{\phi}_f \text{Cov}(R_i, f) + \alpha \bar{\phi}_g \text{Cov}(R_i, g) + \omega \text{Cov}(R_i, z).$$

Since aggregate dividends load only on f , $W_0 \text{Cov}(R_i, R_p^{\text{mkt}}) = \bar{\gamma} \text{Cov}(R_i, f)$, which yields equation (OA.38).

Because f , g , and z are mean zero and mutually orthogonal, equation (OA.38) holds if and only if there is a linear SDF that satisfies equation (OA.39) (Cochrane, 2009, Section 6.3).

It remains to derive the risk-free rate. Let $C_{0,\tau}$ denote date-0 consumption for type τ , and let $C_0 = \mu C_{0,S} + (1 - \mu) C_{0,B}$ denote aggregate date-0 consumption. The risk-free bond

Euler equation for type τ is

$$1 = \rho R_f \mathbb{E} \left[e^{-\alpha(C_1^\tau - C_{0,\tau}) - \omega z} \right].$$

Using the idiosyncratic-risk factoring in equation (OA.30), this becomes

$$1 = \rho R_f e^{\alpha C_{0,\tau} - \alpha \Gamma_{0,\tau}} \mathbb{E} \left[e^{-\alpha(\tilde{Y}_\tau + W_0 R_p^\tau) - \omega z} \right].$$

Portfolio market clearing implies $\mathbf{x}^\tau = \mathbf{x}^{\text{mkt}} + \bar{\boldsymbol{\delta}} - \boldsymbol{\delta}^\tau$. Therefore, using equation (OA.42), the type-specific hedging demand offsets the type- τ adjusted labor-income loading, while the aggregate hedging demand adds back the population-weighted adjusted loading:

$$\tilde{Y}_\tau + W_0 R_p^\tau = \bar{Y}_0 + \phi_{f,\tau} f + \phi_{g,\tau} g + (1 + \bar{\gamma} f) + (\bar{\phi}_f - \phi_{f,\tau}) f + (\bar{\phi}_g - \phi_{g,\tau}) g \quad (\text{OA.44})$$

$$= 1 + \bar{Y}_0 + (\bar{\gamma} + \bar{\phi}_f) f + \bar{\phi}_g g. \quad (\text{OA.45})$$

Put differently, both types hold portfolios which deliver the same loading of total wealth plus risk adjusted labor income on both f and g . More explicitly, because f and g are the normalized factor realizations and returns are affine in (f, g, z) (and the fact that tradable assets span f and g as noted in Remark 1), the portfolios with weights $\boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, f)$ and $\boldsymbol{\Sigma}^\dagger \text{Cov}(\mathbf{R}, g)$ are the Moore-Penrose factor-mimicking portfolios for the f and g payoff innovations, respectively. The pseudoinverse selects the minimum-norm replicating portfolio; adding any portfolio in the null space of $\boldsymbol{\Sigma}$ leaves the payoff innovation unchanged. Up to deterministic components absorbed by the risk-free position, this implies

$$W_0(\bar{\boldsymbol{\delta}} - \boldsymbol{\delta}^\tau)'(\mathbf{R} - R_f \mathbf{1}) = (\bar{\phi}_f - \phi_{f,\tau}) f + (\bar{\phi}_g - \phi_{g,\tau}) g, \quad (\text{OA.46})$$

which justifies equation (OA.45). Hence, by normality and orthogonality,

$$1 = \rho R_f \exp \left[-\alpha(1 + \bar{Y}_0) + \alpha C_{0,\tau} - \alpha \Gamma_{0,\tau} + \frac{1}{2}(\alpha^2(\bar{\gamma} + \bar{\phi}_f)^2 \sigma_f^2 + \alpha^2(\bar{\phi}_g)^2 \sigma_g^2 + \omega^2 \sigma_z^2) \right].$$

Since the same risk-free rate must satisfy this equation for both types, $\alpha C_{0,\tau} - \alpha \Gamma_{0,\tau}$ must be

common across τ . Date-0 goods market clearing across types pins down this common term:

$$\alpha C_{0,\tau} - \alpha \Gamma_{0,\tau} = \mu(\alpha C_{0,S} - \alpha \Gamma_{0,S}) + (1 - \mu)(\alpha C_{0,B} - \alpha \Gamma_{0,B}) = \alpha C_0 - \alpha \bar{\Gamma}_0.$$

Substituting this expression into the Euler equation and rearranging yields equation (OA.40).

Finally, using equation (OA.32), equation (OA.36), and $\lambda_B = 1$,

$$\begin{aligned} \alpha \bar{\phi}_g &= \alpha[\mu(\kappa_S \lambda_S + \Gamma_{g,S}) + (1 - \mu)(\kappa_B + \Gamma_{g,B})] \\ &= \alpha[\mu \kappa_S \lambda_S + (1 - \mu) \kappa_B] + \alpha[\mu \Gamma_{g,S} + (1 - \mu) \Gamma_{g,B}], \end{aligned}$$

which is equation (OA.41).

■

Remark 3 (Connection to a model in which private firms are non-traded) *Consider an alternative model extension (without non-traded labor income) in which private firms are non-traded with dividends that follow the same factor structure as in the baseline model (i.e., with exposure to both f and g), but also with idiosyncratic risk as described in Subsection OA.4.3. Note that in this alternative model setup, systematic risk (f and g) is still spanned by tradable assets (see Remark 1). Hence, agents can freely choose any linear exposure to aggregate risk via their portfolio choices with respect to tradable public firms. In this model, our assumption that all agents initially hold the same tradable financial endowment is essentially a normalization rather than a substantive restriction for the pricing results. Specifically, if S -workers were initially endowed with tradable claims to S firms and B -workers with tradable claims to B firms, unconstrained date-0 trading will allow agents to trade to undo heterogeneous initial exposures to aggregate risk. With CARA preferences, this change in the initial allocation of tradable claims affects type-specific date-0 wealth and risk-free bond positions for the two agent types, but not final dollar risky positions or equilibrium risk premia. In this alternative model where private firm ownership claims are non-tradable, we could collect differences between actual endowments of unlisted firms and our assumed pro-rata starting point into the nontraded endowment (similar to labor income in the main model extension above, or any other non-traded asset). However, implications for post-trade*

consumption would be identical to the current formulation (which proceeds as if claims publicly and privately listed firms' profits are tradable). The pricing conclusions about g below continue to hold, after moving the relevant cash flows out of the financial portfolio and into labor income.

The key role played by additional non-tradable income is twofold: 1) to capture additional sources of aggregate resources available to consume which are not captured by a claim to firm ownership (proxied for by the BMF as defined in our baseline model) and 2) to capture uninsurable risk. Thus, a model in which private firms are segmented from public firms (and non-tradable) would deliver the same implications as our model extension that incorporates non-tradable labor income.

OA.4.8 Conditions under which g is not priced ($\bar{\phi}_g = 0$)

$\bar{\phi}_g$ in equation (OA.41) indexes departures from our baseline model in which g is not priced. The systematic channel reflects the aggregate g -loading of labor income through the type-specific scales κ_τ (i.e., when $\mu\kappa_S\lambda_S + (1-\mu)\kappa_B \neq 0$, aggregate labor income loads on g). The idiosyncratic-risk channel reflects the aggregate certainty-equivalent effect of state-dependent idiosyncratic risk on g -exposure (i.e., when $\mu\Gamma_{g,S} + (1-\mu)\Gamma_{g,B} \neq 0$, aggregate *risk-adjusted* labor income loads on g). In this subsection, we explore conditions under which g is not priced in the extended model.

Assumption 1 (Labor income uncorrelated with g) *Suppose that the following condition holds:*

$$\mu\kappa_S\lambda_S + (1-\mu)\kappa_B = 0. \tag{OA.47}$$

Under Assumption 1, aggregate labor income, $\mu Y_S + (1-\mu)Y_B$, does not load on g .

One simple setup which satisfies Assumption 1 is where both types of workers have similar exposures to the *per-worker* cash flow shock that their employers experience. Specifically, suppose

$$\kappa_S = \kappa \frac{\theta}{\mu}, \quad \kappa_B = \kappa \frac{1-\theta}{1-\mu}.$$

Then, using $\lambda_S = -(1 - \theta)/\theta$ and $\lambda_B = 1$,

$$\mu\kappa_S\lambda_S + (1 - \mu)\kappa_B = \kappa\theta\lambda_S + \kappa(1 - \theta)\lambda_B = 0.$$

Two conditions are likely to hold in practice. First, a smaller fraction of workers work at small firms relative to the fraction of small firms in the economy, so $\mu < \theta$. Second, cash flow shocks experienced by small firms are spread out across fewer workers. Assumption 1 ensures that these two forces offset each other so that the aggregate non-tradable labor income endowment does not depend on g . A similar intuition likely also applies to exposures to f , but we do not need to take a stand on this for our main results regarding pricing of g .

Assumption 2 (Idiosyncratic certainty equivalents uncorrelated with g) *With respect to the redistributive shock g , the certainty-equivalent loadings satisfy*

$$\mu\Gamma_{g,S} + (1 - \mu)\Gamma_{g,B} = 0. \tag{OA.48}$$

Assumption 2 imposes an intuitive restriction on state-dependent idiosyncratic risk. A sufficient condition is that Assumption 1 holds and

$$\Gamma_{g,\tau} = \chi\kappa_\tau\lambda_\tau \tag{OA.49}$$

for a scalar χ that does not depend on τ .

The literature on state-dependent tail risk often emphasizes the idea that aggregate shocks are concentrated on a subset of agents rather than spread evenly across the population (e.g., Mankiw, 1986; Guvenen et al., 2014; Catherine, 2022; Schmidt, 2025). In this setting, that concentration means the idiosyncratic risk faced by workers of type τ responds to the same type-specific shock that moves their systematic labor income. Equation OA.49 imposes that the certainty-equivalent effect of this state-dependent idiosyncratic risk is proportional to the loading of Y_τ on g , with a common proportionality constant across types. Additionally, Assumption 2 automatically holds when idiosyncratic risk does not load on g (i.e., $\Gamma_{g,\tau} = 0$). Assumptions 1 and 2 jointly imply that aggregate risk-adjusted labor income does not load on g and the aggregate hedging demands against aggregate and state-dependent idiosyncratic

risk to g net out to zero. This property leads to the following corollary to Proposition 1, the proof of which is immediate.

Corollary 1 (Benchmark equilibrium under which $\bar{\phi}_g = 0$) *Under Assumptions 1 and 2, equation (OA.41) simplifies such that $\bar{\phi}_g = 0$.*

Thus, Assumptions 1 and 2 yield an economically interpretable benchmark in which the systematic and idiosyncratic risk channels are both zeros, so the two channels jointly imply $\bar{\phi}_g = 0$. This generates an isomorphism with our baseline model, which we summarize in the following corollary to Proposition 1.

Corollary 2 (Isomorphism between the extended and baseline models) *Suppose Assumptions 1 and 2 hold, so that $\bar{\phi}_g = 0$. The extended model is isomorphic to the baseline model with α and ρ replaced by*

$$\tilde{\alpha} \equiv \alpha \frac{\bar{\gamma} + \bar{\phi}_f}{\bar{\gamma}} \tag{OA.50}$$

and

$$\tilde{\rho} \equiv \rho \exp[\tilde{\alpha}(1 - D_0) - \alpha(1 + \bar{Y}_0 - C_0) - \alpha\bar{\Gamma}_0], \tag{OA.51}$$

where D_0 is the date-0 consumption term in the baseline risk-free-rate expression in equation (OA.3). Under this mapping, g is unpriced, so the IFF has zero risk premium. All beta-gap formulas, the VMF/BMF decomposition, the size-anomaly mechanism, and the IFF alpha results therefore hold with the same functional forms as in the baseline model.

The key point is that Assumptions 1 and 2 ensure that aggregate risk-adjusted labor income does not load on g . The remaining effect of nontraded labor income and state-dependent idiosyncratic risk is to rescale the price of aggregate f -risk – which is isomorphic to changing risk aversion to $\tilde{\alpha}$ – and to shift the risk-free rate – which is isomorphic to changing the rate of time preference parameter to $\tilde{\rho}$. Hence, the asset pricing predictions of the baseline model are observationally equivalent to the predictions of the extended model under Assumptions 1 and 2.

OA.4.9 Summary of conditions that lead to priced g ($\bar{\phi}_g \neq 0$)

The restrictions in Assumptions 1 and 2 likely do not hold exactly in the data. Equation (OA.41) in Proposition 1 makes clear that Assumption 1 shuts down the systematic channel of $\bar{\phi}_g$, and Assumption 2 shuts down the idiosyncratic risk channel. Departures from either assumption are sufficient to deliver $\bar{\phi}_g \neq 0$ and a priced g . Both routes share a common economic origin: each creates an asymmetry in how the redistributive shock g affects workers' effective (risk-adjusted) income or utility, so that g no longer “washes out” of aggregate risk-adjusted labor income \tilde{Y} . Below, we argue that empirically-relevant parameterizations associated with relaxing each of these assumptions both point to a negative price of g risk (i.e., $\bar{\phi}_g < 0$).

Relaxing Assumption 1: $\mu\kappa_S\lambda_S + (1 - \mu)\kappa_B \neq 0$ (**systematic channel**). If the per-worker scales κ_τ depart from the cancellation benchmark $\kappa_S = \kappa\theta/\mu$, $\kappa_B = \kappa(1 - \theta)/(1 - \mu)$ described under Assumption 1, then aggregate labor income $\mu Y_S + (1 - \mu)Y_B$ loads on g directly and the systematic term in equation (OA.41) is nonzero. Since $\lambda_S = -(1 - \theta)/\theta < 0 < \lambda_B = 1$, the systematic channel is negative whenever S-workers' per-worker exposure to S-firm cash-flow shocks, κ_S , is sufficiently large relative to B-workers' exposure, κ_B . Such a parameterization is empirically natural if small firms are more labor-intensive than large firms, so that their workers' incomes are more tightly tied to firm outcomes, or if workers at small firms bear cash-flow shocks more directly because small firms offer less internal wage smoothing. Under this parameterization, a positive g realization—which reallocates cash flows from S-firms toward B-firms—lowers aggregate risk-adjusted worker welfare on net and produces a negative price of g -risk, $\bar{\phi}_g < 0$. This parameter restriction is consistent with empirical evidence from the labor and asset pricing literatures. For instance, Autor et al. (2020) find that small firms tend to have higher labor shares than large firms (see also Hartman-Glaser et al., 2019). Furthermore, Donangelo et al. (2019) show that firms with higher labor shares have operating profits that are more sensitive to economic shocks and earn higher expected equity returns.

Relaxing Assumption 2: $\mu\Gamma_{g,S} + (1 - \mu)\Gamma_{g,B} \neq 0$ (**idiosyncratic-risk channel**). If the population-weighted average of state-dependent idiosyncratic-risk certainty-equivalent loadings $\Gamma_{g,\tau}$ is not zero, then the idiosyncratic-risk channel contributes to $\bar{\phi}_g$. One force

which could generate such a deviation is if state-dependent idiosyncratic risk is stronger for S-workers than for B-workers—as could happen if smaller firms are less able to provide insurance to their workers. These smaller firms are less likely to have internal labor markets and more likely to be financially constrained, and therefore may pass on more risk to their workers in response to bad shocks. In such a case, a positive g realization, which reallocates cash flows away from S-firms, would lower the average certainty equivalent of idiosyncratic labor-income risk on net. Equivalently, one would expect $\mu\Gamma_{g,S} + (1 - \mu)\Gamma_{g,B} < 0$, which generates a negative price of g -risk ($\bar{\phi}_g < 0$). This parameter restriction is consistent with empirical evidence from the labor and asset pricing literatures. For instance, Davis and Haltiwanger (1992) document that private firms (which are more likely to be small) have more volatile growth rates than public firms. Small firms are also less likely to diversify this risk for their workers: Guiso et al. (2005) use matched employer-employee data to show that more volatile firms and those with more credit risk (i.e., firms that are also more likely to be small) do not insure workers against idiosyncratic risk as much as less volatile or less credit constrained firms.

Both cases discussed above are directionally consistent with the negative point estimate of the IFF premium (see Subsection 6.1). The fact that our estimated IFF premium is economically small and statistically indistinguishable from zero suggests that any such violations of the cancellation restrictions in Assumptions 1 and 2 are modest in practice.

OA.5 BMF and IFF construction methodology: Additional details

OA.5.1 Methodology: Details

In this section, we describe our methodology for constructing the BMF in more detail. We observe the returns, sales, and market values of *public* firms, and the VMF is a value-weighted index of the returns on these firms. We would like to construct an index whose return reflects the value-weighted return on *all* firms, public and private, which we call R_{all} (a scaled version of the BMF). Reweighting the public sample can deliver a consistent estimator for the R_{all} if we adjust for: 1. the probability a firm is listed (i.e., selection), and 2. the economic scale

we want to represent (i.e., market value).

Thus, our goal is to create a methodology that allows us to estimate the value-weighted return on the population of firms (public and private). The fundamental issue in constructing this quantity is that we do not observe market values and returns for private firms. In this section, we describe a procedure that produces an unbiased estimate of R_{all} (under mild assumptions), which we will use as the basis for constructing our BMF.

Let the universe of firms (public and private) be indexed by $i = 1, \dots, N$ and be binned into portfolios indexed by $p(i) \in \{1, \dots, P\}$. In our case, we will consider binning firms into portfolios based on conditional double sorts first on industry and then, within an industry, on sales. The industry bins are defined using the SIC 2-digit industry code grouping from KMZ. Let S_i denote the sales of firm i . Let $k(i) \in \{1, \dots, K\}$ denote the firm type, with each type distinguished by its valuation multiple, α_k (in our case, market equity divided by sales). Thus, the market value of firm i is given by $V_i = \alpha_{k(i)} S_i$. Note that while $\alpha_{k(i)}$ is observable for publicly listed firms, it is unobservable for private firms. Thus, we treat $\alpha_{k(i)}$ like it is unobserved for the purposes of our analysis.

The return on all firms can be expressed as

$$R_{all} = \frac{\sum_{i=1}^N V_i R_i}{\sum_{i=1}^N V_i} \quad (\text{OA.52})$$

where R_i is the return on firm i . Let \mathcal{L} denote the set of publicly-listed firms and \mathcal{U} denote the set of unlisted firms. Our goal is to build a consistent estimator for R_{all} , \hat{R}_{all} such that $\mathbb{E}[\hat{R}_{all}] = R_{all}$, using only observed data on $(S_i, V_i, R_i)_{i \in \mathcal{L}}$ and total sales by size and industry among all firms (public and private) from KMZ, S_p^{all} . Note that we can use the observed data on public firms to compute the total sales within each bin among public firms, S_p^{pub} . In principle, we cannot recover R_{all} exactly because we do not observe the firm type across all firms. Pragmatically, to circumvent this issue, we assume that,

1. conditional on $p(i)$, firm sales S_i is independent of latent type $k(i)$, and
2. conditional on $p(i)$ and $k(i)$, returns R_i are i.i.d.

The first assumption allows us to recover the correct share of market value associated with

firms of each type within portfolio $p(i)$, while the second assumption ensures that observed returns on listed firms are representative of the unobserved returns on unlisted firms of the same type within portfolio $p(i)$. These assumptions allow us to recover a consistent estimator for R_{all} , \hat{R}_{all} , as we describe below.

Sampling process: Assume a firm's listing probability factors as

$$Pr \{i \in \mathcal{L}|p, k\} = \rho_{p,k} = f_p \lambda_k, \quad \lambda_k = \frac{\alpha_k}{\delta_p}, \quad 0 < f_p \leq 1, \quad \delta_p > 0 \quad (\text{OA.53})$$

where f_p is a bin-specific scaling factor and δ_p is a positive scaling constant. Thus, higher-multiple firms are more likely to be public conditional on belonging to bin p .⁴³ This proportionality will significantly simplify the reweighting scheme, as shown below. A key implication of this proportionality is that the unknown multiple $\alpha_{k(i)}$ cancels in any expression of the form $V_{k(i)}/\rho_{p,k}$ because $V_{k(i)} \equiv \alpha_{k(i)}S_i$ and only the ratio δ_p/f_p matters (i.e., information that is only bin specific, but unrelated to the unobserved type k).

A simple microfoundation: Suppose that, within each bin p , the manager who runs each firm i trades off the costs versus the benefits of being a public firm. In particular, she compares her idiosyncratic cost of being public $u_i \sim U[0, \bar{u}_p]$ against the benefit of listing, $b_p \alpha_k$.⁴⁴ This structure delivers the result that

$$Pr \{i \in \mathcal{L}|p, k\} = Pr \{b_p \alpha_k \geq u_i\} = \rho_{p,k} = \frac{\alpha_k}{\bar{u}_p b_p} \equiv \frac{f_p}{\delta_p} \alpha_k. \quad (\text{OA.54})$$

Thus, the assumption that the listing probability is proportional to the valuation multiple in equation (OA.53) can be thought of as being consistent with the notion that firms with higher valuation multiples (i.e., firms with more valuable investment opportunities) receive higher benefits from access to public markets.

⁴³We do not have a characteristic in our model related to a firm's valuation ratio, which we treat as unobserved there. Thus, as a stand-in, in our model calibration we assume that a firm's probability of going public is related to its size. This assumption is not only consistent with our empirical motivation from Section 3.2, but it is also consistent with the logic here as long as size is correlated with valuation multiples, for which there is empirical evidence—small firms tend to be value firms and large firms tend to be growth firms.

⁴⁴Within each bin, p , we assume that costs and benefits are both proportional to sales, so u_i and $b_p \alpha_k$ can be thought of as cost and benefit per dollar of sales. We do this for simplicity of exposition. However, we do not need to impose this proportionality restriction across bins to generate this microfoundation.

Step-by-step construction of \hat{R}_{all} : Our methodology for constructing \hat{R}_{all} is as follows:

1. **Compute a within-bin sales-weighted mean multiple:** We would like to estimate $\mathbb{E}_p[\alpha_k]$, the expected value of α_k for *all* firms conditional on the firms being in portfolio p . We proxy for this quantity using observable data from publicly listed firms as:

$$\bar{\alpha}_p = \frac{\sum_{i \in \mathcal{L}, p(i) \in p} S_i \alpha_i}{\sum_{i \in \mathcal{L}, p(i) \in p} S_i} = \frac{\sum_{i \in \mathcal{L}, p(i) \in p} V_i}{\sum_{i \in \mathcal{L}, p(i) \in p} S_i} \equiv \frac{V_p^{pub}}{S_p^{pub}}.$$

2. **Compute the bin-specific listing probability:** The average listing probability in bin p is

$$\bar{\rho}_p = \frac{S_p^{pub}}{S_p^{all}},$$

which also follows from the assumption that S_i is conditionally independent of type $k(i)$ within each portfolio. Note that this implies that

$$\frac{\delta_p}{f_p} = \frac{\bar{\alpha}_p}{\bar{\rho}_p},$$

hence we can capture the quantities we need to adjust for selection into public markets using two observable sufficient statistics for each bin p : 1) We can measure $\bar{\rho}_p$ from the KMZ data, and 2) we can measure $\bar{\alpha}_p$ using data on listed firms.

3. **Compute inverse-probability weights:** Apply the generic Horvitz-Thompson (inverse sampling probability) weight, $V_i/\rho_{p,k}$, to a public firm

$$V_i^* \equiv \frac{V_i}{\rho_{p,k}} = \frac{\alpha_{k(i)} S_i}{f_{p(i)} \alpha_{k(i)} / \delta_{p(i)}} = \frac{\delta_{p(i)}}{f_{p(i)}} S_i.$$

After normalizing across the public sample, each firm receives a portfolio weight of

$$w_i^* = \frac{V_i^*}{\sum_{i \in \mathcal{L}} V_i^*}.$$

4. Aggregate the return

$$\hat{R}_{all} = \sum_{i \in \mathcal{L}} w_i^* R_i,$$

which is a consistent estimate of R^{all} when the number of firms is large because $\frac{\sum_{i \in \mathcal{L}} V_i^* R_i}{\sum_{i \in \mathcal{L}} V_i^*} \xrightarrow{p} \frac{\sum_{i=1}^N V_i R_i}{\sum_{i=1}^N V_i}$ given our two assumptions enumerated above.

Note that $\bar{\alpha}_p$ is an asymptotically biased estimator for $\mathbb{E}_p[\alpha_k]$ due to the fact that selection of firm i into the public market is a function of $\alpha_{k(i)}$. This causes our estimator $\bar{\alpha}_p/\bar{\rho}_p$ to be related to δ_p/f_p as (in the limit as the number of firms in portfolio p , N_p , goes to infinity):

$$\lim_{N_p \rightarrow \infty} \frac{\bar{\alpha}_p}{\bar{\rho}_p} = \frac{\delta_p}{f_p} \left(\frac{\mathbb{E}_p[\alpha_k^2]}{\mathbb{E}_p[\alpha_k]^2} \right) = \frac{\delta_p}{f_p} \left(1 + \frac{\text{VAR}_p[\alpha_k]}{\mathbb{E}_p[\alpha_k]^2} \right). \quad (\text{OA.55})$$

Thus, the within-portfolio bias is related to the squared coefficient of variation of α_k 's within the portfolio of all firms. Unfortunately, we do not observe these α_k 's in the portfolio of all firms to estimate these quantities (indeed, if we could observe them, we could estimate the $\mathbb{E}_p[\alpha_k]$ directly). However, given our assumptions above, one can show that the expected value of α_k across all firms in portfolio p can be expressed as

$$\mathbb{E}_p[\alpha_k] = \frac{1}{\mathbb{E}_{p,\mathcal{L}} \left[\frac{1}{\alpha_k} \right]}, \quad (\text{OA.56})$$

where $\mathbb{E}_{p,\mathcal{L}}$ reflects an expectation with respect to the set of *publicly listed* firms. An unbiased estimator of $\mathbb{E}_{p,\mathcal{L}} \left[\frac{1}{\alpha_k} \right]$ (i.e., the expectation across all public firms in portfolio p in a given period) is

$$\frac{1}{\alpha_k} = \frac{1}{N_p} \sum_{i \in \{p,\mathcal{L}\}} \frac{S_i}{V_i}. \quad (\text{OA.57})$$

This tends to be a very noisy estimator since market values and sales fluctuate dramatically (their ratio more so), and thus it is unlikely to be an accurate approximation of $\mathbb{E}_{p,\mathcal{L}} \left[\frac{1}{\alpha_k} \right]$ in finite samples of firms in the portfolios we construct. On the other hand, if the coefficient of variation in equation (OA.55) is the same across portfolios within each time period, this

bias does not affect our reweighting scheme because it affects each portfolio in the same way. Thus, our preferred specification uses the biased but (likely) less volatile estimator for $\bar{\alpha}_p$ in equation (OA.55), which is exactly valid for our purposes under the assumption that the coefficient of variation is the same across portfolios in each period.⁴⁵

Interpretation: Within-bin sales weights neutralize the over-representation of high- α types because the sampling probability is proportional to the valuation multiple α (i.e., as opposed to using the more conventional value-weights, $V_i \equiv \alpha_{k(i)} S_i$). The factor $\bar{\alpha}_p / \bar{\rho}_p$ then rescales each bin to its full-economy value. The $1/\bar{\rho}_p$ adjusts aggregate sales for the fact that private firms are missing from the sample, while the valuation multiple $\bar{\alpha}_p$ converts sales to a market value.

Putting it all together, each firm receives the following (unscaled) weight in the construction of \hat{R}_{all} , expressed using directly-observable quantities:

$$w_i^* = \frac{\bar{\alpha}_p}{\bar{\rho}_p} S_i = \underbrace{\frac{V_p^{pub}}{S_p^{pub}}}_{\text{valuation mult.}} \underbrace{\frac{S_p^{all}}{S_p^{pub}}}_{\text{selection adj.}} \underbrace{S_i}_{\text{size}}, \quad (\text{OA.58})$$

which clarifies each component of the adjustment in an intuitive manner. The ratio S_p^{all}/S_p^{pub} as an adjustment for the selection of public firms to scale up the observed firm's sales S_i . The valuation multiple V_p^{pub}/S_p^{pub} then translates the scaled sales into an implied value weight. Whereas the above procedures formalize the assumptions necessary to interpret the resulting weight as an implied value weight reflecting all firms (public and private) and a microfoundation for this weight, equation (OA.58) is a more intuitive expression to help understand what our reweighting procedure does (and the one we cite in our main text in equation (13)).

⁴⁵For completeness, we do construct a version of the BMF using the estimator for $\bar{\alpha}_p$ in equation (OA.57) (but otherwise following the same procedure we describe in Section 5.1 that we use to construct our main BMF). This alternative BMF has an average return (7.32%) very close to that of our main BMF (7.30%) and the two are highly correlated at 99.2%. We also find results based on this alternative BMF are both qualitatively and quantitatively similar to our main results. However, this alternative BMF construction depends on a number of implementation choices to deal with the noisiness of the estimator in equation (OA.57). Most notably, we winsorize the stock-level S_i/V_i at the 1% level in each cross-section and average this ratio using a rolling 5-year window. Since results using this alternative BMF are similar to those using our main BMF, and to avoid needing to take a stand on these implementation choices, we use the more stable estimator in equation (OA.55).

We can use the expression in equation (OA.58) to aggregate implied value weights across the universe of firms in any dimension along which we can sort public stocks, although we focus on the sales and industry dimensions. Furthermore, we can use it to aggregate implied sales weights across the universe of firms along any dimension as well by using only the sampling adjustment and size terms (i.e., not pre-multiplying by the valuation multiple) to get the firm-level implied sales among the universe of firms implied by the current firm i .

OA.5.2 Interpretation of the redistributive shock g in the context of the BMF construction methodology

The microfoundation discussed above in Subsection OA.5.1 provides a rationale for why some firms are more likely to be public than others, and is closely connected to the interpretation of the redistributive shock g in our theoretical model (a scaled version of the IFF) that we discuss in Subsection 4.1 of the main paper. Firms with high listing benefits in our microfoundation are precisely the firms that are overrepresented in the public market, and g captures forces that redistribute resources between these firms and the rest of the economy in our model. We briefly discuss three firm characteristics that are related to a firm’s probability of going public—size, growth options/valuation multiples, and industry—for which there is empirical and theoretical evidence of real-side forces that redistribute resources across firms that vary along these dimensions in a manner consistent with the role of g .

Size. A growing literature documents that a small number of large firms account for a disproportionate and rising share of aggregate economic activity, profits, and market capitalization—often referred to as the rise of “superstar firms” (see, e.g., Autor et al., 2020; Kwon et al., 2024). Because these large firms are overrepresented in the public market, real-side forces that redistribute economic activity between smaller and larger firms generate exactly the kind of redistributive effects that g captures in our model. Gertler and Gilchrist (1994) provide one example of such redistribution. They find that small manufacturing firms experience larger output declines in response to monetary policy shocks compared to large manufacturing firms. Furthermore, they find that bank lending to small firms contracts whereas that to large firms initially rises. Finally, Crouzet and Mehrotra (2020) find that the top 1% of firms by size are less cyclically sensitive than the bottom 99%, and that these

differences in cyclical exposure only have modest impacts on macroeconomic aggregates.

Growth options/valuation multiples. Public firms tend to have more valuable growth opportunities and higher valuation multiples than private firms, which is itself an explanation for why high-multiple firms select into public markets (see, e.g., Kim and Weisbach, 2008; Chemmanur et al., 2010; Doidge et al., 2017). Because growth firms are overrepresented in the public market, g also captures redistribution of economic activity between value and growth firms, though the correlation with standard growth proxies is much weaker than with size. Papanikolaou (2011) and Kogan et al. (2017) provide examples of how shocks to growth options can have redistributive effects. In particular, they emphasize how investment-specific technology shocks can reallocate resources between growth firms and value firms (because growth firms derive more of their value from future investment opportunities, while value firms derive more of their value from assets in place).

Industry. The incentives to go public can vary across industries based on product market or growth characteristics (among others), and can lead to differences in equilibrium industry composition between public and private firms (see, e.g., Maksimovic and Pichler, 2001; Chemmanur et al., 2010; Chod and Lyandres, 2011). For example, capital-intensive industries such as manufacturing rely more heavily on external equity financing and may also have a higher optimal scale (e.g., due to fixed costs) than service-sector industries. As such, g partially reflects an aggregate of industry-specific shocks that reallocate resources towards industries that are overrepresented in publicly-traded firms. Papanikolaou (2011) describes a mechanism through which (good) investment-specific technology shocks benefit a sector that produces investment-goods (growth firms that are more likely to go public) relative to a sector that produces consumption-goods (value firms that are less likely to go public). However, results in Doidge et al. (2017) imply that within-industry firm size and growth characteristics are likely more first order for determining whether a firm within a certain industry goes public than the industry itself. Thus, the size and growth characteristics we discuss above are likely more first-order for understanding the reallocative shocks embodied by g . This logic is also consistent with our finding that our results are similar whether we construct our BMF using the size-industry portfolio sorting approach described in Subsection 5.2 or the size-only portfolio sorting approach described in Subsection 5.3.

OA.5.3 Data and empirical implementation: Details

The CRSP and Compustat data provide nearly all the data necessary to implement the methodology for constructing our BMF described in Section 5.1 equation (13) ($V_{p(i)}^{pub}$, $S_{p(i)}^{pub}$, and S_i). The only additional data we need is the total sales among all public and private firms in a given industry-sales bin, S_p^{all} . For this, we use the data from Kwon et al. (2024) (KMZ).

The Internal Revenue Service (IRS) collects sales, assets, and net income data on all firms (both public and private) in the US. The IRS also provides industry-level tabulations until 2013 when a rule change (IRS Publication 1075) precluded this level of granularity from being released. KMZ provide a digital version of this data tabulated by industry and size, and we use their tabulated sales to estimate S_p^{all} for all firms associated with each of our sales-industry portfolios.

KMZ also provide aggregates based on net income, assets, and employment. We do not use net income because the data ends in 1974. We use sales as opposed to assets because sales should be more reflective of the overall economic contribution of firms regardless of their particular mix of assets and labor used for production. We do not use employment for a similar reason. Regardless, in unreported results we find that using either assets or employment leads to similar results as using sales in our aggregation procedure. KMZ also provide aggregates based on sales-only groups. The former data spans 1963–2013 and the latter from 1963–2018.

Within a given industry, KMZ aggregates total firm sales by bins sorted on firm sales breakpoints. In general, these KMZ sales breakpoints do not align with the NYSE sales breakpoints used to construct our portfolios of public firms. To address this issue, for each industry-sales portfolio we identify which KMZ industry-sales bins are spanned by the NYSE breakpoints and add those KMZ sales totals together. When an NYSE breakpoint falls between two KMZ sales breakpoints, we assume the number of firms within that KMZ sales bin follows a Pareto distribution and use the estimated distribution to compute the implied total sales of public and private firms in the public firm sales range that falls within the KMZ bin.⁴⁶ This allows us to estimate the only additional piece of information needed to

⁴⁶When needed, we estimate the Pareto distribution shape and scale parameters by imposing two conditions: 1. That the implied CDF match the fraction of firms below the KMZ sales breakpoint immediately below the current NYSE breakpoint, and 2. That the implied CDF match the fraction of firms above the

implement the methodology from Section 5.1, which is total firm sales (private and public) for a given industry-size bin, S_p^{all} . KMZ also provide aggregated data based on a sales-only aggregation (i.e., not conditional on being within certain industries). We use this data and follow the same procedure as described above (without first sorting into industries) to construct our alternative “sales-only” BMF.

Thus, we can use the data described above to compute our empirical proxy for monthly returns that reflect returns to the value-weighted portfolio of all firms in the economy, \hat{R}_{all} .

OA.6 PCA: Additional details

In this section, we discuss some additional details related to our PCA described in Section 5.3.

OA.6.1 PCA details

It is well known that equity returns follow a strong factor structure (see, for example, Kozak et al., 2018). When returns follow a factor structure, average returns can be described by a linear relationship with factor loadings (Ross 1976, 1977). Given these insights, Chamberlain and Rothschild (1983) and Connor and Korajczyk (1986, 1988) pioneered work related to using PCs to explain average asset returns. Using PCs as factors can be thought of as a statistical and complementary method to what has become the *de facto* method of constructing factors from portfolios sorted on stock characteristics.⁴⁷ Indeed, Kozak et al. (2018) conclude that there is nothing special about such factor models and their ability to explain average returns and note that a model with a small number of PCs does about as well explaining average returns as any of the extant factor models they study. Most importantly, the asymptotic properties of these estimators are well understood, and one can obtain consistent

KMZ sales breakpoint immediately above the current NYSE breakpoint. In cases where the NYSE breakpoint is above the highest KMZ breakpoint, we instead identify the shape and scale parameters by imposing the following two conditions: 1. That the implied CDF match the fraction of firms below the KMZ sales breakpoint immediately below the current NYSE breakpoint (i.e., the highest KMZ breakpoint in this case), and 2. That the implied average firm sales value in this range match that in the KMZ data.

⁴⁷Some notable recent extensions of the basic PCA framework include Kelly et al. (2019), who develop an instrumented PCA that helps provide a link between factors constructed from characteristics to those constructed from PCs. Lettau and Pelger (2020) show how imposing restrictions via their “risk premium PCA” allows their constructed factors to better fit the cross-section of expected returns in addition to explaining time series variation in returns as with a regular PCA. Giglio and Xiu (2021) show how PCs can be used to consistently estimate risk premia on both traded and nontraded factors.

estimates of common factors with a large number of portfolios and time periods.⁴⁸

We choose the 382 portfolios represented by the portfolio sorts described in Online Appendix Table OA.3 for our PCA to balance a few tradeoffs. First, we would like to use a broad cross-section of portfolios. Second, we need both equal-and value-weighted versions of the portfolios. The equal-weighted portfolios are used for the PCA, and the value-weighted portfolios are used in our asset pricing tests (as per standard convention). Third, we would like the portfolios to be as standard and easily accessible as possible. Most portfolios are based on 5x5 double sorts. However, we include three sets of single-sorted portfolios (size, beta, and book-to-market) because of the historical importance of the anomalies represented by these portfolios. Note all these portfolios are a superset of the equity portfolios used for the PCA analysis in Giglio and Xiu (2021).

OA.6.2 Projected BMF details

Here we provide more details on how we construct our PCA-based (or IPCA-based) BMFs as robustness checks on our main BMF as described in Section 5.3. We describe a general method for projecting our main BMF onto potentially multiple PCs, then using this projection as a proxy for \hat{R}_{all} to ultimately identify alternative BMFs according to equation (14). However, when we only project onto one PC this methodology recovers exactly the same BMF as if we were to use that PC directly as a proxy for \hat{R}_{all} . Pragmatically, this is how we construct our alternative BMFs (because we only use the first PC or IPC), but the procedure described herein provides the logic as to why this is equivalent to using PCs to identify our main BMF's exposure to these PCs, and how these alternative BMF constructions represent projections of our main BMF onto these PCs.

In general, we can construct our alternative BMFs by projecting our main BMF (identified

⁴⁸Bai and Ng (2006) proved that errors associated with estimating factors do not affect the limiting distribution of factor-augmented vector autoregression (FAVAR) estimators. In other words, estimates of factors are consistent, and subsequent regression estimates that use estimated PCs have the same asymptotic properties as if the true unobserved factors were used instead. This property holds when $\frac{T}{N} \rightarrow 0$ as $N, T \rightarrow \infty$. The sample sizes that are used to construct estimates of factors and loadings usually involve a number of stocks that is an order of magnitude larger than the number of time periods, so this condition approximately holds. Even when this condition fails, the estimated factors can be treated as data, and \sqrt{T} -consistent estimates of second-stage regression parameters obtain when $\frac{\sqrt{T}}{N} \rightarrow 0$, but standard errors need to be corrected to reflect the fact that a generated regressor is used. Stock and Watson (2002) provide conditions under which factors are consistently estimable even when loadings exhibit moderate time variation.

in equation (14)) onto one or more PCs (or IPCs from Kelly et al. (2019)), then use this as an alternative proxy for \hat{R}_{all} and follow the projection procedure described in Section 5.2 as usual to identify the alternative BMFs and IFFs. To be clear, we project our BMF onto PCs using the following full-sample regression

$$BMF_{t+1} = a_{PC} + b'_{PC} \cdot PC_{t+1} + \varepsilon_{BMF,t+1}, \quad (\text{OA.59})$$

where BMF_{t+1} is our main BMF and PC_{t+1} is a vector of PCs/IPC. Note that we use excess portfolio returns to construct our PCs. Furthermore, PCs are constructed by applying the PCA transformation matrix directly to our 382 excess return portfolios so that the resulting PCs are weighted combinations of tradable portfolios. We can view the component projected onto the PCs as another scaled version of our BMF and treat it as an alternative proxy for \hat{R}_{all} as

$$\hat{R}_{all,t+1}^{alt} - R_f \equiv b'_{PC} \cdot PC_{t+1}. \quad (\text{OA.60})$$

We then simply follow the procedure described in equation (14) (substituting in $\hat{R}_{all,t+1}^{alt} - R_f$ for $\hat{R}_{all,t+1} - R_f$) to give this alternative BMF the same scaling with respect to the VMF as our main BMF.

OA.7 Local projection impulse response analysis: Additional details

Commentary in this section refers to results discussed at a high level in Section 6.3 and to various panels from Figure 5.

OA.7.1 LPIR macro-variable choices

We begin with a description of the macro-aggregates and other return series we use in the LPIRs and their links to SDFs from standard macro-finance models.

Representative agent models: Panels A–D show LPIRs for measures of aggregate risk that appear in common representative agent-based models with complete markets. Consumption growth (see, e.g., Lucas, 1978, and the many papers that followed) responds

strongly and positively to the BMF but has no statistically significant response to the IFF (if anything, it is negative). This is also true of related aggregate consumption risk measures including GDP growth, total income growth, and (the negative of) real uncertainty (see, e.g., Bansal et al., 2014, for a discussion of the pricing of uncertainty). Note that the mean-reverting response of real uncertainty is expected since this is a stationary time series. We also expect this mean reversion for all other stationary time series we investigate.

Limited participation models: Given that a high fraction of aggregate stock market wealth (and overall wealth, for that matter) is owned by a small subset of high-income/high-net worth investors, we also consider measures that are more relevant in models that emphasize the concentration of aggregate risks onto market participants (see, e.g., Basak and Cuoco, 1998). Panels E–F show LPIRs for some common measures of aggregate risk that appear in limited participation models such as the income growth of top earners (see, e.g., Malloy et al., 2009; Parker and Vissing-Jørgensen, 2009; Campbell et al., 2016) and proprietor’s income as in Heaton and Lucas (2000). In both cases, the BMF generates a strong and positive response in the variables (albeit with some reversion in the latter case), whereas the IFF generates no response.

Incomplete markets models: In the presence of imperfect risk sharing and state-dependent distributions of idiosyncratic shocks to agents’ consumption and wealth, higher moments of shocks that redistribute consumption across different agents enter the stochastic discount factor (see, e.g., Panageas et al., 2020, for a survey). Panels G–I show LPIRs for some common measures related to undiversifiable risk in incomplete markets models including the skewness measure from Schmidt (2025), unemployment, and the common idiosyncratic volatility (CIV) measure from Herskovic et al. (2016). In all cases, the BMF generates a strong and positive response, whereas there is no response (or a negative response in the case of CIV) to the IFF. The positive coefficients on the BMF are expected in a model in which good news about the BMF signals lower idiosyncratic risk (higher skewness and lower volatility). Although these models are silent on IFF sign predictions, the zero or negative coefficients we estimate show the IFF has a relationship with these idiosyncratic risk measures distinct from that of the BMF.

Intermediary risk models: Panels J–K show LPIRs when using tradable factors as-

sociated with intermediary risk, including the intermediary risk factor from He et al. (2017) and the common fund flows factor from Dou et al. (2024). In both cases, the BMF generates strong and positive responses in the variables (consistent with negative BMF realizations being bad states for intermediaries). In contrast, the IFF fails to generate statistically significant responses.

Other asset classes: According to Roll’s (1977) critique, one reason the VMF may lead to a rejection of the CAPM is that it does not contain other sources of investors’ wealth held outside the public US equity market. These omitted sources of wealth can drive a wedge between the VMF and the theoretical ideal market index, even in a world in which CAPM assumptions hold. It is therefore informative to understand the comovement of both the BMF and the IFF with other sources of financial wealth, several key components of which are included in Panels L–O. First, Panel L considers an important component of the return on risky debt, showing that (the negative of) credit spreads respond strongly and positively to the BMF but do not respond to the IFF. Panel M shows that the (negative of) bond returns are also positively related to the BMF, but not related to the IFF. The real interest rate is the expected growth in the SDF, so that realized returns on bonds are negatively correlated with innovations in expected SDF growth (approximately, since we use nominal returns here). Thus, positive bond returns reflect good news in terms of marginal utility.⁴⁹ Finally, returns on other asset classes that reflect other components of wealth (such as hedge funds and REITs in Panels N and O, respectively) respond strongly and positively to the BMF, whereas hedge funds show no response to the IFF (with REITs show a long-term negative response to the IFF).

OA.7.2 LPIR estimation

Suppose the BMF or IFF goes up by one standard deviation unit. By how much would an econometrician revise her forecast about future macroeconomic conditions or returns on other asset classes on the basis of that information? To answer this question, we follow the method in Schmidt (2025), which is itself related to Jordà (2005).⁵⁰ While we refer the

⁴⁹Consistent with this interpretation, Lustig et al. (2013) show that bond returns are positively correlated with their measure of returns on wealth

⁵⁰The identifying assumptions proposed in Schmidt (2025) rely on correct specification of the mean of returns and allow the conditional expectation of the macro variable Y_{t+h} to potentially be misspecified. This

reader to those papers for further details, the main result is that if we have a variable Y_t , we can estimate a quantity analogous to an impulse response in a VAR to a shock, z_{t+1} , at horizon $h = 1, \dots, H$ by running the following regression:

$$Y_{t+h} = \alpha_h + \beta_h z_{t+1} + \gamma_h Y_t + \phi_h(L)Y_t + \epsilon_{t+h} \quad (\text{OA.61})$$

where, in our context, z_{t+1} is the market return component of interest (i.e., the BMF or the IFF), $\phi(L)$ is a lag polynomial, and ϵ_{t+h} is the forecast error. The LPIR is simply $\sum_{h=1}^H \beta_h$.

OA.7.3 BMF versus VMF LPIRs

In this section, we compare the LPIRs for the BMF with those for the VMF, which we plot in Figure OA.8 (where both are scaled to have unit standard deviation, similar to results in Figure 5). Both sets of LPIRs are similar. However, those for the BMF tend to be slightly above those for the VMF (although the differences are not statistically significant). This result is expected, and the difference in magnitudes can be explained by a “back of the envelope” calculation. Ignoring control variables described in Online Appendix OA.7.2, one can relate macroeconomic variable loadings on the VMF to those on the BMF approximately as

$$\beta^{VMF} = \beta^{BMF} \frac{\sigma_{BMF}^2}{\sigma_{VMF}^2} + \beta^{IFF} \frac{\sigma_{IFF}^2}{\sigma_{VMF}^2}, \quad (\text{OA.62})$$

where β^i is a macroeconomic variable’s loading on factor i for $i \in \{BMF, VMF, IFF\}$ (note the slight abuse of notation—these betas are not the same as our VMF and BMF betas in the main paper) and σ_i^2 reflect variances. Since we present our main LPIR results as responses with respect to the z -transformed time series, we can also derive the approximate relationship between the betas for these transformed series as

$$\beta_{\sigma_{VMF}}^{VMF} = \beta_{\sigma_{BMF}}^{BMF} \frac{\sigma_{BMF}}{\sigma_{VMF}} + \beta_{\sigma_{IFF}}^{IFF} \frac{\sigma_{IFF}}{\sigma_{VMF}}, \quad (\text{OA.63})$$

is advantageous given that returns are much closer to random walks than many of the macro variables in the consideration set.

where $\beta_{\sigma_i}^i$ reflects the beta with respect to the z-transformed time series i for $i \in \{BMF, VMF, IFF\}$. As we show in our main draft (Figure 5), for many response time series we have $\beta_{\sigma_{IFF}}^{IFF} \approx 0$. Thus, we expect attenuation bias to be the dominant force that causes a wedge between $\beta_{\sigma_{VMF}}^{VMF}$ and $\beta_{\sigma_{BMF}}^{BMF}$ in Figure OA.8. Furthermore, the strength of that bias should be related to the ratio of the BMF's to the VMF's standard deviation. Table 1 implies that this ratio is about 0.9. Thus, we expect LPIRs for the VMF to be approximately 90% of those for the BMF, which is approximately true in the plot (remembering that we simplify the back-of-the-envelope calculation by not considering the effect of control variable or multiple forecast horizons).

What about cases where the LPIR with respect the IFF is negative (for instance, in Panel I for Common Idiosyncratic Volatility)? In these cases, we expect the additive bias term to increase the wedge between BMF and VMF betas even further, which is exactly what we observe. In cases where the IFF beta is positive (albeit never statistically significant, such as in Panel G for Income Risk, Panel K for Common Mutual Fund Flow, or Panel H for Hedge Fund Return), we expect the gap between BMF and IFF betas to narrow, which is exactly what we observe.

So, the direction and size of the gaps between BMF and VMF LPIRs varies as we expect based on the intuition from equation (OA.63), and its magnitude is typically small because the standard deviation of the VMF is similar to that of the BMF. To reiterate, the similarity between the LPIRs with respect to the BMF and VMF is expected. The main message that we want to convey with results in Figure 5 is that the LPIRs for the BMF compared to the IFF display dramatically different behavior as discussed in Section 6.3. In particular, there are strong responses with respect to the BMF (as there are for the VMF, which we show here), but that the component of the VMF that we remove (the IFF) does not generate such strong responses and even generates responses in the “wrong” direction in some cases.

OA.8 The IFF and factors beyond size

Although size is the most salient dimension over which public and all firms differ (based on quantities we can measure for both), other characteristics likely differ between these two sets of firms as well. Thus, other standard characteristics-based pricing factors beyond size

factors may also load on the BMF and IFF as described by equation (9) in our model. It is therefore possible that other factors could also be spanned by the BMF (or, at least to be better priced by the BMF than the VMF), similar to our finding for size factors in Section 8.2. We investigate this possibility in this section.

We begin with an investigation specifically related to value factors. Intuitively, access to public equity markets is likely to be more useful for firms that have valuable growth opportunities. Consistent with this logic, our BMF construction methodology assumes that, within industry-size bins, firms with higher market value-to-sales ratios are more likely to go public, and derives weights to adjust for this selection. If, empirically, the VMF tends to overweight growth-like firms with high multiples, one might wonder whether the BMF can also explain the value anomaly and value factors.

Under such a hypothesis, we would expect that, even conditional on size, growth firms with low book-to-market ratios might load positively on the IFF relative to value firms, resulting in a gap between their VMF and BMF betas. Figure OA.15 shows beta gaps for 25 portfolios double-sorted by size and book to market quintiles (a subset of our 382 value-weighted portfolios described in Table OA.3). Clusters capture different size quintiles, and within-cluster columns capture different within-size-quintile book-to-market ratio bins. Consistent with other results presented in the paper, beta gaps are especially large within the smallest size quintile and shrink as we move across the size distribution towards larger firms. Consistent with value-like firms being underrepresented (similar to small-caps), beta gaps are typically the most negative for the highest book-to-market portfolios within each size quintile (represented by the green bars). The spread across book-to-market sorted portfolios is particularly pronounced within the top two size quintiles where constituents are more likely to be overrepresented in the public market, and we note that VMF betas only exceed BMF betas for the largest quintile of firms in the bottom three book-to-market quintiles. Within the bottom three size quintiles, the beta gap spread across book-to-market quintiles is smaller, and patterns are sometimes non-monotonic, which reflects the fact that low book-to-market firms often have higher market betas (using either BMF or VMF) and therefore are subject to a larger multiplicative attenuation effect, a force which tends to make beta gaps larger (see equations (10) and (11)).

Qualitatively these differences by book-to-market ratios go in the direction of helping to explain the value premium (i.e., the beta gaps for growth firms tend to be smaller and less negative than the beta gaps for value firms, implying value-like firms load more on the BMF than the VMF), but quantitative impacts are smaller relative to the impact of size (i.e., beta gap differences are much larger across size quintiles than across value quintiles). Thus, we expect the BMF to explain some of the value effect, but less so than the size effect.

To further investigate the extent to which the BMF can explain the value effect, we run factor spanning tests for value factors in models that explicitly include such factors (i.e., the FF3, FF5, and FF5C models) when either the VMF or the BMF is used as a market factor in Table OA.7. This analysis is analogous to our size factor spanning tests in Table 5. Table OA.7 shows that, in our full sample, the HML factor in the FF3 model is not spanned regardless of whether we use the VMF or the BMF as the market factor. However, it is spanned in both the FF5 and FF5C models. The latter result is not surprising as Fama and French (2015) note that HML is spanned in their 5-factor model when using a similar sample period. In the early sample period when the value premium was relatively strong (1963–1992), HML is not spanned in any of the tests. Interestingly, in the more recent sample period when the value premium is relatively weak (1993–2021), HML is spanned by the other FF3 model factors whether we use VMF or BMF as the market factor. However, it is not spanned (and has a negative alpha) under the FF5 and FF5C models whether we use VMF or BMF as the market factor. In all cases, spanning test alphas are similar whether we use the VMF or the BMF as the market factor. Thus, BMF-based models performs similarly to VMF-based models when it comes to pricing value factors and (typically) neither are able to span these factors, which contrasts analogous results we show for size factors.

Next, we generalize the above factor spanning tests to explore whether the BMF can price factors associated with other anomalies by running factor spanning tests for each factor in each of the nine models we investigate (see Section 6.2 for a description of these models). We report results in the form of a scatterplot of factor alphas from BMF-based models plotted against factor alphas from VMF-based models in Figure OA.16 Panel A using our full sample of data (1963–2021). Results are similar in the earlier 1963–1992 and later 1993–2021 samples, so we omit these for brevity. The main message from this figure is that, as we

saw for HML factors in Table OA.7, replacing the VMF with the BMF does little to alter factor alphas for any factors other than size factors.

Why are these other factors not spanned by the BMF? Our model implies that factors (such as SMB) constructed by going long firms that are more likely to go public and short firms that are less likely to go public are spanned by the BMF and the IFF and, additionally, that 100% of their variation can be explained by exposure to these factors (see equation (9)). Of course, this exact relationship holds only in the simplified world of the model. In the real world, the composition of public and all firms may differ along many dimensions and we may not expect the BMF and IFF to explain all variation among stocks with exposure to one particular dimension. However, if a factor of interest has a large exposure to the IFF, then its pricing will be drastically impacted depending on whether one uses the VMF or the BMF as a pricing factor in spanning tests.

To investigate this issue, we plot R-squared values from bivariate regressions of each model factor on the BMF and the IFF in OA.16 Panel B. It is clear that all size factors have large exposures to the IFF and, furthermore, are the factors that have the highest total R-squared values (i.e., most of their variation can be explained by exposure to the BMF and the IFF). Other factors have lower total R-squared values and, importantly, a lower amount of their overall variation can be explained by exposure to the IFF. Thus, these factors are not particularly exposed to the IFF and we do not expect material pricing differences for these factors whether we use the VMF or the BMF as the market factor in the related factor models. This is exactly what we observe in Figure OA.16 Panel A.

It is worth noting that the share of total R-squared arising from IFF exposure is relatively high for factors associated with value (for instance, HML from the FF3 and FF5 models, IA from the Q model, and CMA from the FF5 model). However, the exposures for these value-related factors to the IFF are much lower than the comparable exposures of size factors. As highlighted in Table OA.7 and Figure OA.16 Panel A, these smaller exposures are not large enough to materially affect pricing under the VMF versus the BMF.

OA.9 GARCH models: Additional details

We estimate parameters in the GARCH volatility forecasting models by numerical maximum likelihood estimation (MLE). After estimating parameters, we calculate their asymptotic standard errors by approximating the first and second derivatives of the log-likelihood function at maximum likelihood estimates. Reported standard errors are computed using the “outer product” method. Details for numerical MLE estimation and calculation of asymptotic standard errors can be found in Hamilton (Time Series Analysis, 1994, pp. 133–148).

The non-trivial difference between the BMF- and the VMF-based volatility forecasts supports our conjecture that the IFF adds noise to the lagged forecasting errors used to generate conditional variance forecasts when using the VMF. Additional evidence for this conjecture can be found in Byun (2016), who shows that cross-sectional dispersion in the returns of different stocks, a measure of aggregate idiosyncratic risk, does not help forecast volatility of the S&P 500 index when used as an additional explanatory variable in GARCH forecasting models. See Byun (2016) for a detailed explanation and also for an alternative approach through which cross-sectional dispersion measures can improve volatility forecasts indirectly.

OA.9.1 Alternative GARCH models

This section provides more details related to the GJR-GARCH and E-GARCH models. In a study of the relationship between expected excess stock returns and volatility, Glosten et al. (1993) introduced the GJR-GARCH model, which allows for larger feedback from prior squared negative returns relative to positive returns. Denoting by I_t an indicator variable for a positive forecasting error in period t , the conditional variance is given by

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 u_t^2 + \alpha_2 \sigma_t^2 + \alpha_3 I_t u_t^2, \quad (\text{OA.64})$$

where $I_t = 1$ for $u_t > 0$ and 0 otherwise.

Next, Nelson (1991) proposed the EGARCH model, which specifies the conditional variance in logarithmic form as

$$\ln(\sigma_{t+1}^2) = \alpha_0 + \alpha_1 (|\tilde{u}_t| - E|\tilde{u}_t|) + \alpha_2 \ln(\sigma_t^2) + \alpha_3 \tilde{u}_t, \quad (\text{OA.65})$$

where $\tilde{u}_t = u_t/\sigma_t$ is a standardized forecast error in period t , and $E|\tilde{u}_t| = 2/\pi$ for a normally distributed \tilde{u}_t . Note that the above asymmetric GARCH models capture the leverage effect through α_3 : with $\alpha_3 < 0$, the conditional variance (σ_{t+1}^2) is higher for a negative than for a positive forecasting error (u_t).⁵¹

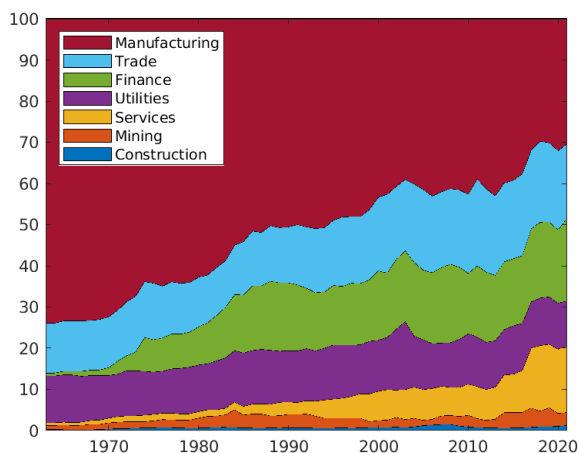
OA.10 Bootstrap procedure for estimating p-values

We use the following bootstrap procedure to estimate p-values whenever indicated in our tables. We begin by assuming that our PCs (in cases where we use PCs to construct alternative BMFs), the BMF, and the IFF are fixed within a given sample of data (i.e., we do not sample from portfolio returns and then recompute the PCA, the BMF, and the IFF within each simulated sample). We then sample the same number of periods of data as in our sample of interest (with replacement) 1,000 times, recomputing the statistic of interest each time. P-values are the number of simulated statistics that satisfy the testing criterion in which we are interested (as a fraction of the total number of simulations).


⁵¹The recognition of the leverage effect goes back to Black (1976), who noted a negative correlation between current returns and future volatility, and is further investigated by Christie (1982). According to the leverage effect, a reduction in the equity value would raise the debt-to-equity ratio, hence raising the riskiness of the firm's equity as manifested by an increase in future volatility.

Figure OA.1: Industry-sorted portfolio sales shares in the aggregate market


Panel A: Sales shares (public)




Panel B: Sales shares (all)



Panel C: Sales shares (public, top 2.5%)




Panel D: Sales shares (all, top 2.5%)




Note: This figure plots the sales shares in the aggregate market for public firms and all firms (public and private) among seven industry portfolios based on industry definitions from Kwon et al. (2024). Panel A plots sales shares over time among public firms. Panel B plots sales shares over time among all firms. Panel C plots sales shares over time among the top 2.5% of public firms by market capitalization using NYSE breakpoints. Panel D plots sales shares over time among the top 2.5% of all firms by market capitalization using NYSE breakpoints. Sales shares for all firms are calculated according to the methodology described in Section 5 and by double-sorting firms into portfolios based on industry then within-industry sales using NYSE breakpoints (i.e., according to our main BMF specification). Data are from 1963–2021. See Sections 3.2 and OA.2 for commentary and Section 5 for additional details on the construction of these measures.

Figure OA.2: Implied value weights of individual stocks in the BMF

Panel A: Distribution of individual stock implied value weights in the BMF




Panel B: Effective sample size ($1/HHI$)



Note. This figure plots information about the distribution of implied value weights on individual stocks over time according to our BMF construction methodology described in Section 5 using equation (13). Implied value weights are based on our main BMF specification that double-sorts stocks into portfolios based on industry then within-industry sales. Panel A plots the mean, 10th percentile, and 90th percentile of implied value weights in the cross-section of stocks each month. Note that the mean across all value weights (blue line) is the same as the weight that would be applied to each stock to construct an equal-weighted portfolio, and that there is considerable spread in the implied value weights around this line. Panel B plots the inverse Herfindahl-Hirschman Index (HHI) for 1) an equal-weighted index, 2) the implied value weights used to construct our BMF, and 3) the actual value weights in the VMF each month in the cross-section of stocks. Note that the inverse HHI (under some assumptions) is equivalent to the effective sample size of firms in an index. Thus, our implied value weights produce a BMF that has an effective sample size that is more similar to that of the VMF than to that of an equal-weighted index of stocks. Data are from 1963–2021. See Section 5 for additional details on the construction of these measures and Section 5.2 for specific commentary.


Figure OA.3: Sales-sorted portfolio shares in the aggregate market (sales-only aggregation)




Note. This figure plots average sales and value shares for 10 sales-sorted portfolios of either public or all firms that use NYSE breakpoints. The red bars plot value shares among portfolios of public firms in the aggregate market for public firms. The blue bars plot sales shares among portfolios of public firms in the aggregate market for public firms. The yellow bars plot sales shares among portfolios of all firms in the aggregate market for all firms. The green line plots the fraction of sales accounted for by public firms in each bin relative to the total sales among all firms in the bin. The all-firm sales shares (in yellow) are calculated using data from Kwon et al. (2024) according to the sales-only single-sorting methodology described in Section 5.3 with more details provided in Online Appendix OA.5.3. Data are from 1963–2021. See Section 3.2 for additional commentary.

Figure OA.4: Industry-sorted portfolio sales shares in the aggregate market (sales-only aggregation)


Panel A: Sales shares (public)




Panel B: Sales shares (all)



Panel C: Sales shares (public, top 2.5%)




Panel D: Sales shares (all, top 2.5%)




Note: This figure plots the sales shares in the aggregate market for public firms and all firms (public and private) among seven industry portfolios based on industry definitions from Kwon et al. (2024). Panel A plots sales shares over time among public firms. Panel B plots sales shares over time among all firms. Panel C plots sales shares over time among the top 2.5% of public firms by market capitalization using NYSE breakpoints. Panel D plots sales shares over time among the top 2.5% of all firms by market capitalization using NYSE breakpoints. Sales shares are calculated according to the methodology described in Section 5 and by single-sorting firms into portfolios based on sales using NYSE breakpoints (i.e., according to our alternative bottom-up BMF specification described in Section 5.3). Data are from 1963–2021. See Section 5 for additional details on the construction of these measures and Section 5.3 for specific commentary.

Figure OA.5: Implied value weights of individual stocks in the BMF (sales-only aggregation)

Panel A: Distribution of individual stock implied value weights in the BMF




Panel B: Effective sample size ($1/\text{HHI}$)




Note. This figure plots information about the distribution of implied value weights on individual stocks over time according to our BMF construction methodology described in Section 5 using equation (13). Implied value weights are based on our alternative BMF specification described in Section 5.3 that single-sorts stocks into portfolios based on sales using NYSE breakpoints. Panel A plots the mean, 10th percentile, and 90th percentile of implied value weights in the cross-section of stocks each month. Note that the mean across all value weights (blue line) is the same as the weight that would be applied to each stock to construct an equal-weighted portfolio, and that there is considerable spread in the implied value weights around this line. Panel B plots the inverse Herfindahl-Hirschman Index (HHI) for 1) an equal-weighted index, 2) the implied value weights used to construct our alternative BMF, and 3) the actual value weights in the VMF each month in the cross-section of stocks. Note that the inverse HHI (under some assumptions) is equivalent to the effective sample size of firms in an index. Thus, our implied value weights produce a BMF that has an effective sample size that is more similar to that of the VMF than to that of an equal-weighted index of stocks. Data are from 1963–2021. See Section 5 for additional details on the construction of these measures and Section 5.3 for specific commentary.

Figure OA.6: Value shares in the aggregate market over time (sales-only aggregation)


Panel A: Sales-sorted portfolios 1 and 10 (public)




Panel B: Sales-sorted portfolios 1 and 10 (all)



Panel C: Industry portfolios (public)




Panel D: Industry portfolios (all)




Note: This figure plots the value shares in the aggregate market for public firms and all firms (public and private) over time. Panels A and B plot the value shares of sales-sorted portfolios 1 and 10 (using NYSE breakpoints) among public and all firms, respectively. Panels C and D plot the value shares by industry among public and all firms, respectively. Value shares for all firms are calculated according to the methodology described in Section 5 and by single-sorting firms into portfolios based on sales using NYSE breakpoints (i.e., according to our alternative BMF specification described in Section 5.3). In particular, stock-level implied value weights are calculated according to equation (13) and then aggregated by sales (Panel B) or industry (Panel D). Data are from 1963–2021. See Section 5 for additional details on the construction of these measures and Section 5.3 for specific commentary.

Figure OA.7: BMF and IFF cumulative return time series

Panel A: BMFs

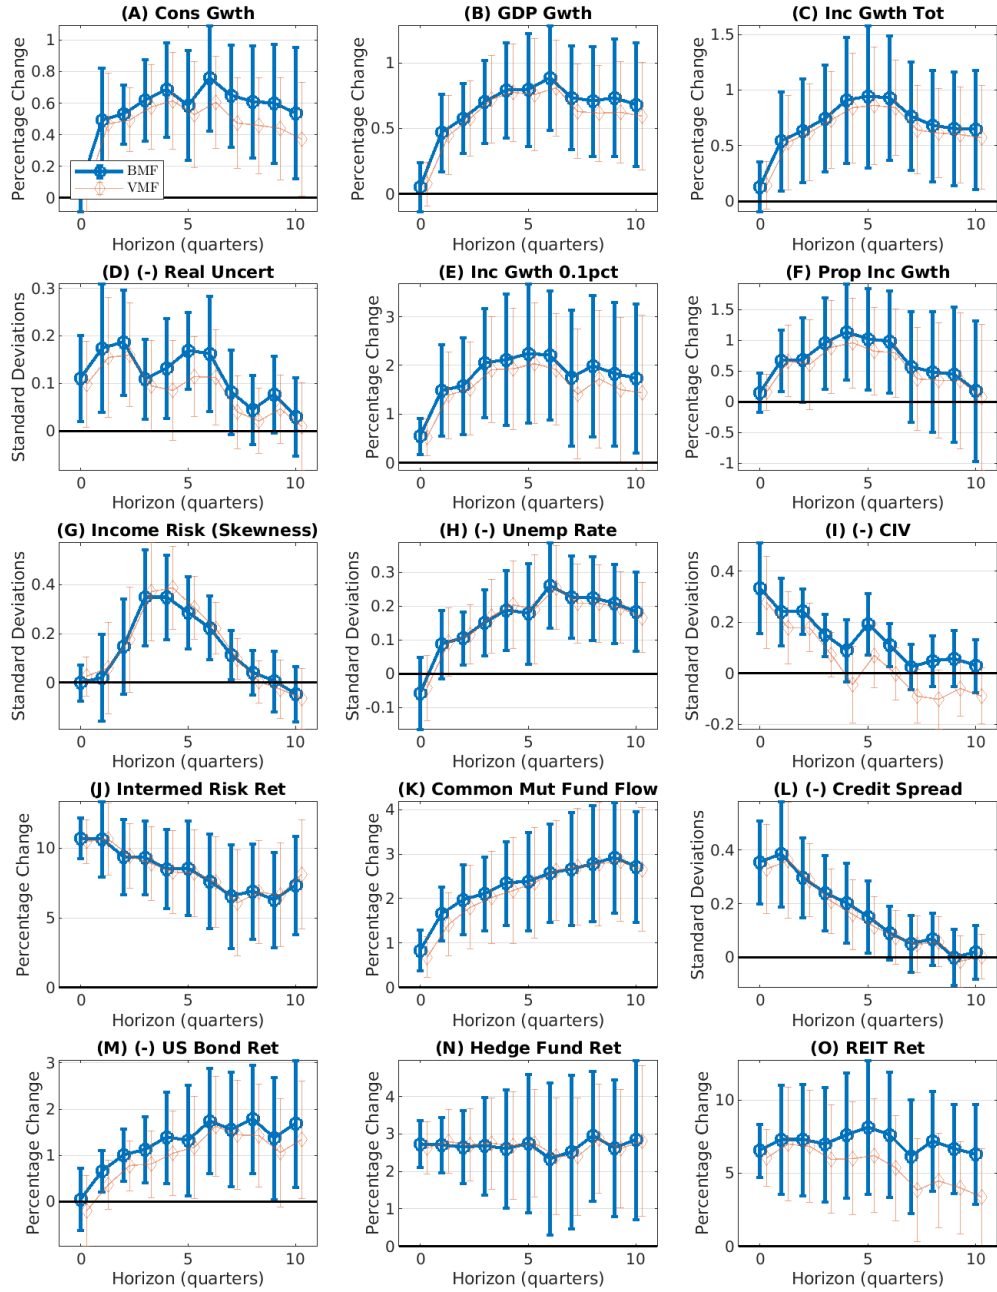


Panel B: IFFs



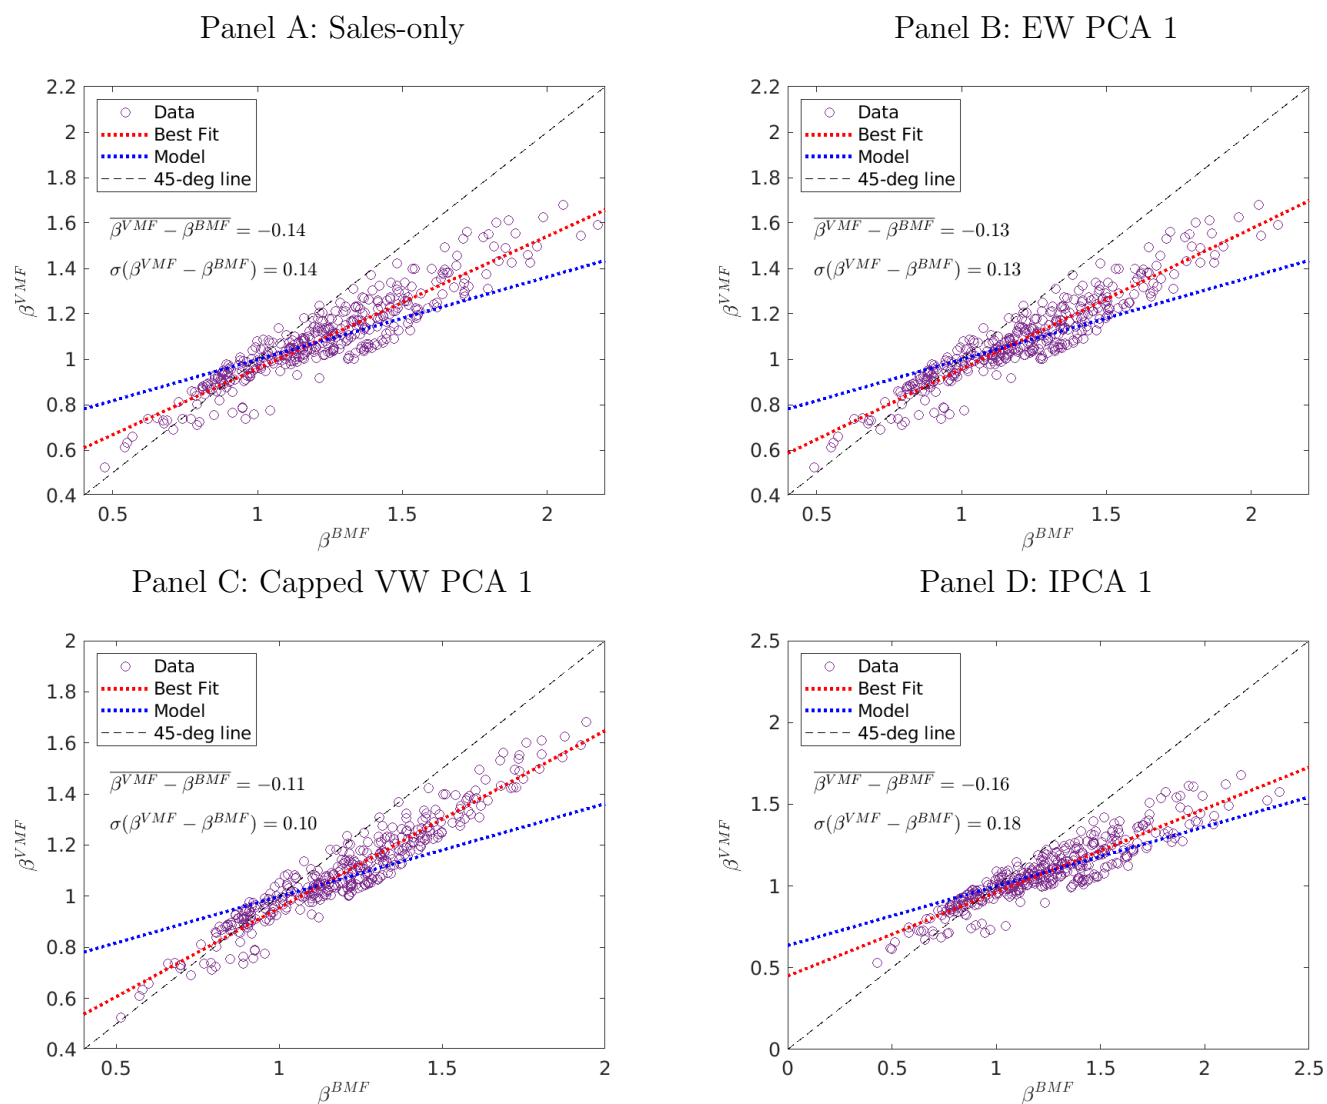
Note. This figure plots cumulative return time series for our main BMF and alternative BMFs considered for robustness as described in Section 5.3 (Panel A). Panel B plots the respective IFF cumulative return time series. Shaded grey regions indicate NBER recession periods. Data are from 1963–2021 (except for the IPCA 1-based BMF and IFF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019); in this case, we fill missing return data with that from our main BMF or IFF for the purposes of plotting). See Section 6.1 for more commentary on BMF and IFF summary statistics and performance.

Figure OA.8: Local projection impulse responses to the BMF and VMF




Note. This figure plots local projection impulse responses of selected macroeconomic aggregates and financial return time series to the BMF and VMF. The indexes are defined in Section 6.3 and are signed so that increases in the variable value correspond to decreases in marginal utility in standard models. A “(-)” in the panel title indicates that the negative value of a time series is used to conform to this logic. Monthly returns are compounded to quarterly frequency to match that of the macroeconomic aggregates. The point estimates represent the impact of a standard deviation increase in the predictor variable on each response variable. Response variables are in units of percent in the case of return or growth time series (e.g., consumption growth, US bond returns, etc.) or units of standard deviation in the case of stationary time series (e.g., unemployment rate, real uncertainty, etc.). The initial point estimates (i.e., when the horizon is equal to zero) correspond to the contemporaneous response of each macroeconomic aggregate to each predictor variable. The error bars indicate 95% confidence intervals computed using Newey-West standard errors with 10 lags. Data are from 1963–2021 (conditional on variable availability). See Section 6.3 and Online Appendix OA.7.3 for additional details.

Figure OA.9: Betas under the VMF versus BMF (robustness)



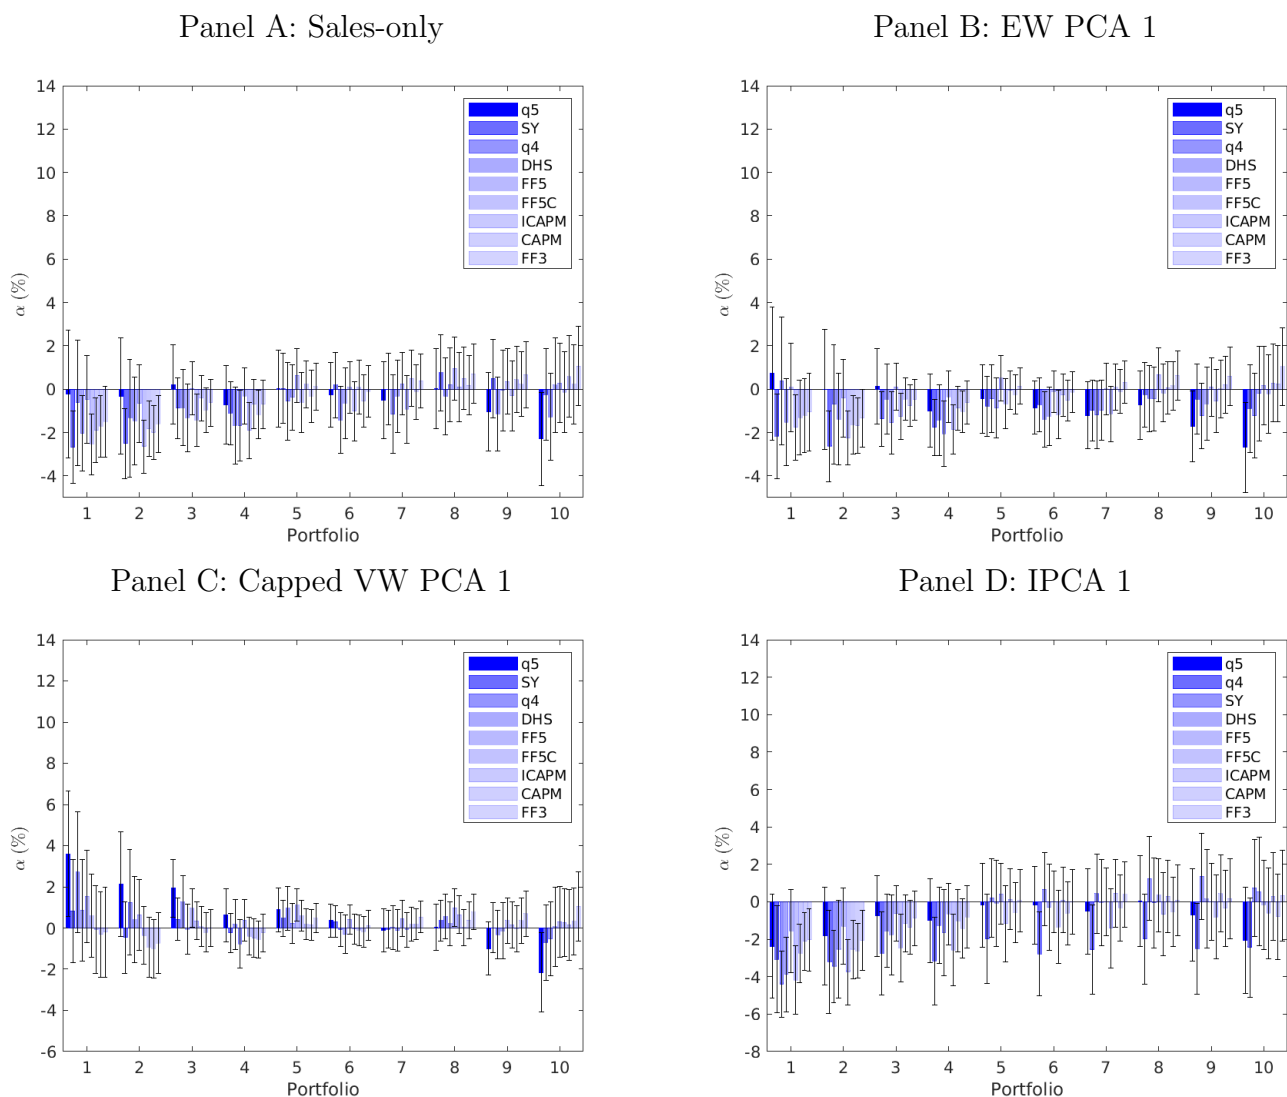
Note. This figure plots univariate VMF betas against univariate BMF betas estimated for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3. The different panels plot results using alternative versions of our main BMF considered for robustness and described in Section 5.3. Dotted lines reflect 45-degree lines to help with the comparison. The dotted red lines are the OLS best-fit lines for the data. The dotted blue lines are the OLS best-fit lines for our theoretical model's implied VMF and BMF betas (based on results reported in Figure 3 Panel A). Data are from 1963–2021 except that related to the IPCA-based BMF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019). Results can be compared to model implications from Figure 3 Panel A, and empirical results based on our main BMF in Figure 6 Panel B. See Section 7.2 for additional details.

Figure OA.10: VMF versus BMF alphas (robustness)



Note. This figure plots VMF alphas against BMF alphas estimated for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3. It also reports their mean absolute values. The different panels plot results using alternative versions of our main BMF considered for robustness and described in Section 5.3. Data are from 1963–2021 except that related to the IPCA-based BMF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019). Results can be compared to model implications from Figure 3 Panel B, and empirical results based on our main BMF in Figure 6 Panel C. See Section 7.1 for additional details.

Figure OA.11: Size-sorted decile portfolio alphas using BMF-based models (robustness)





Note. This figure plots BMF alphas for value-weighted size-sorted decile portfolios (a subset of our set of 382 characteristics-sorted portfolios described in Online Appendix Table OA.3) for each of the nine factor models described in Section 6.2. We drop size factors from the six models with size factors (FF3, FF5, FF5C, q4, SY, and q5) in this analysis because, consistent with our theoretical model and our empirical results, size factors help correctly price size-sorted portfolios when the VMF is used as the market factor, but add no pricing power when the BMF is used. The different panels plot results using alternative versions of our main BMF considered for robustness and described in Section 5.3. The error bars represent 90% confidence intervals on each alpha estimate based on Newey and West (1987, 1994) standard errors. Data are from 1963–2021 (conditional on factor model availability) except that related to the IPCA-based BMF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019). Results can be compared to model implications from Figure 3 Panel C, and empirical results based on our main BMF in Figure 10 Panel B. See Section 8.1 for additional details.

Figure OA.12: Size-sorted decile portfolio alphas (1993–2021)

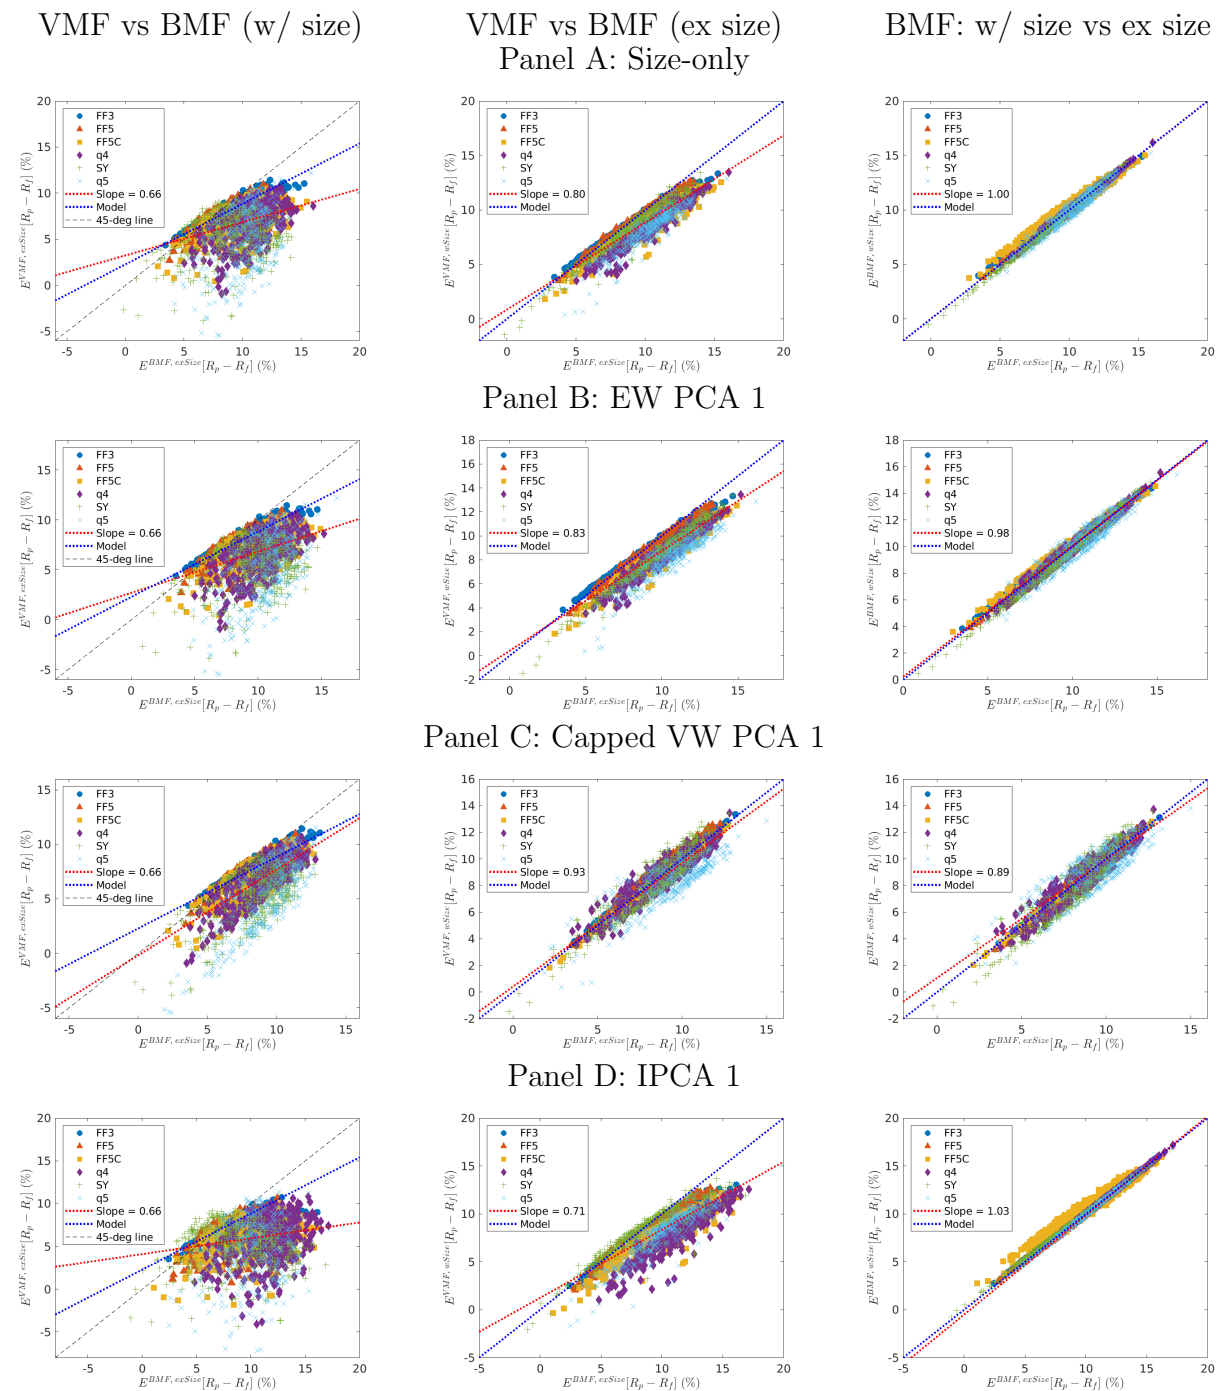
Panel A: VMF-based models

Panel B: BMF-based models



Note. This figure plots alphas for value-weighted size-sorted decile portfolios (a subset of our set of 382 characteristics-sorted portfolios described in Online Appendix Table OA.3) for each of the nine factor models described in Section 6.2 for the 1993–2021 subsample. We drop size factors from the six models with size factors (FF3, FF5, FF5C, q4, SY, and q5) in this analysis because, consistent with our theoretical model and our empirical results, size factors help correctly price size-sorted portfolios when the VMF is used as the market factor (Panel A), but add no pricing power when the BMF is used (Panel B). The error bars represent 90% confidence intervals on each alpha estimate based on Newey and West (1987, 1994) standard errors. Data are from 1993–2021 (conditional on factor model availability). See Section 8.1 for additional details.


Figure OA.13: Size factors and VMF- versus BMF-based model-implied returns (robustness)




Note. This figure plots model-implied expected excess returns for our set of 382 value-weighted characteristics-sorted portfolios described in Online Appendix Table OA.3 for the six standard VMF or BMF as the market factor, and models that either include or exclude the size factors. Comparisons are made between models that use either the VMF or BMF as the market factor, and models that either include or exclude the size factors. The different panels plot results using different versions of our BMF described in Section 5.3. The first column plots results for VMF- versus BMF-based models when size factors are excluded from the models. The middle column plots results for VMF- versus BMF-based models with size factors included in the VMF-based models but excluded from the BMF-based models. The last column plots results for BMF-based models that either include size factors (y axis) or exclude size factors (x axis). The dotted red lines are OLS best-fit lines across all models and portfolios. The dotted blue lines are relationships implied by our theoretical model. For the first column of figures, we estimate this relationship as an OLS best-fit line for the theoretical model's expected returns documented in Figure 3 Panel D. For the middle and last columns, the model implies that the VMF and BMF-based expected returns fall on the 45-degree line, which is what we plot in these cases. Data are from 1963–2021 (conditional on factor availability) except that related to the IPCA-based BMF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019). Results can be compared to model implications from Figure 3 Panel D, and empirical results based on our main BMF in Figure 12 Panels A–C. See Section 8.3 for additional details.

Figure OA.14: VMF and BMF conditional volatility comparison

Panel A: Volatilities

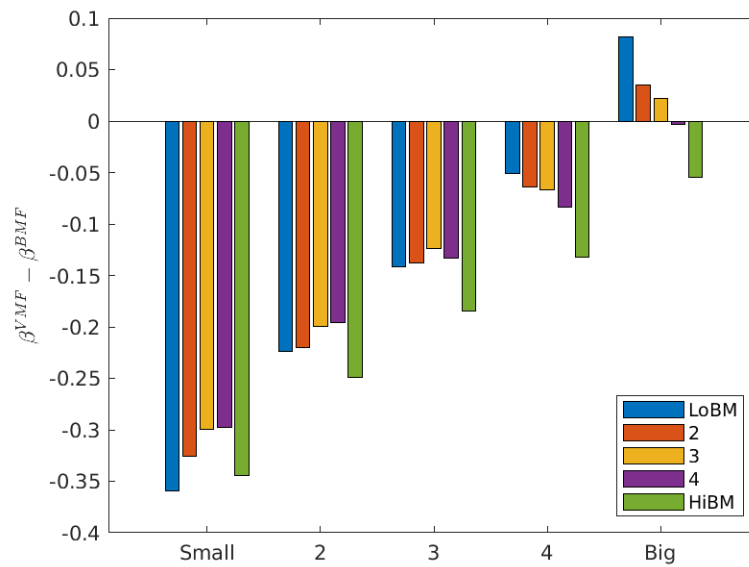


Panel B: Differences in volatilities



Note. This figure compares VMF and BMF conditional annualized volatilities estimated using a GARCH(1,1) model on monthly data. Panel A plots in-sample forecasts for the conditional volatility of each index (on a log scale for easier comparison). The dotted red line represents the conditional volatility forecast estimated for the VMF, whereas the solid blue line represents the forecast obtained for the BMF. Panel B plots the difference between the forecasts in Panel A. Data are from 1963–2021. See Section 9.2 for additional details.

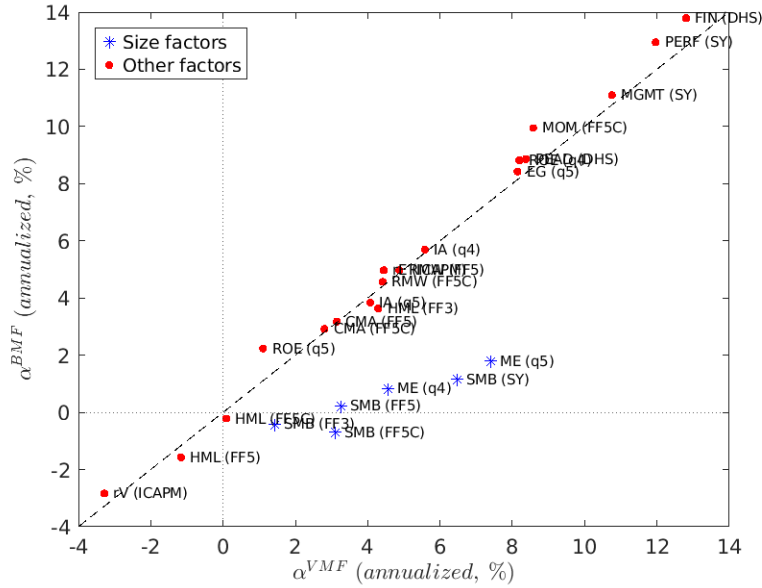
Figure OA.15: Size x book-to-market sorted portfolio beta gaps




Note. This figure plots beta gaps (i.e., the differences in VMF- versus BMF-based CAPM betas, $\beta^{VMF} - \beta^{BMF}$) for value-weighted size x book-to-market sorted portfolios (a subset of our 382 portfolios described in Online Appendix Table OA.3). Portfolios are grouped into size quintile clusters on the x axis. Within the size clusters, portfolios vary in their book-to-market values. Data are from 1963–2021 (conditional on factor availability). See Online Appendix OA.8 for additional details.

Figure OA.16: The IFF and factor spanning

Panel A: Factor spanning test alphas from BMF- versus VMF-based models



Panel B: R-squared values from factors regressed on BMF and IFF



Note. Panel A plots factor spanning test alphas from models that use the BMF as the market factor against those from models that use the VMF as the market factor (i.e., alphas from regressing each factor on other model factors). Alphas for all factors in each model are reported and labeled using the convention “Factor name (model)” for all nine factor models described in Section 6.2. Panel B plots R-squared values from a bivariate regression of each factor on the BMF and the IFF, with contributions from the BMF in blue and contributions from the IFF in red. Data are from 1963–2021 (conditional on factor availability). See Online Appendix OA.8 for additional details.

Table OA.1: Model parameter summary

Parameter	Value	Description
ρ	0.999	Time discount rate (calibrated ex ante)
C_0	0.993	Initial consumption (calibrated given one moment condition and other params)
α	2.946	Agg. abs. risk aversion (match moments)
ω	4.119	Price of risk on z (calibrated given one moment condition and other params)
σ_f	0.020	Std. dev. of f shock (match moments)
σ_g	0.250	Std. dev. of g shock (match moments)
σ_z	0.020	Std. dev. of z shock (calibrated ex ante)
γ_S	2.307	S loading in f (dividends) (match moments)
γ_B	12.988	B loading in f (dividends) (match moments)
θ	0.957	Share of S trees (by number) (calibrated ex ante)
\bar{d}	7.672	Dividend shifter for big versus small firms (match moments)
p_S^{pub}	0.125	Probability of small firm going public (calibrated ex ante)
p_B^{pub}	0.807	Probability of big firm going public (calibrated ex ante)
w_S	0.617	Weight of S-type firms in “small” portfolio (matched moments)
w_B	0.778	Weight of B-type firms in “big” portfolio (matched moments)

Note: This table summarizes all primitive parameters in our theoretical model described in Section 4 and reports their calibrated values. The description also explains (briefly) whether the parameters are calibrated ex ante based on directly observed values in the data or through our over-identified moment matching procedure described in Online Appendix Section OA.3.2.

Table OA.2: Model moment summary

Moment	Data	Model	Simulation
R_f	4.39	4.39	4.39
$\mathbb{E}[R_w]$	14.37	15.77	15.77
$\text{VAR}[R_w]^{1/2}$	19.03	19.78	19.61
$\mathbb{E}[VMF]$	6.79	7.86	7.86
$\text{VAR}[VMF]^{1/2}$	15.36	15.24	15.15
β_{VMF}^w	0.73	0.69	0.69
$\text{VAR}[IFF]^{1/2}$	6.34	6.76	6.69
$\mathbb{E}[R_S]$	14.36	15.69	15.69
$\text{VAR}[R_S]$	20.18	19.64	19.47
$\mathbb{E}[R_B]$	12.07	12.39	12.39
$\text{VAR}[R_B]$	15.17	15.34	15.24
β_S^w	1.03	0.99	0.99
β_S^{VMF}	1.17	1.16	1.16
β_S^{IFF}	0.04	0.02	0.04
β_B^w	0.71	0.70	0.70
β_B^{VMF}	0.97	1.01	1.01
β_B^{IFF}	0.96	0.96	0.97
β_{SMB}^{BMF}	0.44	0.42	0.42
β_{SMB}^{IFF}	-0.93	-0.94	-0.94

Note: This table reports moments from our over-identified system of moments used for calibrating our theoretical model described in Section 4. The “Data” column corresponds to moments we measure in the data. We use \hat{R}_{all} (described in Section 5.1) as an empirical proxy for R_w . See Online Appendix Section OA.3.2 for a description of how we construct the small and big portfolios corresponding to R_S and R_B . The “Model” column corresponds to moments according to model-implied expressions for these moments provided in Online Appendix Section OA.3.2 and given our calibrated parameters reported in Online Appendix Table OA.1. The “Simulation” column corresponds to moment values measured using our simulated data described in Section 4.5.

Table OA.3: Portfolios used in cross-sectional tests

Sorting variables	Start year	No. ports.
Size/Book-to-market	1926	25
Size/Operating profitability	1963	25
Size/Investment	1963	25
Book-to-market/Investment	1963	25
Operating profitability/Investment	1963	25
Size/Momentum	1927	25
Size/Short-term reversal	1926	25
Size/Long-term reversal	1931	25
Size/Accruals	1963	25
Size/Market beta	1963	25
Size/Net share issues	1963	35
Size/Variance	1963	25
Size/Residual variance	1963	25
Industry	1926	17
Size	1926	10
Market beta	1963	10
Book-to-market	1926	10

Note: This table lists the portfolio sets from Kenneth French’s website we use for one of our PCA-based BMF constructions (equal-weighted versions of the portfolios) and our cross-sectional asset pricing tests (value-weighted versions of the portfolios). This provides a total of 382 portfolios over our main sample period, which spans 1963–2021.

Table OA.4: BMF regressions onto PCs and IPCs

	I	II	III	IV	V
Panel A: PCs from EW PCA					
α (%)	0.10 [0.38]	0.12 [0.46]	0.09 [0.33]	0.02 [0.06]	-0.16 [-0.63]
PC1	4.00*** [123.26]	4.00*** [127.75]	4.00*** [125.74]	4.00*** [123.08]	4.00*** [124.77]
PC2		-0.16*** [-5.31]	-0.16*** [-5.18]	-0.16*** [-5.09]	-0.16*** [-5.13]
PC3			0.01 [0.38]	0.01 [0.37]	0.01 [0.37]
PC4				0.12*** [4.69]	0.12*** [4.81]
PC5					0.08** [2.46]
$adjR^2$	0.98	0.98	0.98	0.99	0.99
Panel B: PCs from capped VW PCA					
α (%)	0.58 [1.54]	0.90** [2.57]	0.69* [1.83]	1.24*** [3.29]	1.39*** [4.10]
PC1	3.95*** [99.75]	3.95*** [104.61]	3.95*** [101.23]	3.95*** [106.23]	3.95*** [118.74]
PC2		-0.44*** [-5.27]	-0.44*** [-5.42]	-0.44*** [-7.48]	-0.44*** [-10.04]
PC3			0.09 [1.41]	0.09* [1.84]	0.09** [2.15]
PC4				-0.29*** [-4.36]	-0.29*** [-6.22]
PC5					-0.24*** [-5.21]
$adjR^2$	0.96	0.97	0.97	0.97	0.98
Panel C: IPCs					
α (%)	-0.64 [-1.38]	-0.31 [-0.89]	-0.63* [-1.77]	-0.71** [-1.97]	0.30 [0.55]
IPC1	4.00*** [59.43]	3.08*** [79.95]	1.97*** [47.17]	1.15*** [31.64]	1.04*** [31.41]
IPC2		2.65*** [77.41]	-3.36*** [-100.67]	3.72*** [103.30]	3.69*** [99.84]
IPC3			1.16*** [33.84]	1.08*** [31.54]	-0.97*** [-33.21]
IPC4				0.64*** [13.47]	0.46*** [12.07]
IPC5					0.49*** [10.16]
$adjR^2$	0.95	0.98	0.98	0.98	0.98

Note: This table reports results from regressions of our main BMF (constructed according to the methodology described in Section 5 and by double-sorting firms into portfolios based on industry then within-industry sales using NYSE breakpoints) onto various sets of statistical factors. Panel A reports results from regressing our main BMF onto PCs from a PCA on equal-weighted versions of our set of 382 characteristics-sorted portfolios described in Online Appendix Table OA.3. Panel B reports results from regressing our main BMF onto PCs from a PCA on capped value-weighted versions of our set of 382 characteristics-sorted portfolios described in Online Appendix Table OA.3. We construct capped value-weighted portfolios by assigning stock weights based on market values winsorized at the eightieth NYSE percentile as in Jensen et al. (2023). See footnote 19 for more details. Panel C reports results from regressing our main BMF onto instrumented principal components (IPCs) from Kelly et al. (2019). We use IPCs based on their model that does not include intercepts. Note that each set of factors has its own model, so the factor in the one-factor model is distinct from the first factor in the two-factor model and so forth. See footnote 20 for additional details. All factors are signed to have positive average returns and scaled to have unit standard deviation. T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. PCA data are from 1963–2021 and IPCA data are from 1964–2014.

Table OA.5: IFF spanning tests (robustness)

	CAPM	FF3	FF5	FF5C	q4	SY	DHS	q5	ICAPM
Panel A: VMF as the market factor (ex size factors) [Alternate interpretation: BMF spanning test]									
$\alpha_{Size-ind}^{VMF, exSize}$	-1.67**	-1.10	-2.18***	-3.00***	-3.53***	-4.24***	-3.51***	-5.19***	-2.20***
	[-2.04]	[-1.38]	[-2.73]	[-3.52]	[-3.86]	[-4.51]	[-3.20]	[-5.68]	[-2.76]
$\alpha_{Size-only}^{VMF, exSize}$	-1.89**	-1.29	-2.45***	-3.33***	-3.94***	-4.60***	-4.07***	-5.62***	-2.58***
	[-2.23]	[-1.59]	[-2.90]	[-3.78]	[-4.20]	[-4.70]	[-3.54]	[-5.86]	[-3.16]
$\alpha_{EW PCA1}^{VMF, exSize}$	-1.60**	-1.02	-2.12***	-2.92***	-3.41***	-4.36***	-3.83***	-5.22***	-2.24***
	[-2.00]	[-1.34]	[-2.71]	[-3.69]	[-4.06]	[-5.12]	[-3.68]	[-6.21]	[-3.20]
$\alpha_{Cap VW PCA1}^{VMF, exSize}$	-1.07	-0.48	-1.21*	-1.58**	-1.97***	-2.79***	-2.38***	-3.60***	-1.64**
	[-1.55]	[-0.77]	[-1.87]	[-2.45]	[-2.88]	[-3.97]	[-2.97]	[-5.18]	[-2.41]
$\alpha_{IPCA1}^{VMF, exSize}$	-2.98***	-2.65**	-4.24***	-5.23***	-6.34***	-6.32***	-6.31***	-7.50***	-3.72***
	[-2.93]	[-2.51]	[-4.01]	[-4.51]	[-5.12]	[-5.07]	[-4.06]	[-5.84]	[-4.00]
Panel B: BMF as the market factor (ex size factors) [Alternate interpretation: VMF spanning test]									
$\alpha_{Size-ind}^{BMF, exSize}$	-0.51	0.19	-0.19	-0.86	-1.46	-0.64	-0.44	-2.00**	-0.15
	[-0.65]	[0.24]	[-0.24]	[-1.02]	[-1.62]	[-0.69]	[-0.44]	[-2.13]	[-0.19]
$\alpha_{Size-only}^{BMF, exSize}$	-0.55	0.19	-0.12	-0.83	-1.56*	-0.37	-0.45	-1.89*	-0.16
	[-0.66]	[0.23]	[-0.15]	[-0.93]	[-1.65]	[-0.37]	[-0.42]	[-1.88]	[-0.18]
$\alpha_{EW PCA1}^{BMF, exSize}$	-0.52	0.17	-0.25	-0.92	-1.50*	-1.07	-0.89	-2.32**	-0.46
	[-0.68]	[0.23]	[-0.32]	[-1.14]	[-1.75]	[-1.18]	[-0.88]	[-2.51]	[-0.56]
$\alpha_{Cap VW PCA1}^{BMF, exSize}$	-0.35	0.29	0.04	-0.19	-0.58	-0.54	-0.34	-1.53**	-0.46
	[-0.51]	[0.45]	[0.05]	[-0.29]	[-0.84]	[-0.74]	[-0.44]	[-2.07]	[-0.65]
$\alpha_{IPCA1}^{BMF, exSize}$	-1.31	-0.32	-0.46	-1.28	-2.73**	0.37	-0.18	-1.65	-0.38
	[-1.16]	[-0.27]	[-0.40]	[-1.01]	[-2.00]	[0.28]	[-0.12]	[-1.15]	[-0.32]

Note: Table continued on next page.

Table OA.5: IFF spanning tests (robustness, continued)

	CAPM	FF3	FF5	FF5C	q4	SY	DHS	q5	ICAPM
Panel C: VMF as the market factor (with size factors) [Alternate interpretation: BMF spanning test]									
$\alpha_{Size-ind}^{VMF, wSize}$	-	-0.42	-0.66	-1.55***	-1.61**	-1.24**	-	-2.15***	-
	-	[-1.06]	[-1.32]	[-3.28]	[-2.21]	[-2.13]	-	[-3.12]	-
$\alpha_{Size-only}^{VMF, wSize}$	-	-0.56	-0.79	-1.74***	-1.88***	-1.35**	-	-2.34***	-
	-	[-1.49]	[-1.60]	[-4.02]	[-2.59]	[-2.40]	-	[-3.57]	-
$\alpha_{EW PCA1}^{VMF, wSize}$	-	-0.33	-0.53	-1.41***	-1.43**	-1.26***	-	-2.08***	-
	-	[-0.98]	[-1.27]	[-4.24]	[-2.34]	[-2.93]	-	[-4.12]	-
$\alpha_{Cap VW PCA1}^{VMF, wSize}$	-	0.11	0.19	-0.24	-0.15	-0.03	-	-0.72**	-
	-	[0.37]	[0.72]	[-1.22]	[-0.39]	[-0.08]	-	[-2.48]	-
$\alpha_{IPCA1}^{VMF, wSize}$	-	-1.17**	-1.79***	-2.91***	-3.46***	-2.49***	-	-3.42***	-
	-	[-2.28]	[-2.61]	[-4.05]	[-3.35]	[-3.00]	-	[-3.38]	-
Panel D: BMF as the market factor (with size factors) [Alternate interpretation: VMF spanning test]									
$\alpha_{Size-ind}^{BMF, wSize}$	-	-0.03	-0.06	-1.26**	-1.05	0.00	-	-1.11	-
	-	[-0.07]	[-0.11]	[-2.27]	[-1.23]	[0.00]	-	[-1.38]	-
$\alpha_{Size-only}^{BMF, wSize}$	-	-0.19	-0.20	-1.60***	-1.35	0.01	-	-1.26	-
	-	[-0.38]	[-0.34]	[-3.01]	[-1.40]	[0.02]	-	[-1.38]	-
$\alpha_{EW PCA1}^{BMF, wSize}$	-	-0.07	-0.18	-1.44***	-1.11	-0.58	-	-1.51*	-
	-	[-0.17]	[-0.34]	[-3.42]	[-1.28]	[-1.06]	-	[-1.92]	-
$\alpha_{Cap VW PCA1}^{BMF, wSize}$	-	0.29	0.43	-0.10	0.16	0.52	-	-0.30	-
	-	[0.79]	[1.26]	[-0.41]	[0.31]	[1.24]	-	[-0.63]	-
$\alpha_{IPCA1}^{BMF, wSize}$	-	-0.58	-0.66	-2.56***	-2.83**	0.03	-	-1.50	-
	-	[-0.75]	[-0.82]	[-2.75]	[-2.26]	[0.03]	-	[-1.28]	-

Note: This table reports IFF factor spanning test results (alphas) for each of our nine standard factor models described in Section 6.2. Panel A reports results when the VMF is used as the market factor in each model with size factors excluded from the spanning regressions (for models that have size factors). Panel B reports results when the BMF is used with size factors excluded from the spanning regressions (for models that have size factors). Panel C reports results when the VMF is used as the market factor in each model only for models that include size factors with size factors included in the spanning regressions. Panel D reports results when the BMF is used as the market factor in each model only for models that include size factors with size factors included in the spanning regressions. The first set of rows in each panel reproduces results using our main BMF from Table 2 for comparison. The remaining rows report results for alternative versions of our BMF described in Section 5.3. Note that there are multiple rows associated with VMF-based spanning tests because each alternative BMF implies an alternative IFF. An alternative interpretation of the results in Panels A and C are that they reflect VMF spanning tests, whereas those from Panels B and D reflect BMF spanning tests (with alphas from those tests having the opposite sign but same t-statistics as those reported). Alphas are annualized and in percent. T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. Data are from 1963–2021 (conditional on factor availability) except that related to the IPCA-based BMF, which is limited to 1964–2014 based on data availability from Kelly et al. (2019). See Section 6.2 for additional details.

Table OA.6: Size factor spanning tests (1993–2021)

	Factor model					
	FF3	FF5	FF5C	q4	SY	q5
Panel A: VMF as the market factor						
α^{VMF}	0.25	3.06	2.75	3.44	4.84**	6.35**
	[0.13]	[1.50]	[1.43]	[1.36]	[2.40]	[2.41]
β^{VMF}	0.17***	0.06	0.07	0.08	0.06	0.02
	[4.83]	[1.04]	[1.48]	[1.06]	[1.36]	[0.27]
$adjR^2$	0.10	0.24	0.24	0.14	0.06	0.18
Panel B: BMF as the market factor						
α^{BMF}	-1.72	-0.27	-1.59	-0.40	-0.53	1.21
	[-0.99]	[-0.14]	[-0.91]	[-0.16]	[-0.30]	[0.46]
β^{BMF}	0.40***	0.32***	0.39***	0.37***	0.43***	0.33***
	[9.03]	[6.16]	[8.60]	[5.48]	[7.76]	[4.57]
$adjR^2$	0.29	0.35	0.38	0.26	0.24	0.27

Note: This table reports size factor spanning test results using the six standard factor models described in Section 6.2 that include size factors (i.e., results from regressing size factors on other model factors) for the 1993–2021 subsample. Panel A reports results when the VMF is used as the market factor in each model. Panel B reports results when the BMF is used. α^{VMF} (α^{BMF}) is the alpha when the VMF (BMF) is used (annualized and in percent). β^{VMF} (β^{BMF}) is the size factor beta with respect to the VMF (BMF). *Size factor alphas under the VMF are similar in magnitude to those reported for the 1963–2021 sample in Table 5, though they are statistically insignificant in all but the SY and q5 models. This is likely due to the shorter sample period as opposed to indicating size factors are unimportant for pricing during this sample since their VMF alphas are similar to those from the 1963–2021 period reported in Table 5.* T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. Data are from 1993–2021 (conditional on factor availability). See Section 8.2 for additional details.

Table OA.7: HML factor spanning tests

	Factor model (1963–2021)			Factor model (1963–1992)			Factor model (1993–2021)		
	FF3	FF5	FF5C	FF3	FF5	FF5C	FF3	FF5	FF5C
Panel A: VMF as the market factor									
α^{VMF}	4.29**	-0.82	0.39	5.99***	2.95**	3.62***	1.88	-4.04**	-2.89
	[2.53]	[-0.60]	[0.31]	[3.86]	[2.34]	[2.73]	[0.65]	[-2.09]	[-1.49]
β^{VMF}	-0.11**	0.06	0.03	-0.22***	-0.06**	-0.06*	-0.02	0.18***	0.13**
	[-2.36]	[1.34]	[0.86]	[-5.26]	[-2.07]	[-1.95]	[-0.23]	[3.11]	[2.35]
$adjR^2$	0.05	0.46	0.50	0.13	0.58	0.59	0.05	0.48	0.51
Panel B: BMF as the market factor									
α^{BMF}	3.63**	-1.45	-0.26	5.90***	2.71**	3.45**	0.64	-4.88***	-3.97**
	[2.11]	[-1.10]	[-0.21]	[3.52]	[2.12]	[2.57]	[0.23]	[-2.79]	[-2.11]
β^{BMF}	0.00	0.15***	0.10**	-0.20***	-0.04	-0.04	0.14	0.30***	0.24***
	[0.02]	[2.84]	[2.31]	[-3.06]	[-0.93]	[-0.99]	[1.54]	[5.01]	[3.68]
$adjR^2$	0.03	0.49	0.51	0.06	0.58	0.59	0.07	0.52	0.53

Note: This table reports HML factor spanning test results using the three standard factor models described in Section 6.2 that include HML factors (i.e., results from regressing HML factors on other model factors). Panel A reports results when the VMF is used as the market factor in each model. Panel B reports results when the BMF is used. The leftmost columns report results for the 1963–2021 sample, the middle columns report results for the 1963–1992 subsample, and the rightmost columns report results for the 1993–2021 subsample. α^{VMF} (α^{BMF}) is the alpha when the VMF (BMF) is used (annualized and in percent). β^{VMF} (β^{BMF}) is the HML factor beta with respect to the VMF (BMF). T-statistics are in square brackets and are estimated according to Newey and West (1987, 1994). */**/** represent statistical significance at 90%, 95%, and 99%, respectively. See Online Appendix OA.8 for additional details.