DISCUSSION PAPER SERIES

DP16109

Have Risk Premia Vanished?

Simon Smith and Allan Timmermann

FINANCIAL ECONOMICS

ISSN 0265-8003

Have Risk Premia Vanished?

Simon Smith and Allan Timmermann

Discussion Paper DP16109
Published 04 May 2021
Submitted 02 May 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Have Risk Premia Vanished?

Abstract

We apply a new methodology for identifying pervasive and discrete changes (``breaks'') in cross-sectional risk premia and find empirical evidence that these are economically important for understanding returns on US stocks. Size and value risk premia have fallen off to the point where they are insignificantly different from zero at the end of the sample. The market risk premium has also declined systematically over time but remains significant and positive as does the momentum risk premium. We construct a new instability risk factor from cross-sectional differences in individual stocks' exposure to time-varying risk premia and show that this factor earns a premium comparable to that of commonly used risk factors. Using industry- and characteristics-sorted portfolios, we show that some breaks to the return premium process are broad-based, affecting all stocks regardless of industry- or firm characteristics, while others are limited to stocks with specific style characteristics. Moreover, we identify distinct lead-lag patterns in how breaks to the risk premium process impact stocks in different industries and with different style characteristics.

JEL Classification: N/A

Keywords: N/A

Simon Smith - simon.c.smith@frb.gov
Federal Reserve Board

Allan Timmermann - atimmermann@ucsd.edu
UCSD and CEPR

Acknowledgements
We are grateful for suggestions on an early version of this paper from Bryan Kelly, Dacheng Xiu, and seminar participants at the 2018 SoFiE Summer School at University of Chicago and the March 2019 UK Inquire Group meeting in Windsor. We are also grateful for comments from Gurdip Bakshi and Pierre Perron as well as conference and seminar participants at the North American Winter Meeting of the Econometric Society 2021, Boston University, and Temple University. Linyan Zhu provided excellent research assistance. Any remaining errors are our own.

Have Risk Premia Vanished?

Simon C. Smith[a], Allan Timmermann[b]

[a]*Federal Reserve Board*
[b]*University of California, San Diego*

May 2, 2021

Abstract

We apply a new methodology for identifying pervasive and discrete changes ("breaks") in cross-sectional risk premia and find empirical evidence that these are economically important for understanding returns on US stocks. Size and value risk premia have fallen off to the point where they are insignificantly different from zero at the end of the sample. The market risk premium has also declined systematically over time but remains significant and positive as does the momentum risk premium. We construct a new instability risk factor from cross-sectional differences in individual stocks' exposure to time-varying risk premia and show that this factor earns a premium comparable to that of commonly used risk factors. Using industry- and characteristics-sorted portfolios, we show that some breaks to the return premium process are broad-based, affecting all stocks regardless of industry- or firm characteristics, while others are limited to stocks with specific style characteristics. Moreover, we identify distinct lead-lag patterns in how breaks to the risk premium process impact stocks in different industries and with different style characteristics.

Keywords: Cross-sectional variation in risk premia, instability risk factor, industry and style portfolios, Bayesian analysis

JEL classifications: G10, C11, C15

Email address: `atimmermann@ucsd.edu`. `Phone: (858) 5340894. Fax: (858) 5340745. Address: Rady School of Management, Otterson Hall, Room 4S146, 9500 Gilman Drive 0553, La Jolla, CA 92093-0553, USA. Corresponding author.` (Allan Timmermann)

We are grateful for suggestions on an early version of this paper from Bryan Kelly, Dacheng Xiu, and seminar participants at the 2018 SoFiE Summer School at University of Chicago and the March 2019 UK Inquire Group meeting in Windsor. We are also grateful for comments from Gurdip Bakshi as well as conference and seminar participants at the North American Winter Meeting of the Econometric Society 2021, Boston University, and Temple University. Linyan Zhu provided excellent research assistance. Any remaining errors are our own.

1. Introduction

Equity risk premia play a key role for investment strategies in the stock market. Empirical findings that stock characteristics such as book-to-market value, market cap, and return momentum are associated with sizeable risk premia have profoundly impacted the investment industry with countless mutual funds specializing in investment styles such as small caps, growth, value, or momentum stocks.[1] The attractiveness of such investment strategies hinges critically not only on the magnitude of the associated risk premia, but also on their stability over time. For example, high allocations to value or small-cap stocks will be notably less attractive if the risk premia associated with these types of stocks have been significantly reduced over time. Shifts in risk premia also introduce an additional source of risk for investors – particularly if their impact varies across industries and firm characteristics.

Recent empirical evidence suggests that cross-sectional risk premia associated with a broad array of firm-level characteristics vary considerably through time, reaching unusually high levels of volatility during economic crises and periods with elevated distress in financial markets.[2] Moreover, such instability is mirrored across a broad range of asset classes and investment styles: using a century of data on six asset classes, Ilmanen et al. (2019) find considerable evidence of time variation in single-factor returns and volatility for value, momentum, carry, and defensive investment strategies.

Recognizing the need to formally test for shifts in risk premia, Fama and French (2021) report evidence of a substantial decline in the value risk premium but are unable to reject the null hypotheses that the value premium (i) is constant across pre- and post-1992 subsamples and (ii) is zero in the post-1992 subsample. However, their test uses just a handful of portfolios and likely has low power given the inherent noise in monthly premia. Moreover, the use of portfolios may mask the risk-return tradeoff in underlying stocks (Lewellen et al. 2010). Finally, they do not consider if their break date (1992) is the break location supported

[1] Several studies have found that firm characteristics are priced, e.g., Fama and French (1993); Berk et al. (1999); Carlson et al. (2004); Zhang (2005); Carlson et al. (2006); Novy-Marx (2013).

[2] Freyberger et al. (2020) and Gu et al. (2020) document significant time variation in the mapping from a variety of firm-level predictors to expected returns. Gagliardini et al. (2016) find that risk premia are large and volatile in crisis periods and deviate considerably from the path implied by a constant-parameter model. Ang and Kristensen (2012) use a nonparametric approach to estimate and track time variation in the factor loadings of conditional CAPM or multi-factor models. Adrian et al. (2015) propose regression-based estimators of dynamic asset pricing models that capture time-variation in beta loadings and risk premia.

by the data or if there are more than a single break.

In this paper, we propose a novel approach to test for and model instability in risk premia which exploits information in large cross-sections of individual stock returns. Our approach is very flexible and does not assume that the dates of any breaks or even the number of breaks is known in advance. Using cross-sectional information turns out to be key to our ability to accurately estimate the location and magnitude of shifts to risk premia.[3] In turn, more accurate estimates of risk premia enhance our ability to test hypotheses such as constant risk premia, zero risk premia at the end of the sample, or even a monotonically declining pattern in risk premia, allowing us to sharpen the conclusions about risk premia in Fama and French (2021).

Using monthly returns data on a sample of more than 23,000 stocks from 1950 to 2018, we find strong evidence of four breaks in a four-factor model that allows for breaks in the intercept ("alpha"), risk premium coefficients, and idiosyncratic volatility. The break dates are located at July 1972, October 1981, June 2001, and October 2008, thus coinciding with the oil price shocks of the early seventies, the change in the Fed's monetary policy regime, the crash of the Tech bubble, and the Global Financial Crisis (GFC).

We find that the equity risk premium, value premium, and size premium all vary significantly over time and have declined systematically over the nearly seven decades covered by our sample, with particularly large declines observed for the size and value premia. Conversely, after an initial decline in the early seventies, the momentum risk premium has recovered and is back to a level close to its value in the 1950s. Tests conducted on the final (post-GFC) regime do not reject the null hypothesis that the size and value risk premia have fallen to zero. We also cannot reject the null that these risk premia have declined monotonically over the last seven decades. Conversely, we strongly reject that the market and momentum risk premia are zero in the last regime or that they have declined uniformly over time. Our empirical evidence suggests that all four breaks are broad-based and affect both the risk premium coefficients as well as individual stock alphas and idiosyncratic volatility parameters.[4]

[3]Data on individual stocks for improved estimation of risk premia has also recently been used to deal with the errors-in-variables bias by Jegadeesh et al. (2019).

[4]Evidence of mispricing is much stronger during the early part of our sample, declining significantly after 2001. It is also far greater for the less liquid microcaps compared to larger stocks.

We next explore the cross-sectional asset pricing implications of instability in the risk premium process. Stocks with different (style) characteristics have different exposures to variation in risk premia and should, therefore, also vary in how vulnerable they are to instability in risk premia. Stocks with greater exposure to instability risk should therefore earn a greater instability risk premium provided that instability risk is priced in the cross-section.

To see if this prediction holds, we construct a break risk factor using the difference between forecasts of individual stock returns from models with and without breaks. We use this break risk factor to explore whether individual stocks with the largest sensitivity to the break risk factor earn higher returns than stocks with lower break sensitivity. We find that returns on break sensitivity-sorted portfolios increase monotonically with the high-sensitivity quintile of stocks earning a statistically significant 3.4% higher annual return than the low-sensitivity quintile of stocks. Similarly, Fama-MacBeth regressions that control for other stock characteristics such as size, value, and prior return performance, show that the break characteristic obtains a similar level of significance as the book-to-market ratio and in turn is more significant than both size and momentum.

To better understand the portfolio implications of instability in risk premia, we next explore which types of stock characteristics – e.g., industry and investment style – are associated with high exposure to instability risk. To this end, we use industry and characteristics-sorted portfolios to dissect differences in break sensitivity. Across industries, we find that Telecommunication, Utility, Oil, Business Equipment, and Financial stocks exhibit the greatest break sensitivity. Conversely, stocks in the Wholesale, Textile, Mining, Books, and Meals industries exhibit the smallest break sensitivity. Small firms' returns are more sensitive to breaks while big firms are the least sensitive. Conditional on size, value firms are more sensitive to breaks than growth firms and loser stocks are more sensitive than winner stocks.

Next, we explore the economic drivers of breaks by generalizing the common break framework introduced by Smith and Timmermann (2021) to allow breaks to be noncommon, possibly hitting any subset of series in the cross-section at different times. This analysis, which uses the methodology developed by Smith (2018a), allows us to (i) differentiate between market-wide and style-specific breaks; and (ii) identify whether certain industry or characteristic-sorted portfolios are affected earlier or later in the breakpoint cycle. Some breaks (e.g., 1973, and 2008) are very broad and affect stocks across multiple industries and investment styles. Other breaks are more specific to individual styles or industries and so do not have the same broad-based effects.

Inspecting the speed at which different portfolios are affected by breaks, we find that Financials, Telecommunication, Retail, Services, Steel, Chemicals, Oil, and Construction are generally among the first industries to be affected by breaks to risk premia. Moreover, the lead-lag relation varies across breaks with Financials playing a leading role during the 1929 market crash and Global Financial Crisis, while Telecommunication stocks were leading during the dotcom crash and Oil stocks were leading in 1973. The speed of information diffusion has increased over time as the lead-lag delay time between the first and last affected industries has clearly been reduced.

Style portfolios also differ in how rapidly they are affected by breaks: momentum portfolios are generally among the earliest to be affected with loser stocks leading winner stocks. Similarly, large stocks tend to be affected earlier by breaks than small stocks and growth stocks generally move earlier than value stocks which are among the last ones to be affected by breaks.

Finally, we conduct an out-of-sample analysis which shows that our panel break model can be used to generate more accurate return forecasts than alternative constant-parameter and time-varying parameter benchmarks. When these forecasts are used by a moderately risk averse mean-variance investor to form portfolios, this leads to a rotation out of industry portfolios that are hit early in the breakpoint cycle, such as oil after 1973, telecommunications after 2001, and financials after 2008, and results in gains in annual certainty equivalent returns around two percent. Over the seven decades covered by our sample, the major shift we identify in risk premia is associated with a substantial decline in the optimal allocation to small caps and value stocks.

The outline of the paper is as follows. Section 2 introduces our methodology, including the return regressions and prior specifications. Section 3 presents our data and empirical evidence of breaks. Section 4 constructs our cross-sectional break risk factor and compares it with existing risk factors from the finance literature. Section 5 focuses on the timing and effect of breaks in return regressions conducted for different portfolios of stocks sorted by industry or investment style. Section 6 studies out-of-sample return predictability and economic gains from exploiting it, while Section 7 concludes.

2. Methodology

This section introduces our panel regression approach to modeling discrete and pervasive shifts in the risk premium process. We justify our assumption of discrete, pervasive shifts or "breaks" in return premia in three ways. First, a key feature of our approach is that it allows us to identify economically large and long-lasting regime shifts as opposed to smaller and more local variation in risk premia. Focusing on breaks that are pervasive allows us to fully exploit the rich information available in the cross-section of stock returns. Second, and consistent with the idea of discrete shifts in risk premia, the changes that we identify empirically are associated with important economic events and coincide with large shifts in aggregate valuation measures such as the dividend-price ratio of the market portfolio. From an asset pricing perspective, large movements in valuation ratios is exactly what one would expect when risk premia shift. Third, we use economically motivated priors to ensure that the variation in risk premia falls within ranges that are economically plausible.[5]

One could estimate risk premia using a two-stage approach in which firm-specific betas are estimated in the first step while risk premia are estimated in the second step. This, however, introduces an error-in-variables problem and increases parameter estimation error. As an alternative, we therefore follow recent studies and estimate risk premia directly in a single step from regressions of firms' stock returns on a set of stock or firm characteristics.

We next explain the details of our Bayesian panel break approach which builds on the framework of Fama and French (2020) who demonstrate that stacking Fama-Macbeth regressions across time gives rise to a factor model representation.[6] We generalize this framework, however, to allow for structural breaks that capture factor risk premia which

[5] Pástor and Stambaugh (2001) identify breaks in the equity premium process and use transition regimes to link adjacent regimes.

[6] A related literature finds evidence of breaks in expected equity returns. For example, Pástor and Stambaugh (2001) find 15 structural breaks in estimates of the U.S. equity premium from a data set spanning approximately 150 years. Bekaert et al. (2002) identify common breaks in return models and link them to global equity market integration. van Binsbergen et al. (2020) identify a structural break in the magnitude of systematic risk around 2000, after which U.S. crash risk declines against the global average.

undergo occasional shifts at unknown times.

2.1. *Estimating time-varying risk premia*

Suppose we observe a panel of monthly stock returns r_{it}, measured in excess of a risk-free rate, on $i = 1, \ldots, N_t$ firms over a sample $t = 1, \ldots, T$.[7] Moreover, let X_{it-1} denote a vector of firm or stock characteristics for firm i observed at time $t-1$. Characteristics could include observable features such as firm size and book-to-market ratio or estimated stock characteristics such as (factor) betas or return momentum.

Fama and French (2020) demonstrate that, when stacked across t, cross-sectional regressions of returns on lagged firm characteristics become factor models that can be estimated using time-series information. Building on this insight, consider the regression model

$$r_{it} = \alpha_i + r_{zt} + \lambda_t' X_{it-1} + \epsilon_{it}, \qquad \epsilon_{it} \sim N(0, \sigma_i^2). \tag{1}$$

From Fama and French (2020), the slope estimates λ_t are portfolio returns that can be interpreted as factors with pre-specified time-varying factor loadings (characteristics) and r_{zt} is the month-t return on a regular portfolio comprising the left-hand-side assets with weights summing to one when all explanatory variables are set to zero. This return component is therefore common to all stocks.[8] Finally, α_i captures any mispricing of asset i.[9]

The model in Equation (1) and conventional time series factor models both attempt to explain variation in returns. However, there are also important differences between the two approaches. The time series approach uses factors that are prespecified, e.g., from sorts of stocks on book-to-market equity, size, or prior returns and optimizes over the factor loadings which are assumed to be time-invariant. Conversely, estimates of Equation (1) optimize over the common return component (r_{zt}) and the factor returns λ_t so as to minimize the sum of squared residuals given the prespecified time-varying factor loadings.

[7] Our panel approach can easily accommodate variation in the number of stocks at time t, N_t.

[8] To obtain this component, we employ the common correlated effects framework of Pesaran (2006), effectively extracting r_{zt} from the cross-sectional average return.

[9] In a model without the intercept, α_i, Fama and French (2020) note that the time series average of ϵ_{it} will capture mispricing of asset i. To enable us to capture shifts in mispricing, we explicitly include α_i and impose that ϵ has mean zero.

The time-series average return on a factor is often used to estimate it's risk premium. For example, the historical mean (excess) return on the market portfolio is commonly used as an estimate of the equity risk premium. However, if risk premia remain constant within certain blocks of time ("regimes") but can shift across regimes, then risk premia should be computed only on the data from the same regimes.[10] To capture possible time variation in risk premia, we therefore generalize the model in Equation (1) to allow any subset of the factor risk premia, mispricing parameters (alphas), and volatilities to shift an unknown number of times (K) at unknown locations $\tau = (\tau_1, \ldots, \tau_K)$ which give rise to $K+1$ separate regimes.

We initially assume that the breaks are common and affect all assets at the same time, but subsequently relax this assumption. The assumption that breaks to return premia have a pervasive effect on the cross-section of stock returns effectively allows us to use the full cross-section of returns to identify breaks in the risk premium process, vastly increasing the power of our approach. It also ensures that we only identify breaks to the risk premium process that are truly common.

Our panel break model for stock returns thus takes the following form:

$$r_{it} = \alpha_{ik} + r_{zt} + \lambda_k' X_{it-1} + \epsilon_{it}, \qquad \epsilon_{it} \sim N(0, \sigma_{ik}^2), \qquad t = \tau_{k-1}+1, \ldots, \tau_k. \tag{2}$$

Here λ_k denotes the risk premia and α_{ik} the degree of mispricing of asset i in state ("regime") k. Our baseline model uses four lagged characteristics (X_{it-1}) – market beta, size, value, and momentum – and estimates variation in the associated risk premia across regimes. Given our large cross-section of stocks, estimating a full covariance matrix in each regime is not possible. We therefore adopt a common factor structure to absorb dependence across stocks and assume that the remaining residuals in Equation (2) are uncorrelated (Pesaran 2006). While this may seem a strong assumption, the model that we take to the data includes five common factors (r_{zt}, λ_k') which, in practice, should be sufficient to absorb most of the common variation in returns. For instance, Bai and Ng (2002) estimate that just two factors is sufficient to capture variation in the cross-section of U.S. stock returns.[11] In fact, as we

[10] Pástor and Stambaugh (2001) estimate time variation in the U.S. equity risk premium as the average of market excess returns within regimes that are separated by structural breaks.

[11] Stock and Watson (2002) report that just six factors are required to capture most of the variation in

demonstrate below, the assumption of uncorrelated residuals – conditional on including five factors – is supported by the data. Moreover, even if the residuals were significantly correlated, we could further include latent factors to absorb additional correlation.

To capture changes to cross-sectional risk premia in Equation (2), we use a Bayesian panel break methodology that accounts for uncertainty about breaks.[12] Our approach builds on and extends that of Smith and Timmermann (2021) who examine how breaks in the coefficients of a single predictor affect time-series predictability of returns. Conversely, building on Fama and French (2020) our analysis here identifies breaks to pooled cross-sectional risk premia that load on firm-specific characteristics.

2.2. Prior distributions

Before continuing with the analysis, we next explain our choice of priors which follows conventional practice and specifies Gaussian distributions over the slope coefficients and conjugate inverse gamma priors over the residual variances.[13]

The choice of priors should be guided by asset pricing theory and reflect what is economically plausible in terms of the magnitude of any deviations from the underlying factor pricing model. Throughout our paper, benchmark returns are either excess returns or returns on zero-investment (long-short) portfolios. In this case, conventional asset pricing models imply that $\alpha_k = (\alpha_{1k}, \ldots, \alpha_{Nk}) = 0_N$ in the kth regime (Huberman et al. 1987). Centering α_k a priori at zero, the specification of σ_α reflects the prior belief that the pricing model holds. Setting $\sigma_\alpha = 0$ corresponds to a dogmatic belief that the pricing model holds with absolutely no mispricing. Conversely, setting $\sigma_\alpha = \infty$ reflects a prior belief that any degree of mispricing is equally likely. Small values of σ_α reflect prior beliefs that are skeptical about the existence of mispricing but do not rule it out entirely; larger values reflect stronger prior beliefs that there may be some mispricing.

Further, we choose our prior to ensure that an economically unreasonable high Sharpe ratio is unlikely since this would give rise to an approximate arbitrage opportunity by

215 different financial and macroeconomic time series.

[12]Frequentist approaches, such as Bai and Perron (1998) and Baltagi et al. (2016), ignore break uncertainty and may therefore compromise small-sample inference; see Pástor and Stambaugh (2001).

[13]For a detailed description of the prior choices, see Appendix B.

generating high expected returns without being exposed to much risk (Shanken 1992).[14] This scenario could arise if a high intercept estimate, α_{ik}, coincides with a low idiosyncratic volatility, σ_{ik}. Our prior places very little weight on this scenario by linking the intercept to the residual volatility; see MacKinlay (1995); Pástor and Stambaugh (1999); Pástor (2000).[15] Following Pástor and Stambaugh (1999), our baseline analysis adopts a moderate prior belief by setting σ_α equal to 5%. We apply the same prior belief that the α_{ik} values are centered at zero across all regimes, i.e., that the degree of mispricing is constant. This does not rule out that some assets may be more mispriced in one regime and less mispriced in another because residual volatilities are allowed to vary across regimes.

Finally, our prior assumes that breaks occur, on average, every twenty years. The prior on the slope coefficients λ_k is Gaussian. The prior hyperparameter σ_λ controls the degree of shrinkage applied: the smaller this hyperparameter, the more the slopes get pulled toward zero. We specify a moderate degree of shrinkage by setting σ_λ equal to 0.08 (Wachter and Warusawitharana 2009).

3. Instability in risk premia

This section introduces our returns data and presents empirical evidence on the presence of pervasive breaks to the risk premia of the standard Fama-French factors and momentum. We also explore evidence of shifts in the mispricing parameters (alphas) and in the return volatility parameters.

[14] Dybvig (1983) and Grinblatt and Titman (1983) use residual variances to study how much any given asset can depart from a factor model. Shleifer and Vishny (1997) argue that high volatility can introduce limits to arbitrage and thus cause a given asset to be mispriced.

[15] The Gaussian prior on the intercept is conditional on the residual volatility and thus the variance of the intercept combines the residual variance and the prior variance σ_α^2. Since the prior on α_{ik} is centered at zero, a low residual variance will shrink the intercept estimate towards zero, making a value far from zero highly unlikely. As the residual variance increases, the intercept is pulled less strongly toward zero and thus intercept estimates further away from zero become more likely.

3.1. Data

We use monthly data on a total of $N = 23,664$ stocks observed between January 1950 and June 2018 sourced from CRSP, Compustat, and I/B/E/S. Our sample includes stocks listed on the NYSE, AMEX and NASDAQ. Stocks are only included if they have a market value on CRSP at the end of the previous month and a value for common equity in the firm's financial statement.

Data are compiled on 94 firm characteristics detailed in Green et al. (2017). Table A1 of the Web Appendix lists the variables and the corresponding acronyms.[16] We relate stock returns to characteristics measured at the end of the previous month and assume that annual (quarterly) characteristics are available in month $t-1$ if the firm's fiscal year (quarter) ended at least six (four) months before month $t-1$.[17]

3.2. Break Locations

Our empirical analysis focuses on a four-factor model obtained by regressing firm-level excess stock returns on an intercept, market beta ($\hat{\beta}$), size ($SIZE$), book-to-market value (BM), and momentum (MOM):

$$r_{it} = \alpha_{ik} + r_{zt} + \lambda_{MKT,k}\hat{\beta}_{it-1} + \lambda_{SIZE,k}SIZE_{it-1} + \lambda_{BM,k}BM_{it-1} + \lambda_{MOM,k}MOM_{it-1} + \epsilon_{it}. \tag{3}$$

Measurement of the four characteristics follows Green et al. (2017) so that market beta is estimated using weekly returns and equal-weighted market returns for the three-year period ending in month $t-1$ (with at least 52 weeks of returns), size is the natural logarithm of market capitalization measured at the end of month $t-1$, book-to-market value is the book value of equity divided by the prior fiscal year-end market capitalization, and momentum

[16]This table corresponds to Table 1 of Green et al. (2017) and is only included for reference. We are grateful to Jeremiah Green for making available on his website SAS code to extract the data from CRSP, Compustat and I/B/E/S.

[17]A more detailed explanation of the characteristics is provided in the Appendix of Green et al. (2017). Characteristics are cross-sectionally winsorized at the 1st and 99th percentiles of their monthly observations. The I/B/E/S statistical period date and CRSP monthly end date are used to align I/B/E/S and CRSP data in calendar time.

is computed as the 11-month cumulative return from month $t-12$ through month $t-2$.

This four-factor model is widely used in empirical work which makes it important to investigate the stability of the associated risk premia. Subsequently, we also consider evidence of instability in the expected return premia of a much larger model that includes all 94 characteristics from the data set of Green et al. (2017).

Figure 1 displays the posterior probabilities of the number (top window) and location of breaks (bottom window) affecting the parameters of the four-factor model. Approximately 75% of the posterior weight is assigned to a model with four breaks with the remaining 25% roughly evenly distributed among models with three and five breaks, respectively. Given the strong evidence of four breaks, our empirical analysis focuses on this model, but it is important to bear in mind that our Bayesian approach accounts for uncertainty about both the number of breaks and their location. Detailed discussion of our formal definition of breaks is provided in Appendix D.

The location of each of the four breaks is estimated quite accurately. The four posterior mode break dates are July 1972, October 1981, June 2001, and October 2008 with around 75% of the probability assigned to one particular month.

3.3. Breaks in expected return premia

We next consider how the risk premia vary across the five regimes identified by the four breaks displayed in Figure 1. To this end, the top and middle rows in Figure 2 display the evolution in the equity, value, size, and momentum risk premia, i.e., the values of the λ_k parameters in Equation (3). In addition, the bottom left panel shows the equity risk premium obtained from a single-factor (CAPM). In this model, the equity risk premium varies from 6 to 6.5% in the two regimes prior to 1981, declines to a slightly lower range between 5.4 and 5.7% in the next two regimes, before falling to 3.5% after 2008. These are economically plausible values and suggest a marked decline in the equity risk premium after the GFC. The equity risk premium obtained from the four-factor model (top left panel) evolves along a similar path.[18]

[18] Unlike Pástor and Stambaugh (2001) we do not impose a smoothness condition which imposes that the equity risk premium gradually transitions between regimes.

Next, consider the evolution in the risk premium associated with the book-to-market ratio (top right corner). This declines monotonically from a level above 3.5% per year prior to 1972 to a little over 1% in the period after the GFC. Hence, over the course of our sample, the value risk premium has declined by two-thirds of its initial level, suggesting a sizeable erosion in the amount by which returns on value stocks have outpaced growth stocks.

The size premium (middle left panel) shows a similar, if even starker erosion from a level above 4% per year prior to 1972 to less than 1% after 1981, followed by a sharp reduction to nearly zero in 2001 where it hovers for the remainder of our sample. Hence, the size premium seems largely to have disappeared after 2001.

The momentum premium (middle right panel) behaves very differently from the other risk premia. Starting at 4% per annum in the first regime, this premium drops markedly to a level near 1% before gradually increasing in the remaining part of the sample, dominated by a particularly sharp increase from 2% to 3.5% in 2008.

We conclude from these findings that the equity risk premium, value, and size premium all have undergone secular declines over the nearly seven decades covered by our sample. The reductions are largest for the size and value premia which, at the end of our sample, are either close to zero (size) or markedly smaller (value premium) than in the early sample. Conversely, after declining sharply in the early seventies, the momentum risk premium has subsequently risen steadily and is now close to its original value in the early sample.

The bottom right panel displays the corresponding total risk premium estimated from the baseline model (black line) and the same model without breaks (red line) which yields an average risk premium just above 6%. The combined four-factor risk premium starts out at a little under 7% in the early part of the sample, only to decline near-monotonically to a level close to 5% in the final regime.

To formally evaluate the empirical validity that our factor model leaves no cross-sectional dependence among the idiosyncratic shocks, we estimate average pairwise correlations between residuals and compute the test for cross-sectional dependence (CD) proposed by Pesaran (2004). Under the null of no dependence, the CD statistic has a standard Normal distribution. For our data, the CD statistic is 1.89 so we cannot reject the null hypothesis of no cross-sectional dependence remaining in the residuals.[19]

3.4. Formal tests for time-varying and declining risk premia

Fama and French (2021) find that the value premium has diminished considerably since 1991.[20] Constructing six portfolios sorted on size and book-to-market, they report that the annualized value premium fell from 4.3% (1963-1991) to 0.6% (1992-2019) for large caps and from 7% to 4% for small caps. They cannot reject the null hypothesis that the risk premium is zero in the second subsample, but also cannot reject that the value premium is constant across the two subsamples. However, their tests likely have low power as they use just a handful of portfolios and monthly risk premia tend to be highly volatile.

Exploiting information in a large cross-section of individual stocks, as we do here, circumvents this problem and increases our ability to detect shifts to risk premia. To examine whether our estimates imply that risk premia have vanished, the upper panel of Table 1 displays the final regime's four-factor risk premium estimates (expressed as annualized percentages) and corresponding t-statistics (in brackets below) from our panel break model that regresses firm-level excess returns on market beta, size, value, and momentum as displayed in Equation (3). In the final regime (2008-2018) the value premium (1.28%) is not significantly different from zero. Similarly, at 0.36% per year, we cannot reject the null that the size premium has gone to zero in the final regime. Conversely, with t-statistics of 4.8 and 3.3, respectively, we strongly reject the null that the equity risk premium (4.6%) and momentum risk premium (3.6%) equal zero in the last regime. These results demonstrate that our tests have the ability to identify which risk premia have vanished over time versus which ones remain significant.

To more directly compare our findings to those in Fama and French (2021), we next impose a single break at the same time (1991) as that assumed by Fama and French (2021) and then use our methodology to estimate risk premiums. The results, displayed in the middle panel of Table 1 show that the value premium declined from an annualized 3.36%

[19]The CD test might also be viewed as a test against weak dependence. For large panels ($N > 10$) like ours, weak dependence is unlikely to cause any serious problems for inference (Pesaran 2015).

[20]Schwert (2003) and Linnainmaa and Roberts (2018) also report that the value premium has declined over time.

(1950-1991) to 1.53% (1992-2018). Moreover, in contrast to the results in Fama and French (2021), we find overwhelming evidence in favor of a significant change in the value premium before and after 1991.[21] This demonstrates the added power that comes from using the full cross-section to test for changes in risk premia.

Alquist et al. (2018) report that the size effect diminished shortly after its publication. Performing the same test for a single break in the size premium occurring at 1981, again we find overwhelming evidence in favor of the break. Specifically, the size premium declined from an annualized 4.20% (1950-1981) to 0.65% (1982-2018).

These tests show that risk premia have changed over time but do not reveal whether there has been a systematic downward trend. To explore whether this holds, we separately test whether each of the four risk premia monotonically decline over the five regimes. To do this, we use the Monotonic Relation test developed by Patton and Timmermann (2010) which is nonparametric, does not require a functional form (i.e. linear), and is easy to implement using bootstrap methods. Under the null, the risk premium is constant or weakly increasing across regimes, while under the alternative it is monotonically decreasing. When the bootstrap p-value is less than 0.05, we conclude that the risk premium is significantly monotonically decreasing.

Results from this test are displayed in the lower panel of Table 1. There is clear evidence of a significant monotonically decreasing value risk premium (p-value below 0.05) across our five regimes, and borderline evidence of a monotonically decreasing size premium and total risk premium. However, the equity and momentum risk premia are not significantly monotonically decreasing, in line with Figure 2.[22]

3.5. Breaks vs. time-varying parameters

Our approach assumes that changes in model parameters are rare but discrete. This perspective allows us to more sharply identify the locations at which the largest changes took

[21] As we subsequently explain, the Bayes factor of 179.87 reported in the table represents very strong evidence against the null of unchanged risk premia, see Kass and Raftery (1995).

[22] Studies that suggest the equity premium has declined over time include Blanchard (1993), Jagannathan et al. (2001), and Fama and French (2002).

place.[23] Depending on which events led to the change in the parameters, at other times we might expect parameter changes to be more gradual.

To test whether a time-varying parameter model with smoothly-evolving parameters might better approximate the underlying data generating process compared with our breakpoint approach, we estimate the following specification:[24]

$$r_{it} = \alpha_{it} + r_{zt} + \lambda_{MKT,t}\hat{\beta}_{it-1} + \lambda_{SIZE,t}SIZE_{it-1} + \lambda_{BM,t}BM_{it-1} + \lambda_{MOM,t}MOM_{it-1} + \epsilon_{it} \quad (4)$$

with $\epsilon_{it} \sim N(0, \sigma^2)$, and the parameters, $\theta_t = (\alpha_t, r_{zt}, \lambda_t)$, follow a random walk

$$\theta_t = \theta_{t-1} + u_t, \quad (5)$$

in which $u_t \sim N(0, Q)$ and $Q = \mathrm{Diag}(\phi_1, \ldots, \phi_6)$ is a diagonal matrix so the state innovations are conditionally independent. We further assume that the initial value is Normally distributed $\theta_0 \sim N(\theta, Q)$.

To measure the strength of evidence in favor of our breakpoint specification relative to this time-varying parameter specification, we next compute a Bayes factor. The Bayes factor (86.42) suggests strong evidence in favor of our discrete break specification using standard threshold values.[25]

That model parameters sometimes change very rapidly in a way that is well approximated by discrete breaks is confirmed by inspecting five-year rolling window average estimates of factor risk premia. For example, fluctuations in the size premium tend to be quite sharp, rather than slow moving. Moreover, the risk premia estimated from the time-varying parameter model sometimes change very sharply, e.g., by 150 basis points over one or two months for the momentum factor. These observations support our formal Bayes factor test which strongly favors breakpoints as opposed to time-varying parameters.

[23] Jochmann et al. (2013) also find that the parameters of their return prediction models sometimes change very rapidly.

[24] See Cogley and Sargent (2005) and Primiceri (2005) for popular time-varying parameter specifications.

[25] Bayes factors are constructed from the marginal likelihood of each model computed using the method of Chib (1995) and are the preferred Bayesian model comparison approach as they integrate over all parameters in the model and inherently penalize model complexity. Bayes factors between 1 and 3 are inconclusive, values between 3 and 20 indicate positive evidence in favour of our baseline model, while values greater than 20 indicate strong evidence (Kass and Raftery 1995).

3.6. *Aggregate and idiosyncratic volatility*

The top panel of Figure 3 graphs the aggregate volatility obtained from our Bayesian panel breakpoint model in Equation (3) and estimated as the standard deviation of r_{zt} in each regime. Aggregate volatility starts just below 16% per year, rises to a level above 17% in 1972, before monotonically declining throughout the remainder of the sample, reaching 14% in the final regime (2008-2018), its lowest value of the sample.

Our approach also allows the volatility of the idiosyncratic error term ϵ_{it} to vary across regimes. To see how the average idiosyncratic volatility evolves over time, the lower panel of Figure 3 graphs the value-weighted average of firm-level residual volatility estimates through our sample, expressed as an annualized percentage. In the first regime (1950-1972), idiosyncratic volatility is very low, amounting to less than 10% per year. Average idiosyncratic volatility then nearly triples in 1972, before further rising above 30% per year in 1981 and to a level close to 40% in 2001.[26] Following the GFC, idiosyncratic volatility comes down substantially, declining to a level just below 25%.

3.7. *Mispricing*

To gain insights into how any mispricing has evolved over time, Table 2 evaluates the cross-sectional distribution of α estimates. For each of the five regimes, we report the average posterior mean and standard deviation along with various percentiles of the α estimates from regressions of firm-level stock returns on market beta, size, value, and momentum as displayed in Equation (3). The final column reports the proportion of individual firm-level alpha estimates that are significantly different from zero at the 5% level, using a two-sided test. In each panel, the bottom row shows the same statistics obtained from a model without breaks fitted to the full sample. All alpha estimates use full-sample information and are reported in annualized percentage terms. These values could therefore not have been exploited in real time for improved investment performance and, as they use shorter samples, are more strongly affected by estimation error than the constant-parameter estimates shown

[26] Consistent with these findings Campbell et al. (2001) report that firm volatility has increased markedly from 1962 to 1997.

in the bottom row. In practice, this means that the cross-sectional range of alpha estimates within each regime is somewhat wider than is normally the case.

With this caveat in mind, first consider the top panel (all stocks). In the full sample, the mean alpha estimate is 0.41%, or 41 basis points (bps) per annum with an inter-quartile range from -0.71% to 2.25% and a standard deviation of 2.87. Moving to the individual regimes, we find stronger evidence of mispricing in the early parts of our sample: the mean alpha estimate is around 2.5% per year in both the first (1950-1972) and third (1981-2001) regimes and the 75th percentile is more than twice as high in these regimes as its average, full-sample value. In these early regimes, sizeable proportions (25% and 20%) of the alpha estimates are significantly different from zero.

Evidence of mispricing in individual stocks has markedly reduced over time, however, and the mean alpha estimates are negative, at -0.30% and -0.62% per year, in the final two regimes. Although the range of alpha estimates is wider in these regimes than they are in the full sample, this can to a large extent be attributed to the greater effect of sampling error in the shorter-lived regimes. Indeed, the proportion of stocks whose alpha estimates are significantly different from zero is much smaller in the final two regimes – eight and six percent, respectively – than in the full sample (18%).

We would expect to find stronger evidence of mispricing in the four-factor model among the smallest, most illiquid stocks that are harder to trade. To see if this is indeed the case, the middle and bottom panels of Table 2 show separate results for larger stocks and micro caps. Consistent with our expectation, we find a far wider interquartile range of alpha estimates for micro caps (-19.85%; 12.74%) than for the larger stocks (-0.58%; 1.95%). The percentage of stocks with significant alpha estimates is also larger for micro caps than for the larger stocks.

We conclude from these findings that there is substantial *ex-post* evidence of time-variation in mispricing for individual stocks during our sample and that (i) the mispricing is much stronger during the early parts of our sample, declining significantly after 2001; and (ii) mispricing is stronger for micro caps than for large stocks.

3.8. Which model parameters are affected by instabilities?

Our empirical analysis up to this point uncovers strong evidence that alphas, risk premia, and idiosyncratic volatilities change across the five regimes identified by our model. However, while we have inspected the magnitude of the shifts in these parameters across regimes, we have not formally tested whether all parameters change at the break dates or whether they are unaffected by regime shifts.

To address this point, we next conduct formal hypothesis tests that disentangle which parameters are most affected by instabilities. Specifically, we estimate several restricted versions of the baseline model that allow for breaks in (i) mean coefficients (α and λ) only, (ii) idiosyncratic volatility (σ) only, (iii) α only, and (iv) λ only. To gauge the strength of evidence in favor of our general baseline model relative to each restricted model, we again compute Bayes factors. The results, displayed in Table 3 for the full sample, i.e., across all breaks, as well as on a break-by-break basis, show overwhelming evidence that all four breaks are broad-based and affect both the mean and volatility parameters. Focusing on the mean coefficients, there is also strong evidence that all four breaks hit both the risk premia (λ) and pricing errors (α).

We conclude that there is strong support for discrete regime shifts in the parameters of the simple four-factor return regression model in Equation (3). Moreover, these shifts are broad-based, economically large, and highly statistically significant.

3.9. Breaks and Macroeconomic Risks

Studies such as Lettau et al. (2008) argue that variation in macroeconomic risk helps explain movements in the equity risk premium. Using quarterly data from 1952:1 to 2002:4, these authors identify a structural break in 1992 at which point volatility declines, and they find a striking correlation between movements in macroeconomic risk and the stock market.

To see if a similar relationship holds for our data, we next examine if low frequency movements in macroeconomic risk are related to low frequency movements in the total risk premium identified by our model. Computing the average real uncertainty measure from Jurado et al. (2015) and Ludvigson et al. (2021) within the regimes identified by our baseline model, along with our total risk premium estimate (the black line in the bottom-right panel of Figure 2), we find a 0.74 correlation (across regimes) between the two series.[27]

Our total risk premium estimate is also highly correlated with low frequency movements in the dividend-price ratio. The average dividend-price ratio within regimes identified by our model has a 0.65 correlation with our total risk premium estimate.[28]

We next address whether exposure to such regime shifts is itself a source of risk that is priced in the cross-section of equity returns.

4. Break Risk Factor

The empirical evidence in the previous section shows that risk premia associated with stock or firm characteristics such as market betas, size, book-to-market value, and return momentum are affected by pervasive and economically large breaks. Exposure to this type of instability in risk premia introduces a separate source of risk in individual asset returns as well as returns on portfolios focusing on particular investment styles. For example, investors holding small value stocks will be exposed to the risk that the risk premia associated with size and value will change in a manner that makes their return distribution more difficult to estimate and predict than if risk premia were constant. Break risk matters particularly to long-term buy-and-hold investors who do not rotate their portfolio allocations very frequently, but can also be important to short-term investors because of the challenges associated with detecting breaks and updating estimates of risk premia in real time.

These arguments suggest that instability in the risk premium process is itself a source of risk that could give rise to a break risk factor. This is economically plausible because the breaks identified by our approach occur during economic and financial crises. Stocks more exposed to major macroeconomic events and financial crises might plausibly be expected to earn higher returns as compensation for risk exposure to "bad states".

[27]Moreover, recursive real-time estimates of our break probabilities are positively correlated with the monthly real (0.16 correlation), macroeconomic (0.27), and financial (0.35) uncertainty measures taken from Jurado et al. (2015) and Ludvigson et al. (2021). These uncertainty measures tend to spike around our posterior mode break dates, as do our real-time break probability estimates.

[28]Our total risk premium estimate has a similar positive, albeit less pronounced, correlation with the earnings-price ratio (0.34).

4.1. Individual stocks' exposure to instability risk

To establish whether instability risk is economically important, we must demonstrate that (i) regime shifts are pervasive and affect the returns of multiple stocks or portfolios; (ii) exposure to instability risk is priced in the cross-section and stocks with greater exposure to this type of risk earn higher returns, on average, than stocks with low exposure, assuming that instability risk does not hedge against other sources of risk.

The first point (pervasiveness) is indirectly established by the fact that we use a panel regression approach to identify common breaks in style risk premia. Because our approach penalizes large models with many parameters, it is highly unlikely to identify regime shifts that only affect a small subset of stocks. To further strengthen this point, we provide formal evidence in Section 5 that a wide set of industry and style-sorted portfolios are affected by changes in regimes.

To address the second point, we need a measure of how much individual stocks are affected by breaks which we can use to sort stocks into portfolios with high and low break sensitivities. Moreover, we need to be able to compute this measure in real time before performing the portfolio sorts.

To measure individual stocks' sensitivity to instability risk, we build on a literature that links large changes in consumption growth and heightened macroeconomic uncertainty, both features of the break dates identified by our empirical analysis, to variation in aggregate valuation measures such as the price-dividend ratio. For example, disaster risk models such as Barro (2009), Gabaix (2012), Martin (2013), and Wachter (2013) imply that assets whose prices fall when a disaster occurs have a higher expected return because of their higher exposure to disasters. This is similar to the mechanism in our analysis where stocks with a higher exposure to break risk earn a higher risk premium. Similarly, Berkman et al. (2011) find that their crisis severity index is positively correlated with the earnings-price ratio and dividend yield, while Lettau and Van Nieuwerburgh (2008) show that breaks to the steady state dividend growth rate can lead to parameter instability in regressions of returns on the lagged dividend-price ratio.

Using these insights, we estimate a panel break model that relates individual stock returns to the lagged value of the aggregate log dividend-price ratio, dp_{t-1}:[29]

$$r_{it} = \alpha_{ik} + \beta_{ik} dp_{t-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k. \tag{6}$$

Next, we generate out-of-sample return forecasts from Equation (6) estimated with and without breaks.[30] For each stock, i, and each month in the sample, t, we then compute the difference between forecasts from the panel model with breaks ($\hat{r}_{it,Brk}$) and without breaks ($\hat{r}_{it,NoBrk}$):

$$BRK_{it} = \hat{r}_{it,Brk} - \hat{r}_{it,NoBrk}, \qquad i = 1, \ldots N, \qquad t = 121, \ldots, T. \tag{7}$$

BRK_{it} is larger for stocks with greater exposure to break risk, and we refer to this as stock i's break risk characteristic (at time t). Finally, as we next describe, we examine if differences in such exposures translate into differences in risk premia.

4.2. Fama-MacBeth Regressions

We evaluate the ability of our break risk measure in Equation (7) to explain the cross-section of returns by estimating cross-sectional regressions each month

$$r_{it} = r_{zt} + \lambda_{BRK,t} BRK_{it-1} + \lambda'_{2t} X_{it-1} + \epsilon_{it}. \tag{8}$$

Here X_{it-1} contains log-size, log book-to-value, and prior one-year return performance, excluding the most recent month, for the ith stock at time $t-1$.[31] Next, following the Fama-MacBeth methodology, we use the time-series estimates of $\lambda_{BRK,t}$ and λ'_{2t} to evaluate the mean and standard deviation of these slope coefficients.

The first column of the top panel of Table 4 displays the results. The break risk factor obtains nearly the same significance as the book-to-market ratio in explaining the cross-

[29] Empirically, Paye and Timmermann (2006) and Rapach and Wohar (2006) find evidence of breaks in the slope coefficient of the dividend-price ratio in return regressions such as (6). Smith and Timmermann (2021) also provide evidence of breaks in the relation between stock returns and the lagged dividend-price ratio using data on individual stock returns but do not address whether these breaks are more important to particular types of stocks ("styles").

[30] We use a 10-year warm-up estimation period.

[31] Following Novy-Marx (2013), we do not include market beta as a control variable.

section of returns and its t-statistic is approximately one-and-a-half times larger than that of the size and momentum variables. Average returns are also higher for firms highly exposed to break risk than for those with the smallest exposure.

To corroborate that our results are not overly sensitive to the proposed measure of break risk exposure, columns 2-5 in Table 4 present results using alternative proxies of the break risk factor. Our second measure uses the root-squared difference between forecasts produced by panel models fitted with and without breaks. The third, fourth and fifth columns use the difference at each point in time in the intercept, slope and volatility parameters, respectively, estimated from panel models with and without breaks. All five measures are highly statistically significant.[32]

Following Novy-Marx (2013), the bottom panel of Table 4 reports results from the same analysis on break risk measures that have been demeaned by industry. The results are broadly similar, except the t-statistic of every break risk measure is increased, so that adjusting the risk measure by industry obtains even more power to explain the cross-section of expected returns.

These results demonstrate the robustness of our findings. From herein we focus on the break risk factor measured by the difference between the forecasts produced by the panel models with and without breaks in Equation (7).

4.3. Sorts on break sensitivity

Running Fama and MacBeth (1973) regressions on individual stocks places considerable emphasis on micro-cap stocks that make up a sizable share of the number of stocks but only account for a small fraction of the total market capitalisation. Such regressions may also be sensitive to outliers and impose a potentially misspecified parametric relation between the variables, compromising subsequent inference.

To alleviate this concern, we next construct value-weighted portfolios sorted according to our instability risk factor and provide a nonparametric test of the hypothesis that exposure

[32] All results use Newey and West (1987) heteroskedasticity-adjusted t-statistics. The third measure (based on the intercept) has the least power to explain the cross-section of expected returns but still obtains a significant t-statistic of 2.62.

to break risk predicts average returns in the cross-section. Table 5 displays results for these portfolios sorted on our break risk factor. The first row ("Low") shows results for the bottom quintile of stocks ranked by break sensitivity, while the fifth row ("High") shows results for the stocks most sensitive to breaks. Column one reports the average monthly return earned by each quintile portfolio, followed by the alpha and slope coefficients obtained from time-series regressions of the portfolio returns on the three factors of Fama and French (1993) – market (MKT), size (SMB) and value (HML) – with t-statistics reported in brackets below.

Returns on the break-sorted portfolios increase monotonically with our risk factor and the high-sensitivity quintile portfolio earns a 0.28% higher average monthly return than the low-sensitivity portfolio, equivalent to an annualized return premium of 3.36% which is statistically significant at the 5% level with a t-statistic of 2.33.

Turning to the risk-adjusted performance from the three-factor regressions, once again we see monotonically increasing values of alpha as we move from the least to the most break-sensitive stocks. Moreover, the alpha estimate of both the least break-sensitive stocks (at -0.17% per month) and the most break-sensitive stocks (at 0.20%) are both significantly different from zero. At 0.37% per month or more than 4% annualized, this difference is also economically large.

To alleviate concerns about transaction costs raised by Novy-Marx and Velikov (2015) and Hou et al. (2020), we follow Chordia et al. (2017) and perform the same analysis omitting all stocks with a price below $3 or a market capitalisation below the 20th percentile of the NYSE capitalisation distribution. The bottom panel of Table 5 displays the results which, while marginally weaker, tell the same basic story.

These results provide further cross-sectional evidence of the existence of an economically important break risk factor. Stocks whose expected return processes are most sensitive to the instability in risk premia identified by our methodology earn both higher average returns (about 3% per year) and higher risk premia (about 4% per year) than stocks with the lowest sensitivity to breaks.

4.4. Break risk and other risk factors

The past two decades has seen an explosion in the number of factors that reportedly explain the cross-section of expected returns, a phenomenon labeled the 'factor zoo' by Cochrane (2011).[33] Amidst this plethora of factors, it is important to address whether our proposed break risk factor remains significant even after accounting for the presence of other candidate risk factors. To this end we first consider the relation between the break risk factor and existing risk factors. The upper panel in Table 6 reports pairwise correlations among a number of factors, including the market, book-to-market, size, momentum, and break risk factor. Our break risk factor is relatively weakly correlated with the conventional four risk factors, with correlations ranging from -0.26 (momentum) to 0.28 (market).

The middle panel reports the maximum as well as the 10th, 25th, 50th, 75th and 90th percentiles of the correlations between our break risk characteristic in Equation (7) and the other 94 characteristics. These characteristic correlations are computed for each series in the cross-section and the table reports the average over the cross-section. The median (maximum) correlation is 0.10 (0.42), consistent with no other single characteristic or factor being able to explain the majority of the variation in break risk.

The five characteristics most strongly correlated with our break risk characteristic are, in descending order, idiosyncratic return volatility, return volatility, volatility of liquidity (share turnover), cash flow to debt, and cash flow volatility. Evidently break risk contains information related to both return and cash flow volatility. Interestingly, not even a combination of these five characteristics explains much of the variation in break risk. A regression of break risk on the five characteristics produces an R^2 of 0.24, supporting our claim that break risk contains genuinely new information that is not spanned by existing characteristics or risk factors.

[33] Using a high t-statistic threshold of three Harvey et al. (2016) identify approximately 150 factors.

4.5. Risk Factors in Individual Regimes

Different risk factors and characteristics may be important at different points in time and our panel break approach is ideally suited for addressing time variation in which risk factors help explain cross-sectional variation in expected returns. We do this by estimating the model

$$r_{it} = \alpha_{ik} + r_{zt} + \lambda_{BRK,k} BRK_{it-1} + \lambda_{2,k}' X_{it-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k, \tag{9}$$

where $\lambda_{BRK,k}$ denotes the risk premium on our break risk factor and $\lambda_{2,k}$ captures the risk premia estimates on the remaining 94 characteristics in the kth regime.[34]

Table 7 reports the outcome of estimating Equation (9) on our panel of firm-level stock returns. For each regime identified by our model, we show the characteristics that earn significant risk premia using a t-statistic threshold of three as proposed by Harvey et al. (2016). The total number of selected characteristics in each regime is reported at the bottom of the table.[35]

In total, 24 different factors (out of 95) get selected at least once in our sample. Only the market risk factor gets selected in every regime. The size (market value) and book-to-market risk factors both get selected in the first three regimes, but not in the final, consistent with our findings in Figure 2 that risk premia on these factors are waning. The momentum risk factor is selected in the third and fourth regimes, again consistent with Figure 2.

Our proposed break risk factor gets selected in the last three regimes, i.e., the period from 1981-2018. This is strong evidence that the break risk factor is important in explaining cross-sectional variation in stock returns. In fact, besides the three Fama-French risk factors and our break risk factor, none of the other risk factors gets selected in more than a single regime, indicating that the explanatory power of these factors is not stable over time.

[34] Smith (2018b) performs Bayesian model selection of the 94 characteristics, allowing for model uncertainty and multiple breaks in the set of characteristics that independently inform the cross-section of expected returns. Here, we further include our proposed break risk characteristic to evaluate whether it holds information about the cross-section of returns that is not spanned by the 94 characteristics.

[35] The break dates are aligned with those identified in our earlier four-factor model.

Interestingly, the number of factors that gets selected in an individual regime peaks at 12 during 2001-2008 before dropping sharply to only three factors in the final regime after the GFC. Only the market risk factor, momentum, and our break risk factor get selected after the financial crisis.[36]

We conclude from this evidence that only five factors – market risk, size, book-to- value, momentum, and our new break risk factor – have consistent power over cross-sectional variation in stock returns for the majority of the sample. This is a new finding and illustrates the kind of insights our approach can be used to provide. In fact, as shown in the bottom panel of Table 7, a constant-parameter approach that uses the full data sample to select factors chooses 16 factors, failing to separate out the many factors whose effect on the cross-section of stock returns is mainly confined to short sub-samples from those factors with a more robust effect.

5. Pervasiveness and timing of Breaks to Industry and Characteristics-sorted Portfolios

Cross-sectional returns data on individual stocks, the main focus up to this point, can be used to boost the power of our ability to detect breaks. Conversely, returns on more broadly diversified portfolios formed along industry, characteristics or "style" lines can be used to understand whether certain types of firms are more affected by break risk than others, helping us better interpret the economic sources and investment consequences of exposure to break risk.

Pursuing this idea, this section estimates our panel break model on a set of industry and characteristics-sorted portfolios. Next, using these portfolios, we introduce the noncommon breakpoint procedure developed by Smith (2018a) which allows breaks to hit any subset of series in the cross-section and at different times. This approach enables us to accomplish

[36] Green et al. (2017) acknowledge that the assumption of time invariance implicit in the majority of firm characteristic studies is unlikely to hold after 1980 because of "changes in the volume, nature, and costs of trading in stocks that occurred from 1980 to 2014, including Reg. FD, the decimalization of trading quotes, Sarbanes-Oxley, accelerated SEC filing requirements, auto quoting, and computerized long/short quantitative investment". Without using a formal test, they identify instability in the number of selected characteristics which falls from 12 to two after 2003. Our approach finds a similar reduction from 12 to three factors slightly later (after 2008).

three tasks: (i) distinguishing between market-wide and industry or style-specific breaks; (ii) evaluating whether particular assets are hit earlier or later in the break cycle; and (iii) evaluating whether lead-lag relations vary through time. For instance, one might expect that the oil industry played a leading role during the 1970s, telecommunications during the early-2000s, and financials/real estate during the GFC.

5.1. Break Risk for Industry and Style Portfolios

We start by estimating our panel break model using monthly returns on 30 value-weighted industry portfolios, a set of 5×5 portfolios sorted on size and book-to-market, and a set of 5×5 portfolios sorted on size and momentum. All returns are computed in excess of a one-month T-bill rate. Our data on portfolio returns is a bit longer than that used on the individual stocks and run from July 1926 through December 2019. This longer sample of portfolio returns provides a way to cross-validate the robustness of our findings on the effect of breaks on individual firms' returns. Data are sourced from Ken French's website along with aggregate data on the three factors of Fama and French (1993).

To identify differences and similarities in how breaks affect different types of stocks, our analysis is undertaken separately for the three sets of test portfolios using the specification in Equation (6). This allows us to address whether breaks are specific to particular investment styles or industries, or whether they are more pervasive and affect most or all portfolios.

First consider the evidence of breaks in the model fitted to the 30 industry portfolio returns. For the 1926-2019 sample, the mode (and mean) for the number of breaks is six, with approximately 88% of the probability mass distributed between five and six breaks, corresponding to a break occurring roughly once every fifteen years. The timing for most of the breaks is well defined with posterior probabilities concentrated around 1929, 1973, 2001, and 2008, thus coinciding with major economic events such as the Great Depression, the oil price shocks of the 1970s, the dotcom crash in the early-2000s, and the Global Financial Crisis. Reassuringly, in the sub-sample that overlaps with the individual stock returns data (1950-2018), the break dates identified for the industry portfolio returns are either the same or very close. Compared to the results for the individual stocks, the posterior probability mass for the break locations is more disperse, indicating that the effect of breaks on different industry portfolios was not confined to a single month but diffused gradually through time.

A similar number of breaks is identified for the 25 portfolios sorted on size and either book-to-market or momentum. For example, the model fitted on the portfolios sorted on size and book-to-market identifies seven breaks with similar locations to those for the industry portfolios.[37]

Having established the similarity in both the number and location of breaks across different portfolios, we next analyze which portfolios exhibit the greatest sensitivity to breaks. To this end, we rank the portfolios by their sensitivity to breaks as measured by the mean squared difference between the forecasts from models estimated with and without breaks.[38]

The top panel in Table 8 shows break sensitivity results for the top and bottom quintile of industries. Returns on telecommunication stocks exhibit the greatest sensitivity to breaks, followed by the utilities, oil, business equipment, and financial industries. Stocks in the wholesale, mining, textile, books and meals industries are least sensitive to breaks. Cyclical industries thus appear to be more sensitive to breaks than non-cyclical industries and the break sensitivities of the first group tend to be three to four times greater than those of the latter group of industries.

Among the 25 portfolios sorted on size and book-to-market ratio (middle panel in Table 8), small firms' returns are most sensitive to breaks and big firms least sensitive. Differences in break sensitivity are economically large with small firms' break sensitivity being six to seven times larger than that of large firms. Though size matters more to break sensitivity than book-to-market value does, there is also a clear relation between firms' book-to-market ratios and their break sensitivity. Conditional on firm size, value firms are more sensitive to breaks than growth firms and there is a near-monotonically decreasing relation between book-to-market ratio and break sensitivity.

The bottom panel in Table 8 shows similar findings for the stocks sorted on size and momentum. Conditional on firm size, "loser" stocks with the smallest prior returns are

[37]This finding does *not* follow automatically since we estimate our panel break model separately for the industry returns and the two sets of 5×5 characteristics sorted portfolios and so the break detection could be very different.

[38]Our results are robust to using other sensitivity measures such as the standard deviation of the estimated intercept, slope coefficient or residual variance across regimes.

more sensitive to breaks than "winner" stocks with an almost monotonically decreasing relation between prior returns and break sensitivity.

These findings suggest that firms normally thought of as being riskier (small firms and value stocks) also have greater exposure to breaks in their return processes. Firms with poor prior-year return performance also tend to be more exposed to break risk which could be related to the occasional resurgence in the returns of "loser" stocks documented by Daniel and Moskowitz (2016).

5.2. Market-wide versus Characteristics-specific breaks

We next evaluate whether the breaks are market-wide or specific to certain industries or styles such as size, value, and momentum using (excess) returns on 30 value-weighted portfolios (10 univariate decile sorts on each of size, value, and momentum).

Adopting the methodology developed by Smith (2018a), we allow any subset of assets $1 \leq N_k \leq N$ to be affected by the kth break occurring at the common time τ_k. This is accomplished by generalizing Equation (6) to

$$r_{it} = \alpha_{ik} + \beta_{ik} dp_{t-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k \tag{10}$$

in which $\beta_{ik+1} = \beta_{ik}$ for those portfolios that are not hit by the kth break. Conversely, the common break assumption in the baseline model in Equation (6) restricts all portfolios to be hit by breaks ($N_k = N$ for all k).

Starting with the style-sorted portfolios, Figure 4 displays the estimated break dates for the model in Equation (10). In ranked order, portfolios 1 through 10 track decile portfolios ranked on return momentum (winners followed by losers), portfolios 11 and 20 represent the lowest and highest decile of book-to-market-sorted portfolios (value and growth, respectively), and decile portfolios 21 and 30 contain the smallest and biggest firms sorted on market capitalization. Our sample period goes back to 1926 so that, in addition to the four post-war breaks identified in the baseline analysis, we detect a further three breaks in 1929, 1933, and 1940.

The figure nicely illustrates that some breaks affect all style-sorted portfolios while other breaks are driven by a single style. For example, three breaks (in 1929, 1973, and 2008) are common across all three investment styles; another two breaks are specific to the size (1933) and book-to-market (1982) sorted portfolios; finally, two breaks (1940 and 2001) affect the size and momentum-sorted portfolios but not the book-to-market-sorted portfolios.

Conditional on a break affecting a given investment style, almost all of the decile portfolios within that style are affected by the break. This shows that the breaks we identify genuinely are linked to style characteristics. Our findings further demonstrate that some breaks such as those observed in 1929, 1973, and 2008 are very broad, affecting stocks regardless of their characteristics, while others are more style-specific.

This finding has potentially important economic implications. First, it suggests that the risk premia of portfolios that are exposed to the value or growth risk factors are not as sensitive to break risk as are portfolios exposed to size- or momentum risk. Second, it suggests that when the risk-premium process changes for the book-to-market, size, and momentum factors, it has a pervasive effect across stocks at different ends of the spectrum sorted on these variables.

Applying the same approach to the industry portfolios, we find that almost all industries are affected by each of the breaks. Moreover, while the two earliest breaks affect 25 and 26 of the industries, respectively, the last two breaks affect 29 and 30 of the industry portfolios, suggesting that the breaks have become more pervasive over time. Firms in different industries are likely to have non-zero loadings on the style factors which helps explain why the vast majority of industries are affected by each of the breaks even when some style portfolios are not impacted by all breaks.

5.3. Speed of adjustment to breaks

Studying the speed with which different types of stocks react to breaks can provide insights into the underlying economic drivers of such breaks. Indeed, stocks with different style- or industry characteristics may react more or less rapidly to breaks due to the gradual dissemination of information about breaks which is likely to take time to uncover and process. Hou (2007) reports that slow information diffusion across sectors is a primary driver of lead-lag dynamics in return predictability, causing the lead-lag relation between big and small firms to occur primarily within industries. The effect is caused by a slow reaction to negative information. The lead-lag effect is larger for firms that are smaller, less competitive, and neglected. Hong et al. (2007) find that the returns of industries such as retail, services, commercial real estate, metal, and petroleum lead the aggregate market by up to two months. Similarly, Croce et al. (2019) report evidence that the lead-lag relation across firms varies through time.[39]

These findings suggest that information diffusion across markets is gradual and that the aggregate stock market responds to information in industry returns with a lag. Generalizing the model in Equation (10) to allow the timing of breaks to vary across assets, we have

$$r_{it} = \alpha_{ik} + \beta_{ik} dp_{t-1} + \epsilon_{it}, \qquad t = \tau_{k_i - 1} + 1, \ldots, \tau_{k_i} \tag{11}$$

where now τ_{k_i} denotes the time at which the ith portfolio is hit by the k_ith break.[40] The common break assumption in the baseline model in Equation (6) restricts all portfolios to be hit at the same time as $\tau_{k_i} = \tau_k$ for all i and k. By relaxing this assumption, the noncommon break model in Equation (11) captures the possibility of shifts in the lead-lag pattern in which individual return series are affected by breaks.

Figure 5 displays the timing of the noncommon breaks across the 30 industries for four of the most economically interesting break dates, namely 1929, 1973, 2001, and 2008. The leading industries identified by our approach are broadly aligned with those identified by Hong et al. (2007) as Financials, Telecommunication, Retail, Services, Steel, Chemicals, Oil, and Construction are the first industries to be affected by breaks to the return process. Some of the leading industries, such as Oil, Financials, and Telecommunications are also most sensitive to risk as can be seen from Table 8.

Allowing the lead-lag relations to vary through time turns out to be empirically important. For instance, Financials had a leading role during the 1929 Wall Street Crash (top left window) and the Global Financial Crisis (bottom right), while Telecommunication stocks were the first to be affected by the break associated with the dotcom crash (bottom left),

[39] Croce et al. (2019) find that the telecommunications industry became more leading from 1995 to 2000, real estate during the early-2000s, and finance after 2005. Consumer goods leads national output by about one month, manufacturing lags by about two months, and business equipment lags consumer goods by nearly three quarters.

[40] For full details of the model and estimation we refer the reader to Smith (2018a).

and Oil stocks were affected earlier than other sectors by the break associated with the oil price shock of 1973 (top right).

The speed of information diffusion across different industries, as measured by the delay between the first and final industry hit by a break, has increased over time. The average lead-lag delay across the first four industry breaks is 8.25 months while it equals 3 months across the final three breaks.

We next undertake a similar analysis across the 30 style-sorted portfolios, i.e., 10 univariate sorts on each of size, value, and momentum. Using the same methodology and focusing on the same four breaks as in Figure 5, Figure 6 reveals several interesting patterns. First, momentum portfolios tend to be among the earliest to be affected by breaks, with "loser" stocks moving before "winner" stocks. Second, size-sorted portfolios tend to be affected before stocks sorted on book-to-market ratio, with large stocks moving before small stocks. Stocks sorted on book-to-market tend to move slowest, with growth stocks generally moving before value stocks which are lagging or even not being hit altogether in the case of the break associated with the Dotcom bubble.

In summary, our analysis uncovers a number of new insights. First, we show that far from being stable, the lead-lag patterns in portfolio returns vary considerably over time and are related to the cause of the event triggering the break. Second, we show that stocks with low prior-year returns tend to be affected before stocks with high prior-year returns, that large caps are affected before small caps (consistent with Lo and MacKinlay (1990)), and that value stocks tend to be affected later than growth stocks.

5.4. Industry timing premium

Our finding that industries and investment styles are affected at different speeds by breaks in risk premia begs the question whether firms that are hit earlier by breaks earn a "timing premium" relative to those that are hit later. Stocks whose returns move earlier tend to be more important for the price discovery process and should be more highly correlated with the market, justifying a positive timing premium. Consistent with this, Croce et al. (2019) find that firms in leading industries pay an annualized return that is 4% higher on average than that paid by firms in lagging industries, with 1.5-2% being a pure timing premium on advance information. Similarly, Savor and Wilson (2016) show that firms scheduled to report earnings earlier in the cycle earn an abnormal return of almost ten percent per year.

To examine this point we recursively estimate the noncommon breaks model in Equation (11) on the 30 industry portfolio returns. Next, we sort the industry portfolios into quintiles based on the timing of the final breakpoint detected. A zero-cost investment strategy that goes long in the top (leading) and short in the bottom (lagging) quintile portfolios earns an annualized return of 1.4% which is statistically significant even after controlling for the market, size, value, and momentum factors.

6. Out-of-sample Return Forecasts and Portfolio Implications

Our final section analyzes the out-of-sample accuracy of the return forecasts generated by our panel break model and examines investment implications of these forecasts.

6.1. Accuracy of out-of-sample return forecasts

We begin by evaluating the out-of-sample forecast accuracy of our panel break model and comparing it to a range of alternative specifications that either are simpler versions of our general specification – allowing us to identify the features of our model that are particularly important – or use a different approach to capture time variation in expected returns. Specifically, we compare our approach to four benchmarks: a univariate time series break model, a constant-parameter panel model, a time-varying parameter model featuring small changes to the parameters every period, and the historical average.[41]

Using a warm-up period of ten years, forecasts are generated by recursively estimating each month using only historically available data our model and the benchmark forecasting models based on the specification in Equation (6). Forecasts from our model incorporate any uncertainty surrounding the number and timing of breaks as well as parameter uncertainty. Market portfolio forecasts are constructed as the value-weighted average of the portfolio-

[41] Consistent with our panel break model, the time-varying parameter model (not shown here) shows considerable evidence of parameter instability and a notable downward drift in the risk premium estimates of the equity, value, and size premia.

level forecasts.

To evaluate whether any improved predictive accuracy is statistically significant, we use the test statistic of Clark and West (2007) that accounts for our forecasting models being nested which can lead conventional test statistics to have nonstandard distributions. Against the four benchmarks, we find that the panel break model performs significantly better out-of-sample at the 10% critical level for between 25 and 27 of the 31 industry portfolios (including the market portfolio). Our panel break model also produces significantly better return forecasts for between 20 and 22 of the 26 portfolios sorted on size and value and for between 22 and 23 of the 26 portfolios sorted on size and momentum. Across all 83 portfolios and four benchmark models (332 cases), return forecasts from our model never significantly underperform.

6.2. Investment Implications

We next explore the economic significance of our model's return forecasts for a risk-averse mean-variance investor who allocates her portfolio every month between the riskless asset and a risky portfolio constructed from each set of test portfolios. In each month t, the risky portfolio is constructed as the vector of weights ω_t chosen to maximize the expected utility from the return on the risky portfolio next month, $r_{p,t+1}$:

$$E[U(r_{p,t+1} \mid A)] = r_{f,t} + \omega_t' \hat{r}_{t+1} - \frac{A}{2} \omega_t' \hat{S}_t \omega_t. \tag{12}$$

Here $r_{f,t}$ denotes the risk-free rate in month t, $\hat{r}_{t+1}$ denotes the vector of return forecasts for month $t+1$ computed using information available at month t, $\hat{S}_t$ denotes the covariance matrix that is estimated using the residuals from the return prediction model at month t, and A denotes the risk aversion coefficient which is set equal to three following Campbell and Thompson (2008). We constrain the portfolio weights to sum to one and rule out any short selling or leverage.

Compared to the optimal portfolio weights based on historical averages of the moment estimates, the average industry allocations based on our out-of-sample panel break return forecasts are substantially higher for the smoke, telecommunications, services, and financial industries. Conversely, the weights are lower for beer, healthcare, autos, and business equipment.[42] Certainty equivalent returns of the panel break model are about 2% per annum higher than that of the alternative benchmarks.

Undertaking a similar investment exercise on the 25 (5×5) portfolios sorted on size and value, we compute portfolio allocations across the five portfolios comprising (i) the smallest stocks and (ii) stocks with the highest book-to-market ratios, on average, for the first and final decades of our out-of-sample period. Average allocations to the five smallest stock portfolios declined from 40% to just 6% from the first to the final decade. Similarly, average allocations to the highest book-to-market ratio (value) stock portfolios declined from 42% to 4%. These shifts in allocations are driven by the systematic decline in the size and value premia identified in our empirical analysis.[43]

For both sets of 5×5 sorted portfolios, we find utility gains in the neighborhood of 2% per annum relative to the four benchmarks. The panel break model could therefore have been used in real time to generate return forecasts that, when implemented in a simple investment strategy, produce sizeable economic gains.

6.3. Rotation of Portfolio Allocations

To better understand what generates the utility gains associated with our model's return forecasts, we next consider how the portfolio weights change around break points. To this end, Table A3 of the Web Appendix reports the average allocation in three-year windows before and after the three most recent breaks in our sample. For each set of test assets, we limit the results to the five portfolios whose portfolio allocations are most strongly affected by these three breaks. For the industry portfolios (top panel), weights were significantly reduced for oil and chemical stocks after the 1973 oil price shock while the allocations to financial, services, and telecommunication stocks came down significantly following the break associated with the end of the dotcom bubble. Finally, financial, services and oil stocks all saw reduced allocations after the break associated with the GFC. Figure 7 complements these findings by showing 36-month trailing moving average estimates of the portfolio allo-

[42] A full set of results is presented in Web Appendix Table A2.

[43] We find a similar shift away from the smallest stocks for the 25 portfolios sorted on size and momentum.

cations to the industry whose portfolio allocation is most strongly affected by each of the breaks.

A similar rotation is seen among the style-sorted portfolios (middle and bottom panels of Table A3 of the Web Appendix): Following all three breaks, we see a large reduction in the allocation to large stocks with low prior returns. Moreover, the 1973 break induces a sharp reduction in large growth stocks, while conversely the 2001 and 2008 breaks lead to a significant decline in the allocation to large value stocks.

These results demonstrate significant rotation in the optimal portfolio weights around the time of the breaks identified by our panel break methodology.

7. Conclusion

We present new evidence of instability in the mapping from characteristics such as firm size, book-to-market ratio, and return momentum to expected returns, with the market equity risk, size, and value premia undergoing marked reductions over time. The breaks we identify line up closely with major economic shocks, including the oil price shocks in the seventies and the Global Financial Crisis.

We show that individual firms display very different degrees of sensitivity to instability in the risk premium process and use this to form a break risk factor that goes long in the most break-sensitive stocks and shorts the least break-sensitive stocks. This break risk factor obtains similar or even stronger significance than conventional size, value, and momentum factors in Fama-MacBeth regressions.

Our evidence reveals that the impact and lead-lag timing of instability risk vary significantly across firms in different industries and with different size, value, or momentum characteristics. Stocks with poor past returns ("losers"), large market capitalization, and low book-to-market ratios tend to be affected earlier by breaks than stocks with high past returns, small market capitalization, and high book-to-market ratios. Firms in the telecommunication, utility, oil, business equipment and financial sectors are most affected by break risk, while firms in the meals, books, textiles, mining and wholesale industries are least impacted. Similarly, small value stocks are more strongly affected by break risk than large growth stocks as are small stocks with low prior-year returns compared with large stocks with high prior-year returns.

The breaks we uncover are all associated with major economic shocks and financial market distress which thus appear to have a long-lasting impact and give rise to new regimes with significantly altered risk premia. Notably, size and value risk premia are insignificantly different from zero in the period after the Global Financial Crisis. This pattern is quite different from the mechanism in disaster risk models in which risk premia settle back to their historical mean once the disaster probability returns to normal levels. Similarly, compared with long-run-risk models, our results suggest that the risk premium process can be quite stable for long periods of time but is interrupted by large, pervasive shifts triggered by episodes of economic and financial distress. Although these episodes are relatively rare, their long-lasting impact on risk premia means that they have an important effect on investment performance and portfolio choice. In particular, the waning size and value risk premia identified by our analysis implies a significant reduction in the optimal portfolio allocation to small caps and value stocks between the first and last decades in our sample.

References

Adrian, T., Crump, R. K., and Moench, E. (2015). Regression-based estimation of dynamic asset pricing models. *Journal of Financial Economics*, 118(2):211–244.

Alquist, R., Israel, R., and Moskowitz, T. (2018). Fact, fiction, and the size effect. *Journal of Portfolio Management*, 45(1):34–61.

Ang, A. and Kristensen, D. (2012). Testing conditional factor models. *Journal of Financial Economics*, 106(1):132–156.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–78.

Baltagi, B. H., Feng, Q., and Kao, C. (2016). Estimation of heterogeneous panels with structural breaks. *Journal of Econometrics*, 191(1):176–195.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64.

Bekaert, G., Harvey, C. R., and Lumsdaine, R. L. (2002). Dating the integration of world equity markets. *Journal of Financial Economics*, 65(2):203–247.

Berk, J. B., Green, R. C., and Naik, V. (1999). Optimal investment, growth options, and security returns. *Journal of finance*, 54(5):1553–1607.

Berkman, H., Jacobsen, B., and Lee, J. B. (2011). Time-varying rare disaster risk and stock returns. *Journal of Financial Economics*, 101(2):313–332.

Blanchard, O. J. (1993). The vanishing equity premium. In *Finance and the International Economy*. Citeseer.

Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2001). Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk. *Journal of Finance*, 56(1):1–43.

Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531.

Carlson, M., Fisher, A., and Giammarino, R. (2004). Corporate investment and asset price dynamics: Implications for the cross-section of returns. *Journal of Finance*, 59(6):2577–2603.

Carlson, M., Fisher, A., and Giammarino, R. (2006). Corporate investment and asset price dynamics: Implications for seo event studies and long-run performance. *Journal of Finance*, 61(3):1009–1034.

Chib, S. (1995). Marginal likelihood from the gibbs output. *Journal of the American Statistical Association*, 90(432):1313–1321.

Chordia, T., Goyal, A., and Saretto, A. (2017). p-hacking: Evidence from two million trading strategies. *Unpublished manuscript*.

Clark, T. E. and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1):291–311.

Cochrane, J. H. (2011). Presidential address: Discount rates. *Journal of Finance*, 66(4):1047–1108.

Cogley, T. and Sargent, T. J. (2005). Drifts and volatilities: monetary policies and outcomes in the post wwii us. *Review of Economic Dynamics*, 8(2):262–302.

Croce, M. M., Marchuk, T., and Schlag, C. (2019). The leading premium. Technical report, National Bureau of Economic Research.

Daniel, K. and Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics*, 122(2):221–247.

Dybvig, P. H. (1983). An explicit bound on individual assets' deviations from apt pricing in a finite economy. *Journal of Financial Economics*, 12(4):483–496.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Fama, E. F. and French, K. R. (2002). The equity premium. *Journal of Finance*, 57(2):637–659.

Fama, E. F. and French, K. R. (2020). Comparing cross-section and time-series factor models. *Review of Financial Studies*, 33(5):1891–1926.

Fama, E. F. and French, K. R. (2021). The value premium. *Review of Asset Pricing Studies*, 11(1):105–121.

Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636.

Freyberger, J., Neuhierl, A., and Weber, M. (2020). Dissecting characteristics nonparametrically. *Review of Financial Studies*, 33(5):2326–2377.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly journal of economics*, 127(2):645–700.

Gagliardini, P., Ossola, E., and Scaillet, O. (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica*, 84(3):985–1046.

Green, J., Hand, J. R., and Zhang, X. F. (2017). The characteristics that provide independent information about average us monthly stock returns. *Review of Financial Studies*, 30(12):4389–4436.

Green, P. J. (1995). Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732.

Grinblatt, M. and Titman, S. (1983). Factor pricing in a finite economy. *Journal of Financial Economics*, 12(4):497–507.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68.

Hong, H., Torous, W., and Valkanov, R. (2007). Do industries lead stock markets? *Journal of Financial Economics*, 83(2):367–396.

Hou, K. (2007). Industry information diffusion and the lead-lag effect in stock returns. *Review of Financial Studies*, 20(4):1113–1138.

Hou, K., Xue, C., and Zhang, L. (2020). Replicating anomalies. *The Review of Financial Studies*, 33(5):2019–2133.

Huberman, G., Kandel, S., and Stambaugh, R. F. (1987). Mimicking portfolios and exact arbitrage pricing. *Journal of Finance*, 42(1):1–9.

Ilmanen, A., Israel, R., Moskowitz, T. J., Thapar, A. K., and Wang, F. (2019). Do factor premia vary over time? a century of evidence. *SSRN Working Paper*.

Jagannathan, R., McGrattan, E. R., and Scherbina, A. (2001). The declining us equity premium. Technical report, National Bureau of Economic Research.

Jegadeesh, N., Noh, J., Pukthuanthong, K., Roll, R., and Wang, J. (2019). Empirical tests of asset pricing models with individual assets: Resolving the errors-in-variables bias in risk premium estimation. *Journal of Financial Economics*, 133(2):273–298.

Jochmann, M., Koop, G., Leon-Gonzalez, R., and Strachan, R. W. (2013). Stochastic search variable selection in vector error correction models with an application to a model of the uk macroeconomy. *Journal of Applied Econometrics*, 28(1):62–81.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Lettau, M., Ludvigson, S. C., and Wachter, J. A. (2008). The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies*, 21(4):1653–1687.

Lettau, M. and Van Nieuwerburgh, S. (2008). Reconciling the return predictability evidence. *Review of Financial Studies*, 21(4):1607–1652.

Lewellen, J., Nagel, S., and Shanken, J. (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial economics*, 96(2):175–194.

Linnainmaa, J. T. and Roberts, M. R. (2018). The history of the cross-section of stock returns. *Review of Financial Studies*, 31(7):2606–2649.

Lo, A. W. and MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction? *Review of Financial Studies*, 3(2):175–205.

Ludvigson, S. C., Ma, S., and Ng, S. (2021). Uncertainty and business cycles: exogenous impulse or endogenous response? *forthcoming, American Economic Journal: Macroeconomics*.

MacKinlay, A. C. (1995). Multifactor models do not explain deviations from the capm. *Journal of Financial Economics*, 38(1):3–28.

Martin, I. W. (2013). Consumption-based asset pricing with higher cumulants. *Review of Economic Studies*, 80(2):745–773.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

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

Novy-Marx, R. and Velikov, M. (2015). A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1):104–147.

Pástor, L. (2000). Portfolio selection and asset pricing models. *Journal of Finance*, 55(1):179–223.

Pástor, L. and Stambaugh, R. F. (1999). Costs of equity capital and model mispricing. *Journal of Finance*, 54(1):67–121.

Pástor, L. and Stambaugh, R. F. (2001). The equity premium and structural breaks. *Journal of Finance*, 56(4):1207–1239.

Patton, A. J. and Timmermann, A. (2010). Monotonicity in asset returns: New tests with applications to the term structure, the capm, and portfolio sorts. *Journal of Financial Economics*, 98(3):605–625.

Paye, B. S. and Timmermann, A. (2006). Instability of return prediction models. *Journal of Empirical Finance*, 13(3):274–315.

Pesaran, M. H. (2004). General diagnostic tests for cross section dependence in panels. *Unpublished Manuscript.*

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4):967–1012.

Pesaran, M. H. (2015). Testing weak cross-sectional dependence in large panels. *Econometric Reviews*, 34(6-10):1089–1117.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852.

Rapach, D. and Wohar, M. (2006). Structural breaks and predictive regression models of aggregate us stock returns. *Journal of Financial Econometrics*, 4(2):238–274.

Savor, P. and Wilson, M. (2016). Earnings announcements and systematic risk. *Journal of Finance*, 71(1):83–138.

Schwert, G. W. (2003). Anomalies and market efficiency. *Handbook of the Economics of Finance*, 1:939–974.

Shanken, J. (1992). The current state of the arbitrage pricing theory. *Journal of Finance*, 47(4):1569–1574.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Smith, S. (2018a). Noncommon breaks. *Unpublished Manuscript.*

Smith, S. (2018b). Type-s errors, multiple testing, and the factor zoo: A time-varying bayesian approach. *Unpublished Manuscript.*

Smith, S. C. and Timmermann, A. (2021). Break risk. *Review of Financial Studies*, 34(4):2045–2100.

Stock, J. H. and Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, 20(2):147–162.

van Binsbergen, J., Hua, S., and Wachter, J. A. (2020). Is the united states a lucky survivor: A hierarchical bayesian approach.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68(3):987–1035.

Wachter, J. A. and Warusawitharana, M. (2009). Predictable returns and asset allocation: Should a skeptical investor time the market? *Journal of Econometrics*, 148(2):162–178.

Zhang, L. (2005). The value premium. *Journal of Finance*, 60(1):67–103.

Appendix A. Likelihood function

This appendix specifies the likelihood function used to estimate our model. To this end, we introduce some notations. Our panel break approach allows the intercepts, slope coefficients, and variances to shift following a break. Recalling that τ_k refers to the date for the kth break, the duration of the kth regime is denoted $l_k = \tau_k - \tau_{k-1}$ and consists of observations $\tau_{k-1}+1, \ldots, \tau_k$. Let $\alpha_k = (\alpha_{1k}, \ldots, \alpha_{Nk})$, $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_{K+1})$, $\lambda_k = (\lambda_{k,1}, \ldots, \lambda_{k,J})$, $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_{K+1})$, $\sigma_k^2 = (\sigma_{k1}^2, \ldots, \sigma_{kN}^2)$, $\boldsymbol{\sigma^2} = (\sigma_1^2, \ldots, \sigma_{K+1}^2)$ denote the parameters in the individual regimes and collect all parameters in $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\lambda}, \boldsymbol{\sigma^2})$. Finally, let $\boldsymbol{X_{t-1}}$ denote the observations on the J characteristics for the N stocks at time $t-1$ and define $\boldsymbol{X} = (\boldsymbol{X_1}, \ldots, \boldsymbol{X_{T-1}})$. The likelihood of the data can then be written as[44]

$$p(\boldsymbol{r} \mid \boldsymbol{X}, \boldsymbol{\theta}, \tau) = \prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi\sigma_{ik}^2)^{\frac{l_k}{-2}} exp \left[\sum_{t=\tau_{k-1}+1}^{\tau_k} \frac{(r_{it} - \alpha_{ik} - r_{zt} - \lambda_k' X_{it-1})^2}{-2\sigma_{ik}^2} \right]. \tag{A.1}$$

Appendix B. Priors

Next, we provide details of the prior distributions used by our model.

Appendix B.1. Prior on the regime durations

Following Smith and Timmermann (2021), we place a Poisson prior distribution over the regime durations

$$p(l_k \mid \gamma_k) = \frac{\gamma_k^{l_k} e^{-\gamma_k}}{l_k!}, \qquad k = 1, \ldots, K+1, \tag{B.1}$$

in which the Poisson intensity parameter γ_k has a conjugate Gamma prior distribution

$$p(\gamma_k) = \frac{d^c}{\Gamma(c)} \gamma_k^{c-1} e^{-d\gamma_k}, \qquad k = 1, \ldots, K+1. \tag{B.2}$$

[44]For expositional ease we suppress r_{zt} herein.

A prior belief that a break occurs, on average, every 20 years is achieved by setting c=480 and d=2.

Appendix B.2. Priors on regression coefficients

For regimes $k = 1, \ldots, K+1$ and firms $i = 1, \ldots, N$, we specify an inverse gamma prior on the idiosyncratic residual variances

$$p(\sigma_{ik}^2) = \frac{b^a}{\Gamma(a)} {\sigma_{ik}^2}^{-(a+1)} exp\left(-\frac{b}{\sigma_{ik}^2}\right), \tag{B.3}$$

and a Gaussian prior on the intercepts, conditional on the variances

$$p(\alpha_{ik} \mid \sigma_{ik}^2) = 2\pi^{\frac{-1}{2}} (\sigma_{ik}^2)^{\frac{-1}{2}} (\sigma_{\alpha}^2)^{\frac{-1}{2}} exp\left(\frac{\alpha_{ik}^2}{-2\sigma_{ik}^2 \sigma_{\alpha}^2}\right), \tag{B.4}$$

in which a and b are the prior hyperparameters of the residual variance and σ_{α}^2 reflects the prior belief about the degree of mispricing. To achieve a prior residual variance equal to the variance of the return data, the prior hyperparameter a is set equal to 2 and b is set equal to the variance of the return data across all i and t.[45]

Risk premium estimates have a Gaussian distribution. For regimes $k = 1, \ldots, K+1$

$$p(\lambda_k) = \left(2\pi^{\frac{-J}{2}} \mid V_\lambda \mid^{\frac{-1}{2}}\right) exp\left(\frac{\lambda_k' V_\lambda^{-1} \lambda_k}{-2}\right), \tag{B.5}$$

in which $V_\lambda = 1_J \sigma_\lambda^2$.

Multiplying the likelihood function by the priors yields the posterior distribution. Inference is performed on the posterior distribution which is approximated using Markov chain Monte Carlo methods.

[45]For the out-of-sample analysis, for this calculation we use only the return data available at the time the model is estimated to avoid look-ahead bias.

Appendix C. Estimating the model

Model estimation comprises three steps. First, the parameters in regimes $k = 1, \ldots, K+1$ are estimated from their full conditional distributions using a Gibbs step

$$\begin{aligned} \sigma_{ik}^2 \mid \cdot &\sim IG(\tilde{a}_{ik}, \tilde{b}_{ik}), \qquad i = 1, \ldots, N, \\ \alpha_{ik} \mid \cdot &\sim N(\rho_{ik}, s_{ik}^2), \qquad i = 1, \ldots, N, \\ \lambda_k \mid \cdot &\sim N(\mu_k, \Sigma_k), \end{aligned} \tag{C.1}$$

in which

$$\Sigma_k^{-1} = V_\lambda^{-1} + \sum_{t=\tau_{k-1}+1}^{\tau_k} \boldsymbol{X_{t-1}} \boldsymbol{X_{t-1}}', \tag{C.2}$$

$$\mu_k = \Sigma_k \sum_{t=\tau_{k-1}+1}^{\tau_k} \boldsymbol{X_{t-1}} r_t,$$

$$s_{ik}^{-2} = \sigma_\alpha^{-2} + l_k, \qquad i = 1, \ldots, N$$

$$\rho_{ik} = s_{ik}^2 \sum_{t=\tau_{k-1}+1}^{\tau_k} r_{it}, \qquad i = 1, \ldots, N$$

$$\tilde{a}_{ik} = a + l_k/2, \qquad i = 1, \ldots, N$$

$$\tilde{b}_{ik} = \frac{1}{2}\left(2b + \sum_{t=\tau_{k-1}+1}^{\tau_k} r_{it}^2 - \mu_k' \Sigma_k^{-1} \mu_k\right), \qquad i = 1, \ldots, N$$

where r_t denotes the excess stock returns on the N firms at time t. The second and third steps estimate the break locations and number of breaks, respectively, in the same manner as in Smith and Timmermann (2021) but use Equation (C.2) to compute the acceptance probabilities.

Appendix D. Formal definition of breaks

Our model is estimated using a reversible jump Markov chain Monte Carlo algorithm (Green 1995). This approach repeatedly attempts to 'jump' between models with different numbers of breaks. With a sufficient number of iterations, the posterior model probabilities and corresponding break locations are approximated by the proportion of iterations spent at each number and timing of breaks.

We now formally define what constitutes a breakpoint. For each jump, whether to accept the move (and thus introduce a different number of breaks) is determined by a Bayes factor, the preferred Bayesian model comparison method.

Suppose we attempt to jump from K to K^* breaks. The Bayes factor is a likelihood ratio of the model with K^* breaks and the model with K breaks. The posterior probability of model K, M_K, having observed the data (r, X) is

$$\Pr(M_K \mid r, X) = \frac{\Pr(r, X \mid M_K)\Pr(M_K)}{\Pr(r, X)}, \tag{D.1}$$

the elements of which can be approximated using the marginal likelihood approach of Chib (1995).

The probability of accepting the jump from K breaks to K^* breaks is reflected in the Bayes factor

$$BF_{M_K, M_{K^*}} = \frac{\int \Pr(\theta_{K^*} \mid M_{K^*})\Pr(r, X \mid \theta_{K^*}, M_{K^*}) d_{\theta_{K^*}}}{\int \Pr(\theta_K \mid M_K)\Pr(r, X \mid \theta_K, M_K) d_{\theta_K}} = \frac{\Pr(M_{K^*} \mid r, X)\Pr(M_K)}{\Pr(M_K \mid r, X)\Pr(M_{K^*})}. \tag{D.2}$$

Assuming equal prior model probabilities, $\Pr(M_K) = \Pr(M_{K^*})$, the Bayes factor will equal the ratio of posterior probabilities of the respective models.

Two advantages of the Bayes factor approach are, first, that it automatically penalizes model complexity to guard against overfitting, and thus does not rely on ad hoc penalty terms. Second, it does not depend on a single set of parameters as it integrates over all parameters in each model with respect to their priors, thus accounting for parameter uncertainty.

Table 1: Risk premium tests

	Equity	Value	Size	Momentum	Total
Positive risk premium tests (Final regime: 2008-2018)					
Risk premium	4.61	1.28	0.36	3.59	
t-stat	(4.77)	(0.85)	(0.28)	(3.29)	
Single breakpoint tests					
Bayes factor (1991)		179.87			
Risk premium (1950-1991)		3.36%			
Risk premium (1992-2018)		1.53%			
Bayes factor (1981)			162.43		
Risk premium (1950-1981)			4.20%		
Risk premium (1982-2018)			0.65%		
Monotonic relation tests					
	0.18	0.01	0.05	0.23	0.06

Table 1: Risk Premium Tests. The upper panel of this table displays the final regime's risk premium estimates (expressed as annualized percentages) and corresponding t-statistics (in brackets below) from the Bayesian panel break approach when regressing firm-level excess returns on market beta, value, size, and momentum as displayed in Equation (3). The middle panel displays the results of two separate single breakpoint tests: a break in the size premium at 1981 and a break in the value premium at 1991. Bayes factors express the strength of evidence in favor of the break – values greater than 150 represent overwhelming evidence in favor of the break Kass and Raftery (1995). We also report pre- and post-break risk premium estimates. The lower panel displays p-values from Patton and Timmermann (2010)'s Monotonic Relation tests when testing separately whether each of the four factor risk premia – and the total risk premium – monotonically decline across the five regimes identified by the baseline model. p values lower than 0.05 imply significant evidence in favor of monotonically declining risk premia.

Table 2: Cross-sectional Distribution of α Estimates

Regime	Mean	St.dev.	5%	10%	25%	Median	75%	90%	95%	sig.
					All stocks					
1950:01-1972:07	2.67	1.41	-3.49	-1.86	0.47	2.46	4.63	7.85	9.58	0.25
1972:08:-1981:10	-1.42	2.27	-14.22	-7.97	-2.83	-0.03	1.81	3.96	5.74	0.17
1981:11-2001:06	2.52	3.34	-11.39	-4.36	0.50	2.79	5.88	9.62	13.77	0.20
2001:07-2008:10	-0.30	5.09	-20.12	-10.45	-2.19	1.38	3.97	8.33	13.11	0.08
2008:11-2018:06	-0.62	5.67	-20.14	-9.75	-1.74	0.95	3.30	7.77	12.75	0.06
Full sample	0.41	2.87	-6.97	-4.85	-0.71	0.40	2.25	4.57	7.08	0.18
					Larger stocks (Micro-caps excluded)					
1950:01-1972:07	2.59	0.95	-2.82	-1.38	0.55	2.40	4.47	7.03	8.50	0.24
1972:08:-1981:10	-1.02	1.35	-10.98	-6.96	-2.64	-0.04	1.71	3.58	4.79	0.16
1981:11-2001:06	2.71	1.59	-7.55	-3.13	0.63	2.79	5.69	8.75	11.56	0.19
2001:07-2008:10	0.21	2.14	-14.64	-8.44	-1.90	1.38	3.81	7.27	10.18	0.08
2008:11-2018:06	0.03	2.07	-14.66	-7.56	-1.55	0.95	3.13	6.65	9.91	0.06
Full sample	0.68	1.08	-5.89	-3.62	-0.58	0.40	1.95	4.03	6.27	0.17
					Micro-caps					
1950:01-1972:07	4.10	4.77	-27.58	-13.91	-10.58	-6.47	14.42	20.02	31.86	0.33
1972:08:-1981:10	-9.30	8.01	-48.81	-44.99	-29.94	-23.13	10.81	17.40	30.11	0.23
1981:11-2001:06	-1.30	13.22	-64.72	-47.92	-29.50	-20.14	25.33	39.27	50.98	0.36
2001:07-2008:10	-10.61	20.15	-107.81	-83.05	-53.23	-33.65	31.86	51.56	84.80	0.11
2008:11-2018:06	-15.18	21.01	-146.35	-104.29	-54.26	-34.32	30.10	54.47	64.13	0.08
Full sample	-4.75	8.23	-46.20	-35.97	-19.85	-9.53	12.74	19.86	32.95	0.23

Table 2: Cross-sectional distribution of α estimates. The top panel of this table displays, for each of the five regimes, the average posterior mean and the standard deviation of the posterior means, as well as the 5th, 10th, 25th, median, 75th, 90th, and 95th percentiles of the α estimate from our Bayesian panel break approach when regressing firm-level excess returns on market beta, size, value, and momentum as displayed in Equation (3). All values are in annualized percentage terms. The final column reports, for each regime, the proportion of stocks that have α_i estimates that are significantly different from zero at the 5% level using a two-sided test. The final row of each panel displays corresponding results for the full sample using the constant-parameter model. All results use a prior standard deviation of α of 5%. The middle and lower panels report results without micro-caps and for only micro-caps. Micro-caps are defined as stocks with a price less than $3 or a market capitalization below the 20th percentile of the NYSE capitalization.

Table 3: Which Model Parameters are Affected by Breaks? Bayes Factors

	BF_{mean}	BF_{σ}	BF_{λ}	BF_{α}
All breaks	160.23	210.37	172.89	230.44
Jul 1972	138.73	232.50	197.91	174.37
Oct 1981	164.21	239.53	188.66	228.05
Jun 2001	187.60	200.11	132.41	267.54
Oct 2008	211.14	189.74	97.95	215.02

Table 3: Parameters affected by breaks: Bayes factors. This table displays Bayes factors that indicate the strength of evidence in favor of our baseline model relative to each of four restricted models, including models that allow breaks only in (i) mean coefficients, that is, α and λ (corresponding Bayes factor is denoted BF_{mean}), (ii) volatility (BF_{σ}), (iii) risk premia (BF_{λ}), and (iv) α (BF_{α}). Results are displayed for the full sample, that is across all breaks (top row), and for each individual break (rows 2-5). Bayes factors are computed from the marginal likelihood of our baseline model and that of the restricted model. Marginal likelihoods are computed using the method of Chib (1995). The strength of evidence in favor of our baseline model relative to the restricted model is evaluated using the standard thresholds detailed in Kass and Raftery (1995): values greater than 20 indicate strong evidence in favor of the baseline model.

Table 4: Fama-Macbeth Regressions of Returns on Break Risk Factor

Independent variable	Break risk measures				
	(1)	(2)	(3)	(4)	(5)
	Slope coefficients ($\times 10^2$) and (test-statistics)				
BRK	0.64	0.54	0.30	0.51	0.53
	(4.72)	(4.27)	(2.62)	(4.13)	(4.22)
log(B/M)	0.37	0.27	0.25	0.28	0.31
	(5.08)	(5.01)	(4.87)	(5.14)	(5.33)
log(ME)	-0.16	-0.11	-0.11	-0.09	-0.09
	(-3.20)	(-3.20)	(-3.19)	(-3.07)	(-3.07)
$PR1YR$	0.89	0.64	0.59	0.63	0.66
	(3.31)	(3.29)	(3.19)	(3.23)	(3.31)
	Results demeaned by industry				
BRK	0.92	0.67	0.35	0.65	0.66
	(5.50)	(5.14)	(2.99)	(4.95)	(5.08)
log(B/M)	0.41	0.28	0.27	0.29	0.31
	(5.20)	(5.08)	(4.99)	(5.12)	(5.23)
log(ME)	-0.18	-0.11	-0.10	-0.09	-0.07
	(-2.97)	(-3.02)	(-2.97)	(-2.91)	(-2.88)
$PR1YR$	0.44	0.61	0.55	0.63	0.69
	(2.89)	(3.20)	(3.10)	(3.22)	(3.40)

Table 4: Fama-Macbeth regressions of returns on break risk factor. This table displays the coefficients and Newey and West (1987) heteroscedasticity-adjusted test-statistics (in brackets below) from Fama-Macbeth regressions of firms' returns on our break risk factor (BRK). The first measure of the break risk factor (column 1) is computed at each time for each firm as the difference between forecasts produced from the Bayesian panel models with and without breaks using the dividend-price ratio as the predictor. The second measure (column 2) is the root squared difference between these forecasts. The third, fourth, and fifth measures (columns 3-5) are the difference at each point in time between the intercept, slope, and volatility estimates, respectively, from the panel models with and without breaks. We control for book-to-market [log(B/M)], size [log(ME)] and past performance measured over the previous year ($PR1YR$). The bottom panel presents results from the same analysis in which the break risk measure has been demeaned by industry.

Table 5: Return Performance of Portfolios of Stocks Sorted on Break Sensitivity

Portfolio	r	α	MKT	SMB	HML
			All stocks		
Low	0.22	-0.17	1.07	0.02	0.03
	(2.12)	(-2.09)	(21.10)	(1.51)	(3.95)
2	0.26	-0.08	0.96	0.00	0.11
	(2.06)	(-2.15)	(31.01)	(1.73)	(2.85)
3	0.34	-0.02	0.99	0.01	-0.04
	(2.54)	(-1.52)	(32.87)	(2.04)	(-1.03)
4	0.44	0.07	1.00	0.06	0.12
	(1.98)	(0.99)	(23.69)	(1.34)	(2.08)
High	0.50	0.20	1.07	-0.01	-0.01
	(2.33)	(2.12)	(22.16)	(-1.65)	(-2.64)
High-low	0.28	0.37	0.00	-0.03	-0.04
	(2.33)	(3.05)	(1.06)	(-1.86)	(-1.30)
			Without micro-caps		
Low	0.15	-0.13	0.92	0.00	0.00
	(2.38)	(-2.11)	(18.30)	(1.58)	(2.76)
2	0.22	-0.10	1.01	0.01	0.16
	(2.17)	(-2.40)	(33.02)	(1.74)	(2.92)
3	0.26	-0.04	0.93	0.02	-0.02
	(2.52)	(-1.50)	(31.05)	(2.17)	(-1.21)
4	0.35	0.06	1.06	0.04	0.07
	(2.17)	(1.03)	(22.75)	(1.31)	(1.99)
High	0.39	0.17	0.90	-0.02	-0.01
	(2.38)	(2.10)	(21.85)	(-2.08)	(-3.05)
High-low	0.24	0.30	-0.02	-0.02	-0.01
	(2.30)	(2.92)	(-1.04)	(-1.83)	(-1.40)

Table 5: Return performance of portfolios of stocks sorted on break sensitivity. This table displays monthly value-weighted average excess returns to quintile portfolios sorted according to our break risk factor measured through the difference in the forecasts from the panel models with and without breaks using the dividend-price ratio as the predictor. We also report coefficients and test-statistics (in brackets below) estimated from time-series OLS regressions of quintile portfolio returns on the Fama and French market (MKT), size (SMB) and value (HML) factors. The bottom panel presents results for the same analysis removing all stocks with a price less than $3 or a market capitalisation below the 20th percentile of the NYSE capitalisation.

Table 6: Break Risk Correlations

	mrkt	bm	mve	mom	brk	
			Correlations with Factors			
mrkt	1	0.24	0.32	-0.34	0.28	
bm		1	0.13	-0.42	0.24	
mve			1	-0.15	0.18	
mom				1	-0.26	
brk					1	
			Correlations with Characteristics			
	10%	25%	50%	75%	90%	max
brk	0.01	0.05	0.10	0.25	0.36	0.42
R^2	0.24					

Table 6: Break risk correlations. The upper panel of this table displays the correlations amongst a number of factors, namely, the market (mrkt), book-to-market (bm), size (mve), momentum (mom), and our break risk factor (brk). The middle panel reports the maximum and the 10th, 25th, 50th, 75th, and 90th percentiles of the correlations between our break risk characteristic and the 94 characteristics considered by Green et al. (2017). The lower panel reports the R^2 from a regression of our break risk characteristic on the five characteristics with which it is most strongly correlated, namely, in descending order: idiosyncratic return volatility, return volatility, volatility of liquidity (share turnover), cash flow to debt, and cash flow volatility.

Table 7: Characteristics that are Significant in Different Regimes

	1980:01-1981:10	1981:11-2001:06	2001:07-2008:11	2008:12-2018:06	
			Panel break model		
	beta	beta	beta	beta	
		brk	brk	brk	
	mve	mve	mve		
	bm	bm	bm		
			mom1m	mom1m	
	sgr				
	retvol				
	turn				
	baspread				
	aeavol				
	agr				
		rdmve			
		roaq			
		cashpr			
		lgr			
		gma			
		hire			
			herf		
			ps		
			salerec		
			std˙dolvol		
			ear		
			chcsho		
			chatoia		
Total	9	10	12	3	
			Constant-parameter model		
	beta	brk	mve	bm	mom1m
	retvol	baspread	aeavol	agr	rdmve
	roaq	lgr	hire	herf	ear
	chatoia				
Total	16				

Table 7: Characteristics that are significant for the cross-section of expected returns in different regimes. The upper panel of this table reports, for each regime identified by our panel break model, the characteristics that are significant using a t-statistic threshold of three when regressing firm-level excess stock returns on the 94 characteristics of Green et al. (2017) and our break risk factor. The total number of selected characteristics is reported at the bottom of the table. The posterior mode break dates occur at October 1981, July 2001, and November 2008. The characteristic definitions correspond to those in Table A1 except for brk which denotes our break risk factor. The lower panel reports which characteristics are selected from the constant-parameter model using a t-statistic threshold of three.

Table 8: Portfolios Most and Least Affected by Break Risk

Portfolio	Size of break rank	MSFD	Portfolio	Size of break rank	MSFD
		Industries			
Telcm	1	0.0222	Whlsl	25	0.0059
Util	2	0.0169	Mines	26	0.0052
Oil	3	0.0145	Textls	27	0.0045
Buseq	4	0.0141	Books	28	0.0033
Fin	5	0.0139	Meals	29	0.0028
Hlth	6	0.0137	Other	30	0.0022
		Size and book-to-market			
SMALL HiBM	1	0.0528	ME3 LoBM	21	0.0067
ME2 HiBM	2	0.0461	ME4 LoBM	22	0.0059
SMALL BE4	3	0.0399	BIG BE3	23	0.0049
ME2 BE4	4	0.0368	BIG BE2	24	0.0036
SMALL BE3	5	0.0290	BIG LoBM	25	0.0033
		Size and momentum			
SMALL LoPRIOR	1	0.0205	BIG PRIOR2	21	0.0031
ME2 LoPRIOR	2	0.0189	ME3 PRIOR4	22	0.0027
SMALL PRIOR2	3	0.0165	ME4 HiPRIOR	23	0.0025
ME2 PRIOR2	4	0.0148	BIG HiPRIOR	24	0.0018
ME2 PRIOR3	5	0.0141	BIG PRIOR4	25	0.0016

Table 8: Portfolios most and least affected by break risk. This table lists the upper and lower twenty percent of portfolios according to the magnitude of the total impact of breaks on their respective return forecasts (with 1 denoting the largest impact) for each of our three test assets. This magnitude is captured by the mean squared forecast difference ('MSFD') between panel models with and without breaks. We report results for industry portfolios (top panel), 5×5 portfolios sorted on size and book-to-market (middle), and 5×5 portfolios sorted on size and momentum (bottom).

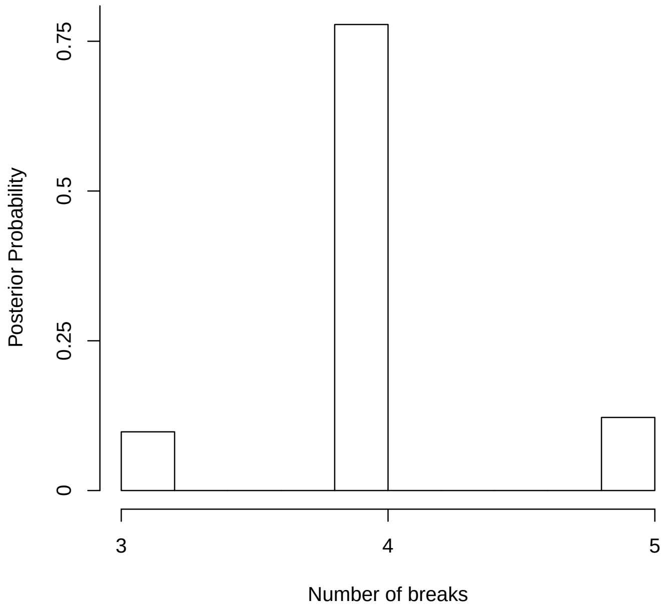

(a) Posterior model probabilities

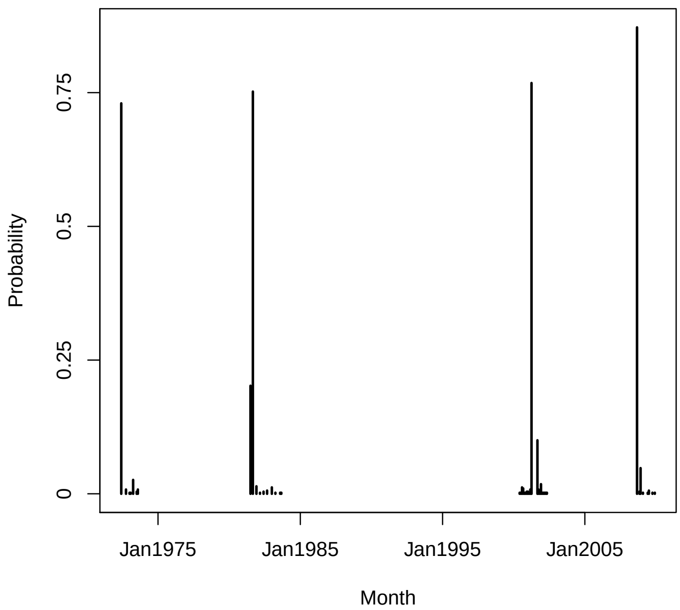

(b) Posterior break location probabilities

Figure 1: This figure displays the posterior distribution of (i) the number of breaks and (ii) break locations estimated from our Bayesian panel break model when regressing firm-level excess stock returns on lagged market beta, size, value, and momentum as displayed in Equation (3).

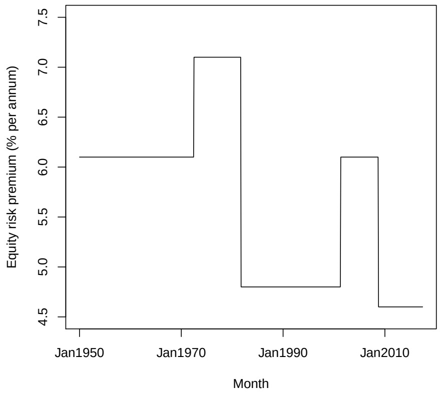

(a) Equity premium (4-factor)

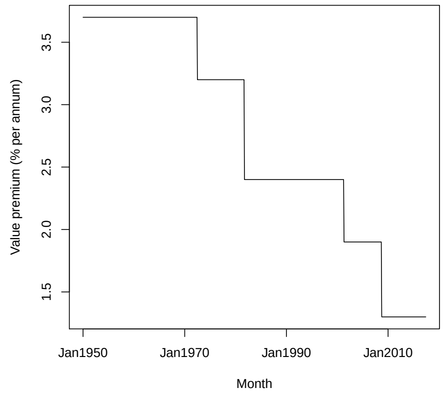

(b) Value premium

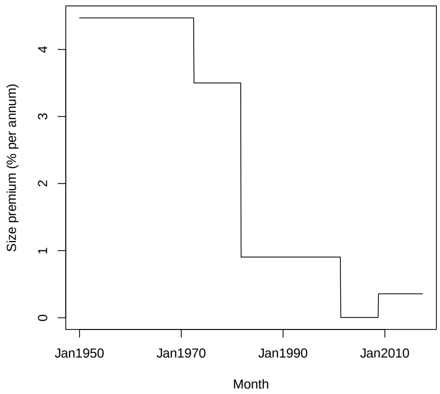

(c) Size premium

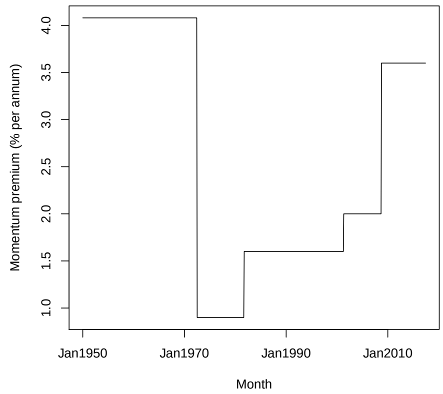

(d) Momentum premium

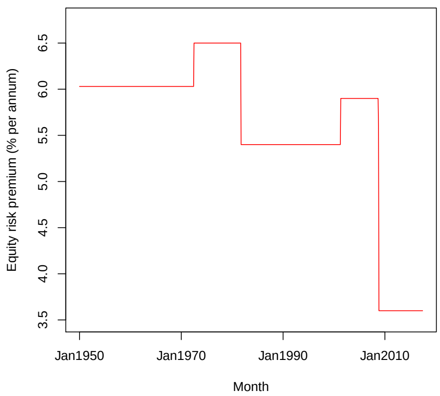

(e) Equity premium (CAPM)

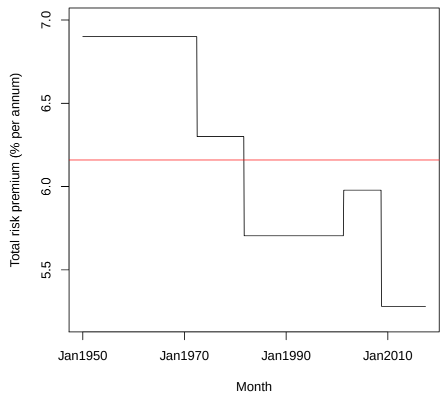

(f) Total risk premium

Figure 2: The top and middle panels of this figure graph the posterior mean estimates of time-varying risk premia from our Bayesian panel break model when regressing firm-level excess stock returns on lagged market beta, size, value, and momentum in a multivariate regression as displayed in Equation (3). The lower left panel shows the time-varying risk premium from a corresponding CAPM panel break regression that only includes market betas as regressors. The lower right panel graphs the total risk premium estimated from the four-factor model with (black line) and without (red line) breaks.

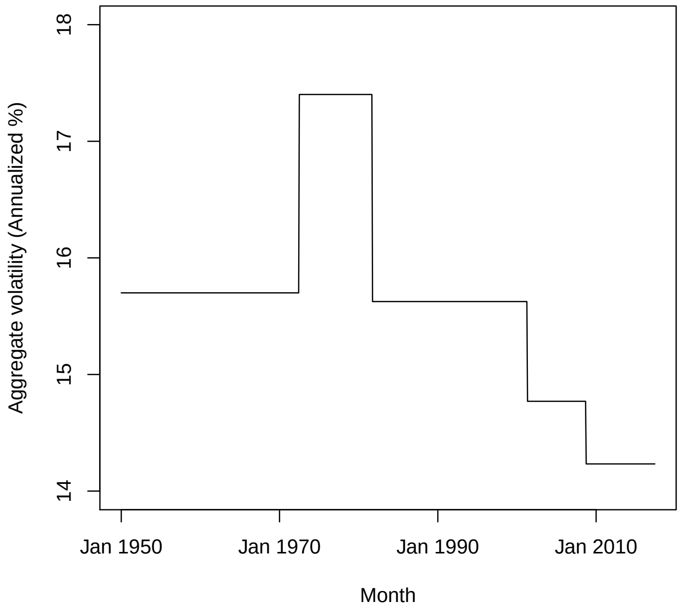

(a) Aggregate volatility

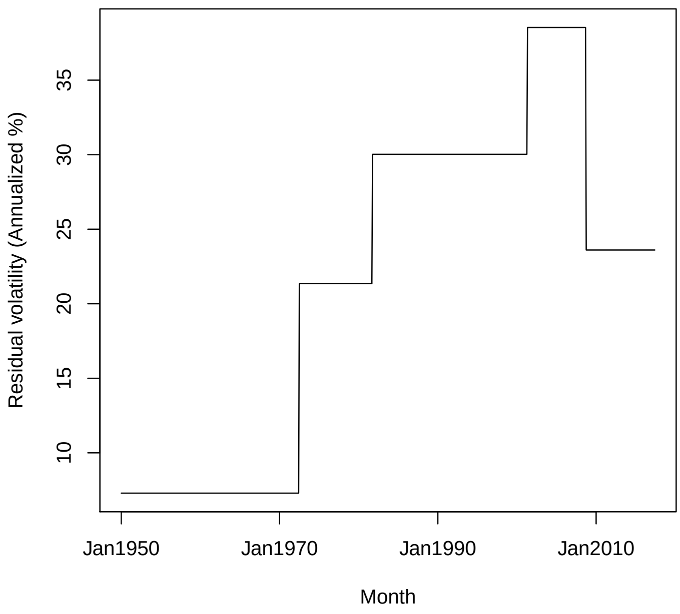

(b) Idiosyncratic volatility

Figure 3: The top panel of this figure graphs the aggregate volatility estimates from our Bayesian panel break approach when regressing firm-level excess stock returns on lagged market beta, size, value, and momentum as displayed in Equation (3). The aggregate volatility is estimated as the standard deviation of R_{zt} in each regime, expressed as an annualized percentage. The lower panel graphs the value-weighted average of firm-level posterior mean residual volatility estimates (expressed as an annualized percentage) from the same model.

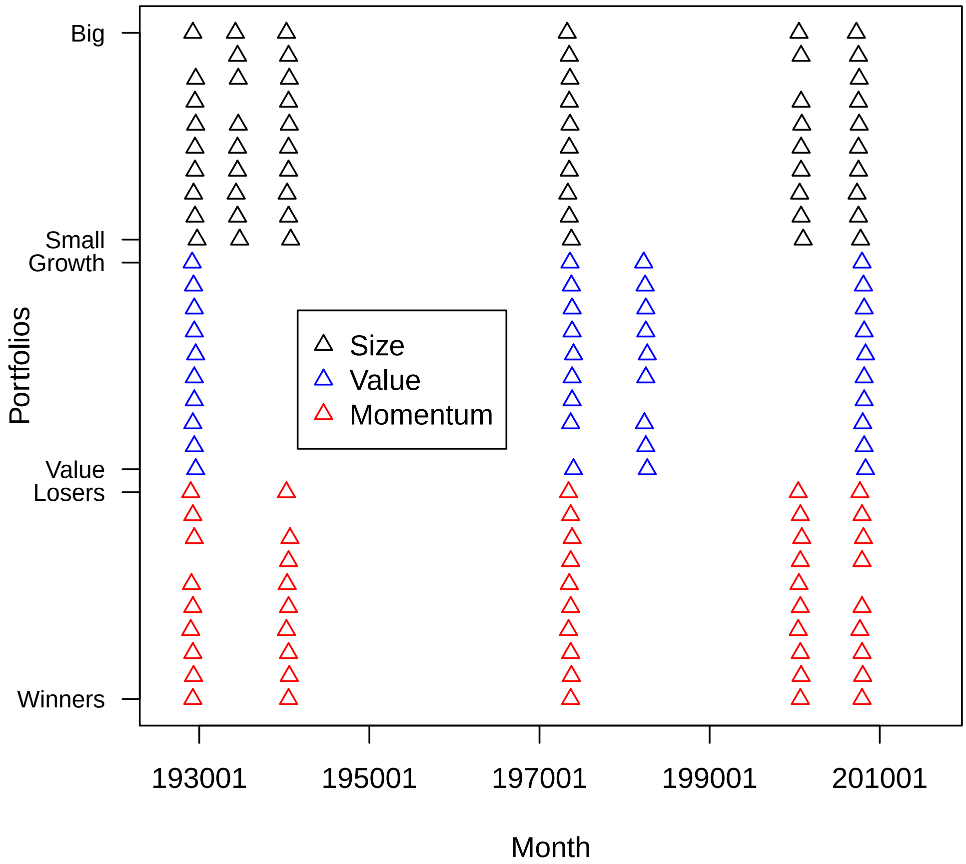

Figure 4: This figure displays the posterior mode break dates estimated from the Bayesian panel break model when regressing the excess returns on 30 portfolios on the lagged aggregate dividend-price ratio as displayed in Equation (10) by applying the methodology developed by Smith (2018a) that allows for any subset of series in the cross-section to be hit by breaks. The 30 portfolios include 10 univariate sorts on each of size (black triangles), book-to-market (blue), and momentum (red).

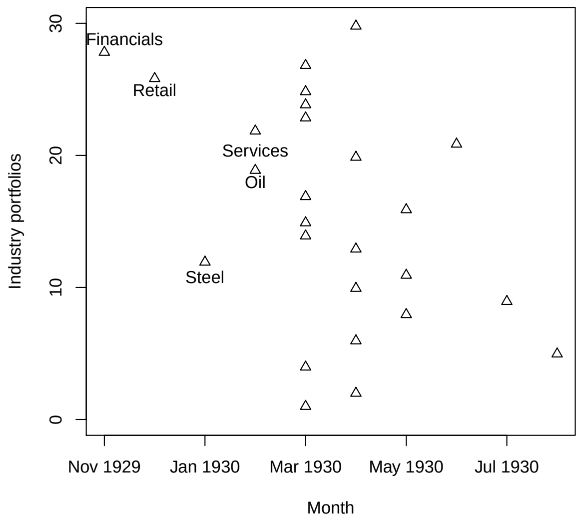

(a) Wall Street Crash

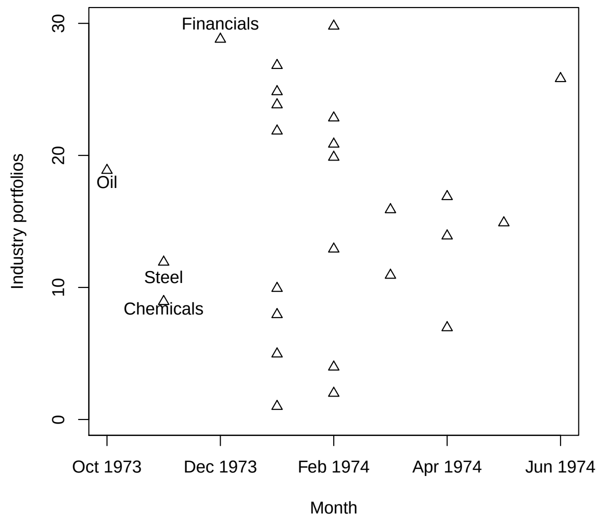

(b) Oil price shock

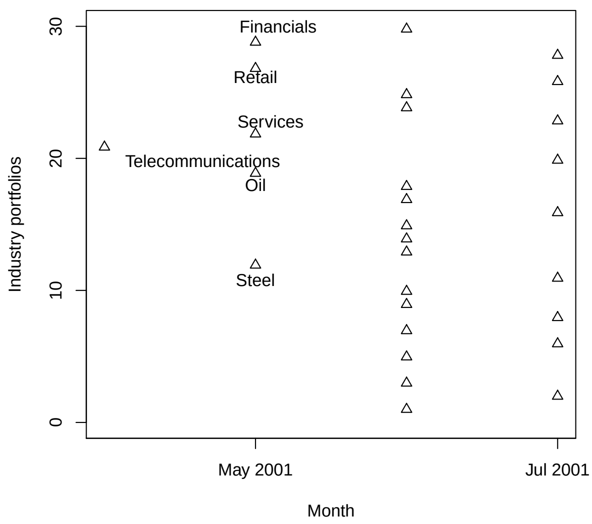

(c) Dotcom bubble

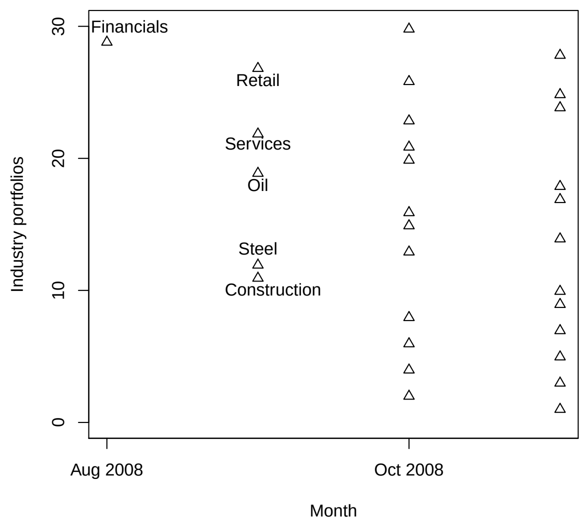

(d) Global Financial Crisis

Figure 5: This figure displays the posterior mode break dates estimated from the panel break model when regressing excess returns on 30 industry portfolios on the lagged aggregate dividend-price ratio by applying the methodology developed by Smith (2018a) that allows for any subset of series in the cross-section to be hit by breaks at different times as displayed in Equation (11). Industry portfolio orderings follow Ken French thus portfolio 1 is Food and 30 is Other. We display the timing for four of the most economically interesting break dates: 1929, 1973, 2001, and 2008.

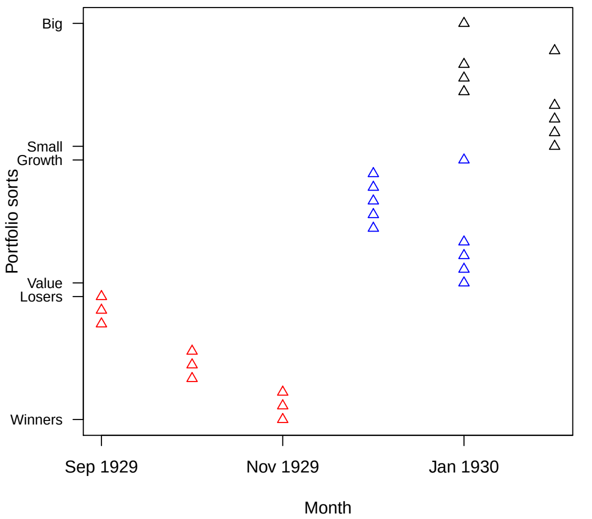

(a) Wall Street Crash

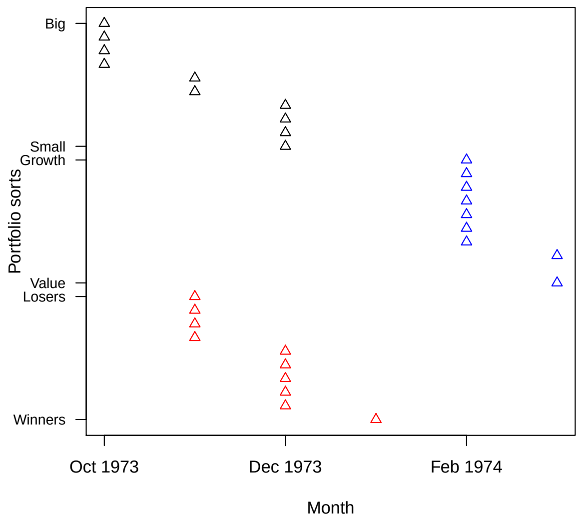

(b) Oil price shock

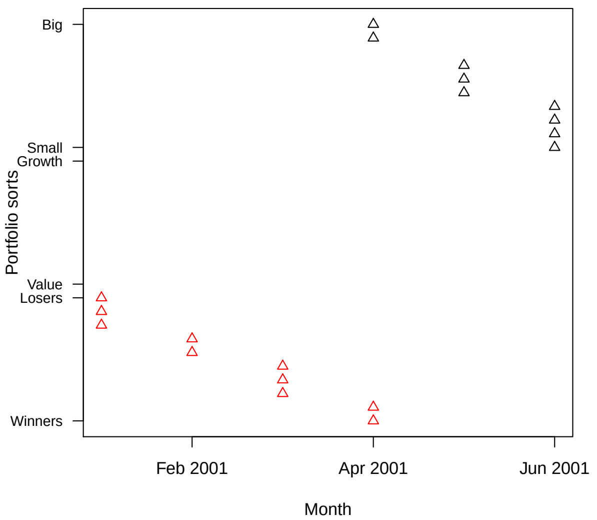

(c) Dotcom bubble

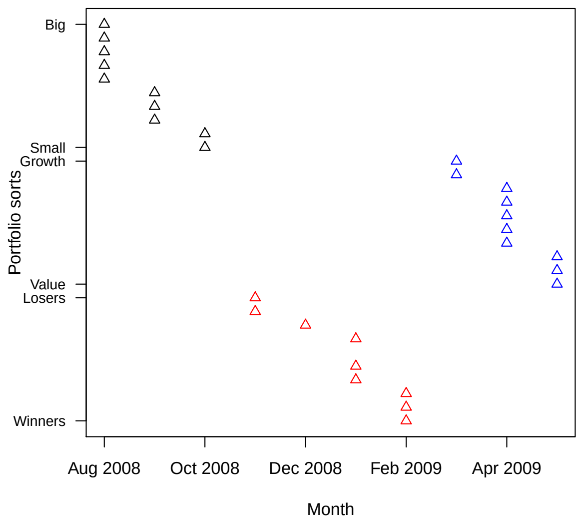

(d) Global Financial Crisis

Figure 6: This figure displays the posterior mode break dates estimated from the panel break model when regressing excess returns on 30 style portfolios – 10 univariate sorts on each of size (black triangles), book-to-market (blue), and momentum (red) sourced from Ken French's website – on the lagged aggregate dividend-price ratio by applying the methodology developed by Smith (2018a) that allows for any subset of series in the cross-section to be hit by breaks and at different times as displayed in Equation (11). We display the timing for four of the most economically interesting break dates: 1929, 1973, 2001, and 2008.

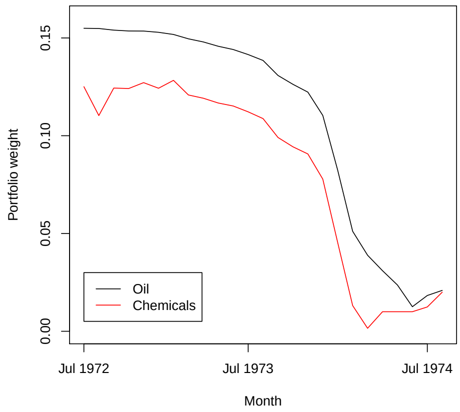

(a) Oil price shock

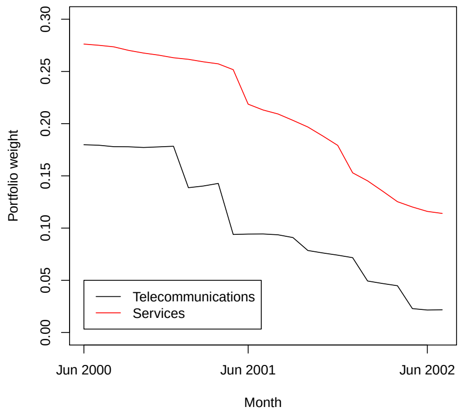

(b) Dotcom bubble

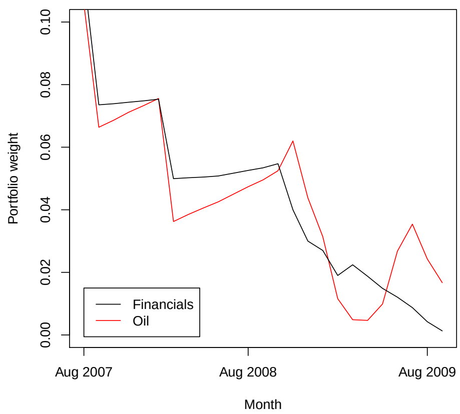

(c) Global financial crisis

Figure 7: This figure displays portfolio weights for a subset of industries around the break dates identified in 1973 (top window), 2001 (middle), and 2008 (lower). Specifically, we graph the 36-month trailing moving average of real-time monthly portfolio weights that are allocated between the 30 industries in the multi-asset portfolio. Allocations are generated from recursively estimating the panel breakpoint model specification in Equation (6) using only data available at the time each forecast is made. Allocations across the 30 portfolios are constrained such that they sum to one and any short selling or leverage is precluded. We display results for the industries whose portfolio allocations are most affected by each break.

Table A1: Firm Characteristic Acronyms and Definitions

Acronym	Definition	Acronym	Definition
absacc	Absolute accruals	***mom1m***	1-month momentum
acc	Working captial accruals	***mom36m***	36-month momentum
aeavol	Abnormal earnings announcement volume	***ms***	Financial statement score
age	no. years since first Compustat coverage	***mve***	Size
agr	Asset growth	***mve_ia***	Industry-adjusted size
baspread	Bid-ask spread	***nanalyst***	Number of analysts covering stocks
beta	Beta	***nincr***	Number of earnings increases
bm	Book-to-market	***operprof***	Operating profitability
bm_ia	Industry-adjusted book-to-market	***orgcap***	Organisational capital
cash	Cash holdings	***pchcapx_ia***	Industry-adjusted Δ% in capital exps.
cashdebt	Cash flow to debt	***pchcurrat***	Δ% in current ratio
cashpr	Cash productivity	***pchdepr***	Δ% in depreciation
cfp	Cash-flow-to-price ratio	***pchgm_pchsale***	Δ% in gross margin - Δ% in sales
cfp_ia	Industry-adjusted cash-flow-to-price ratio	***pchsale_pchinvt***	Δ% in sales - Δ% in inventory
chatoia	Industry-adjusted Δ in asset turnover	***pchsale_pchrect***	Δ% in sales - Δ% in A/R
chcsho	Δ in shares outstanding	***pchsale_pchxsga***	Δ% in sales - Δ% in SG&A
chempia	Industry-adjusted change in employees	***pchsaleinv***	Δ% sales-to-inventory
chfeps	Δ in forecasted EPS	***pctacc***	Percent accruals
chinv	Δ in inventory	***pricedelay***	Price delay
chmom	Δ in 6-month momentum	***ps***	Financial statements score
chnanalyst	Δ in number of analysts	***rd***	*R&D* increase
chpmia	Industry-adjusted Δ in profit margin	***rd_mve***	*R&D* to market capitalisation
chtx	Δ in tax expense	***rd_sale***	*R&D* to sales
cinvest	Corporate investment	***realestate***	Real estate holdings
convind	Convertible debt indicator	***retvol***	Return volatility
currat	Current ratio	***roaq***	Return on assets
depr	Depreciation/*PP&E*	***roavol***	Earnings on volatility
disp	Dispersion in forecasted EPS	***roeq***	Return on equity
divi	Dividend initiation	***roic***	Return on invested capital
divo	Dividend omission	***rsup***	Revenue surprise
dy	Dividend to price	***salecash***	Sales to cash
ear	Earnings to announcement return	***saleinv***	Sales to inventory
egr	Growth in common shareholder equity	***salerec***	Sales to receivables
ep	Earnings to price	***secured***	Secured debt
fgr5yr	Forecasted growth in 5-year EPS	***securedind***	Secured debt indicator
gma	Gross profitability	***sfe***	Scaled earnings forecast
grCAPX	Growth in capital expenditures	***sgr***	Sales growth
gr1tnoa	Growth in long-term net operating assets	***sin***	Sin stocks
herf	Industry sales concentration	***sp***	Sales to price
hire	Employee growth rate	***std_dolvol***	Volatility of liquidity ($ trading volume)
idiovol	Idiosyncratic return volatility	***std_turn***	Volatility of liquidity (share turnover)
ill	Illiquidity	***stdcf***	Cash flow volatility
indmom	Industry momentum	***sue***	Unexpected quarterly earnings
invest	Capital expenditures	***tang***	Debt capacity / firm tangibility
IPO	New equity issue	***tb***	Tax income to book income
lev	Leverage	***turn***	Share turnover
mom12m	12-month momentum	***zerotrade***	Zero trading days

Table A1: Firm characteristic acronyms and definitions. This table provides acronyms and definitions for the 94 firm characteristics considered in our study, and corresponds to Table 1 of Green et al. (2017).

Table A2: Allocations across portfolio sorts

Portfolio	Brk					Hist avg				
					Industries					
food	0.00					0.01				
beer	0.16					0.23				
smoke	0.15					0.09				
books	0.02					0.00				
hlth	0.00					0.06				
chems	0.06					0.14				
elceq	0.01					0.02				
autos	0.00					0.07				
oil	0.06					0.04				
telcm	0.06					0.03				
servs	0.34					0.15				
buseq	0.07					0.11				
paper	0.00					0.02				
fin	0.06					0.00				
					Size and book-to-market					
	SMALL	ME2	ME3	ME4	BIG	SMALL	ME2	ME3	ME4	BIG
LoBM	0.04	0.01	0.00	0.00	0.00	0.03	0.03	0.03	0.03	0.02
BE2	0.03	0.02	0.02	0.00	0.00	0.02	0.04	0.04	0.03	0.02
BE3	0.07	0.04	0.04	0.02	0.00	0.05	0.04	0.04	0.04	0.03
BE4	0.07	0.07	0.05	0.03	0.02	0.05	0.05	0.05	0.04	0.03
HiBM	0.14	0.13	0.10	0.06	0.04	0.07	0.06	0.05	0.05	0.05
					Size and momentum					
	SMALL	ME2	ME3	ME4	BIG	SMALL	ME2	ME3	ME4	BIG
LoPRIOR	0.11	0.13	0.08	0.05	0.05	0.05	0.02	0.01	0.01	0.00
PRIOR2	0.09	0.07	0.05	0.05	0.02	0.07	0.04	0.04	0.02	0.01
PRIOR3	0.04	0.03	0.03	0.02	0.02	0.07	0.04	0.03	0.03	0.02
PRIOR4	0.05	0.03	0.02	0.01	0.00	0.07	0.05	0.04	0.04	0.03
HiPRIOR	0.02	0.01	0.02	0.00	0.00	0.08	0.06	0.06	0.06	0.04

Table A2: Allocations across portfolio sorts. The top panel of this table reports the weight allocations, averaged across the out-of-sample period, to the 30 industry portfolios. We display allocations obtained from our panel break model (Brk) model displayed in Equation (6) and the prevailing mean (Hist avg). Industries that are assigned less than 0.01 weight by both models are omitted. The middle panel displays the allocations across the 25 portfolios sorted on size and book-to-market. The lower panel displays the allocations across the 25 portfolios sorted on size and momentum.

Table A3: Portfolio Allocations Around Breaks

Portfolio	1973		2001		2008	
	pre	post	pre	post	pre	post
Industries						
fin	0.04	0.02	0.07	0.03	0.07	0.02
servs	0.24	0.27	0.28	0.16	0.18	0.12
telcm	0.07	0.08	0.16	0.03	0.05	0.05
oil	0.14	0.01	0.04	0.05	0.08	0.02
chems	0.09	0.02	0.01	0.02	0.03	0.00
Size and momentum						
BIGLoPRIOR	0.07	0.02	0.01	0.00	0.03	0.00
BIGPRIOR2	0.00	0.00	0.05	0.02	0.04	0.00
BIGPRIOR3	0.04	0.02	0.08	0.03	0.05	0.00
ME4LoPRIOR	0.11	0.07	0.23	0.12	0.13	0.01
ME4PRIOR2	0.14	0.01	0.06	0.02	0.04	0.01
Size and book-to-market						
BIGLoBM	0.03	0.02	0.04	0.01	0.01	0.00
BIGBE2	0.00	0.00	0.05	0.02	0.01	0.00
BIGBE3	0.01	0.00	0.02	0.00	0.04	0.02
BIGBE4	0.08	0.08	0.14	0.05	0.08	0.03
ME4BE2	0.19	0.04	0.08	0.03	0.03	0.00

Table A3: Portfolio allocations pre- and post-breaks. This table displays real time allocations to various portfolio sorts averaged across the 36 months before and after the full sample posterior mode break dates in 1973, 2001, and 2008. Allocations are generated from recursively estimating the panel breakpoint model specification in Equation (6) using only data available at the time each forecast is made. Forecasts are generated separately for the three test assets: 30 industry portfolios (top panel), 5 $\times$ 5 sorts on size and momentum (middle), and 5 $\times$ 5 sorts on size and book-to-market (bottom). For each of the three test assets, allocations across the 30 (or 25) portfolios are constrained such that they sum to one and any short selling or leverage is precluded. For each of the three test assets, we report results for the five portfolios whose allocations are most affected by the breaks in 1973, 2001, and 2008.