Break Risk

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

[a] *USC Dornsife INET, Department of Economics, USC, 3620 South Vermont Ave., CA, 90089-0253, USA*
[b] *University of California, San Diego, La Jolla, CA 92093-0553, USA*

October 27, 2018

Abstract

We propose a new approach to forecasting stock returns in the presence of breaks that simultaneously affect a large cross-section of stocks. Exploiting information in the cross-section enables us to detect breaks in return prediction models with little delay and to generate out-of-sample return forecasts that are significantly more accurate than those from existing approaches. Moreover, we find that firms whose equity risk premium processes are most affected by breaks earn significantly higher average returns than firms with lower break sensitivity, suggesting the existence of a break risk factor that is priced in the cross-section.

Keywords: Predictability of stock returns, break risk, cross-sectional variation in returns, panel data, structural breaks, Bayesian analysis, cash flow growth

JEL classifications: G10, C11, C15

Email addresses: simonsmi@usc.edu (Simon C. Smith), atimmermann@ucsd.edu (Allan Timmermann)

The suggestions of Michael Brennan, Wayne Ferson, Bryan Kelly, Arthur Korteweg, Ross Valkanov, Dacheng Xiu, Nancy Xu (discussant at the SFS Cavalcade) and seminar participants at Arizona State University, Rutgers University, UCSD, USC, SFS Cavalcade NA 2018, SoFiE 2018, and the 2018 SoFiE Summer School have been helpful. Any remaining errors are our own.

1. Introduction

Attempts to forecast stock market returns are plagued by instability in the underlying prediction models as evidenced by a large empirical literature. For example, Pástor and Stambaugh (2001) identify multiple breaks in a model linking equity risk premiums to changes in stock market volatility. Similarly, Lettau and Van Nieuwerburgh (2008), Pettenuzzo and Timmermann (2011), Dangl and Halling (2012) and Johannes et al. (2014) find evidence of unstable parameters in the relation between stock market returns and the lagged dividend-price ratio.[1]

Such instability is plausibly a defining feature of return predictability. Indeed, we would expect predictable patterns to 'self-destruct' as investors attempt to exploit them. Schwert (2003), Green et al. (2011), and McLean and Pontiff (2016) test this idea and find evidence that abnormal returns tend to diminish after they become public knowledge. This mechanism is less relevant to the extent that return predictability reflects a time-varying risk premium. However, even in this case, changes in institutions, regulations, and public policy can shift the correlation between observable predictor variables and the underlying risk factors and cause model instability. For example, firms may shift away from paying dividends towards repurchasing shares if taxes on dividends rise, leading to changes in the relation between dividend yields and future stock returns.

Model instability poses severe challenges to attempts at successfully predicting stock market returns. Using the full historical sample to estimate the parameters of a return forecasting model is not an attractive option if the parameters change over time since the resulting estimates may be severely biased. Conversely, using a shorter window of time (possibly after a break has occurred) leads to larger estimation error and less accurate forecasts.

Strategies for modeling the dynamics in the parameters of the return prediction model face two key challenges as pointed out by Lettau and Van Nieuwerburgh (2008). First, investors may have difficulty detecting breaks in real time. Second, and equally importantly, if a break is detected with little delay, only few observations from the current regime are available to estimate the parameters of the forecasting model, potentially leading to volatile and inaccurate return forecasts. Overcoming

[1] Paye and Timmermann (2006) and Rapach and Wohar (2006) undertake a series of econometric tests for model instability and find significant evidence of breaks in the relation between aggregate stock market returns and a variety of predictor variables proposed in the finance literature.

these challenges has proven difficult. Indeed, in their empirical analysis Lettau and Van Nieuwerburgh (2008) find that regime shifts in the dividend-price ratio cannot be exploited to improve out-of-sample forecasts of stock returns.

In this paper we propose an approach that addresses each of these concerns, in the process uncovering new insights into the sources of model instability and its economic consequences. We address the first challenge (slow detection of breaks) by exploiting information in the cross-section of stock returns, enabling breaks to be detected relatively quickly in real time. We address the second challenge (imprecise model estimates) by adopting a Bayesian approach that uses economically motivated priors to shrink the parameters towards sensible values that rule out shifts that are implausibly large in an economic sense. Specifically, following Pástor and Stambaugh (1999), we specify a prior on the intercept of the return equation which does not imply implausible Sharpe ratios. Moreover, following Wachter and Warusawitharana (2009) the prior on the slope coefficient of the predictor is centered on zero with a relatively tight variance implying that investors are skeptical about the existence of predictability. If a break has been recently detected, the few data points from the current regime will be unable to shift the slope estimate far from zero. However, as the length of the regime increases, the degree of shrinkage towards zero is reduced.

The key identifying assumption in our analysis is that the timing of breaks is relatively homogeneous across stocks. This assumption allows us to exploit the benefits from pooling cross-sectional and time-series information. To the extent that information dissemination across different segments of the market is relatively efficient, we would expect that the power of different predictor variables should carry over from the aggregate stock market to individual stocks or portfolios. This suggests that instability in return prediction models can be more effectively detected and estimated in the context of a panel that pools return information across multiple stock portfolios. For example, if a predictor ceases to predict returns on the aggregate stock market portfolio, we would expect to find a similar effect on industry portfolios at approximately the same time. Exploring the simultaneous timing of breaks may allow us to both increase our ability to detect breaks and accurately determine their timing.[2]

While we assume that any breaks affect all stocks at the same time, we allow the intercept, slope, and variance parameters to differ across stocks, thus capturing any

[2] Bekaert et al. (2002) estimate a single common structural break in Vector Autoregressive models to date world equity market integration. They also find that using multiple series reduces the confidence interval around the estimated break date.

heterogeneity in the equity premium and volatility characteristics of individual stocks. Exploiting cross-sectional information to estimate shifts in model parameters turns out to be crucial to our ability to detect breaks in real time and generate forecasts that use information since the most recent break.[3]

Our main analysis focuses on a return prediction model that uses the lagged dividend-price ratio as a predictor variable. We jointly model predictability on 30 industry portfolios using monthly returns data over the 90-year period 1926-2015. Market forecasts can then be constructed as a weighted sum of the individual industry forecasts. Empirically, we find evidence of ten breaks corresponding to a little more than one break on average per decade. Our approach also identifies secular shifts in return volatility.[4]

To help frame the question addressed in this paper, consider an investor who was using the dividend-price ratio to predict stock returns during the financial crisis in 2008-2009. As the crisis grew deeper, investors would plausibly have questioned whether the ability of the dividend-price ratio to predict future returns had deteriorated. Such concerns would have been well founded. Figure 1a plots month-by-month snapshots of the estimated (posterior) probability that a new break occurred in a panel return forecasting model that uses the dividend-price ratio. The likelihood that a break has occurred increases smoothly from the end of 2007 to the fall of 2008 before stabilizing in early 2009. This increase in the likelihood that a break has occurred has an important effect on the estimated slope coefficient of the dividend-price ratio (shown in Figure 1b) which declines from 0.25 prior to the crisis to 0.08 in early 2009. This example shows how, in real time, our approach would have detected the reduced predictability of stock returns from the dividend-price ratio and, accordingly, have adjusted the sensitivity of the forecasts to this predictor variable.

Following earlier studies such as Campbell and Thompson (2008), Goyal and Welch (2008) and Rapach et al. (2010), we assess the predictive accuracy of our return forecasts using a variety of statistical and economic performance measures. For the market portfolio, we find that the return forecasts from the panel break model are significantly more accurate than those produced by the historical average (Goyal and Welch 2008), a time-series model, or a panel model with no breaks. Specifically, our

[3]Polk et al. (2006), Hjalmarsson (2010) and Bollerslev et al. (2018) also consider predictability of stock returns and volatility in a panel setting.

[4]Conventional approaches to model time-varying volatility tend to capture more short-lived periods of volatility clustering. See Andersen et al. (2006) for a review of the literature on volatility forecasting.

panel-break approach generates significantly more accurate out-of-sample forecasts with improvements in the R^2 value for the market portfolio exceeding 0.5% against all of the three benchmarks.[5] Moreover, an out-of-sample asset allocation analysis for a modestly risk averse investor with mean-variance utility suggests that the return forecasts from the panel break model generate certainty equivalent returns around 2% per annum relative to the benchmarks.

Having identified breaks as being important to the return processes for a set of industry portfolios, we next explore whether break risk exposure has implications for the cross-section of stock returns. In particular, we compute the break sensitivity of individual stocks' returns by comparing return forecasts from (panel) models with and without breaks. Firms whose risk premium process is most strongly affected by breaks are found to earn significantly higher returns than firms whose risk premia are less sensitive to breaks. This finding is robust to adjusting for exposure to market, size and value risk factors and the effect is economically large as average risk-adjusted returns of stocks with high break sensitivity exceed those of stocks with low break sensitivity by more than four percent per annum. Thus, we find strong evidence that exposure to break risk is priced in the cross-section.

Return predictability can arise either from predictability in risk premia or from predictability in cash flow growth. While risk premia are unobservable, we can proxy for cash flows through dividends. We therefore undertake a separate analysis of dividend growth predictability and explore whether any breaks separately identified in the dividend process line up with the breaks found in the returns data. We find that, indeed, the vast majority of breaks in stock returns are preceded by breaks to dividend growth. This suggests that investors' awareness of breaks in the underlying dividend growth process is a driver of breaks in stock market returns.

The remainder of the paper is set out as follows. Section 2 lays out our panel-break approach and compares it to existing methods from the literature on return predictability. Section 3 conducts our empirical analysis and reports evidence of structural breaks. Section 4 evaluates the return forecasts of a set of industry portfolios and the market portfolio. Section 5 conducts a cross-sectional analysis of break risk premia, while Section 6 analyzes breaks in the dividend process. Section 7 performs robustness checks, and Section 8 concludes.

[5] For the 30 industry portfolios we find that our approach generates significantly more accurate forecasts in between 24 and 26 cases measured relative to the three benchmarks without a single case in which our forecasts are significantly worse than the benchmark forecasts.

2. Methodology

This section reviews alternative approaches to capturing instability in return prediction models and introduces our panel data approach to estimating breaks that simultaneously affect multiple return series. Our main specification is a heterogeneous panel model with an unknown number of breaks occurring at unknown times. While we allow the *magnitude* of shifts to parameters to vary across portfolios, we assume that the *timing* of the breaks is common in the cross-section.

Our approach differs from conventional return prediction models in two regards: first, it uses panel data, as opposed to the more conventional single-equation time-series approach used throughout the literature; second, it allows for breaks.

To quantify the importance of each of these differences, we compare our approach to (i) a pure time-series approach that allows for breaks, thus highlighting the importance of using cross-sectional (panel) information; and (ii) a constant-parameter panel model that uses the same cross-sectional information as our approach, allowing us to gauge the importance of allowing for breaks. We explain the basic methodology below. Throughout the analysis, we assume a cross-section of N return series and T time-series observations.[6]

2.1. Portfolio-specific Breaks and Parameters

The most general return prediction model we consider assumes that both the model parameters and breaks are unit-specific and so allows for the maximum degree of flexibility in how the individual return series are modeled. This yields a time-series model which is applied to the cross-section of the N return series on a unit-by-unit basis. Following standard practice in the return predictability literature, we focus on prediction models that include an intercept and a single predictor which can either be specific to each portfolio, X_{it}, or be the same (market-wide) predictor, X_t. In each case, excess returns on the ith asset at time t, r_{it}, is our dependent variable.

Suppose the data generating process is time-varying and subject to an unknown number of stock- or portfolio-specific breaks, K_i, which split the sample into $K_i + 1$ distinct regimes for the ith portfolio. Moreover, let $\tau_i = (\tau_{i1}, \ldots, \tau_{iK_i})$ denote a K_i-vector of breakpoints for the ith series. The time-series model fitted to each return

[6]For a more detailed and formal exposition of the econometric properties of the methods described in this section, see Smith et al. (2017).

series in the cross-section takes the form[7]

$$r_{it} = \mu_{ik_i} + \beta_{ik_i} X_{t-1} + \epsilon_{it}, \qquad t = \tau_{ik_i - 1} + 1, \ldots, \tau_{ik_i}, \qquad k_i = 1, \ldots, K_i + 1, \quad (1)$$

where μ_{ik_i} and β_{ik_i} denote the intercept and slope coefficients in the k_ith regime and the error term is assumed to be Normally distributed $\epsilon_{it} \sim N(0, \sigma^2_{ik_i})$ for $k_i = 1, \ldots, K_i + 1$, and $t = \tau_{ik_i-1} + 1, \ldots, \tau_{ik_i}$.

Following existing studies such as Pástor and Stambaugh (2001), we estimate this break model to each individual return series using the algorithm of Chib (1998).

2.2. Pooled Breaks and Portfolio-specific Parameters

With both portfolio-specific parameters and break dates, the model in equation (1) assumes that each cross-sectional unit is independent. However, increased power in break detection can be achieved by combining information from the cross-section of portfolio returns. Our panel approach therefore estimates breaks by pooling the information from the cross-section to identify the *timing* of the K common breaks, while still estimating the parameters for each individual series

$$r_{it} = \mu_{ik} + \beta_{ik} X_{t-1} + \epsilon_{it}, \qquad i = 1, \ldots, N, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k, \qquad k = 1, \ldots, K+1. \quad (2)$$

Again, we assume that the error-term is Normally distributed with unit-specific variance $\epsilon_{it} \sim N(0, \sigma^2_{ik})$ in the (common) kth regime and $\tau = (\tau_1, \ldots, \tau_{K+1})$ for all i.[8]

Popular model specifications from the finance literature can be obtained as special cases of equation (2). In particular, the historical average or prevailing mean model of Goyal and Welch (2008) is obtained by setting $K = 0$ and omitting X_{t-1}. Similarly, a conventional panel model with no breaks is obtained when $K = 0$.

Stock returns typically exhibit high levels of correlation due to their loadings on common factors. While these correlations do not directly affect our forecasts of expected stock returns, ignoring them can reduce the increased break detection power obtained by using panel data rather than the individual time series of returns (Kim 2011; Baltagi et al. 2016) and so it is important to address this point. Note that a strategy of directly estimating the full covariance matrix of the residuals is infeasible

[7] For convenience we assume that τ_{i0}=0 and $\tau_{iK_i+1} = T$ for all i.

[8] The posterior distribution and estimation of this model are detailed in Appendices A and B, respectively.

since it severely delays break point detection.[9] Instead, we assume that correlations across residual returns are generated by a set of common factors. For example, in the case of a single market return factor, $r_{Mkt,t}$, we have

$$\begin{aligned} r_{it} &= \mu_{ik} + \beta_{ik}' X_{t-1} + \epsilon_{it}, \qquad t = \tau_{k-1}+1,\dots,\tau_k, \qquad k = 1,\dots K+1, \\ \epsilon_{it} &= \gamma_{ik} r_{Mkt,t} + \nu_{it}, \end{aligned} \tag{3}$$

where γ_{ik} denotes the factor loading for the ith asset in regime k, and ν_{it} denotes the idiosyncratic residuals. The theoretical properties of this approach are presented in Smith et al. (2017).[10]

2.3. Out-of-sample Return Forecasts

At each point in time, we generate out-of-sample return forecasts for the $i = 1,\dots,N$ industry portfolios by loading the slope estimate on the predictor variable from the final regime and adding the intercept estimate:

$$\hat{r}_{i,t+1} \mid K = \hat{\mu}_{iK+1} + \hat{\beta}_{iK+1} X_t. \tag{4}$$

Note that this step incorporates uncertainty surrounding the break locations but conditions on the number of breaks K.[11] To handle uncertainty about the number of breaks, let K_{min} and K_{max}, respectively, denote the lowest and highest number of breaks that are assigned a nonzero posterior probability by our estimation procedure. We then apply Bayesian Model Averaging to integrate out uncertainty about K:

[9] To see this, note that a cross-sectional dimension of only $N = 30$ requires estimating 525 parameters in each regime, consisting of $3N = 90$ regression parameters and $N_\rho = (N^2 - N)/2 = 435$ correlations. A regime duration shorter than $525/N \approx 18$ periods would therefore require estimating more parameters than we have observations within that regime. In our empirical application, every single break is detected with a considerably shorter delay than this.

[10] Empirically, capturing correlation in the residuals of the individual asset return series through their exposure to the common market factor appears to work very well. For example, the absolute value of the pairwise correlation averaged across 30 industry portfolios in our industry portfolio application is reduced from 0.74 to 0.13 after accounting for the market factor, while the cross-sectional dependence test statistic of Pesaran (2004) is reduced from 168.84 to 2.39, which is no longer significant at the 1% level.

[11] For simplicity, we do not formally state the Bayesian Model Averaging that is done over the break locations. Avramov (2002) reports that Bayesian Model Averaging improves performance when forecasting stock returns in the presence of model uncertainty.

$$\hat{r}_{i,t+1} = \sum_{K=K_{min}}^{K_{max}} p(K \mid \mathbf{r}, X)\hat{r}_{i,t+1} \mid K \tag{5}$$

in which $\mathbf{r}$ denotes the excess returns on the N portfolios across the T time periods and X denotes the T observations on the predictor.

Next, using a bottom-up approach, we forecast the aggregate market return as the value-weighted average of the underlying N forecasts

$$\hat{r}_{Mkt,t+1} = \sum_{i=1}^{N} w_{it}\hat{r}_{it+1}, \tag{6}$$

where $w_t = (w_{1t}, \ldots, w_{Nt})$ denotes the vector of (predetermined) value weights on the N assets at time t.

Our out-of-sample return forecasts are generated recursively with an initial "warm-up" sample of ten years. Hence, the initial parameters of each model are estimated using data from July 1926 through June 1936 and a forecast is made at June 1936 for July 1936. We then expand the estimation period by one month and estimate the parameters of each model using data from July 1926 through July 1936 and produce a return forecast for August 1936. This process is repeated until finally we estimate the parameters of each model using data from July 1926 through November 2015 and generate the forecast for December 2015.

2.4. Prior Distributions

Our Bayesian methodology combines information in the data transmitted through the likelihood function with prior information. Essentially, we assume conventional conjugate Normal priors over the regression coefficients and inverse gamma priors on the variance parameters within each regime.[12] The hyperparameters that determine the frequency of breaks to the coefficients are set so that a break occurs on average roughly once per decade.

Importantly, we let the key prior parameters be economically motivated. Given evidence of weak return predictability such as Goyal and Welch (2008), we center our prior for β at zero. Moreover, inspired by Wachter and Warusawitharana (2009), we explore an economically motivated prior distribution that allows investors to have

[12] Further details of the shape of the priors are provided in Appendix A.

different views regarding the degree to which excess returns are predictable. In the absence of breaks, if the slope coefficient β on the predictive variable is equal to zero, this implies no predictability, and the predictive regression is simply the 'no predictability' benchmark model, i.e., the historical average. A Bayesian analysis allows many different degrees of predictability reflecting the scepticism of the investor as to whether excess returns are predictable. For instance, if β is normally distributed with zero mean and variance σ_β^2, then setting $\sigma_\beta^2 = 0$ implies a dogmatic prior belief that excess returns are not predictable, while $\sigma_\beta^2 \to \infty$ specifies a diffuse prior over the value of β implying that all degrees of predictability (and hence values of the R^2 from the predictive regression) are equally likely. An intermediate view suggests the investor is sceptical about predictability but does not rule it out entirely.

As noted by Wachter and Warusawitharana (2009) it is undesirable to place a prior directly on β_i since a high variance of the predictor σ_X^2 might lower the prior on β_i whereas a large residual variance σ_i^2 might increase it. To address this point, we scale β_i to account for these two variances, placing instead the prior over this 'normalised beta'

$$\eta_i = \beta_i \frac{\sigma_X}{\sigma_i}. \tag{7}$$

Our prior on η_i is

$$p(\eta_i) \sim N(0, \sigma_\eta^2), \tag{8}$$

which by (7) is equivalent to placing the following prior on β_i

$$p(\beta_i) \sim N\left(0, \frac{\sigma_{\eta_i}^2}{\sigma_X^2}\sigma_i^2\right). \tag{9}$$

We compute σ_X^2 as the empirical variance of the predictor variable over the full sample available at the time the recursive forecast is made.[13]

Linking the prior distribution of β_i to σ_X and σ_i is an attractive feature because it implies that the distribution on the R^2 from the predictive regression is well-defined. In population, for a single risky asset the proportion of the total variance that originates from variation in the predictable component of the return is

$$R_i^2 = \frac{\beta_i^2\sigma_X^2}{\beta_i^2\sigma_X^2 + \sigma_i^2} = \frac{\eta_i^2}{\eta_i^2 + 1}, \qquad i = 1, \ldots, N \tag{10}$$

[13] Computing σ_X^2 using only data available in the most recent regime is less robust due to the possibility of very short regimes.

which implies that no risky asset can have an R_i^2 that is 'too large'.

The informativeness of the prior is determined by σ_η which is constant across all i. We refer to Wachter and Warusawitharana (2009) for a full explanation but provide the main results here for completeness. When $\sigma_\eta = 0$ the investor assigns all probability to an R_i^2 value of zero for all i. Figure 2 displays how investors assign more weight to a positive R_i^2 as σ_η increases. Specifically, when $\sigma_\eta = 0.02$, investors assign 0.0003 probability to R_i^2 values greater than 0.005. When $\sigma_\eta = 0.04$, investors assign 0.075 probability to R_i^2 values greater than 0.005. When $\sigma_\eta = 0.06$, investors assign 0.235 probability to R_i^2 values greater than 0.005. For large values of σ_η, investors assign approximately equal probabilities to all values of R_i^2. Our main empirical analysis considers a moderate degree of predictability by setting $\sigma_\eta = 0.04$ following Wachter and Warusawitharana (2009), but we also explore the robustness of the results when this parameter is adjusted.

It may also be desirable to specify that high Sharpe Ratios are a priori unlikely. A high absolute value of the intercept term combined with a low residual variance would imply a high Sharpe Ratio. In the spirit of Pástor and Stambaugh (1999) we multiply the prior standard deviation of the intercept term σ_μ, by the corresponding estimated residual standard deviation in the kth regime for the ith portfolio σ_{ik}. Because the intercept term has a prior mean of zero, a low residual variance reduces the overall variance of the intercept, thereby making a large absolute intercept value and hence a high Sharpe Ratio improbable. As the residual variance increases, the probability assigned to large absolute intercept values increases accordingly. Following Pástor and Stambaugh (1999), we adopt a moderate prior belief in the empirical analysis by setting the prior intercept variance σ_μ equal to 5%.[14]

3. Empirical Results: Evidence of Breaks

This section introduces our return data and predictor variables and presents empirical evidence on the location and number of breaks identified by our approach.

[14]See also Avdis and Wachter (2017) who report that maximum likelihood estimation that incorporates information about dividends and prices results in an economically meaningful reduction in the equity premium estimate that is more reliable relative to the commonly used sample mean.

3.1. Data

As our dependent variable, we use monthly returns on 30 value-weighted industry portfolios from July 1926 through December 2015 sourced from Kenneth French's website, all computed in excess of a one-month T-bill rate. We also source monthly returns excluding dividends from French's website, and the 5×5 portfolios sorted on size and book-to-market or on size and momentum. We also source monthly aggregate data on the three factors of Fama and French (1993).

Our lead predictor is the aggregate dividend-price ratio using 12-month moving sums of dividends on the $S\&P$ 500, but we also consider predictors such as the one-month Treasury-bill rate, the term spread (the difference between the long term yield on government bonds and the Treasury-bill rate), and the default spread (the yield spread between BAA- and AAA-rated corporate bonds), all sourced from Amit Goyal's website.

3.2. Evidence of Breaks

We first consider the evidence of breaks in the return prediction model. To this end, the top panel in Figure 3 plots the posterior probability distribution for the number of breaks estimated on the full sample of 90 years of data for the model that uses the lagged dividend-price ratio as a predictor. The mode (and mean) for the number of breaks is 10, with approximately 90% of the probability mass distributed between 9 and 10 breaks. These estimates suggest a break occurring roughly once per decade.

The lower panel in Figure 3 plots the posterior probability for the location of the breaks. The timing for most of the breaks appears to be quite well defined with clear spikes in the posterior probabilities in 1929, 1933, 1972, 1998, and 2008. Thus, the break dates coincide with major economic events such as the Great Depression, the oil price shocks of the 1970s, the Asian Financial crisis and the bailout of LTCM, and the financial crisis of 2008. Interestingly, the posterior probability mass is quite disperse during the recent financial crisis, indicating that its effect on different industry portfolios was not confined to a single month but diffused gradually through time. Note also that there are long periods without any evidence of model instability, e.g., the twenty-year period from 1950 to 1970.[15]

[15]The estimated break dates do not change if we include the size and value factors alongside the market factor in our panel regression.

The breakpoints identified by our panel approach are very different from those obtained from the breakpoint algorithm of Chib (1998) applied to the univariate time series of returns on the individual industry portfolios. In fact, for each of the industry portfolios the univariate breakpoint model fails to detect a single break, always favoring the model with zero breaks which receives, on average, 91.21% of the posterior model probability. This suggests that the univariate tests have too weak power to identify breaks off individual return series.[16]

Our heterogeneous panel break model allows the parameters of the equity premium processes for different industry portfolios to be affected more or less severely by breaks and it can be insightful to analyze which industries exhibit the greatest sensitivity to breaks. To explore this point, for each of the industry return series we compute the standard deviation of the estimated intercept, slope and volatility parameters across the 11 regimes. The top panel in Table 1 reports the standard deviation in parameters across regimes (our measure of break sensitivity) for the top and bottom five industries as ranked by the sensitivity of the slope coefficients to breaks. The table shows that, indeed, not all industries are equally affected by breaks. Returns on the oil industry portfolio exhibit the greatest sensitivity to breaks, followed by financials and telecommunication firms. Least sensitive to breaks are the returns on stocks in the services and wholesale industries.

3.3. Evolution in Return Forecasts

Figure 4 shows out-of-sample forecasts of market returns from the heterogeneous panel model with (dotted red line) and without (dashed purple line) breaks and the prevailing mean model (solid black line). The forecasts generated from the prevailing mean model are much smoother than the other ones. Return forecasts from the two panel models display higher volatility than the prevailing mean model and are also quite different from each other, indicating the importance of allowing for breaks.

3.4. Real-time Detection of Breaks

A key challenge when generating return forecasts in a setting that accounts for instability is how quickly the model is able to identify breaks in real time. Severe delays

[16] Pástor and Stambaugh (2001) identify breaks in returns based on assumptions about joint movements in the mean and variance of returns.

in breakpoint detection can lead to poor forecasting performance, particularly if the distance between breaks is relatively short, causing some regimes to be overlooked altogether. Conversely, if shifts to parameter values can be identified with little delay, this opens the possibility of improved forecasting performance. The ability to detect breaks in real time is therefore of central importance to investors seeking to re-allocate their portfolios in a timely manner.

To shed light on this issue, Figure 5 plots the break dates estimated in real time. The real-time breakpoint detection performance of the model with pooled breaks and portfolio-specific parameters works as follows. The initial model is estimated using the first ten years of data. Next, the estimation window is expanded by one month and the model is re-estimated until we reach the end of the sample, recording the break dates identified at each point in time. The vertical line in the figure marks the first period at which the model is estimated given the initial training window of ten years (120 monthly observations) while the 45 degree line (to the right of the vertical line) marks the points at which a break could first be detected, corresponding to a delay of zero. Circles on the graph mark the break dates as estimated in real time with horizontal bands of circles indicating that an initial break date estimate is confirmed to have occurred as subsequent data arrive. The figure is dominated by these bands whose initial points start with only a short delay from the 45 degree line, demonstrating the ability of the procedure to rapidly detect the onset of a break. Conversely, initial break estimates that are not supported by subsequent data, appear as isolated circles outside the horizontal bands and are indicative of "false alarms". There are not too many instances in which the approach detects what subsequently turns out to be spurious breaks.

Lettau and Van Nieuwerburgh (2008) and Viceira (1997) find evidence of instability in time-series predictive regressions of the aggregate market return on the dividend-price ratio. They also find that such instability cannot be exploited to generate more accurate out-of-sample return forecasts because their univariate method is unable to detect breaks in real time. Figure 5 shows that, by incorporating cross-sectional information from multiple return series, our panel break procedure has increased break detection power relative to the time-series approach.

To further highlight this point, Figure 6 plots the number of months before a break was first detected in real time, measured relative to the full-sample (ex-post) estimate of the break date. The majority of breaks in the dividend-price ratio model were detected within five to seven months of their occurrence, with the longest delay being 9 months.

The ability of our panel break approach to identify breaks with relatively little delay stands in marked contrast to the long delays typically associated with breakpoint modeling in the context of univariate time-series.

4. Evaluation of Return Forecasts

This section compares the predictive performance of our heterogeneous panel break model with a univariate time-series break model, a heterogeneous panel model without breaks, and the simple historical average, the latter serving as a 'no predictability' benchmark. We report both statistical and economic measures of forecasting performance, the latter based on how a risk averse mean-variance investor would utilize the forecasts from the different return prediction models.

4.1. Out-of-sample Forecasting Performance

To evaluate the accuracy of the return forecasts, Figure 7 plots the cumulative sum of squared error differences ($CSSED$) obtained by subtracting the sum of squared errors produced by our panel break forecasts from the sum of squared errors generated by each of the benchmark models:

$$CSSED_t = \sum_{\tau=1}^{t}(e^2_{Bmk,\tau} - e^2_{Pbrk,\tau}), \qquad (11)$$

in which $e_{Bmk,\tau}$ and $e_{Pbrk,\tau}$ denote the respective forecast time t errors from the benchmark and our panel break model. Positive and rising values of the CSSED measure represent periods where the panel break model outperforms the benchmark, while negative and declining values suggest that the panel break model is underperforming. Moreover, if the performance of the panel break model measured against the benchmark is dominated by a few observations, this will show up in the form of sudden spikes in these graphs. In contrast, a smooth, upwardsloping graph indicates more stable outperformance of the panel break model measured against the benchmark.

Figure 7 presents plots of the CSSED values for the market portfolio and three representative industries (oil, financials and telecommunications). The plots show that the heterogeneous panel break model consistently outperforms its competitors over the 80-year sample. For the market portfolio (top left hand corner), the CSSED curve for the panel model with breaks measured relative to the prevailing mean model rises throughout the out-of-sample period with no long spells of underperformance. The strong performance against the historical average is particularly impressive given that this benchmark has been found by Goyal and Welch (2008) to be very difficult to beat out-of-sample. A similarly strong performance is seen for the panel break model measured against the panel model without breaks or against the univariate market model that allows for breaks.

Similar improvements in predictive accuracy from the panel break model are seen in the plots for the three industry portfolios displayed in Figure 7. The plots continue to show clear and consistent improvements against the prevailing mean and univariate time-series model while the improvements against the no-break panel model are more concentrated towards the last 15 years of the sample for the oil and telecommunications industries.

We evaluate the forecasting performance of the panel break model relative to the benchmark using the out-of-sample R_i^2 measure of Campbell and Thompson (2008):

$$R^2_{OoS} = 1 - MSPE_{Pbrk}/MSPE_{Bmk}. \tag{12}$$

Here $MSPE_{Pbrk}$ and $MSPE_{Bmk}$ denote the mean squared prediction error (MSPE) for the panel break and benchmark models, respectively. A positive R^2_{OoS} value indicates that the panel break model outperforms the benchmark, while a negative value indicates it underperforms.

Figure 8 plots histograms of the R^2_{OoS} values for each of the thirty industry portfolios and the market portfolio based on comparisons of the forecasting performance of our proposed panel breaks model relative to the three benchmark models. For the 31 portfolios our method outperforms all three benchmarks 29 times. Moreover, many of the R^2_{OoS} values are economically large: Campbell and Thompson (2008) estimate that even an R^2_{OoS} value as small as one-half of one percent on monthly data is economically large for a mean-variance investor with moderate risk aversion.

Table 2 uses the test statistic of Diebold and Mariano (1995) to more formally evaluate the statistical significance of the relative performance of the panel break model against the three benchmarks. The table shows that the panel break model performs significantly better than the benchmark at the 10% level for 25, 24, and 26 of the 30 industry portfolios including the market index compared to the predictive performance of the heterogeneous panel model with no breaks, the prevailing mean, and the time-series break model, respectively. Using the test statistic of Clark and West (2007) which allows for nested models, this outperformance is significant at the

10% level for 27, 26, and 27 of the 31 portfolios relative to the no-break panel model, the prevailing mean model and the univariate time-series model, respectively. Conversely, the panel break model does not underperform relative to these benchmarks at the 10% level for any of the 31 portfolios.

These findings underline that the improvements in predictive accuracy that we observe for the panel break model is not simply a result of expanding the information set from a univariate time-series setting to a panel setup that incorporates cross-sectional information. Conversely, allowing for breaks in a univariate setting also does not produce nearly the same gains in predictive accuracy as the panel break model. Rather, it is the joint effect of using cross-sectional information in a panel setting and allowing the return forecasts to account for breaks that generates improvements in predictive accuracy.

The results also demonstrate that our panel model with breaks has the ability to adapt to breaks, and thus handle model instability, while simultaneously reducing the effect of estimation error which has so far plagued real-time (out-of-sample) return forecasts, see Lettau and Van Nieuwerburgh (2008). Key to the improved forecasting performance is our ability to detect breaks to return prediction models with little delay, combined with our use of economically-motivated priors which dampen the adverse effect of estimation error which tends to greatly reduce the accuracy of return forecasts inside new regimes.

4.1.1. Forecasting Performance in the Aftermath of Breaks

To the extent that pooling cross-sectional information helps the panel break model speed up learning, we would expect forecasting performance to be particularly good in the immediate aftermath of a break, particularly if the break is large in magnitude.

To see if this holds, Figure 9 graphs the cumulative difference in the sum of squared errors as a function of the time since the initial break detection, measured in months, i.e., in break point 'event time'. Specifically we compute the squared forecast errors each month following the detection of each break in the out-of-sample period and then take the mean of the squared errors in each period across the breaks. For example, in the first month following break detection, we average across the squared forecast errors in the period immediately following each of the breaks that are detected over the out-of-sample forecasting window. Our panel-break method outperforms the competing benchmarks by the largest margin in the short period after a break is detected, demonstrating the value from using our panel procedure to detect the onset of a break more quickly in real time.

4.2. Economic Utility from Return Forecasts

In addition to evaluating the statistical performance of forecasts from our panel break model relative to a variety of benchmarks, it is important to evaluate their economic performance. For each of the 31 portfolios (including the market) we therefore compute the utility gain to a mean-variance investor who at each period allocates his portfolio between a single risky asset and risk-free T-bills.[17] In this one-period-ahead forecasting exercise, at time t the mean-variance investor allocates a portion of his portfolio to equities in period $t+1$ based on forecasts of the mean and variance of excess returns denoted $\hat{r}_{t+1}$ and $\hat{\sigma}^2_{t+1}$, both computed using only information available at time t[18]

$$w_{Bmk,t} = \frac{1}{A} \frac{\hat{r}_{t+1}}{\hat{\sigma}^2_{t+1}}, \tag{13}$$

Figure 10 shows the distribution of utility gains across the 31 portfolios that we consider here. Allocations are based on the out-of-sample return forecasts that use the dividend-price ratio as a predictor. We show annualized certainty equivalent (CER) values for the panel breaks model relative to the utility obtained from the three benchmark specifications, i.e., a panel model with no breaks, a prevailing mean model computed for each portfolio, and time-series forecasts fitted separately to the individual industry portfolios, in each case allowing for breaks.

The plots show that the panel break model generates positive CER values for at least 28 of the 31 portfolios, regardless of the benchmark. Moreover, the estimated utility gain from using the panel-break forecasts is generally economically large. For the market portfolio it is 2% per annum when measured against the forecasts from the no-break panel model or the time series model with breaks, and it exceeds 1.5% relative to the prevailing mean model.

If breaks in the model parameters do not strongly affect a particular industry portfolio, it is unlikely that a model that accounts for such breaks can significantly outperform a model that ignores breaks. To see if this holds, Table 3 explores the

[17] A small set of studies that explore the utility gain to a mean-variance investor include Campbell and Thompson (2008), Goyal and Welch (2008) and Rapach et al. (2010).

[18] Following Campbell and Thompson (2008), we use a five year rolling window of monthly stock returns to estimate the variance of stock returns, assume a risk aversion coefficient of $A = 3$ and restrict the portfolio weights to fall between 0% and 150% to rule out short-selling and highly leveraged positions.

relation between the magnitude of the break, as measured by the mean squared difference between the forecasts from the panel models with and without breaks, and the utility gains for that portfolio, again measured using the panel models with and without breaks. The calculations assume a mean-variance investor with a coefficient of risk aversion of three. To keep it simple, we show only results for the upper and lower quartile of industries, ranked by mean squared forecast difference.

We find that those industries for which breaks have the biggest effect on the forecasts (upper quartile) generally lead to higher utility gains both in absolute and relative terms, while industries whose return forecasts are least affected by breaks (lower quartile) are associated with the smallest utility gains.

4.2.1. Industry Allocation Analysis

We next explore the utility gain of a mean-variance investor who each period allocates his wealth between the risk-free rate and a risky portfolio constructed from the 30 industry portfolios (see e.g. Avramov and Wermers (2006) and Banegas et al. (2013)). Let $r_{p,t+1}$ denote the return on the risky portfolio at time $t+1$ in excess of the risk-free rate $r_{f,t+1}$. The return on the risky portfolio is constructed from the returns on the 30 industry portfolios at time $t+1$, r_{t+1}, and the corresponding portfolio value weights w_{t+1}. Using our panel break model at each time t we determine the weight vector ω_t to allocate among the 30 industry portfolios in the next period, i.e., we solve for the ω_t that maximizes the expected utility

$$E[U(r_{p,t+1} \mid A)] = r_{f,t} + \omega_t' \hat{r}_{t+1} - \frac{A}{2}\omega_t' \hat{S}_t \omega_t, \qquad (14)$$

subject to the summability constraint $\sum_{i=1}^{N} \omega_{it} = 1$, and $\omega_{it} \in [0,1]$ for $i = 1, \ldots, N$ to preclude any leverage or short selling of individual industries.[19] The covariance matrix, $\hat{S}_t$, is estimated using the residuals from the return prediction model up to time t. This process is repeated for each time period out-of-sample.

Table 4 shows the time-series average industry allocations under the four different return forecasts.[20] There are some notable differences across the different forecasts. For example, the average allocation to services is only 16% under the historical average

[19]Imposing constraints on the portfolio weights is akin to applying shrinkage on the variance-covariance estimates which can lead to performance improvements in mean-variance analysis. See Jagannathan and Ma (2003) and DeMiguel et al. (2007).

[20]To preserve space, we omit industries that have an allocation of less than 0.01 for every model.

model, but close to 40% under the three other approaches. Conversely, the two panel models allocate substantially more (13% and 24%, for the break and no-break models, respectively) to the smoke industry than the historical average (9%). In turn, the univariate time-series model allocates most (23%) to telecommunications, with the allocation from the panel break model (5%) in a distant second place.

The top panel of Table 5 reports the resulting out-of-sample utility gains from these optimized allocations across industry portfolios. Specifically, the table shows the annualized CER values for the panel break model measured relative to the three benchmarks, in each case using the dividend-price ratio as a predictor (top line). Relative to the historical average forecasts, the CER value of the panel break model is 2.19% per annum. Average gains in the CER value of the panel break model remain large–approximately 2% per annum–when measured against the univariate time-series and no-break panel models.

Furthermore, the improved predictive performance in the immediate aftermath of a break being detected translates into even larger utility gains during these periods. Table 5 also reports utility gains computed using only those time periods that occur within two years of a break first being detected ('After breaks'). The annualized CER value of 3.02% is even higher reflecting the ability of our approach to exploit the rapid detection of breaks for utility gains.

These results suggest that the panel-break forecasts of returns on the individual industry portfolios could have been used out-of-sample to generate economically meaningful improvements over forecasts from the three benchmarks.

5. Is there a Break Risk Factor?

Our panel breakpoint methodology can detect breaks that are pervasive in the sense that they simultaneously affect the returns of multiple stocks or portfolios. We next explore whether these breaks can be characterized as a risk factor that is priced in the cross-section, that is, whether stocks with greater exposure to break risk earn higher returns.

To evaluate this hypothesis, we first create out-of-sample forecasts from the heterogeneous panel models with and without breaks, using monthly data on individual US stocks traded on the NYSE, AMEX or NASDAQ stock exchanges at some point during our sample period from July 1926 through December 2015. We use CRSP data on 7,299 different stocks. For each stock, we use the lagged dividend-price ratio on the market as the predictor. Further, we use an initial warm-up estimation period of ten years so the out-of-sample analysis begins in 1936.[21]

For each stock, i, and each month in the sample, t, we then compute the difference between the forecast from the panel break model ($\hat{r}_{it+1,PBrk}$) and the forecast from the panel model without breaks ($\hat{r}_{it+1,NoBrk}$):

$$BRK_{it+1} = \mid \hat{r}_{it+1,PBrk} - \hat{r}_{it+1,NoBrk} \mid, \qquad i = 1, \ldots N, \qquad t = 120, \ldots, T-1. \quad (15)$$

We expect BRK_{it+1} to be larger for stocks with greater exposure to break risk. We next investigate whether such exposures translate into higher risk premia.

5.1. Fama and MacBeth Regressions

We evaluate the ability of our break risk measure to explain the cross-section of returns by estimating the following cross-sectional regression each month

$$r_{it} = \lambda_{0t} + \lambda_{1t} BRK_{it} + \lambda_{2t}' Z_{it} + e_{it} \qquad (16)$$

in which Z_{it} contains three control variables, namely size (log(ME)), value (log(B/M)) and return performance measured over the previous year excluding the most recent month ($PR1YR$), for the ith stock at time t. Next, following the conventional Fama-MacBeth methodology, we use the time-series estimates of λ_{1t} and λ_{2t} to evalute the mean and standard deviation of these slope coefficients.

The top panel of Table 6 displays the results. The break risk factor has nearly as much power as book-to-market in explaining the cross-section of stock returns and has approximately one-and-a-half times the power of the market size and momentum according to the test statistics. Average returns are also higher for firms exposed to break risk than those not exposed.

We also present results of alternative proxies of the break risk factors (columns 2-5 in Table 6). Our second measure uses the root-squared difference between the forecasts produced by the panel models with and without breaks. The third, fourth and fifth measures use the absolute difference at each time point in the intercept, slope and volatility parameters, respectively, estimated from the panel models with and without breaks. All five measures are statistically significant using Newey and West (1987) heteroscedasticity-adjusted t-statistics, although the third measure (based on

[21]The unbalanced panel this introduces is readily handled by our methodology.

the intercept) has the least power (t-statistic of 2.62) to explain the cross-section of expected returns.

Following Novy-Marx (2013), the bottom panel of Table 6 reports results from the same analysis performed using the break risk measures that have been demeaned by industry using the 49 industries of Fama and French (1997). The results are broadly similar except the t-statistic of every break risk measure is increased, suggesting that adjusting the risk measure by industry obtains even more power to explain the cross-section of expected returns.

From herein we focus on the break risk factor measured by the absolute difference between the forecasts produced by the panel models with and without breaks presented in equation (15) because it is the measure that has the most power in explaining the cross-section of expected returns.

5.2. Sorts on Break Sensitivity

Running Fama and MacBeth (1973) regressions using individual stocks places a lot of emphasis on nano- and micro-cap stocks that make up a considerable share of the number of stocks but only account for a small fraction of the total market capitalisation. The regressions may also be sensitive to outliers and impose a potentially misspecified parametric relation between the variables. Any subsequent inference may therefore be compromised. To alleviate this concern, we construct value-weighted portfolios sorted according to our instability risk factor and provide a nonparametric test of the hypothesis that exposure to break risk predicts average returns in the cross-section.

Table 7 displays results for the 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 that are most sensitive to breaks. Column one shows 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 to the break-sorted portfolios are monotonically increasing with our risk factor with the high-sensitivity portfolio earning a 0.27% higher average monthly return than the low-sensitivity portfolio, corresponding to an annualized return premium of 3.24% for the quintile of firms with the highest break exposure compared with the quintile of firms least exposed to this source of risk. This premium is significant at the 5% level with a t-statistic of 2.18.

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 -0.18% per month) and the most break-sensitive stocks (at 0.17%) are both significantly different from zero, as is their difference which, at 0.35% per month or more than 4% annualized, is economically large.

To alleviate concerns about transaction costs raised by Novy-Marx and Velikov (2015) and Hou et al. (2017) 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 7 displays the results. The results, while marginally weaker, tell the same basic story.

These results provide cross-sectional evidence of the existence of a break risk factor. They suggest that stocks whose equity premium processes are most sensitive to the type of breaks 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.

5.3. Break Exposure and Company Characteristics

The alpha estimates in Table 7 account for exposure to the Fama-French market, size and book-to-market factors. This is important because break risk exposure could well be correlated with size or book-to-market characteristics at the firm level. Whether such a relation exists is what we next explore.

To this end, we first re-estimate a separate set of panel break models using returns on a set of 5×5 portfolios sorted either on size and book-to-market or on size and momentum. We then rank the 25 portfolios by the sensitivity of their (dividend-price ratio) slope coefficient computed as the standard deviation of the estimated slopes across the different regimes. Results from this analysis are shown in Table 8.

Looking at the 25 portfolios sorted on size and book-to-market ratio (top panel), small firms' return processes are seen to be the most sensitive to breaks, while big firms are the least sensitive. Moreover, differences are large as the small firms have two to three times as large a sensitivity as large firms do.

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. In fact, conditional on firm size there is a monotonically decreasing relation between book-to-market ratio and break sensitivity as returns on value firms are more sensitive to breaks than returns on growth firms.

The bottom panel in Table 8 shows similar findings for the stocks sorted on size and momentum. Once again, small firms' returns are more sensitive to breaks than large firms. Moreover, conditional on firm size there is a monotonically increasing relation between prior returns and break sensitivity as "loser" stocks with the smallest prior returns are more sensitive to breaks than are "winner" stocks.

These findings suggest that firms normally thought of as being riskier (small firms and value firms) 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 explained by occasional large resurgences in the returns of 'losers' documented by Daniel and Moskowitz (2016).

6. Breaks in Dividend Growth and Return Predictability

Return predictability can arise from two principal sources, namely predictability of cash flow growth or predictability of equity risk premia. Because they are not directly observable, inference on variation in equity risk premia is dependent on how these are modeled. In contrast, we can obtain good proxies for cash flows. This section therefore explores whether the breaks identified in the return prediction models are linked to shifts in the underlying dividend growth process.[22]

Predictability of dividend growth is still being contested. For example, Cochrane (2007) argues that there is little evidence that dividend growth can be predicted. Conversely, van Binsbergen et al. (2010) report annual out-of-sample R^2 values of 13.9-31.6% for dividend growth rates, using a present value filtering approach. Chen (2009) and Kelly and Pruitt (2013) also present evidence of dividend growth predictability.

Such disagreement regarding dividend growth predictability could be caused by time-variation in predictability, i.e., dividends could be highly predictable in certain regimes while largely unpredictable in other ones. To address this question, we run a predictive regression with the dividend growth series for the 30 industry portfolios

[22] Our analysis is also relevant for past work on the predictive power of the dividend-price ratio (or dividend yield) over stock returns. Lettau and Ludvigson (2005) find that forecasts of dividends and forecasts of stock returns covary over the business cycle, implying that positively correlated fluctuations in expectations of both dividend growth and returns have counterbalancing effects on the log dividend-price ratio.

as the dependent variable and an intercept, an autoregressive term and a lagged predictor (the dividend-price ratio) on the right-hand-side.[23]

Table 9 displays the posterior mean and standard deviation of the estimated intercept, AR(1) slope, dividend-price ratio slope, and the volatility obtained from our heterogeneous panel break model. The predicted dividend growth rate varies in a wide range that spans high-growth states with a large positive intercept and AR(1) coefficient (regime eight) and states with negative expected dividend growth (regimes one and three). The AR(1) coefficient is highly significant and positive in nine out of ten regimes. Similarly, the estimated slope of the dividend-price ratio is negative in nine of ten regimes as we would expect if forecasts of higher future dividend growth lead to higher current prices and thus a smaller dividend-price ratio.[24]

The bottom row of Figure 11 displays the estimated break dates identified by our panel break model. Blue triangles mark the posterior modes of the break dates estimated from the heterogeneous panel break model fitted to dividend growth, while red triangles mark the modes of break dates fitted to the model for excess returns. Note that the 1929 break identified by the return prediction model goes undetected by the dividend growth model. This is likely because our dividend growth sample begins in 1928 and it can be difficult to detect breaks at the very beginning of the sample (Bai and Perron 1998). The remainder of the posterior modes of estimated break dates are very close to their original modes when excess returns are the dependent variable. In fact, every remaining break from the return model is estimated within one year of the original break date estimate except for an additional break identified in the 1990s and the break in 1986 being overlooked. The results across the other three predictive variables listed in the upper rows are even stronger.

Interestingly, the breaks in the dividend growth regression lead the breaks in the excess return regression. Ignoring the first break identified by the excess return regression (since the corresponding break is not detected by the dividend growth

[23]Our dividend measure is constructed as follows. First, we extract a monthly dividend yield for each industry portfolio as the difference between the monthly CRSP returns with and without dividends $D_t/P_{t-1} = (D_t + P_t)/P_{t-1} - P_t/P_{t-1}$. Next, we construct a monthly price index for each industry portfolio using the corresponding returns without dividends P_t/P_{t-1}. Multiplying the dividend yield by the price index gives us a monthly dividend series D_t for each industry portfolio which we use to compute 12-month cumulative dividends as $D_t^{12} = \sum_{t=-11}^{t} D_t$. Finally, we construct the year-on-year dividend growth rate for each portfolio at each time period $G_t = log(D_t^{12}) - log(D_{t-12}^{12})$.

[24]The only case where the AR(1) and dividend-price ratio coefficients have the wrong sign is in regime 2. This reversal of sign can sometimes happen in short-lived regimes due to collinearity between the regressors.

regression), the break in the dividend growth panel model leads the break in the return model by an average of 23 months. Further ignoring the first of the two breaks in the mid-1990s (for which the break dates from the two models are very different), the dividend growth model detects breaks that lead the return model by 12 months, on average. This evidence is suggestive that prior breaks in the dividend growth rate do, at least in part, explain the observed breaks in the return prediction models.

As a final piece of evidence on the link between breaks to the dividend growth process and breaks to returns, we analyze whether those industry portfolios whose equity premium processes are most strongly affected by breaks tend to be the same industries whose dividend growth rates are most sensitive to breaks. We use the same methodology as in the analysis of break sensitivity for the industry return processes and, once again, rank industries by the standard deviation (across regimes) in the slope coefficient on the dividend yield. The bottom panel in Table 1 shows that those industries whose return equations are most (least) affected by breaks tend to be the same industries whose dividend growth processes are most (least) sensitive to breaks. This evidence provides cross-sectional support to our earlier finding that breaks to the dividend growth process is an important driver of breaks to the return process.

7. Robustness of results

This section undertakes a number of exercises to (i) investigate whether breaks are common across portfolios; (ii) establish the general validity of our empirical findings to other predictor variables from the finance literature; and (iii) explore the sensitivity of our findings to our choice of priors.

7.1. Are breaks common across portfolios?

Our analysis assumes that the timing of the breaks is common although its impact can differ across assets. To investigate whether this assumption is reasonable, we run our estimation procedure using the same model but replacing the excess returns on the 30 industry portfolios with two alternative sets of returns, namely (i) Fama-French 5×5 portfolios sorted on size and book-to-market, and (ii) 5×5 portfolios sorted on size and momentum.[25]

Figure 12 plots the posterior modes for the break locations identified separately using the original industry portfolios, the two sets of 5×5 portfolios sorted on size and either book-to-market or momentum, or the complete set of individual stocks. For the size and book-to-market portfolios, the locations of all but one of the breaks (the 1987 break which was previously identified in 1996) are estimated within one year of the original posterior modes, while the break in 2010 is overlooked altogether. The very similar break dates imply that our common break assumption is reasonable. For the portfolios sorted on size and momentum, our procedure detects one break in 1933 while previously two breaks were detected in 1932 and 1934, respectively. Otherwise the break dates are very similar. For all portfolio sorts, as well as for the cross-section of individual stocks, the results are similar across the other three predictors we consider.

The small magnitude in the differences in break date estimates suggest that breaks affect portfolios with very different characteristics at roughly the same time and indicate that our common break assumption is a reasonably accurate description of our data.

7.2. Results using other predictors

Up to this point we focused our analysis on a return prediction model that uses the dividend-price ratio as a predictor variable. However, a variety of other predictor variables have been proposed in the literature on return predictability so we next consider three additional predictor variables in common use, namely the one-month T-bill rate, the default spread, and the term spread. Robustness of the performance of our panel breaks approach for these additional predictors will increase our confidence in the broader success and applicability of our approach.

We undertake the same analysis as that conducted for the model that uses the dividend-price ratio as a predictor. Again, all forecasts are generated recursively out-of-sample with a ten-year warm-up period.

To preserve space we do not show as many details of the analysis as we did for the model that uses the dividend-price ratio as a predictor. However, it is worth noting

[25] Since the power to detect breaks increases with the size of the cross-section N it is preferable to select portfolio sorts such that N is close to the original value of 30. Portfolio sorts involving investment or profitability begin only in 1963 and thus are not suitable for our robustness check.

that a similar number of breaks is identified for the models that use the three other predictor variables. For example, the model that uses the T-bill rate as a predictor identifies 9 breaks, the locations of which are very similar to those for the model based on the dividend-price ratio (Figure 13a). Moreover, the ability to detect these breaks with little delay continues to hold for this predictor (Figure 13b).

In fact, the ability of our approach to identify breaks in 'real time' holds across all predictors. To illustrate this, Figure 6 shows the distribution of the delays in detecting breaks for all four predictors. The vast majority of breaks are detected with a delay of a few months although a few breaks get detected with a longer delay.

Figure 14 summarizes our findings through plots of the difference in the cumulative sum of squared errors for our panel break approach measured relative to that of the three benchmarks described earlier. For simplicity, we focus our analysis on the market portfolio, but similar plots are obtained for the majority of industry portfolios. We see clear evidence that our panel break approach consistently produces more accurate forecasts than the alternatives.

Table 2 supplements the graphical analysis with more formal results. The panel break model continues to perform well for the three alternative predictors, generating Diebold-Mariano test statistics that indicate significant improvements over the three benchmarks for between 20 and 25 of the 31 industry and market portfolios. We see significant underperformance for only one case out of a total of 372 pair-wise forecast comparisons across predictor variables and benchmarks. Even stronger results are obtained using the Clark-West test statistic for which the panel break model significantly outperforms the three benchmarks for between 22 and 28 of the 31 portfolios.

Turning to the economic value associated with the panel break forecasts, the CER values in Table 5 show that the forecasts from the panel break model, when implemented in a simple mean-variance investment strategy, continue to generate utility gains in the neighborhood of 2% per annum measured relative to an investment strategy based on the historical average, a pure time-series (break) model, or a no-break panel model.

These results corroborate the more detailed analysis of the model that used the dividend-price ratio as a predictor of stock returns and so suggest that our findings are not sensitive to using a particular predictor variable.

7.3. Sensitivity of Results to Priors

Table 10 displays the results of a prior sensitivity analysis. Specifically, we adjust one hyperparameter at a time from our baseline specification and re-estimate the model. First, adjusting the hyperparameter c that controls our prior expected regime duration from 10 years to 5 (20) years results in the detection of 12 (8) breaks. The significant outperformance of the panel break model against each of the three benchmarks for the majority of portfolios is unaffected by adjusting this prior; neither is it affected by reasonable adjustments to the hyperparameter b that controls the prior volatility or σ_μ that controls the prior dispersion of the intercept. Adjusting the hyperparameter σ_η^2 to economically plausible values also has little impact on the forecasts. However, allowing this parameter to become very large which implies implausibly high R^2 values leads to a marked deterioration in the accuracy of the forecasts, corroborating the findings of Wachter and Warusawitharana (2009).

8. Conclusion

A large literature on predictability of stock market returns has found evidence of model instability, suggesting that the parameters of commonly-used return prediction models change markedly over time. Such model instability helps explain why out-of-sample return forecasts often are found to perform poorly compared to a simple constant-expected return benchmark as found by Goyal and Welch (2008). While model instability is, thus, known to affect return forecasts, exploiting it has so far proved largely elusive due to the noisy nature of returns and the low predictive power of return prediction models which makes detecting and quantifying the effects of shifts in parameter values exceedingly difficult using data on individual return series; see Lettau and Van Nieuwerburgh (2008).

In this paper, we develop an approach that exploits cross-sectional information to detect breaks from the joint dynamics of multiple return series. While our approach assumes the timing of the breaks to be the same, the effects of any breaks are allowed to differ across individual return series. Empirically, we find that using cross-sectional information in a panel break model substantially increases our ability to detect breaks in return prediction models. Importantly, our approach has the ability to accurately detect breaks in real time with very little delay. Combined with economically-motivated priors, this means that out-of-sample return forecasts generated by our panel break model are consistently more accurate than forecasts generated by a variety of extant approaches from the literature, with gains in predictive accuracy being particularly large shortly after a break has occurred.

Investors use return predictions as an input into their portfolio allocation decisions and so the possibility of breaks to the process generating expected returns can itself be viewed as a risk factor with the potential to have a sizeable impact on investment performance. Consistent with this, we find strong cross-sectional evidence in support of the presence of a break risk factor. Firms whose returns are most sensitive to break risk tend to earn higher risk-adjusted returns than firms that are less sensitive to this source of risk, suggesting that break risk is priced in the cross-section.

References

Andersen, T. G., Bollerslev, T., Christoffersen, P. F., and Diebold, F. X. (2006). Volatility and correlation forecasting. *Handbook of Economic Forecasting*, 1:777–878.

Avdis, E. and Wachter, J. A. (2017). Maximum likelihood estimation of the equity premium. *Journal of Financial Economics*, 125(3):589–609.

Avramov, D. (2002). Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64(3):423–458.

Avramov, D. and Wermers, R. (2006). Investing in mutual funds when returns are predictable. *Journal of Financial Economics*, 81(2):339–377.

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.

Banegas, A., Gillen, B., Timmermann, A., and Wermers, R. (2013). The cross section of conditional mutual fund performance in european stock markets. *Journal of Financial Economics*, 108(3):699–726.

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.

Bollerslev, T., Hood, B., Huss, J., and Pedersen, L. H. (2018). Risk everywhere: Modeling and managing volatility. *Review of Financial Studies*, 31(7):2729–2773.

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.

Chen, L. (2009). On the reversal of return and dividend growth predictability: A tale of two periods. *Journal of Financial Economics*, 92(1):128–151.

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

Chib, S. (1998). Estimation and comparison of multiple change-point models. *Journal of Econometrics*, 86(2):221–241.

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. (2007). The dog that did not bark: A defense of return predictability. 21(4):1533–1575.

Dangl, T. and Halling, M. (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics*, 106(1):157–181.

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

DeMiguel, V., Garlappi, L., and Uppal, R. (2007). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13:253–263.

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. (1997). Industry costs of equity. *Journal of Financial Economics*, 43(2):153–193.

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

Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508.

Green, J., Hand, J. R., and Soliman, M. T. (2011). Going, going, gone? the apparent demise of the accruals anomaly. *Management Science*, 57(5):797–816.

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

Hjalmarsson, E. (2010). Predicting global stock returns. *Journal of Financial and Quantitative Analysis*, 45(1):49–80.

Hou, K., Xue, C., and Zhang, L. (2017). Replicating anomalies. *Unpublished Manuscript.*

Jagannathan, R. and Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance*, 58(4):1651–1683.

Johannes, M., Korteweg, A., and Polson, N. (2014). Sequential learning, predictability, and optimal portfolio returns. *Journal of Finance*, 69(2):611–644.

Kelly, B. and Pruitt, S. (2013). Market expectations in the cross-section of present values. *Journal of Finance*, 68(5):1721–1756.

Kim, D. (2011). Estimating a common deterministic time trend break in large panels with cross sectional dependence. *Journal of Econometrics*, 164(2):310–330.

Koop, G. and Potter, S. M. (2007). Estimation and forecasting in models with multiple breaks. *Review of Economic Studies*, 74(3):763–789.

Lettau, M. and Ludvigson, S. C. (2005). Expected returns and expected dividend growth. *Journal of Financial Economics*, 76(3):583–626.

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

McLean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? *Journal of Finance*, 71(1):5–32.

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. 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.

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, University of Cambridge*.

Pettenuzzo, D. and Timmermann, A. (2011). Predictability of stock returns and asset allocation under structural breaks. *Journal of Econometrics*, 164(1):60–78.

Polk, C., Thompson, S., and Vuolteenaho, T. (2006). Cross-sectional forecasts of the equity premium. *Journal of Financial Economics*, 81(1):101–141.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies*, 23(2):821–862.

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

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

Smith, S. C. and Timmermann, A. (2017). Detecting breaks in real time: A panel forecasting approach. *Unpublished Manuscript, University of California San Diego*.

Smith, S. C., Timmermann, A., and Zhu, Y. (2017). Variable selection in panel models with breaks. *Unpublished Manuscript, University of California San Diego*.

van Binsbergen, Jules, H., and Koijen, R. S. (2010). Predictive regressions: A present-value approach. *Journal of Finance*, 65(4):1439–1471.

Viceira, L. M. (1997). Testing for structural change in the predictability of asset returns. *Manuscript, Harvard University*.

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.

Appendix A. Priors

We now provide details of the prior distributions used by our return prediction models.

Appendix A.0.1. Prior on the Regime Durations

Unit-specific Breaks: Chib (1998)'s method restricts the regime duration of each time-series in the cross-section to follow a geometric prior distribution

$$p(l_{ik_i} \mid p_{k_i k_i}) = p_{k_i k_i}^{l_{ik_i}-1}(1 - p_{k_i k_i}), \qquad k_i = 1, \ldots, K_i + 1, \tag{A.1}$$

in which the prior nontransition probability $p_{k_i k_i}$ follows a conjugate beta distribution

$$p(p_{k_i k_i}) = \frac{\Gamma(g+h)}{\Gamma(g)\Gamma(h)} p_{k_i k_i}^{g-1}(1 - p_{k_i k_i})^{h-1}, \qquad k_i = 1, \ldots, K_i. \tag{A.2}$$

Common Breaks: Koop and Potter (2007) note that such a monotonically decreasing geometric prior on the regime durations enforced by Chib (1998)'s algorithm may be unrealistic and therefore suggest specifying a Poisson distribution instead. In the panel break model we develop the regime durations have a Poisson prior distribution

$$p(l_k \mid \lambda_k) = \frac{\lambda_k^{l_k} e^{-\lambda_k}}{l_k!}, \qquad k = 1, \ldots, K + 1, \tag{A.3}$$

where the Poisson intensity parameter λ_k has a conjugate Gamma prior distribution

$$p(\lambda_k) = \frac{d^c}{\Gamma(c)} \lambda_k^{c-1} e^{-d\lambda_k}, \qquad k = 1, \ldots, K + 1, \tag{A.4}$$

in which c and d are the hyperparameters of $\lambda = (\lambda_1, \ldots, \lambda_{K+1})$.

Appendix A.0.2. Priors on Parameters β and σ

The estimation of the panel break model is simplified by specifying conjugate priors on the error-term variances σ^2 and on the regression coefficients θ conditional on the error-term variances σ^2.[26] For consistency we use the equivalent prior specification

[26] For ease of exposition, let $\boldsymbol{X}$ denote a $(2 \times T)$ matrix that combines a unit vector of length T with the predictor X. This results in $\theta = (\mu, \beta)$ containing estimates of both the intercept and slope

for the model with unit-specific breaks and parameters.

For the model with common breaks and unit-specific parameters we specify an inverse gamma prior over the regime-specific variances for $i = 1, \ldots, N$

$$p(\sigma_{ik}^2) = \frac{b^a}{\Gamma(a)} \sigma_{ik}^{2^{-(a+1)}} exp\left(-\frac{b}{\sigma_{ik}^2}\right), \qquad k = 1, \ldots, K+1. \tag{A.5}$$

and a Normal prior with zero mean is placed over the regression coefficients conditional on the variances

$$p(\theta_{ik} \mid \sigma_{ik}^2) = 2\pi^{-\kappa/2}(\sigma_{ik}^2)^{-\kappa/2} \mid V_\beta \mid^{-1/2} exp\left(-\frac{1}{2\sigma_{ik}^2}\theta_{ik}' V_\theta^{-1}\theta_{ik}\right), k = 1, \ldots, K+1,$$
$$V_\theta = \begin{pmatrix} \sigma_\mu^2 & 0 \\ 0 & \sigma_\eta^2/\sigma_X^2 \end{pmatrix}. \tag{A.6}$$

Appendix A.1. Prior Elicitation

For the empirical application we assume that a break occurs approximately every decade for the heterogeneous panel break model we develop and the benchmark time-series break model. This is in line with findings in earlier studies such as Pástor and Stambaugh (2001). Specifically, we set our hyperparameter values as $d = h = 2$ and $c = g = 240$ to give a prior expected regime duration of 120 periods. We further set $a = 2$ and $b = 0.0049$ to give a prior expected error-term variance equal to 0.0049 which is the average of the variance of the excess returns across all the portfolios using the full sample. The choice of hyperparameter values for σ_μ and σ_η, and the scaling of σ_η with the empirical variance of the predictive variable using the full sample available at the time each forecast is made σ_X, have been explained in Section 2.4.

Appendix A.2. Posterior Distribution

Inference is performed on the posterior distribution which combines information in the data with prior information supplied by the user. We now explicitly detail the posterior distributions of the different models we consider.

coefficient.

Unit-specific Breaks and Parameters: Letting $\Theta_i = (\theta_i, \sigma_i^2)$, the posterior distribution of the model with unit-specific breaks and parameters is

$$p(\Theta_i \mid \boldsymbol{r_i}, \boldsymbol{X}, \tau_i) = \left(\prod_{k_i=1}^{K_i+1} (\sigma_{ik_i}^2)^{-((l_{ik_i}+\kappa)/2+a+1)} (2\pi)^{-(\kappa+l_{ik_i})/2} \mid V_\theta \mid^{-1/2} \frac{b^a}{\Gamma(a)} \right)$$
$$\times\, exp\left[-\frac{1}{2} \sum_{k_i=1}^{K_i+1} \left(\frac{\theta'_{ik_i} V_\theta^{-1} \theta_{ik_i} + 2b + \left(r'_{ik_i} r_{ik_i} - r'_{ik_i} X'_{k_i} \theta_{ik_i} - \theta'_{ik_i} X_{k_i} r_{ik_i} + \theta'_{ik_i} X_{k_i} X'_{k_i} \theta_{ik_i} \right)}{\sigma_{ik_i}^2} \right) \right], \tag{A.7}$$

in which $r_{ik_i} = (r_{i\tau_{k_i-1}+1}, \dots, r_{i\tau_{k_i}})$ and

$$\underset{(2\times l_{ik_i})}{X_{k_i}} = \begin{pmatrix} 1 & \dots & 1 \\ X_{\tau_{ik_i-1}+1} & \dots & X_{\tau_{ik_i}} \end{pmatrix}.$$

Pooled Breaks and Unit-specific Parameters The posterior distribution of the model with pooled breaks and unit-specific parameters is

$$p(\boldsymbol{r} \mid \boldsymbol{X}, \tau) = \prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi)^{-l_k/2} \frac{\mid \Sigma_k \mid^{1/2}}{\mid V_\theta \mid^{1/2}} \frac{b^a}{\Gamma(a)} \frac{\Gamma(\tilde{a}_k)}{\tilde{b}_{ik}^{\tilde{a}_k}}. \tag{A.8}$$

in which, for regimes $k = 1, \dots, K+1$

$$\begin{aligned} \Sigma_k^{-1} &= V_\beta^{-1} + X_k X'_k, \\ \rho_{ik} &= \Sigma_k X_k r_{ik}, \qquad i = 1, \dots, N \\ \tilde{a}_k &= a + (l_k)/2, \\ \tilde{b}_{ik} &= \frac{1}{2} \left(2b + r'_{ik} r_{ik} - \rho'_{ik} \Sigma_k^{-1} \rho_{ik} \right), \qquad i = 1, \dots, N, \end{aligned} \tag{A.9}$$

in which $r_{ik} = (r_{i\tau_{k-1}+1}, \dots, r_{i\tau_k})$ and

$$\underset{(2\times l_k)}{X_k} = \begin{pmatrix} 1 & \dots & 1 \\ X_{\tau_{k-1}+1} & \dots & X_{\tau_k} \end{pmatrix}.$$

Appendix B. Estimation of the Model

This appendix provides details of the procedures used to estimate the different models considered in the paper. The model with unit-specific breaks and parameters is repeatedly estimated for each time-series in the cross-section using the multiple breakpoint model of Chib (1998). This procedure estimates a series of models each with a different number of breaks and subsequently uses the marginal likelihood approach of Chib (1995) to derive the posterior model probabilities and determine the optimal number of breaks. Given the popularity of Chib (1998)'s algorithm along with the desire to save space we do not present the details of the algorithm here.

In contrast, the models with common breaks analyse the entire cross-section at once using an alternative estimation procedure that introduces the number of breaks as a parameter in the model and performs inference over this parameter by jumping between different numbers of breaks. The proportion of the Markov chain Monte Carlo run that is spent at each number of breaks approximates the posterior model probabilities (Green 1995). Our estimation approach has a range of desirable properties relative to Chib (1998); we refer the reader to Smith and Timmermann (2017) for a thorough discussion.

Briefly, we simulate the breakpoint vector τ in two steps. First, a global movement is provided by attempting either to add or remove a breakpoint on each sweep of the Markov chain Monte Carlo run. Second, to ensure the estimated breakpoint locations converge to their true values all that is required is a small perturbation of each breakpoint delivered by a random-walk Metropolis-Hastings step. Finally the parameters can be sampled from their full conditional distributions.

Table 1: Ranking of Industries by break magnitudes: Returns and dividend growth

		s.d. across regimes		
Industry	Size of break rank	Intercept	Slope	Volatility
		Returns		
Oil	1	0.091	0.130	0.077
Fin	2	0.090	0.115	0.054
Telcm	3	0.085	0.112	0.054
Beer	4	0.087	0.099	0.053
Util	5	0.083	0.090	0.049
Books	26	0.025	0.066	0.037
Other	27	0.025	0.064	0.032
Fabpr	28	0.020	0.060	0.034
Whlsl	29	0.017	0.057	0.032
Servs	30	0.015	0.052	0.028
		Dividend growth		
Oil	1	0.061	0.095	0.069
Telcm	2	0.059	0.084	0.062
Fin	3	0.060	0.082	0.059
Util	4	0.058	0.077	0.051
Buseq	5	0.046	0.074	0.048
Whlsl	26	0.018	0.035	0.019
Servs	27	0.011	0.032	0.017
Other	28	0.010	0.030	0.015
Books	29	0.009	0.027	0.013
Meals	30	0.008	0.024	0.010
		Rank correlation		
		0.89	0.92	0.90

Table 1: Break magnitude for returns and dividend growth. This table reports the standard deviation of the regime-specific intercept, slope and volatility posterior mean estimates from our heterogeneous panel break model when the dependent variable is excess returns (top panel) or dividend growth (bottom panel) and the predictive variable is the dividend-price ratio. The portfolios are ranked in terms of break size according to the standard deviation of the regime-specific slope estimates. The bottom line reports the correlation between industry ranks based on movements across regimes in the intercept, slope and volatility estimates for the return and dividend growth models.

Table 2: Statistical significance of gains in predictive accuracy

Predictor	DM				CW			
	t <-1.64	-1.64< t <0	0< t <1.64	t >1.64	t <-1.64	-1.64< t <0	0< t <1.64	t >1.64
				No break panel				
dp	0	2	4	25†	0	2	2	27†
tbl	0	3	5	23†	0	3	2	26†
tms	0	4	6	21†	0	4	3	24†
dfs	0	6	5	20†	0	6	3	22†
				Industry prevailing mean				
dp	0	2	5	24†	0	2	3	26†
tbl	0	2	4	25†	0	2	1	28†
tms	1	4	5	21†	0	5	3	23†
dfs	0	7	3	21†	0	7	2	22†
				Time series break				
dp	0	2	3	26†	0	2	2	27†
tbl	0	4	4	23†	0	4	1	26†
tms	0	5	5	21†	0	5	4	22†
dfs	0	6	4	21†	0	6	2	23†

Table 2: Statistical significance of forecast improvements. This table reports the statistical significance of the gains in predictive accuracy for our panel break model relative to the heterogeneous panel model with no breaks (No break panel), the industry prevailing mean and the time series model with breaks applied to each portfolio in turn (Time series break) when forecasting with the dividend-price ratio (dp), the treasury-bill rate (tbl), the term spread (tms), or the default spread (dfs). Significance is evaluated using the Diebold-Mariano test (DM) and the procedure of Clark and West (2007) (CW). For each procedure the table displays the number of portfolios for which our method produces significantly worse, insignificantly worse, insignificantly better, and significantly better forecasts at the 10% level. † indicates the particular bin in which the *t*-statistic for the market portfolio lies.

Table 3: Magnitude of break by portfolio

Industry	Size of break rank	MSFD	Utility Gain rank	Utility gain
		Upper quartile		
telcm	1	0.0220	1	2.48
util	2	0.0172	13	1.63
oil	3	0.0148	7	1.79
buseq	4	0.0144	6	1.94
fin	5	0.0142	2	2.40
hlth	6	0.0140	3	2.31
beer	7	0.0129	5	1.97
		Lower quartile		
fabpr	24	0.0059	16	1.45
whlsl	25	0.0055	29	- 0.34
textls	26	0.0050	22	1.04
mines	27	0.0041	26	0.67
books	28	0.0033	28	0.15
meals	29	0.0029	30	- 0.47
other	30	0.0024	27	0.61

Table 3: Magnitude of break by industry. This table lists in descending order the upper and lower quartile portfolios according to the magnitude of the total impact of breaks on their respective forecasts (with 1 denoting the largest and 30 denoting the smallest impact). This magnitude is captured by the mean squared forecast difference ('MSFD') between the panel models with and without breaks. The table also reports the ranking of the utility gain (certainty equivalent return), expressed as an annualised percentage, for a mean-variance investor with a risk aversion coefficient of three when forecasting with the panel break model relative to the panel model without breaks using the dividend-price ratio as the predictive variable.

Table 4: Industry allocations

Industry	Hist avg	No brk	ts	Pbrk
food	0.01	0.00	0.00	0.00
beer	0.23	0.00	0.11	0.19
smoke	0.09	0.24	0.07	0.13
books	0.00	0.00	0.00	0.02
hlth	0.06	0.02	0.01	0.00
chems	0.14	0.06	0.13	0.04
txtls	0.00	0.01	0.00	0.00
elceq	0.02	0.01	0.00	0.01
autos	0.07	0.02	0.04	0.00
oil	0.04	0.00	0.00	0.00
telcm	0.03	0.01	0.23	0.05
servs	0.16	0.45	0.37	0.42
buseq	0.11	0.17	0.01	0.07
paper	0.02	0.00	0.01	0.00
fin	0.00	0.00	0.01	0.05

Table 4: Industry allocations. This table reports the weight allocated to each of the thirty industries, averaged across the out-of-sample period, using four alternative approaches to predict stock returns. We omit industries which have an allocation of less than 0.01 for each model. 'Hist avg' denotes the industry prevailing mean, 'No brk' denotes the heterogeneous panel model without breaks, 'ts' denotes the time series break model, and 'Pbrk' denotes the heterogeneous panel model with breaks.

Table 5: Utility gains from portfolio investment strategies

	Full sample			After breaks		
Predictor	hist avg	no brk	ts	hist avg	no brk	ts
	Investment in 30 industry portfolios					
dp	2.19	2.02	1.97	3.02	2.43	2.72
tbl	2.04	2.10	2.34	2.61	2.80	3.06
tms	1.99	2.42	1.86	2.21	2.57	2.37
dfs	2.02	1.92	2.29	2.89	2.53	2.72
	Market portfolio					
dp	1.59	1.92	2.03	2.50	2.46	2.63
tbl	1.85	2.02	1.69	2.42	2.64	2.18
tms	2.12	1.62	1.84	2.59	2.20	2.41
dfs	1.71	1.47	1.89	2.02	1.59	2.21

Table 5: Utility gains. The top panel of this table reports the out-of-sample utility gain (certainty equivalent return) for a mean-variance investor with a risk aversion of three who at each period allocates wealth between a risk-free asset (T-bills) and an optimal risky portfolio that is constructed from the 30 industry portfolios. We report the utility gain measured relative to each of the three benchmark models, namely, the prevailing mean (hist avg), the panel model with no breaks (no brk), and the time series break model (ts). The utility gain is computed first using the full sample and second using only the observations that fall within 24 months of a break being detected without counting any observation twice. Results are presented for the four predictors we consider: the dividend-price ratio (dp), the T-bill rate (tbl), the term spread (tms), and the default spread (dfs). The reported certainty equivalent returns are expressed as annualised percentages. The bottom panel reports the utility gain for a mean-variance investor with a risk aversion of three who allocates his wealth between T-bills and the market portfolio.

Table 6: Fama-Macbeth regressions of returns on break risk factor

Independent variable	Slope coefficients ($\times 10^2$) and (test-statistics)				
	(1)	(2)	(3)	(4)	(5)
	Break risk measures				
BRK	0.59	0.54	0.30	0.51	0.53
	(4.60)	(4.27)	(2.62)	(4.13)	(4.22)
log(B/M)	0.30	0.27	0.25	0.28	0.31
	(5.19)	(5.01)	(4.87)	(5.14)	(5.33)
log(ME)	-0.08	-0.11	-0.11	-0.09	-0.09
	(-3.02)	(-3.20)	(-3.19)	(-3.07)	(-3.07)
$PR1YR$	0.57	0.64	0.59	0.63	0.66
	(3.17)	(3.29)	(3.19)	(3.23)	(3.31)
	Break risk measures demeaned by industry				
BRK	0.71	0.67	0.35	0.65	0.66
	(5.62)	(5.14)	(2.99)	(4.95)	(5.08)
log(B/M)	0.28	0.28	0.27	0.29	0.31
	(5.09)	(5.08)	(4.99)	(5.12)	(5.23)
log(ME)	-0.07	-0.11	-0.10	-0.09	-0.07
	(-2.89)	(-3.02)	(-2.97)	(-2.91)	(-2.88)
$PR1YR$	0.53	0.61	0.55	0.63	0.69
	(3.07)	(3.20)	(3.10)	(3.22)	(3.40)

Table 6: 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 is computed at each time for each firm as the absolute difference between forecasts produced from our heterogeneous panel model with and without breaks. The second measure is the root mean squared difference between these forecasts. The third, fourth and fifth measures 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 and the industries are the 49 industry portfolios of Fama and French (1997).

Table 7: Return performance of portfolios of stocks sorted on break sensitivity

Portfolio	r	α	MKT	SMB	HML
			All stocks		
Low	0.26	-0.18	1.02	0.01	0.05
	(1.98)	(-2.04)	(22.21)	(1.43)	(3.93)
2	0.32	-0.06	0.98	0.03	0.13
	(2.19)	(-1.99)	(31.98)	(1.69)	(2.80)
3	0.44	-0.01	0.94	0.00	-0.09
	(2.25)	(-1.60)	(32.04)	(2.06)	(-1.21)
4	0.46	0.02	1.01	0.06	0.14
	(1.98)	(1.01)	(24.09)	(1.31)	(2.02)
High	0.53	0.17	0.96	-0.03	-0.03
	(2.58)	(2.04)	(23.26)	(-1.99)	(-3.78)
High-low	0.27	0.35	-0.06	-0.04	-0.08
	(2.18)	(2.97)	(-1.05)	(-2.03)	(-1.55)
			Without micro-caps		
Low	0.26	-0.13	0.97	0.01	0.02
	(1.99)	(-2.00)	(19.89)	(1.44)	(2.87)
2	0.28	-0.10	0.99	0.02	0.14
	(2.10)	(-2.01)	(31.68)	(1.69)	(2.82)
3	0.39	-0.05	0.90	0.00	-0.04
	(2.18)	(-1.71)	(32.00)	(2.03)	(-1.02)
4	0.47	0.04	0.98	0.02	0.10
	(2.01)	(1.11)	(23.21)	(1.01)	(2.00)
High	0.49	0.16	0.88	-0.05	-0.02
	(2.44)	(2.03)	(21.06)	(-2.03)	(-3.59)
High-low	0.23	0.29	-0.09	-0.06	-0.04
	(2.10)	(2.78)	(-1.24)	(-2.04)	(-1.60)

Table 7: Excess returns to portfolios 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 absolute difference in the forecasts from the heterogeneous panel models with and without breaks when the dividend-price ratio is 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 factors, i.e., the market (MKT), size (SMB) and value (HML). 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 8: Break sensitivity and stock characteristics

Portfolio	Size of break rank	s.d. across regimes Intercept	Slope	Volatility
		Size and Value		
SMALL.HiBM	1	0.067	0.141	0.088
ME1BE4	2	0.059	0.132	0.078
ME1BE3	3	0.058	0.129	0.078
ME1BE2	4	0.051	0.111	0.069
SMALL.LoBM	5	0.041	0.107	0.067
BIG.HiBM	20	0.032	0.052	0.054
ME5BE4	21	0.028	0.050	0.048
ME5BE3	23	0.024	0.044	0.039
ME5BE2	24	0.021	0.038	0.029
BIG.LoBM	25	0.019	0.030	0.024
		Size and Momentum		
SMALL.LoPRIOR	1	0.049	0.091	0.087
ME1PRIOR2	2	0.043	0.080	0.082
ME1PRIOR3	3	0.032	0.062	0.075
ME1PRIOR4	4	0.031	0.060	0.071
SMALL.HiPRIOR	5	0.027	0.056	0.052
BIG.LoPRIOR	20	0.029	0.042	0.046
ME5PRIOR2	22	0.020	0.036	0.042
ME5PRIOR3	23	0.012	0.025	0.038
ME5PRIOR4	24	0.009	0.022	0.026
BIG.HiPRIOR	25	0.008	0.020	0.022

Table 8: Break sensitivity and stock characteristics. This table reports the standard deviation across regimes of the regime-specific intercept, slope and volatility posterior mean estimates, respectively, from our heterogeneous panel break model fitted to portfolios sorted on size and value (top panels) and size and momentum (bottom panels) when the predictive variable is the dividend-price ratio. The portfolios are ranked by the standard deviation of the regime-specific slope estimates.

Table 9: Parameter estimates for the dividend growth model

Regime	Break dates	Intercept		AR(1)		Dp slope		Volatility	
		Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
1	Feb 1931	-0.038	(0.005)	0.135	(0.039)	-0.082	(0.023)	0.201	(0.003)
2	May 1933	0.030	(0.004)	-0.193	(0.031)	0.057	(0.024)	0.156	(0.003)
3	Aug 1939	-0.127	(0.016)	0.188	(0.043)	-0.034	(0.066)	0.216	(0.010)
4	Mar 1945	0.021	(0.002)	0.308	(0.022)	-0.029	(0.018)	0.113	(0.001)
5	Oct 1968	0.019	(0.004)	0.515	(0.032)	-0.156	(0.022)	0.163	(0.003)
6	Jan 1987	0.027	(0.002)	0.215	(0.029)	-0.267	(0.014)	0.147	(0.002)
7	Dec 1998	0.002	(0.005)	0.126	(0.045)	-0.305	(0.021)	0.228	(0.004)
8	Sep 2007	0.072	(0.004)	0.430	(0.041)	-0.169	(0.019)	0.204	(0.003)
9	May 2009	0.019	(0.009)	0.307	(0.069)	-0.205	(0.022)	0.327	(0.006)
10		0.043	(0.010)	0.834	(0.112)	-0.187	(0.019)	0.545	(0.007)

Table 9: Dividend growth parameter estimates. This table displays the posterior mean and standard deviation (s.d.) of the intercept, the slope on the AR(1) term and the slope on the lagged dividend-price ratio (Dp slope) obtained from the heterogeneous panel break model in each regime it identifies. The reported values are value-weighted averages across the parameter estimates on the 30 industry portfolios. We also report the mean and standard deviation of the volatility parameter. The posterior modes of the identified break dates are also reported.

Table 10: Sensitivity of results to priors

Hyp. value	K	DM				CW			
		t <-1.64	-1.64< t <0	0< t <1.64	t >1.64	t <-1.64	-1.64< t <0	0< t <1.64	t >1.64
					No-break panel				
c=120	12	0	2	6	23†	0	2	4	25†
c=480	8	0	4	5	22†	0	4	4	23†
b=1	10	0	4	5	22†	0	4	4	23†
b = 10	10	0	3	6	22†	0	3	4	24†
σ_η = 0.02	10	0	2	5	24†	0	2	3	26†
σ_η = 0.06	10	1	3	6†	21	1	3	4	23†
σ_η = 100	16	14†	8	5	4	14†	8	5	4
σ_μ = 10%	10	1	4	6	20†	1	4	4	22†
					Industry prevailing mean				
c=120		0	4	5	22†	0	4	2	25†
c=480		0	4	6	21†	0	4	4	23†
b=1		0	5	6	20†	0	5	5	21†
b = 10		0	5	6†	20	0	5	5†	21
σ_η = 0.02		0	4	3	24†	0	4	2	25†
σ_η = 0.06		0	4	3	24†	0	4	2	25†
σ_η = 100		20†	5	3	3	21†	4	3	3
σ_μ = 10%		1	3	5	22†	0	4	3	24†
					Time series break				
c=120		0	4	5	22†	0	4	3	24†
c=480		0	3	7	21†	0	3	5	23†
b=1		0	3	6	22†	0	3	3	25†
b=10		0	5	6	20†	0	5	4	22†
σ_η = 0.02		0	5	6†	20	0	5	3†	23
σ_η = 0.06		0	2	6	23†	0	2	4	25†
σ_η = 100		15†	4	5	7	16†	3	5	7
σ_μ = 10%		0	5	6	20†	0	5	4	22†

Table 10: Sensitivity of results to priors. This table reports results when forecasting excess returns on the 31 portfolios (including the market) using the heterogeneous panel break model and the dividend-price ratio as the predictive variable. We adjust one hyperparameter at a time so the hyperparameter value displayed in the table is used alongside all the remaining values detailed in Section Appendix A.1. The three panels display the statistical significance of outperformance or underperformance of our model relative to the three benchmarks we consider. Significance is evaluated using the Diebold-Mariano test (DM) and the procedure of Clark and West (2007) (CW). For each procedure the table displays the number of cases for which our method produces significantly worse, insignificantly worse, insignificantly better, and significantly better forecasts at the 10% level. † indicates the particular bin in which lies the t-statistic for the market return forecasts. In each case, the number of breaks K detected by our model is also reported.

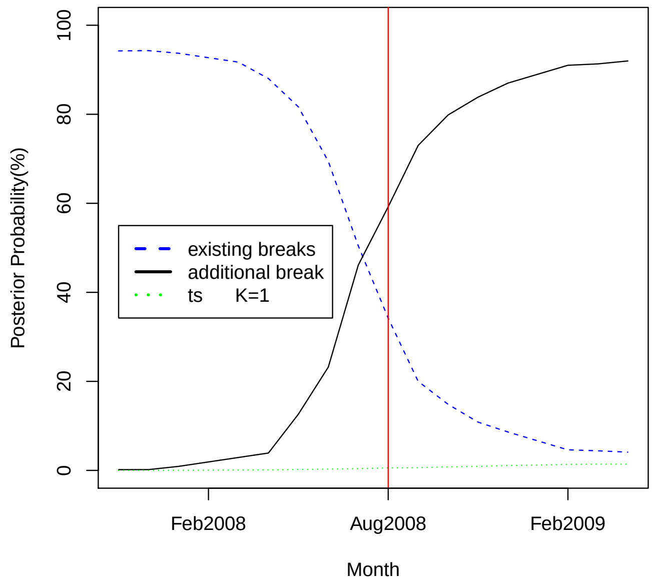

(a) Posterior Break Probabilities

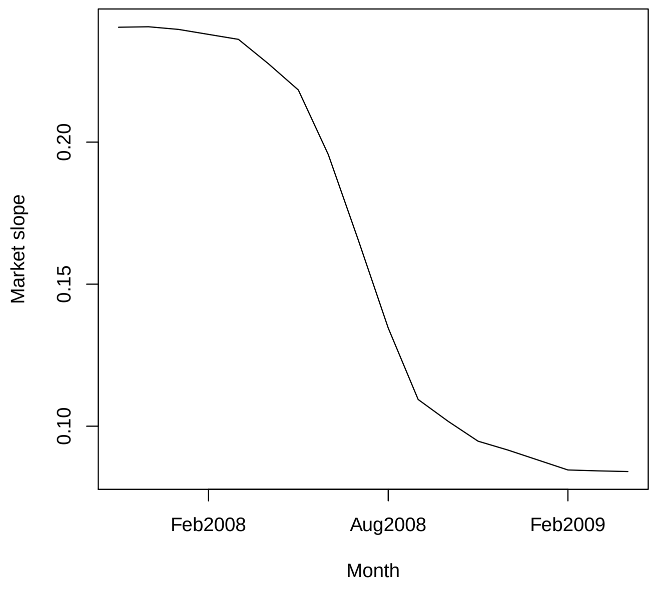

(b) Market Slope

Figure 1: This figure displays how an investor could have updated her belief that the global financial crisis had caused a permanent shift in the return prediction model (top window) and how she could have used Bayes' rule to update the slope coefficient on the lagged dividend-price ratio in the return prediction model (bottom window) using our panel break model. The green dotted line in the top window graphs period-by-period snapshots of the estimated probability of a break being detected by the time series model with breaks (ts), averaged across the 30 industry portfolios. The solid black line (dashed blue line) denotes the correspondingly updated posterior model probability assigned to an additional break by the panel break model (Pbrk). The vertical red line denotes the time at which the posterior mode includes a new break.

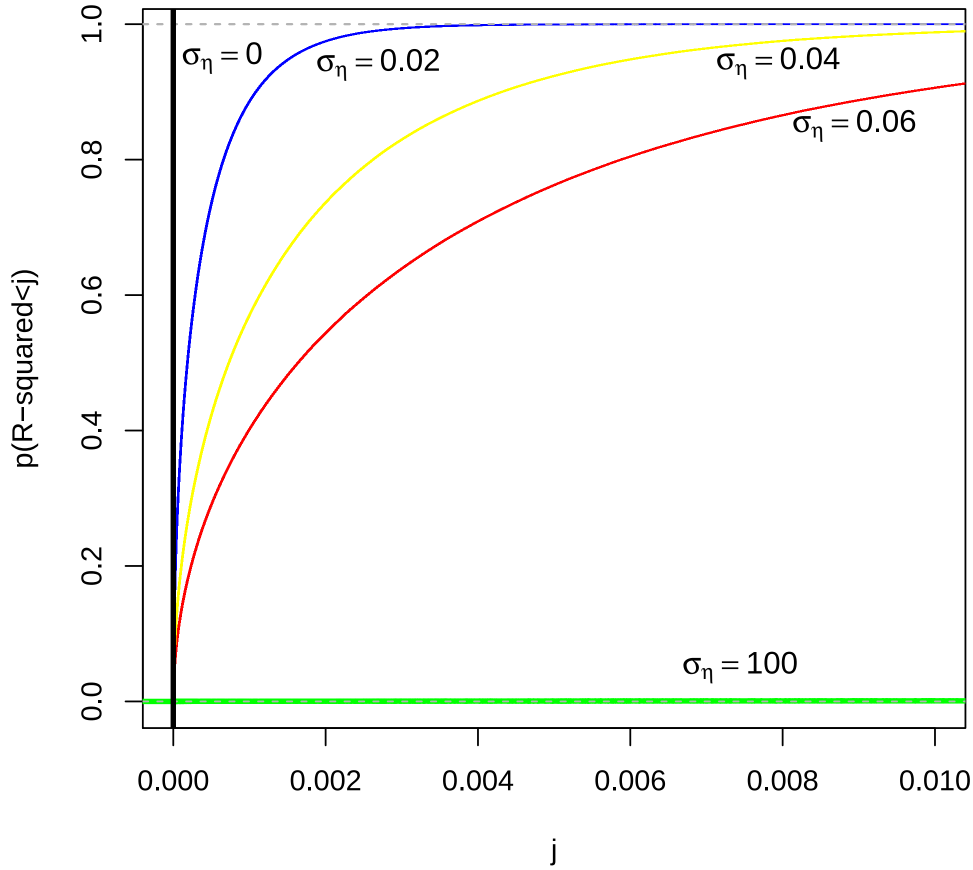

Figure 2: This figure displays the prior probability that the R-squared of a predictive regression lies below a certain value j, ranging from 0 to 0.01, for different degrees of scepticism regarding predictability. The investor's degree of scepticism is captured by the prior standard deviation of the normalised slope coefficient σ_η. A value of 0 denotes a dogmatic prior, a value of infinity denotes a diffuse prior, and intermediate values denote scepticism about the existence of return predictability.

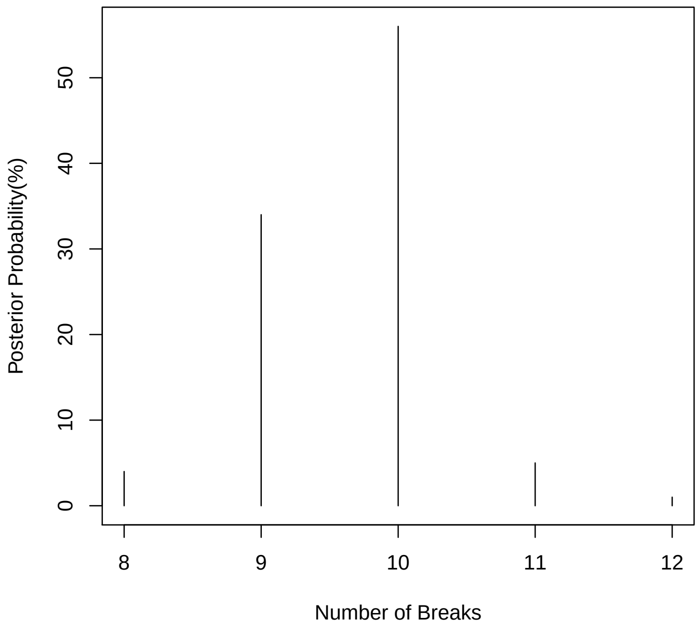

(a) Posterior Model Probability

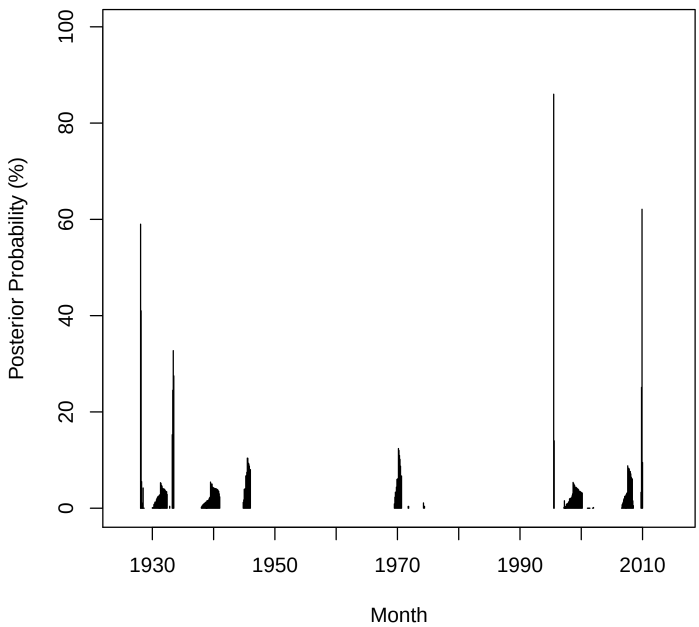

(b) Posterior Break Locations

Figure 3: This figure displays the posterior probabilities for the number of breaks (top panel) and the posterior break dates (bottom panel) obtained from a panel break model that regresses industry portfolio returns on the lagged dividend-price ratio.

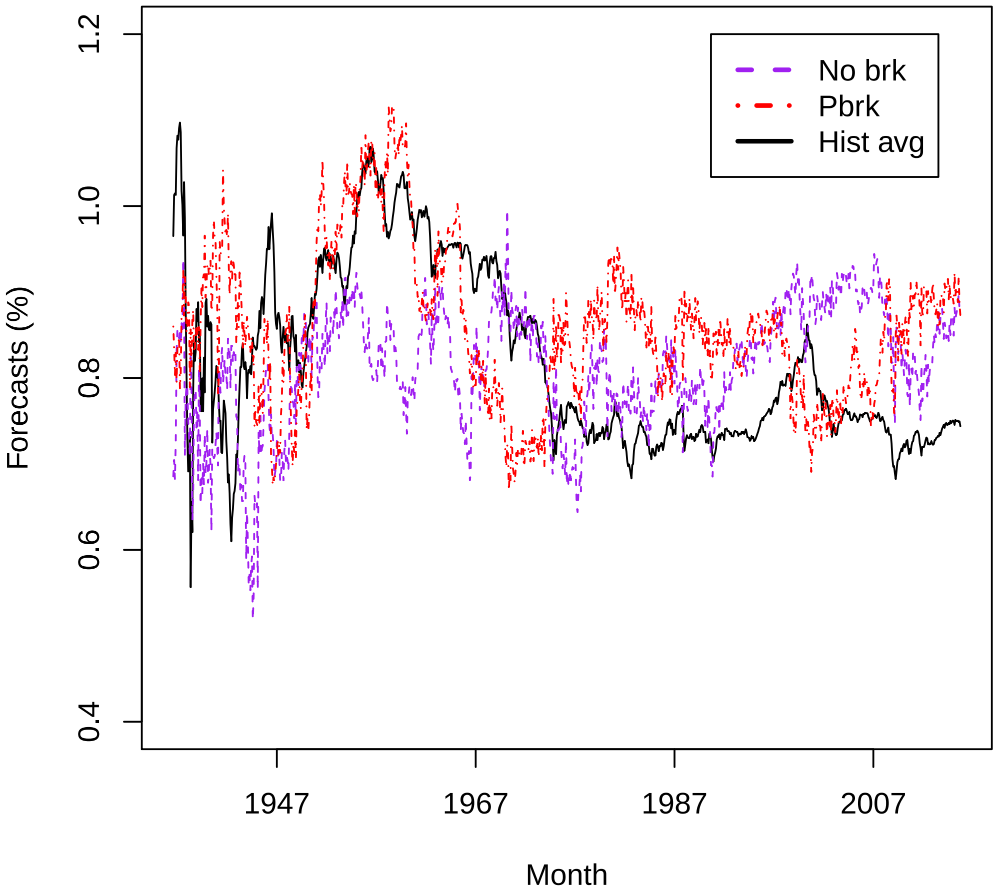

Figure 4: This figure graphs out-of-sample forecasts of returns on the market portfolio produced by the panel model with (Pbrk; red dotted line) and without breaks (No brk; purple dashed line) and the historical average (Hist avg; solid black line) using the dividend-price ratio as the predictive variable.

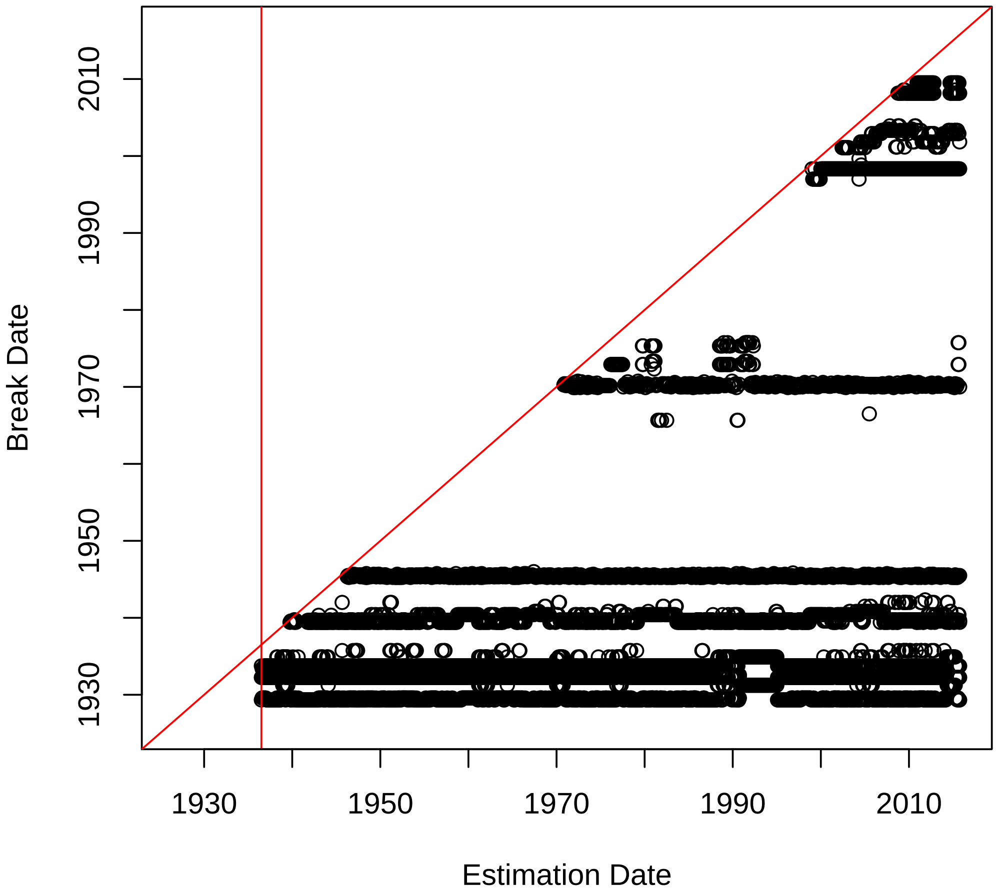

Figure 5: This figure displays the real-time break detection obtained from our panel break model fitted to 30 industry portfolio return series. The vertical red line denotes the initial estimation period and the 45 degree line (to the right of the vertical line) denotes the date at which a break could first be identified. Black circles mark the estimated break dates.

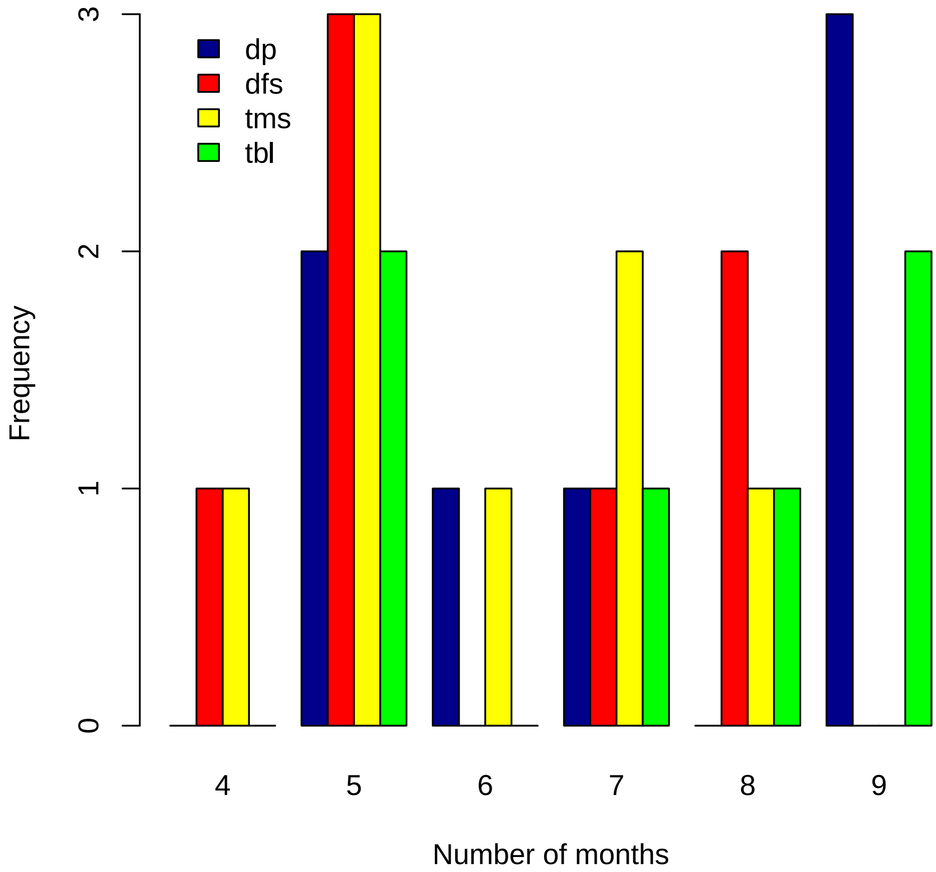

Figure 6: This figure displays the number of months it took to first detect each of the breaks that occurred after the initial estimation period when predicting with each of the four predictor variables we consider, namely, the dividend-price ratio (dp), the T-bill rate (tbl), the term spread (tms), and the default spread (dfs).

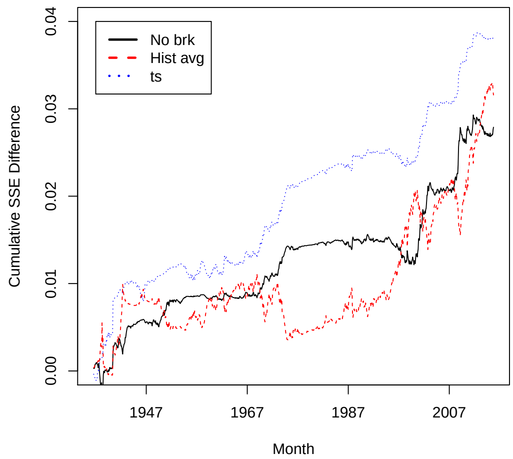

(a) Market portfolio

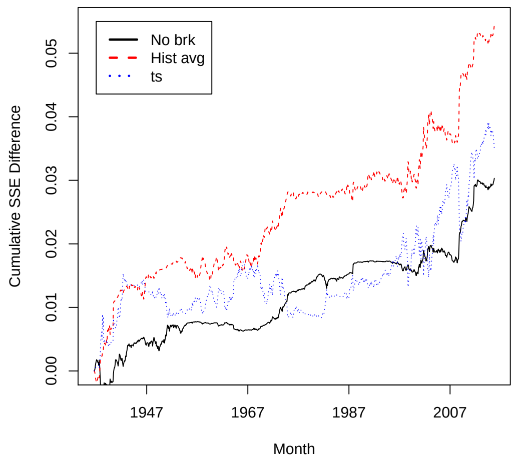

(b) Oil

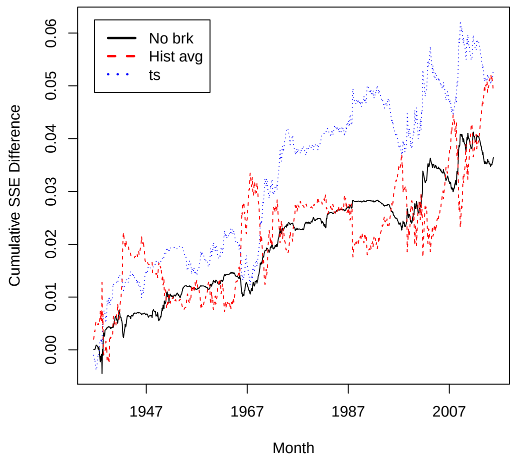

(c) Financials

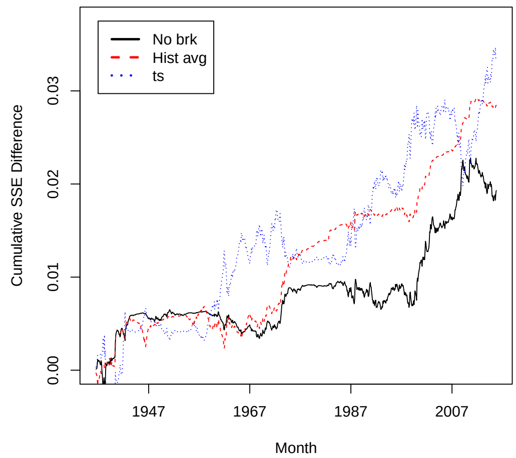

(d) Telecommunications

Figure 7: This figure graphs the cumulative difference in the sum of squared errors for the portfolio in question obtained from our heterogeneous panel break model relative to each of the benchmark models. The benchmark models are the heterogeneous panel model with no breaks ('No brk'), the industry prevailing mean ('Hist avg') and the time-series model with breaks ('ts') estimated using the algorithm of Chib (1998) applied to each portfolio in turn. The dividend-price ratio is the predictive variable and the portfolio being forecast is detailed in the subcaption of each subfigure.

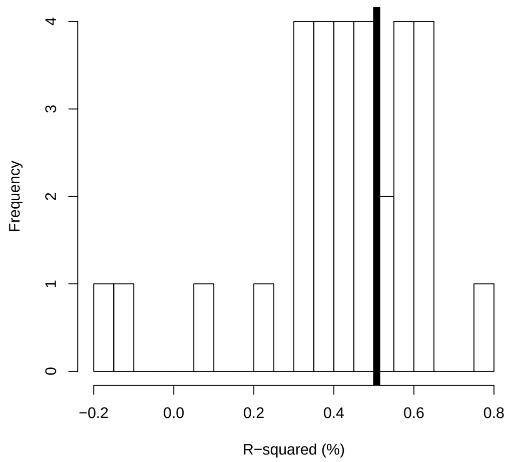

(a) Heterogeneous panel no break model

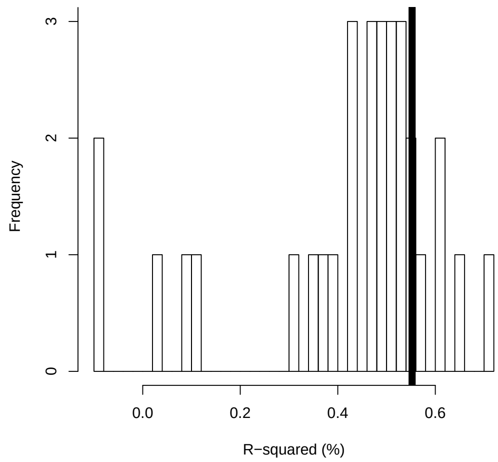

(b) Industry prevailing mean

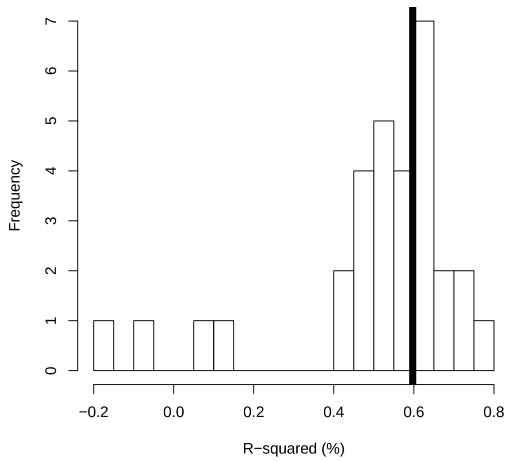

(c) Time-series break model

Figure 8: This figure displays the out-of-sample R-squared values obtained when comparing the forecasting performance of our heterogeneous panel break model with the benchmark model in question for each of the thirty industry portfolios and the market portfolio using the dividend-price ratio as the predictive variable. The benchmark model for each subfigure is described in the subcaption. The thick black vertical line marks the out-of-sample R-squared value for the market portfolio.

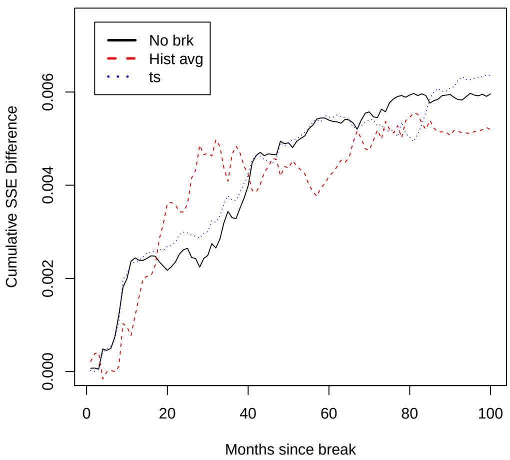

(a) Market portfolio

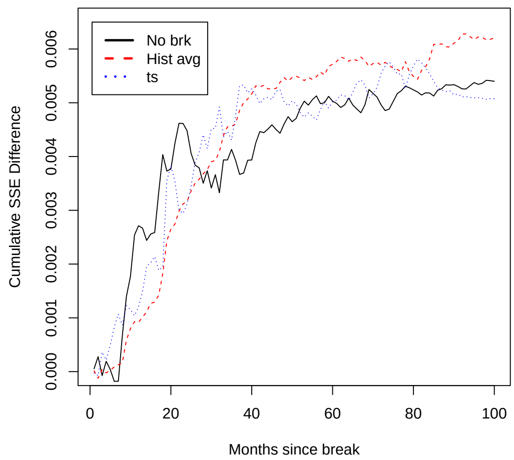

(b) Oil

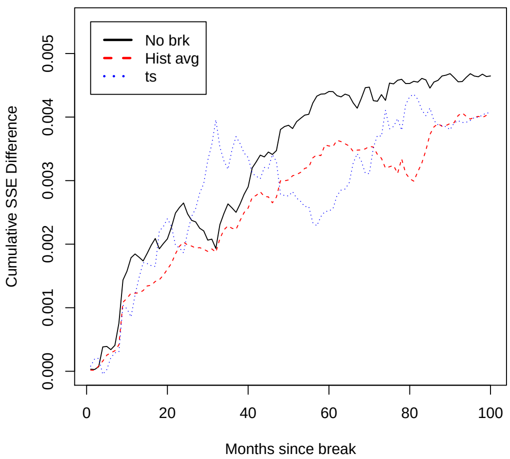

(c) Financials

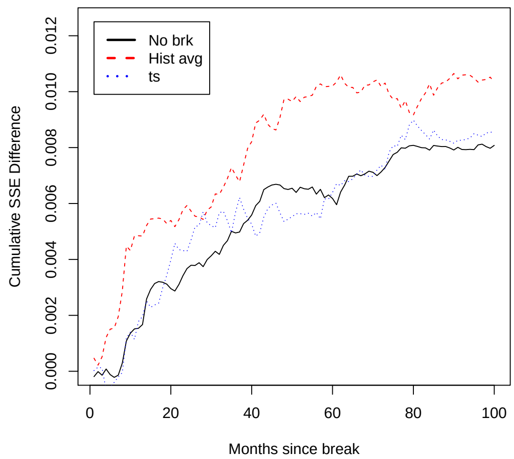

(d) Telecommunications

Figure 9: This figure graphs the cumulative difference in the sum of squared errors for the portfolio in question obtained from our heterogeneous panel break model relative to each of the benchmark models as a function of the time since initial detection of each break. The squared forecast error is computed as a function of the time since the break was initially detected for each of the breaks that onsets over the out-of-sample period, and the average is taken across these breaks. The benchmark models are the heterogeneous panel model with no breaks ('No brk'), the industry prevailing mean ('Hist avg'), and the time-series model with breaks ('ts') estimated using the algorithm of Chib (1998) applied to each portfolio in turn. The dividend-price ratio is the predictive variable and the portfolios being forecast are detailed in the subcaption of each subfigure.

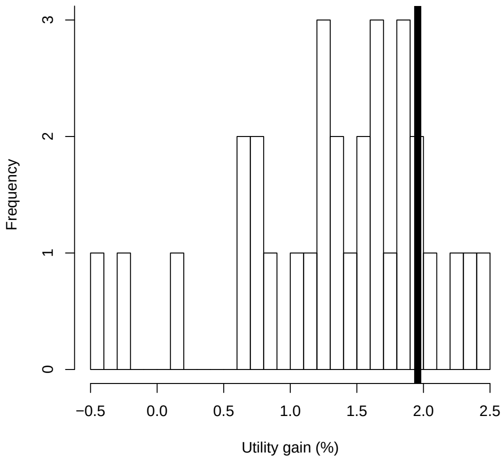

(a) Heterogeneous panel no break model

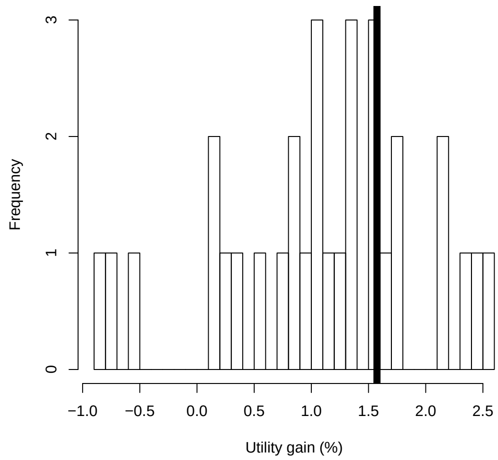

(b) Industry prevailing mean

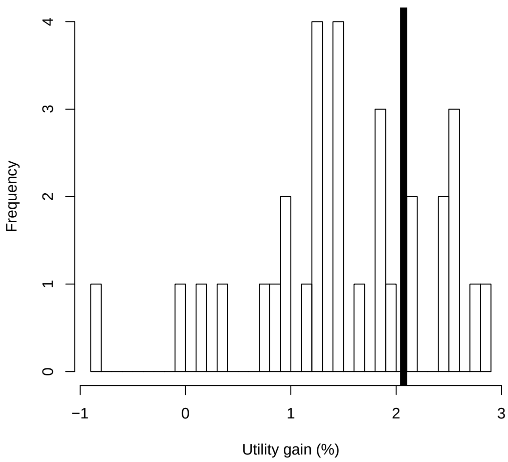

(c) Time-series break model

Figure 10: This figure displays the out-of-sample utility gain (certainty equivalent return) to a mean-variance investor who allocates his wealth between the portfolio in question and the risk-free rate. Utility gains are reported as annualised percentages obtained when comparing the forecasting performance of our heterogeneous panel break model with the benchmark model in question for each of the thirty industry portfolios and the market portfolio when using the dividend-price ratio as the predictive variable. The benchmark model for each subfigure is described in the subcaption. The thick vertical black line marks the utility gain for the market portfolio.

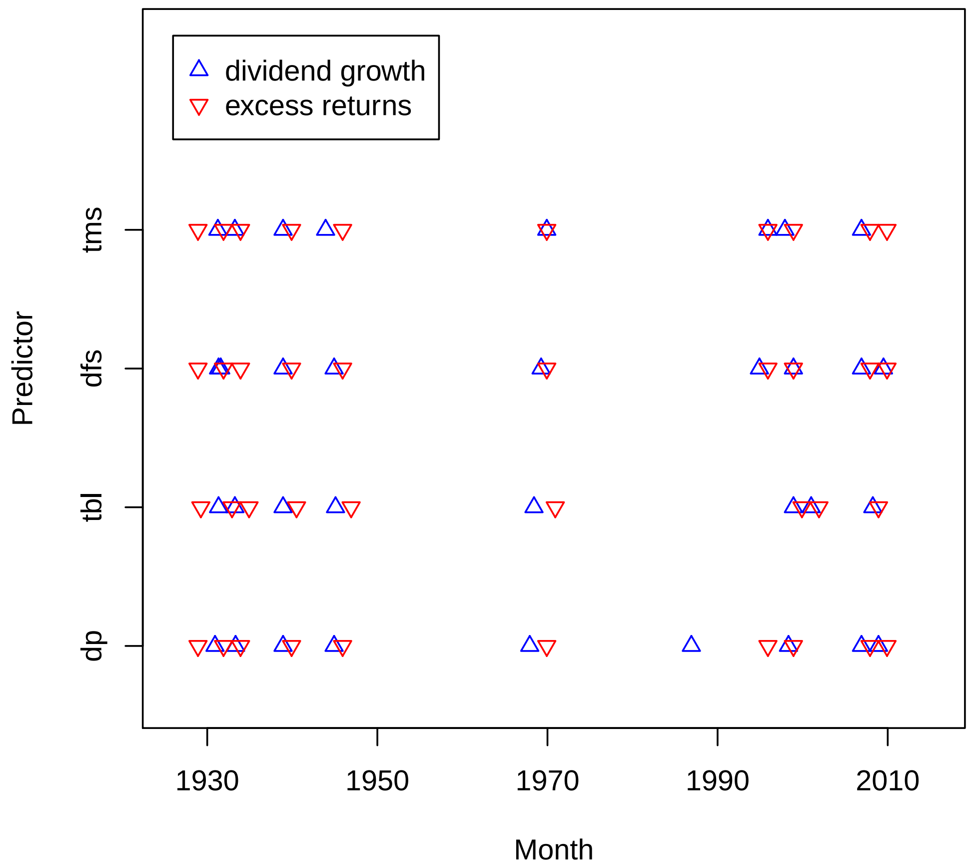

Figure 11: This figure compares the posterior modes of the break dates identified by our panel break model fitted either to excess returns on 30 industry portfolios (red triangles) or to dividend growth for the same set of industries (blue triangles). The four rows show results for different predictors, namely, the dividend-price ratio (dp), the T-bill rate (tbl), the default spread (dfs), or the term spread (tms).

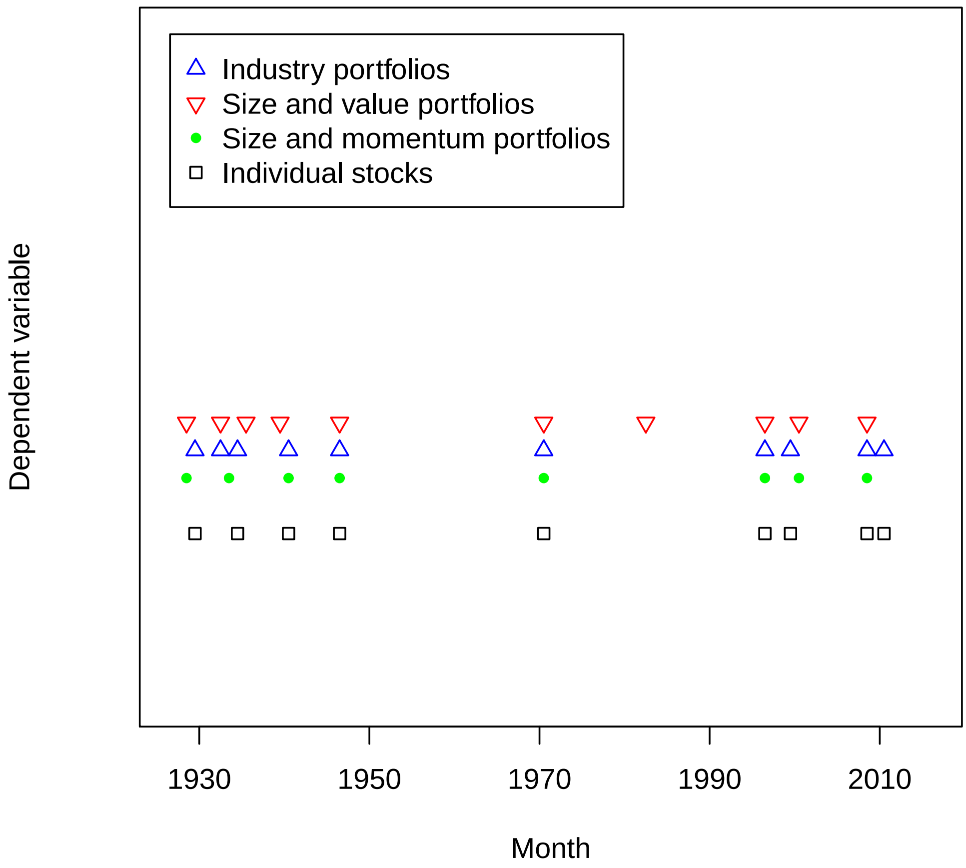

Figure 12: This figure displays the posterior break modes estimated from the heterogeneous panel break model when regressing the dependent variable on an intercept term and the lagged dividend-price ratio. The dependent variable is either the excess returns on the 30 industry portfolios (blue triangles), the excess returns on the 5×5 portfolios sorted on size and value (red triangles), the excess returns on the 5×5 portfolios sorted on size and momentum (green circles), or the 7,299 individual stocks (black squares).

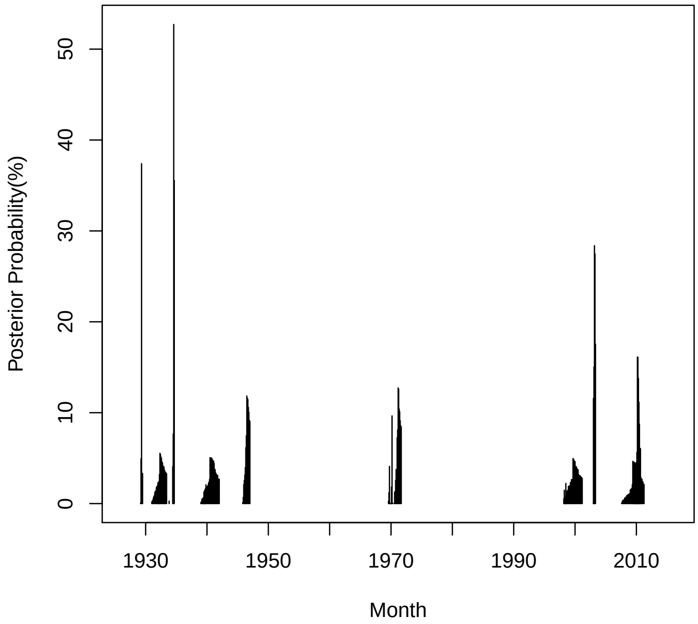

(a) Posterior Break Dates

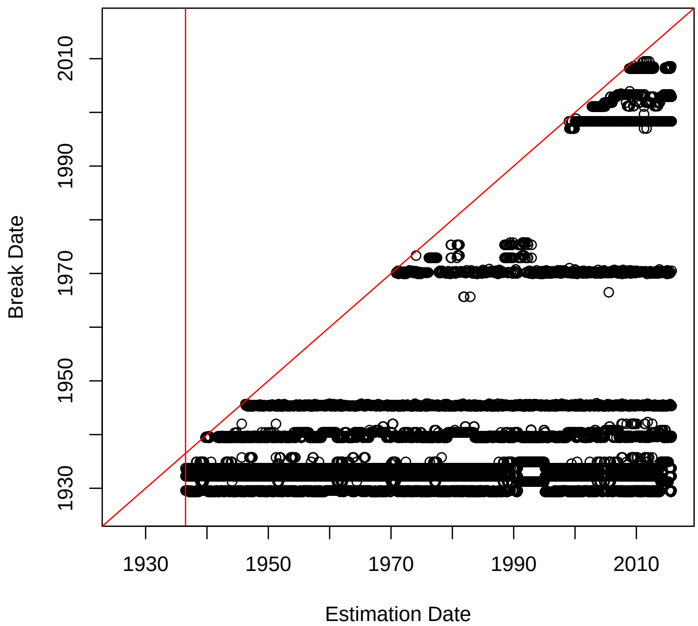

(b) Recursive breaks

Figure 13: This figure displays the posterior probability of the break dates (top panel) and the real-time break detection (bottom panel) obtained from regressing returns on 30 industry portfolios on the lagged treasury-bill rate. In the bottom panel, empty black circles denote the estimated break dates and the vertical red line marks the initial estimation period of ten years. The 45 degree line (to the right of the vertical red line) marks the time at which a break could first be detected.

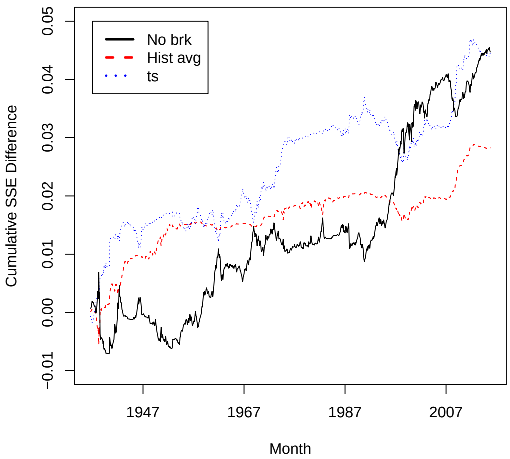

(a) T-bill rate

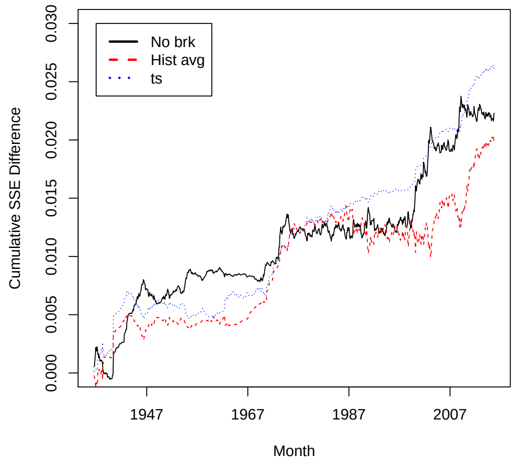

(b) Term spread

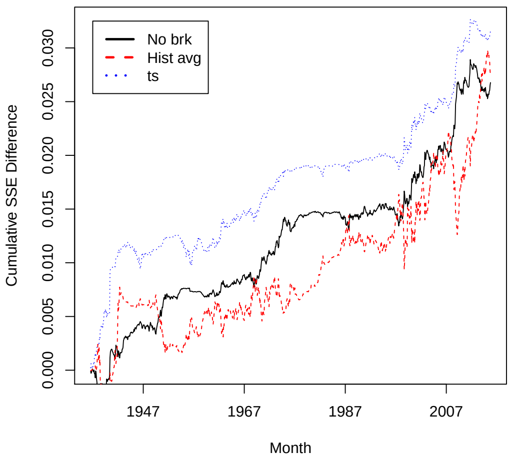

(c) Default spread

Figure 14: This figure graphs the cumulative sum of squared error differences for return forecasts of the market portfolio obtained from our panel break model measured relative to three benchmark models. The benchmark models are the heterogeneous panel model with no breaks ('No brk'), the industry prevailing mean ('Hist avg') and the time-series break model with breaks ('ts') applied to each portfolio in turn. The predictive variable in question is detailed in the subcaption of each subfigure.