Break Risk

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

[a] *USC Dornsife INET, Department of Economics, University of Southern California, Los Angeles, CA 90089, USA*

[b] *University of California, San Diego, La Jolla, CA 92093-0553, USA*

March 20, 2020

Abstract

We develop a new approach to modeling and predicting 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. To identify the economic sources of breaks, we explore the asset pricing restrictions implied by a present value model which links breaks in return predictability to breaks in the cash flow growth and discount rate processes.

Keywords: Predictability of stock returns, break risk, present value model, time-varying risk premia, cash flow growth, panel data, Bayesian analysis

JEL classifications: G10, C11, C15

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

We are grateful to two anonymous referees and the editor (Stijn Van Nieuwerburgh) for numerous suggestions that greatly improved the paper. We also acknowledge the comments of Elena Andreou (discussant at 13th Annual Hedge Fund conference at Imperial), Michael Brennan, Wayne Ferson, Antonio Gargano, John Handley, Bryan Kelly, Arthur Korteweg, Federico Nardari, Andreas Neuhierl (discussant at AFA), Ross Valkanov, Dacheng Xiu, Nancy Xu (discussant at SFS Cavalcade) and seminar participants at National University of Singapore, University of Warwick, Queen Mary University of London, Arizona State University, Richmond Fed, Rutgers University, UCSD, USC, University of Manchester, University of Melbourne, SFS Cavalcade NA 2018, SoFiE 2018, 13th Annual Hedge Fund Conference at Imperial College, 2018 SoFiE Summer School, Inquire 2019 Spring Residential Conference, and AFA 2019. Yida Peng and Ritong Qu provided excellent research assistance. Any remaining errors are our own.

1. Introduction

Attempts to forecast stock market returns are plagued by instability in the underlying prediction models as documented in 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]

Instability in return predictability models can be caused by a variety of factors. The present value model suggests that return predictability may be unstable due to changes in either the cash flow growth or discount rate processes which, in turn, could reflect shifts in investors' risk aversion, assets' exposures to systematic risk factors, or the quantity of risk. Large changes in the macroeconomic environment or changes in institutions, regulations, and public policy can also shift the relation between observable predictor variables and the underlying risk factors and cause model instability. Alternatively, instability can be related to the learning process of investors as we would expect predictable patterns in returns to 'self-destruct' once investors attempt to exploit them.[2] Given these diverse explanations, it is important to gain a better understanding of "break risk" as this helps shed light on the mechanism and economic sources underlying return predictability.

Model instability is economically important because it poses severe challenges to investors' attempts at successfully predicting stock market returns. Using a long historical sample to estimate the parameters of a return forecasting model is not an

[1] Paye and Timmermann (2006) and Rapach and Wohar (2006) conduct 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. Bekaert et al. (2002) use estimates of common structural breaks to date world equity market integration.

[2] 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.

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) can lead to large estimation error and imprecise forecasts.

An alternative strategy is to specify a model for how the parameters of the return prediction model change. However, this approach faces 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 model parameters, leading to volatile and inaccurate return forecasts. Overcoming 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 develop a new 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, effectively enabling breaks to be detected rapidly in real time.[3] 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 economically implausible values. Following Pástor and Stambaugh (1999), we link the prior on the intercept in the return prediction model to the residual volatility to ensure that little weight is assigned to implausibly high Sharpe ratios. Moreover, following Wachter and Warusawitharana (2009), we center the prior on the slope coefficient of the predictor on zero with a relatively tight variance so as to rule out implausibly large degrees of return predictability and high volatility in the return forecasts. This combination of a multivariate (panel) break model and economically motivated (Bayesian) priors is

[3] Bai et al. (1998), Bai (2010), and Baltagi et al. (2016) show that using multiple series in the cross-section results in more precise estimates of break dates relative to using a single time series.

what makes our approach work well empirically.

The key identifying assumption in our analysis which allows us to exploit the benefits from pooling cross-sectional and time-series information on returns is that the timing of breaks is relatively homogeneous across stocks. This assumption is justified when breaks occur in the risk premia of common risk factors or when the quantity of risk shifts at the same time. For example, if a variable ceases to predict returns on the aggregate stock market portfolio, we would expect to find a similar effect on individual stocks or stock portfolios at approximately the same time. Exploring the simultaneous timing of breaks increases our ability to detect breaks and accurately determine their timing.

Our empirical analysis focuses on individual stock returns. Using 90 years of monthly returns data across 14 predictor variables from Goyal and Welch (2008), we find evidence of eight breaks corresponding to roughly one break on average per decade. Breaks affect the returns of many different firms, regardless of firm or industry characteristics, indicating that the breaks we identify are broad-based and systematic in nature, although break sensitivities can vary widely across stocks.

To narrow down the economic sources of breaks to return predictability, we adopt a simple present value asset pricing approach. The present value model implies that return predictability is linked to predictability of dividend growth and/or the dividend-price ratio through a tight set of cross-equation restrictions. Breaks in the coefficients of the return prediction model must therefore be matched by breaks to the coefficients of the dividend growth rate or dividend-price ratio models. To explore if this holds, we use Bayes factors to test the present value model restrictions. We find strong evidence that the three variables experience breaks at the same time and that the restrictions implied by the present value model are supported by the data. Moreover, we provide further evidence on which parts of the return prediction model are affected by individual historical breaks - factor loadings (betas) or risk premia (or both) - through a set of specification tests.

Following earlier studies such as Campbell and Thompson (2008), Rapach et al. (2010), and Dangl and Halling (2012), we assess the predictive accuracy of our return forecasts using a variety of statistical and economic performance measures. We find that out-of-sample 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 with breaks, a time-varying parameter model with slowly moving parameters, a rolling window approach, and a panel model with no breaks. Specifically, our panel-break approach generates significantly more accurate out-of-sample forecasts with improvements in the out-of-sample R^2 value for the market portfolio exceeding 0.5% against all of these benchmarks. An asset allocation analysis for a modestly risk averse mean-variance investor suggests that the return forecasts from the panel break model generate certainty equivalent returns approximately 2% per year higher than each of the benchmarks.

The remainder of the paper is set out as follows. Section 2 motivates and introduces our panel break approach. Section 3 reports empirical evidence of breaks. Section 4 investigates the economic sources of breaks, while Section 5 evaluates the economic and statistical significance of return predictability for our panel break model and compares our approach to a set of alternative methods. Section 6 concludes.

2. Methodology

This section provides economic motivation for our panel break approach to capturing instability in return prediction models, introduces our approach more formally, and explains how we estimate the model parameters. Technical details are presented in a set of appendices. Throughout the analysis, we assume that there are $i = 1, ..., N$ return series and $t = 1, ..., T$ time series observations.[4]

[4] For simplicity, our notation ignores that not all asset returns are observed in all time periods since we are dealing with an unbalanced panel.

2.1. Economic Sources of Instability in Return Predictability

We begin with some economic motivation for our analysis of instability in return prediction models. Suppose the return on asset i in period $t+1$, r_{it+1}, measured in excess of a risk free rate, is affected by a set of common risk factors, f_{t+1}, with factor loadings b_{it} and conditional risk premia, λ_t. Standard conditional asset pricing models imply that $E_t[r_{it+1}] = b'_{it}\lambda_t$ and so we can decompose r_{it+1} as follows:

$$r_{it+1} = b'_{it}\lambda_t + \epsilon_{it+1}, \tag{1}$$

where $\epsilon_{it+1} = b'_{it}(f_{t+1} - E_t[f_{t+1}]) + u_{it+1}$ is an unpredictable return component and u_{it+1} is an idiosyncratic (asset-specific) shock with zero mean, so $E_t[\epsilon_{it+1}] = 0$. This simple model suggests that return predictability can arise either from time-variation in betas or from time-varying risk premia.

Next, suppose that variation in the common risk premia, λ_t, is driven by a set of observable state variables, x_t, with coefficients θ_λ, so $\lambda_t = \theta'_\lambda x_t$. State variables, x_t, commonly used by the finance literature to track time variation in risk premia include the dividend-price ratio, short-term interest rates, and default spreads. The relation between these variables and risk premia (θ_λ) is unlikely to remain constant over time. For example, the relation between short interest rates and risk premia was presumably very different during the zero-lower bound state than in more normal times. Similarly, studies such as Rapach et al. (2010), Henkel et al. (2011), and Dangl and Halling (2012) find that the relation between predictor variables and expected stock returns is very different in recession and expansion states. Building on these observations, it is important to allow the mapping from predictors, x_t, to risk premia, λ_t, to vary over time. We can achieve this by allowing this mapping to change at a set of break dates τ_k for $k = 1, ..., K$ with $\theta_{\lambda t} = \theta_{\lambda k}$ for $\tau_{k-1} < t \leq \tau_k$, so that

$$\lambda_t = \theta'_{\lambda k} x_t \qquad \tau_{k-1} < t \leq \tau_k. \tag{2}$$

This approach is flexible and allows for repeated states (e.g., recessions and expansions) or small breaks every period (a time-varying parameter model) as special cases. The assumption that the timing of breaks, τ_k, is common across stocks is a natural one in this context as breaks to the risk premium of the common factors will affect the vast majority of stocks.

Turning to the factor loadings, a common approach for capturing time-variation is to let the betas be a (linear) function of a set of observable stock or firm characteristics, c_{it}, with coefficients θ_b, so that $b_{it} = \theta'_b c_{it}$, see, e.g., Gu et al. (2018). Again we can capture possible time-variation in the mapping from firm characteristics to factor loadings through a step function, $\theta_{bt} = \theta_{bk}$ for $\tau_{k-1} < t \leq \tau_k$:

$$b_{it} = \theta'_{bk} c_{it} \qquad \tau_{k-1} < t \leq \tau_k. \tag{3}$$

The specifications in (2) and (3) are quite general and so we next consider some special cases of particular economic interest. First, suppose that firm characteristics, c_{it}, can vary widely in the cross-section but only move slowly over time so that $c_{it} \approx c_i$.[5] Using equation (1), the case with static betas and time-varying risk premia leads to the return model

$$r_{it+1} = \beta'_{ik} x_t + \epsilon_{it+1} \qquad \tau_{k-1} < t \leq \tau_k, \tag{4}$$

where $\beta'_{ik} = b'_i \theta'_{\lambda k}$, so the mapping from observable predictors, x_t, to expected returns can vary across stocks and breakpoints, τ_k. Breaks in return predictability identi-

[5] Assuming that we can sort stocks into portfolios, p, with similar characteristics for all stocks, $i \in p$, we have $b_{it} \approx b_p$ and so can conduct a similar analysis for portfolios of stocks.

fied by this specification are, thus, driven by changes in the aggregate risk premium process. This model is economically interesting and a good candidate for identifying breaks since it seems a priori plausible to expect the aggregate risk premium to be different in economic states such as the global financial crisis.

Another interesting special case arises when risk premia are constant, $\lambda_t = \lambda$, while stock betas are allowed to vary over time according to equation (3) so that

$$r_{it+1} = \beta_{ik}' c_{it} + \epsilon_{it+1}, \qquad \tau_{k-1} < t \leq \tau_k, \tag{5}$$

where $\beta_{ik}' = \lambda' \theta_{ik}'$. Here, the mapping from firm characteristics, c_{it}, to expected returns can vary over time in a way that allows us to identify common breaks.

By jointly modeling returns on individual stocks, our panel approach is well-suited for identifying breaks that simultaneously affect a broad cross-section of stocks. Moreover, recalling the definition of the return shocks, $\epsilon_{it+1} = b_{it}'(f_{t+1} - E_t[f_{t+1}]) + u_{it+1}$, breaks to factor loadings ("betas") will typically produce breaks in both expected returns and in the variance-covariance of returns. Breaks in omitted risk factors will have the same effect, whereas breaks to the risk premia of the included risk factors should only affect expected returns.

2.2. A Panel Break Model for Stock Returns

We next propose a model that builds on the above economic motivation. Identifying breaks from a single time series tends to be very hard. Conversely, assuming that the timing of breaks is the same across individual stocks can greatly increase the power of break tests and, thus, enhance our ability to estimate the location and size of any breaks. Provided that the individual return series are not too strongly correlated, each return series yields valuable additional information that helps us identify breaks.[6] Our panel approach therefore estimates breaks by pooling the information from the cross-section to identify the *timing* of the common breaks, while still estimating the parameters for each individual return series:[7]

$$r_{it+1} = \mu_{ik} + \beta_{ik} X_t + \epsilon_{it+1}, \qquad \tau_{k-1} < t \leq \tau_k, \qquad k = 1, \ldots, K+1, \tag{6}$$

where ϵ_{it+1} is a zero-mean unexpected return component. Allowing for an intercept in (6) means that our model nests the constant expected return ("no predictability") benchmark model of Goyal and Welch (2008) as a special case. Note that our model is flexible enough to allow for partial breaks such that any subset of the regression coefficients in (6) may shift following an identified break. For example, the intercept may be the only parameter to shift following the first break, while all coefficients could shift following the second break.

Stock returns typically exhibit high levels of cross-sectional covariance due to their loadings on common risk factors. Indeed, breaks to factor loadings, b_{it}, or breaks in the coefficients of omitted common factors can lead to simultaneous breaks in the covariance of the residuals from the panel return prediction model. While these covariances do not directly affect our forecasts of expected stock returns, ignoring them can reduce our ability to detect breaks with panel data (Kim 2011; Baltagi et al. 2016). We therefore allow the covariance structure of return shocks, including the variance of ϵ_{it+1}, to vary across return series and break locations $\tau = (\tau_1, \ldots, \tau_K)$.

In practice, once we allow for breaks to the covariance of returns, estimating the full covariance matrix of residuals becomes an unattractive strategy - not only because such estimates are likely to be imprecise, but also because it severely delays

[6] Bai et al. (1998) make a similar point in the context of a small-scale vector autoregression. We show that their result is even stronger for panel data with a large cross-sectional dimension.

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

break point detection.[8] Instead, we model correlations across residual returns as generated by a small set of factors. In our empirical application, we use a single (market) factor to capture correlation in the residuals of the individual asset return series. Empirically, this works very well. For example, the absolute value of the pairwise correlation averaged across the individual stocks in our empirical application is reduced from 0.82 to 0.16 after allowing for a single factor, while the cross-sectional dependence test statistic of Pesaran (2004) is reduced from 193.74 to 2.27, which is no longer significant at the 1% level. For robustness, we also conduct results allowing for three or five Fama-French factors and obtain similar results throughout the analysis. Appendix C contains further discussion of this point.

2.3. Prior Distributions

To estimate our model, we adopt a Bayesian methodology which 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.[9] The hyperparameters that determine the frequency of breaks to the coefficients are set so that a break is expected to occur roughly once per decade, but we also conduct an analysis with more conservative priors implying breaks that, on average, occur every 20 years.

Importantly, we let the key prior parameters be economically motivated. Given empirical evidence of weak return predictability such as Goyal and Welch (2008), we center our prior for β in (6) at zero. Moreover, inspired by Wachter and Warusaw-

[8] To see this, note that having only $N = 30$ return series 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.

[9] Further details of the shape of the priors are provided in Appendix B.

itharana (2009), we explore an economically motivated prior distribution that allows investors to have 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. Bayesian analysis allows different degrees of predictability, reflecting the scepticism of the investor as to whether 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 that the investor is sceptical about predictability but does not rule it out entirely.

Wachter and Warusawitharana (2009) note that 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. We therefore 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 sample available at the time the recursive forecast is made.

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.

For a single risky asset the proportion of the total variance that originates from variation in the predictable component of the return is

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

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 1 displays how investors assign more weight to a positive R_i^2 as σ_η increases. Specifically, when $\sigma_\eta = 0.02, 0.04$, and 0.06, investors assign probabilities to R_i^2 exceeding 0.005 of 0.0003, 0.075, and 0.235, respectively. 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 our results when this parameter is adjusted.

It is also desirable to specify that high Sharpe Ratios are a priori unlikely. A high absolute value of the intercept term μ_{ik} 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 stock 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 standard deviation σ_μ equal to 5%.[10]

[10]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.

2.4. Posterior Distribution and Estimation

Combining the priors and likelihood, described in Appendix A, we obtain the posterior distribution. A methodological contribution of our paper is to show how specifying a fully conjugate model allows the parameters to be marginalised from the posterior. This enables the breakpoints to be estimated independently of the parameters, as we prove formally in Proposition 1 in Appendix B. This result considerably reduces the computational burden of estimation of our panel break model (Appendix D), which is important when using panel data involving large values of N and T.

3. Empirical Evidence of Breaks

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

3.1. Data

Our analysis of firm-level returns uses monthly CRSP data on individual US stocks traded on the NYSE, AMEX or NASDAQ stock exchanges at some point during the sample period from July 1926 through December 2015. In total, we have data on 24,743 different stocks. Returns are computed in excess of a one-month T-bill rate.

Our predictors mostly come from Amit Goyal's website and are constructed following Goyal and Welch (2008). The aggregate dividend-price ratio (dp) uses 12-month moving sums of dividends on the $S\&P$ 500. The other predictors are the one-month Treasury-bill rate (tbl), term spread (tms), default spread (dfy), earnings-price ratio (ep), dividend payout ratio (de), book-to-market ratio (bm), default return spread (dfr), long term yield (lty), inflation (infl), long term return (ltr), stock variance (svar), corporate issuing activity (ntis), and cross-sectional premium (csp). Data on the cross-sectional premium (csp) run from January 1953 through December 2015 and are constructed following the src factor described in Polk et al. (2006).

3.2. Break numbers and locations

Figure 2 displays our findings for the number of breaks and their location using either the unrestricted return forecasting model fitted to firm-level data with the aggregate dividend-price ratio as a predictor and priors implying a break frequency of 10 years (left panels) or the present value model fitted to the industry portfolio returns (right panels). The location for most of the breaks is well defined with clear spikes in the posterior probabilities shown in the middle windows. The present value model identifies five breaks occurring in 1934, 1939, 1969, 1998, and 2010. The return prediction model detects eight breaks. Five of these breaks occur at approximately the same time as for the present value model with the additional three breaks occurring at 1929, 1948 and 2002. Imposing the cross-equation restrictions in the present value model tempers the break detection relative to the unconstrained return prediction model.[11]

Our plots identify long periods without any evidence of model instability, e.g., the twenty eight-year period from 1940 to 1968. Such periods are displayed by the horizontal parts of the line in the bottom window of Figure 2 which plots the cumulative posterior break probability.

For comparison, we also estimate break dates by applying the frequentist panel break approach of Baltagi et al. (2016) to regressions of firm-level returns on the dividend-price ratio. This model identifies six breaks - displayed as red triangles in the middle panels of Figure 2 - at times similar to our Bayesian approach although it does not detect breaks in 1940 or 1998.

[11] Using a 20-year prior on the break frequency for the unconstrained baseline return prediction model leads to five breaks being detected at approximately the same dates as for the present value model as the threshold for detecting new breaks is raised.

In turn, the breaks identified by these panel approaches are very different from those obtained from the breakpoint algorithm of Chib (1998) applied to the univariate time series of firm-level stock returns: the univariate breakpoint model fails to detect a single break for 82% of the stocks, instead favoring the model with zero breaks, suggesting that the univariate tests have too weak power to identify breaks off individual return series.

For comparison, Lettau and Van Nieuwerburgh (2008) find evidence of breaks in 1954 and 1995. The latter is close to our break date. They use annual data from 1927-2004, whereas we use monthly data up to 2016. Viceira (1997) finds evidence of structural instability in long-horizon return regressions after the second world war. Paye and Timmermann (2006) report evidence of breaks in 1962 and 1974 in monthly stock return regressions on the T-bill rate and also detect a break in the relation between the dividend yield and stock returns in 1994. Rapach and Wohar (2006) also find evidence of breaks in monthly return equations. Both studies document that the break dates vary substantially depending on which predictor variables are used. This is to be expected since the breaks reflect time variation in the covariance between the predictor and stock returns. Pástor and Stambaugh (2001) identify 16 different stable regimes in the equity premium process corresponding to 15 breaks or transitions, during a 160 year sample from 1834-1999. Although their sample is quite different from ours, they identify breaks around some of the dates for which we find breaks, including the Great Depression (1932), 1940s, and the early seventies and mid-nineties. The latter is a bit earlier than the break identified by us which can be explained by the earlier end-date of their sample.

3.3. Real-time Detection of Breaks

A key challenge for return predictability is how quickly the model is able to identify breaks in real time. Severe delays in break 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 3a plots the break dates estimated in real time. The figure uses the lagged dividend-price ratio as the predictor but similar plots are obtained using other predictor variables. The real-time breakpoint detection performance 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 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”, of which there are not too many instances.

Figure 3b 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. Across four common predictor variables, all breaks are detected within three to eight months of their occurrence.

This ability of our panel break approach to identify breaks with little delay again stands in marked contrast to the long delays typically associated with break modeling for univariate time-series. This is an important point as emphasized by Lettau and Van Nieuwerburgh (2008) and Viceira (1997) who find that it is difficult to exploit model instability to generate more accurate out-of-sample return forecasts.

3.4. Variation in slope coefficients and the market risk premium

Having identified both the number and location of breaks in the return prediction model, we consider the magnitude of the breaks to the slope coefficient of the dividend-price ratio along with the resulting effect on the expected market risk premium.

Figure 4 illustrates the magnitude of the breaks in the slope coefficient of the dividend-price ratio. The top-left window plots the recursively estimated value-weighted posterior mean of the slope coefficient from panel break regressions of firm-level returns on the aggregate dividend-price ratio. These estimates could have been used by investors in real time to generate return forecasts. With estimates ranging between zero and 1.5 and centered close to one, the values of the slope coefficients are economically sensible. We observe a notable downward trend in the slope coefficient during the 1990s followed by a marked increase during the last few years of the sample.[12]

A common approach to handle parameter instability is to use a rolling estimation window and so it is useful to compare our panel break estimates to results from this approach. To this end, the bottom left window of Figure 4 plots the value-weighted average of the slope coefficients from rolling 10-year time series OLS regressions of firm-level stock returns on the lagged aggregate dividend-price ratio. The posterior mode of the break dates estimated from our baseline Bayesian panel break model are marked by the vertical lines in the figure. Interestingly, many of the large, discrete shifts in the rolling-window estimates are closely lined up with our break estimates.

[12] Note that the parameter estimates undergo some periods in which they are slow-moving while also experiencing rapid spikes that resemble breaks. Reassuringly, these spikes correspond closely to the breaks identified by our baseline model.

However, variation in the rolling-window estimates is nearly an order of magnitude larger than for our panel break estimates and the rolling window estimates turn negative both in the late 1990s and during the Global Financial Crisis.

The out-of-sample estimates of expected returns from the panel break model (top right panel) vary in a range between -0.7% and 3.9% per month. The most notable feature is the high volatility prior to WW2 and the systematic drop in expected returns during the 1950s, after which expected returns are quite stable around 0.5% per month except for a brief spike during the late 1970s. In contrast, the expected returns implied by the rolling window estimates (bottom right panel) are twice as volatile (the forecasts have a standard deviation of 0.012 compared with 0.006), cover a much wider range between plus and minus 7%, and are very spiky.

4. Economic Sources of Breaks

This section attempts to identify the economic sources of the breaks. First, we use a simple reduced-form present value model that links breaks in the return prediction model to breaks in the cash flow and discount rate processes. Next, we investigate whether breaks in the return prediction model can be attributed to breaks in factor loadings, breaks in risk premia, or both. Finally, we summarize some of the possible economic sources of breaks.

4.1. A present value model analysis

To understand the economic sources of breaks in an asset pricing (return predictability) context, it is worth recalling that the conventional linear, constant-coefficient prediction model relating asset returns to a set of state (predictor) variables can be derived under a set of basic assumptions: (i) asset prices are determined by a standard no-arbitrage Euler equation; (ii) the aggregate state of the economy can be represented by a stationary (constant coefficient) vector autoregressive process; either (iiia) risk prices are constant and risk quantities can be time-varying but the conditional moment generating function of the state variables determining their dynamics has an affine structure; or (iiib) risk prices are time-varying with multivariate normally distributed and homoskedastic innovations. These assumptions hold for a broad class of asset pricing models, including models in which agents have Epstein-Zin recursive preferences and aggregate consumption growth is an affine function of the vector of state variables, as well as models with time-varying risk-aversion.[13]

Breaks to the regression coefficients of return prediction models can, therefore, arise as a result of breaks in the processes determining either (i) risk prices; (ii) risk quantities; or (iii) the dynamics (e.g., persistence) of the underlying state variables. Breaks to risk premia fall in the first category, while breaks in betas or in volatility parameters fall in the second category, reflecting a shift in the amount of systematic risk in the economy. Breaks to the mean, persistence, or volatility of the dividend growth process fall in the third category although they can also affect risk prices and quantities.

To make the implications of these broad points more concrete in a format that lends itself to empirical testing, we next conduct an analysis of breaks in the context of a simple reduced-form present value model which is widely used throughout the asset pricing literature. Specifically we use the asset pricing restrictions implied by this model to trace back breaks in the return prediction model to breaks in the underlying dividend growth and dividend-price equations.

Before proceeding further, it is worth pointing out some limitations to our analysis. Most importantly, in line with the vast majority of empirical studies on return predictability, we take a partial equilibrium perspective and do not account for the general equilibrium implications that breaks to return predictability could have. In particular, we do not model how agents account for how such breaks can affect their

[13]For further discussion and derivations, see, e.g., Lustig et al. (2013) and Farmer et al. (2019).

own future posterior beliefs and so ignore the effect of breaks on agents' future learning. Asset pricing models that account for such effects can be solved only for a narrow class of models, see, e.g., Collin-Dufresne et al. (2016), and addressing this point - while clearly very interesting - falls beyond the scope of our analysis.

We first introduce some notations. Let p_{t+1} denote the logarithm of the stock price at the end of month $t+1$, while d_{t+1} is the logarithm of the dividends during month $t+1$. Further, let $\Delta d_{t+1} = d_{t+1} - d_t$ denote the growth rate in dividends, while $d_t - p_t$ is the log dividend-price ratio, and r_{t+1} denotes the log-return in period $t+1$.

The approximate log-linearized present value model proposed by Campbell and Shiller (1988) relates stock returns to the dividend-price ratio and dividend growth:

$$r_{t+1} \approx \kappa - \rho(d_{t+1} - p_{t+1}) + \Delta d_{t+1} + (d_t - p_t), \tag{11}$$

where ρ and κ are log-linearization constants.

State variables that can forecast Δd_{t+1} and $d_{t+1} - p_{t+1}$ should, according to (11) also help forecast returns. To analyze if this holds, we generalize the analysis in Cochrane (2008) to allow for return predictability not only from $d_t - p_t$, but also from the lagged growth rate of dividends, Δd_t, through the simple predictive equations

$$\begin{aligned} d_{t+1} - p_{t+1} &= a_0 + a_1(d_t - p_t) + a_2\Delta d_t + \epsilon_{dp,t+1} \\ \Delta d_{t+1} &= b_0 + b_1(d_t - p_t) + b_2\Delta d_t + \epsilon_{d,t+1}. \end{aligned} \tag{12}$$

From the present value identity (11) and equation (12), we have

$$r_{t+1} = c_0 + c_1(d_t - p_t) + c_2\Delta d_t + \epsilon_{r,t+1}, \tag{13}$$

where the coefficients in (13) are subject to the cross-equation restrictions

$$
\begin{aligned}
c_0 &= \kappa - \rho a_0 + b_0, \\
c_1 &= 1 + b_1 - \rho a_1, \\
c_2 &= b_2 - \rho a_2.
\end{aligned} \tag{14}
$$

The present value model in (13) implies that variables that can predict dividend growth or the log dividend-price ratio should also, in general, be expected to forecast returns with coefficients that are subject to the restrictions in (14).

Because the cross-equation restrictions in (14) must hold at each point in time, breaks in the c coefficients in the return prediction model should be aligned with breaks to the corresponding a or b coefficients in the dividend-price ratio and dividend growth models. This suggests modifying (12) to allow for breaks:

$$
\begin{aligned}
d_{t+1} - p_{t+1} &= a_{0k} + a_{1k}(d_t - p_t) + a_{2k}\Delta d_t + \epsilon_{dp,t+1}, \qquad \tau_{k-1} < t \leq \tau_k, \\
\Delta d_{t+1} &= b_{0k} + b_{1k}(d_t - p_t) + b_{2k}\Delta d_t + \epsilon_{d,t+1}, \qquad \tau_{k-1} < t \leq \tau_k,
\end{aligned} \tag{15}
$$

where

$$
r_{t+1} = c_{0k} + c_{1k}(d_t - p_t) + c_{2k}\Delta d_t + \epsilon_{r,t+1}, \qquad \tau_{k-1} < t \leq \tau_k \tag{16}
$$

subject to a set of cross-equation restrictions that account for breaks:

$$
\begin{aligned}
c_{0k} &= \kappa - \rho a_{0k} + b_{0k}, \\
c_{1k} &= 1 + b_{1k} - \rho a_{1k}, \\
c_{2k} &= b_{2k} - \rho a_{2k}.
\end{aligned} \tag{17}
$$

Using these restrictions, we can proceed to relate breaks in dividend growth and/or the dividend-price ratio processes to breaks in the return prediction model. Before doing so, note that dividend growth data tends to be very lumpy and irregular at the firm level, see, e.g., Pettenuzzo et al. (2019). We therefore aggregate dividend data to form 30 industry portfolios and estimate the prediction models for the dividend-price ratio and dividend growth variables in (15) and the return series in (16) subject to the cross-equation restrictions in (17). The resulting panel involves 90 time series.

To explore whether breaks to the dividend growth prediction equation in (15) are a source of breaks in the return equation, first consider the results from estimating this model without imposing the cross-equation restrictions in (17), using the lagged dividend growth rate and the lagged dividend-price ratio as predictors. The top panel in Table 1 displays the posterior mean and standard deviation of the estimated intercept, slopes on the dividend-price ratio and lagged dividend growth, and the volatility parameter obtained from our panel break model fitted to the 30 industry portfolios. We find evidence of nine breaks during the sample. The predicted dividend growth rate varies in a wide range that spans high-growth states with a large positive intercept and persistence coefficient (regime eight) and states with negative expected dividend growth (regimes one and three). The coefficient on lagged dividend growth 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.[14]

These results are interesting in their own right as predictability of dividend growth is still being contested. For example, Cochrane (2008) argues that there is little evidence that dividend growth can be predicted. Conversely, Chen (2009), van Binsbergen et al. (2010), and Kelly and Pruitt (2013) present evidence of dividend growth predictability. Our results suggest that disagreement regarding dividend growth predictability might be linked to time variation in predictability, i.e., dividend growth is

[14]The only case where the lagged dividend growth and dividend-price ratio coefficients have unexpected signs is in regime 2. This reversal of signs can sometimes happen in short-lived regimes due to collinearity between the regressors.

predictable in certain regimes while largely unpredictable in other ones.

The bottom panel in Table 1 displays results from the prediction model fitted to the dividend-price ratio. For this variable, we find evidence of six breaks with dates that closely mirror those separately detected from the return and dividend growth predictive regressions. The slope on the lagged dividend-price ratio remains relatively high across breaks, with values ranging between 0.70 and 0.90, while the slope on the lagged dividend growth term is negative and significant, except for in the short-lived period ending in June 1929.

To get an initial sense of whether breaks to predictability of dividend growth or the dividend-price ratio can help explain breaks to the return prediction equation, Figure 5 displays posterior mode break dates estimated separately for the return (black triangles), dividend growth (blue triangles) and dividend-price ratio (green triangles) prediction models. The three sets of break dates are clearly related although some break dates show up for the return model but not for the dividend growth model, and vice versa.

To see the effect of imposing the economic cross-equation restrictions (17) implied by the present value model, the bottom row in Figure 5 shows (red triangles) the break dates identified from the model that jointly estimates the present value prediction equations subject to the cross-equation restrictions. We now detect five breaks in the sample which occur in 1934, 1939, 1969, 1998, and 2010. These dates are similar to the breaks identified by the return prediction model except there is no break identified in 1929, 1948 or 2002.

This plot shows that imposing the economic constraints implied by the present value model has the effect of reducing the number of breaks. By jointly estimating our panel break model for the dividend growth, dividend-price ratio, and return series, breaks that are not supported by the data in the sense that they are not affecting all three series at the same time are less likely to be selected by our procedure, and so the asset pricing restrictions implied by the present value model serve the role of reducing the number of breaks from eight to five.

To more formally test the present value model in the presence of breaks, we next use Bayes Factors to determine whether (i) the three present value prediction equations experience common breaks and (ii) the regression coefficients satisfy the present-value model restrictions in (17). For the former, we estimate the present value model with common breaks and compare this to the same model that allows breaks across the three equations to be non-common (but with the restriction that if one equation is hit then all 30 industries are hit at the same date) using the method of Smith (2018). Next, we compute marginal likelihood values for both models using the procedure of Chib (1995). Finally, we construct Bayes Factors (BFs) for the common break model relative to the non-common model. Bayes Factors are the preferred Bayesian model comparison method. Values between 1 and 3 are inconclusive, values between 3 and 20 indicate positive evidence in favour of the restricted model, while values between 20 and 150 indicate strong evidence in support of the restricted model (Kass and Raftery 1995).

Empirically, we find a Bayes Factor of 69.44. This represents strong evidence in favour of the breaks to the prediction equations (15) - (16) being common against the (less restrictive) alternative of non-common breaks. This suggests that the timing of breaks affecting the dividend growth, dividend-price ratio and return processes is the same.

Next, we construct a Bayes Factor to determine whether the coefficients of the three prediction equations satisfy the present-value model cross-equation restrictions. To achieve this, we estimate two models with common breaks: the first imposes the present value restrictions in (17), while the second does not. Again, we construct a Bayes Factor for the model with the coefficient restrictions imposed relative to the model without these restrictions. We obtain a Bayes Factor of 73.89 which indicates strong evidence in support of the validity of the coefficient restrictions in (17).[15]

[15] We also perform the same test of the present value model, but without allowing for breaks. This

Having verified that the cross-equation restrictions implied by the present value model are supported by the data, we display the posterior mean and standard deviations of the resulting parameter estimates in Table 2. The table shows that a_1 is significant in regimes 3, 4 and 6, while b_1 and c_1 are significant in regimes 3, 4, 5, and 6. Volatility peaks during the Great Depression and Global Financial Crisis. Both the dividend-price ratio and the dividend growth processes are quite persistent with coefficients that fall in the 0.70-0.90 range. The coefficient on the dividend-price ratio in the dividend growth equation is always negative - consistent with a higher dividend price ratio being driven by lower stock prices reflecting worsening dividend growth prospects.

4.1.1. Breaks in risk aversion

While we can construct good proxies for cash flows (dividends), we do not directly observe discount rates. However, the present value model implies that breaks in the dividend-price ratio can be traced back to breaks in either the dividend growth or the discount rate processes (Lettau and Van Nieuwerburgh 2008).

To explore if breaks to discount rates play a role in explaining breaks in return predictability, we consider a simple approach. Investors' coefficient of risk aversion is of course unobserved, but following Merton (1980) we can obtain an estimate of it by regressing stock returns on a proxy for the realized stock market variance which is one of the predictor variables used by Goyal and Welch (2008). Figure 6 graphs the implied risk aversion coefficient (solid black line) obtained as the posterior mean of the pooled slope coefficient estimated from a panel regression of stock returns on the lagged variance which imposes the five break dates estimated from the present value model. The red dashed line graphs the t-statistic of the slope coefficient. The coefficient of risk aversion takes on economically plausible values, falling in a range

generates a Bayes Factor of 80.19 indicating strong evidence in favour of the coefficient restrictions.

between one and two, with t-statistics that are highly statistically significant in all regimes other than between 1934 and 1939. The implied risk aversion coefficient is quite high (exceeding two) between 1927 and 1934, before dropping to a level below 1.5 between the breaks in 1939 and 1969. Risk aversion reaches a low point near unity between 1970 and 1998, before rising to levels near 1.5 and 2 following the breaks in 1998 and 2010.[16]

4.2. Breaks in factor loadings versus risk premia

Whether breaks to return prediction models simultaneously affect both the price and quantity of risk rather than only one of these determinants of expected returns is not clear from asset pricing theory. To formally evaluate whether each of the five breaks identified by the present value model is associated with a break in (i) beta loadings, (ii) risk premia, or (iii) both, we estimate three models that allow for breaks in (i) betas only, (ii) risk premia only and (iii) both. Each of these three models is estimated K times, each time imposing one of the known break dates from the baseline model. The models do not allow for any further breaks to be detected. Next, we construct two Bayes Factors (BFs) that quantify the strength of evidence in favour of a break in only betas or in risk premia. This is repeated for the K breaks, thus giving us a set of $2 \times K$ Bayes Factors.

For each break, if there is a break in betas only, the corresponding BF will be large and the other small. The same is true for a break only to the risk premia. If both BFs are small, both betas and risk premia have experienced a break.

We estimate three separate regressions using our panel break model. First, we regress firm- or portfolio-level stock returns on their corresponding book-to-market

[16] In line with previous studies (Rapach et al. 2010; Henkel et al. 2011; Dangl and Halling 2012), we find evidence of stronger return predictability in recessions than in expansions. This supports the assertion of Campbell and Cochrane (1999) and others that time-variation in risk aversion is an important determinant of risk premia and a source of return predictability.

ratios to analyze breaks in beta loadings:

$$r_{it} = \mu_{ik} + \beta_{ik} BM_{it-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k. \tag{18}$$

Second, we regress firm- or portfolio-level returns on the aggregate dividend-price ratio to test for breaks in risk premia (λ):

$$r_{it} = \mu_{ik} + \lambda_{ik} dp_{t-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k. \tag{19}$$

Third, we regress firm- or portfolio-level returns on both the book-to-market ratio and the aggregate dividend-price ratio, allowing for breaks in both the betas and risk premia:

$$r_{it} = \mu_{ik} + \beta_{ik} BM_{it-1} + \lambda_{ik} dp_{t-1} + \epsilon_{it}, \qquad t = \tau_{k-1} + 1, \ldots, \tau_k. \tag{20}$$

We then compute Bayes Factors of the marginal likelihoods from equations (18) and (19) relative to that from (20). We perform this separately for each of the K breaks identified by the present value model that we impose on a model that regresses industry portfolio returns on the lagged aggregate dividend-price ratio and industry-level book-to-market ratios. The first and second rows in each panel of Table 3 displays the Bayes factors followed by the regime-specific β and λ coefficients with t-statistics (in brackets below). The two breaks associated with the Great Depression (1934) and Global Financial Crisis (2010), as well as the break in 1969, affect both the factor loadings and the risk premia of the model, while the breaks in 1939 and 1998 affect only the betas. Moreover, risk premia are notably higher coming out of the Great Depression (1934) and after the Global Financial crisis (2010-2015). The lower panel of Table 3 shows that we arrive at similar conclusions when we use firm-level stock returns.[17]

[17] An analysis of the time variation in rolling window coefficient estimates of returns on individual

4.3. Sources of breaks: Summary and Discussion

All five break dates identified by our present value model (1934, 1939, 1969, 1998, and 2010) are closely aligned with major economic crises/recessions or periods of major military conflicts: The break in 1934 corresponds to the end of the Great Depression, 1939 to the breakout of the Second World War, 1969 marks the end of Lyndon Johnson's escalation of the Vietnam War and the beginning of so-called 'Vietnamization' and the withdrawal of US troops along with the beginning of a period with elevated inflation, 1998 marks the collapse of LTCM and the Asian Financial Crisis, while 2010 marks the end of the Global Financial Crisis.

It is intuitively plausible to associate large macroeconomic or regulatory/political shocks with breaks in the return prediction model and Branch and Evans (2010) rigorously establish this link in the context of an equilibrium asset pricing model with incomplete learning. Investors in their model learn about the underlying state variables in the economy using underparameterized forecasting models to predict cash flows which, in turn, gives rise to multiple equilibria ("regimes"). Changes between equilibria may be triggered by large exogenous shocks which get transmitted to stock returns through investors' prediction models and can affect both the price and quantity of risk. Whether a particular exogenous shock triggers a regime shift depends on a set of gain parameters which measure the rate at which investors update their beliefs and thus determine how fast their learning process adapts to the shocks.[18]

The five breaks are also linked to major regulatory changes. For instance, the break in 1934 is closely aligned with the Glass-Steagall Act of 1933 which created the Federal Deposit Insurance Corporation which regulated interest rates on deposits and partitioned commercial and investment banking. The 1939 break is close to

size- and book-to-market-sorted portfolios and the market portfolio lines up with the break dates in the risk premium and beta parameters that we uncover here. We are grateful to a referee for suggesting this robustness check.

[18] Similar effects arise in the model proposed by Hong et al. (2007) due to discrete paradigm shifts in investors' beliefs.

the Investment Company Act and Investment Advisers Act of 1940 which required institutions whose business primarily involved trading stocks to provide full disclosure regarding their investment objectives as well as minimizing conflicts of interest. It also restricted the size of mutual funds, thus affecting flows into the stock market. The break in 1969 is aligned with the Securities Investor Protection Act of 1970 that used broker-dealer fees to create a fund to protect consumers of failed brokerage firms. The 1998 break is aligned with the Gramm-Leach-Bliley Act of 1999 that repealed major parts of the Glass-Steagall Act and the Bank Holding Act of 1956. The 2010 break coincides with the Dodd-Frank Act that tightened financial regulation following the Global Financial Crisis.[19] Some of these regulatory shifts - particularly the major ones such as the Glass-Steagall Act and the Dodd-Frank Act - affected investors' ability and willingness to hold stocks and so are likely to have affected both the price and the quantity of risk.

Turning to the three sources of risk, i.e., the quantity and price of risk as well as the dynamics of the cash flow growth process, Table 3 shows evidence of large shifts in the beta parameters capturing movements in the quantity of risk in the economy along with the parameters of the risk premium process. Further, the largest economic crises, notably the Great Depression and Global Financial Crisis, cause breaks in betas and risk premia whereas military conflicts such as the breakout of the Second World War and smaller economic shocks such as the collapse of LTCM and the Asian Financial crisis may lead to a break in betas only and so have a more limited effect on asset prices. Likewise, the most significant regulatory changes such as the Glass-Steagall Act and Dodd-Frank Act cause a break in both the price and quantity of risk whereas smaller changes (and also deregulatory changes such as the Gramm-Leach-Bliley Act) may affect only betas.

Similarly, the break point estimates in Table 1 show that each of the breaks in the

[19] Note that breaks in stock returns often occur prior to the final approval of regulatory changes which typically are preceded by a lengthy consultation process during which financial markets become aware of any future policy changes.

present value model is closely aligned with a break in the dividend growth process. For example, the break in 1968 identified in the dividend growth equation is associated with a large increase in the predictive power of the dividend-price ratio over dividend growth as well as a decrease in the persistence of the dividend growth process (from 0.51 to 0.21). Predictability of future dividend growth from past dividend growth picks up notably after 2009, as the coefficient on lagged dividend growth increases from 0.307 to 0.834.

Taken together, these results suggest that all three sources of breaks played a role during our sample although the extent to which each component was affected varied across historical break events.

As a final exercise to help us better understand the economic sources of breaks to return predictability, Table 4 shows how the values of a broad range of state variables change across the six regimes identified for the present value model. We track the variation in the mean and volatility of value-weighted stock returns (top two rows) and the dividend growth rate (rows 3 and 4) followed by the coefficient of risk aversion (estimated as explained in Section 4.1.1) (row 5) and five measures of economic, financial and political uncertainty (rows 6-10). We observe marked shifts in the mean and volatility of stock market returns, corroborating the finding from our return predictability analysis that the distribution of stock returns is fundamentally different across the six regimes identified by the present value model.

The table reveals similarly large changes across regimes in the mean and volatility of dividend growth. This is, again, consistent with the finding from Table 1 that each of the breaks in the present value model are associated with a break in the dividend growth equation occurring either in the same or in a neighboring year. For example, at almost 7% per annum, the dividend growth rate from 1939-1969 was more than twice as high as the growth rate from 1926-1934 (3% per annum). Table 4 also reveals that the estimated coefficient of risk aversion was almost twice as high between 1926 and 1934 and, again, after 2010 compared to its value between 1939 and 1969. Finally, the table shows that levels of macroeconomic and financial market uncertainty were generally quite modest from 1939-1998 but increased significantly between 1998 and 2010 before declining again up to the end of the sample.

Ultimately, a more elaborate asset pricing analysis is required to pin down the economic sources of breaks. However, our analysis points in some promising directions. Our empirical results suggest that major regulatory changes and economic crises such as the Great Depression (1934) and Global Financial Crisis (2010) can have a more universal impact and affect both the price and quantity of risk. Large regulatory shifts can change the risk-bearing capacity of the financial markets and so will plausibly affect the equilibrium price of risk. Similarly, large exogenous macro shocks can affect the uncertainty surrounding the underlying cash flow process and, thus, risk quantities in the economy as well as investors' willingness to carry risk.

5. Evaluation of Return Forecasts

We next evaluate the statistical and economic significance of our out-of-sample return forecasts.[20] Our panel break approach differs from conventional return prediction models in two main regards. First, it pools cross-sectional and time-series data, as opposed to the more conventional single-equation time-series approach used throughout most of the literature. Second, it allows for breaks. To quantify the importance of each of these differences, we compare the predictive accuracy of our model to four alternatives that differ along one or more dimensions: (i) a pure time-series approach that allows for breaks, thus highlighting the importance of using cross-sectional (panel) information; (ii) a constant-parameter panel model that uses the same cross-sectional information as our approach, allowing us to gauge the importance of breaks; (iii) a time-varying parameter model that allows for small changes to the parameters every

[20] To avoid spurious inference, we perform our forecasting analysis using only stocks with at least 60 out-of-sample observations of which there are 11,210.

period, thus letting us evaluate whether parameter instability is best modeled as a sequence of slow changes or as the outcome of few, but potentially larger, discrete breaks; and (iv) the simple historical average which serves 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.

5.1. Out-of-sample Return Forecasts

Before inspecting the empirical results, we first explain how our return forecasts are computed. At each point in time, t, we generate out-of-sample forecasts for the $i = 1, \ldots, N$ stock return series by loading the slope estimate on the predictor variable from the final regime at that time $(K_t + 1)$ and adding the intercept estimate:

$$\hat{r}_{i,t+1} \mid K_t = \hat{\mu}_{iK_t+1} + \hat{\beta}_{iK_t+1} X_t. \tag{21}$$

This step incorporates uncertainty surrounding the break locations but conditions on the number of breaks K_t. To handle uncertainty about the number of breaks, let $K_{t,min}$ and $K_{t,max}$, denote the lowest and highest number of breaks that are assigned a nonzero posterior probability by our estimation procedure at time t. We then apply Bayesian Model Averaging to integrate out uncertainty about K_t:

$$\hat{r}_{i,t+1} = \sum_{K_t = K_{t,min}}^{K_{t,max}} p(K_t \mid \mathbf{r}, X) \hat{r}_{i,t+1} \mid K_t, \tag{22}$$

where $\mathbf{r}$ denotes the returns on the N stocks across the t time periods and X denotes the t observations on the predictor.[21]

In turn, using a bottom-up approach, we forecast the aggregate market return as

[21] Avramov (2002) reports that Bayesian Model Averaging improves forecasting performance for stock returns in the presence of model uncertainty.

the value-weighted average of the underlying N forecasts from equation (22)

$$\hat{r}_{Mkt,t+1} = \sum_{i=1}^{N} w_{it}\hat{r}_{it+1}, \tag{23}$$

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

To implement these methods, 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 the end of the sample (December 2015.)

5.2. Statistical Measures of Forecasting Performance

To visually assess the accuracy of the out-of-sample return forecasts, we inspect plots of 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:

$$CSSED_t = \sum_{\tau=1}^{t} (e^2_{Bmk,\tau} - e^2_{Pbrk,\tau}), \tag{24}$$

in which $e_{Bmk,\tau}$ and $e_{Pbrk,\tau}$ denote the respective forecast errors from the benchmark and our panel break model at time τ. 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. 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 plots the CSSED values for the market portfolio over the full out-of-sample period (1936-2015, top panel) as well as over the subsample 1990-2000 (lower panel).[22] To keep the plots relatively easy to interpret, we only show the performance of our panel break forecasts against forecasts from the prevailing mean model, but similar results are obtained against the other benchmarks. The top window shows that our panel break model outperforms the prevailing mean over the 80-year evaluation sample. This strong performance against the historical average - which holds regardless of whether our model is trained on the full sample or only on post-WWII data - is particularly impressive given that this benchmark has been found by Goyal and Welch (2008) to be very difficult to beat out-of-sample.

The CSSED curve for the panel model with breaks measured relative to the prevailing mean rises through much of the out-of-sample period, although with a period of sustained underperformance from December 1990 until April 2000. The reason for our model's poor performance during this period is that it predicts relatively low returns compared with actual (realized) returns.[23] The low average forecasts from our panel break model during this period are primarily driven by the dividend-price ratio being approximately half (0.02) of its full sample average (0.04). In fact, our panel break model partially accommodates the marked shift in the relation between the dividend-price ratio and stock returns during the nineties by reducing the slope coefficient on this predictor relative to the constant-coefficient model. This means that our panel break model produces return forecasts that, during the nineties, are less accurate than those from the prevailing mean model but more accurate than those from a constant-coefficient model that uses the dividend-price ratio as a predictor.[24]

[22] The plots are based on the panel break model with a 10-year prior on the break frequency, but very similar results are obtained with a 20-year prior for this model or using the timing of the breaks recursively identified by the present value model.

[23] For this period, the mean forecast from our model is 0.002 compared with 0.007 for the prevailing mean and a mean realized excess return of 0.006.

[24] During the nineties, the average value of the recursively estimated slope coefficient from our panel break model (0.47) is below the estimate obtained from a constant-coefficient model (0.57).

From March 2000 until October 2002, the CSSED value rises from 0.007 to 0.011. This is driven by our forecasts being on average lower (0.001) than those from the prevailing mean (0.007) while realized excess returns were low (-0.014). Our low returns are driven by both the dividend-price ratio (0.01) and our recursive slope estimate (0.20) being on average low. From October 2007 through March 2009, the CSSED value rises from 0.009 to 0.012. Outperformance results from our average forecast being low (0.002) due to a low dividend-price ratio (0.02) and slope estimate (0.19).

We observe a particularly strong rise in the CSSED curve during the final 15 years (bottom panel). This strong performance occurs, first, because our average forecast for this period is 0.0035, close to the average realized excess return of 0.0032, while the average forecast from the prevailing mean is almost twice as large (0.0066). Second, our forecasts have a positive correlation of 0.098 with the realized excess returns while the prevailing mean forecasts have a negative correlation of -0.177 during this period.

The Great Depression experienced high volatility in both dividend growth and stock returns and so may have a disproportionate influence on our empirical results. To explore whether the forecasting performance of our method is equally strong during the post-WWII sample. We start our sample in January 1946 and once again use a ten-year warm-up sample such that the first forecast is generated for January 1956 from a model that is estimated using data from January 1946 through December 1955. The results are equally strong for this period with our model generating significantly more accurate forecasts for the majority of stocks and for the market portfolio relative to the four benchmarks.

These plots visualize the evolution in the predictive accuracy of the return forecasts from our model relative to a set of common alternatives. We next summarize the average forecasting performance of the panel break model relative to the benchmarks using the out-of-sample R^2 measure of Campbell and Thompson (2008):

$$R^2_{i,OoS} = 1 - MSPE_{i,Pbrk}/MSPE_{i,Bmk}. \quad (25)$$

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

Figure 8 plots histograms of the R^2_{OoS} values for individual stocks and the market portfolio (thick vertical black line) based on comparisons of the forecasting performance of our proposed panel break model relative to the four benchmarks.[25] For the vast majority of stocks, our method outperforms all four benchmarks and many of the R^2_{OoS} values are economically large.[26]

To more formally evaluate the statistical significance of the relative performance of the panel break model against the four benchmarks, Table 5 uses the test statistic of Clark and West (2007) that accounts for nested models. First consider the results based on the lagged dividend-price ratio prediction (top row). The table shows that the panel break model performs significantly better than all four benchmarks at the 10% critical level for between 7,908 and 9,332 of the 11,211 stocks (right column) including the market portfolio, denoted by the † symbol. Conversely, the return forecasts from our panel break model only significantly underperform the benchmarks for 2-5% of the stocks.

Turning to the other predictor variables, the results in Table 5 show that our panel break approach continues to significantly outperform the four benchmark models for the majority of stocks for thirteen of the fourteen predictor variables. Our model

[25] Here and for most of the ensuing return predictability results, we use the panel break model with a 10-year prior on the break frequency. However, very similar results are used for this model with a 20-year prior or if we use the timing of the breaks identified by the present value model.

[26] Campbell and Thompson (2008) estimate that an R^2_{OoS} value of one-half of one percent on monthly data is economically large for a mean-variance investor with moderate risk aversion.

only experiences relatively poor performance for the corporate equity issuance (ntis) predictor (bottom row). Moreover, the results are robust to using a different (post-war) sample period that excludes the Great Depression and starts in 1946.

These findings underline that the improvements in predictive accuracy observed for the panel break model are not simply a result of expanding the information set from a univariate time series to a panel setup that incorporates cross-sectional information. 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.[27]

The results also demonstrate that our panel break model has the ability to adapt to breaks while simultaneously reducing the effect of estimation error which has so far plagued real-time (out-of-sample) return forecasts (Lettau and Van Nieuwerburgh 2008). Key to the improved forecasting performance is our ability to detect breaks in return predictability with little delay, combined with our use of economically-motivated priors. This reduces the effect of estimation error which adversely affects the accuracy of return forecasts inside new regimes.

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. Our panel-break method outperforms the four competing benchmarks by the largest margin in the three-year period after a break is detected before flattening out, demonstrating the value from using our panel procedure to detect the onset of a break more quickly in real time.

[27]Reassuringly, the proportion of stocks for which our panel break model generates significantly more accurate return forecasts remains high if we use a more conservative hyper parameter which implies that breaks are expected to occur every 20 years instead of every 10 years. For this prior, using the dividend-price ratio as a predictor, we find significant improvements relative to the no break panel for 80.14% of stocks, 76.50% for the prevailing mean model, 85.02% for the time-series break model, and 73.39% for the TVP model.

5.3. Relation to Present Value Model

The analysis in Section 4.1 uses industry portfolios to impose restrictions implied by the present value model whereas the out-of-sample forecasting results use individual stock returns.[28] Hence, it is important to link these two analyses as we next do.

First, we note that our empirical findings show that (a) the five breaks identified by the present value model which imposes a set of asset pricing constraints are closely in tune with breaks identified by our panel break model applied to the return prediction equation familiar from the literature on return predictability provided that we apply a conservative prior on the break frequency (20 years) as opposed to an expected break frequency of 10 years; (b) our choice of return data (industry portfolios versus individual stocks) makes little difference to the empirical break point findings.

To explain the mechanics of how imposing the present value model restrictions leads to a reduction in the number of breaks identified by our approach, note that the period 1998 - 2010 is deemed to be a single regime by the present value model. Conversely, the dividend growth model splits this period into two shorter regimes, namely 1998-2007 and a very short regime, 2007-2009 (see Table 1). Similarly, the model for the dividend-price ratio splits this period into one regime between 1997-2002 and another between 2002-2008. Because the middle break locations in the dividend growth and dividend-price ratio models (2007 and 2002, respectively) are not aligned in time, they get dropped by the present value model which requires that the breaks should affect the return, dividend growth and dividend-price ratio processes at the same time. Imposing this constraint, thus, leads to a model with a single long regime 1998-2010 as opposed to two or three shorter ones for this period.

Applying a more conservative prior on the expected break frequency is equivalent to being more skeptical about breaks occurring. This implies setting the hurdle

[28] Note that it is not an option for us to apply the present value approach to individual stocks given the lumpy and irregular nature of the dividend process at the individual stock level; see, e.g., Pettenuzzo et al. (2019).

higher for accepting new breaks and so has a similar effect of detecting only the largest (common) breaks.

Second, to establish the link even more directly, we use the posterior break probabilities from the present value model in the forecasting exercise for the individual stocks. In particular, we evaluate the out-of-sample forecasting results for the individual stock return data when imposing on our baseline return prediction model the break dates that are estimated recursively from the present value model using only the information available at the time the forecast is made. Note that the model parameters continue to be fitted to the individual stock return series - it is only the posterior break probabilities that we utilize from the present value model fitted to the 30 industry portfolios. This exercise, thus, directly bridges the present value analysis from Section 4.1 conducted on the 30 industry portfolios with the out-of-sample return forecasts conducted for the individual stocks. We continue to obtain strong out-of-sample return prediction results when applying this approach (see Tables 6 and 8).

Our first point helps to explain this finding since it demonstrates the similarity between the break dates identified using the present value model fitted to the 30 industry portfolios and the break dates identified by a model fitted to individual stocks but using a relatively conservative prior on the break frequency.[29]

The effect of using the break locations identified by the present value model is particularly strong for the return prediction model based on the ntis variable. The baseline return prediction model that uses the ntis variable produces inaccurate forecasts relative to all the benchmarks, in large part because it identifies 12 breaks, many of which are close together, leading to large estimation error. In turn, this causes the slope coefficient to frequently switch sign and undergo sharp reversals. Modifying the model to instead use the five breaks identified by the present value model, we find a

[29]To verify this point we repeated the out-of-sample analysis using the more conservative prior on break frequencies without imposing the present value model break dates on the benchmark model. Reassuringly, we find very similar evidence of out-of-sample return predictability.

notable improvement in forecast accuracy as can be seen in Tables 6 and 8.

5.4. Comparison with alternative approaches

To place our results in context, we next compare our findings to those obtained using some of the leading forecasting approaches in the return predictability literature. Existing studies differ in terms of sample periods, forecasting approaches and predictor variables. To facilitate a direct comparison based on the same data, we compare our results to those obtained from four approaches, namely (i) imposing economic constraints on the return prediction model (Campbell and Thompson 2008; Pettenuzzo et al. 2014); (ii) a rolling-window estimator; an (iii) the equal-weighted average forecast combination approach (Rapach et al. 2010), and (iv) a regime switching model similar to that used by Henkel et al. (2011).

Using forecasts of returns on the market portfolio constructed from firm-level return forecasts that include the aggregate dividend-price ratio as a predictor, our panel break model generates out-of-sample R^2 values that are 0.34%, 0.30%, 0.16%, and 0.52% higher than the values obtained under the sign-constrained approach of Campbell and Thompson (2008), a five-year rolling window model estimated by OLS, the forecast combination approach of Rapach et al. (2010) that uses the dividend-price ratio, T-bill rate, default and term premia, and a regime switching model similar to that of Henkel et al. (2011). These improvements in forecast accuracy are significant at the 1%, 5%, 10%, and 1% levels, respectively.

5.5. Economic Utility from Return Forecasts

In addition to comparing the statistical performance of forecasts from our panel break model to a set of benchmarks, it is important to evaluate their economic performance. We construct decile portfolios of the individual stocks using the mean squared difference between the forecasts from the models with and without breaks as our sensitivity measure. Portfolios are rebalanced each period using this measure.[30] For each of the portfolios we next compute the utility gain to a mean-variance investor who each period allocates his portfolio between a single risky asset and risk-free T-bills.[31] At time t the investor allocates a portion of his portfolio to equities in period $t+1$, $w_t = (1/A) \times \hat{r}_{t+1}/\hat{\sigma}^2_{t+1}$, 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.[32]

If breaks in the model parameters do not strongly affect a particular stock, it is unlikely that a model that accounts for such breaks can significantly outperform a model that ignores breaks for this stock or portfolio. To see if this holds, Table 7 explores the 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 (third column), and the utility gains for that portfolio, assuming a mean-variance investor with a coefficient of risk aversion of three (fifth column).

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

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 individual stocks. Using our panel break model, at each time t we determine the weight vector ω_t to allocate among stocks in the next period, i.e., we solve for the ω_t that maximizes the expected utility for the return on the risky portfolio at time $t+1$,

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

[31]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).

[32]Following Campbell and Thompson (2008), we use a five year rolling window of monthly 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 high leverage.

$r_{p,t+1}$,

$$E[U(r_{p,t+1} \mid A)] = r_{f,t} + \omega_t' \hat{r}_{t+1} - \frac{A}{2} \omega_t' \hat{S}_t \omega_t, \tag{26}$$

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 stocks.[33] 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.[34]

The top panel of Table 8 reports out-of-sample utility gains from these optimized allocations across stocks. Specifically, the table shows the annualized CER values for the panel break model measured relative to the four benchmarks, with different rows tracking the results for the individual predictor variables. For example, for the lagged dividend-price ratio (top row), the CER value of the panel break model is 2.79% per annum relative to the historical average forecasts. Average gains in the CER value of the panel break model remain large (2-3% per annum) when measured against the univariate time-series break, time-varying parameter and no-break panel models.[35]

Overall, for 13 of the 14 predictor variables, we obtain very similar estimates of CER gains around 2% per annum for our panel break return forecasts measured relative to the four benchmarks. Only for the corporate equity issuance (ntis) variable do we find negative estimates.

The improved predictive performance in the immediate aftermath of a break being detected translates into even larger utility gains during these periods. Computing utility gains using only those time periods that occur within two years of a break first being detected ('After breaks'), for the lagged dividend-price ratio predictor, the annualized CER values are now even larger (2.35-3.45%). This reflects the gains in

[33] 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).

[34] Detailed results for portfolio allocations are available on request.

[35] Using the more conservative priors that imply a break every 20 years, the annualized CER values for the panel model that uses the dividend-price ratio as a predictor are 2.87% (relative to the no-break panel), 2.58% (prevailing mean), 2.39% (time series break), and 1.97% (TVP model).

predictive accuracy resulting from our approach being able to detect breaks quickly in real time.

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

Our findings are competitive with or improve on many popular approaches. Campbell and Thompson (2008, Table 4) report average utility values ranging from 0.4% to 1.4% per annum as a result of imposing a positive intercept and a bounded slope in regressions of stock returns on various valuation ratio predictors. Dangl and Halling (2012) report certainty equivalent values ranging from -1.74% to 0.88% per annum (for post-1947 data), although forecast combination in the form of Bayesian Model Averaging (BMA) raises this to 2.57% per annum. Rapach et al. (2010) report CER values around 2% from simple forecast combinations.

5.6. Multivariate prediction models

So far our analysis has focused on breaks to return prediction models with a single predictor variable. This focus is in line with the return predictability literature, see, e.g., Goyal and Welch (2008), Lettau and Van Nieuwerburgh (2008), and Pettenuzzo et al. (2014). However, it is also of interest to explore evidence from multivariate regressions that include more than one predictor. To this end, we study a bivariate model that includes the dividend-price ratio and the default spread as predictors along with a four-variable model that includes the most commonly used and popular return predictors (dividend-price ratio, short T-bill rate, default and term premia).

The results (reported in Table 9) show that the bivariate panel break model continues to perform better than the no-break panel and time-series break and time-varying parameter models, all extended to include two predictors. However, the performance of the bivariate models deteriorates notably relative to the more parsimonious prevailing mean model. We find similar results for the four-variable panel break model, although the effects are more pronounced for this larger model.

To help explain these findings, the top row in Figure 9 displays the break dates (posterior modes) for the four-variable forecasting model. Using a 10-year prior on the break frequency, this multivariate model detects 12 breaks. Each of the univariate break models (shown in rows 2-5) uncover seven or eight breaks. However, because the break dates for the univariate return prediction models are not perfectly aligned in time, the multivariate model detects more breaks than the individual univariate models. This increase in the number of breaks detected by the multivariate model results in larger estimation errors in the slope coefficients on the predictors. For example, the figure shows that the dividend-price ratio in the univariate model is stable from 1970 through 1997. However, because the other predictors experience breaks during this period, the multivariate model detects breaks at 1973, 1979 and 1992. In the multivariate model, the slope on the dividend-price ratio is thus estimated separately from 1970-73, 1974-79, 1980-92, and 1993-97 instead of estimating it from 1970-97. As a result, the predictive information in the dividend-price ratio is reduced in what is already a low signal-to-noise environment.

One approach to overcome this issue is to first estimate univariate panel-break models separately for each of the predictor variables, i.e., models that include one predictor variable at a time. This step effectively uses all data to estimate the slope coefficients for the individual predictors during the time spans where these relations are deemed to be stable. In a second step, we form a simple equal-weighted combination of the forecasts from the 14 univariate panel break return prediction models. Table 9 shows that this procedure works very well. For example, relative to the prevailing mean this approach generates significantly more accurate forecasts for 81% of stocks, including the market portfolio for which we obtain an out-of-sample R^2 value of 0.82%. Using the resulting return forecasts in an asset allocation exercise similar to that described above, we obtain an annualized utility gain of 3.47%.

To disentangle the effect of including information on multiple variables versus accounting for breaks, note that our univariate panel break model based on the dividend-price ratio predictor generates an out-of-sample R^2 value of 0.16% relative to the forecast combination approach of Rapach et al. (2010) which uses the most common four predictors. However, relative to an approach that uses all 14 predictors, our univariate panel break approach underperforms with an out-of-sample R^2 value of -0.20%. Comparing the performance of an equal-weighted forecast combination across our 14 univariate panel break prediction models relative to the equal-weighted combination that uses all 14 predictors, we generate a positive, but more modest out-of-sample R^2 value of 0.08%. Note that the underlying forecasting models used in this comparison are very different: Our approach fits 14 univariate Bayesian panel break models to the individual stock returns while the approach of Rapach et al. (2010) uses 14 univariate models fitted to aggregate stock returns. A more direct comparison that uses similar models is provided by examining the equal-weighted combination of the 14 univariate forecasts generated by our panel break model versus univariate forecasts from a no-break panel model fitted to individual stock returns. This comparison is provided in the second row of the bottom panel of Table 9. We find that the equal-weighted combination of 14 univariate return forecasts from our panel-break model is significantly more accurate than the equal-weighted combination of 14 univariate return forecasts from the no-break model for 95% of the individual stocks. Taken together, our results suggest that while accounting for breaks in a panel context can lead to improved forecast accuracy, incorporating multivariate information can also lead to important gains in forecasting performance.

Finally, it is worth highlighting that the asset pricing restrictions from the present value model help reduce the increased parameter estimation error resulting from adding more predictor variables. To see this, recall from Table 1 that the bivariate models fitted to the dividend growth rate and dividend-price ratio data identify nine and six breaks, respectively. However, once we impose the cross-equation restrictions from the present value model, the number of breaks identified by the bivariate model gets reduced to five. Hence, imposing the present value restrictions effectively help us deal with the potential for break proliferation due to the inclusion of additional predictor variables.

To further explore this point, we estimate the present value model using univariate models that include only one predictor variable at a time, namely either the dividend growth rate or the dividend-price ratio. In both cases, we find that the break dates are closely aligned with those of the bivariate present value model. Hence, the bivariate model does not detect additional breaks relative to the univariate models and so estimation error does not negatively impact the bivariate model in the same way we experience in the unconstrained multivariate return prediction models. The key difference is that we impose cross-equation restrictions in the present value model while we do not do this in the return prediction models. The asset pricing restrictions help us better identify the break dates and thus avoid large estimation error.

6. Conclusion

A large literature on predictability of stock 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 (Goyal and Welch 2008). While model instability is, thus, known to affect return forecasts, exploiting it has so far largely proved elusive due to the noisy nature of returns and the low power of return prediction models which makes detecting and quantifying the effects of changing 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 common across stocks, the effects of breaks are allowed to differ across individual stocks. The break dates identified by our approach, such as the Great Depression, World War II, the oil price shocks of the seventies and, most recently, the Global Financial Crisis, suggest that return predictability can undergo large changes in the presence of such major events.

Empirically, we find that pooling cross-sectional and time-series information substantially increases our ability to detect breaks in return prediction models with very little delay. Combined with economically-motivated priors, this means that out-of-sample return forecasts from our panel break model are consistently more accurate than forecasts from a variety of alternative approaches, with gains in predictive accuracy being particularly large shortly after a break has occurred.

We explore the economic sources of breaks using several approaches. First, we use the asset pricing restrictions implied by a simple reduced-form present value model to trace back breaks in return predictability to breaks in the dividend growth and/or discount rate processes. We develop an approach for rigorously testing if the timing and magnitude of breaks in the return prediction model can be linked to concurrent breaks in dividend growth and dividend-price ratio equations and find that this – along with the asset pricing restrictions implied by the present value model – is supported empirically. We also develop ways for determining whether breaks in return predictability can be attributed to breaks in factor loadings (betas), breaks in risk premia, or breaks in both. Our empirical results suggest that although some breaks reflect shifts in either betas only or both betas and risk premia, both sources of breaks played a role in our sample.

References

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.

Bai, J. (2010). Common breaks in means and variances for panel data. *Journal of Econometrics*, 157(1):78–92.

Bai, J., Lumsdaine, R. L., and Stock, J. H. (1998). Testing for and dating common breaks in multivariate time series. *Review of Economic Studies*, 65(3):395–432.

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

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

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.

Branch, W. A. and Evans, G. W. (2010). Asset return dynamics and learning. *Review of Financial Studies*, 23(4):1651–1680.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Campbell, J. Y. and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228.

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.

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. (2008). The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575.

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, 106(3):664–98.

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

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.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Farmer, L., Schmidt, L., and Timmermann, A. (2019). Pockets of predictability. *Unpublished Manuscript, MIT, University of Virginia and UCSD.*

Fearnhead, P. (2006). Exact and efficient bayesian inference for multiple changepoint problems. *Statistics and Computing*, 16(2):203–213.

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.

Gu, S., Kelly, B., and Xiu, D. (2018). Empirical asset pricing via machine learning. Technical report, National Bureau of Economic Research.

Henkel, S. J., Martin, J. S., and Nardari, F. (2011). Time-varying short-horizon predictability. *Journal of Financial Economics*, 99(3):560–580.

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

Hong, H., Stein, J. C., and Yu, J. (2007). Simple forecasts and paradigm shifts. *Journal of Finance*, 62(3):1207–1242.

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.

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

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.

Koop, G. and Potter, S. M. (2009). Prior elicitation in multiple changepoint models. *International Economic Review*, 50(3):751–772.

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

Lustig, H., Van Nieuwerburgh, S., and Verdelhan, A. (2013). The wealth-consumption ratio. *Review of Asset Pricing Studies*, 3(1):38–94.

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

Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8(4):323–361.

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.*

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

Pettenuzzo, D., Sabbatucci, R., and Timmermann, A. (2019). High-frequency cash flow dynamics. *forthcoming Journal of Finance.*

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

Pettenuzzo, D., Timmermann, A., and Valkanov, R. (2014). Forecasting stock returns under economic constraints. *Journal of Financial Economics*, 114(3):517–553.

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. (2018). Noncommon breaks. *Unpublished Manuscript, University of Southern California.*

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. Likelihood function

This appendix specifies the likelihood function of our return prediction model in (6) whose intercept μ_i, slope coefficient β_i, and error-term variance σ_i^2 are subject to breaks.

First, we introduce some notations. Let $l_k = \tau_k - \tau_{k-1}$ denote the duration of the kth regime which contains the observations $\tau_{k-1}+1, \ldots, \tau_k$, and let $l = (l_1, \ldots, l_{K+1})$. Next, let $\mu_i = (\mu_{i1}, \ldots, \mu_{iK+1})$, $\beta_i = (\beta_{i1}, \ldots, \beta_{iK+1})$, $\sigma_i^2 = (\sigma_{i1}^2, \ldots, \sigma_{iK+1}^2)$, $\theta_i = (\mu_i, \beta_i)$, and $\Theta_i = (\mu_i, \beta_i, \sigma_i^2)$. Further, let r_i denote the $(T \times 1)$ vector of excess returns on the ith asset, while $\boldsymbol{X}$ denotes the $(\kappa \times T)$ matrix which combines a unit vector with the observations on the predictor.[36]

Assuming that the errors follow a Gaussian distribution with zero mean and regime-specific variance σ_{ik}^2, the likelihood function for the model with pooled breaks and unit-specific parameters is

$$\begin{aligned} p(\boldsymbol{r} \mid \boldsymbol{X}, \Theta_i, \tau) &= \prod_{i=1}^{N} \prod_{k=1}^{K+1} \prod_{t=\tau_{k-1}+1}^{\tau_k} p(r_{it} \mid X_{t-1}, \Theta_{it}), \\ &= \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi\sigma_{ik}^2)^{-l_k/2} \right) exp \left[-\frac{1}{2} \sum_{i=1}^{N} \sum_{k=1}^{K+1} \sum_{t=\tau_{k-1}+1}^{\tau_k} \frac{(r_{it} - X_{t-1}'\theta_{ik})^2}{\sigma_{ik}^2} \right] \end{aligned} \tag{A.1}$$

in which $\mathbf{r} = (r_1, \ldots, r_N)$. For future use, define $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)$, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_N)$, $\boldsymbol{\sigma^2} = (\sigma_1^2, \ldots, \sigma_N^2)$, $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$, and $\boldsymbol{\Theta} = (\Theta_1, \ldots, \Theta_N)$.

[36] Making the first of the $\kappa = 2$ covariates a unit vector results in the first element of $\theta_i = (\mu_i, \beta_i)$ being an estimate of the intercept μ_i for the ith return.

Appendix B. Priors and posteriors

This appendix provides details of the prior and posterior distributions used by our return prediction models. Following Koop and Potter (2007), we assume that regime durations have a Poisson prior distribution

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

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

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

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

Appendix B.1. Priors on parameters β and σ

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.

We specify an inverse gamma prior over the regime-specific variances

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

A Gaussian prior with zero mean is placed over the regression coefficients conditional on the variances[37]

$$p(\theta_{ik} \mid \sigma^2_{ik}) = 2\pi^{-\kappa/2}(\sigma^2_{ik})^{-\kappa/2} \mid V_\theta \mid^{-1/2} exp\left(-\frac{1}{2\sigma^2_{ik}}\theta'_{ik}V_\theta^{-1}\theta_{ik}\right), k = 1, \ldots, K+1,$$
$$V_\theta = \begin{pmatrix} \sigma^2_\mu & 0 \\ 0 & \sigma^2_\eta/\sigma^2_X \end{pmatrix}. \tag{B.4}$$

Appendix B.2. Prior Elicitation

The prior used in our empirical application assumes that a break occurs approximately every decade. 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. 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, σ_X, have been explained in Section 2.3.

Appendix B.3. Derivation of Posterior Distribution

We next show how we can simplify the posterior distribution by marginalizing the parameters. This significantly reduces the computational burden of estimating our panel break model. First, multiply the likelihood function by the prior distributions

[37] For consistency, we use the equivalent prior specification for the time series break model.

to obtain the posterior

$$
\begin{aligned}
p(\boldsymbol{\Theta} \mid \boldsymbol{r}, \boldsymbol{X}, \tau) &= \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi\sigma_{ik}^2)^{-l_k/2} \right) exp \left[-\frac{1}{2} \left(\sum_{i=1}^{N} \sum_{k=1}^{K+1} \sum_{t=\tau_{k-1}+1}^{\tau_k} \frac{(r_{it} - X_t'\theta_{ik})^2}{\sigma_{ik}^2} \right) \right] \\
&\times \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} \frac{b^a}{\Gamma(a)} (\sigma_{ik}^2)^{-(a+1)} \right) exp \left(\sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{-b}{\sigma_{ik}^2} \right) \\
&\times \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi\sigma_{ik}^2)^{-\kappa/2} \mid V_\theta \mid^{-1/2} \right) exp \left[-\frac{1}{2} \left(\sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{\theta_{ik}' V_\theta^{-1} \theta_{ik}}{\sigma_{ik}^2} \right) \right] \\
&= \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} (\sigma_{ik}^2)^{-((l_k+\kappa)/2+a+1)} (2\pi)^{-(\kappa+l_k)/2} \mid V_\theta \mid^{-1/2} \frac{b^a}{\Gamma(a)} \right) \\
&\times exp \left[-\sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{2b + r_{ik}' r_{ik}}{2\sigma_{ik}^2} \right] \times exp \left[-\frac{1}{2} \sum_{i=1}^{N} \sum_{k=1}^{K+1} \left(\frac{\rho_{ik}' \Sigma_k^{-1} \rho_{ik}}{\sigma_{ik}^2} \right) \right] \\
&\times exp \left[-\frac{1}{2} \sum_{i=1}^{N} \sum_{k=1}^{K+1} \left(\frac{(\theta_{ik} - \rho_{ik})' \Sigma_k^{-1} (\theta_{ik} - \rho_{ik})}{\sigma_{ik}^2} \right) \right] \qquad \text{(B.5)}
\end{aligned}
$$

where $r_{ik} = (r_{i\tau_{k-1}+1}, \ldots, r_{i\tau_k})$. Multiplying and dividing by $2\pi^{\kappa/2}|\Sigma_k|^{1/2}(\sigma_{ik}^2)^{\kappa/2}$ and integrating out $\boldsymbol{\theta}$, we obtain

$$
\begin{aligned}
p(\boldsymbol{\sigma^2} \mid \boldsymbol{r}, \boldsymbol{X}, \tau) &= \int p(\boldsymbol{\Theta} \mid \boldsymbol{r}, \boldsymbol{X}, \tau) d\boldsymbol{\theta} \\
&= \left(\prod_{i=1}^{N} \prod_{k=1}^{K+1} (\sigma_{ik}^2)^{-((l_k+\cancel{\kappa})/2+a+1)} (2\pi)^{-(\cancel{\kappa}+l_k)/2} \mid V_\theta \mid^{-1/2} \frac{b^a}{\Gamma(a)} \cancel{(2\pi)^{\kappa/2}} |\Sigma_k|^{1/2} \cancel{(\sigma_{ik}^2)^{\kappa/2}} \right) \\
&\times exp \left[-\sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{2b + r_{ik}' r_{ik}}{2\sigma_{ik}^2} \right] \times exp \left[-\frac{1}{2} \sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{(\rho_{ik}' \Sigma_k^{-1} \rho_{ik})}{\sigma_{ik}^2} \right] \\
&\cancel{\times \int \prod_{k=1}^{K+1} \prod_{i=1}^{N} (2\pi\sigma_{ik}^2)^{-\kappa/2} |\Sigma_k|^{-1/2} exp \left[\sum_{k=1}^{K+1} \sum_{i=1}^{N} -\frac{((\theta_{ik} - \rho_{ik})' \Sigma_k^{-1} (\theta_{ik} - \rho_{ik}))}{2\sigma_{ik}^2} \right] d\theta_{ik}.} \\
& \qquad \text{(B.6)}
\end{aligned}
$$

The integral of the final term is equal to 1 since θ_{ik} has a multivariate normal (full conditional) distribution with mean ρ_{ik} and covariance matrix $\Sigma_k \sigma_{ik}^2$ for $k = 1, \ldots, K+1$ and $i = 1, \ldots, N$.

Finally, multiplying and dividing by $\Gamma(\tilde{a}_k)/\tilde{b}_{ik}^{\tilde{a}_k}$ enables us to marginalise $\boldsymbol{\sigma^2}$ from the posterior:

$$
\begin{aligned}
p(\boldsymbol{r} \mid \boldsymbol{X}, \tau) &= \int p(\boldsymbol{\sigma^2} \mid \boldsymbol{r}, \boldsymbol{X}, \tau) d\boldsymbol{\sigma^2} \\
&= \prod_{i=1}^{N} \prod_{k=1}^{K+1} (2\pi)^{-l_k/2} \mid V_\theta \mid^{-1/2} \frac{b^a}{\Gamma(a)} \frac{\Gamma(\tilde{a}_k)}{\tilde{b}_{ik}^{\tilde{a}_k}} |\Sigma_k|^{1/2} \\
&\times \prod_{i=1}^{N} \prod_{k=1}^{K+1} \int \frac{\tilde{b}_{ik}^{\tilde{a}_k}}{\Gamma(\tilde{a}_k)} (\sigma_{ik}^2)^{-(\tilde{a}_k+1)} exp \left[-\sum_{i=1}^{N} \sum_{k=1}^{K+1} \frac{\tilde{b}_{ik}}{\sigma_{ik}^2} \right] d\sigma_{ik}^2,
\end{aligned}
\tag{B.7}
$$

where the integral of the final term is equal to 1 since σ_{ik}^2 has an inverse gamma distribution with hyperparameters $\tilde{a}_k$ and $\tilde{b}_{ik}$ for $k = 1, \ldots, K+1$ and $i = 1, \ldots, N$. Marginalising θ and σ^2 from the posterior therefore leaves us with the following result:

Proposition 1. Assuming inverse gamma priors on the error-term variances, σ^2, and Gaussian priors on the regression coefficients, μ and β, conditional on σ^2, the posterior distribution of the heterogeneous panel model with breaks takes the form

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

in which[38]

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

and $X_k = (X_{\tau_{k-1}+1}, \ldots, X_{\tau_k})$.

[38]The bold face used to denote matrices is omitted for Σ_k throughout for expositional ease. Here V_θ denotes the prior covariance matrix, which is constant across regimes, of the intercept μ and the slope coefficient β.

Appendix C. Accounting for Cross-sectional Dependencies

The most direct way of accounting for cross-sectional dependencies is to estimate the variance-covariance matrix in each regime. Since we allow for regimes as short as one period, such an approach may require estimating more parameters than observations available in a regime. This will occur if the length of the kth regime $l_k < (3N + (N^2 - N)/2)/N$. For the industry portfolio application, $N = 30$ and so this would arise if any regime duration is shorter than 18 months.[39] Restricting the minimum regime duration is not appealing since one of the key benefits of our approach is the ability to detect breaks quickly in real time: we always detect breaks with delays shorter than 18 months.

An alternative method for allowing for cross-sectional dependencies is to introduce a common factor that absorbs the correlations (Bai and Ng 2002; Pesaran 2006). We therefore introduce the excess return on the U.S. aggregate market portfolio as an observed common factor on the right-hand-side of our panel regression models with portfolio-specific factor loadings[40]

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

where f_t denotes the market factor in our baseline results and the three- or five-factors of Fama-French in the robustness checks. The factor loadings for the ith asset in regime k are denoted γ_{ik} and ν_{it} denotes the idiosyncratic residuals.

[39] This problem is exacerbated for individual stocks where the cross-section may run into the thousands and the panel may be unbalanced.

[40] The baseline results are robust to extending the setup to include the three or five factors of Fama and French (2015) as observed common factors.

Appendix D. 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 break-point 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 details of this algorithm here.

In contrast, our panel break model analyses 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).

We next formally explain estimation of the breakpoints and the parameters under this new approach. Before explaining how the breakpoints are determined, we first characterize the full conditional distributions from which the parameters of each of the models we consider are sampled.

Appendix D.1. Estimation of Parameters

Using the earlier notations, for a given collection of breakpoints, the full conditional distributions of the parameters take the form (for $k = 1, \ldots, K+1, i = 1, \ldots, N$)

$$\sigma_{ik}^2 \mid \cdot \sim IG(\tilde{a}_k, \tilde{b}_{ik}), \tag{D.1}$$

$$\theta_{ik} \mid \cdot \sim MVN(\rho_{ik}, \Sigma_k \sigma_{ik}^2), \tag{D.2}$$

where the corresponding values of $\rho, \Sigma, \tilde{a}$, and $\tilde{b}$ are computed from (B.9). Given the estimated number of breakpoints and their locations, it is straightforward to estimate the parameters from their full conditional distributions using the Gibbs sampler.

Appendix D.2. Estimation of Breakpoint Locations

For each of the $k = 1, \ldots, K$ breakpoints, we perturb τ_k by an integer u sampled uniformly from the interval $[-s, s]$ such that $\tau_{k^*} = \tau_k + u$. If $u = 0$ the proposal is rejected.[41] Let the proposed breakpoint vector be denoted τ^*, which is simply τ with τ_{k^*} substituted for τ_k.

Next, we compute the updated regime durations $l_{k^*} = \tau_{k^*} - \tau_{k-1}$, $l_{k^*+1} = \tau_{k+1} - \tau_{k^*}$, and using (B.9) we compute $\tilde{a}_{k^*}$, $\tilde{a}_{k^*+1}$, $\Sigma_{k^*}^{-1}$, $\Sigma_{k^*+1}^{-1}$, and, for $i = 1, \ldots, N$, ρ_{ik^*}, ρ_{ik^*+1}, $\tilde{b}_{ik^*}$, and $\tilde{b}_{ik^*+1}$. The proposal is accepted with probability $min(1, \alpha)$ where α is

$$\alpha = \frac{l_k!}{l_{k^*}!} \frac{l_{k+1}!}{l_{k^*+1}!} \frac{\Gamma(l_{k^*} + c)}{\Gamma(l_k + c)} \frac{\Gamma(l_{k^*+1} + c)}{\Gamma(l_{k+1} + c)} \prod_{i=1}^{N} \frac{|\Sigma_{k^*}|^{1/2}}{|\Sigma_k|^{1/2}} \frac{|\Sigma_{k^*+1}|^{1/2}}{|\Sigma_{k+1}|^{1/2}} \frac{\Gamma(\tilde{a}_{k^*})}{\Gamma(\tilde{a}_k)} \frac{\Gamma(\tilde{a}_{k^*+1})}{\Gamma(\tilde{a}_{k+1})} \frac{\tilde{b}_{ik}^{\tilde{a}_k}}{\tilde{b}_{ik^*}^{\tilde{a}_{k^*}}} \frac{\tilde{b}_{ik+1}^{\tilde{a}_{k+1}}}{\tilde{b}_{ik^*+1}^{\tilde{a}_{k^*+1}}}. \tag{D.3}$$

All but the two regimes separated by the breakpoint that is being perturbed remain unchanged and thus do not enter the acceptance calculation.

If τ_{k^*} is rejected, we simply discard it and then attempt to perturb the next breakpoint. If accepted, we replace τ_k with τ_{k^*} and replace the existing values of l, $\tilde{a}$, $\tilde{b}$, ρ, and Σ^{-1} for the current regimes k and $k+1$ with the corresponding accepted proposed values in regimes k^* and k^*+1.

[41]The value of s is tuned to achieve the desired acceptance ratio of 0.25.

Appendix D.3. Estimating the Number of breakpoints

In many economic applications the number of breaks is not known in advance and so must be determined according to some objective function. For example, one can compare a model with $K = 2$ breaks to a model with a single break ($K = 1$) through the ratio

$$BF_{21} = \frac{P(K = 2 \mid r)/P(K = 1 \mid r)}{P(K = 2)/P(K = 1)} = \frac{p(r \mid K = 2)}{p(r \mid K = 1)}, \tag{D.4}$$

which captures the change in the odds in favour of model 2 relative to model 1 as we move from the prior to the posterior. In the absence of an informative prior model probability, the marginal likelihoods are proportional to the posterior model probabilities which are the key quantities of interest. There are a number of approaches for estimating the posterior model probabilities and associated Bayes factors for competing models, each corresponding to a different number of breaks.

Chib (1998)'s algorithm fixes the number of breaks in advance, which leads to a nonuniform prior distribution on the breakpoint locations (Koop and Potter 2009) while also restricting the model to a monotonically decreasing geometric prior on the regime durations. Koop and Potter (2007) develop a procedure that overcomes these problems. Their hierarchical hidden Markov (HHM) model approach can, however, be undesirable since it may be sensitive to how the chain is initialised.[42]

To overcome these problems, we adopt the reversible jump Markov chain Monte Carlo approach that includes the number of breaks K as a parameter in the model and explores both the model and parameter space jointly by 'jumping' between different numbers of breaks. The proportion of time spent at each number of breaks is equal to the posterior model probabilities. Using conjugate priors on the regression parameters μ, β and σ allows us to marginalise them from the posterior (Proposition 1) and

[42]If the chain is incorrectly initialised, the HHM approach may not be able to escape certain parts of the parameter space and may lead to poor mixing and ultimately spurious inference as demonstrated by Fearnhead (2006).

explore the model space alone, greatly reducing the complexity of the algorithm.

To estimate the number of breakpoints, we begin the sampler at a given number of breakpoints. With equal probability we enter a birth or death move which, respectively, attempt to introduce a new, or remove an existing, breakpoint. We then compute the acceptance probability which ensures that detailed balance is maintained across the entire parameter space. If the move is accepted, the breakpoint vector is updated. Otherwise the proposal is discarded.

Birth Move: With probability $b_K = 0.5$ a birth move is entered.[43] This move attempts to increase K to $K+1$ and hence introduces a new breakpoint, τ_{k^*}, that is sampled uniformly from the discrete time series sample $\tau_{k^*} \sim U[1, T]$. If an existing breakpoint is proposed ($\tau_{k^*} \in \tau$) the move is immediately rejected.

Otherwise, let τ^* denote the proposed breakpoint vector which consists of τ and τ_{k^*}. Let k^c denote the existing regime we are attempting to split, i.e., $\tau_{k^c-1} < \tau_{k^*} < \tau_{k^c}$, and let k^* and k^*+1 denote the two new proposed regimes. Regime k^* contains observations $\tau_{k^c-1}+1, \ldots, \tau_{k^*}$, so that $l_{k^*} = \tau_{k^*} - \tau_{k^c-1}$. Regime k^*+1 contains observations $\tau_{k^*}+1, \ldots, \tau_{k^c}$ and hence $l_{k^*+1} = \tau_{k^c} - \tau_{k^*}$. We now compute $\tilde{a}_{k^*}$, $\tilde{a}_{k^*+1}$, $\Sigma_{k^*}^{-1}$, $\Sigma_{k^*+1}^{-1}$, and, for $i = 1, \ldots, N$, ρ_{ik^*}, ρ_{ik^*+1}, $\tilde{b}_{ik^*}$, and $\tilde{b}_{ik^*+1}$ using (B.9).

The birth move is accepted with probability $min(1, \alpha)$, where α equals

$$\alpha = \frac{p(\boldsymbol{r} \mid \boldsymbol{X}, \tau^*)}{p(\boldsymbol{r} \mid \boldsymbol{X}, \tau)} \times \frac{p(\tau^*)}{p(\tau)} \times \frac{T}{K+1} \times \frac{2}{2}, \tag{D.5}$$

where the final term 2/2 cancels because $b_k = 0.5$, T corresponds to the sampling of τ_{k^*}, and $(K+1)^{-1}$ corresponds to the uniform sampling of an existing breakpoint to

[43] If the sampler is at $K = T-1$, corresponding to a break occurring every period, we set $b_{T-1} = 0$.

return to K breakpoints if we had $K+1$ breakpoints.[44] We find that α simplifies to

$$\alpha = \frac{d^c}{\Gamma(c)} \frac{l_{k^c}!}{l_{k^*}!\, l_{k^*+1}!} \frac{\Gamma(l_{k^*+1}+c)}{\Gamma(l_{k^c}+c)} \frac{\Gamma(l_{k^*}+c)}{(d+1)^c} \frac{T}{K+1} \frac{b^{aN}}{\Gamma(a)^N} \times \prod_{i=1}^{N} \frac{\Gamma(\tilde{a}_{k^*})}{\tilde{b}_{ik^*}^{\tilde{a}_{k^*}}} \frac{\Gamma(\tilde{a}_{k^*+1})}{\tilde{b}_{ik^*+1}^{\tilde{a}_{k^*+1}}} \frac{\tilde{b}_{ik^c}^{\tilde{a}_{k^c}}}{\Gamma(\tilde{a}_{k^c})} \frac{|\Sigma_{k^*}|^{1/2} |\Sigma_{k^*+1}|^{1/2}}{\mid V_\theta \mid^{1/2} |\Sigma_{k^c}|^{1/2}}. \tag{D.6}$$

All but the existing regime we are proposing to split remain unchanged and thus do not enter into the acceptance probability.

If the move is rejected, we drop the proposal altogether. Otherwise we substitute τ^* for τ, while the corresponding values of l, Σ^{-1}, ρ, $\tilde{a}$, and $\tilde{b}$ are updated by removing their values for regime k^c and adding their values for regimes k^* and k^*+1.

Death Move: With probability $d_k = 1 - b_k = 0.5$ a death move is entered.[45] This move proposes to eliminate a breakpoint and move from K to $K-1$. The breakpoint τ_{k^c} we propose to eliminate is sampled uniformly from the set of existing breakpoints, that is, $\tau_{k^c} \sim U[\tau_1, \tau_K]$. Let τ^* denote the proposed breakpoint vector which is equal to τ with τ_{k^c} removed. Let k^c and k^c+1 denote the two existing regimes (divided by the breakpoint τ_{k^c}). We propose to replace regimes (k^c and k^c+1) with one longer regime (k^*). The length of the proposed regime is therefore the sum of the two existing regimes $l_{k^*} = l_{k^c} + l_{k^c+1}$. In similar fashion to the birth move, we compute $\tilde{a}_{k^*}$ and $\Sigma_{k^*}^{-1}$, and, for $i = 1, \ldots, N$, we compute ρ_{ik^*} and $\tilde{b}_{ik^*}$ using (B.9).

We accept the death move with a probability $min(1, \alpha)$, where α equals

$$\alpha = \frac{p(\boldsymbol{r} \mid \boldsymbol{X}, \tau^*)}{p(\boldsymbol{r} \mid \boldsymbol{X}, \tau)} \times \frac{p(\tau^*)}{p(\tau)} \times \frac{K}{T} \times \frac{2}{2}, \tag{D.7}$$

where the final term cancels because $b_k = d_k$, K corresponds to the uniform sam-

[44]The acceptance probability consists of proposal densities (i) from our current position (K) to the proposed position (K+1) and (ii) from our proposed position (K+1) back to our current position (K). This is required to ensure every move is 'reversible' and thus our algorithm does not move to a part of the parameter space it cannot return from which would cause the algorithm to stop mixing.

[45]If the sampler is at $K = 0$, $d_0 = 0$ since there is no existing breakpoint to remove.

pling of the breakpoint we are attempting to eliminate from the set of existing K breakpoints, and T^{-1} corresponds to the uniform sampling of a new breakpoint from the interval $U[1,T]$ if we were at $K-1$ breakpoints and attempting a birth move to return to our current position of K breakpoints. We simplify α to

$$\alpha = \frac{l_{k^c}!\, l_{k^c+1}!}{l_{k^*}!\, d^c} \frac{\Gamma(l_{k^*}+c)}{\Gamma(l_{k^c}+c)} \frac{(d+1)^c \Gamma(c)}{\Gamma(l_{k^c+1}+c)} \frac{K}{T} \frac{\Gamma(a)^N}{b^{aN}} \times \prod_{i=1}^{N} \frac{\tilde{b}_{ik^c}^{\tilde{a}_{k^c}}}{\tilde{b}_{ik^*}^{\tilde{a}_{k^*}}} \frac{\Gamma(\tilde{a}_{k^*})}{\Gamma(\tilde{a}_{k^c})} \frac{|\Sigma_{k^*}|^{1/2} |\ V_\theta\ |^{1/2}}{|\Sigma_{k^c}|^{1/2} |\Sigma_{k^c+1}|^{1/2}} \frac{\tilde{b}_{ik^c+1}^{\tilde{a}_{k^c+1}}}{\Gamma(\tilde{a}_{k^c+1})}. \tag{D.8}$$

If the move is accepted, we substitute τ^* for τ, and we update l, Σ^{-1}, ρ, $\tilde{a}$, and $\tilde{b}$ by removing their values for the existing regimes k^c and k^c+1 and adding their values for the new regime k^*.

Table 1: Parameter estimates for the dividend growth and dividend-price ratio models

Regime	Break dates	b_0: Intercept		b_1: Slope on $d_t - p_t$		b_2: Slope on Δd_t		Volatility	
		Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
				Dividend growth					
1	Feb 1931	-0.038	(0.005)	-0.082	(0.023)	0.135	(0.039)	0.201	(0.003)
2	May 1933	0.030	(0.004)	0.057	(0.024)	-0.193	(0.031)	0.156	(0.003)
3	Aug 1939	-0.127	(0.016)	-0.034	(0.066)	0.188	(0.043)	0.216	(0.010)
4	Mar 1945	0.021	(0.002)	-0.029	(0.018)	0.308	(0.022)	0.113	(0.001)
5	Oct 1968	0.019	(0.004)	-0.156	(0.022)	0.515	(0.032)	0.163	(0.003)
6	Jan 1987	0.027	(0.002)	-0.267	(0.014)	0.215	(0.029)	0.147	(0.002)
7	Dec 1998	0.002	(0.005)	-0.305	(0.021)	0.126	(0.045)	0.228	(0.004)
8	Sep 2007	0.072	(0.004)	-0.169	(0.019)	0.430	(0.041)	0.204	(0.003)
9	May 2009	0.019	(0.009)	-0.205	(0.022)	0.307	(0.069)	0.327	(0.006)
10	Dec 2015*	0.043	(0.010)	-0.187	(0.019)	0.834	(0.112)	0.545	(0.007)
Regime	Break dates	a_0: Intercept		a_1: Slope on $d_t - p_t$		a_2: Slope on Δd_t		Volatility	
		Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
				Dividend-price ratio					
1	Jun 1929	-0.020	(0.021)	0.745	(0.545)	0.185	(0.129)	0.412	(0.025)
2	Dec 1940	0.029	(0.009)	0.695	(0.417)	-0.224	(0.095)	0.323	(0.020)
3	May 1948	0.077	(0.018)	0.807	(0.239)	-0.150	(0.057)	0.285	(0.019)
4	Oct 1997	0.015	(0.005)	0.736	(0.242)	-0.065	(0.028)	0.145	(0.012)
5	May 2002	0.038	(0.015)	0.640	(0.230)	-0.173	(0.075)	0.120	(0.009)
6	Nov 2008	0.029	(0.007)	0.892	(0.441)	-0.261	(0.078)	0.227	(0.016)
7	Dec 2015*	0.023	(0.008)	0.700	(0.327)	-0.144	(0.060)	0.221	(0.014)

Table 1: Dividend growth and dividend-price ratio parameter estimates. This table displays the posterior mean and standard deviation (s.d.) of the intercept, volatility and the slopes on the lagged dividend-price ratio and lagged dividend growth rate when the growth rate (top panel) or the dividend-price ratio (bottom panel) is the dependent variable. Regime-specific estimates are from the heterogeneous panel break model, value-weighted averages across the 30 industry portfolios. The posterior modes of the identified break dates, the final time periods in the corresponding regimes, are also reported.* denotes that Dec 2015 is not a break date but the end of the sample.

Table 2: Parameter estimates for the Present Value model

Regime	Break dates	a_1	b_1	c_1	a_2	b_2	c_2	Volatility
1	Mar 1934	0.803	-0.011	0.255	-0.012	0.843	0.853	0.247
		(0.770)	(0.068)	(0.187)	(0.009)	(0.762)	(0.640)	(0.061)
2	Nov 1939	0.733	-0.055	0.223	-0.042	0.792	0.829	0.240
		(0.671)	(0.059)	(0.156)	(0.038)	(0.548)	(0.659)	(0.029)
3	Nov 1969	0.801	-0.101	0.203	-0.083	0.744	0.818	0.127
		(0.287)	(0.038)	(0.061)	(0.032)	(0.315)	(0.209)	(0.018)
4	Jul 1998	0.750	-0.167	0.139	-0.052	0.779	0.825	0.105
		(0.286)	(0.058)	(0.040)	(0.022)	(0.300)	(0.276)	(0.017)
5	Jun 2010	0.810	-0.120	0.158	-0.080	0.830	0.902	0.266
		(0.529)	(0.031)	(0.064)	(0.064)	(0.402)	(0.786)	(0.029)
6	Dec 2015*	0.778	-0.082	0.227	-0.016	0.861	0.875	0.184
		(0.109)	(0.030)	(0.071)	(0.007)	(0.342)	(0.196)	(0.021)

Table 2: Parameter estimates for the Present Value model. This table displays the regime-specific posterior mean (and standard deviation in brackets below) of various parameters for the cross-equation restricted present value model using our panel break modeling approach. The reported values are value-weighted averages across the parameter estimates on the 30 industry portfolios. The posterior modes of the identified break dates, the final time periods in the corresponding regimes, are also reported. * denotes that Dec 2015 is not a break date but the end of the sample.

Table 3: Bayes Factors for Break Type

			Industry Portfolios			
			Break Date			
	1934:03	1939:11	1969:11	1998:07	2010:06	
BF_{beta}	1.35	**172.40**	1.65	**104.12**	1.87	
BF_{rp}	2.32	2.08	1.45	2.39	1.19	
			Regime span			
	1926:07-1934:03	1934:04-1939:11	1939:12-1969:11	1969:12-1998:07	1998:08-2010:06	2010:07-2015:12
β	0.10	0.16	0.27	0.02	0.12	0.19
	(2.17)	(0.99)	(2.47)	(2.53)	(0.48)	(1.14)
λ	0.05	1.01	0.21	2.22	0.41	2.73
	(0.64)	(0.37)	(2.21)	(2.46)	(0.21)	(3.86)
			Individual Stocks			
			Break Date			
BF_{beta}	2.27	**118.44**	1.48	**72.17**	2.08	
BF_{rp}	1.65	1.42	2.17	2.70	1.25	
			Regime span			
	1926:07-1934:03	1934:04-1939:11	1939:12-1969:11	1969:12-1998:07	1998:08-2010:06	2010:07-2015:12
β	0.07	0.19	0.42	0.12	0.24	0.17
	(2.32)	(0.90)	(2.76)	(2.23)	(0.68)	(1.12)
λ	0.23	1.34	0.55	2.48	0.72	2.69
	(1.25)	(0.22)	(2.30)	(2.85)	(0.26)	(3.35)

Table 3: Bayes Factors for Break Type. The top two rows in each panel of this table display Bayes Factors that provide the strength of evidence in favour of each break affecting only betas (BF_{beta}) or risk premia (BF_{rp}). Both Bayes Factors being small implies that the break affects both betas and risk premia. Bayes Factors that present strong evidence in favour of breaks only in betas (risk premia) display BF_{beta} (BF_{rp}) in bold font. Bayes Factors are constructed from the marginal likelihood of each model which is computed using the procedure of Chib (1995). Marginal likelihoods are computed for three models: one that allows for breaks only in (i) risk premia (BF_{rp}), (ii) beta loadings (BF_{beta}) or (iii) both. We also display the posterior mean estimate, with t-statistics in brackets below, for the β and λ estimates in each regime. The upper (lower) panel shows results when using 30 industry portfolio (individual stock) return series and imposing the five breaks identified by the present value model.

Table 4: Average values of state variables across breaks

Variable			Regime			
	1926:07-1934:03	1934:04-1939:11	1939:12-1969:11	1969:12-1998:07	1998:08-2010:06	2010:07-2015:12
Mean excess return	8.39	11.08	10.03	6.74	-0.35	15.24
St.dev. of excess return	0.11	0.08	0.04	0.04	0.05	0.04
Mean div. growth rate	3.00	4.08	6.84	5.40	4.44	4.08
St.dev. of div. growth rate	0.53	0.80	1.31	1.12	0.89	0.82
Risk aversion	2.08	1.82	1.40	1.09	1.51	1.90
Economic Policy Uncertainty				106.39	90.17	106.39
Monetary policy				102.70	96.20	74.99
Three component uncertainty				102.76	100.82	103.67
Macro uncertainty			0.59	0.67	0.70	0.62
Financial uncertainty			0.85	0.91	0.98	0.85

Table 4: Average values of state variables across breaks. This table displays in each of the six regimes identified by the present value breakpoint model the mean and volatility of (i) excess returns and (ii) the dividend growth rate. We also report the risk aversion estimate in each regime. The mean excess return and mean dividend growth rates are expressed as annualized percentages. Finally, we show the mean of various economic uncertainty data that are sourced from Sydney Ludvigson's website.

Table 5: Statistical significance of gains in predictive accuracy

Predictor	t <-1.64		-1.64< t <0		0< t <1.64		t >1.64	
				No break panel				
dp	321	(2.86%)	726	(6.48%)	1,021	(9.11%)	9,143†	(81.55%)
tbl	186	(1.66%)	1,346	(12.01%)	789	(7.04%)	8,890†	(79.29%)
tms	451	(4.02%)	1,755	(15.65%)	952	(8.49%)	8,053†	(71.83%)
dfs	89	(0.79%)	2,089	(18.63%)	921	(8.22%)	8,112†	(72.36%)
ep	275	(2.45%)	881	(7.86%)	563†	(5.02%)	9,492	(84.67%)
de	579	(5.16%)	450	(4.01%)	809	(7.22%)	9,373†	(83.61%)
bm	300	(2.68%)	1,103	(9.84%)	780	(6.96%)	9,028†	(80.53%)
dfr	269	(2.39%)	1,207	(10.77%)	421	(3.76%)	9,314†	(83.08%)
lty	177	(1.58%)	1,506	(13.43%)	608†	(5.42%)	8,920	(79.56%)
infl	404	(3.60%)	407	(3.63%)	347	(3.09%)	10,053†	(89.67%)
ltr	289	(2.58%)	880	(7.85%)	459	(4.09%)	9,583†	(85.48%)
svar	333	(2.97%)	1,542	(13.75%)	2,019	(18.91%)	7,317†	(65.27%)
csp	282	(2.52%)	911	(8.13%)	859	(7.66%)	9,159†	(81.69%)
ntis	1,090†	(9.72%)	4,562	(40.69%)	2,414	(21.53%)	3,145	(28.05%)
				Prevailing mean				
dp	237	(2.11%)	1,331	(11.87%)	718	(6.40%)	8,925†	(79.61%)
tbl	442	(3.94%)	976	(8.71%)	781†	(6.97%)	9,012	(80.39%)
tms	89	(0.79%)	1,489	(13.28%)	178†	(1.59%)	9,455	(84.34%)
dfs	125	(1.11%)	2,201	(19.63%)	437	(3.89%)	8,448†	(75.35%)
ep	368	(3.28%)	1,209	(10.78%)	2,743	(24.467%)	6,891†	(61.47%)
de	890	(7.94%)	566	(5.05%)	2,256	(20.12%)	7,499†	(66.89%)
bm	290	(2.59%)	896	(7.99%)	2,128	(18.98%)	7,897†	(70.44%)
dfr	387	(3.45%)	2,010†	(17.93%)	247	(2.20%)	8,567	(76.42%)
lty	333	(2.97%)	1,355	(12.09%)	147	(1.31%)	9,376†	(83.63%)
infl	622	(5.55%)	1,109	(9.89%)	1,155	(10.30%)	8,325†	(74.26%)
ltr	245	(2.19%)	1,678	(14.97%)	498	(4.44%)	8,790†	(78.41%)
svar	410	(3.66%)	1,245	(11.11%)	1,762	(15.72%)	7,794†	(69.52%)
csp	373	(3.33%)	1,204	(10.74%)	1,385	(12.35%)	8,249†	(73.58%)
ntis	3,121†	(27.84%)	5,576	(49.74%)	2,097	(18.70%)	417	(3.72%)
				Time series break				
dp	522	(4.67%)	1,233	(10.99%)	124	(1.11%)	9,332†	(83.24%)
tbl	235	(2.09%)	1,457	(12.99%)	1,144	(10.20%)	8,375†	(74.70%)
tms	321	(2.86%)	2,090	(18.64%)	322	(2.87%)	8,478†	(75.62%)
dfs	346	(3.09%)	1,671	(14.91%)	193	(1.72%)	9,001†	(80.29%)
ep	420	(3.75%)	890	(7.94%)	2,136	(19.05%)	7,765†	(69.26%)
de	76	(0.68%)	1,467	(13.09%)	2,778	(24.78%)	6,890†	(61.45%)
bm	98	(0.87%)	2,293	(20.45%)	1,500	(13.38%)	7,320†	(65.29%)
dfr	156	(1.39%)	765	(6.82%)	1,981	(17.67%)	8,309†	(74.11%)
lty	444	(3.96%)	1,290	(11.51%)	1,585	(14.14%)	7,892†	(70.39%)
infl	300	(2.68%)	1,578	(14.08%)	477†	(4.25%)	8,856	(78.99%)
ltr	202	(1.80%)	1,220	(10.88%)	757†	(6.75%)	9,032	(80.56%)
svar	187	(1.67%)	567	(5.06%)	3,566	(31.81%)	6,891†	(61.47%)
csp	612	(5.46%)	1,194	(10.65%)	1,217	(10.86%)	8,188†	(73.04%)
ntis	1,548	(13.81%)	2,187†	(19.51%)	3,209	(28.62%)	4,267	(38.06%)
				Time-varying parameter				
dp	230	(2.05%)	465	(4.15%)	2,608	(23.26%)	7,908†	(70.54%)
tbl	121	(1.08%)	1,090	(9.72%)	1,578	(14.08%)	8,422†	(75.12%)
tms	357	(3.18%)	899	(8.02%)	2,599	(23.18%)	7,356†	(65.61%)
dfs	400	(3.57%)	1,351	(12.05%)	1,369	(12.21%)	8,091†	(72.17%)
ep	509	(4.54%)	905	(8.07%)	295	(2.63%)	9,502†	(84.76%)
de	397	(3.49%)	498	(4.44%)	1,073	(9.57%)	9,243†	(82.45%)
bm	128	(1.14%)	2,010	(17.93%)	2,293	(20.45%)	6,780†	(60.48%)
dfr	98	(0.87%)	1,109	(9.89%)	2,148	(19.16%)	7,856†	(70.07%)
lty	75	(0.67%)	855	(7.63%)	2,046†	(18.25%)	8,235	(73.45%)
infl	256	(2.28%)	469	(4.18%)	3,162	(28.20%)	7,324†	(65.33%)
ltr	420	(3.75%)	778	(6.94%)	823†	(7.34%)	9,190	(81.97%)
svar	330	(2.94%)	1,091	(9.73%)	1,357	(12.10%)	8,433†	(75.22%)
csp	240	(2.14%)	764	(6.81%)	2,255	(20.11%)	7,952†	(70.93%)
ntis	3,890†	(34.69%)	6,244	(55.69%)	121	(1.08%)	956	(8.53%)

Table 5: 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 prevailing mean, the time series model with breaks applied to each stock in turn (Time series break), and the time-varying parameter model when forecasting firm-level stock returns with either the dividend-price ratio (dp), treasury-bill rate (tbl), term spread (tms), default spread (dfs), earnings-price ratio (dp), dividend payout ratio (de), book-to-market ratio (bm), default return spread (dfr), long term yield (lty), inflation (infl), long term return (ltr), stock variance (svar), corporate equity issuance (ntis), or cross-sectional premium (csp). Significance is evaluated using the procedure of Clark and West (2007). For each procedure the table displays the number of stocks (with the percentage in brackets on the right-hand side) for which our method produces significantly worse, insignificantly worse, insignificantly better, and significantly better forecasts at the 10% level. The market portfolio is constructed as the value-weighted average of the individual return forecasts. † indicates the particular bin in which the market portfolio lies. We only include stocks that have at least 60 out-of-sample observations to robustify our results. This leaves 11,210 stocks.

Table 6: Statistical significance of gains in predictive accuracy: PV breaks

Predictor	t <-1.64		-1.64< t <0		0< t <1.64		t >1.64	
				No break panel				
dp	622	(5.55%)	1,354	(12.08%)	1,589	(14.17%)	7,646†	(68.20%)
tbl	289	(2.58%)	1,370	(12.22%)	583	(5.20%)	8,969†	(80.00%)
tms	1,050	(9.37%)	1,663	(14.83%)	744	(6.64%)	7,754†	(69.16%)
dfs	176	(1.57%)	659	(5.88%)	1,854	(16.54%)	8,522†	(76.01%)
ep	851	(7.59%)	2,043	(18.22%)	378	(3.37%)	7,939†	(70.81%)
de	480	(4.28%)	1,562	(13.93%)	1,332†	(11.88%)	7,837	(69.90%)
bm	376	(3.35%)	1,244	(11.09%)	1,271†	(11.34%)	8,320	(74.21%)
dfr	1,085	(9.68%)	987	(8.80%)	662	(5.90%)	8,477†	(75.61%)
lty	460	(4.10%)	1,879	(16.76%)	1,258	(11.22%)	7,614†	(67.92%)
infl	65	(0.58%)	1,201	(10.71%)	482	(4.29%)	9,463†	(84.41%)
ltr	335	(2.99%)	894	(7.97%)	642	(5.73%)	9,340†	(83.31%)
svar	421	(3.76%)	1,286	(11.47%)	1,105	(9.86%)	8,399†	(74.92%)
csp	1,285	(11.46%)	1,053	(9.39%)	650	(5.79%)	8,223†	(73.35%)
ntis	720	(6.42%)	2,118	(18.89%)	1,582	(14.11%)	6,791†	(60.57%)
				Prevailing mean				
dp	289	(2.58%)	875	(7.80%)	2,205	(19.67%)	7,842†	(69.95%)
tbl	175	(1.56%)	1,001	(8.93%)	1,652†	(14.74%)	8,383	(74.77%)
tms	560	(4.99%)	1,124	(10.02%)	752†	(6.71%)	8,775	(78.27%)
dfs	388	(3.46%)	2,573	(22.95%)	652	(5.82%)	7,598†	(67.77%)
ep	957	(8.54%)	1,684	(15.02%)	1,785	(15.92%)	6,785†	(60.52%)
de	851	(7.59%)	393	(3.51%)	947	(8.45%)	9,020†	(80.46%)
bm	294	(2.62%)	1,490	(13.29%)	2,008	(17.91%)	7,419†	(66.18%)
dfr	121	(1.08%)	1,357†	(12.10%)	897	(8.00%)	8,836	(78.82%)
lty	385	(3.43%)	999	(8.91%)	628	(5.60%)	9,199†	(82.05%)
infl	768	(6.85%)	985	(8.79%)	1,739	(16.26%)	7,508†	(66.97%)
ltr	1,120	(9.99%)	1,833	(16.35%)	520	(4.64%)	7,738†	(69.02%)
svar	243	(2.17%)	1,542	(13.75%)	1,186	(9.03%)	8,096†	(72.21%)
csp	481	(4.29%)	2,003	(17.87%)	756	(6.74%)	7,971†	(71.01%)
ntis	1,285	(11.46%)	1,959	(17.47%)	1,598†	(14.25%)	6,369	(56.81%)
				Time series break				
dp	559	(4.99%)	1,362	(12.15%)	997	(8.89%)	8,293†	(73.97%)
tbl	1,088	(9.70%)	2,120	(18.91%)	246	(2.19%)	7,757†	(69.91%)
tms	89	(0.79%)	2,308	(20.59%)	456	(4.07%)	8,358†	(74.55%)
dfs	374	(3.34%)	1,451	(12.94%)	822	(7.33%)	8,564†	(76.39%)
ep	289	(2.58%)	642	(5.73%)	1,645	(14.67%)	8,635†	(77.02%)
de	419	(3.74%)	1,325	(11.82%)	1,280	(11.42%)	8,187†	(73.03%)
bm	307	(2.74%)	2,451	(21.86%)	825	(7.36%)	7,628†	(68.04%)
dfr	62	(0.55%)	1,830†	(16.32%)	1,457	(12.99%)	7,862	(70.13%)
lty	261	(2.33%)	1,952	(17.41%)	1,354	(12.08%)	7,644†	(68.18%)
infl	1,217	(10.86%)	487	(4.34%)	355†	(3.17%)	9,152	(81.63%)
ltr	863	(7.69%)	1,354	(12.08%)	64	(0.57%)	8,930†	(79.65%)
svar	208	(1.86%)	1,409	(12.57%)	1,757	(15.67%)	7,837†	(69.90%)
csp	1,205	(10.75%)	1,036	(9.24%)	647	(5.77%)	8,323†	(74.24%)
ntis	1,873	(16.71%)	1,255	(11.19%)	2,904	(25.90%)	5,179†	(46.19%)
				Time-varying parameter				
dp	559	(4.99%)	1,021	(9.11%)	1,483	(13.23%)	8,148†	(72.68%)
tbl	266	(2.37%)	954	(8.51%)	1,638	(14.61%)	8,353†	(74.51%)
tms	672	(5.99%)	800	(7.14%)	1,914	(17.07%)	7,825†	(69.79%)
dfs	145	(1.29%)	1,682	(15.00%)	1,365	(12.18%)	8,019†	(71.53%)
ep	775	(6.91%)	1,690	(15.07%)	881	(7.86%)	7,865†	(70.15%)
de	324	(2.89%)	1,006	(8.97%)	1,485	(13.25%)	8,396†	(74.89%)
bm	286	(2.55%)	2,206	(19.68%)	1,375	(12.26%)	7,344†	(65.51%)
dfr	56	(0.49%)	1,858	(16.57%)	461	(4.11%)	8,836†	(78.82%)
lty	792	(7.06%)	1,325	(11.82%)	1,650	(14.72%)	7,444†	(66.39%)
infl	373	(3.33%)	800†	(7.14%)	1,499	(13.37%)	8,539	(76.17%)
ltr	500	(4.46%)	469	(4.18%)	804†	(7.17%)	9,438	(84.19%)
svar	321	(2.86%)	1,309	(11.68%)	558	(4.98%)	9,023†	(80.48%)
csp	773	(6.89%)	860	(7.67%)	1,352†	(12.06%)	8,226	(73.37%)
ntis	1,845	(16.46%)	1,240	(11.06%)	1,341	(11.96%)	6,785†	(60.52%)

Table 6: Statistical significance of forecast improvements imposing the break dates from the Present Value model. This table reports the statistical significance of the gains in predictive accuracy for our panel break return prediction model that imposes the break dates estimated from the present value model using only the data available at the time the forecast is made. Forecasts perform Bayesian Model Averaging across any uncertainty surrounding the number and timing of breaks. Performance is reported relative to the heterogeneous panel model with no breaks (No break panel), the prevailing mean, the time series model with breaks applied to each stock in turn (Time series break), and the time-varying parameter model when forecasting firm-level stock returns with either the dividend-price ratio (dp), treasury-bill rate (tbl), term spread (tms), default spread (dfs), earnings-price ratio (dp), dividend payout ratio (de), book-to-market ratio (bm), default return spread (dfr), long term yield (lty), inflation (infl), long term return (ltr), stock variance (svar), corporate equity issuance (ntis), or cross-sectional premium (csp). Significance is evaluated using the procedure of Clark and West (2007). For each procedure the table displays the number of stocks (with the percentage in brackets on the right-hand side) for which our method produces significantly worse, insignificantly worse, insignificantly better, and significantly better forecasts at the 10% level. The market portfolio is constructed as the value-weighted average of the individual return forecasts. † indicates the particular bin in which the market portfolio lies. We only include stocks that have at least 60 out-of-sample observations to robustify our results. This leaves 11,210 stocks.

Table 7: Magnitude of break by decile portfolio

Portfolio	Size of break rank	MSFD	Utility Gain rank	Utility gain
High	1	0.0331	1	2.77
9	2	0.0241	2	2.54
8	3	0.0197	5	2.20
7	4	0.0166	3	2.39
6	5	0.0158	4	2.37
5	6	0.0150	7	2.00
4	7	0.0112	6	2.11
3	8	0.0098	8	1.94
2	9	0.0072	9	1.85
Low	10	0.0070	10	1.81

Table 7: Magnitude of break by decile portfolios. This table lists in descending order the decile portfolios constructed according to our break risk measure: the mean of the squared differences (MSFD) between the forecasts generated for the entire sample by the baseline panel models with and without breaks. Portfolios are rebalanced each month using stocks with non-missing observations in that month based on our break risk measure. 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 aggregate dividend-price ratio as the predictive variable. The mean-variance investor at each period allocates his wealth between the risk-free rate (T-bills) and the risky (decile) portfolio.

Table 8: Utility gains from portfolio investment strategies

Predictor	hist avg	no brk	ts	tvp
		Baseline Model		
dp	2.79	3.02	2.37	1.84
tbl	3.12	2.43	2.31	1.91
tms	1.78	3.30	2.44	1.89
dfs	1.77	2.00	2.45	1.90
ep	1.96	3.09	2.55	1.78
de	3.03	2.33	2.08	2.85
bm	3.00	1.94	2.56	2.32
dfr	1.99	2.65	1.82	1.74
lty	3.21	1.86	2.09	1.97
infl	2.22	2.97	2.04	3.05
ltr	1.78	2.35	3.01	2.22
svar	2.05	2.34	1.88	1.79
csp	2.63	1.89	3.22	2.47
ntis	-0.75	-0.43	-0.25	-1.03
		20-year regime duration prior		
dp	2.15	2.59	2.34	1.95
tbl	2.64	2.77	2.00	2.15
tms	1.65	2.19	2.57	1.90
dfs	1.97	1.82	2.96	1.89
ep	2.14	2.03	3.05	2.32
de	3.11	2.64	1.88	2.50
bm	3.00	2.16	2.44	2.28
dfr	1.74	2.20	1.80	1.68
lty	3.42	1.90	1.65	2.43
infl	2.38	2.89	2.02	2.52
ltr	1.82	2.33	2.49	2.47
svar	2.74	2.10	1.63	1.42
csp	3.04	2.87	2.53	2.13
ntis	-0.11	-0.33	0.02	-0.27
		Imposing Breaks from PV Model		
dp	2.25	2.64	1.95	2.18
tbl	3.31	2.43	1.92	2.27
tms	1.96	2.32	2.55	1.94
dfs	1.75	2.04	3.11	1.90
ep	2.32	3.19	2.42	1.75
de	2.86	2.32	2.03	1.99
bm	3.24	1.96	2.28	3.14
dfr	2.01	1.85	1.84	1.37
lty	2.27	1.95	1.32	2.49
infl	2.41	2.90	1.85	2.75
ltr	1.89	2.55	1.88	2.29
svar	2.42	1.86	1.52	1.44
csp	3.00	2.85	2.20	2.34
ntis	1.12	0.85	1.43	0.99

Table 8: Utility gains. The top panel of this table reports the out-of-sample utility gain (certainty equivalent return) across the entire out-of-sample period 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 individual stocks when forecasting with the baseline model. We report the utility gain measured relative to each of the four benchmark models, namely, the prevailing mean (hist avg), the panel model with no breaks (no brk), the time series break model (ts), and the time-varying parameter model (tvp). The middle panel reports the same results when using a 20-year prior expected regime duration and the lower panel for when the recursively estimated break dates from the present value model are imposed on the baseline return prediction model. Results are presented for the 14 predictors we consider: the dividend-price ratio (dp), T-bill rate (tbl), term spread (tms), default yield spread (dfs), earnings-price ratio (ep), dividend payout ratio (de), book-to-market (bm), default return spread (dfr), long term yield (lty), inflation (infl), long term return (ltr), stock variance (svar), corporate equity issuance (ntis), and cross-sectional premium (csp). The reported certainty equivalent returns are expressed as annualised percentages.

Table 9: Statistical significance of gains in predictive accuracy from multivariate models

Predictor	t <-1.64		-1.64< t <0		0< t <1.64		t >1.64	
Multivariate models								
Dividend-price ratio and default spread								
hist avg	3,769	(33.62%)	2,044	(18.23%)	2,771	(24.72%)	2,627†	(23.43%)
no brk	1,897	(16.92%)	1,435	(12.80%)	4,807†	(42.88%)	3,072	(27.40%)
ts brk	980	(8.74%)	2,851	(25.43%)	3,173	(28.31%)	4,207†	(37.53%)
tvp	2,605	(23.24%)	3,001	(26.77%)	2,892†	(25.79%)	2,713	(24.20%)
Dividend-price ratio, default spread, term spread, and T-bill rate								
hist avg	6,455†	(57.58%)	2,474	(22.07%)	1,505	(13.43%)	777	(6.93%)
no brk	1,627†	(14.51%)	2,830	(25.25%)	4,101	(36.58%)	2,653	(23.67%)
ts brk	592	(5.28%)	3,454†	(30.81%)	3,309	(29.52%)	3,856	(34.39%)
tvp	1,922	(17.15%)	3,371	(30.07%)	3,221†	(28.73%)	2,697	(24.06%)
Equal-weighted forecast combinations from univariate models								
Dividend-price ratio and default spread								
hist avg	1,548	(13.81%)	1,246	(11.12%)	1,492	(13.31%)	6,925†	(61.78%)
no brk	290	(2.59%)	642	(5.73%)	895	(7.98%)	9,384†	(83.71%)
ts brk	388	(3.46%)	1,085	(9.68%)	189	(1.69%)	9,549†	(85.18%)
tvp	521	(4.65%)	1,100	(9.81%)	892	(7.96%)	8,698†	(77.58%)
Dividend-price ratio, default spread, term spread, and T-bill rate								
hist avg	963	(8.59%)	1,392	(12.42%)	2,461	(21.95%)	6,395†	(57.04%)
no brk	248	(2.21%)	547	(4.88%)	842	(7.51%)	9,574†	(85.39%)
ts brk	159	(1.42%)	962	(8.58%)	245	(2.19%)	9,845†	(87.82%)
tvp	389	(3.47%)	909	(8.11%)	1,087	(9.69%)	8,826†	(78.73%)
All 14 predictors								
hist avg	352	(3.14%)	771	(6.88%)	1,001	(8.93%)	9,087†	(81.05%)
no brk	122	(1.09%)	471	(4.20%)	443	(3.95%)	10,175†	(95.58%)
ts brk	202	(1.80%)	310	(2.77%)	356	(3.18%)	10,343†	(92.26%)
tvp	267	(2.38%)	631	(5.63%)	853	(7.61%)	9,460†	(84.38%)

Table 9: Statistical significance of forecast improvements from multivariate models. The top panel of this table reports the statistical significance of any gains in predictive accuracy for our multivariate panel break model that uses both the dividend-price ratio and default spread (top half of top panel) or the dividend-price ratio, default spread, term spread, and T-bill rate (lower half of top panel) relative to either the prevailing mean model (hist avg), our panel model without breaks (no brk), the univariate time series breakpoint model (ts brk) or the time-varying parameter model (tvp). Significance is evaluated using the procedure of Clark and West (2007). The table displays the number of cases (stocks) for which our method produces significantly worse, insignificantly worse, insignificantly better, and significantly better forecasts at the 10% level, with the percentage of total stocks reported in brackets on the right-hand-side. The market portfolio is constructed as the value-weighted average of the individual forecasts. † indicates the particular bin in which the t-statistic for the market portfolio lies. The lower panel reports results when equally-weighting the forecasts from our univariate panel breakpoint models that use either the dividend-price ratio and default spread (top third) or the dividend-price ratio, default spread, term spread, and T-bill rate (middle third) or all 14 predictors (lower third).

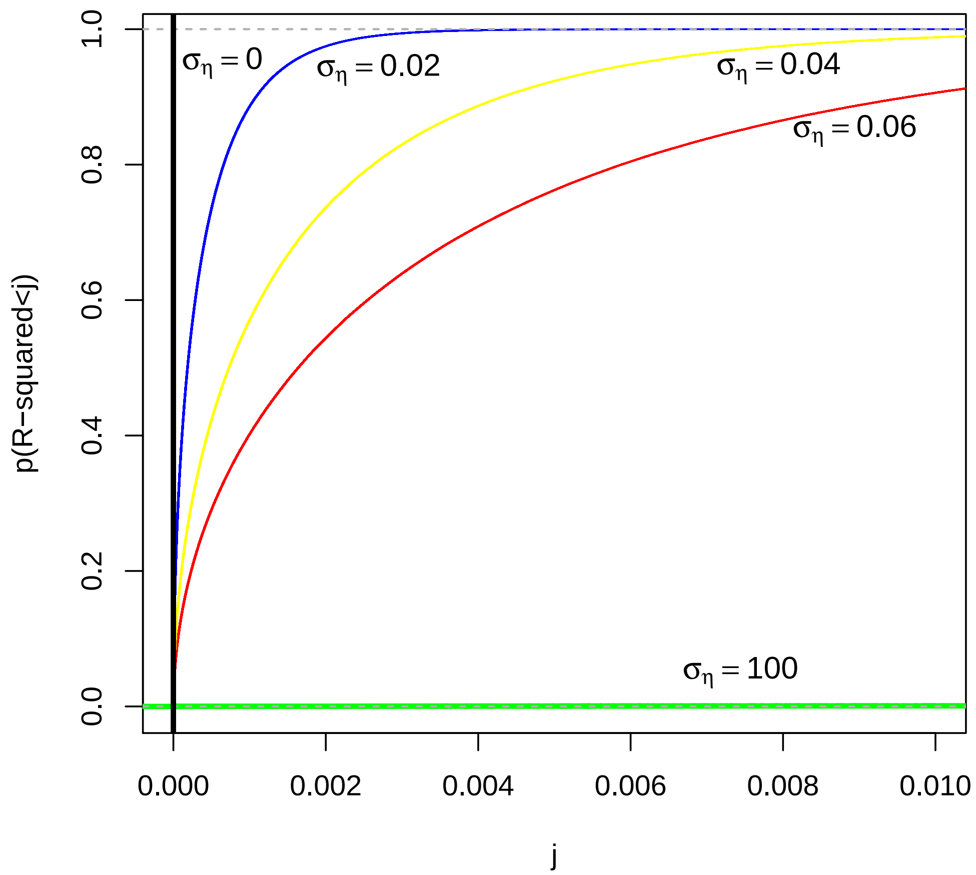

Figure 1: This figure displays the prior probability that the R-squared of a predictive regression lies below a certain value j, ranging from 0 to 0.01, for different degrees of scepticism regarding predictability. The investor's degree of scepticism is captured by the prior standard deviation of the 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.

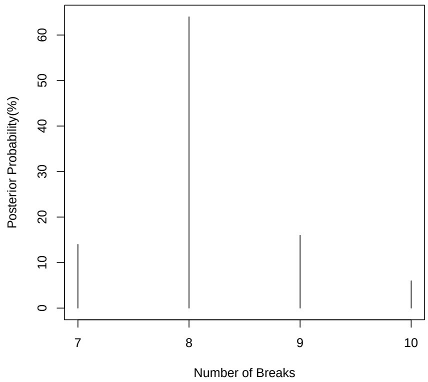

(a) Posterior Model Probabilities

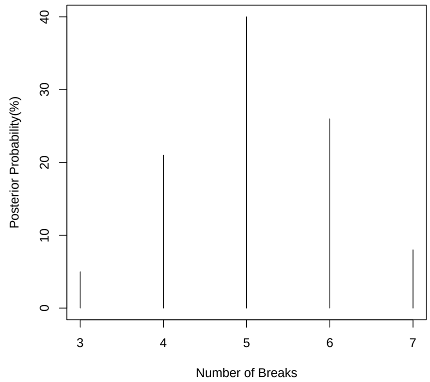

(b) Posterior Model Probabilities

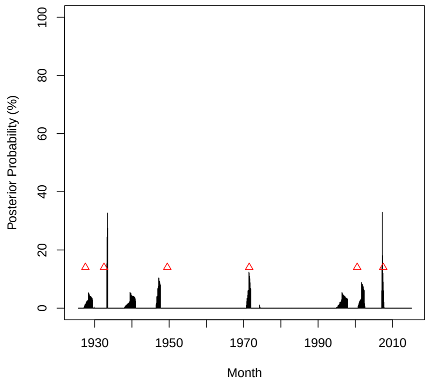

(c) Posterior Break Locations

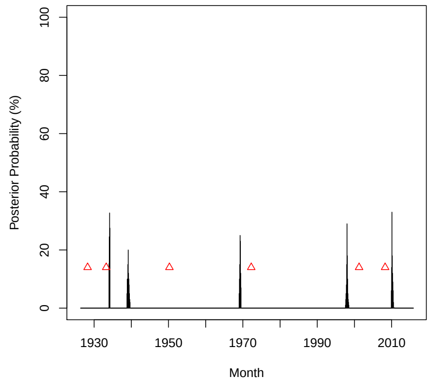

(d) Posterior Break Locations

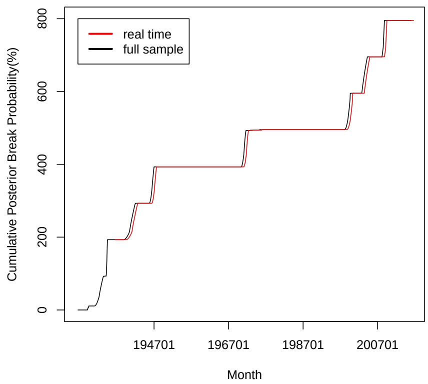

(e) Cumulative Break Probability

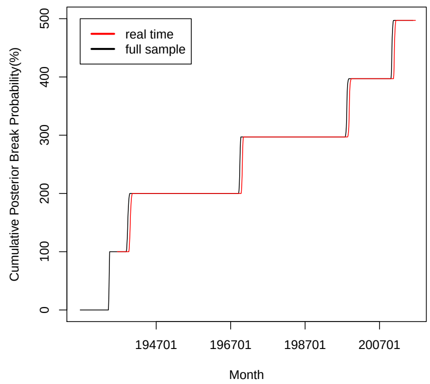

(f) Cumulative Break Probability

Figure 2: This figure displays the posterior probabilities for the number of breaks (top panel), the posterior break dates (middle) and the cumulative posterior break probability (bottom) when the prior expected regime duration is ten years estimated from our Bayesian panel break model that regresses firm-level returns on the lagged aggregate dividend-price ratio (left panel) or from the Present Value model that uses 30 industry portfolios and imposes the cross-equation restrictions (right panel). The red triangles in the middle panel mark the break dates estimated by the frequentist panel breakpoint approach of Baltagi et al. (2016). The red line in the bottom panel denotes the real time estimates.

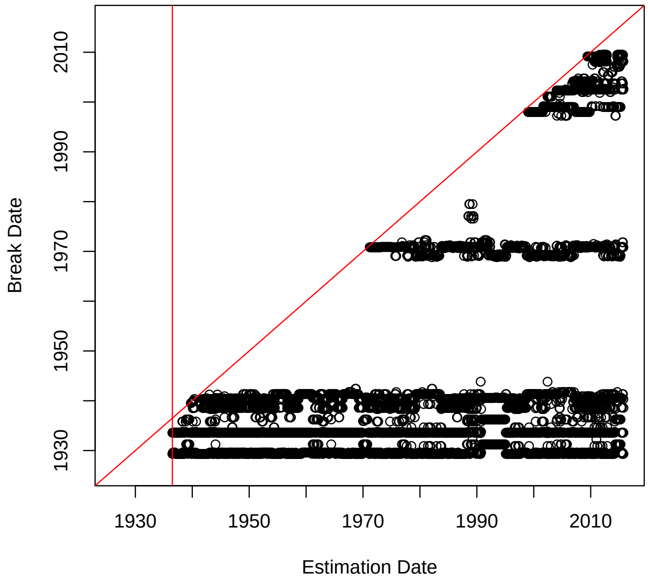

(a) Recursively Estimated Break Dates

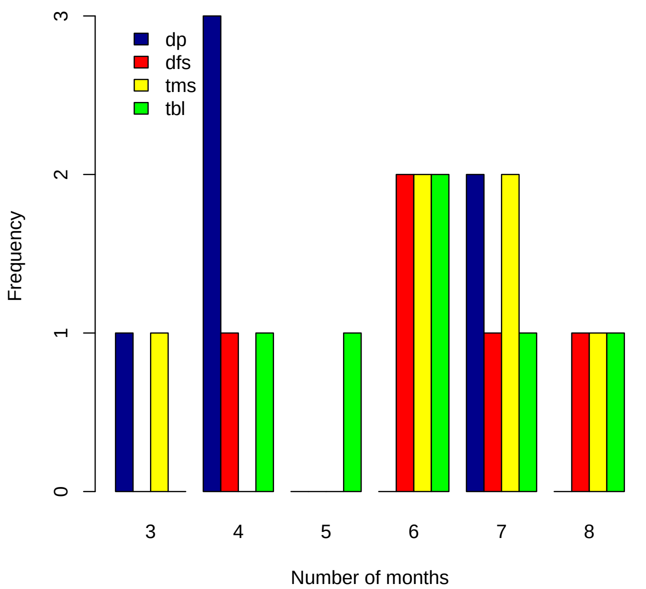

(b) Real-time Break Detection Delay

Figure 3: The top panel of this figure displays the real-time break detection obtained from our panel break model fitted to individual stock returns. 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. The bottom window displays the number of months it took to first detect each of the breaks that occurred after the initial estimation period when predicting with four predictors: the dividend-price ratio (dp), T-bill rate (tbl), term spread (tms), and default yield spread (dfs).

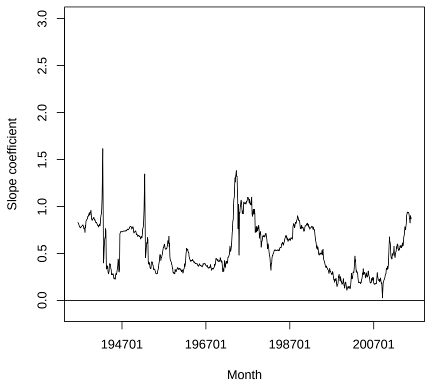

(a) Slope coefficient

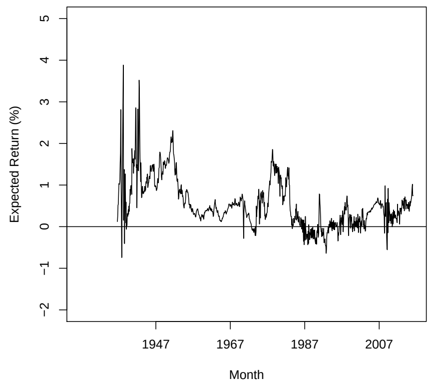

(b) Expected Return

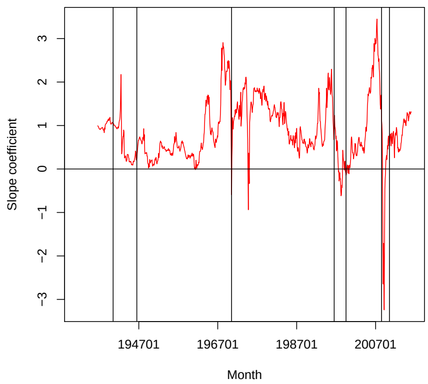

(c) Slope coefficient (10-year)

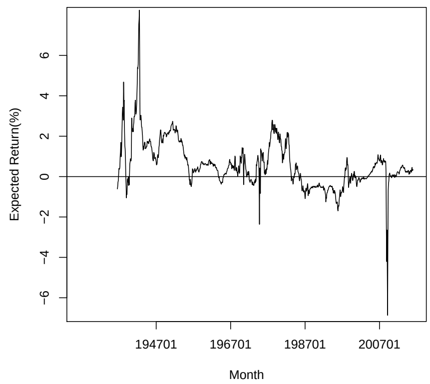

(d) Expected Return (10-year)

Figure 4: The top left window of this figure graphs the value-weighted posterior mean slope coefficient real time estimate from our Bayesian breakpoint model that regresses firm-level returns on the lagged aggregate dividend-price ratio. The top right window graphs the corresponding expected return real time estimate, computed as the posterior mean of the intercept plus the posterior mean of the slope multiplied by the value of the predictor in that month. The bottom windows display the value-weighted slope coefficient and expected return estimates from a 10-year rolling window OLS regression. The vertical lines in the bottom left window mark the posterior mode break dates estimated from our breakpoint model that occur after the first month in which the 10-year rolling window model is estimated: June 1936.

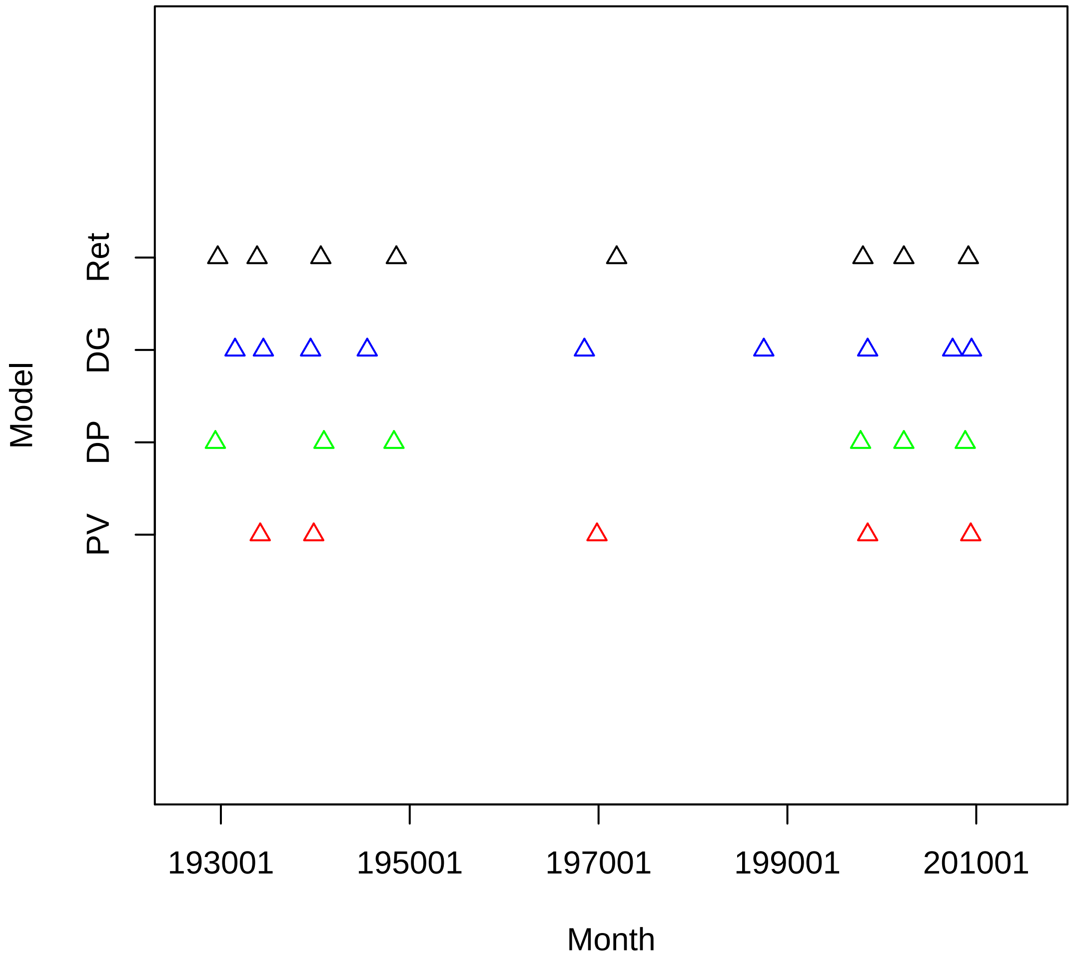

Figure 5: This figure displays the posterior mode break dates estimated by our panel breakpoint model for a series of models. First, the baseline return prediction model that uses the aggregate dividend-price ratio as the predictor (black triangles). Second, the dividend growth model that regresses industry-level dividend growth rates on lagged industry-level dividend growth rates and lagged industry-level dividend-price ratios (blue triangles). Third, the dividend-price ratio model that regresses industry-level dividend-price ratios on lagged industry-level dividend growth rates and lagged industry-level dividend-price ratios (green triangles). Fourth, the present value model (red triangles) that imposes the cross-equation restrictions in (17). Each model uses 30 industry portfolios.

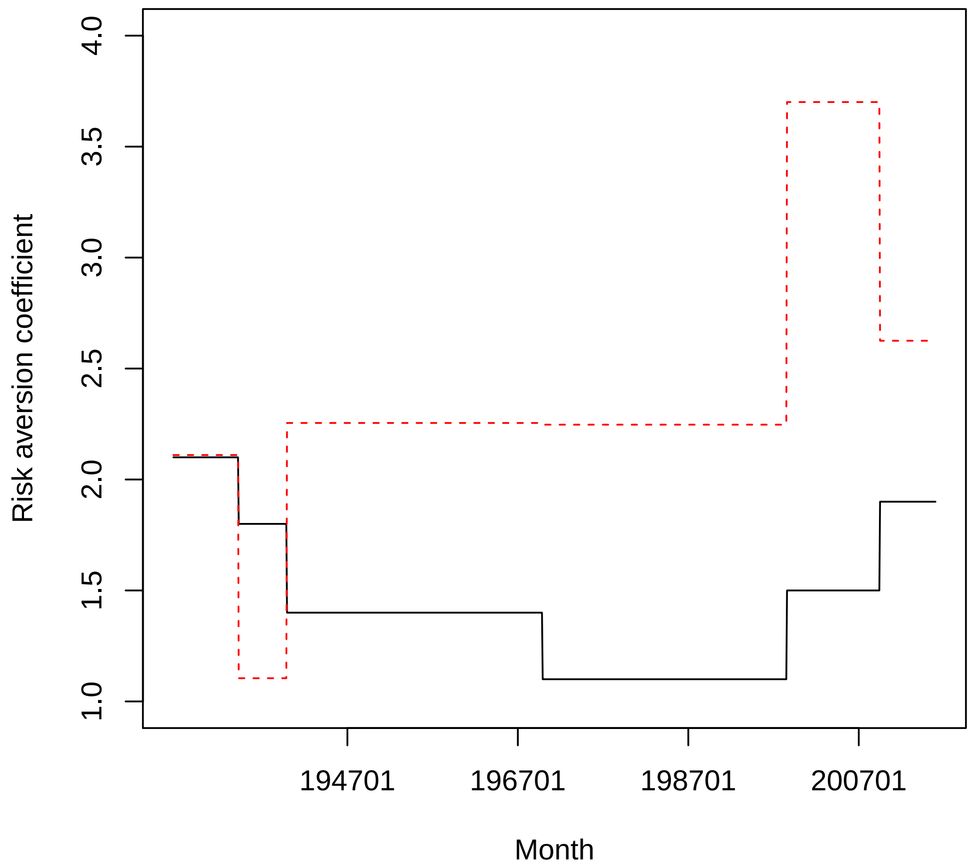

Figure 6: This figure graphs the monthly risk aversion coefficient (solid black line) which is the posterior mean of the pooled slope coefficient estimated using our Bayesian panel breakpoint model when firm-level stock returns are the independent variable and the predictive variable is stock variance (svar) when imposing the five break dates estimated from the PV model. The red dashed line graphs the t-statistic of the slope coefficient.

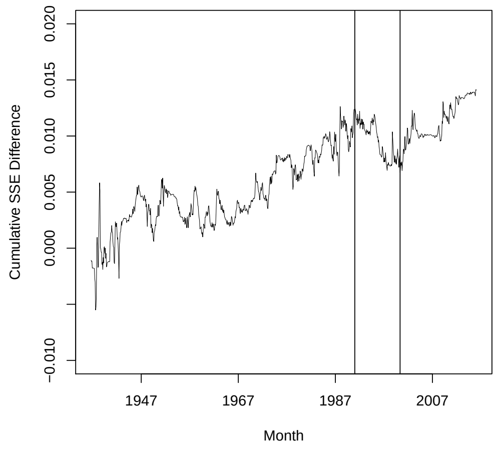

(a) 1936-2015

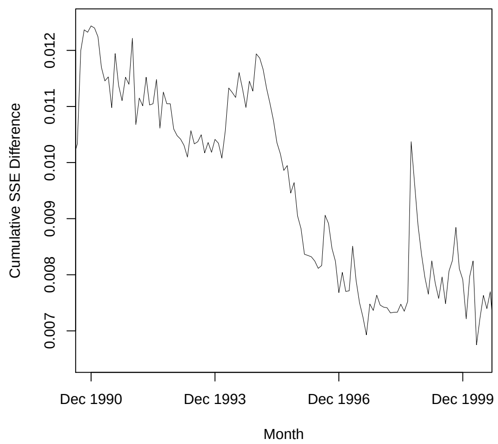

(b) 1990-2000

Figure 7: The top panel of this figure graphs the cumulative sum of squared forecast error differences from July 1936 through December 2015 between those generated by the prevailing mean model and our panel breakpoint model for the market portfolio. The breakpoint model forecasts are a value-weighted average of individual stocks using the dividend-price ratio as the predictive variable. The vertical lines mark the main period of underperformance: December 1990 - April 2000. The lower panel zooms in on the period from 1990 through 2000.

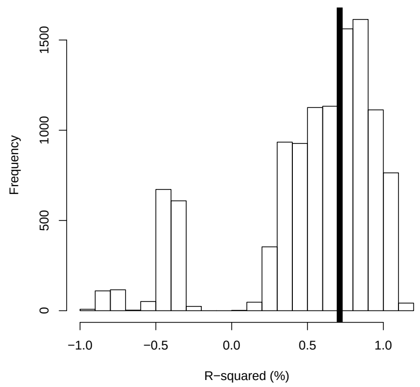

(a) Prevailing Mean

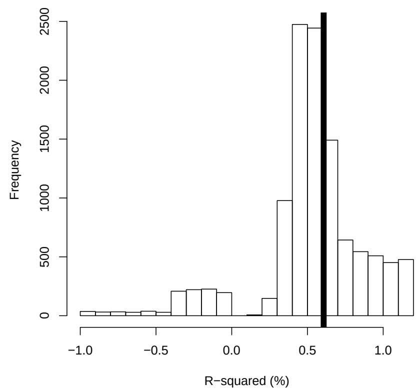

(b) No Break Panel

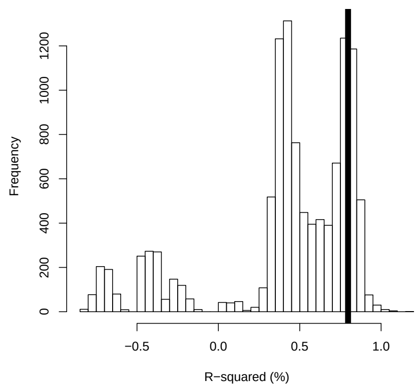

(c) Time series Breakpoint

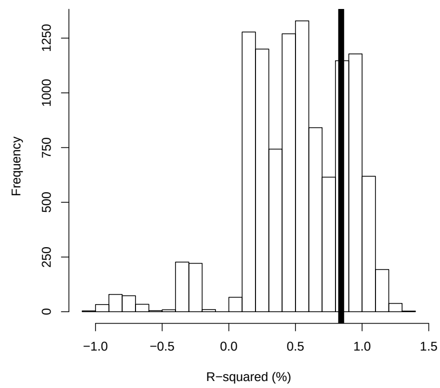

(d) TVP

Figure 8: This figure displays the out-of-sample R^2 values (expressed as percentages) obtained when comparing the forecasting performance of our heterogeneous panel break model with the prevailing mean model (top left window), no break panel (top right), time series breakpoint model (bottom left), and the time-varying parameter model (bottom right) for the firm-level stock returns. The thick vertical black line denotes the market portfolio that is constructed as the value-weighted average of the firm-level forecasts. The aggregate dividend-price ratio is the predictive variable.

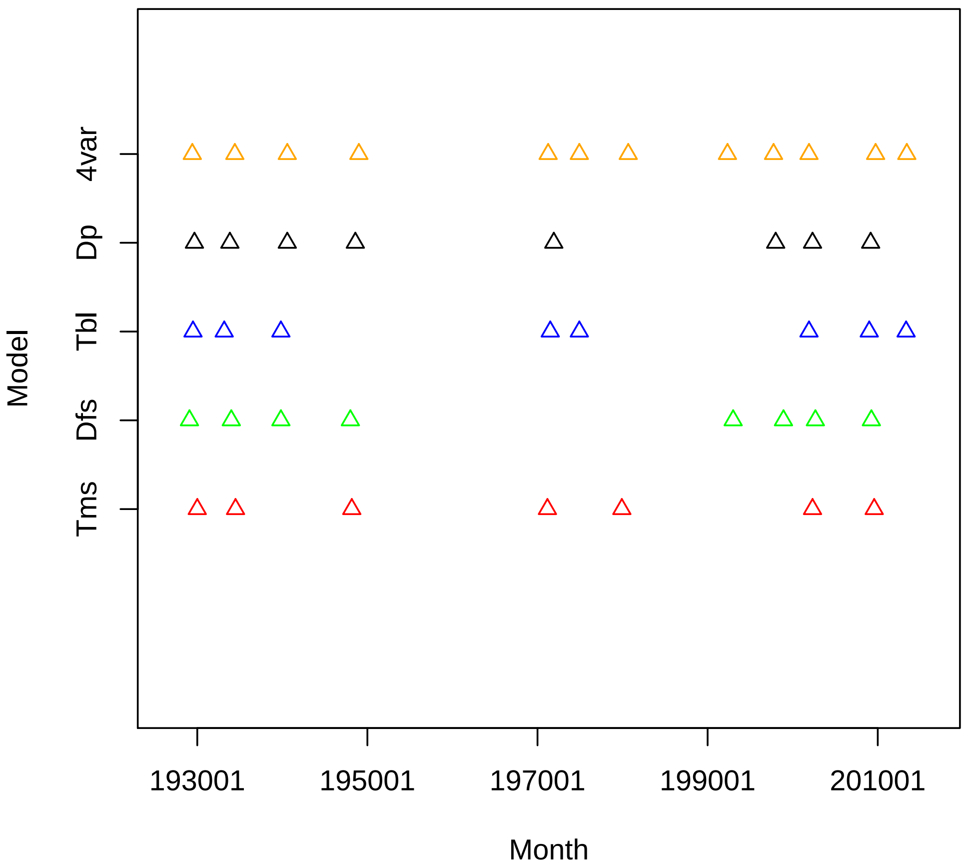

Figure 9: This figure displays the posterior mode break dates estimated from our Bayesian breakpoint model for four univariate return prediction models that use either the dividend-price ratio (black triangles), T-bill rate (blue), default spread (green), or term spread (red) as the predictive variable. We also plot the break modes from the multivariate breakpoint return prediction model that uses all four of these variables (orange).