Breaking Bad: Parameter Uncertainty Caused by Structural Breaks in Stocks*

Lukas Salcher†, Sebastian Stöckl‡

This version: December, 2022

Abstract

Estimating parameter inputs for portfolio optimization has been shown to be notoriously difficult and gets further complicated by structural breaks and regime shifts in financial data. We argue that these structural breaks ultimately result in parameter uncertainty, to which investors are averse. On an aggregate market level, this ambiguity-aversion gives rise to a premium for parameter uncertainty as stocks with high (low) parameter uncertainty are avoided/sold (more attractive/bought). We propose a novel measure called break-(adjusted stock-) age that proxies for parameter uncertainty and is based on detecting structural breaks in stock returns using unsupervised machine learning techniques. Our measure reveals (i) that break-age is priced significantly in the cross-section of stock returns and (ii) that break-age is a powerful proxy for parameter uncertainty.

Keywords: Stock-age, structural breaks, parameter uncertainty

JEL classification: G12, G14, G17

*We would like to thank Michael Hanke for very helpful comments and inputs. Furthermore, we are grateful to all participants of the 36^{th} workshop of the Austrian Working Group on Banking and Finance. Finally, parts of this paper have been inspired by a Master thesis supervised by Sebastian Stöckl.

†University of Liechtenstein, Institute for Finance, Fürst-Franz-Josef-Strasse, 9490 Vaduz, Liechtenstein; email: lukas.salcherl@uni.li; tel: +423 265 1348 (Corresponding author).

‡University of Liechtenstein, Institute for Finance, Fürst-Franz-Josef-Strasse, 9490 Vaduz, Liechtenstein; email: sebastian.stoeckl@uni.li; tel: +423 265 1153.

1 Introduction

Most financial models that deal with portfolio allocation require some form of parameter input like expectations about an asset's future performance and risk (i.e. expected return, variances and covariances) in order to derive optimal portfolio weights. While this task seems straightforward, empirical evidence suggests otherwise. **?** for example, demonstrate that estimating reliable parameter inputs for portfolio optimization requires vast amounts of historical data and is notoriously difficult as the required length of historical observations is non-existent. Consequently, investors are required to estimate future returns, variances and covariances of assets with less data which increases estimation errors. Adding to that, portfolio selection models typically treat parameter inputs with absolute certainty, resulting in well-known issues related to mean-variance optimization such as poor out-of-sample performance, lacklustre diversification as well as unintuitive allocations (**?**).

While there are approaches that try to mitigate or circumvent these issues, such as, for example, **?** or **?**, research shows that financial data is affected by outliers (i.e. **?**). Furthermore, **?** argue that the behaviour of financial markets can change abruptly and frequently which in turn dramatically alters the "mean, volatility and correlation patterns in stock returns" and increases the uncertainty of parameter estimates (p. 1057). Typically, such abrupt changes are caused by changes of technological, legislative or institutional nature as well as economic shocks or policy shifts (**?**). The presence of time-series breaks and regime shifts in financial data further complicates estimating/forecasting parameters and gives rise to two problems: First, using estimates based on long historical samples is no longer appropriate since these estimates are likely biased. Second, estimating input parameters with only few historical observations after a break has been detected is also highly problematic due to estimation errors (**?**). Consequently, it is reasonable to state that financial data is of a time-varying nature and in the presence of breaks, not all historical observations are suitable for estimation and prediction tasks (i.e. **?**).

These findings suggest that parameter uncertainty is inherent to both, long- and short-term estimates and amplified by structural breaks. From the perspective of an investor, these regime shifts are not desirable as they further diminish the reliability of available data and ultimately increase parameter uncertainty. **?** set out to investigate how investors react to parameter uncertainty and, according to **?**, discover that investors allocate less wealth to stocks with high parameter uncertainty. Based on this observation, they further deduce that investors are averse to parameter uncertainty which in turn renders stocks with high parameter uncertainty less attractive for investors relative to stocks with a lower degree of parameter uncertainty. This implies that investors with an aversion to parameter uncertainty (i.e. "ambiguity-aversion") withdraw their capital from stocks with high or increasing parameter uncertainty and re-allocate it to stocks with low or decreasing parameter uncertainty. On an aggregate market level, this ambiguity-aversion gives rise to a premium for parameter uncertainty as stocks with high (low) parameter uncertainty are avoided/sold (more attractive/bought).

Given these observations, we propose a novel measure called *break*-(adjusted stock-) *age* that proxies for parameter uncertainty and is based on detecting structural breaks in stock returns using unsupervised machine learning techniques. Although change point detection models are typically very technical in nature, associating breaks in stock returns with periods of elevated parameter uncertainty seems to be straightforward. This is especially true since breaks can hint at a change in the underlying return generating process which, in the worst case, implies that previous return observations are no longer reliable when it comes to analysing the risk-return dynamics of a stock, thereby increasing the uncertainty with which parameters have to be estimated. Over time, more and more post-break returns can be observed, increasing the reliability of estimated parameters from an investor's perspective. Based on this notion, the objective of our paper is to demonstrate (i) that there is a premium for parameter uncertainty in the cross-section and (ii) that the proposed measure which calculates stock-age on the basis of structural breaks (i.e. break-age) is a suitable proxy for parameter uncertainty as it captures the premium associated with it.

Our results indicate that there is a substantial premium for assets with recent breaks in their time-series regardless of the statistical test used to identify such breaks. Furthermore, we discover that this premium, which we link to parameter uncertainty, is strongest for breaks in variance and for breaks in the mean-variance relationship. Moreover, breaks that have been detected by parametric test statistics that assume Gaussian-distributed returns carry a higher premium than their non-Gaussian counterparts which indicates that investors implicitly assume that returns are normally distributed. Finally, the premium for break-age and, by extension, the premium for parameter uncertainty is more pronounced for smaller stocks. We argue that this is the case because smaller stocks are less well researched and have less media and analyst coverage thereby prolonging the resolution of parameter uncertainty following a structural break.

The remainder of this paper is structured as follows. Section 2 briefly reviews the most important findings related to parameter uncertainty, structural breaks in financial data and firm-age. Thereafter, Section 3 outlines the dataset and methodology we use to investigate whether or not break-age is an appropriate proxy for parameter uncertainty and how this is priced in the cross-section of stock returns. In Section 4 we present and discuss our most important results. Finally, Section 5 concludes and outlines possible future research.

2 Literature Review

Our research is located at the intersection of breakpoint detection models, firm-age and financial performance literature as well as research dedicated to the consequences of parameter uncertainty in financial markets. As breakpoint detection models are typically very technical in nature, most of the significant advances in this field can be found in literature on signal processing, statistical quality control and unsupervised machine learning which is why we skip illuminating developments from that area and rather focus on their importance for the field of Finance. Financial applications of breakpoint detection models have been considered only recently by, for example, **?** who document structural breaks in the equity risk premium ranging from four to six percent. Moreover, **?** detect switching regimes in financial markets and thereby capture financial markets' tendency to abruptly change their behaviour. Furthermore, they document a link between regime switches and periods with different regulations and economic policies. Finally, **?** as well as **?** investigate the role of structural breaks in the cross-section when it comes to forecasting and the detection of market-wide breaks.

Moving on from the literature on structural breaks in the equity premium, research on the link between firm-age and financial performance (i.e. stock returns) is scarce and mainly focuses on performance in the context of IPOs. Here, the most significant findings indicate that stocks tend to outperform on their day of initial public offering (IPO) due to underpricing, followed by a long-run underperformance starting on the post-IPO day and lasting up to 60 months thereafter (i.e. **??**). Furthermore, the findings of **?** indicate that mature firms (i.e. companies aged 12 to 35 years) tend to outperform both, young and old firms (i.e. companies that are younger than 12 years or older than 35 years respectively) in terms of stock returns. Furthermore, **?** investigate the role of ambiguity in the context of IPOs and find that, on a theoretical basis, the IPO puzzle can be explained if the underwriter prices systematic and idiosyncratic ambiguity while diversified investors only price systematic ambiguity. More recently, **?** proposed a novel measure for parameter uncertainty based on financial turbulence as defined by **?** and discovers that parameter uncertainty can be used to predict aggregate stock returns. Apart from that, there is no research that links stock-age and parameter uncertainty despite the recent increase in academic interest (**?**). Furthermore, stock- or firm-age related research and investment strategies typically focus on IPOs as starting points and do not consider structural breaks. Finally, the literature does – so far - not relate parameter uncertainty to structural breaks in time-series (relationships).

Parameter uncertainty and the consequences thereof have also been subject to research. Starting with **?** who first distinguished between uncertainty and risk in an economic framework by defining a situation with risk as one in which outcomes are unknown but following a known probability distribution. In contrast to that, an uncertain situation is characterized by unknown outcomes and unknown probability distributions (i.e. so-called "Knightian Uncertainty"; **?**). Furthermore, since the formulation of the so-called "Ellsberg Paradox" it is known that agents prefer risk (i.e. "probabilized" situation) over uncertainty (i.e. "non-probabilized" situation; **?**).

More recently, **?** set out to investigate how investors react to parameter uncertainty and thereby provided support for a theory of "ambiguity-aversion". **?** later pick up on the notion of parameter uncertainty and describe its implications on an aggregate market level. Apart from that, **?** examine the consequences of uncertainty for financial markets and argue that the trading behaviour of investors changes as trading happens less frequently and there are larger no-trade intervals. Furthermore, **?** find that in an attempt to avoid uncertainty in stock markets, traders limit their stock market participation. In addition to that, **?** discover that the level of various risk premia is influenced by uncertainty. More precisely, they argue that risk premia tend to increase alongside the number of investors that are averse to uncertainty. Several researchers have explored how the presence and level of uncertainty influences the risk-taking preferences of investors. According to **?**, **?**, as well as **?** allocations to risky assets decrease with increasing levels of parameter uncertainty representing an increasing preference for safe assets. Finally, **?** observe that investors avoid unfamiliar securities especially in periods with elevated levels of parameter uncertainty.

3 Data & Methodology

In the following section we describe the dataset and methodology used to investigate how well break-age proxies for parameter uncertainty and whether or not there is a significant premium in the cross-section for stocks with recent breaks. Therefore, the first subsection briefly outlines the dataset underlying our examination. After that, subsection two describes the actual methodology we use in more detail.

3.1 Data

To evaluate how well break-age proxies for parameter uncertainty and how breaks are priced in the cross-section of stock returns, we use the following datasets: We focus on the US-stock market and use monthly delisting-adjusted stock returns from CRSP beginning in 1925. Based on this dataset we have 33'460 individual PERMNO's resulting in a total of 4.4 million observations. Furthermore, we download factor returns and risk-free rates from Kenneth French's website[1].

3.2 Methodology

Before digging deeper into which breakpoint detection methods we use, we first want to clarify some terminology related to breakpoint detection in general. According to **?**, breakpoint detection methods are defined as machine learning techniques that aim at identifying sudden changes in a given time-series $(Y_t)_{t=1,...}$ as represented by the dataset $y_{1:T} := \{y_1, ...y_T\}$. In this context, a breakpoint at time τ is identified "if the statistical properties of the sub-sequences $\{y_1, y_2, ..., y_\tau\}$ and $\{y_{\tau+1}, y_{\tau+2}, ..., y_T\}$ differ" (**?**, p. 9). In the case of multiple breakpoints, a partition $\rho_{1:T}$ of the set $\{1, 2, ..., T\}$ needs to be identified such that the time-series' statistical properties within each of the sub-sequences remain the same while they change from one sub-sequence to another (**?**).

Apart from defining what we mean with structural break and breakpoints in general, it is also necessary to further specify the amount of data we use for detecting breakpoints. In general, there are two possible methods regarding the timing of detecting breakpoints: First of all, there are so-called offline detection methods that use the entire available dataset $y_{1:T}$ to check for potential breakpoints. Alternatively, this type of method is also known as "batch change detection" or "phase I change detection". From a Finance perspective, such break detection methods are similar to in-sample test. Although we use a setting that is closer to out-of-sample tests for our actual investigation, a sound understanding of how offline methods work is essential for our paper. Thus, we further explore and illustrate the mechanics of offline detection methods using the following example: Assume that the exemplary time-series under investigation has at most one breakpoint. If there are no breakpoints, our time-series is independently and identically distributed according to F_0. In contrast to that, if our time-series does indeed have a breakpoint at any time τ, then the series follows the distribution F_0 prior to τ and thereafter follows the distribution F_1. Note that $F_0 \neq F_1$ (**?**). If we now want to test whether or not a break occurs at time τ, we can do so by testing the following hypotheses:

$$\begin{aligned} H_0 &: Y_i \sim F_0(y;\theta_0), \quad i = 1, ..., T \\ H_1 &: Y_i \sim \begin{cases} F_0(y;\theta_0) & i = 1, 2, ..., \tau \\ F_1(y;\theta_1) & i = \tau+1, \tau+2, ..., T \end{cases} \end{aligned} \tag{1}$$

[1]See https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

where θ_i refers to the parameters of the distribution that might be unknown. According to **?** such a problem can easily be solved with a two-sample hypothesis test. Of course, which test statistic should be used is determined by the assumptions that are made with respect to the distribution of the time-series and which parameter is expected to experience breaks. In the case of Gaussian distributed data and expected breaks in the mean a suitable test statistic is given by the Student-t test. Once an appropriate test statistic $D_{\tau,T}$ has been chosen, we can determine its value and, given that $D_{\tau,T}$ exceeds a threshold value $h_{\tau,T}$, we can reject the null hypothesis of identical distributions and therefore conclude that a breakpoint has occurred at time τ (**?**).

In reality however, we cannot know in advance when a breakpoint occurs (i.e. τ is unknown to us). Hence, we have to compute $D_{\tau,T}$ for every point in time between $1 < \tau \leq T$ and chose the maximum value from there. This basically means that our time-series is split up into every possible combination of two sub-sequences and the two-sample test statistic is calculated for every possible splitting point. Following **?** the test statistic then becomes:

$$D_T = \max_{\tau=2,...T-1} D_{\tau,T} = \max_{\tau=2,...,T-1} \left| \frac{\tilde{D}_{\tau,T} - \mu_{\tilde{D}_{\tau,T}}}{\sigma_{\tilde{D}_{\tau,T}}} \right| \tag{2}$$

Consequently, we can reject the null hypothesis of no breakpoints if $D_T > h_T$ for an appropriately chosen threshold h_T that is determined by an acceptable probability for Type I errors α (i.e. the probability of falsely detecting a breakpoint although there is none). In terms of locating a breakpoint, the best estimate is given by τ, i.e. the value that maximizes D_T (**?**):

$$\hat{\tau} = \arg\max_{\tau} D_{\tau,T}. \tag{3}$$

The second type of detection method is referred to as online detection methods, "sequential change detection" or "phase II detection methods". In contrast to offline methods, online methods try to detect breakpoints at time τ, using only the data $y_{1:t}$ that is available up to t, where t is $0 < \tau \leq t \leq T$ (**?**). Drawing parallels to Finance terminology, online detection methods correspond to out-of-sample tests. So whenever a novel observation y_t is added to the time-series, D_T is calculated again using the aforementioned offline methodology. In line with offline methods, a breakpoint is detected if $D_T > h_T$ which implies rejecting the null hypothesis of no breakpoints as long as the threshold value is chosen appropriately. If the threshold is not exceeded, another observation y_{T+1} is added and the procedure is repeated.

According to **?**, the online setting allows the threshold value h_T to be chosen such that the probability of false positives (i.e. Type I error α) remains constant over time. Under the null hypothesis this implies that:

$$\begin{aligned} P(D_1 > h_1) &= \alpha \\ P(D_T > h_T | D_{T-1} \leq h_{T-1}, ..., D_1 \leq h_1) &= \alpha, \quad T > 1 \end{aligned} \tag{4}$$

After briefly describing the logic behind breakpoint detection models, we use the remainder of this section to further outline the methodology used to illustrate how break-age (i.e. the time since the last breakpoint has been detected) captures parameter uncertainty and how increased levels of parameter uncertainty are priced in the cross-section of stock returns. We therefore opted for a case-study-like setting. The reasoning behind this choice is that our measure for parameter uncertainty is dependent on long time-horizons as new (post-break) observations are only added on a monthly basis. While such an approach is perfectly suited for capturing how parameter uncertainty is slowly resolved over time, it is less suitable for building the typical long-short portfolios that are used to illustrate how statistically and economically significant a newly discovered premium is. Furthermore, building long-short portfolios first requires building quantile portfolios of equal size. Breakpoints are however highly individual and do not allow for such an equal allocation to different portfolios which further supports our choice. Thus, as a first step, we compute abnormal returns using the Fama-French-Carhart four-factor model based on 12-month rolling regressions:

$$\begin{aligned} AR_{i,t} = & R_{i,t} - \beta_{MKT,i,t} \cdot R_{MKT,t} - \beta_{SMB,i,t} \cdot R_{SMB,t} - \beta_{HML,i,t} \cdot R_{HML,t} \\ & - \beta_{MOM,i,t} \cdot R_{MOM,t} \end{aligned} \tag{5}$$

where $R_{i,t}$ is the return of stock i at time t, $R_{MKT,t}$, $R_{SMB,t}$, $R_{HML,t}$, $R_{MOM,t}$ denote market and factor-mimicking returns at time t and the various $\beta_{i,t}$ refer to the regression coefficients obtained from regressing these factor returns onto individual stock returns.

These abnormal returns from (5) are then used to check, whether structural breaks in the mean and/or variance can be detected. For this purpose, we use multiple different online breakpoint detection methods. More precisely, we apply the CPM framework of **?** which provides standard (non-)parametric tests for single and multiple breakpoint detection. In total, we consider five separate parametric tests: three are based on the assumption of Gaussian distributed returns and two are tests that do not require this assumption. In detail, the tests we consider are the following:

- Gaussian:
  - Student-t test statistic for break in mean returns

  - Bartlett test statistic for break in variance of returns
  - Generalized likelihood ratio test statistic for breaks in mean and variance of returns
- Non-Gaussian:
  - Mann-Whitney test statistic for break in location of a series
  - Mood test statistic for break in scale of a series

All of these different tests are designed to detect single and multiple breakpoints and are implemented using the R-package *cpm*.

Once the breaks in each individual time-series of abnormal returns are detected with the help of the aforementioned tests, we calculate the break-age of each stock and based upon this measure continue with computing equally- and value-weighted cumulative abnormal returns for each month after break detection and benchmark them against returns of IPO stocks following equations (6) and (7):

$$AR_t = \sum_{i=1}^{N} w_{i,t} AR_{i,t} \tag{6}$$

$$CAR_t = \prod_{k=0}^{t} (1 + (AR_t^{BP} - AR_t^{IPO})) - 1 \tag{7}$$

Our choice of IPO stocks as a benchmark for our stocks with recent breaks is motivated by the fact that the IPO date of stock is typically used when referring to the age of a stock and therefore serves as a natural alternative to our break-age. Furthermore, and in contrast to stocks with recent breaks, IPO stocks are very well researched upon their IPO and are therefore less exposed to parameter uncertainty than conventional stocks.

4 Results

In the following we present and discuss our results. Thus, the first sub-section explores breakpoint statistics. Thereafter, sub-section two dives into the break-age premium we discovered based on equally-weighted cumulative abnormal returns while the results presented in sub-section three also focus on the break-age premium but are based on value-weighted cumulative abnormal returns.

4.1 Breakpoint statistics

Table 1 displays breakpoint detection statistics. In particular, we present the percentage of stocks for which we have detected at least one breakpoint alongside the median number of breakpoints per stocks, the median time to detect a breakpoint in months as well as the median time between breakpoints also measured in months.

[PLACE TABLE 1 HERE.]

First of all, we focus on the percentage of stocks for which the methods we use has identified at least one breakpoint. Depending on the used test statistic approximately 30% to 73% of the time-series under investigation have experienced a break in mean returns or variance. Consequently, we can conclude that a fairly large number of stocks experiences at least one breakpoint which further implies that our results are not driven by extremely rare events. In addition to that, we notice that the test statistics which assume that observations are Gaussian-distributed discover breaks for more stocks compared to their non-Gaussian counterparts. If we compare the frequency of breaks in mean with the frequency of breaks in variance we can also state that stocks more often experience breaks in variance compared to breaks in mean (i.e. percentages for Mann-Whitney and Student-t test statistics are smaller than percentages for Mood and Bartlett). Moving on to the second column of Table 1, we can see that the median number of breakpoints per stock is two respectively three. Again this suggests that the return and risk dynamics of stocks abruptly change much more frequently than assumed. Finally, the last two columns of Table 1 highlight the timing of breakpoints. According to our results, the median time to detect a breakpoint in months ranges from 24 months for the Bartlett test statistic up to 36 months for the test statistic of Mood. Interestingly, the median time between breakpoints also denoted in months is just slightly shorter than the time to detect a breakpoint, regardless of which test statistic is being used. This measure ranges from 20 months to 33 months. Based on this observation, we conclude that, in some cases a breakpoint might be detected only after a second breakpoint is already in the data (but not yet detected). All in all, we want to stress once more that our breakpoint statistics suggest that breakpoints occur for a broad range of stocks, more frequently than assumed and are more often related to breaks in variance.

4.2 Equally-weighted break-age premium

Our main results are summarized in Figure 1, which depicts cumulative abnormal returns benchmarked against the returns of IPO stocks over the full sample period. Furthermore, breakpoints are only included here if their break-date was larger than 1, meaning that the first month after a IPO has not been considered to be a break-point.

[PLACE FIGURE 1 HERE.]

At a first glance, we see that regardless of which test statistic is being used, cumulative abnormal returns are positive and increasing from the detection month up until approximately 25 months. Thereafter, the cumulative abnormal returns begin to decrease for the next 50 to 75 months. We take this observation as a confirmation of our hypothesis that break-age is a suitable proxy for parameter uncertainty and that parameter uncertainty is priced in the cross-section of stock returns. Furthermore, the observed pattern confirms that the expected abnormal returns associated with parameter uncertainty slowly build up over time as uncertainty resolves. From an economic perspective, we interpret the depicted hump shape as evidence that investors can realize abnormal returns only slowly after a breakpoint has been detected as parameter uncertainty unwinds step-by-step with the release of new information.

Next, we want to focus on the differences between breaks in mean returns and breaks in variance. If we therefore focus on the Gaussian test statistics, we can see that changes in the variance of a time-series come with a higher premium than changes in the mean of a time-series (i.e. the CARs for Bartlett are higher than the CARs for Student-t). The same observation can be made for non-Gaussian test statistics as the CARs for Mood (i.e. break in scale) are also higher than the CARs for Mann-Whitney (i.e. break in location). Based on these observations, we argue that breaks in the variance of a variable increase parameter uncertainty more strongly compared to breaks in the mean of a variable and therefore are associated with a higher premium.

Now, we focus on the differences between Gaussian and non-Gaussian test statistics. Therefore, we directly compare the CARs of the Gaussian test statistics (i.e. Bartlett and Student-t) with their non-Gaussian counterparts (i.e. Mood and Mann-Whitney). As can be seen in Figure 1, the CARs of breaks detected by the Bartlett test statistics come with a higher premium than the breaks detected by the Mood test statistics. The same can be said if we compare the CARs of breaks detected by the Student-t test statistics with the CARs based on Mann-Whitney test statistics. Admittedly, it seems easier to detect breaks if assumptions are made about the distribution of a variable. Nonetheless, on the basis of this observation we argue that investors either care for normally distributed returns or simply assume that returns are normally distributed and therefore require higher compensation for increasing uncertainty due to abrupt breaks in the first two distributional moments.

Finally, we investigate how changes in the mean and variance of a time-series are priced. For this purpose, we focus on the CARs of the generalized likelihood ratio model (i.e. the blue line in Figure 1). The pattern of the GLR test statistic is very similar to that of the Bartlett test statistic. According to our interpretation, breakpoints detected by the GLR test statistic basically indicate breaks in the Sharpe ratio of returns. Given this interpretation, it is not surprising that GLR-breaks are associated with a very high premium as investors are certainly concerned about changes in the risk-return structure of their investments.

[PLACE FIGURE 2 HERE.]

Moving on, Figure 2 also depicts cumulative abnormal benchmarked returns. This time however, the CARs are put into perspective and compared with alternative CARs. More precisely, the dashed lines are CARs that include IPOs for breakpoint detection. As already mentioned, IPO stocks are typically very well researched which is part of the whole procedure leading up to the IPO and therefore are exposed to less parameter uncertainty compared to stocks that recently experienced a break. Comparing the dashed lines with the solid lines (i.e. the base case which does not include IPO stocks) certainly confirms this notion as the dashed CARs are considerably lower than the solid CARs. Furthermore, the dotted lines represent CARs that start immediately after the break date and therefore ignore the time it takes to detect that a break has occurred. Again these lines are considerably lower than their base-case counterparts which we attribute to the fact that a change in return and/or risk dynamic takes time to manifest itself. Therefore, immediately after the break date, investors cannot know that a break occurred which is why CARs stay muted. Additionally, the dashed and dotted lines further indicate that breakpoints in the mean detected by Gaussian test statistics gain importance. Finally, the dashed and dotted CARs are flatter compared to the solid lines and start decreasing earlier.

[PLACE FIGURE 3 HERE.]

[PLACE FIGURE 4 HERE.]

Figure 3 and Figure 4 depcit the same results as Figure 1 and Figure 2 but based on a shorter time-period starting in January 1980. From our perspective, the interpretation of these results remains mostly unchanged. Again, breaks in variance come with a higher premium than breaks in mean returns and breaks based on test statistics that assume Gaussian-distributed returns again deliver higher CARs than their non-Gaussian counterparts. Overall, the main differences to the results based on the full sample are that the premium for the short sample seems to be higher and deteriorate less quickly.

4.3 Value-weighted break-age premium

In addition to discussing the results for the equally-weighted break-age premium, we now consider the break-age premium on the basis of value-weighted cumulative abnormal returns. For this purpose Figure 5 depicts the same results as Figure 2 but for value-weighted CARs. What is striking, is the fact that overall, the CARs in Figure 5 are much lower compared to their equally-weighted counterparts. Our interpretation is that of course size is a major contributor when it comes to parameter uncertainty. The reason behind that is the fact that the stocks of smaller (and probably less known) companies are less well researched and less prominent in indices resulting in an overall higher level of parameter uncertainty compared to large-cap stocks that often end up dominating indices as well as media and analyst coverage.

[PLACE FIGURE 5 HERE.]

By and large, the patterns discovered for the importance of breaks in variance versus breaks in means as well as the importance of Gaussian versus non-Gaussian test statistics remain unchanged. The premium for GLR-based breaks, essentially representing breaks in Share ratios has, however, diminished on a relative basis.

[PLACE FIGURE 6 HERE.]

Finally, as depicted in Figure 6, the value-weighted CARs based on the shorter time-frame only confirm the observations made for the entire sample.

5 Conclusion

Estimating parameter inputs for portfolio optimization has been shown to be notoriously difficult resulting in disappointing out-of-sample performance (**??**). The procedure of estimating parameters is further complicated by breaks and regime shifts in financial data caused by, for example, corporate actions and events such as mergers, acquisitions or earnings announcements (**???**). These abrupt changes in time-series ultimately result in parameter uncertainty, to which investors are averse (i.e. "ambiguity-aversion"; **?**). On an aggregate market level, this ambiguity-aversion gives rise to a premium for parameter uncertainty as stocks with high parameter uncertainty are avoided and sold while stocks with low parameter uncertainty become more attractive and are therefore bought.

Against this backdrop, we propose a novel measure that proxies for parameter uncertainty and is called *break*-(adjusted stock-)*age*. Our measure is based on detecting structural breaks in stock returns using unsupervised machine learning techniques. Despite the sophisticated models required to detect such breaks, we argue that associating breaks in stock returns with periods of elevated parameter uncertainty is very plausible from an economic perspective. Based on this notion, our paper demonstrates (i) that there is a significant premium for parameter uncertainty in the cross-section and (ii) that the proposed measure which calculates stock-age on the basis of structural breaks is a suitable proxy for parameter uncertainty. More precisely, we discover that there is a substantial premium for assets with recent breaks in their time-series regardless of the test statistic used to identify breakpoints. Furthermore, this premium which we link to parameter uncertainty is strongest for breaks in variance and for breaks in the mean-variance relationship. Additionally, breaks that have been detected by parametric test statistics that assume Gaussian-distributed returns carry a higher premium than their non-Gaussian counterparts which indicates that investors implicitly assume that returns are normally distributed. Finally, the premium for break-age and, by extension, the premium for parameter uncertainty is more pronounced for smaller stocks. We argue that this is the case because smaller stocks are less well researched and have less media and analyst coverage thereby prolonging the resolution of uncertainty following a breakpoint. Altering the time-frame does not substantially change our results.

Going forward, we suggest focusing on alternative breakpoint detection models that are also capable of detecting breakpoints in regression coefficients. Such models would increase our understanding of what specifically drives parameter uncertainty in stock returns. Furthermore, we are currently working on defining a more suitable benchmark that considers for example industry affiliation and the sensitivity to certain risk premia instead of using stock-age as the only matching variable. Additionally, distinguishing between market-wide and stock-specific breaks could shed more light on which events actually create uncertainty for stocks. Moreover, the role of market capitalization in connection to parameter uncertainty needs further investigation. Finally, we would like to confirm our predictions related to parameter uncertainty and the trading behaviour of investors by considering how trading volume changes following the detection of a breakpoint.

A Appendix

Figures

Figure 1: Cumulative abnormal benchmarked returns (vs. IPOs), equally-weighted, full sample, only breakpoints (break-date >1).

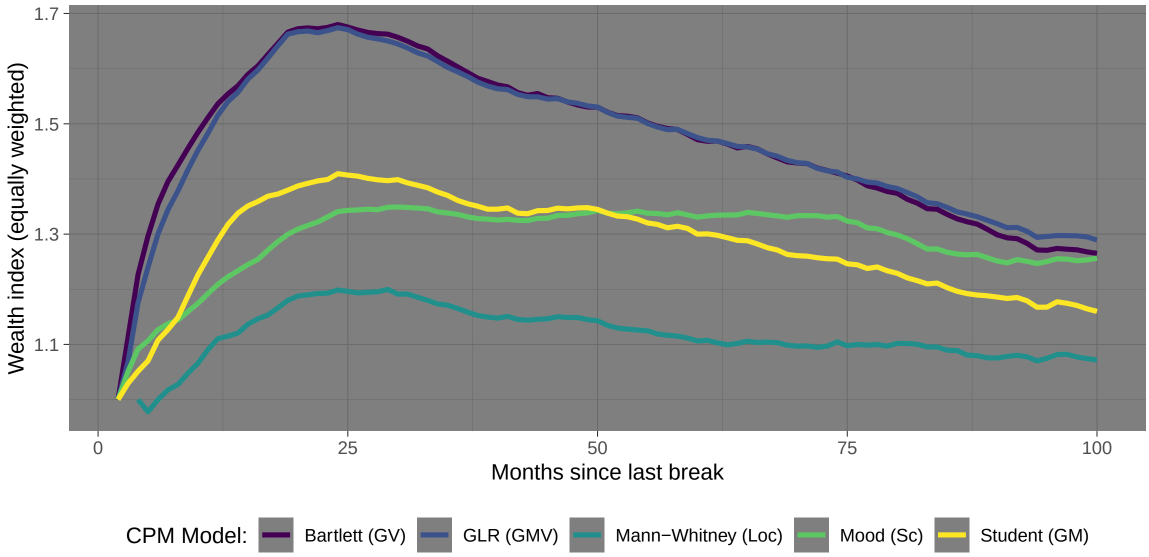

Figure 2: Cumulative abnormal benchmarked returns (vs. IPOs), equally-weighted, full sample (solid line: only breakpoints, dashed line: including IPOs, dotted line: including breakpoints as of break date).

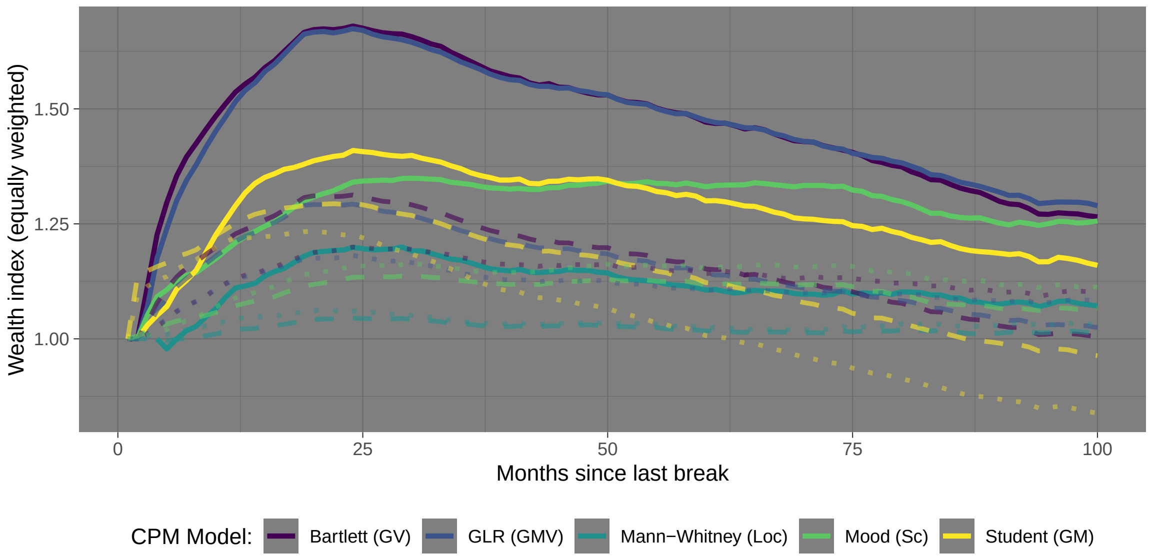

Figure 3: Cumulative abnormal benchmarked returns (vs. IPOs), equally-weighted, short sample, only breakpoints (break-date >1).

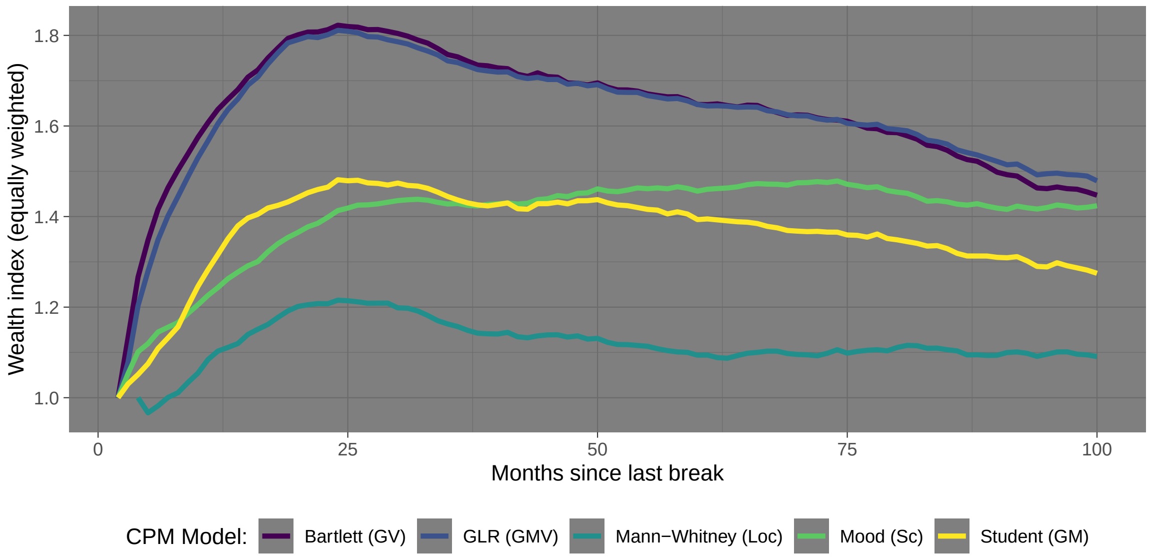

Figure 4: Cumulative abnormal benchmarked returns (vs. IPOs), equally-weighted, short sample (solid line: only breakpoints, dashed line: including IPOs, dotted line: including breakpoints as of break date).

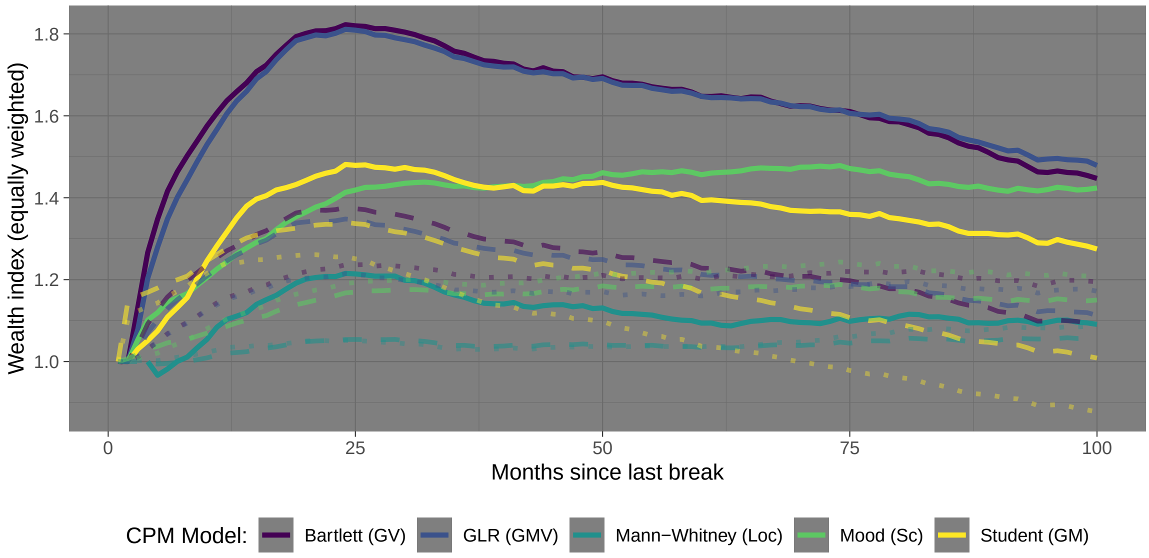

Figure 5: Cumulative abnormal benchmarked returns (vs. IPOs), value-weighted, full sample (solid line: only breakpoints, dashed line: including IPOs, dotted line: including breakpoints as of break date).

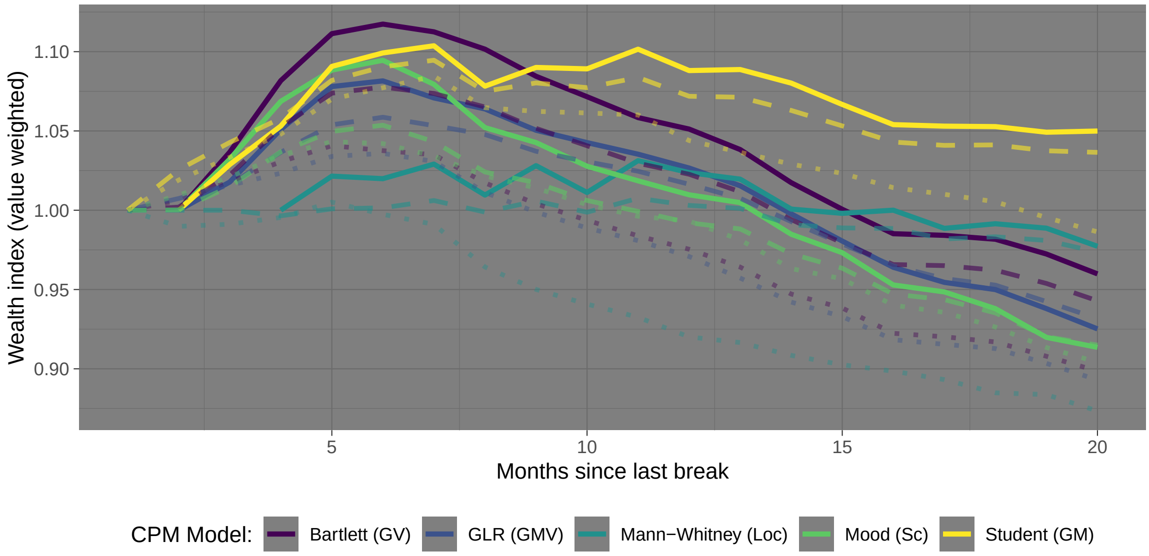

Figure 6: Cumulative abnormal benchmarked returns (vs. IPOs), value-weighted, short sample (solid line: only breakpoints, dashed line: including IPOs, dotted line: including breakpoints as of break date).

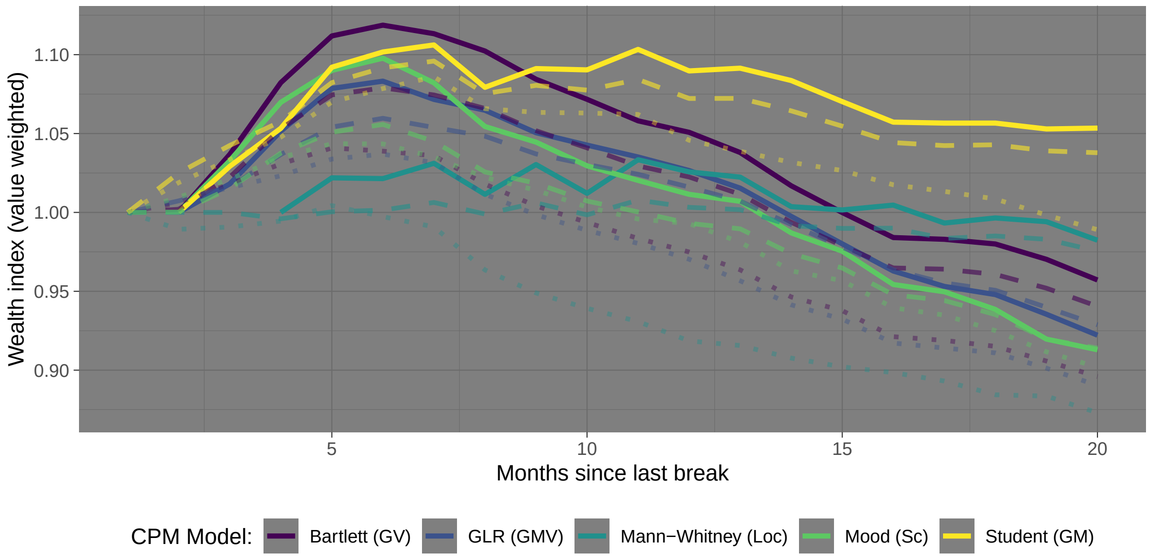

Tables

	% of stocks with BPs	Median no of BPs per stock	Median time to detect BP	Median time between BPs
Mann-Whitney	30.28	2.00	33.00	32.00
Mood	53.32	2.00	36.00	33.00
Student-t	57.80	2.00	34.00	32.00
Bartlett	73.02	3.00	24.00	20.00
GLR	68.30	3.00	27.00	21.00

Table 1: breakpoint detection statistics depicting percentage of stocks with detected breakpoints, the median number of breakpoints per stock, the median breakpoint detection time as well as the median time between breakpoints.