Essays on Idiosyncratic Risk and Return Predictability

by

Daniel Mantilla-García

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

[EDHEC Business School]

November 2011

Dissertation Committee

Raman Uppal, PhD

Ravi Bansal, PhD

René Garcia, PhD

EDHEC BUSINESS SCHOOL

APPROVAL FORM

CANDIDATE'S NAME: Daniel MANTILLA-GARCIA

TITLE OF DISSERTATION: Essays in Idiosyncratic Risk and Return Predictability

APPROVAL:

Raman UPPAL Raman Uppal *Chair*

Ravi BANSAL Ravi Bansal *External Examiner*

René GARCIA René Garcia *Member*

DATE: October 21st, 2011

Abstract

The first paper provides formal arguments and empirical evidence that justifies the use of the cross-sectional variance as a measure of average idiosyncratic volatility. The observability at any frequency of this measure allows new results on the relation of idiosyncratic risk and future returns. The paper shows that the cross-sectional variance predicts the return of the equally-weighted market portfolio over short horizons and that the predictability power of idiosyncratic risk is further increased when adding a measure of cross-sectional skewness to the cross-sectional variance factor in the predictive regressions. Finally, it provides evidence that average idiosyncratic volatility is a positively rewarded risk factor.

The second paper proposes a method to estimate the structural breaks in the mean of the dividend-price ratio. This bayesian technique incorporates the uncertainty about the location and magnitude of the breaks and yields the current-regime mean of this classic stock return's predictor. Adjusting the dividend-price ratio by its current regime mean, improves the explanatory power of the dividend-price ratio of future returns in-sample, as well as its out-of-sample forecasting ability to a very significant extent.

The third paper decomposes the growth rate of the standard portfolio insurance strategy and unveils the (perhaps) surprising role that the correlation between the underlying assets plays on the performance of this type of investment strategy. The paper also introduces the growth optimal portfolio insurance strategy, which combines the growth-rate maximization objective with the constraint of insuring a fixed proportion of the portfolio, expressed in terms of the value of a given stochastic benchmark. The results suggest that the growth optimal strategy outperforms the equivalent standard parametrization of the strategy over long horizons.

Alla promessa della mia felicitá, Lucia.

Contents

Acknowledgements

I would like to thank my thesis directors, Professors René Garcia and Lionel Martellini for giving me the opportunity to work with them and to be one of the first students to join the PhD program they started to launch almost four years ago. The support and guidance they offered me throughout these years were crucial to complete this work and to learn and think as a researcher. Their clear vision on the development of research helped me to pick and choose the topics of this thesis. Working with such great researchers was an extremely enriching experience. I also thank Professors Ravi Bansal and Raman Uppal for their insightful feedback and for assessing my thesis.

I would also like to thank *generous* Professor Abraham Lioui for the countless moments he spent listening to my ideas and for providing thought-provoking and inspiring ones. I also thank EDHEC's faculty for their overwhelming energy and commitment with the students and the program. Their courses vastly broaden my intellectual horizons and gave me a deeper understanding of several issues and methods in finance.

Special thanks to my old friend and working companion Vijay Vaidyanathan. Working with such a brilliant and extraordinary person was enlightening and fun. Vijay's hard work, drive and creativity were determinants for the completion of our paper and also helped me to overcome other multiple challenges encountered throughout these years. At the end, having close an expert in programming and cinema proved to be extremely useful for completing a PhD in finance!

I also thank Gideon, Samuel, Michelle, Kai, Sanjay, Rachid, Jasmine, Kelvin and all other participants of the PhD, from whom I learned a lot about "the real world" as well as methods to solve stochastic differential equations. I enjoyed sharing the challenging adventure of learning and doing research with such a brilliant and diverse group.

I thank Maud Gauchon for the great job she did on the logistics of the program as well as Vincent Milhau, Patrice Retkowsky and all my other colleagues at EDHEC-Risk. A team of exceptional researchers and hard working people, of which I feel proud to have been part of.

I would like to thank Jean René Giraud, from Koris International, for his support and for encouraging me to work on the subject of the third paper of this thesis. A fascinating topic in which I plan to keep working in the future.

I also thank Professors David Schröder and Jean Dante Tosca, who's attitude towards life made these years of hard work a much funnier experience.

I deeply thank my parents, Alvaro and Maria Clemencia and my brother Juan Manuel, which I felt close all the way through. This is the culmination of the heritage they wanted to leave me: "the best education we can give you". The three of them were of great inspiration and support.

I'm also very thankful with Lino, Elvira, Elena and the Murri family, for their endless kindness, generosity and hospitality.

Finally, I would like to thank "the secret behind the PhD" (Vaidyanathan, 2011). Lucia's enormous support, patience and love kept me up and running every time, no matter how bad I would fall, and she learned to appreciate perhaps the most foreign language she had ever heard: finance.

List of Figures

List of Tables

Chapter 1

Idiosyncratic Risk and the Cross-Section of Stock Returns

René Garcia, Daniel Mantilla-García and Lionel Martellini*

Idiosyncratic volatility has received considerable attention in the recent financial literature. Whether average idiosyncratic volatility has recently risen, whether it is a good predictor for aggregate market returns and whether it has a positive relationship with expected returns in the cross-section are still matters of active debate. We revisit these questions from a novel perspective, by taking the cross-sectional variance of stock returns as a measure of average idiosyncratic variance. Two key advantages of this measure are its model-free nature and its observability at any frequency, which allows us to present new results on the properties of daily idiosyncratic volatility series. Through central limit arguments, we formally show that the cross-sectional dispersion of stock returns can be regarded as a consistent and asymptotically efficient estimator for idiosyncratic volatility. We empirically confirm that the cross-sectional measure provides a very good proxy for average idiosyncratic risk as implied by standard asset pricing models and that it predicts well aggregate returns, especially at the daily frequency. The predictability power of idiosyncratic risk is further increased when adding a measure of cross-sectional skewness to the cross-sectional variance factor. We finally provide evidence that idiosyncratic risk is a positively rewarded risk factor.

Keywords: Idiosyncratic Risk, Cross-sectional Variance, Asset Pricing.

*We thank Michel Crouhy, Robert Engle, Stéphane Gregoir, Harrison Hong, Abraham Lioui, Frederic Smadja, Georges Tauchen, Stéphane Tyc, Vijay Vaidyanathan, Volker Ziemann, as well as seminar and conference participants at EDHEC Business School, the Financial Econometrics Conference (Toulouse), the Financial Econometrics Workshop (SMU), the Conference in Asset and Risk Management in the Aftermath of the Financial Crisis (HEC Lausanne), for useful discussions and comments. This research has benefited from the support of the Chair "Produits Structurés et Produits Dérivés", Fédération Bancaire Française. René Garcia is Professor of Finance at EDHEC Business School and a research fellow at CIREQ and CIRANO (rene.garcia@edhec.edu). Lionel Martellini is Professor of Finance at EDHEC Business School and Scientific Director of EDHEC-Risk Institute (lionel.martellini@edhec.edu). Daniel Mantilla-García is PhD candidate at EDHEC Business School and Head of Research & Development at Koris International (daniel.mantilla@koris-intl.com).

1.1 Introduction

The recent financial literature has paid considerable attention to idiosyncratic volatility. Campbell et al. (2001) and Malkiel and Xu (2002) document that idiosyncratic volatility increased over time, while Brandt et al. (2009) show that this trend completely reversed itself by 2007, falling below pre-1990s levels and suggest that the increase in idiosyncratic volatility through the 1990s was not a time trend but rather an "episodic phenomenon". Bekaert et al. (2008) confirm that there is no trend both for the United States and other developed countries. A second fact about idiosyncratic volatility is also a source of contention. Goyal and Santa-Clara (2003) put forward that idiosyncratic volatility has forecasting power for future excess returns, while Bali et al. (2005) and Wei and Zhang (2005) find that the positive relationship is not robust to the sample chosen. Finally, while some economic theories suggest that idiosyncratic volatility should be positively related to expected returns, Ang et al. (2006) find that stocks with high idiosyncratic volatility have low average returns.

An underlying issue in all these studies is the measurement of idiosyncratic volatility. Campbell et al. (2001) use a value-weighted sum of individual firm idiosyncratic variances, computed as the variances of residuals of differences between individual firm returns and the return of an industry portfolio to which the firm belongs.[1] In addition to this measure, Bekaert et al. (2008) use also the individual firm residuals of a standard Fama and French three-factor model to compute a value-weighted aggregate idiosyncratic volatility.[2]

We revisit the issues regarding the dynamics and forecasting power of idiosyncratic variance by using instead the cross-sectional dispersion of stock returns. Through central limit arguments, we provide the formal conditions under which

[1] This amounts to imposing unit beta restrictions in an industry-market model.
[2] This is also the approach followed in Ang et al. (2006).

the cross-sectional variance (CSV) of stock returns asymptotically converges towards the average idiosyncratic variance.[1] One key advantage of this measure is obviously its observability at any frequency, while the previous approaches have used monthly measures based on time series of daily returns. A second important feature is that this measure is model-free, since we do not need to obtain residuals from a particular model to compute it.

We confirm empirically that the cross-sectional variance is an excellent proxy for the idiosyncratic variance obtained from the CAPM or the Fama-French models, as done in the previous literature. Correlations between the CSV measure and the model-based measures estimated monthly, are always above 99%, whether we consider equally-weighted or capitalization-weighted measures of idiosyncratic variance. We also estimate a regime-switching model for CSV time series at both daily and monthly frequencies and find remarkably coherent results in terms of parameter estimates. If we were to build a daily series of model-based idiosyncratic variance, we will roll a window of one-month of daily data, which will result in a very persistent time series. We construct such a daily series but could not find any regimes. This reinforces the usefulness of the CSV to capture idiosyncratic volatility at high frequency.

The regime-switching model indicates clearly that the CSV is counter-cyclical, the dispersion of returns being high and quite variable when economic growth subsides. We analyze further the relation between CSV and economic and financial variables. In particular, we find that there exists a substantial correlation between the equal-weighted CSV and consumption growth volatility. This is consistent with Tédongap (2010) who provides strong evidence that consumption volatility

[1] Goyal and Santa-Clara (2003) argue informally that their measure can be interpreted as a measure of cross-sectional dispersion of stock returns, but do not establish a formal link between the two. In the practitioners' literature (see DiBartolomeo (2006)), cross-sectional dispersion of returns is called *variety* and is used in risk management and performance analysis.

risk explains a high percentage of the cross-sectional dispersion in average stock returns for the usual set of size and book-to-market portfolios that have been used in tests of asset pricing models. In intertemporal asset pricing models of Bansal and Yaron (2004), Bollerslev et al. (2009) and Bollerslev et al. (2009), consumption growth volatility is a measure of economic uncertainty, which is a priced risk factor that affects returns, therefore providing a rationale for the observed correlation between CSV and consumption growth volatility.

On the debate about predictability of aggregate returns by the idiosyncratic variance, we first verify empirically that the CSV measure leads to the same conclusions that other studies (in particular Goyal and Santa-Clara (2003) and Bali et al. (2005)) have reported at the monthly frequency. Then, we report new results at the daily frequency. Specifically, we show that the predictive power of idiosyncratic volatility is much stronger both quantitatively and statistically at the daily frequency than at the monthly frequency. This relationship is robust to the inclusion of return variance and option-implied variance as additional variables in the predictive regressions.

We find that the relation is much stronger and stable across periods between the equally-weighted measure of aggregate idiosyncratic volatility and the returns on the equally-weighted index than for the market-cap weighted equivalents. Economic sources of heterogeneity between firms, as diverse as they can be, are better reflected in an equally weighted measure, all other things being equal. This argument is consistent with previous findings in Bali et al. (2005), who argue that the relationship between equal-weighted average idiosyncratic risk and the market-cap weighted index on the sample ending in 1999:12 is mostly driven by small stocks traded in the NASDAQ. Of course, when the bubble burst, the market capitalization of dot.com small firms was relatively more affected causing the relationship to break down in 2000 and 2001. This effect is not prevalent in an equally-weighted index, for which the relationship remains strong.

However, the frequency at which predictive regressions are run has an important impact on the results, since at lower frequencies we find little evidence of predictability for the equally-weighted measure of CSV. At quarterly and annual frequencies, we find that the capitalization-weighted measure of CSV is a very strong predictor of the aggregate value-weighted returns. When using CSV^{CW} alone as a predictor we obtain remarkable R^2s of 4% and 26% at quarterly and annual frequencies, respectively. Adding the implied variance brings the R^2s to almost 19% and 29%. In all these predictability regressions, the sign of the CSV^{CW} variable is negative. We relate these results to potential explanations in terms of missing factors, Guo and Savickas (2008), or dispersion of investors' opinions, Cao et al. (2005).

Finally, we unveil an asymmetry in the relationship between idiosyncratic variance and returns and show that the predictive power of specific risk is substantially increased when a cross-sectional measure for idiosyncratic skewness is added as explanatory variable. In fact, this is yet another key advantage of our measure that it lends itself to straightforward extensions to higher-order moments.

The statistical significance of the moments of the cross-sectional distribution in these predictive regressions of future returns is not the same as the cross-sectional pricing of stocks or portfolios. However, as emphasized in Goyal and Santa-Clara (2003), the two pieces of evidence are related. Using a Fama-MacBeth procedure with several sets of portfolios, we find support for a positive and significant price of risk for the exposure to the idiosyncratic variance risk. Theoretical rationalizations of a positive relation between idiosyncratic risk and expected returns can be found in the asset pricing literature. Levy (1978), Merton (1987) and Malkiel and Xu (2002) pricing models relate stock returns to their beta with the market and their beta to market-wide measures of idiosyncratic risk. In these models, an important portion of investors' portfolios may differ from the market. Their holdings may be affected by corporate compensation policies, borrowing constraints, heterogeneous beliefs and include *non-traded* assets that add background risk to their traded portfolio decisions (e.g. human capital and private businesses). These theoretical predictions are also in line with Campbell et al. (2001)'s argument that investors holding a limited number of stocks hoping to approximate a well-diversified portfolio would end up being affected by changes in idiosyncratic volatility just as much as by changes in market volatility. More recently, Guo and Savickas (2008) argue that changes in average idiosyncratic volatility provide a proxy for changes in the investment opportunity set and that this proxy is closely related to the book-to-market factor[1].

Ang et al. (2006) and Ang et al. (2009) find results that are opposite to our findings and to these theories since stocks with high idiosyncratic volatility have low average returns but cannot fully rationalize this result. However, Huang et al. (2009) find that the negative sign in the relationship between idiosyncratic variance and expected returns at the stock level becomes positive after controlling for return reversals. Similarly, Fu (2009) documents that high idiosyncratic volatilities of individual stocks are contemporaneous with high returns, which tend to reverse in the following month.

The rest of the paper is organized as follows. In Section 1.2, we provide a formal argument for choosing the cross-sectional variance of returns as a measure of average idiosyncratic volatility, explore its asymptotic and finite-distance properties, as well as the assumptions behind its use, and compare it to other measures formerly selected in the literature. Section 1.3 provides an empirical implementation

[1] Alternative explanations of the relation between idiosyncratic risk and return are the firm's assets' call-option interpretation by Merton (1974) where equity is a function of total volatility as in Black and Scholes (1973) as well as Barberis et al. (2001) prospect theory asset pricing model with loss aversion over (owned) individual stock's variance.

of the concept, again in comparison with other measures, by studying its time-series behavior, outlining the presence of regimes and a counter-cyclical property. In Section 1.4, we provide new results on the predictability of returns by idiosyncratic volatility, and we also extend the analysis to idiosyncratic skewness. Section 1.5 focuses on the analysis of the cross-sectional relationship between idiosyncratic risk and expected returns. Section 1.6 concludes and a technical appendix collects proofs and more formal derivations.

1.2 The Cross-sectional Variance as a Measure of Idiosyncratic Variance

Let N_t be the total number of stocks in a given universe at day t, and assume with no loss of generality a conditional single factor model for excess stock returns.[1] That is, we assume that for all $i = 1, ..., N_t$, the return on stock i in excess of the risk-free rate can be written as:

$$r_{it} = \beta_{it} F_t + \varepsilon_{it}. \tag{1.1}$$

where F_t is the factor excess return at time t, β_{it} is the beta of stock i at time t, and ε_{it} is the residual, with $E(\varepsilon_{it}) = 0$ and $cov\,(F_t, \varepsilon_{it}) = 0$. We assume that the factor model under consideration is a strict factor model, that is $cov\,(\varepsilon_{it}, \varepsilon_{jt}) = 0$ for $i \neq j$.[2]

[1] Assuming a single factor structure is done for simplicity of exposure only and the results below can easily be extended to a multi-factor setting.

[2] This assumption is made in the single index or diagonal model of Sharpe (1963) and in the derivation of the APT in Ross (1976). It implies that all commonalities are explained by the factor model in place. One should notice that the very definition of idiosyncratic risk relies precisely on the assumption of orthogonal residuals: assuming that the model is the "true" factor model implies that the"true" idiosyncratic risk is the one measured with respect to that model, which in turn implies that no commonalities should be left after controlling for the common factor exposure.

Given T observations of the stock returns and the factor return, one can use the residuals of the regression to obtain a measure of the idiosyncratic variance of asset i by: $\sigma_i^2 = \frac{1}{T}\sum_{t=1}^{T}\varepsilon_{it}^2$. An average measure of idiosyncratic variance over the T observations (say a month) can be obtained by averaging across assets such individual idiosyncratic variance estimates. This is the approach that has been followed by most related papers with observations of the returns at a daily frequency to compute monthly idiosyncratic variances.

We propose instead to measure at each time t the cross-sectional variance of observed stock returns. Using formal central-limit arguments, we show that, under mild simplifying assumptions, this cross-sectional measure provides a very good approximation for average idiosyncratic variance. In contrast with most previous measures of average idiosyncratic variance, the CSV offers two main advantages: it can be computed directly from observed returns, with no need to estimate other parameters such as betas, and it is readily available at any frequency and for any universe of stocks.

1.2.1 Measuring the cross-sectional variance

To see this, first let $(w_t)_{t\geq 0}$ be a given weight vector process. The return on the portfolio defined by the weight vector process (w_t) is denoted by $r_t^{(w_t)}$ and given by:

$$r_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it} r_{it}. \tag{1.2}$$

We restrict our attention to non-trivial weighting schemes, ruling out situations such that the portfolio is composed by a single stock. We also restrict the weights to be positive at every given point in time. Hence, a weighting scheme (w_t) is a vector process which satisfies $0 < w_{it} < 1 \; \forall \; i, t$.

The cross-sectional variance measure is defined as follows.

Definition (***CSV***): The ***cross-sectional variance*** measure under the weighting scheme (w_t), denoted by $\boldsymbol{CSV}_t^{(w_t)}$, is given by

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it} \left(r_{it} - r_t^{(w_t)} \right)^2 . \tag{1.3}$$

A particular case of interest is the ***equally-weighted CSV*** (or EW CSV), denoted by CSV_t^{EW} and corresponding to the weighting scheme $w_{it} = 1/N_t \ \forall \ i, t$:

$$CSV_t^{EW} = \frac{1}{N_t} \sum_{i=1}^{N_t} \left(r_{it} - r_t^{EW} \right)^2 , \tag{1.4}$$

where r_t^{EW} is the return on the equally-weighted portfolio.

Another weighting scheme of interest is the cap weighting scheme. If we denote by c_{it} be the market capitalization of stock i at the beginning of the month corresponding to day t, $C_t = \sum_{i=1}^{N_t} c_{ti}$ the total market capitalization and r_t^{CW} the return on the market capitalization-weighted portfolio, the ***cap-weighted (CW)*** (or CW CSV) is defined as:

$$CSV_t^{CW} = \sum_{i=1}^{N_t} w_{it}^{CW} \left(r_{it} - r_t^{CW} \right)^2 , \tag{1.5}$$

where $w_{it}^{CW} = \sum_{i=1}^{N_t} \frac{c_{it}}{C_t}$.

For any given weighting scheme (in particular EW or CW), the corresponding cross-sectional measure is readily computable at any frequency from observed returns. This stands in contrast with the previous approaches that have used monthly measures based on time series regressions on daily returns. The second important feature of the CSV is its model-free nature, since we do not need to specify a particular factor model to compute it.[1]

[1]While Goyal and Santa-Clara (2003) and Wei and Zhang (2005)consider the equally-weighted CSV in conjunction with other measures, they do not provide a thorough discussion

1.2.2 A Formal Relationship between CSV and Idiosyncratic Variance

The following proposition establishes a formal link between CSV and idiosyncratic variance. It is an asymptotic result ($N_t \to \infty$) obtained under the assumptions of homogeneous betas and residual variances across stocks, i.e. $\beta_{it} = \beta_t = 1 \ \forall \ i$, $E(\varepsilon_{it}^2) = \sigma_\varepsilon^2 (t) \ \forall i$. These assumptions will be relaxed below.

Proposition 1 ***(CSV as a proxy for idiosyncratic variance - asymptotic results)****:*
Assume $\beta_{it} = \beta_t = 1 \ \forall \ i$ (homogeneous beta assumption) and $E(\varepsilon_{it}^2) = \sigma_\varepsilon^2 (t) \ \forall i$ (homogeneous residual variance assumption), then for any strictly positive weighting scheme, we have that:

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it} \left(r_{it} - r_t^{(w_t)} \right)^2 \underset{N_t \to \infty}{\longrightarrow} \sigma_\varepsilon^2 (t) \ \ \textit{almost surely.} \tag{1.6}$$

Proof See Appendix A.1.

This result is important because it draws a formal relationship between the dynamics of the cross-sectional dispersion of realized returns and the dynamics of idiosyncratic variance. Note that this asymptotic result $CSV_t^{(w_t)} \longrightarrow \sigma_\varepsilon^2 (t)$ holds for any weighting scheme that satisfies $0 < w_{it} < 1 \ \forall \ i, t$. Of course, at finite distance, different weighting schemes will generate different proxies for idiosyncratic variance. In the empirical analysis that follows, we shall focus on the equally-weighted scheme, while also considering the cap-weighted scheme for comparison purposes. Formal justification for our focus on the equally-weighted scheme is provided in the next section, where we show that the EW CSV is the

about the conditions under which it can be interpreted as a proxy for idiosyncratic variance nor their empirical validity in the data, as we provide in this paper.

best estimator for idiosyncratic variance within the class of CSV obtained under a strictly positive weighting scheme.

1.2.3 Properties of CSV as an Estimator for Idiosyncratic Variance

First, we derive in Proposition 2 the bias and the variance of the CSV as an estimator of idiosyncratic variance. Then we study their asymptotic limits as the number of firms grows large and conclude that the equally-weighted CSV is the best among all-positively-weighted estimators.

Proposition 2 ***(Bias and variance of CSV)***:
Maintaining the homogenous beta assumption ($\beta_{it} = \beta_t = 1 \ \forall \ i, t$) and the homogeneous residual variance assumption ($E(\varepsilon_{it}^2) = \sigma_\varepsilon^2 \, (t) \ \forall i$), for any strictly positive weighting scheme, we have that:

$$E\left[CSV_t^{(w_t)}\right] = \sigma_\varepsilon^2 \, (t) \left(1 - \sum_{i=1}^{N_t} w_{it}^2\right) \tag{1.7}$$

To analyze the variance of the CSV estimator, we further make the assumption of multi-variate normal residuals $\varepsilon \sim N(0, \Sigma^\varepsilon)$, where Σ^ε denotes the variance covariance matrix of the residuals. Under this additional assumption, we obtain:

$$Var\left[CSV_t^{(w_t)}\right] = 2\sigma_\varepsilon^2 \, (t) \left(\left(\sum_{i=1}^{N_t} w_{it}^2\right)^2 + \sum_{i=1}^{N_t} w_{it}^2 - 2\sum_{i=1}^{N_t} w_{it}^3\right) \tag{1.8}$$

Proof See Appendix A.2 for a proof in the slightly more general case when the homogeneous specific variance assumption has been relaxed.

Hence the CSV is a biased estimator for idiosyncratic variance, with a bias given by the multiplicative factor $\left(1 - \sum_{i=1}^{N_t} w_{it}^2\right)$, which can be easily corrected for since it is available in explicit form. In the end, the bias and variance of the CSV appear to be minimum for the EW scheme, which corresponds to taking $w_{it} = 1/N_t$ at each date t. It is easy to see, that this bias disappears and the variance tends to zero for the equally-weighted scheme when the number of stocks grows large, as explained in the following proposition.

Proposition 3 ***(Properties of the equally-weighted CSV)***

The bias and variance of the EW CSV as an estimator for specific variance disappear in the limit of an increasingly large number of stocks:

$$E\left[CSV_t^{EW}\right] \underset{N_t\to\infty}{\longrightarrow} \sigma_\varepsilon^2\left(t\right).$$
$$Var\left(CSV_t^{EW)}\right) \underset{N_t\to\infty}{\longrightarrow} 0.$$

Proof See Appendix A.2 for a proof in the slightly more general case when the homogeneous specific variance assumption has been relaxed..

The equally-weighted CSV thus appears to be a consistent and asymptotically efficient estimator for idiosyncratic variance. As such, it is the best estimator in the class of CSV estimators defined under any positive weighting scheme, and it dominates in particular the cap-weighted CSV as an estimator for idiosyncratic variance. If we relax the homogeneous residual variance assumption, we obtain that:

$$E\left[CSV_t^{EW}\right] \underset{N_t\to\infty}{\longrightarrow} \frac{1}{N_t}\sum_{i=1}^{N_t}\sigma_{\varepsilon_{it}}^2.$$

Hence, the assumption of homogenous residual variances comes with no loss of generality. In the general case with non-homogenous variances, the CSV simply appears to be an asymptotically unbiased estimator for the *average* idiosyncratic variance of the stocks in the universe. We also have:

$$Var\left(CSV_t^{EW}\right) < 2\bar{\sigma}_{\varepsilon}^4\left(t\right)\left(\frac{1}{N_t}\right) \underset{N_t \to \infty}{\longrightarrow} 0.$$

where the quantity $\bar{\sigma}_{\varepsilon}^2\left(t\right)$ is an upper bound for the individual idiosyncratic variances (see Appendix A.2).

We now discuss the impact on these results of relaxing the homogeneous beta assumption.

1.2.4 Relaxing the Homogeneity Assumption for Factor Loadings

Relaxing the homogenous beta assumption involves a bias that remains strictly positive even for an infinite number of stocks and an equal-weighting scheme. We characterize this bias in the next proposition in order to gauge its magnitude for given models of returns.

Proposition 4 ***Bias of CSV as an estimator for average idiosyncratic variance in the presence of heterogenous betas***: *Relaxing the assumptions $\beta_{it} = \beta_t = 1\ \forall\ i, t$ (homogeneous beta assumption) we have, for any strictly positive weighting scheme:*

$$E\left[CSV_t^{(w_t)}\right] = \sum_{i=1}^{N_t} w_{it}\sigma_{\varepsilon_i}^2\left(t\right) - \sum_{i=1}^{N_t} w_{it}^2\sigma_{\varepsilon_i}^2\left(t\right) + E\left[F_t^2 CSV_t^{\beta}\right], \tag{1.9}$$

where CSV_t^{β} denotes the cross-sectional variance of stock betas:

$$CSV_t^{\beta} = \sum_{i=1}^{N_t} w_{it}\left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right)^2.$$

Proof See Appendix A.2.3.

The first term $\sum_{i=1}^{N_t} w_{it}\sigma_{\varepsilon_i}^2(t)$ in equation (1.9) represents the average idiosyncratic variance of stocks within the universe under consideration. The second term $-\sum_{i=1}^{N_t} w_{it}^2\sigma_{\varepsilon_i}^2(t)$ is the negative bias that was also present even in the presence homogenous beta assumptions. If we focus on the equally-weighted scheme, the sum of these two terms is equal to $\frac{1}{N_t}\sum_{i=1}^{N_t} \sigma_{\varepsilon_i}^2(t)\left(1-\frac{1}{N_t}\right)$ so that the bias disappears in the limit of an increasingly large number of stocks. The third term $E\left[F_t^2 CSV_t^{\beta}\right]$ in equation (1.9) represents, on the other hand, an additional (positive) bias for the CSV as an estimator of average idiosyncratic variance, which is introduced by the cross-sectional dispersion in betas, and which does not disappear in the limit of a large number of stocks.

Using the explicit expression provided here, in section 1.3.1 we directly measure this *beta dispersion bias* using the CAPM and the Fama and French three-factor model as benchmark factor models. As we will see, although the cross-sectional dispersion of betas has a non-negligible magnitude, once it is multiplied by the square of the return of the market portfolio its relative size with respect to the level of idiosyncratic risk becomes very small. An extensive analysis of the CSV in the empirical section suggests that the homogeneous beta assumption does not represent a material problem for the CSV as an estimator of idiosyncratic variance as implied by standard asset pricing models (i.e. CAPM and Fama-French).

1.2.5 Competing Measures of Idiosyncratic Risk

In this section, we describe measures that have been used in the literature, and which will be used for comparison purposes in subsequent sections of the paper. The standard approach consists of considering idiosyncratic variance either relative to the CAPM and or to the Fama-French (FF) model (Fama and French (1993)):

$$r_{it} = b_{0it} + b_{1it}XMKT_t + b_{2it}SMB_t + b_{3it}HML_t + \varepsilon_{it}^{FF} \tag{1.10}$$

where r_{it} denotes the excess return at time t of stock i, $XMKT$ is the excess return on the market portfolio, SMB is the size factor and HML is the value factor. The idiosyncratic variance for asset i is the variance of the residuals of the regression, that is, $\sigma^2(\varepsilon_{it}^{FF})$. To obtain an estimate for average idiosyncratic variance, Bekaert et al. (2008) and Wei and Zhang (2006) use a market capitalization weighting:

$$FF_t^{CW} = \sum_{i=1}^{N_t} w_{it}\sigma^2(\varepsilon_{it}^{FF}). \tag{1.11}$$

For comparison purposes we also look at the equally-weighted average of FF idiosyncratic variance in what follows. An alternative approach to average (*mostly*) idiosyncratic risk estimation has been suggested by Goyal and Santa-Clara (2003), with a measure given by:

$$GS_t^{EW} = \frac{1}{N_t}\sum_{i=1}^{N_t}\left[\sum_{d=1}^{D_t} r_{id}^2 + 2\sum_{d=2}^{D_t} r_{id}r_{id-1}\right], \tag{1.12}$$

where r_{id} is the return on stock i in day d and D_t is the number of trading days in month t.[1]

Campbell et al. (2001) propose yet an alternative measure of average idiosyncratic variance, under a very particular setting that allows one to avoid running regressions each period.[2] However, their measure is not instantaneous since a window of data is still needed to estimate individual variances. In what follows, we do not repeat the analysis with this measure because Bekaert et al. (2008) have shown that it is very closely related to the measure obtained from standard asset pricing models. In particular, Bekaert et al. (2008) find that the measure of Campbell et al. (2001) and the FF-based one have a correlation of 98% and share

[1]As in Goyal and Santa-Clara (2003), when the second term makes the estimate negative, it is ignored. This measure has been originally used in French et al. (1987).

[2]They assume that all betas are equal to one and substract industry returns in addition to market returns to control for risk.

most of the same structural breaks.

1.3 Empirical Implementation

In order to perform an empirical analysis of our measure for idiosyncratic risk, we collect daily US stock returns (common equity shares only) and their market capitalization from CRSP data base. Our longest sample runs from July 1963 to December 2006. We also extract the FF factors and the one-month Treasury bill from Kenneth French web-site data library for the same sample period. Each month, we drop stocks with missing returns and with non-positive market capitalization at the beginning of the month. The number of firms varies between 377 and 7293, and remains greater than one thousand 75% of the time. The maximum number of stocks is reached during the .com bubble. Then, we estimate every month the cap-weighted idiosyncratic variance as in equation (1.11), as well as the equal-weighted version.[1] Similarly, we estimate the cap-weighted and equal weighted average idiosyncratic variance relative to the CAPM. We also estimate the GS average variance measure as in equation (1.12) and its cap-weighted version. Finally, we estimate on a daily basis the equal and cap-weighted versions of the CSV as in equations (1.4) and (1.5). In order to construct the monthly series for our cross-sectional measures, we estimate the average of the daily series at the end of each month. For comparison purposes we also estimate the FF-based average idiosyncratic variance (EW and CW) on a daily basis using a rolling window sample of one month. We annualize all figures in order to compare daily and monthly measures. Following Bekaert et al. (2008), we fit a regime-switching model to the monthly and daily series in order to further compare the different measures. Last, we look at the relation between the CSV measures of idiosyncratic

[1] We use previous period market capitalization and assume it is constant within the month.

variance and selected economic and financial variables.

1.3.1 Measuring the CSV bias

Some of the previous research on idiosyncratic volatility has been conducted under the assumption of homogeneous betas across stocks (see Campbell et al. (2001) and Goyal and Santa-Clara (2003) in particular). As illustrated in Proposition 4 and discussed in Appendix A.2.3, the presence of non-homogeneous betas introduces a positive bias on the CSV as an estimator for average idiosyncratic variance, which is given by the first term in equation (1.9). We now measure the impact of this bias with respect to the CAPM as a benchmark model.

First, we compute the bias $E\left[F_t^2 CSV_t^{\beta}\right]$ for every month in the sample using beta estimates for each stock with both the equal-weighted and the cap-weighted market returns. To gauge its importance, we divide it by the average idiosyncratic variance, also measured with respect to the CAPM.[1]

Table D.1 presents a summary of the distribution of the time series of cross-sectional dispersion of betas, its product with the squared return of the market portfolio (hence the bias itself) and the proportion of this bias with respect to the average idiosyncratic variance at the end of every month. Although the cross-sectional dispersion of betas is sizable, once it is multiplied by the squared return of the market portfolio, the size of the bias remains relatively small. The median of the distribution of $\frac{F_t^2 CSV_t^{\beta}}{\sigma^2_{\varepsilon_t}}$, is 0.348% for the equal-weighted scheme and 0.351% for the cap-weighted measure, computed over the whole sample (July 1963 to December 2006). The 97.5 quantiles are 3.24 and 3.47 respectively.

On the other hand, the formal discussion about the properties of the CSV as a measure of idiosyncratic variance on section 1.2.4 also uncovered the fact

[1] This is measured as in equation (1.11) with just the market returns with both weighting schemes.

that another bias (but negative in sign) coming from the CSV weighting scheme concentration is also introduced. Proposition 2 predicts two properties about this weighting bias: first, it should be negative and minimal for an equally-weighted scheme. Second, it should be very small for a high number of stocks. The beta-bias then is more likely to dominate the concentration-bias when using an equal-weight scheme.

Using the explicit expression for this bias provided in Proposition 4 we estimate the proportion of the size of this weights-concentration bias with respect to the average idiosyncratic variances implied by the CAPM.[1] In the last line of the upper and lower panels of Table D.1 we report quantiles of the distribution of this bias for both weighting schemes. The corresponding medians are 0.030% and 0.426% for the EW and CW schemes respectively. Since the bias is of opposite sign to the beta cross-sectional dispersion bias, we need to assess the resulting overall bias.

We measure the total bias as the intercept of a regression of the CSV on the average idiosyncratic variance estimated with respect to the CAPM or the Fama-French three-factor model:

$$CSV_t^{w_t} = bias + \psi \sigma^2_{model} \left(w_t \right) + \zeta_t, \tag{1.13}$$

where w_t refers to the weighting scheme (equal-weight or market-cap) and model stands for either the CAPM or the Fama-French three-factor model.

Table D.2 reports summary statistics for regression (1.13). The bias of the CSV measured with respect to standard asset pricing models is small in magnitude for both weighting schemes (in the order of 10^{-5}). While it remains statistically significant, we can safely consider that the impact of the bias remains immaterial

[1] As noted earlier, it would be straightforward to remove the impact of this bias by dividing the CSV measure by the factor $\left(1 - \sum_{i=1}^{N_t} w_{it}^2\right)$, equal to $\left(1 - \frac{1}{N_t}\right)$ in the EW case.

for any practical purposes. Another interesting finding is the sign of the bias. For the equal-weighted quantities, the sign of the bias is positive, while it is negative for the cap-weighted ones. Therefore, the beta bias dominates the weighting bias for equal-weighted averages in both models. This is consistent with the prediction made by the theoretical analysis regarding the relative impact of the weighting-bias for different weighting schemes. Regarding the model, the bias is larger when the idiosyncratic variance is measured with respect to the Fama-French model instead of the CAPM for both weighting schemes, as expected, but its magnitude remains negligible.

1.3.2 Comparison with Other Measures

In this section we compare the CSV measure to the afore-mentioned, more conventional, measures of idiosyncratic risk (i.e., FF-based, CAPM-based and GS). To obtain these other measures, we need to re-estimate the relevant factor model using a rolling window of one-month worth of daily data to allow for time-variation in beta estimates (or total-variance variation for the GS). In Table D.3, we report summary statistics for the monthly time series of annualized idiosyncratic variances based on 516 observations from January 1964 to December 2006.[1]

On the monthly series, the annualized means of the equally-weighted CSV, FF-based and CAPM-based measures are 38.4%, 38.3% and 38.7%, respectively, while the EW GS variance is 34.2%. The standard deviations are 8.5%, 8.6%, 8.7% for the CSV, FF-based and CAPM-based measures and 7.0% for the GS measure. For the cap-weighted version, the CSV, FF and CAPM idiosyncratic variance measures have an annualized mean of 8.5%, 7.6%, 8.0%, respectively and the GS measure mean is 11.2%. The standard deviations are also closer for the CSV, FF and

[1]In this section of the paper, we start the sample period in January 1964 to allow for direct comparison with Bekaert et al. (2008). In the predictability section, we instead start the sample in July 1963.

CAPM measures than for GS. Although GS argue that their measure fundamentally constitutes a measure of idiosyncratic risk, with the idiosyncratic component accounting for about 85% of the total EW average measure, it is strictly speaking an average of total stock variance. Our CSV measure is very close to idiosyncratic variance measures derived from traditional asset pricing models, confirming that the assumption about beta homogeneity is not a major problem.

The cross-correlation matrix reported in Table D.3 provides further evidence on the closeness of the CSV to the other model-based measures. Correlations are very high between CSV^{EW} and $CAPM^{EW}$ (99.93%) and FF^{EW} (99.75%, as well as between CSV^{CW} and $CAPM^{CW}$ (99.48%) and CSV^{CW} and FF^{CW} (98.56%). The high correlations between the CAPM and the FF measures (99.88% and 99.18% for EW and CW respectively) also indicate that adding factors does not drastically affect the estimation of idiosyncratic variance. Correlations between the GS measures and the other measures are always smaller but remain close to 90% when considering the same weighting scheme. Correlations between measures for different weighting schemes are much lower, irrespective of the estimation method, indicating that the choice over the weighting scheme is fundamentally important for estimating idiosyncratic variance, as stressed in our theoretical analysis in section 1.2.

Table D.4 provides mean and standard-deviation estimates for the daily average idiosyncratic variance measures. The mean of the EW CSV is 38.4%, practically equal to the mean of EW idiosyncratic variance based on the FF model. For the cap-weighted measures, the CSV has a slightly higher mean than the FF-based one. For the CSV daily series, the standard deviation is higher than for the FF-based measure for both weighting schemes. This is due to the different nature of the two series. The CSV only includes information from the cross-section of realized returns, while the FF idiosyncratic variance is a persistent, overlapping, rolling-window estimate. Each daily estimate of idiosyncratic variance for the FF model differs from the previous one by only two observations out of the approximately 21 trading days included in a month (the first and last days).

The smoothness of the idiosyncratic variance estimates obtained with the rolling-window methodology is illustrated in Figures D.1 and D.2, which plot daily CSV and FF idiosyncratic variances for each weighting scheme respectively. It should also be noted that the estimation of the FF-measure is computationally much more expensive than for the CSV measure, which is based on observable quantities.

The lower panel of Table D.4 presents cross-correlations for the daily series of idiosyncratic variance measures. Although the coefficients are smaller than for the monthly series, the relationship remains strong provided the comparison is done for the same weighting scheme: 82.6% and 73.9% for EW and CW measures respectively. The difference with the monthly series correlations may again be explained by the presence of the smoothed estimation procedure inherent to the FF-based measure. Overall, it appears that the CSV measure is extremely close to CAPM or FF-based measures at the monthly frequency, when the latter measures suffer from no particular bias, and that the CSV measure appears to be a good and instantaneous proxy for idiosyncratic variance at the daily frequency, when the standard measures are subject to artificial smoothing due to overlapping data.

1.3.3 Extracting Regimes in Idiosyncratic Risk

Bekaert et al. (2008) fit a Markov regime-switching model with a first-order auto-correlation structure (see Hamilton (1989b)) for the monthly series of idiosyncratic variance based on the FF model. In this section, we want to estimate this model with our CSV measure both at the monthly and daily frequencies. While we expect that the fit will be close to Bekaert et al. (2008) for the monthly series given our previous results on the similarity of the series, we want to verify whether such a model provides a similar fit for the daily series.

In this model, two regimes are indexed by a discrete state variable, s_t, which follows a Markov-chain process with constant transition probabilities. Let the current regime be indexed by i and the past regime by j and x_t be the original idiosyncratic variance. In this parsimonious model, x_t follows an AR(1) model:

$$x_t - \mu_i = \phi(x_{t-1} - \mu_j) + \sigma_i e_t,\ \ i, j \in \{1, 2\} \tag{1.14}$$

The transition probabilities are denoted by $p = P[s_t = 1|s_{t-1} = 1]$ and $q = P[s_t = 2|s_{t-1} = 2])$. The model involves a total of 7 parameters, $\{\mu_1, \mu_2, \sigma_1, \sigma_2, \phi, p, q\}$.

We first verify that the CSV and the FF-based measures give the same results for the monthly series. The estimation results for the monthly series of the FF^{CW}, CSV^{CW}, FF^{EW} and CSV^{EW} are reported in the upper panel of Table D.5. For corresponding weighting schemes, the parameters in both regimes are similar between the two measures. For both measures the low-mean, low-variance regime presents a higher probability of remaining in the same state.

We then fit the same model to the daily time series and present the parameter estimates in the lower panel of Table D.5. It should be stressed that for our CSV measure, the parameter values of the average level of idiosyncratic variance μ in both regimes are found to be quite close to the values obtained with the monthly series. This result suggests that the process observed at the daily frequency is not just a noisy series, but actually captures the same underlying process observed at the monthly frequency. This stands in sharp contrast with the FF-based measure, for which the maximum-likelihood estimation procedure could not recognize two regimes when daily data is used, as evidenced by the fact that the parameter values for the mean level of idiosyncratic variance are basically the same for the two regimes. This problem, combined with an autocorrelation parameter very close to one, is likely caused by the overlapping data problem present in the daily FF measure, which corresponds to the *smoothing* effect mentioned in the previous section.

In Figures D.3 and D.4 we plot the filtered probabilities (conditional on information up to time t) of remaining in state 1 (high-mean and high variance regime), as well as the monthly CSV and FF average idiosyncratic variance time series for the CW and EW weighting schemes, respectively.[1] At the monthly frequency, our measure and the FF-based measure appear to be remarkably close for both the equal-weighted and cap-weighted schemes. Also, we find that the dates of regime changes, marked by the filtered probabilities, are the same most of the times for the cap-weighted and the equal-weighted measures.[2] We also find that periods in the higher-mean and higher-variance regime are more persistent for the equally-weighted measure compared to the cap-weighted measure (except during the tech bubble period). Overall, our filtered probability series resembles closely the one presented in Bekaert et al. (2008) for the cap-weighted FF and Campbell et al. (2001) measures.[3]

The shaded areas in Figures D.3 and D.4, which time stamp the NBER recession periods, indicate that the peaks in the probability of remaining in the high-mean high-variance regime coincide most of the times with the contraction periods. Therefore, the CSV measure is counter-cyclical, the dispersion of returns

[1]These are estimates of the transition probabilities conditional to information up to time t given all sample data.

[2]One notable exception is the regime change of 1980 : 05, which is present for the cap-weighted measure and absent for the equally-weighted one.

[3]The small difference might come from the fact that Bekaert et al. (2008) fit a model with two different autocorrelation coefficients (one for each regime) as opposed to one. However, they find the two coefficients to be fundamentally equal in both regimes, which supports using a more parsimonious model.

being high and quite variable when economic growth subsides. In the next section, we want to analyze further the relation between the CSV and other economic and financial variables.

1.3.4 CSV Relation with Economic and Financial Variables

To put this analysis in the proper context, we should go back to the very nature of idiosyncratic risk. In an asset pricing model, it represents the risk that belongs specifically to an individual firm, after accounting for the sources of risk that are common to all firms. In the previous sections, we have shown that the cross-sectional variance of returns provides a very good measure of this idiosyncratic risk, even if it ignores the risk exposures to the usual common risk factors such as the market return or the Fama-French factors. Yet we concluded our time series analysis of CSV by stressing its strong counter-cyclical behavior. To pursue this analysis further we need therefore to rely on equilibrium models that link returns to economic fundamentals. Recently, Bansal and Yaron (2004) have revived consumption-based asset pricing models by showing that two sources of long-run risk — expected consumption growth and consumption volatility as a measure of economic uncertainty — determine asset returns. Further, Tédongap (2010) provides strong evidence that consumption volatility risk explains a high percentage of the cross-sectional dispersion in average stock returns for the usual set of size and book-to-market portfolios that have been used in tests of asset pricing models. Another strand of literature based on the intertemporal CAPM or the conditional CAPM has linked the cross-section of expected returns to other economic or financial variables such as the term spread, default spread, implied or realized measures of aggregate returns variance, and many others.

While our CSV measure is based on the cross-sectional dispersion of realized returns over the whole universe of traded stocks, as opposed to the cross-sectional dispersion of average returns of a limited number of size and book-to-market portfolios, the same theoretical implications should prevail. Therefore, we present below a simple correlation and graphical analysis of the relation between the CSV and some of these key variables. For the economic variables, we chose consumption-growth volatility as a measure of economic uncertainty. Following Bansal and Yaron (2004) and Tédongap (2010), we filter consumption-growth volatility with a GARCH model. For consumption, we used FRED's personal consumption expenditures of non-durables and services monthly series, divided by the consumer price index and the population values to obtain a per-capita real consumption series. We then compute its growth rate from July 1963 to 2006.[1] The second economic variable we consider is inflation volatility, which we filter also with a GARCH process.[2] For the financial variables we use Welch and Goyal (2008)'s data for corporate bond yields on BAA and AAA-rated bonds, long-term government bond yield and 3-months T-bill rate to estimate the credit spread and term spread (as the difference between the first and the second rate in both cases).[3] In Table D.6 we report the correlations between the equally-weighted and cap-weighted measures of cross-sectional variance and the five economic and financial variables during the 1990-2006 period. We also explore some potential asymmetries by computing the CSV^{EW} for the positive and negative returns.

[1]The series IDs at the FRED's webpage are, PCEND and PCES for "Personal Consumption Expenditures: Nondurable Goods" and "Personal Consumption Expenditures: Services", CPIAUCNS for "Consumer Price Index for All Urban Consumers: All Items" and POP for "Total Population: All Ages including Armed Forces Overseas". Bansal and Yaron (2004) used the Bureau of Economic Analysis data available at www.bea.gov/national/consumer_spending.htm on real per-capita annual consumption growth of nondurables and services for the period 1929 to 1998. The series is longer but is available only at annual and quarterly frequencies.

[2]For space considerations, we do not report parameter estimates for the two AR(1)-Garch(1,1) we estimate. They are available upon request from the authors.

[3]Data available at Amit Goyal's webpage: http://www.bus.emory.edu/AGoyal/Research.html

The highest correlation (0.401) is obtained between consumption growth volatility and the equally-weighted measure CSV^{EW}. In Figure D.5 we plot the two series for the period 1990 to 2006. While the CSV series is much noisier than consumption-growth volatility, the coincident movements between the two series are quite remarkable. After a high volatile period just before 2000, both series show a marked downward trend after the turn of the century. A reasonable explanation for this strong correlation is to think about a common factor (aggregate economic uncertainty) affecting the idiosyncratic variance of each security. Aggregating over all securities will make the CSV a function of economic uncertainty. In intertemporal asset pricing models of Bansal and Yaron (2004), Bollerslev et al. (2009) and Bollerslev et al. (2009), economic uncertainty is a priced risk factor that affects returns, therefore providing a fundamental rationale for the observed correlation between CSV and consumption growth volatility. This suggests that CSV should appear to be priced when a Fama-MacBeth procedure is applied to a set of portfolios. We explore this issue in Section 1.5. The correlation of the cap-weighted CSV with consumption growth volatility is not as high (0.241) since it puts more weight on large cap securities, which are in general less affected by economic uncertainty. Looking at the split between CSV^{EW+} and CSV^{EW-}, we see that the correlation is higher for the CSV when conditioning on the negative returns (0.346). This suggests that return dispersion in bear periods is relatively more affected by economic uncertainty.

The next most highly negatively correlated variable is inflation volatility (-0.367). Since 1998, inflation volatility seems to have been on an upward trend, while the cross-sectional variance of returns has been sharply declining. This is clearly apparent in Figure D.6. In presence of higher inflation uncertainty, investors will move towards allocating more to stocks relative to bonds in their portfolios, generating a general increase in stock returns that reduces their cross-sectional variance. The T-bill rate is also relatively highly correlated with CSV^{EW} (0.302). In the type of equilibrium models we have referred to, the risk-free rate, proxied here by the T-bill, will be a function of consumption growth volatility, hence its positive relation with the cross-sectional variance.

For the financial variables (credit spread and term spread), it is interesting to note that the higher correlations are with the cap-weighted measures of the cross-sectional variance. The signs are intuitive. Credit risk affects differently individual firm returns and therefore tends to increase CSV, while a pervasive term spread risk will reduce dispersion by being common to many securities due to a move of investors away from bonds into the stock market.

Given that the cross-sectional variance is significantly linked with economic and financial factors that have been shown to predict returns, we explore in the next section the predictive power of CSV for aggregate returns at various frequencies, especially at daily frequencies, since our measure of idiosyncratic variance allows us to measure CSV at any frequency without any artificial smoothing effect. This is a main advantage over other methods of recovering this idiosyncratic variance.

1.4 New Evidence on the Predictability of the Market Return

There is an ongoing debate on the predictive power of average idiosyncratic variance for average (or aggregate) stock market returns. Goyal and Santa-Clara (2003) find a significantly positive relationship between the equal-weighted average idiosyncratic stock variance and the cap-weighted portfolio returns for the period 1963:07 to 1999:12. They find that their measure of average idiosyncratic (in fact total) variance has a significant relationship with next month return on the cap-weighted portfolio. The regression in GS is as follows:

$$r_{t+1}^{CW} = \alpha + \beta \nu_t^{EW} + \varepsilon_{t+1}, \tag{1.15}$$

where ν_t^{EW} corresponds to GS_t^{EW}. In a subsequent analysis, Bali et al. (2005) argue that this relationship disappeared for the extended sample 1963:07 to 2001:12, and attribute the relationship observed in GS to high-tech-bubble-type stocks (i.e., stocks traded on the NASDAQ) and a liquidity premium. In a similar way, Wei and Zhang (2005) find that the significance of the relationship found by GS disappeared for their sample 1963:07 to 2002:12 and argue that the presumably temporary result of GS was driven mainly by the data in the 1990s. Wei and Zhang (2005) criticize the fact that GS looked at the relationship between an equally-weighted average stock variance and the return on a cap-weighted average stock return, as opposed to an equally-weighted portfolio return. Moreover, both Bali et al. (2005) and Wei and Zhang (2005) find no significant relationship between the cap-weighted measures and the cap-weighted portfolio return in all three sample periods (ending in 1999, 2001 and 2002, respectively).

1.4.1 Monthly Evidence

In this section we confirm existing results and extend them in a number of dimensions, including a longer sample period. The first panel in Table D.7 presents the predictability regression of equally-weighted variance measures on the cap-weighted return as in Goyal and Santa-Clara (2003) and Bali et al. (2005) for their sample periods, as well as the extended sample up to 2006:12. The regression is as in equation 1.15, where ν_t^{EW} corresponds to the EW CAPM-based measure and the CSV.[1] For comparison purposes we start the sample period in

[1] As explained before, the monthly CSV is the average of its daily values during the month.

this section in 1963:07, as in Goyal and Santa-Clara (2003), Bali et al. (2005) and Wei and Zhang (2005).

For the monthly series, we confirm that there is a significant positive relationship in the first sample, and also that it weakens for the subsequent extended samples.[1] The Newey and West (1987) autocorrelation corrected t-stat for 12 lags of the β coefficient of both CSV and the CAPM-based measures goes from 3.5 for the first sample period down to 0.9 for the largest sample. Consequently, the adjusted R^2 goes from 1.3% down to 0.04%. This result confirms the findings of Bali et al. (2005) and Wei and Zhang (2005) for the further extended sample. In section 1.4.4 we propose a possible explanation for this puzzling result.

In the second panel of Table D.7 we present the results of the regression between the equally-weighted average return with the lagged equally-weighted idiosyncratic variance measure, as given by:

$$r_{t+1}^{EW} = \alpha + \beta \nu_t^{EW} + \varepsilon_{t+1} \tag{1.16}$$

where ν_t^{EW} is taken as the CAPM-based average idiosyncratic variance or as the CSV measure. In contrast with the former regression, the relationship is found to be significantly positive for the three sample periods for both measures.[2]

In the third panel of Table D.7 we present the results for the three sample periods of the one-month-ahead predictive regression of the cap-weighted market

[1] We found a similar result using the GS measure of equally-weighed average variance. We do not present these regression results for the sake of brevity given that they generate a similar picture, which has also been confirmed in Bali et al. (2005) and Wei and Zhang (2005).

[2] Wei and Zhang (2005) find a significantly positive relation between the equal-weighted GS measure and the equal-weighted market return for the initial sample. They also test the robustness of the relation by using an equally-weighted cross-sectional variance of monthly returns. They found a significantly positive coefficient for predicting the equal-weighted portfolio return mainly for the long samples starting in 1928 but not for the sample going from 1963 to 2002. Note that our cross-sectional measures differ. Ours is an average of the daily cross-sectional variances over the month. Theirs is the cross-sectional variance of the returns computed over the month.

portfolio using the cap-weighted idiosyncratic variance return as a predictor. In this case, the beta of the idiosyncratic variance is not significant for all three sample periods. This result confirms the findings of Bali et al. (2005) and Wei and Zhang (2005) for the extended sample.

1.4.2 New Predictability Evidence at Daily Frequency

Prevailing measures used in the literature require a sample of past data to estimate additional parameters, constraining existing evidence to the monthly estimations. Fu (2009) finds that high idiosyncratic volatilities of individual stocks are contemporaneous with high returns, which tend to reverse in the following month. Huang et al. (2009) find that the negative relationship between idiosyncratic variance and expected returns at the stock level uncovered in Ang et al. (2006) and Ang et al. (2009) becomes positive after controlling for the return reversals. This provides additional motivation for looking at the predictability relation at a higher frequency than the monthly basis. Using the CSV as a proxy for aggregate idiosyncratic variance allows us to check this relationship at the aggregate (market) level in a more direct way (without having to control for reversals). Taking advantage of the instantaneous nature of the CSV, we run the same predictability regression (1.16) on the one-day-ahead portfolio return using the average idiosyncratic variance.

The upper panel of Table D.8 shows that at a daily basis, this relationship is much stronger, with (Newey-West corrected) t-stats of coefficients for the average idiosyncratic variance across the three samples ranging between 4 and 4.7.

In the lower panel of Table D.8 we report the results for the one-day-ahead predictive regression on the cap-weighted pairs (CSV and market return) for which we find the relation also to be positive and significant, but with a much more obvious deterioration of the t-stat of the cap-weighted idiosyncratic variance coefficient, going from about 5.91 in the first sample down to 1.97 for the longest sample. For this reason and for brevity, we now focus on the relationship between aggregate idiosyncratic risk and the equal weighted market return.[1]

1.4.3 Interpretation of Predictability results

Given this evidence on the predictability of average aggregate returns by idiosyncratic risk, a natural question to ask would be: why does the relationship between the equal-weighted measure and the cap-weighted differ across different sample periods?

Wei and Zhang (2005), Bali et al. (2005) argue that the relationship between idiosyncratic risk and the market index first found by Goyal and Santa-Clara (2003) on the sample ending in 1999:12 was driven by small stocks traded in the NASDAQ and the data coming from the dot-com bubble period. Although we confirm their empirical findings for our sample period, we disagree with their conclusion that the relationship between average idiosyncratic risk and expected returns disappeared since the end of the dot-com bubble. Even though it appears clear that NASDAQ companies played an important role in the relationship of the equal-weighted average idiosyncratic variance with the average market-capitalization expected return during the end of the 1990s, which (obviously) weakened after the burst of the bubble, we find that the relationship between average idiosyncratic risk and future average market returns is robust to choices of the sample period, provided that adequate weighting schemes and horizons are chosen to test this inter-temporal relationship.

The transitory relationship between the equal-weighted average idiosyncratic variance and the cap-weighted market index observed up to the end of the 1990,

[1] The corresponding results using a market cap-weighted scheme can be obtained from the authors upon request.

can be explained by the heterogeneous and transitory nature of the omitted sources of risk captured by idiosyncratic risk and its relation with the inflated valuation of several NASDAQ companies during that period [1].

Some intuition behind the far more robust relationship between the equally-weighted average idiosyncratic variance and the equally-weighted portfolio comes precisely from the logic of standard asset pricing theory. As discussed in the introduction, there are multiple reasons for which average idiosyncratic risk should be related to average returns, due to the heterogeneous sources that may compose idiosyncratic risk. According to CAPM, only systematic risk should explain future returns. However, if during a certain period of time there exists anomalies of any kind (priced omitted risk factors) that, presumably, are not proportionally reflected in the current market capitalization of the companies carrying these factors, then the omitted sources of risk are more likely to explain the returns of a portfolio where all kinds of firms are represented in a similar manner, such as the EW as opposed to a portfolio where big companies are proportionally better represented than smaller ones.

Along these lines, Pontiff (2006) argues that idiosyncratic risk is the largest holding cost borne by rational arbitrageurs in their pursuit of mispricing opportunities. This theory implies that the current level of idiosyncratic risk should predict returns since it should measure the amount of current mispricing opportunities present in the market. Assuming that the same mispricing opportunities disappear in the long run, it appears more likely to observe this relationship be-

[1] The strongest omitted factors in that period (call it the irrational.com factor), partially captured by the equally weighted idiosyncratic variance, started to be increasingly represented in the market-cap index, due to the suddenly-higher market capitalization of precisely the group of companies carrying this temporarily strong omitted factor. The posterior reversal of the situation (i.e., the burst of the bubble) subsequently explains the sharp fade in the relationship between the average idiosyncratic variance and the market-cap portfolio, precisely due to the posterior sudden deterioration of the market capitalization of most stocks carrying this irrational.com factor, and hence notably reducing their representation in the market-capitalization index.

tween idiosyncratic variance and returns over very short horizons. Moreover, all things being equal, large-cap stocks are less likely to present misspricing and hence the predictability implied by this theory would be more likely to be present on the equal-weighted index return rather than the cap-weighted index return, as we observed in predictive regressions at daily and monthly horizons.[1] The sign of the relationship is not predicted by Pontiff's theory in general, because it depends on whether the average (equal or cap-weighted) portfolio is over- or under-priced (it predicts a positive sign for underpriced stocks and a negative sign for overpriced stocks).

1.4.4 Robustness Checks

In this section, we test further explore the relationship documented in the former section in several dimensions. We first want to place the return predictability by idiosyncratic variance in the context of the literature of the risk-return trade-off. Most of the literature on this topic is based on a linear regression between return and volatility. We want to see if including the return variance in the regression changes the predictability results. Second, we test the robustness of the relationship in the presence of an option implied volatility measure. Third, we further test the predictability relationship at quarterly and annual horizons. Finally, we look for the potential asymmetry in the relationship between idiosyncratic variance and future average returns, when the cross-sectional variance is split in two and is computed for returns above or below the mean. Such an asymmetry often exists for positive and negative returns in the volatility modeling of financial time series. The reported presence of asymmetries will provide us with a motivation for

[1] It is well known that large cap stocks are more liquid than small-cap stocks, which implies a higher number of people trading them and usually a higher number of analysts looking at them. Together with less constraints to short-selling, we expect a higher price efficiency for large cap stocks.

extending the cross-sectional dispersion measure to the third moment and find this measure is related with average idiosyncratic skewness and has strong predictive power of the average market return.

1.4.4.1 Inclusion of Return Variance

In order to check wether the relationship between the market portfolio expected return and the aggregate level of idiosyncratic variance (which we document at the monthly and daily frequency) is robust to the inclusion of the variance of the market portfolio, we run the following joint regression:

$$r_{t+1}^{EW} = \alpha + \beta CSV_t + \vartheta Var\left(r_t^{EW}\right) + \epsilon_{t+1}. \tag{1.17}$$

We also run the univariate regression:

$$r_{t+1}^{EW} = \alpha + \vartheta Var\left(r_t^{EW}\right) + \epsilon_{t+1}. \tag{1.18}$$

For the monthly estimations of $Var\left(r_t^{EW}\right)$ we use the realized sample variance over the month (from daily returns). For daily estimations we fitted an AR(1)-EGARCH(1,1) model on the overall sample.[1] In the first two panels of Table D.9, we report regression results at the monthly and daily frequency of both (1.17) and (1.18). In the latter univariate regression, the variance of the equally-weighted portfolio returns does not appear to be significant in explaining the average future returns at the monthly and daily frequencies.

In the regression from equation (1.17), the coefficient of $Var\left(r_t^{EW}\right)$, ϑ, is negative and non-significant at the monthly frequency. At the daily frequency, the

[1] Using the overall sample to estimate the parameters would only give the portfolio variance an advantage to predict future returns. However, from the results we see that even when using such forward-looking estimates for $Var\left(r_t^{EW}\right)$, the significance of the CSV remains strong.

coefficient ϑ was still found to be negative and (marginally) significant. The significance of the CSV coefficient remains valid for both monthly and daily frequencies, and if anything improves slightly after the inclusion of the equally-weighted portfolio variance.

The latter two panels of Table D.9 present the regression results at the monthly and daily frequency of both (1.17) and (1.18) but using the cap-weighted index and CSV equivalents. The relationship at the daily horizon becomes non significant after the inclusion of the realized variance of the market cap-weighted index. At the monthly horizon the relationship remains non significant.

In Table D.10, where we report the quarterly and annual predictability with and without the market variance, we confirm that the equally-weighted cross-sectional variance does not forecast future average returns at low frequencies. However, for the cap-weighted measure of CSV, we observe predictability over the period 1963 to 2006 when it is joined with market variance. The sign is negative while the market variance enters with a positive sign as predicted by the benchmark risk-return trade-off[1]

One fair remark on the results of the predictability regressions is that the relationship using equal-weighted measures only holds at shorter horizons (i.e. daily and monthly). However, this result is in line with Pontiff (2006)'s interpretation of idiosyncratic risk as a barrier for arbitrageurs and with the evidence presented by Fu (2009) at the stock level, who finds that high idiosyncratic volatilities of individual stocks are contemporaneous with high returns, which tend to reverse in the following month.

[1] See also Guo and Savickas (2008) for similar results.

1.4.4.2 Inclusion of Market Realized Variance and Implied Variance

Other measures of variance have been used in trying to link market returns to a measure of market risk. Implied variance (VIX^2) has been used as a forward-looking measure of market variance in addition to realized variance (the sum of squared returns at higher frequency than the targeted frequency for the measure of variance)[1]. We use these measures in Table D.11 along with both CSV measures for daily and monthly predictability. We repeat the exercise in Table D.12 for quarterly and annual frequencies. For these regressions we start the sample in 1990 for data availability for the implied volatility.

Results are similar to the ones in the previous section with market variance. For CSV^{EW}, we observe predictability at high-frequency but not at low frequency, while it is the opposite for CSV^{CW}. For the daily estimates with CSV^{EW}, we find a R^2 of almost 5% when we include all three measures of variance, and all coefficients are significant. But the remarkable result, undocumented until now to our knowledge, is the very high R^2 obtained at quarterly and annual frequencies for the CSV^{CW} measure. When using CSV^{CW} alone as a predictor we obtain R^2s of 4% and 26% at quarterly and annual frequencies, respectively. Adding the implied variance brings the R^2s to almost 19% and 29%. If instead one uses the realized variance instead of implied variance the R^2s are close to 11% and 34%. In all these predictability regressions, the sign of the CSV^{CW} variable is negative.

Guo and Savickas (2008) argue that average idiosyncratic volatility is negatively related to future stock market returns possibly because of its negative correlation with the aggregate book-to-market ratio.[2] If idiosyncratic volatility is

[1]For example, for the monthly variance, one will sum the daily squared returns, while for the daily variance, it is customary to use five-minute or one-minute squared returns.

[2]The argument starts by considering average idiosyncratic volatility as a proxy for changes in the opportunity set related related to technological shocks. They argue that technological innovations have two effects on the firm's stock price: they tend to increase the level of the firm's stock price because of growth options and they also tend to increase the volatility of the

measured from a CAPM model then it will capture the missing book-to-market factor. This explanation runs counter to our previous findings regarding the very high correlation between the measures of idiosyncratic volatility based on the CAPM and the Fama-French models. The two series were almost identical. A more appealing explanation may be to think of cross-sectional variance as a measure of dispersion of returns reflecting the dispersion of opinions among market participants. The negative sign of this relationship at quarterly and annual horizons in the presence of market variance as the second predictor (and also at monthly horizons in the presence of implied variance as the second predictor) is consistent with the model of Cao et al. (2005), in which dispersion of opinions among investors is positively related to stock market volatility but negatively related to conditional excess stock market returns. Furthermore, one may argue that differences of opinions forge themselves over a period of time and hence this effect is more likely to be present at horizons longer than a day.

More generally, we may interpret the CSV as measuring the hedging terms in an intertemporal CAPM model. In this regard, it is interesting to see that the positive risk-return trade-off at the aggregate level, i.e., the relationship between market volatility and expected returns, becomes significant only when taking into account the presence of the omitted factors as captured by the CSV. It is also interesting to note that the interactions of the CSV with the realized variance of the market take place at longer horizons (quarterly and annual), while its interactions with implied variance (VIX^2) tend to be more important at shorter horizons.

firm's stock price because of the uncertainty about which firms will benefit from the new opportunities. The final argument is to say that the book-to-market ratio captures these investment opportunities

1.4.4.3 Asymmetry in the Cross-Sectional Distribution of Returns

We now explore for a potential asymmetry in the relationship between idiosyncratic variance and future average returns, when the cross-sectional variance is split in two and is computed for returns above or below the mean.

This asymmetry may be the result of the leverage effect put forward by Black (1976) since we are considering individual firms in the cross section. We also mentioned in an earlier section that consumption volatility risk affects differently small and large firms or value and growth firms. Therefore, we explore *i*) whether the predictability power is the same for the CSV of returns to the *left* and *right* of the center of the returns' distribution, *ii*) whether the relationship is driven by one of the sides and *iii*) whether the relationship with both sides would have the same sign on their coefficient. In order to do this, we define the CSV_t^+ as the cross-sectional variance of the returns to the right of the cross-sectional distribution (i.e., meaning the cross-section distribution that includes all stocks such that $r_{it} > r_t^{EW}$) and conversely define the CSV_t^- as the cross-sectional variance of the returns to the left of the cross-sectional distribution (i.e., meaning the cross-section distribution that includes all stocks such that $r_{it} < r_t^{EW}$). Then we run the following regression:

$$r_{t+1}^{EW} = \alpha + \beta^+ CSV_t^+ + \beta^- CSV_t^- + \epsilon_{t+1}. \tag{1.19}$$

Table D.13 presents the results of regression (1.19) for daily, monthly, quarterly and annual estimates, and shows a couple of interesting findings. First, splitting the CSV into right and left sides of the cross-sectional distribution made the adjusted R^2 of the predictive regression jump from 0.8% to 1.17% on monthly data and from 0.6% to 1.36% on daily data. Second, there is an asymmetric relationship between the CSV of the returns to the right and left of the cross-sectional distribution and the expected market return: the coefficient of the CSV_t^+

is positive while the one of CSV_t^- is negative in both daily and monthly regressions. However, the coefficients (of both right and left CSVs) are significant only on the daily regression. The summary statistics of the predictive regression on the cap-weighted index using the equivalent cap-weighted CSV measures, displayed in the lower panel of Table D.13, are qualitatively similar to the results on the equal-weighted measures.

These findings suggest that a measure of asymmetry of the cross-sectional distribution would be relevant in the context of exploring the relationship between market expected returns and aggregate idiosyncratic risk. Another key advantage of the CSV measure is that it can be easily extended to higher-order moments. We consider below the skewness of the cross-sectional distribution of returns and assess its predictive power for future returns. To the best of our knowledge, this additional factor, which appears as a natural extension of the CSV for measuring idiosyncratic risk[1], is entirely new in this context.[2] We follow Kim and White (2004) and use a quantile-based estimate (see Bowley (1920)), generalized by Hinkley (1975), as a robust measure of the skewness of the cross-sectional distribution of returns:[3]

$$RCS = \frac{F^{-1}(1-\alpha_1) + F^{-1}(\alpha_1) - 2Q_2}{F^{-1}(1-\alpha_1) + F^{-1}(\alpha 1)} \tag{1.20}$$

for any α_1 between 0 and 0.5 and $Q_2 = F^{-1}(0.5)$. The Bowley coefficient of skewness is a special case of Hinkley's coefficient when $\alpha_1 = 0.25$ and satisfies the Groeneveld and Meeden (1984)'s properties for reasonable skewness coefficients.

[1] We show formally in an appendix available upon request from the authors that there is a link between idiosyncratic skewness and the skewness of the cross-sectional distribution of returns.

[2] At the stock level, Kapadia (2009) uses cross-sectional skewness to explain the puzzling finding in Ang et al. (2006) that stocks with high idiosyncratic volatility have low subsequent returns.

[3] The usual non-robust skewness measure of the cross-section of returns is highly noisy compared to the proposed robust measure, especially at the daily frequency.

It has upper and lower bounds $\{-1, 1\}$.

In Table D.14, we report the results of predictive regressions at the daily and monthly frequencies where we add the robust measure of the cross-sectional skewness to the equally-weighted CSV. The first observation is that the CSV coefficients are very close to the values estimated with the CSV as the only regressor (0.4 for the daily frequency and 0.25 for the monthly one). The t-stats are also almost identical to the ones found in the CSV regressions. However, skewness appears to be a major contributor to the predictability of returns since the R^2 increases significantly compared to the regressions with CSV alone. At the daily frequency, the adjusted R^2 increases to a value of 5.8%. At the monthly frequency, it is still 4.6%. This large increase in predictability when adding skewness suggests that macroeconomic or aggregate financial shocks affect asymmetrically the distribution of returns.

1.5 Is Average Idiosyncratic Risk Priced?

According to Merton's ICAPM, a factor that predicts stock returns in the cross section should also predict aggregate market returns (see Campbell (1993)). By the reverse argument, motivated by the predictability power of (equal-weighted) cross-sectional variance on the average return in the market, we explore in this section whether the CSV^{EW}, interpreted as a risk factor, is rewarded and commands a premium in the cross-section.

1.5.1 CSV Quintiles' Premium

Using daily excess returns every month we run the following regression for each stock i:[1]

[1]We use stocks with non missing values during the current month.

$$r_{it} = \alpha + \beta_{i,csv} CSV_t^{EW}. \tag{1.21}$$

At the end of every month in the sample, we sort stocks using the CSV^{EW} factor loading, β_{csv}, and form equally-weighted and cap-weighted quintile portfolios. We calculate the average return during the overall period for each quintile and the average return difference (i.e., premium) between the first quintile and each of the other four quintiles.

The results for the equally-weighted quintile portfolios are displayed in the upper panel of Table D.15 and in the lower panel for the cap-weighted quintiles. As we can see from this table, all premia are significantly different from zero and economically meaningful. The difference between the first quintile (the one with higher sensitivity) and the second, third and fourth quintile, is around an annualized 30%, while the difference with the fifth quintile is around 15%. This result suggests that the relationship of the CSV and stock returns might not be best described in the simple linear form, which is in line with the asymmetric effect found in section 1.4.4, with the quantities CSV^+ and CSV^-.

1.5.2 Fama-MacBeth Procedure

In order to use the standard set of assets in the asset pricing literature, we extract daily returns data from Kenneth French data library on their 100 (10x10) and 25 (5x5) size/book-to-market portfolios for the period July 1963 to December 2006. Then we run every calendar month the following regression for each portfolio:[1]

$$r_{it} = \alpha + \beta_{i,xmkt} XMKT_t + \beta_{i,smb} SMB_t + \beta_{i,hml} HML_t + \beta_{i,csv} CSV_t^{EW}. \tag{1.22}$$

[1] As before, XMKT stands for excess market return, SMB and HML are the size and book to market Fama-French factors, also directly extracted from Kenneth French data library.

Using the recorded factor loading, β (monthly) time series, we run the following cross-sectional regression every month on the next month's excess returns and record the γ coefficients:

$$r_{it+1}^{m} = \gamma_0 + \gamma_{xmkt}\beta_{i,xmkt}(t) + \gamma_{smb}\beta_{i,smb}(t) + \gamma_{hml}\beta_{i,hml}(t) + \gamma_{csv}\beta_{i,csv}(t). \quad (1.23)$$

We finally test whether the average γ coefficients are statistically different from zero. In order to take into account possible serial correlation in the coefficients, we compute the t-statistic using Newey and West (1987) standard errors with 4 lags (same number of lags as in Ang et al. (2009)).

We use four sets of assets: 100 (10x10) size/book-to-market equally-weighted portfolios and cap-weighted weighted portfolios, and 25 (5x5) size/book-to-market equally-weighted and cap-weighted portfolios. For each of them, we use the CSV^{EW} as the fourth risk factor. The first two panels of Table D.16 present the corresponding Fama-MacBeth regression results. The table displays the annualized coefficients and standard errors (multiplied by 12 from the original monthly values), as well as their corresponding autocorrelation-corrected t-stat and the average R^2. We find the γ coefficient for CSV^{EW} to be positive and significant when we use the 100 and 25 size/book-to-market Fama-French equally-weighted portfolios. However, it is not significant when we use the 25 market cap-weighted portfolios and marginally significant for the 100 market cap-weighted portfolios (although positive in both cases). This later result, again, is not entirely surprising considering that the cross-sectional variation in returns is reduced through the market-capitalization adjustment.

1.6 Conclusion

In this paper we formally introduce an *instantaneous* cross-sectional dispersion measure as a proxy for aggregate idiosyncratic risk that has the distinct advantage of being readily computable at any frequency, with no need to estimate other parameters. It is therefore a model-free measure of idiosyncratic risk. We extensively show how this measure is related to previous proxies of idiosyncratic variance, such as the Goyal and Santa-Clara (2003) measure and measures relative to the classic Fama and French (1993) and CAPM models, which have been previously shown to be very close to the Campbell et al. (2001) proxy as well. We confirm previous findings of Goyal and Santa-Clara (2003), Bali et al. (2005) and Wei and Zhang (2005) on the monthly predictability regressions for the extended sample period using our cross-sectional measure and more standard measures of idiosyncratic variance. We find that the results are robust across these measures. Thanks to the instantaneous nature of our measure, we are able to extend to daily data the evidence on the predictability power of idiosyncratic variance on the future market portfolio return. We provide a statistical argument to support the choice of an equally-weighted measure of average idiosyncratic variance as opposed to a market-cap weighted and explain why both empirically and theoretically such a measure should forecast better the equal-weighted market return. We also showed that this cross-sectional measure displays a sizable correlation with economic uncertainty, as measured by consumption growth volatility, and with several economic and financial variables. One additional advantage of our measure is that it generalizes in a straightforward manner to higher moments and we showed that the asymmetry of the cross-sectional distribution is a very good predictor for future returns. We leave for further research an exhaustive analysis of the properties of the skewness of cross-sectional return distribution as a measure of average idiosyncratic skewness. We also leave for further research an empirical analysis of the CSV measure using international data.

Chapter 2

Predicting Stock Returns in the presence of Uncertain Structural Changes and Sample Noise

Daniel Mantilla-García and Vijay Vaidyanathan*

The power of the dividend price ratio to predict future stock returns has been the subject of intense scrutiny. Most studies on return predictability assume that predictor variables follow stationary processes with constant long run means. In view of recent evidence of the role of structural breaks in the dividend-price ratio mean, we propose an estimation method that explicitly incorporates the uncertainty about the location and magnitude of structural breaks in the predictor in order to extract the regime mean component of the dividend-price ratio. We find that adjusting for structural changes in the ratio's mean and estimation error improves the predictive explanatory power of the dividend-price ratio in-sample, as well as its out-of-sample forecasting ability to a very significant extent.

Keywords: Return Predictability, Structural Breaks, Bayesian Change Point Analysis, Dividend-Price Ratio, Out-of-sample.

*We thank René Garcia, Abraham Lioui, Lionel Martellini and Gideon Ozik for their valuable feedback and comments, Khanh-Linh Loth for her excellent research assistance, conference participants at the 3rd International Conference on Computational and Financial Econometrics (CFE'09), Limassol, Cyprus and seminar participants at EDHEC-France for useful discussions and comments. Daniel Mantilla-García is PhD in Finance candidate at EDHEC Business School and Head of Research and Development at Koris International (daniel.mantilla@koris-intl.com). Vijay Vaidyanathan is PhD in Finance candidate at EDHEC Business School and President of EDHEC-Risk Indices and Benchmarks - North America (vijay.vaidyanathan@edhec-risk.com).

2.1 Introduction

Although the voluminous literature related to the predictability of returns has a long history (e.g. see Spiegel (2008)), the issue remains the subject of controversy and a very active field of research. The subject of debate in the empirical literature encompasses its very existence and nature, the type of predictors as well as the models and methods used to forecast returns. Perhaps the only uncontroverted fact about return predictability is that it is a central question in financial economics from both theoretical and applied standpoints.

A variety of financial variables have been proposed as predictors, including the dividend price ratio, dividend growth ratio, price earnings ratio, dividend payout ratio, stock variance, book value, book to price ratio etc. (see Goyal and Welch (2008) for a comprehensive list). Of these, the dividend price ratio is seen as particularly promising, in part because of the Campbell and Shiller (1988) linearization of the definition of a return. This identity discussed in Cochrane (2008) states that either the dividend price ratio or dividend growth must predict returns. Goyal and Welch (2003) take a contrarian view. While acknowledging the identity they argue that it need not hold in short horizons, where, in effect, the current dividend price ratio predicts little more than the following period's dividend price ratio.

An overview of this vast literature suggests that the evidence of predictability has failed to be conclusive. Supporting evidence has come in two flavors: one arguing that predictability *must* exist (e.g. Cochrane (1992), Cochrane (2008)), and the other showing *some* positive results that are not strong enough to be considered as sufficient proof by its detractors (e.g. Goyal and Welch (2003), Goyal and Welch (2008)). The sources of skepticism include the scarce out-of-sample prediction results, the marginal statistical significance of regression coefficients, the instability of regression coefficients, and the sample sensitivity of the evidence (e.g. Timmermann (2008)).

The branch of the literature that supports the existence of predictability rejects the classical random-walk paradigm of unpredictable and approximately constant expected returns, and posits one that admits the possibility that expected returns contain a time-varying component, either related to economic-cycles (e.g. Lettau and Ludvigson (2005)) or the existence of other regimes that imply predictability of future returns. In this view, the poor out-of-sample R^2 is not necessarily viewed as invalidating the null hypothesis of predictable returns. For instance, Cochrane (2008) states: (p. 1566) *"One can simultaneously hold the view that returns are predictable, ... , and believe that such forecasts are not very useful for out-of-sample forecasting and portfolio advice, given uncertainties about the coefficients in our data sets"*. Indeed, there is even some question about whether the poor performance of out-of-sample tests implies that in-sample results are spurious (Inoue and Kilian (2004)).

Although out-of-sample predictability has been alleged by some to be weak, a few recent methodological proposals for return forecasting finally appear to be successful. Campbell and Thompson (2008) demonstrate that the performance of the out-of-sample tests can be shown to be significant, for instance, by imposing weak and theoretically sound restrictions on the signs of the model parameters. Arguing model uncertainty and instability, Rapach et al. (2009) find that using combinations of individual forecasts (e.g. average forecast from univariate predictive regressions) produce economically significant out-of-sample predictability evidence of the equity premium for quarterly data.

In another vein, Ferreira and Santa-Clara (2010) propose an entirely new methodology to forecast returns out-of-sample, which seems to be very successful in producing stock returns' forecasts, at both annual and monthly horizons. This so called "sum-of-the-parts method" consists of decomposing returns into dividend yield, earnings growth, and price-earnings ratio growth and separately forecasting the three components to obtain stock market returns forecasts. Interestingly, there is no in-sample test for this approach as there is within the predictive regression framework. One potential extension of their method would be to use the model we propose here for the dividend yield mean to better forecast the dividend yield (instead of returns).

Assuming an AR(1) process with a constant trend for expected returns and starting from the Campbell-Shiller identity, Lacerda and Santa-Clara (2010) derive a predictor composed by the dividend-price ratio series adjusted with a forecast of the dividend growth and the mean of dividend-price ratio itself that seem to be very successful for predicting at annual horizons, out-of-sample. Their adjustment is quite different from ours because the resulting adjusted predictor series seems to be even less stationary than the original dividend-price ratio series. However, since Lacerda and Santa-Clara (2010)'s adjustment uses an estimation of the dividend-price ratio mean, it might be possible to further improve return forecasting accuracy by combining the estimation method we propose to estimate the d-p mean, with their adjustment.

Although the Campbell-Shiller identity would imply a linear relationship between the dividend-price ratio and expected returns, McMillan (2009) uses a framework significantly different from the linear predictive regression called the exponential smooth transition model (ESTR). This non-linear model assumes that the parameters describing the relationship between returns and the d-p ratio change over time, taking different values corresponding to a fixed number of regimes. They propose a model that implies 4 different regimes over which the model parameters migrate over time in a smooth fashion. They find that their model out-performs the random walk model in terms of root mean square prediction error, for monthly data in 3 out of 8 markets over the 1980-2007 period.

In a similar spirit, Dangl and Halling (2009) propose to use predictive regressions with time-varying coefficients to predict the risk premium at monthly horizons and find significant improvements in terms of mean square prediction error with respect to the random walk when combining univariate predictive models. Similar to Rapach et al. (2009), they report that their predictions are significantly better during recession periods. Dangl and Halling (2009) document that uncertainty about the level of time-variation in coefficients and uncertainty about the choice of predictive variables are equally important sources of predictive variance. A priori, structural breaks in the predictors and time varying coefficients may coexist. Hence, models such as Dangl and Halling (2009) or McMillan (2009), and ours can be seen as potentially complementary for return forecasting purposes.

Thus, regardless of the point of view on predictability, there is agreement that the uncertainty about the regression coefficients contributes to weakening the predictive power, both in-sample and out-of-sample. Indeed, there is overwhelming evidence (e.g. Ang and Bekaert (2002), Pastor and Stambaugh (2001)) of the presence of such structural breaks. The existence and importance of such breaks have been emphasized by Paye and Timmermann (2003), Pastor and Stambaugh (2001) and Pettenuzzo and Timmermann (2005).

Recently, Lettau and van Nieuwerburgh (2008) shows that adjusting the d-p ratio for structural breaks in its mean, leads to a significant increase in the significance and magnitude of the (in-sample) predictor's coefficient and the regression R^2. However, using their model in combination with a regime switching model (Hamilton (1989a)), they produce out-of-sample forecasts that do not outperform the return's historical average in terms of mean square error. Lettau and van Nieuwerburgh (2008)'s adjustments derive the location of the break by using the Bai-Perron algorithm (Bai and Perron (2003), Bai and Perron (1998)), which estimates the most likely (using a specified criterion) number and location of break points, and subsequently assumes that the timing of the breaks is certain to estimate the d-p ratio regime mean. Unfortunately, the information on the true location of the break is itself uncertain, and therefore introduces uncertainty in the estimate of the parameters.

In this paper, we propose an alternative mechanism to adjust for structural breaks in the d-p ratio's mean by incorporating the uncertainty about the timing of the structural breaks. In contrast to the Bai-Perron algorithm, we suggest a Bayesian approach for detecting the change points. Instead of using likelihood measures to estimate the breaks, the Bayesian approach we use produces a probability for each observation to be a change point. Thus, we are able to use all observed data to characterize the distribution of change points, and generate a posterior mean for the predictor. We show that incorporating the uncertainty related to the locations of the structural changes in this manner produces an adjusted predictor series that is better behaved (stationary) and significantly improves the level and significance of the regression coefficient and the regression R^2 with respect to Lettau and van Nieuwerburgh (2008)'s adjustment. We also show that these improvements are robust to choices of the sample period as well as alternative indices of the stock market portfolio. In order to estimate the probability of a change point and the posterior mean, we use the BH-BCP algorithm of Barry and Hartigan (1993), Erdman and Emerson (2007). Finally, exploiting a measure of uncertainty produced by the bayesian algorithm, we propose a sample noise shrinking methodology that, in combination with the structural changes adjustment, provides sound evidence of out-of-sample return predictability that could have been exploited in real time.

2.2 Dividend-Price Ratio Predictability Power and Structural Changes

The theoretical motivation to use the dividend price ratio to forecast returns comes from Campbell and Shiller (1988)'s log linearization of the ratio, for which the dp's stochastic component is related to expected returns as follows:

$$dp_t = \bar{dp} + E_t \sum_{j=1}^{\infty} \rho^{j-1} \left[(r_{t+j} - \bar{r}) - (\Delta d_{t+j} - \bar{d}) \right], \tag{2.1}$$

where $\rho = \left(1 + exp(\bar{dp})\right)^{-1}$ is a constant, establishing an inter-temporal relation between future returns and dividends. Cochrane (2008) argues this identity implies that the dividend-price ratio must predict either future returns or dividend growth. Given the lack of predictability power of dividend growth by the dividend price ratio that he and other authors documented "if *both* returns and dividend growth are unforecastable, then the price/dividend ratio is constant, which it obviously is not". However, Cochrane (2008)'s argument stands in contrast to the evidence presented by Goyal and Welch (2008), which find that the predictive power of the dividend price ratio (and a battery of other commonly used predictors) have vanished after the oil crisis in the mid 70's.

One should note that the framework in equation (2.1) implies that the steady-state of the economy is constant over time, meaning that the long-run growth dividend rate $\bar{d}$, the long run dividend price ratio level, $\bar{dp}$ as well as the average long-run return of equity, $\bar{r}$ are fixed. Although this assumption seems to hold reasonably in the data for the dividend growth and returns, the dividend price ratio displays structural changes in its mean, making this assumption unrealistic.

However, Lettau and van Nieuwerburgh (2008) extended the Campbell-Shiller framework to allow variations (and even permanent changes) in the steady-state of the economy. Under mild assumptions[1] about the dividend-price ratio steady state level dynamics, they show that, similar to the case with constant steady-state, the log dividend-price ratio is the sum of the steady-state dividend-price ratio and the discounted sum of expected returns minus expected dividend growth in excess of steady-state growth and returns,

$$dp_t = \overline{dp_t} + E_t \sum_{j=1}^{\infty} \rho_t^{j-1} \left[(r_{t+j} - \bar{r}) - (\Delta d_{t+j} - \bar{d}) \right] \tag{2.2}$$

where $\rho_t = \left(1 + exp(\overline{dp_t})\right)^{-1}$. The difference between (2.2) and (2.1) is that the mean of the log dividend-price ratio not only varies over time but it could also be non-stationary.

On the other hand, a standard specification of forecasting variables used by Stambaugh (1986), Stambaugh (1999), Nelson and Kim (1993), Lewellen (1999), Pástor and Stambaugh (2009) and others, is to assume predictors x_t to be a stationary processes with constant mean,

$$x_t = \mu + u_t \tag{2.3}$$

where the mean of the predictor variable, μ is constant and the stochastic component, u_t is assumed to be stationary. It is also standard to assume a linear relationship[2] between the predictor and expected returns, conditional to available information, $\mathcal{F}$ at time t,

[1] Lettau and van Nieuwerburgh (2008) assume the steady-state log dividend-price ratio to be (approximately) a martingale, $E_t\overline{dp}_{t+j} = \overline{dp_t}$. They note that although steady-state growth, expected return and mean dividend-price ratio must be constant in expectations, the steady-state might shift unexpectedly.

[2] For a predictive framework where this assumption is relaxed see Pástor and Stambaugh (2009). However, their predictive system also assumes stationary processes for the predictors, making our model also relevant for their more general setting.

$$E\left(r_{t+1}|\mathcal{F}_t\right) = \bar{r} + b[x_t - \mu]. \tag{2.4}$$

The presence of structural changes in the steady-state mean of the dividend-price ratio documented by Lettau and van Nieuwerburgh (2008) implies a non-stationary series, which would not be a well-suited predictor variable in this setting[1]. However, the decomposition of the dividend-price ratio into a stationary stochastic term and a non-stationary time-varying steady state mean in (2.2), implies that a stationary predictor variable could be obtained from the difference between the current level of dp_t and its non-stationary component $\overline{dp}_t$. Using this idea, and replacing $\mu_t = \overline{dp}_t$ and $x_t = dp_t$ equations (2.3) and (2.4) turn into:

$$x_t = \mu_t + u_t \tag{2.5}$$

$$E\left(r_{t+1}|\mathcal{F}_t\right) = \bar{r} + b[x_t - \mu_t] = a + b\tilde{x}_t. \tag{2.6}$$

Hence, the adjusted series $\tilde{x}_t$, which should be a stationary one, could be obtained if the timing and magnitudes of shifts in the steady state mean $\overline{dp}$ can be adequately estimated. Lettau and van Nieuwerburgh (2008) used the structural breaks methodology developed by Bai and Perron (1998) to identify the most probable dates when a break happened in the dividend price ratio series. Lettau and van Nieuwerburgh (2008) then treated the blocks (i.e. regimes or sub-periods) implied by the break dates as independent sets of information, from which they estimate the conditional mean of the underlying process by simply calculating the sample average for each block which is assumed to be constant within the sub-period. This methodology to estimate $\overline{dp}_t$ does not take into account the uncertainty intrinsic to the estimation of the regime changes and the parameter

[1] Put here references on problems with persistent predictors

estimation error, and is ill-adapted for out-of-sample return forecasting[1].

Instead, we propose estimating the time-varying mean of the dividend-price ratio by using an algorithm that explicitly incorporates the uncertainty about the time when the regime changes happen. This algorithm, originally developed by Barry and Hartigan (1993), called Bayesian Change Point analysis (BCP), allows us to decompose the dividend price ratio into its persistent current-regime mean and its stationary components. This decomposition, induces significant improvement for return predictability evidence in both in-sample and out-of-sample tests, as we show in the following sections.

In the rest of this section we explain in more detail how the BCP algorithm incorporates the uncertainty about the structural changes in the steady-state mean level of the dividend price ratio, then we discuss how this model compares with other models of time series structural breaks and then we introduce a complementary methodology to adjust for parameter estimation error for out-of-sample return forecasting in this context.

2.2.1 Modeling Structural Change Uncertainty with BCP

Consider a sequence of numbers X, consisting of T observations $X_1, X_2, ..., X_T$. We define a partition of X, denoted by ρ, as a sequence of T indices $U_1, U_2, ..., U_T = 1$, where each U_i is 1 if the ith element of X is the end of a block and is 0 otherwise. A block, in this context, is a sub-sequence of X consisting of contiguous elements from X. Thus a partition splits X into a series of non-overlapping contiguous sub-sequences, or blocks. We denote a block by an index pair of the preceding block ending and the block's ending indices (i, j). Thus, each block (i, j) consists

[1]If one would use Lettau and van Nieuwerburgh (2008)'s methodology to forecast return out-of-sample, just after a break date is declared all past observations are dropped, leaving just one data point with which to estimate the current steady-state level of dp, which is a priori, not a very reliable estimate.

of observations $X_{i+1}, X_{i+2}, ...X_j$.

For example, consider the sequence of length 6, $X = 2, 4, 3, 6, 8, 7$. One possible partition of X is $\rho = 001001$. This partitions X into two blocks, the first is block $(0, 3)$ comprising of the sub-sequence $2, 4, 3$ and the second block is $(3, 6)$ which is the sub-sequence $6, 8, 7$. Another possible partition of X is $\rho = 101001$ which splits X into 3 blocks - the singleton (i.e. of length 1) block $(0, 1)$, and the blocks $(1, 3)$ and $(3, 6)$. Note that we use the index 0 to denote the first index of the first block.

In the context where X is a time series, each block therefore corresponds to a contiguous period in time, and represents a temporal "regime" that serves to group all observations during that regime into a block. Any partition ρ is a T dimensional vector, where the Tth component must be 1, but each of the other $T - 1$ values can either be a 0 or a 1.

The algorithm works by considering the space of all possible partitions $\mathcal{P}$ of the sequence X. Since X is of length T, and each data point (except the last, which must be an end-of-block) represents a possible end-of-block (i.e a change point), there are 2^{T-1} possible partitions of X. However, it is possible to compute an exact solution in polynomial time rather than exponential time, because many of the blocks across the different partitions are identical. Since there are are $\binom{N+1}{2}$ possible blocks, the exact solution is $O(n^3)$, making it computationally taxing. Furthermore, it is possible to use Gibbs sampling methods to obtain an MCMC approximation in linear time.

To see how sampling from $\mathcal{P}$ can yield the parameters of interest, recall that any partition $\rho \in \mathcal{P}$ can be represented by $\rho = U_1, U_2, ..., U_T$ where U_i is zero if i is not the end of a block, and U_i is 1 if i is the end of a block. U_T is fixed at 1. If one knew the distribution of ρ, then $E(U_i)$ (under the prior) would yield the probability that the ith observation represented a change point.

The key insight of the BH-BCP algorithm is that given a partition and distribution parameter, we can obtain the likelihood of that combination of partition, parameters and data.

Thus, instead of enumerating all possible partitions, the algorithm samples from the space of all possible partitions, given the observed data. If we draw as a sample a partition $\rho = U_1, U_2, ..., U_T = 1$, each U_i has a value of 1 at a change point and 0 otherwise. Hence, the mean of the U_i across all the samples yields the posterior probability that i is a change point. In a similar manner, we can estimate the means for each block in the sample partition, and combine block mean estimates from all sampled partitions to estimate the posterior block mean for the series as well as the posterior variance of the block means.

2.2.2 The BH-BCP Gibbs Sampling Algorithm

The algorithm starts with the following assumptions. The ith block is assumed to have a block mean of μ_i and all data points in X are assumed to be drawn from $\mathcal{N}(\mu_i, \sigma^2)$. The probability of any point i being a change point is assumed to be p, independently for each i. BH-BCP then imposes a prior distribution on μ_i to be $\mathcal{N}(\mu_0, \sigma_0^2/l)$ where l_i is the length of block i. This has the effect of giving a "higher probability to small departures from μ_0 in large blocks than it does in small blocks; we can expect to identify small departures if they persist for a long time" (Barry and Hartigan (1993)). An estimate of the block mean μ_i is $\bar{\mu_i} = (1-w)\bar{B_i} + w\mu_0$ where $\bar{B_i}$ is the sample mean of the block B_i, μ_0 is the sample mean of X and w is the ratio $\sigma^2/(\sigma_0^2 + \sigma^2)$. Of course, w also needs to be estimated from the data as do p_i, for which BH-BCP provides a fully bayesian solution with reasonable priors for μ_0, σ^2, p, w.

In order to sample from $\mathcal{P}$, the sampler exploits the fact that the conditional distribution $p_i = f(U_i|S, U_j \forall i \neq j)$ is easily computed, allowing us to sample a value of U_i from this conditional distribution. The sampler begins with $P = 0, 0, 0, ..., 1$. For each step of the markov chain, we sample from $f(U_i|S, U_j \forall i \neq j)$. As Erdman and Emerson (2008) have pointed out, the expressions given in Barry and Hartigan (1993) are numerically unstable, and we use their alternate formulation provided in Erdman and Emerson (2008), given by:

$$\frac{p_i}{1-p_i} = \frac{P(U_i = 1|S, U_j \forall i \neq j)}{P(U_i = 0|S, U_j \forall i \neq j)} = \frac{\int_0^{p0} p^b (1-p)^{n-b-1} dp}{\int_0^{p0} p^{b-1} (1-p)^{n-b} dp} \times \frac{\int_0^{w0} \frac{w^{b/2}}{(W_1 + B_1 w)^{(n-1)/2}} dw}{\int_0^{w0} \frac{w^{(b-1)/2}}{(W_0 + B_0 w)^{(n-1)/2}} dw}$$

where b is the number of blocks, W_0, B_0, W_1, B_1 are the within and between the blocks sum of squares when U_i takes the values of 0 and 1 respectively. The integrals are incomplete beta functions which are easily computed, using numerical procedures dating back to Newton (e.g. DiDonato and Morris Jr (1992)). $p0$ and $w0$ are tuning hyper-parameters that can be set to values of less than 1 in order to impose ad-hoc heuristics to limit possible values of p_i and w_i. However, our implementation makes no ad-hoc impositions, and we consider the full distribution by keeping them both at 1. Carefully choosing ad-hoc limiting values through these parameters marginally improve our results, but we opt to avoid these limits and fix them at the full value of 1.

Given a value for p_i, we can now sample a value for U_i from the Uniform distribution, and then proceed to the next step of the chain. At the end of $T - 1$ steps, we would have a sample P from $\mathcal{P}$ which can be repeated for as many MCMC iterations as needed. The jth MCMC iteration yields a sample partition P_j from $\mathcal{P}$, from which we derive an estimate for w_j and therefore for $\mu_{i,j}$ for $i \in \{1...T\}$. Averaging over all the MCMC iterations, we compute the posterior mean π_i for $i \in \{1...T\}$ as the sample average of each $\mu_{i,j}$ over the j MCMC partitions. Similarly, we also compute the posterior variance $\nu_i \in \{1...T\}$ from the sample. Although the primary object of interest at this stage is the posterior mean, we shall show that the posterior variance plays a key role in allowing us to estimate the precision of the posterior mean, which we exploit in obtaining an optimal shrinkage estimator for out-of-sample forecasting.

2.2.3 Relation with alternative structural change models

The Bai-Perron method uses a dynamic programming algorithm to pick the most likely partition P^* in $\mathcal{P}$. The choice of P^* is made on the basis of user supplied constraints on the number of blocks, and picking amongst feasible solutions using an information criterion such as BIC, AIC or log likelihood. Specifically, the algorithm does not provide guidance on constraints on the number of blocks, and is unable to provide an estimate of posterior means or the variance of the posterior means.

However, the Bai-Perron method yields results that are usually consistent with BH-BCP in the sense that the break points identified by Bai-Perron tend to correspond with points that BH-BCP ascribe high probabilities of being change points. In that sense, the BH-BCP algorithm may be thought of as providing results that are consistent with Bai-Perron while providing additional information on the distribution of other possible partitions, and the implications of that distribution on the posterior means and variances of data points in X.

Another approach to identifying discontinuities in X that has been deployed in the literature is to use a state-space approach that models each block as a state. Similar to the BH-BCP algorithm, Markov Switching models have the appealing ability to yield a sequence of filtered and smoothed probabilities that any point represents a transition to a new state. However, Markov Switching models tend to be difficult to estimate in practice because of the relatively large number of parameters that need to be estimated (as a function of the number of states). Since we do not wish to impose either a fixed number of states, nor do we have any particular motivation in estimating a state transition probability matrix, a Markov State Switching model is not necessarily the most natural choice. However, in our analysis of out-of-sample predictions, we include the results of using a Markov Switching model to predict the current regime mean in Appendix B.1.

However, BH-BCP may be thought as a variation of a Markov Switching model where the number of states is unconstrained and the likelihoods are computed under the distribution implied by the imposed priors. Since the number of states is unconstrained, there are no constraints on the transition probability matrix either. The transition probability matrix reduces to a transition probability sequence representing the probability that any particular data point in the input sequence represents a transition to a new state.

In our approach to return predictability with dividend-price ratio, regression coefficients are determined conditional upon structural breaks on the predictor variable. In this sense, one could see predictive regressions with time-varying coefficients (see Dangl and Halling (2009) for instance) as an alternative approach to adjust for structural changes, where the breaks are captured by allowing variations or structural instability in the coefficients instead of in the predictor's mean. In the time-varying coefficient approach, the size and uncertainty of structural breaks are captured through a presumed variation of (estimated) parameters, which are unobservable.

On the other hand our approach admits the existence of structural breaks on the observed dividend-price ratio's historical data. Furthermore, it is interesting to see that equation (2.6) can also be interpreted as a model with a time varying intercept (given by $\bar{r} - bE_{xt}$) as opposed to a time varying slope coefficient.

2.3 Shrinking Noisy Parameter Estimates

The R^2 of the predictive regression reflects the correlation of realized returns with the estimation of expected returns (forecasts) obtained with parameters fitted to the same sample used to assess predictability accuracy. While expected returns are not observable, realized returns, which are a combination of expected and unexpected returns, are arguably the best approximation we have to assess the forecasting performance of a model for expected returns[1]. On the other hand, if there is "noise" in model parameters estimates (such as adjusted predictors or regression coefficients) for small samples the predictive regression may over fit the parameters to the observed "noisy" sample, producing inaccurate out-of-sample forecasts. If this is the case, then the out-of-sample predictions would not be as robust as implied by the regression statistics. One way to mitigate the over-fitting problem inherent in predictive regressions is parameter shrinkage. In this section we propose a minimum variance shrinkage for the estimated BCP-adjusted dividend price ratio and describe the (mean-squared error) optimal shrinkage for the estimated regression coefficients proposed by Ashley (2006), which we subsequently use for out-of-sample predictability tests.

2.3.1 BCP adjusted dividend price ratio and minimum variance shrinkage

Although the BH-BCP algorithm is best suited to in-sample estimation of the posterior mean, the fact that it also produces posterior variances can be exploited to compute an optimal shrinkage estimator for out of sample forecasts as well.

[1]There might be other ways to measure expected returns, based for example in analysts forecast, but they might be as subject as realized returns to measurement errors. Having said this, from a practical perspective, predicting realized returns, and not expected returns is what "matters".

The typical out-of-sample forecast involves estimating values for the intercept a and slope b parameters of the linear model as well as the latest value of the adjusted predictor, $\tilde{x}_i$. From equation (2.6), the *ith* return forecast is obtained as,

$$\hat{r}_{i+1} = \hat{a} + \hat{b}\hat{\tilde{x}}_i, \tag{2.7}$$

where $\hat{\tilde{x}}_i = x_i - \pi_i$ is an estimate of $\tilde{x}_i = x_i - \mu_i$. Unfortunately, since the regime mean, μ_i is not observable –but it has to be estimated– the adjusted predictor estimated value is an imperfect predictor. Hence, the observed value of the adjusted predictor contains an unobservable predictive component and an unobservable noise component.

The alternative to using the observed estimated value of the predictor is to take the extreme position that the observed predictor value has no predictive component and any deviation from the current value of the posterior mean, π_i is comprised entirely of noise. As equation (2.6) illustrates, only deviations from the latest value of the regime mean predicts deviations from the steady-sate return average. Under this assumption, this model feature squares nicely with the random walk hypothesis, making $\hat{r}_i = \bar{r}$[1].

On the other hand, using the observed adjusted predictor value is the alternate extreme position that the observed value is entirely composed of the predictive component with no noise. Thus we have two different estimates for the value of the predictive component of the predictor variable: (i) the random walk estimate, which would ignore the observed data entirely for forecasting purposes and assume it has zero predictive power (null deviation from the current value of the regime mean), and (ii) the naïve estimate, which would assume that the estimated value of the predictor was fully informational and contained no small-sample noise.

[1]One can see this by replacing $x_i = \pi_i$ in equation (2.6).

Instead of using one or the other estimate, we combine these two approaches using a linear shrinkage estimator to obtain an estimate that is optimal in the sense of it being the minimum variance estimator under the modeled assumptions.

Denote p^{rw} and $\tilde{p}$ as the random walk and the naïve estimates of $\tilde{x}$. Under the random walk hypothesis, the predictor estimate is such that the adjusted value is 0, i.e. $p^{rw} = x_i - \pi_i = 0$ which implies that $x_i = \pi_i$, because any deviation from the posterior mean of the predictor value is noise. Since the regime mean is not observable, the estimation error would be proportional to the variance of its estimator. Hence, we estimate the variance of the random walk predictor estimate $Var(p_{rw})$ as is the mean of the posterior variance ν, which we denote as $\bar{\nu}$.

On the other hand, an estimator for the variance of the naïve estimator $\tilde{p}_i = x_i - \pi_i$ is simply the variance of the estimated adjusted predictor, $Var(\tilde{p}) = Var(\hat{\tilde{x}})$. The minimum variance estimator that combines these two estimates is a weighted sum of the two estimators, weighted by the reciprocal of the variances. Since the predictor value for the random walk is zero, we get:

$$p_i^* = \frac{Var(\tilde{p})}{\bar{\nu} + Var(\tilde{p})} \times p_i^{rw} + \frac{\bar{\nu}}{\bar{\nu} + Var(\tilde{p})} \times \tilde{p}_i$$
$$p_i^* = \frac{Var(\tilde{p})}{\bar{\nu} + Var(\tilde{p})} \times 0 + \frac{\bar{\nu}}{\bar{\nu} + Var(\tilde{p})} \times (x_i - \pi_i).$$

Under the assumptions above, the linear shrinkage estimator predictor with the lowest variance is given by:

$$p_i^* = \frac{\bar{\nu}}{\bar{\nu} + Var(\tilde{p})} \times (x_i - \pi_i),$$

where x_i is the ith observation of the unadjusted predictor and π_i is the ith estimated value of the regime mean. Plugging the shrunk estimator for the adjusted predictor back in equation (2.7), yields the BCP-adjusted out-of-sample forecast:

$$\hat{r}_{i+1} = \hat{a} + \hat{b}p_i^*. \tag{2.8}$$

2.3.2 Optimal shrinkage for Predictive Regression Coefficients

The predictive regression coefficients are also estimated values, that a priori suffer from the same small-sample over-fitting problem previously mentioned. Ashley (2006) shows that the unbiased forecast is no longer squared-error optimal in this setting. Instead, the minimum mean squared error forecast represents a shrinkage of the unbiased forecast toward zero. Similar to the minimum variance shrinkage proposed for the adjusted predictor, the shrinkage target for the slope coefficient also coincides with a prior of no predictability (random walk). Following Connor (1997) we correct the estimated regression coefficients as,

$$b^* = \frac{s}{s+j}\hat{b}$$
$$a^* = \bar{r_s} - b^*\bar{x_s},$$

where $\bar{x_s}$ is the historical mean of the predictor up to time s. As the the slope coefficient is shrunk toward zero, the intercept needs to be adjusted to preserve the unconditional return mean. The shrinkage intensity j (measured in units of time periods) is proportional to the weight given to the prior model of no predictability. Ashley (2006) shows that the mean-squared error optimal shrinkage intensity is given by $j = 1/\rho$, where ρ represents the expected explanatory power of the predictive model, and is defined as the expectation of a function of the regression

R-square:

$$\rho = E\left[\frac{R^2}{1-R^2}\right] \approx E\left[R^2\right].$$

Plugging the shrunk estimator for the adjusted regression coefficients back in equation (2.7), yields an out-of-sample forecast with shrinkage:

$$\hat{r}_{s+1} = a^* + b^* x_s. \tag{2.9}$$

2.4 Empirical implementation

Using annual returns (with and without dividends) for the value-weighted broad market index from the Center for Research in Security Prices (CRSP) data for the 1927−2010 period, we construct the corresponding time series of the dividend-price ratio[1]. Table E.1 reports the first and second order autocorrelation coefficients for the log dividend-price together with an Augmented Dickey Fuller test, testing the null hypothesis of a unit root. The first and second order autocorrelation coefficients are 0.92 and 0.84 respectively and the null hypothesis cannot be rejected. These are clear signs of non-stationarity of the log dividend price ratio (raw) series.

Using the BCP algorithm, we decompose the dividend-price ratio into its steady-state (current-regime) level and its transitory (hopefully stationary) component. Figure E.1 displays in its upper panel the time series of the dp ratio together with the estimated posterior mean. One can see that the posterior mean is a slow moving (persistent) time series with smooth variations, in contrast to the somehow artificial step function implied by the estimation procedure implemented in Lettau and van Nieuwerburgh (2008). A priori, there might be an advantage in using the BCP algorithm since Barry and Hartigan (1993) have shown that their

[1] We chose the start of the sample period as in Goyal and Welch (2008).

method is superior to a number of structural change alternatives, in detecting sharp short-lived changes in the parameters.

If the methodology we propose is able to extract the non-stationary component of the dividend price ratio, the adjusted series $\tilde{dp}$, which is equal to the raw dp series minus its posterior mean, should be a stationary one. If this is the case, the adjusted series would be a better suited variable to forecast returns within the classic linear predictive regression setting. In order to verify this hypothesis, we perform the Augmented Dickey Fuller test and estimate the autocorrelation coefficients on the adjusted series. We also perform this analysis on the dp series adjusted by regime means using the methodology proposed by Lettau and van Nieuwerburgh (2008) with one (1991) and two (1954 and 1994) structural breaks.

In Table E.1 we can see that both, the first and second order autocorrelation coefficients (AC) drops dramatically with respect to the unadjusted series' ones, for quarterly and annual time series. Interestingly, the coefficients also present a significant improvement with respect to the adjusted dp series using Lettau and van Nieuwerburgh (2008)'s method. In annual data (upper panel of Table E.1) the AC(1) went from 0.78 and 0.66 (for one and two breaks adjustment) to 0.02 for the BCP adjusted and the AC(2) from to 0.55 and 0.30 to -0.19 respectively. The lower panel of Table E.1 confirms an important improvement when the same break points and the BCP method is applied, with AC(1) falling from 0.73 and 0.63 (for one and two breaks adjustment) to -0.14 using the BCP adjustment and AC(2) from 0.73 and 0.62 to 0.097 respectively. The Augmented Dickey Fuller test, is also clearer in its rejection of the unit root null hypothesis for the BCP adjusted series than for the ones using the Lettau and van Nieuwerburgh (2008) breaks adjustment. The last line of each panel in Table E.1, shows that the posterior mean is a very persistent series (more than the original dp series).

2.4.1 Evidence from Predictability Regressions

We now turn to run the predictability regression using unadjusted log dp and adjusted series using one or two breaks as in Lettau and van Nieuwerburgh (2008) in addition to the adjustment we propose using the Barry-Hartigan Bayesian Change Point methodology. The upper and lower panels of Table E.2 summarizes the results of a one year ahead and one quarter ahead predictive regressions on the log returns of the CRSP value-weighted broad market index, testing hypothesis (2.6). First, the results confirm Lettau and van Nieuwerburgh (2008)'s finding: "While the statistical significance of the coefficient on the unadjusted dividend-price ratio is marginal, coefficients on the adjusted dividend-price ratios are strongly significant." Second, from the table we shall notice that the the $\widetilde{dp}$ adjusted using the BCP posterior mean presents much higher coefficient values and Newey-West adjusted t-stats of 1.28 (8.02) while the adjusted series with the former methodology have lower coefficient values of 0.23 (4.48) and 0.38 (4.57) for one and two breaks respectively in annual data. The R^2 is also improved for the BCP adjusted series (about 19%) with respect to the alternative one and two breaks adjustment (9% and 15%) and with respect to the raw dp series (4%) on annual data. A similar improvement is observed using quarterly series (lower panel of the table). The higher predictability power of the BCP adjusted series is consistent with the notable relative improvement achieved regarding the non-stationarity correction reported in Table E.1.

Similar to results reported in Lettau and van Nieuwerburgh (2008), the former results use the full sample to estimate the current level of the regime mean of the dp ratio. We now test whether the robust predictability evidence presented above, is robust to the sample period chosen and if it could have been recognized before. Using a growing window of data, we estimate the model parameters (e.g.

posterior mean) and run the predictive regression for every subsample using information available up to time t, just as if the test would have been performed in real time every year. Starting the first calibration period with data up to 1965 (same as in Goyal and Welch (2008)) we repeat the exercise every year until the end of the sample. Figure E.2 displays the Newey-West corrected t-stats for the predictor's coefficients time-series of regression and the corresponding R^2s using the raw dp series and the BCP-adjusted series[1]. We find regression's coefficients, t-stats and R^2 to be remarkably stable over time and broadly consistent with the in-sample result of the whole sample, in every sub period we looked at. The predictor coefficient's Newey-West adjusted t-stat has in fact increased since the first subsample, where it presents a value over 5. These results contrast with the less stable and decreasingly significant regression statistics when using the pure dp ratio as a predictor.

Consistent with identity (2.2), our results indicate that the transitory component of the dividend-price ratio has predictability power at the annual horizon, while its steady-state regime-mean is a very persistent series. For this reason, when the raw series of the d-p ratio is used as the predictor variable, as in (2.4) predictive regression, its predictability is blurred by its persistent component.

2.4.2 Out-of-Sample Predictability Evidence

Goyal and Welch (2008) provided evidence that the random walk model outperforms the predictive regression models using (raw) financial ratios ratios among other common predictor variables in out-of-sample predictability tests. Interestingly, Lettau and van Nieuwerburgh (2008) find that, in spite of the considerably

[1] Lettau and van Nieuwerburgh (2008) report a similar growing-window exercise but they use future information in order to declare the breaks in advance for adjusting the dp series. In unreported results we run the same regression using a rolling window of data 30 points and found that the coefficient is of the BCP adjusted series is always higher than the one obtained with one or two breaks adjustments.

improved in-sample predictability evidence documented using the dividend-price ratio series adjusted for structural breaks in mean (using the methodology they proposed), forecasts produced using the adjusted series could not out-perform the historical average in out-of-sample tests.

Unlike the adjustment for structural breaks proposed in Lettau and van Nieuwerburgh (2008), the forecasting methodology we propose captures the uncertainty in structural breaks of the steady-state level of the dividend-price ratio (see section 2.2) and adjusts for parameter estimation error using the optimal shrinkage introduced in section 2.3. We now turn to assess the forecasting accuracy of this predictive methodology. For comparison purposes we also perform out-of-sample forecasts using the raw dividend price ratio with and without Connor (1997)'s regression coefficients's shrinkage.

Similar to Goyal and Welch (2008), we split the sample in two: an initial calibration sample comprised by data from 1927 up to 1965 and an out-of-sample testing period (comprised by the rest of the available data) used to evaluate the forecast accuracy (and significance) of real time forecasts that could have been produced using the model proposed in former sections.

The standard way to measure the out-of-sample predictive power is a measure based on the mean squared predictive error (MSE) of the predictive model with respect to the MSE of the prevailing historical average (null hypothesis of no predictability). Predictive errors are the difference between the predicted value at time t and the realized return on the market at time $t+1$. A broadly used out-of-sample performance measure in the predictability literature is the R^2_{OS} introduced by Campbell and Thompson (2008), which is given by:

$$R^2_{OS} = 1 - \frac{MSE_{pred}}{MSE_{mean}}. \tag{2.10}$$

Where, MSE_{pred} is the mean squared error of the model predictions and MSE_{mean} is the mean squared error of using the return's historical average. A positive value for R^2_{OOS} means that the predictive model out-performed the historical average in terms of cumulative predictive error during the OOS period (and hence an investor could have exploited its predictive power in real time).

We assess the statistical significance of the out-of-sample forecasting gains of using the predictive model using the MSE-F statistic proposed by McCracken (2007), given by

$$MSE - F = (T - s_0)\frac{(MSE_{mean} - MSE_{pred})}{MSE_{pred}}$$

which tests for equal MSE of the unconditional (historical mean) and conditional forecasts (T stands for the total size of the sample period and s_0 for the initial calibration sample). We indicate statistical significance using asterisks according to their critical values for recursive schemes (growing window tests).

We also consider a related forecasting power analysis chart introduced by Goyal and Welch (2003), in which one plots the cumulative difference in MSE of the forecasting model under scrutiny with respect to the prevailing average. In periods during which the forecasting model provides a better (worse) estimate than the historical average, the line presents a positive (negative) slope and the sign of the plotted value matches the one that the R^2_{OS} would have at each point in time.

We also look at an alternative measure of forecasting accuracy introduced in Mincer and Zarnowitz (1969). The Mincer-Zarnowitz regression is given by:

$$r_{t+1} = \alpha + \beta\hat{r}_{t+1} + \epsilon_t \tag{2.11}$$

where $\hat{r}_{t+1}$ is the return forecast done at time t. Note that the regression R^2_{MZ} measures prediction accuracy only if the coefficient of the forecasted returns β is positive (accurate forecasts should have a positive relationship with the forecasted returns).

In order to measure economic risk-adjusted benefits of the forecasts, we calculate realized utility gains for a mean-variance investor on a real-time basis. Following Campbell and Thompson (2008), and Goyal and Welch (2008) we compute the average utility for a mean-variance investor with relative risk aversion parameter $\gamma = 3$ who allocates her portfolio every period between stocks and risk-free bills using forecasts of the stock market return with allocation limits to the stock index of $[0\%, 150\%]$. We perform simultaneously 4 portfolio strategies using 4 different forecasting models for stock returns based on: (i) the historical average, (ii) the raw dividend price ratio combined with OLS regression coefficients, (iii) the raw dividend price ratio combined with shrunk regression coefficients and, (iv) BCP-adjusted dividend price ratio and shrunk parameter values. This exercise also requires the investor to forecast the variance of stock returns which we approximate simply with the sample estimate (for all of them). A mean-variance investor who forecasts the equity premium will decide at the end of period t to allocate the following share of her portfolio to equities in period $t+1$,

$$w_{k,t} = \left(\frac{1}{\gamma}\right)\frac{\hat{r}^k_{t+1}}{\hat{\sigma}^2_{t+1}},$$

where $\hat{r}^k_{t+1}$ corresponds to the kth model excess return forecast[1], where $k = \{1, 2, 3, 4\}$ corresponding to each of the above mentioned models. Over the out-of-sample period, the investor perceives an average utility level of

$$\mathcal{U}_k = \bar{\theta}_k - \left(\frac{1}{2}\right)\gamma\bar{\sigma}^2_k,$$

[1] We use the latest observed value of the risk free rate as the next period forecast and substract it from the market return forecast of each model to obtain excess return forecasts.

for which, $\hat{\theta}_k$ and $\hat{\sigma}_k^2$ denotes the sample return mean and variance of the portfolio strategy using the kth predictive model.

We report utility gains in excess of the non-predictability random walk model which is the difference between utility gains using the kth predictability model ($k = \{2, 3, 4\}$) minus the utility gains of using historical average to forecast returns ($k = 1$), and multiply this difference by 100 to express it in average annualized percentage return. The utility gain corresponds to the certainty equivalent return, which can be interpreted as the portfolio management fee that an investor would be willing to pay to have access to the additional information available in a predictive model relative to the information in the historical return average (i.e. assuming no predictability).

The upper panel of Table E.3 reports the R^2_{MZ} and R^2_{OS} (in percentage terms) together with excess utility gains (denoted as Δ) obtained using the same calibration sample used in Goyal and Welch (2008) - using data up to 1965 - and perform the test over the rest of the period to determine the out-of-sample performance measures for each of the forecasting models. We obtain an R^2_{OS} over 8% which is significant at the 99% confidence level and positive utility gains, compared to the negative R^2_{OS} and utility losses obtained using the raw dividend-price ratio as a predictor with both simple OLS and shrunk regression coefficients. These results confirm that the predictive power of the (adjusted) dividend price ratio, documented in Section 2.4.1, is not spurious. Additionally, they imply that the forecasting methodology proposed in this paper could have been exploited in real time by investors to obtain consistent benefits with respect to the no-predictability assumption of stock market returns. Although we do not take into account transaction costs to assess the benefits of the strategy, they should not outweigh the benefits in this case, given that portfolio rebalancing transactions occur only once a year.

Goyal and Welch (2008) argues that former predictability evidence was mainly driven by data immediately after the oil price shock of 1975. To assess this concern, we repeat the OOS predictability exercise starting to forecast in 1976. Finally we perform the exercise starting to forecast in 1947, for which we have a calibration sample with a reasonable minimum number of data points to estimate the model parameters. The OOS performance measures, R^2_{OS}, R^2_{MZ} and excess utility gains, for these additional out-of-sample periods are reported in the middle and lower panels of Table E.3. We find again that the methodology proposed provides significant forecasting improvement with respect to the no predictability benchmark model and the other two forecasting alternatives considered.

We also report in Appendix B.1 the results of applying a regime-mean adjustment computed by using Markov Switching models. We find that those models are unable to consistently beat the random walk out of sample.

However, in unreported results[1] we find that the forecasting accuracy of the methodology at quarterly horizon with or without the error shrinkage methodology, is not superior to the prevailing historical average over the sample periods analyzed (i.e. we found negative R^2_{OS}). Furthermore, Figure E.4 shows that the predictive model is not always a better estimate than the prevailing historical average. In particular, during periods in which sharp structural changes took place (such as the internet and the credit bubbles crashes) the forecasting accuracy of the model diminished with respect to the non-predictability hypothesis estimate. On the other hand, this is not entirely surprising because the BCP method is designed to detect structural changes that took place in a time series and not to detect when the next break will happen. In other words, the BCP algorithm can tell us when structural changes did take place and is not intended to do any more. In order to test this statement, we perform a pseudo out-of-sample test, in which

[1] Available upon request from the authors.

we use the full sample to estimate the dividend-price ratio varying mean. In this test, using past adjusted dp series and past returns, we estimate at every point in time the regression parameters to forecast next period's return, record the forecasting error and repeat the exercise during the OOS period 1965 to 2010. For comparison purposes, we use the adjusted dp series adjusted by one and two breaks as estimated by Lettau and van Nieuwerburgh (2008). In Table E.4 we present the same OOS performance measures used for the OOS exercise for annual data and the equivalent result in Table E.5 for quarterly data. In this Pseudo-OOS exercises, we find a remarkable improvement in all three forecasting accuracy measures with respect to the non-predictability benchmark. We also find again a significant improvement with respect to the structural breaks adjustment proposed in Lettau and van Nieuwerburgh (2008). Furthermore, the pseudo out-of-sample predictions consistently outperformed historical average all along the OOS period as we can see from Figure E.5, i.e. the forecasting power analysis chart displayed a (steep) positive slope in almost every sub-period in the sample.

2.5 Robustness

In order to further test the robustness of the predictability evidence using the forecasting methodology outlined in this paper, we repeat the out-of-sample forecasting exercise using different portfolios of assets, namely: the S&P 500 index, the equal-weighted CRSP broad market index (Table E.6), the Fama-French Size portfolios (Table E.7) and Book-to-Market portfolios (Table E.8).

For the other market indices (S&P 500 and EW CRSP broad market index) we find qualitatively similar returns to the ones presented for the CRSP value-weighted broad market index. For the size and book-to-market components, the gains are less clear-cut, yet interesting. We find significant forecasting gains with respect to historical average only for the "big" stocks' portfolio and for low (growth) and high (value) - but not for medium- book-to-market stock portfolios.

2.6 Conclusion

The predictive power of dividend price ratio has been questioned by former empirical studies (e.g. Goyal and Welch (2003) and Goyal and Welch (2008)) that find that its return forecasting ability was temporary – a rather episodic phenomenon– and its relation with future returns, spurious or unstable over time, due to the persistent behavior that the ratio presented over the last few decades.

However, interpreting the observed level of the dividend price ratio as a deviation from a current-regime changing mean level, restores a stationary predictor, with renewed implications for return predictability. Considering that there is now overwhelming evidence that there are structural breaks in the dividend-price ratio mean, estimation of the break points is a crucial step towards extracting the predictive component of the ratio.

Our methodology incorporates uncertainty related to the structural change in prediction parameters. We show, using a Bayesian Change Point algorithm, that incorporating this uncertainty significantly improves virtually every measure of predictability evidence in regression coefficients, as well as providing robust evidence across different sample periods and data sets that were not evident with prior techniques. Taking advantage of one by-product measure of the bayesian algorithm (i.e. posterior variance of the predictor variable), we introduce an estimation error shrinking methodology that, in combination with the structural changes adjustment, provides evidence of out-of-sample return predictability that could have been exploited in real time.

Chapter 3

Growth Optimal Portfolio Insurance and Portfolio Insurance's Growth Rate Characterization

Daniel Mantilla-García*

Through a decomposition of the *growth rate* of the standard portfolio insurance strategy (CPPI) we unveil the (perhaps) surprising role that the correlation between the underlying assets plays on the performance of this type of investment strategy. We also find a close relationship between the growth rate and the long term value of the strategy even under common leverage and short-selling constraints. Then we introduce the growth optimal portfolio insurance strategy (GOPI), which combines the intuitively appealing objective of maximizing the value of the portfolio in the long run and the common constraint of insuring a fixed proportion of the portfolio expressed in terms of the value of a given benchmark. We find that this strategy tends to outperform the equivalent CPPI with the standard multiplier over long horizons. Interestingly, the level of the optimal multiplier turns out to be lower than the one implied by the standard methodology in most scenarios. Hence the outperformance achieved by the GOPI does not come at the cost of a higher risk exposure.

Keywords: Portfolio Insurance, Growth Rate, Growth Optimal Portfolio, Correlation.

*The author would like to thank Jean René Giraud, Lionel Martellini and Vijay Vaidyanathan for very interesting discussions and comments and Kevin Sheppard for making available his useful UCSD Multivariate Garch toolbox. The author is PhD candidate at EDHEC Business School and Head of Research & Development at Koris International. Contact details are: **email**: daniel.mantilla@koris-intl.com, **Address**: Koris International, Parc Technologique de Sophia-Antipolis, 45 allée des Ormes, Bâtiment D, BP 1200, 06254 Mougins Cedex - France. **Phone**: + 33 (0)4.92.28.70.61.

3.1 Introduction

The growth optimal portfolio (GOP) has been the subject of a long debate in the literature of portfolio selection since its discovery by Kelly (1956) and Latane (1959). It stands as an alternative to the utility maximization paradigm of portfolio selection and recently a whole theory of asset pricing based on this portfolio has been developed as an alternative to risk-neutral pricing (see for instance Platen (2005) and Christensen et al. (2005) for a complete review of the role of the GOP in finance today).

The GOP has interesting theoretical properties such as outperforming any other portfolio in the long run in terms of wealth (see for instance Thorp (1971)) and to minimize the time to reach a given level of wealth (see Pestien and Sudderth (1985)). It also has the intuitive appealing of maximizing the expected geometric return mean and the median of wealth[1] in the long run (see Ethier (2004)). These properties gained it the support from several authors, in articles such as Markowitz (1976), Breiman (1961), Luenberger (1998), Long et al. (1990). Authors arguing in favor of the GOP believe that growth optimality is a reasonable investment objective in itself for long horizon investors.

On the other hand criticisms of the GOP, such as Samuelson (1963) and Ophir (1978), steam from the view that the only rational approach to portfolio selection is to maximize expected utility. Although the GOP happens to be as well the result of maximizing the expected utility of terminal wealth for the logarithmic utility function (which is also a special case of the power utility), they argue that it is too much of a stretch to treat every investor as a log-utility maximizing investor (see Christensen (2005) and Hakansson and Ziemba (1995) for a review of the

[1] Maximizes the median of wealth in the long run for a portfolio has no relevance from an utility maximization standpoint. However, for skewed distributions the median is a measure of the most likely outcome, thus it might be an interesting property from a practical perspective.

origin and debate about GOP portfolios).

The concept of utility based portfolio selection, although widely used, has been criticized by the observation that investors may be unaware of their own utility functions or behave in manners that would be in strong contrast with its predictions (see for instance Bossaerts (2002)). Wether or not one has a strong believe on which is the right way to represents investors preferences, there seems to be complete agreement that, no matter how long (finite) horizon the investor has, the GOP can neither proxy nor dominate every other strategy in terms of expected utility, different from the log, because utility based portfolios would be more attractive in terms of their risk profile (see Merton and Samuelson (1974)). On the other hand, the growth optimization approach has a practical advantage over the expected utility maximization one, namely its ability to empirically verify ex-post its performance relative to other investment strategies. The growth rate maximization approach and the long-run growth property are formulated in dollars instead of utility units, thus "*it seems plausible that individuals, who observe their final wealth will not care that their wealth process is the result of an ex-ante correct portfolio choice, when it turns out that the performance is only mediocre compared to other portfolios*" (Christensen et al. (2005)).

Christensen et al. (2005) also notes that, even if the GOP dominates another portfolio with a very high probability, the probability of the outcomes where the GOP performs poorly may still be unacceptable to an investor who is more risk averse than a log-utility investor. In other words, the left tail distribution of the GOP may be too "thick" for an investor who is more risk averse than the log-utility investor. An alternative to introduce an investors' risk aversion within the growth maximization approach are the so-called fractional Kelly strategies. These strategies use a fixed allocation between the risk free asset and the GOP that depends on a risk aversion coefficient, similar to the classic approach of Sharpe's ratio maximization and the fund separation theorem (see Grauer (1981) and Platen (2005) for an alternative derivation of the CAPM and the introduction of the risk-aversion parameter).

However, individual investors might have behavioral attitudes toward risk incompatible with the solely objective of maximizing the probability of attaining the highest possible level of wealth in the long run at all costs (the latter objective being represented in the GOP) and for which a fixed allocation to the riskless asset might not be flexible enough to match these preferences. Furthermore, in spite of their typically long horizons, institutional investors are often subject to short-term regulatory constraints such as limits to their underperformance of a given benchmark or on the portfolio's value relative to the price of their liabilities (funding ratio constraints). In this paper, we introduce a portfolio strategy that addresses this particular need for a well defined risk control relative to a benchmark combined with the intuitively appealing objective of maximizing the growth rate of the portfolio. In order to do so, we set the investment objective as maximizing the growth rate of the portfolio subject to the constraint of insuring a minimum floor value for the portfolio, the former being expressed in terms of the value of a given stochastic benchmark. The result of combining this two objectives is equivalent to maximize the growth rate of the classic portfolio insurance strategy, known as Constant Proportion Portfolio Insurance (CPPI) (see Perold (1986), Black and Jones (1987), Perold and Sharpe (1988) and Black and Perold (1989)).

Although most former research on the properties of the CPPI has focused on the standard case with a riskless constant interest rate, we focus on the more general case with a stochastic Reserve or "Core" asset, because of its particular interest for pension funds, individual investors and portfolio managers[1].

[1] There are at least three practical applications for which portfolio insurance strategies with a reserve asset different from cash are crucial. First, from a pension fund perspective, holding too much cash can be quite risky because its liabilities have typically long terms (duration). In this

Through the decomposition and analysis of the CPPI's growth rate we revisit the properties of the portfolio insurance strategy. The characterization of the portfolio insurance's growth rate has conceptual and practical bi-products that were obscure and/or unavailable with former analytical characterizations of the strategy. In particular, we characterize and isolate the role that the correlation between the reserve and the active (or risky) asset plays in determining the value of the strategy and provide closed-form formulas for the value of the strategy *in dollar terms* for the general case with a stochastically moving reserve asset. More importantly, by maximizing the growth rate of the strategy we derive a growth-optimal multiplier which defines the growth optimal portfolio insurance strategy (GOPI). Our results suggest that the growth optimal strategy outperforms the equivalent standard parametrization of the CPPI over long horizons. Interestingly, the growth-optimal strategy presents a more conservative risk profile since its multiplier is most of the times lower than the one implied by the standard methodology.

Our analysis of the growth rate of the portfolio insurance strategy reveals that the "diversification benefits" of low correlations among the underlying assets of unleveraged fixed-mix portfolios is reversed for this type of strategy: the higher the instantaneous correlation between the reserve and the risky assets, the higher the value of the portfolio insurance strategy, everything else being equal. The

case the reserve asset is usually defined as a portfolio composed of fixed-income securities trying to match the obligations of the fund. Second, individual investors might want to insure defined long term benefits or bequest objectives and/or a stream of future consumption needs (Amenc et al. (2009)). This objective can be addressed in the construction of the reserve asset in a similar way as for the pension fund case by treating the future cash-flow needs as liabilities. Third, asset managers might be given the objective of outperforming a particular Benchmark. One way to comply with this relative performance objective is to define as the Core asset the (presumably) stochastic Benchmark so that the possible underperformance of the portfolio with respect to the Benchmark is limited to a well defined level, while still allowing for increasing upside potential coming from available risk premia and active manager views (for a detailed explanation of the practical advantages of this approach called Dynamic Core-Satellite allocation, see Amenc et al. (2004)).

intuition for this effect is that a higher level of correlation induces less relative return reversals" between the two assets, and hence it diminishes what we call the "rebalancing drag"[1]. The importance and the role of correlation, which grows exponentially with the level of the multiplier, is thoroughly illustrated with a graphical analysis. The positive effect of the correlation in this strategy might be counter intuitive for some readers due to a very common confusion of this measure with the relative trend of assets (see Lhabitant (2011) for a very clear illustration of this misinterpretation).

The very definition of the growth rate of the portfolio is based on specific model assumptions for the dynamics of the assets, continuous rebalancing and unlimited leverage and short-selling. For this reason, we also provide an empirical verification with real data of the close relationship between the growth rate and the portfolio insurance's value with discrete rebalancing and leverage/short-selling constraints.

3.2 Assets' Properties and Portfolio's Growth Rate

Former studies on the growth rate of rebalanced portfolios such as Fernholz (2002) focus on the impact that rebalancing has on the compounded return of unleveraged fixed-mix portfolios with assets that present similar long term returns and volatilities. This approach addresses the particular objective of building equity portfolios usually compared to buy-and-hold or market-cap-weighted benchmarks, such as the S&P 500. In what follows we illustrate that the impact that assets' properties such as volatility, correlations and differences in expected return among assets,

[1] In that special case of the CPPI with a riskless asset in the Core, there is an equivalent effect called "volatility cost" (Black and Perold (1992)) also known as "volatility drag".

have on the performance of portfolio insurance strategies is reversed with respect to their effect on fixed-mix portfolios.

3.2.1 A first Intuition on the impact of Trends and Reversals on (Leveraged) Return Compounding

One way to analyze the properties of the Dynamic Core-Satellite or Portfolio insurance strategy[1] is to consider the return of the "index ratio", i.e. the quotient of the values of the performance-seeking asset (also called Satellite), denoted S and the reserve asset (or Core), denoted R, i.e. $I(t) = S(t)/R(t)$. Changes in the index ratio are driven by the relative performance of the DCS's components. Black and Perold (1992) show that for a return δ in the index ratio, the fractional change in the Cushion is proportionally magnified by m (the multiplier), $\Delta C/C = m\delta$. Hence, the dynamics of the Cushion are equivalent to a leveraged buy-and-hold investment in the index ratio, with a leverage factor equal to m.

[1] We refer as the "standard CPPI" to the particular case of the portfolio insurance strategy allocating wealth between the riskless asset and a risky asset and we use the term "Dynamic Core-Satellite portfolio" (DCS) to refer to the general case of Portfolio Insurance with a stochastic reserve or "Core" asset and a performance-seeking or "Satellite" asset.

Example:

In order to develop some intuition about the impact of trends and volatility in (leveraged) buy-and-hold investments, consider the simplest case of a return series composed of two observations: $r = \{r_1, r_2\}$. Let, r_i take one of two values: u (up), for a positive return and $-d$ (down) for a negative one. An upward (downward) trend exists when $r_1 = r_2 = u$ ($r_1 = r_2 = -d$) and a "volatile" return reversal when $r_1 \neq r_2$. The compounded return of a buy-and-hold (BH) investment with leverage factor m presents the following properties (the same properties hold for a BH investment with no leverage, i.e. $m = 1$):

- *Upward trend: The total return is greater than the sum*

$$r^{1:2} = (1 + mu)^2 - 1 = 2um + (mu)^2$$

- *Downward trend: The total return is less negative than the sum*

$$r^{1:2} = (1 - md)^2 - 1 = -2dm + (md)^2$$

- *"Volatile" returns: The total return is lower (or more negative) than the sum*

$$r^{1:2} = (1 + mu)(1 - md) - 1 = mu - md - udm^2$$

Hence, due to the compounding effect of returns, assets with stronger trends and/or lower "volatility", are better buy-and-hold investments, everything else equal. This effects are magnified with leveraged.

In general, for a return series of N observations the compounded return of a buy-and-hold (BH) investment is equal to $r^{1:N} = \prod_{i=1}^{N}(1 + r_i) - 1$, which is commonly expressed in terms of the geometric average,

$$G = \left(\prod_{i=1}^{N}(1 + r_i)\right)^{\frac{1}{N}} - 1.$$

The following is a well known approximation[1] for the geometric average that

[1] Markowitz (1959) proposes a similar approximation that has very close accuracy according

relates it to the arithmetic average and the variance of returns (see Bernstein and Wilkinson (1997)):

$$G \approx A - \frac{1}{2}\frac{\sigma^2}{(1+A)},$$

where A and σ^2 denote arithmetic average and variance of returns. Since A is small compared to 1, the approximation is sometimes done as (see Booth and Fama (1992))

$$G \approx A - \frac{1}{2}\sigma^2. \tag{3.1}$$

Using the properties of the variance and the arithmetic average, it is straightforward to show that for leveraged investments the following similar relationship holds,

$$G^m \approx mA - \frac{1}{2}\frac{m^2\sigma^2}{(1+mA)}. \tag{3.2}$$

Since the arithmetic average is a measure of "trend", this formula confirms the intuition shown in the two-observations example above, for the general case with multiple observations: there is a volatility cost for holding an asset that is magnified by any present leverage. Conversely, positive "trends" have a positive impact on compounded returns but in a more than proportional amount (see second term in the right hand side of Equation (3.2) for which the trend also diminishes the volatility costs) and its effect is also magnified by leverage. Similarly, negative trends have less impact on the total return for a lower level of volatility, every-

to unreported Monte-Carlo simulations of normally distributed random variables. The approximation in Markowitz is

$$G \approx (1+A)\exp\left(-\frac{1}{2}\frac{\sigma^2}{(1+A)^2}\right) - 1.$$

thing else equal. In other words, there is a tension between the "trend" and the "volatility" in the compounded return of an investment. This tension is affected by leverage, which gives proportionally more weight to the volatility cost than to trend gains.

For a buy and hold strategy in a portfolio with two assets, the intuition follows through. A buy and hold strategy allocates proportionally more wealth, with respect to the previous period, to the asset presenting the highest return over the latest period. If in the next period it occurs a (relative) return reversal, the strategy would have allocated more wealth to the (relative) loser asset. Conversely, if no reversal occurs but returns stay on the current trend (the latest winner asset continues to be the winner asset), the BH portfolio allocation to the winner asset would have been relatively higher, thus being a "winner strategy" in this market configuration.

Another strategy of interest are the kind that rebalances (back) to a fixed set of weights. A strategy aiming to keep a fixed proportion of the two assets, needs to rebalance frequently as prices move. In order to keep a fixed-mix of weights, the strategy needs to sell the latest relative winner asset and buy the relative loser. For this reason, this type of strategies, namely unleveraged fixed-mix portfolios present a "buy-low and sell-high" behavior when return reversals occur (winner strategy) and a "buy-high and sell-low" one (loser strategy) in return trends, which are the opposite effects of trends and reversals observed in simple and leveraged buy-and-hold strategies.

3.2.2 On Portfolio Rebalancing and the Growth Rate

This section introduces the concept of the growth rate through the analysis of two types of portfolio defined according to their rebalancing policy: buy-and-hold and fixed-mix portfolios. We look at buy-and-hold portfolios because their behavior resembles to that of portfolio insurance strategies. We are also concerned with the growth rate of fixed-mix strategies because we use it later on for deriving the growth rate of portfolio insurance.

In what follows, we summarize some previous results about the growth rate of portfolios defined by a vector of weights π, assigned to a constant set of multiple assets. Throughout this analysis we assume continuous rebalancing and abstract from transaction costs.

Consider the simple Black-Scholes model for the return of any risky asset A driven by a brownian motion W as follows

$$dA(t)/A(t) = \mu dt + \sigma dW(t)$$

which has an explicit solution for the asset price given by

$$A(t) = A(0)e^{\gamma t + \sigma W(t)} \tag{3.3}$$

where $W(t) = \sqrt{t}z(t)$ for $z \sim \mathcal{N}(0, 1)$. The term γ is called the "growth rate" of A because, for a long horizon, it is equal to the continuously compounded rate of return of the asset:

$$\gamma = \frac{1}{t} \ln\left(\frac{A(t)}{A_0}\right) \text{ as } t \to \infty. \tag{3.4}$$

When $t \to \infty$, the brownian motion term disappears because $\frac{\sigma W(t)}{t} = \frac{\sqrt{t}\sigma z(t)}{t} \to 0$. Hence, the growth rate is the continuously compounded rate of return but also the continuous time version of the geometric return average, because

$$\gamma = \frac{1}{t} E\left[\ln\left(\frac{A(t)}{A_0}\right)\right]. \tag{3.5}$$

In fact, for assets following a geometric brownian motion, the growth rate, γ and the asset's drift μ (also called trend or mean return) are related as follows:

$$\gamma = \mu - \frac{1}{2}\sigma^2. \tag{3.6}$$

Interestingly, Equation (3.6) is equivalent to the relationship between the geometric and arithmetic return averages of Equation (3.1), derived in discrete time.

Fernholz and Shay (1982) show that the growth rate of a fixed-mix portfolio composed by n securities is given by

$$\gamma_\pi^{FM} = \sum_i^n \pi_i \gamma_i + \gamma_\pi^* \tag{3.7}$$

where,

$$\gamma_\pi^* = \frac{1}{2}\left(\sum_i^n \pi_i \sigma_i^2 - \sum_i^n \sum_j^n \pi_i \pi_j \rho_{ij} \sigma_i \sigma_j \right).$$

The first term in Equation (3.7) is the weighted average of the growth rates of the component assets, and the second term, γ_π^* is called the excess growth rate. Fernholz and Shay (1982) find that for **unleveraged** fixed-mix portfolios (i.e. $0 \leq \pi \leq 1$), the excess growth rate is always positive. This quantity is higher for higher standard deviations of the individual assets and for relatively lower or negative correlations. The intuition for the sign of the excess growth rate of unleveraged fixed-mix portfolios comes precisely from their "buy-low sell-high" behavior in the presence of return reversals[1].

On the other hand, the value process of a Buy and Hold portfolio is given by

$$V_\pi^{BH}(t) = \sum_i^n \pi_i(0) A_i(t)$$

[1]This effect is sometimes called "volatility pumping".

where $\pi_i(0)$ is the proportion[1] invested in asset i a time $t = 0$ and the growth rate of the BH portfolio is:

$$\gamma_\pi^{BH} = \frac{1}{t} \ln \left(\sum_i^n \pi_i(0) \left(\frac{A_i(t)}{A_i(0)} \right) \right) \text{ as } t \to \infty. \tag{3.8}$$

From Jensen's inequality and the concavity of the log the following inequality holds,

$$\sum_i^n \pi_i \gamma_i = \frac{1}{t} \sum_i^n \pi_i \ln \left(\frac{A_i(t)}{A_i(0)} \right) \leq \frac{1}{t} \ln \sum_i^n \pi_i \left(\frac{A_i(t)}{A_i(0)} \right) \quad \forall \ 0 \leq \pi \leq 1. \tag{3.9}$$

If the growth rates of the assets in the portfolio are equal then the inequality becomes an equality, making the (unleveraged) FM portfolio superior to the BH one (this is because there is no excess growth rate term on the BH's growth rate). Conversely, this implies that the BH portfolio presents an increase in return relative to the FM one for higher differences in the growth rates of the component assets (hence in *trends* of the relative performance of the assets).

In discrete time, using the approximation in Equation (3.1), Bernstein and Wilkinson (1997) decompose the geometric average of a fixed-mix portfolio and find the same expression of Fernholz's excess growth rate and call it instead the "diversification bonus". Noting that the weighted average of the geometric mean of the components is not exactly equal to the geometric average return of the buy-and-hold portfolio, Bernstein and Wilkinson (1997) define a related quantity called the "rebalancing bonus", which is the difference between the geometric return of a fixed-mix portfolio, G_π^{FM} and the one of the buy and hold portfolio, G_π^{BH}. The

[1] The weights of the BH portfolio vary with time as follows:

$$\pi_i^{BH}(t) = \frac{\pi_i(0) A_i(t)}{\sum_i^n \pi_i(0) A_i(t)}.$$

latter difference can be approximated as

$$G_\pi^{FM}(t) - G_\pi^{BH}(t) \approx \sum_i^n \pi_i(1+G_i) - \left(\sum_i^n \pi_i(1+G_i)^t\right)^{1/t} + \gamma_\pi^*. \tag{3.10}$$

For unleveraged fixed-mix portfolios Bernstein and Wilkinson (1997) find that the first two terms of the right hand side of (3.10) always give a negative contribution, hence favoring the buy and hold portfolio. This term is large when the differences between the geometric returns G_i across assets are large, and vanishes when they are the same. This confirms the intuition that trends in (relative) returns have a positive impact of BH investments and a negative one on unleveraged FM strategies for the general case with multiple assets.

The geometric average and the value process are intrinsically related, since the former is a trivial transformation of the latter[1]. Assuming a geometric brownian motion for each of the n risky assets, A_i for $1 \leq i \leq n$, the value process of a fixed-mix portfolio can be expressed in terms of their model parameters, i.e. expected returns, volatilities and correlations and the portfolio weights as follows (see for instance Wise (1996), A.3 for the following formula):

$$V_\pi^{FM}(t) = V_0 e^{\gamma_\pi^* t} \prod_{i=1}^n \left(\frac{A_i(t)}{A_i(0)}\right)^{\pi_i}. \tag{3.11}$$

[1]The relationship between the value and the geometric average is given by

$$V(t) = V_0(G+1)^t.$$

In continuous time we have with continuous compounding growth rate, γ the relationship is

$$V(t) = V_0 e^{\gamma t}.$$

Replacing (3.3) in (3.11) we get,

$$
\begin{aligned}
V_\pi^{FM}(t) &= V_0 \exp\left(\gamma_\pi^* t + \sum_{i=1}^{n} \pi_i \gamma_i t + \pi_i \sigma_i W_i(t)\right) & (3.12)\\
&= V_0 \exp\left(\gamma_\pi^{FM} t + \sum_{i=1}^{n} \pi_i \sigma_i W_i(t)\right). & (3.13)
\end{aligned}
$$

Since the weighted average of brownian motions is a brownian motion, Equation (3.13) matches the definition of the growth rate of Equation (3.3) for the portfolio value process. In Equation (3.12) the portfolio's value is written instead as a function of the weighted average of the growth rates of the component assets and its excess growth rate.

3.2.3 Portfolio Insurance's Growth Rate and its Components

Constant Proportion Portfolio Insurance (CPPI) asset allocation strategies split the portfolio between a reserve asset (R) and a performance seeking asset (S), with a dynamic allocation to the risky asset defined by the product of the available risk budget or Cushion (C) at time t and a constant multiplier, m. The Cushion is the difference between the current value of the portfolio and the level of the Floor value to be guaranteed, so the value of the portfolio is

$$
V_m^{PI}(t) = F_t + C_t. \tag{3.14}
$$

Assuming perfect correlation between the floor and the reserve asset, we show below that the Cushion can be interpreted as a leveraged fixed-mix portfolio that allocates $m\%$ to the Satellite and $(1-m)\%$ to the reserve asset. Consider the dynamics of the Cushion:

$$
\begin{aligned}
dC_t &= d(V_t - F_t) \\
&= dV_t - dF_t \\
&= V_t \left(\frac{mC_t}{V_t} \frac{dS_t}{S_t} + \left(1 - \frac{mC_t}{V_t}\right) \frac{dR_t}{R_t} \right) - F_t \frac{dR_t}{R_t} \\
&= mC_t \frac{dS_t}{S_t} + (1-m)\, C_t \frac{dR_t}{R_t} \\
&= C_t \left(m \frac{dS_t}{S_t} + (1-m) \frac{dR_t}{R_t} \right).
\end{aligned}
$$

Perold and Sharpe (1995) performed a graphical analysis of the payoff of standard CPPI portfolios allocating wealth between the riskless asset and a risky asset, as a function of the value of the risky asset. Using the fact that the Cushion process can be interpreted as a leveraged fixed-mix portfolio and Equation (3.11), we can generalize their analysis for the case with a stochastically varying reserve asset. This introduces a new feature to the analysis: the role of the assets' correlation[1]. For the two asset case the expression for the fixed mix portfolio value (3.11) simplifies to

$$
V_\pi^{FM}(t) = V_0 \left(\frac{S(t)}{S(0)} \right)^{\pi} \left(\frac{R(t)}{R(0)} \right)^{1-\pi} e^{\frac{1}{2}\pi(1-\pi)(\sigma_S^2 + \sigma_R^2 - 2\rho\sigma_S\sigma_R)t}. \tag{3.15}
$$

Using the fact that the Cushion dynamics are equivalent to a leveraged fixed mix portfolio, we write the value of the portfolio insurance strategy, (3.14) as,

$$
V_m^{PI}(t) = kV_0 \left(\frac{R_t}{R_0} \right) + (1-k)V_0 \left(\frac{S_t}{S_0} \right)^{m} \left(\frac{R_t}{R_0} \right)^{1-m} e^{\gamma_m^* t}, \tag{3.16}
$$

[1]Another way to perform Perold and Sharpe (1995) graphical analysis' generalization with a stochastic reserve asset, is to do it in two dimensions using the change of numeraire proposed in Black and Perold (1992). However, this approach would not shed light on the role that correlation plays in the strategy, since the value of the portfolio would be expressed in terms of the index ratio instead of the component assets' value.

where the Cushion's excess growth rate $\gamma_m^* = \frac{1}{2}m(1-m)(\sigma_S^2 + \sigma_R^2 - 2\rho\sigma_S\sigma_R)$ is in fact negative for $m > 1$, thus we call it a "rebalancing drag"[1]. Hence, contrary to the case of unleveraged fixed-mix strategies, volatility of component assets has a perversive effect on leveraged fixed-mix and portfolio insurance strategies, which confirms the intuition of Section 3.2.1. More interestingly, the diversification benefits that unleveraged fixed-mix portfolios experience in the presence of low or negative correlation between assets is also reversed for portfolio insurance strategies (and for leveraged fixed-mix ones). A positive correlation decreases the rebalancing drag, having a positive effect on compounded returns. The intuition for this is that, a higher correlation induces fewer relative return reversals.

Former studies on the properties of portfolio insurance strategies largely disregard the role that the correlation between the reserve asset and the performance seeking asset can have in the performance of the strategy, because they tend to focus on the particular case with Cash as the reserve asset. When the reserve asset is not equal to the riskless one, there is a correlation between the two components of the portfolio insurance strategy. Although Black and Perold (1992) also study on the properties of portfolio insurance in the case with a reserve asset different from the theoretical riskless asset with constant risk-free rate ,they write the portfolio's value in terms of the reserve asset value, instead of dollar terms. This change of numeraire yields a strategy in which the reserve asset is again the riskless asset and the Satellite is an an artificial value process called the index ratio, which is the value of the risky asset divided by the reserve asset's value. Our approach instead expresses the value of the portfolio in dollar terms, which illustrates the impact of correlation in the portfolio's value much clearly.

Using Equations (3.15) and (3.16), we illustrate the important impact that

[1] In the particular case with the riskless asset in the Core, this expression simplifies and depends only on the volatility of the Satellite. Hence Black and Perold (1992) calls a discrete-time equivalent of this term the "volatility cost".

correlation may have in a portfolio insurance strategy value. Figure F.1 draws the portfolio's value corresponding to four correlation levels, i.e. $\rho = \{-0.5, 0, 0.5, 0.75\}$ and different combinations of values for $S = [50, 200]$ and $R = [80, 120]$ after 5 years (i.e. for a starting value of 100 dollars these values are equivalent to $[-13\%, 15\%]$ and $[-4\%, 4\%]$ return per annum respectively) with volatilities, $\sigma_S = 0.15$ and $\sigma_R = 0.05$. The black surface draws the end of period value of a CPPI strategy with $m = 4$ and $k = 0.9$ as given by Equation (3.16). The red surface draws the end of period value of a Fixed-Mix Strategy as given by Equation (3.15) with the same initial allocation: $m(1-k) = \pi = 0.4$.

As observed in Figure F.1, for a starting value of 100 dollars, the maximum possible value attained by the CPPI for $\rho = -0.5$ is of 189.9 dollars which is equivalent to 13.7% return per annum, while for $\rho = 0.75$, the CPPI reaches a value of 278.9, or 22.8% return per annum, everything else equal.

On the other hand, this particular FM strategy presents a much more moderate change with correlation: For a starting value of 100 dollars, the maximum possible value attained by the FM strategy for $\rho = 0.75$ is 148.4 or 8.2% per annum, while for $\rho = -0.5$, the FM reaches a value of 150.1, or 8.5% per annum, everything else equal. Thus, the important gain obtained by the CPPI strategy with respect to the FM strategy, observed in Figure F.1 (area of black surface over red surface), as correlation increases comes mostly from the the increase in value of the CPPI strategy and not from the FM's value decrease.

Using expression (3.12), the portfolio insurance value (3.16) can be written in terms of the assets' growth rates and covariances as follows (see appendix C.1),

$$V_m^{PI}(t) = F_0 e^{\gamma_R t + \sigma_R W_R(t)} + C_0 e^{(\gamma_m^* + m\gamma_S + (1-m)\gamma_R)t + m\sigma_S W_S(t) + (1-m)\sigma_R W_R(t)}. \quad (3.17)$$

In Appendix C.1 we show that the growth rate of the PI strategy can be approximated by the weighted average of the Floor and Cushion's growth rates with weights k and $1-k$ respectively and that the growth rate of the Cushion, $\gamma_m^{cushion}$ can also be written in terms of the assets' drifts and covariances:

$$\gamma_m^{cushion} = \gamma_m^* + m(\gamma_S - \gamma_R) + \gamma_R \tag{3.18}$$

$$\gamma_m^{cushion} = \gamma_m^* + m(\mu_S - \mu_R) - \frac{1}{2}m(\sigma_S^2 - \sigma_R^2) + \gamma_R \tag{3.19}$$

$$\gamma_m^{cushion} = m(\mu_S - \mu_R) + \mu_R - \frac{1}{2}(m-1)^2\sigma_R^2 - \frac{1}{2}m^2\sigma_S^2 + m(m-1)\rho\sigma_S\sigma_R. \tag{3.20}$$

The term $m(\mu_S - \mu_R)$ in Equation (3.19) illustrates the impact that the differences in assets' drifts (trends) have on the value of leveraged FM and portfolio insurance strategies: the higher the expected outperformance of the Satellite with respect to the Core asset, the higher the expected return of the portfolio (this effect is also reversed with respect to the case of unleveraged FM strategies). Conversely, the term $-\frac{1}{2}m(\sigma_S^2 - \sigma_R^2)$ indicates that having a Satellite asset with a much higher volatility than the Core asset is not desirable, everything else equal[1].

From Equation (3.20) it is possible to infer the relative importance of the different assets characteristics on the Cushion's growth rate. In order of relative importance per unit of each term, for $m > 1$ we have:

- Core and Satellite covariance: $\rho\sigma_S\sigma_R$ (coefficient: $+m(m-1)$)

[1] In fact, for a Core and a Satellite with the same volatility (i.e. $\sigma_S = \sigma_R = \sigma$), the growth rate simplifies to:

$$\gamma_m^{cushion} = m(\mu_S - \mu_R) + \mu_R - \left(m(m-1)(1-\rho) + \frac{1}{2}\right)\sigma^2.$$

The coefficient of σ^2 is always negative for $m > 1$ and hence is a rebalancing drag of volatility which is likely to be important unless there is a high correlation between the two assets.

- Satellite's variance: σ_S^2 (coefficient: $-\frac{1}{2}m^2$)
- Core's variance: σ_R^2 (coefficient: $-\frac{1}{2}(m-1)^2$)
- Expected overperformance: $\mu_S - \mu_R$ (coefficient: $+m$)
- Core's drift (trend): μ_R (coefficient: $+1$)

The order of importance presented above is indicative in particular of the Cushion growth rate's sensitivity to *changes* in the values of each of these assets' characteristics. However, the relative importance of these characteristics in the Cushion' growth rate not only depends on their coefficient but also on the actual values that the asset's parameters may take.

In this section we uncovered the fact that the value of the portfolio insurance strategy increases with the correlation between its assets. This implies that, a potentially interesting alternative to construct a PI strategy is to use one single asset or portfolio with "good" properties in terms of its growth rate parameters as the Core and a leveraged investment in the same asset as the Satellite. In this particular case, the rebalancing drag is minimized for a given level of volatility because $\rho = 1$. We discuss this alternative in Appendix C.3.

We now turn to verify the relationship of the growth rate with the long term value of the portfolio insurance strategy with real data and short-selling constraints.

3.3 Empirical Test of the Growth Rate

The theoretical characterization presented above draws a one-to-one relationship between the value of the portfolio insurance Cushion's growth rate for a given Core or Reserve asset.

The relationship of the growth rate and the portfolio's value that should hold for long enough horizons was derived under the assumptions of continuous rebalancing, unlimited borrowing or short-selling (leverage) and a simple geometric brownian motion model for the dynamics of the component assets. In this section we test to what extend these relationships hold in applications with real data and real world constraints.

We use monthly returns of the 13 EDHEC-Risk Alternative Investment indices[1] over the period January 1997 to March 2011. In the tables of results, we use the following acronyms for the strategy indices:

- (ConvArb) Convertible Arbitrage
- (CTAs) CTA Global
- (Distress) Distressed Securities
- (EM) Emerging Markets
- (MNeutral) Market Neutral
- (EventD) Event Driven
- (FixArb) Fixed-Income Arbitrage
- (GMacro) Global Macro
- (LSequity) Long-Short Equity
- (MergArb) Merger Arbitrage
- (RelVal) Relative Value
- (ShortS) Short Selling
- (FoF) Fund of Funds

The summary statistics of this set of assets are presented in Table F.2. These indices constitute a diverse set in which correlations across indices range from -76% to 93%, average returns between 2.4% and 10.4% and volatilities between 3% and 18.5%.

In order to verify if the theoretical relationship between growth rate and portfolio's value holds in real data with discrete rebalancing and allocation constraints, we compare the ranking of candidate assets according to their growth rate and the

[1]The data and a complete description of it is available in EDHEC-Risk's website, www.edhec-risk.com.

actual end of period value of the portfolio insurance strategy implemented with monthly data and limits to the Satellite exposure of $[0, 1]$ (i.e. short-selling and leverage are not allowed).

We also consider the ranking of assets according to the "semi-theoretical" estimation of the value of the portfolio insurance given by Equation (3.16), which uses the actual end of period value of the assets considered. This formula assumes continuous rebalancing but it does not assume a geometric brownian motion for the dynamics of the assets[1]. The estimation of (3.16) is denoted V_m while the actual implementation with exposure limits and discrete rebalancing of the more General Portfolio Insurance strategy (or Dynamic Core-Satellite strategy) is denoted DCS_m.

One should notice that, if the relationship is verified, the growth rate implied by parameters' estimates for a set of candidate assets would constitute a selection criteria to choose among candidate assets for the strategy. Of course, having good estimates of this parameters would be crucial for this purpose.

3.3.1 Growth Rate ranking with a riskless Core

First consider the simple case with the riskless asset in the Core, for which there is no correlation between the two components of the strategy. In order to verify if the one-to-one relationship of the growth rate and the value of the portfolio holds, we estimate the Cushion's growth rate $\gamma_m^{cushion}$ for each of the the 13 indices and compare it with their ranking according to the end of period value of the portfolio as given by V_m and DCS_m.

In this case with a single risky asset we also look at the raking implied by the

[1]The estimation of the excess growth rate, in this case called rebalancing drag, does not need the assumption of a geometric brownian motion, but it uses a more general setting for the dynamics of the component assets. In fact Bernstein and Wilkinson (1997) derives the same expression in discrete time. See Fernholz (2002) for details on the derivation of the excess growth rate.

leveraged geometric average G^m provided in Equation (3.2). In order to have a reference of the order of magnitude of the number of matches in the ranking among these criteria, we also look at the ranking of assets as implied by their simple geometric average return (without leverage) and the returns autocorrelation[1].

We consider the theoretical value of the portfolio insurance strategy given by Equation (3.16) because it does not depend on the brownian motion assumption. Thus, it is not just a transformation of the formula with the growth rate (Equation (C.2)). The difference in the ranking between V_m and DCS_m can help us infer wether the differences observed in the ranking of the growth rate and the value of actual implementation DCS_m, come from the geometric brownian motion assumption or from the unlimited leverage constraint.

In the case with a risky free asset (B) with constant rate r with dynamics

$$B(t) = B_0 e^{rt} \tag{3.21}$$

and thus with $\sigma_R = 0$, Equation (3.16) simplifies to

$$V_m^{PI}(t) = kV_0 e^{rt} + (1-k)V_0 \left(\frac{S_t}{S_0}\right)^m e^{(1-m)(r+\frac{1}{2}m\sigma_S^2)t}. \tag{3.22}$$

Assuming an interest rate of $r = 3.5\%$ and using the sample estimate of the volatility of the risky asset we can get the theoretical value of the portfolio at the end of the sample period. Equation (3.22) assumes continuous trading and no short-selling limits. Hence, we also estimate the value of the DCS portfolio while rebalancing after a δ/m move in the exposure to the Satellite[2] with $\delta = 5\%$ and with short-selling constraints (Satellite exposure limits are $0 \leq$e≤ 1). Table F.3 presents the raking of candidate assets according to the different criteria for three

[1]The autocorrelation is an alternative measure of "trendiness". The autocorrelation coefficient is not affected by leverage (just as correlation between two series).

[2]This is equivalent to a δ move in the Index Ratio (see Black and Perold (1992)).

different multiplier values: $m = \{2, 3, 5\}$ and $k = 90\%$ (the latter does not affect the raking among different methodologies).

As we can see from Table F.3, the percentage of assets that has exactly the same ranking according to the Cushion's growth rate and the end-of-period-value of the actual implementation of the portfolio, is 100% (13 out of 13) for $m = 2$. For $m = 3$ the percentage is 85% (11 out of 13) and of 46% (6 out of 13) for $m = 5$. In this latter case, the deviations in the place of the assets are not dramatic (for instance GMacro and LSequity simply swap their place).

The former numbers contrast with the percentage of matches of the DCS_m criterion with the ranking of the simple geometric return average and the auto-correlation coefficient, which ranges from 0% to 53%, suggesting that these two alternative measures of return and trendiness are not as good estimates of the suitability of an asset to enter in the portfolio insurance strategy as the growth rate.

Furthermore, the percentage of assets with the same ranking according to all: the growth rate, the value V_m and the leveraged geometric mean is 100% in all cases. This implies that the differences with the actual implementation of the DCS is coming from the fact that the exposure limits were reached during a significant part of the period and not from the model assumptions on the dynamics of the assets used to derive the growth rate formula.

3.3.2 Ranking with stochastic Core and Satellite

In the previous example we chose to put in the Core the riskless asset. In general, the Core might be an asset that varies in time stochastically, being for instance a portfolio hedging a stream of future consumption needs or replicating a stochastic benchmark. We now turn to verify the relationship between the Cushion's growth rate and the end of period value of the portfolio using the same ranking methodology but using the Fixed-Income Arbitrage index as the Core asset of the portfolio.

Table F.4 presents the results of the ranking of the different Candidate Assets, using the geometric return average, the Cushion growth rate, $\gamma_m^{cushion}$, the value of the CPPI using the estimated rebalancing drag and the actual value of the DCS rebalanced after a δ/m move in the exposure to the Satellite with $\delta = 5\%$ and exposure limited to $[0, 1]$, for $m = \{2, 3, 5\}$.

The percentage of indices with the same ranking according to the growth rate and the actual end of period value of the portfolio with constraints and discrete rebalancing are 100% (12 out of 12 indices), 83% (10 out of 12) and 58% (7 out of 12) for $m = \{2, 3, 5\}$ respectively.

Similar to the previous ranking exercise, the percentage of assets with the same ranking according to the Cushion's growth rate and with the theoretical value V_m is 100% in all cases. This implies that the differences with the actual implementation of the DCS mostly comes from the portfolio' allocation constraints and not from the model assumptions used to derive the growth rate formula. The results also imply that the exposure limits are more likely to be reached for higher values of the multiplier.

The results of this raking exercise confirm that, in spite of the assumptions used to derive the theoretical value of the growth rate, in the absence of parameter estimation error, its one-to-one relationship with the portfolio's value holds in most cases. Hence, if one would have good estimates of the expected volatilities, correlations and over-performance of the a set of candidate assets with respect to the Core, one could rank the candidate assets to be chosen as the Satellite using the cushions' growth rate $\gamma_m^{cushion}$.

3.4 Introducing the Growth-Optimal Portfolio Insurance

A strand of the literature on portfolio selection argues in favor of "growth optimal portfolios" (GOP) (see for instance Markowitz (1976), Breiman (1961), Long et al. (1990)). GOP portfolios maximize expected log-utility of terminal wealth, the expected geometric return average and have the interesting theoretical property of outperforming all other strictly positive portfolios in long enough horizons (see Kelly (1956) and Platen and Heath (2006)).

Motivated by the absence of a risk dimension in the GOP, we propose instead to maximize the growth rate of the portfolio subject to the constraint of preventing its value to fall below a given fraction of the value of a benchmark (this fraction is chosen either by the investor or imposed by regulatory constraints). This is equivalent to maximizing the growth rate of the classic CPPI strategy.

Furthermore, in former sections we document the close relationship between the long term value of portfolio insurance strategies and their growth rate. This result suggests as well that a reasonable objective for a long term investor with short-term constraints is to maximize the growth rate of the portfolio insurance strategy. Given a pair of assets with its corresponding parameter values and the fraction to be the guaranteed k, the multiplier should be chosen such that it maximizes the portfolio's growth rate. Since the Floor process does not depend on the multiplier, maximizing the portfolios' growth rate with respect to the multiplier is equivalent to maximize the growth rate of the Cushion (see Appendix C.2 for details). Taking the partial derivative of the Cushion's growth rate (in Equation (3.18)) with respect to m equating to zero and solving for m yields the *growth optimal* multiplier:

$$m^* = \frac{\gamma_S - \gamma_R + \gamma^*}{2\gamma^*} \tag{3.23}$$

where $\gamma^* = \frac{1}{2}(\sigma_S^2 + \sigma_R^2 - 2\rho\sigma_S\sigma_R)$. Replacing (3.23) in Equation (3.18) yields the optimal Cushion's growth rate value as a function of the component asset's parameters:

$$\gamma^{cushion}(m^*) = \frac{(\gamma_S - \gamma_R + \gamma^*)^2}{2\gamma^*} - \frac{(\gamma_S - \gamma_R + \gamma^*)^3}{4(\gamma^*)^2} + \gamma_R.$$

In order to illustrate the importance of the choice of the multiplier value, Figure F.2 plots the Cushion's growth rate as given by Equation (3.18) for the following parameter values: $\mu_S = 0.08$, $\sigma_S = 0.15$, $\mu_R = 0.03$, $\sigma_R = 0.05$, all possible values for the correlation coefficient, i.e. $\rho = [-1, 1]$ and different multiplier values, i.e. $m = [1, 10]$.

Figure F.2 illustrates how m^* maximizes the Cushion's growth rate due to its concavity[1] with respect to m. This figure also illustrates two interesting interactions between the correlation and the multiplier: i) the gray surface indicates that for uncorrelated or negatively correlated assets, the choice of the multiplier becomes critical in determining the Cushion's growth rate, ceteris paribus and ii) the optimal multiplier (dark line) increases with the correlation between the two assets, everything else being equal. Similarly, Figure F.3 illustrates that for highly volatile Satellite assets, the choice of the multiplier becomes critical (gray surface) and the growth-optimal multiplier decreases with the Satellite's volatility (dark line), everything else equal.

In previous sections the parameters determining the dynamics of the compo-

[1]The value of portfolio insurance strategies is convex with respect to the value of the "index ratio" or the outperformance of the Satellite with respect to the Core asset, precisely because of the leverage role of the multiplier. The concavity with respect to "m" in this case is of course for a given level of performance of the component assets (ceteris paribus condition).

nent assets where assumed to be known and constant. This assumption is mild if the purpose is to perform an ex-post performance attribution exercise using the values of the parameters estimated using the sample period in question. However, it is well documented that expected returns, volatilities and correlations for most assets present significant variations in time. If any of these values varies in time, in order to maintain the optimality condition the multiplier should be adjusted accordingly:

$$m_t^* = \frac{\gamma_S(t) - \gamma_R(t) + \gamma^*(t)}{2\gamma^*(t)}. \tag{3.24}$$

In the particular case with a riskless asset as the reserve asset with instantaneous risk-free rate $\mu_R = r$, with $\sigma_R = 0$ Equation (3.24) yields

$$m_t^* = \frac{\mu_S(t) - r}{\sigma_S^2(t)}. \tag{3.25}$$

Using the standard utility maximization setting, Merton (1971), Grossman and Vila (1992) and Basak (2002) find a remarkably similar solution for an optimal multiplier in the case of a portfolio insurer investor with CRRA preferences[1]. Merton (1971), Grossman and Vila (1992) and Basak (2002)'s optimal multiplier is equal to (3.25) times the inverse of the investor's risk aversion parameter. From a practical standpoint, the problem with this approach is precisely to select the risk aversion parameter value, for which there is not general consensus.

On the other hand, the growth-optimal multiplier does not need a risk aversion parameter, hence there is no ambiguity about its optimal (unique) value. However, it is possible to integrate the risk aversion parameter within the growth optimal

[1] In Basak (2002) solves in the presence of inter-temporal consumption and includes the floor violation restriction embedded in the investor's utility function. The investor's marginal utility smoothly converges towards infinity, as opposed to imposing an exogenous constraint in which marginal utility jumps in a discrete way when wealth hits the floor.

portfolio insurance approach by defining the objective of maintaining a target level of growth rate. We discuss this alternative in Appendix C.4.

3.4.1 Maximum Multiplier

The multiplier, m is a key parameter that determines the behavior of the portfolio insurance strategy. Perhaps the most common way to determine the multiplier of the CPPI in practice is using the maximum value that would allow the Cushion to remain positive even in the "worst case scenario". In general, in order to guarantee that the Cushion remains positive, the multiplier has to satisfy the following condition:

$$\frac{C_{t+1}}{C_t} = m\frac{S_{t+1}}{S_t} + (1-m)\frac{R_{t+1}}{R_t} > 0 \Leftrightarrow \frac{S_{t+1}}{S_t} > \frac{(m-1)}{m}\frac{R_{t+1}}{R_t}$$

or equivalently

$$mr_S(t,t+1) + (1-m)r_R(t,t+1) \geq -1 \tag{3.26}$$

$$m(r_S(t,t+1) - r_R(t,t+1)) \geq -(1 + r_R(t,t+1)) \tag{3.27}$$

For $(r_S(t,t+1) - r_R(t,t+1)) < 0$ the inequality (3.27) gets inverted,

$$m \leq \frac{-(1 + r_R(t,t+1))}{r_S(t,t+1) - r_R(t,t+1)}$$

where t and $t+1$ are any two portfolio rebalancing dates. Hence, the maximum value for the multiplier that could guarantee in general the Cushion's positivity condition (3.27) is:

$$m = \frac{-(1 + min(r_R))}{min(r_S - r_R)}. \tag{3.28}$$

In the particular case where the reserve asset is the riskless one, r_R has a minimum value equal to 0 (assuming positive interest rates) and its values are usually small compared to the magnitude of the returns on stocks. For this reason, condition (3.28) is very commonly approximated by:

$$m = \frac{-1}{min(r_S)}.$$

Although this value of the multiplier would respect the insurance promise of the portfolio strategy, complying with a feasibility condition is not necessarily the optimal choice from an investor's perspective. As we saw in previous sections, the rebalancing drag increases with the multiplier at a rate of $m(m-1)$. In the next section we compare the growth optimal portfolio insurance we introduced, with the equivalent portfolio insurance strategy defined by the aforementioned methodology to determine the multiplier.

3.4.2 Empirical test of the Growth-Optimal Portfolio Insurance

In order to empirically evaluate the performance of the GOPI strategy we use the standard set of assets of the CPPI. We use monthly data of the T-bill rates as the reserve asset and the value-weighted CRSP broad market stock index as the risky asset, available from Kenneth French website from January 1926 until December 2010. We first perform two Backtest exercises over long periods of time and compare the performance of the GOP or Kelly criterion[1]), the GOPI and the CPPI strategy, the latter being the standard portfolio insurance strategy

[1]For the pair of assets in question the GOP formula yields a portfolio with leverage and short selling constraints. In the GOP literature the portfolio has the properties mentioned in the introduction only if all the weights happen to be positive. Thus in this case, what we call the GOP corresponds more precisely to the Kelly criterion, according to the denomination in the academic literature.

that uses the multiplier as given by Equation (3.28). Both portfolio insurance strategies have the constraint of preserving 90% of the value of the reserve asset. The allocation of the GOP to the risky asset in the case with two assets, is equal to the multiplier of the GOPI.

In the first Backtest we use the entire sample to estimate the parameters that determine the optimal and the standard multipliers and the allocation of the GOP. Figure F.4 presents the log of the cumulative performance of the three strategies and of an investment in the reserve asset, i.e. cash. The portfolio with the highest cumulative performance during the entire sample is, unsurprisingly the GOP (dark blue line). On the other hand, the GOP does not respect the constraint that the GOPI and CPPI strategies have and it presents a significantly riskier profile that the other two strategies. In fact the GOP in this case is even more risky than the stock index itself, as illustrated in Table F.1. The upper panel of Table F.1 shows that, the minimum value ever attained by the GOP is 17 dollars, for an initial investment of 100, while the minimum value attained by the GOPI and CPPI strategies are 99.2 and 97.5 respectively. This is not surprising because the GOP or Kelly criterion in this case happens to be a leveraged investment in the stock index. The maximum drawdown of the GOP over the sample is equal to 95.6% compared to a maximum drawdown of around 50% for the GOPI and CPPI strategies.

Interestingly, we find that, although the GOPI multiplier is lower than the standard multiplier of the CPPI, the value of the GOPI strategy is higher than the value of the CPPI at the end of the sample and during most of the period, as observed in Figure F.4. The growth optimal multiplier estimated with the entire sample is equal to 1.6 while the standard multiplier of the CPPI is equal to 3.4[1].

[1]In this case we use monthly data and monthly rebalancing, hence the worst case scenario is given by the worst stock return that happened in the 1930s of almost -30%.

The CPPI strategy happens to lose almost all its risk budget at the beginning of the sample an it stays very close to its floor value during the rest of the period. On the other hand the GOPI strategy “survives” the periods of high volatility and presents an important growth in the long run, with an average geometric return of 7.9% compared to 3.6% for the standard CPPI.

As a second Backtest, we perform an out-of-sample exercise in which we split the sample in two. We use the first half of the sample (i.e., $1926:01-1968:05$) to estimate the parameters of the strategies (multipliers and allocation of the GOP) and the second half to implement the three strategies (i.e., $1968:06-2010:12$). We use a very long period of time to estimate the parameters in order to avoid the estimation sample bias mostly present in expected return estimates. The result of this second Backtest is presented in the lower panel of Table F.1. The conclusions are similar to the former Backtest. The GOP attains the highest value among the three strategies but also a minimum value of around 39.2 dollars, compared to minimum values of 100 and 99.5 dollars for the GOPI and CPPI strategies respectively and 70.3 dollars for the stock index. Although the optimal multiplier (1.7) is again lower than the standard multiplier (3.4) the growth optimal portfolio insurance strategy attains a higher value than the standard CPPI and it remains higher during most of the sample period, as illustrated in Figure F.5.

Portfolio insurance strategies are often perceived as dependent on the market conditions at the inception of the strategy. However, one should notice that the first Backtest, starting in 1926, begins in a bull market, while the second Backtest starts with a bear market.

In order to make a more throughout comparison between the GOPI and the standard CPPI, we now turn to an historical simulations exercise, where we perform N Backtest exercises for which the starting dates are selected randomly from the *available sample* period and performed over 5, 10 and 20 years investment horizons. This methodology is equivalent to perform block-bootstrap simulations using a block-size equal to the length of the investment horizon (i.e., either 5, 10 or 20 years).

The *available sample* for choosing the starting dates of the simulations excludes the multiplier's calibration sample, which is the taken as the first 10 years of data, and the last 5, 10 or 20 years of available data, respectively (the latter condition allows all Backtests to have the same length, either 5, 10 or 20 years). We compare the end of period value of the portfolio insurance strategy using the growth-optimal multiplier and the standard multiplier approach. We perform 100 random simulations for each of the three horizons considered.

We use three different methodologies to estimate the standard multiplier of Equation (3.28): *i*) first we use the whole available time series following the first 10 years of calibration sample to determine the minimum values required by its formula in (3.28), which we denote as m_{all}, *ii*) we estimate the standard multiplier that corresponds to the sample period over which each of the backtests simulations will be performed (hence it is a forward looking estimate), denoted "Perfect" m_{period} and *iii*) we use all data available before the starting date of each of the backtest scenarios to estimate the multiplier, which we denote as "Sample" $\hat{m}_{period}$.

In order to estimate the optimal multipliers, as given by Equation (3.24), we also use three different approaches: *i*) we use sample estimates over the testing period of each of the backtest scenarios (hence it is a forward looking estimate) which we denote as "Perfect" m^*, *ii*) we estimate the optimal multiplier using the latest 10 years of data available at the starting date of each simulated scenario, denoted as "Sample" $\hat{m}^*$ and *iii*) in order to estimate the "Dynamic" optimal multiplier series, $\hat{m}^*_t$, we fit[1] a Dynamic Conditional Correlation model (Engle and Sheppard (2001)) to all available stock returns at the starting date of each

[1] We use Kevin Sheppard's UCSD Multivariate Garch toolbox for estimating the DCC model.

simulated scenario and use the estimated parameters and available "innovations" at each point in time to infer the variance and covariance time series over the testing period. We use a 10 year moving average to estimate expected stock returns and take the latest available interest rate value as its next period forecast.

The standard multiplier estimated using the whole available time series after the first 10 years of calibration sample is equal to $m_{all} = 4.2$, which is higher than the standard multiplier obtained using the "Sample" estimate $\hat{m}^* = 3.4$. This latter presents always the same value (see Table F.5) because the "worst" scenario in the available history occurred during the 1930s (within the first 10 years of data). This contrast with the wide range of values for the "Perfect" foresight standard multiplier m_{period}, which falls between 4.2 and 20.4 depending on the backtest scenario.

Table F.6 presents the range and median values of the three different estimates of the growth-optimal multipliers. In order to have a better comparison between these two types of multiplier, we compute the percentage of scenarios in which the optimal multiplier is less or equal to the standard multipliers corresponding to the same periods. Table F.7 presents the results for each pair of multipliers and each horizon. In every pair, even for the perfect foresight multipliers, we find that the optimal multiplier has a lower level in most scenarios. On average across pairs and horizons the percentage of times for which the optimal multiplier is lower or equal than the standard multiplier is around 70%.

Perhaps the most interesting comparison is the percentage of scenarios on which the portfolio with the optimal multiplier obtains a higher value than the portfolio using the standard multiplier approach, presented in Table F.8. When comparing the "Perfect" foresight estimates of the optimal and standard multipliers, we find that the percentage of scenarios favoring the optimal multiplier portfolio is between 51% and 67% (for 20 and 5 years horizon respectively). This result is more dramatic when compared with the also forward-looking estimate of the standard multiplier, m_{all}, in which case the range of the percentage of scenarios favoring the optimal multiplier portfolio is 66% to 85%. The average geometric return of the portfolios across scenarios is higher for the portfolios using the optimal multiplier on the three horizons considered, as shown in Table F.9.

On the other hand the portfolios with the sample estimate of the optimal multiplier outperforms the portfolio using the sample estimate of the standard multiplier in around 40% of the scenarios. This is not surprising considering that the variations in average returns and volatility from one period to the next can be very dramatic in stock returns. These variations are expected to affect more the optimal multiplier than the standard multiplier, because the latter is simply a lower bound approximation of the return. However, using the "Dynamic" estimates of the parameter values composing the optimal multiplier favors again the corresponding portfolio with respect to the one using the "Sample" standard multiplier estimate $\hat{m}^{period}$. In this case the over performance probability for the three horizons considered are 69%, 72% and 72%, for 5, 10 and 20 horizons respectively. Furthermore, the portfolio insurance strategies constructed using the dynamic estimates of the optimal multiplier outperformed the portfolio using the forward looking standard multiplier estimate m_{all} in 64%, 62%, 67% of the scenarios, for 5, 10 and 20 horizons respectively. The superiority of the dynamic optimal multiplier portfolios is also suggested by the higher average geometric return it obtained for the three horizons considered (see Table F.9).

3.5 Conclusion

We derive the growth rate of the portfolio insurance strategy in the general case with a stochastic reserve asset as a function of the underlying assets' characteristics (i.e. expected returns, volatilities and correlation). The analytical characterization of the growth rate of portfolio insurance strategies sheds light on the largely disregarded role that the correlation between the underlying assets plays on the value of this type of strategy.

Through different empirical exercises we find that the growth rate has a very close relationship with the expected value of the portfolio insurance strategy even when common short-selling and leverage constraints are imposed. We also illustrate how can the growth rate of the strategy be decomposed to nail down the relative importance and effect of the characteristics of the underlying assets on the performance of the strategy.

Finally, we introduce the growth optimal portfolio insurance strategy (GOPI), which combines the pragmatic objective of maximizing the growth rate of the portfolio and the common risk-management objective of preserving a given fraction of the value of the portfolio typically indexed to a given benchmark. The growth rate maximization is achieved by using the growth optimal multiplier, while the investor chooses the benchmark and the fraction of its value to insure as in the standard CPPI case.

Through empirical tests with real data and in the presence of short-selling and leverage constraints, we find that the GOPI strategy tends to outperform a CPPI using the standard multiplier selection methodology while keeping a more conservative risk profile.

References

Amenc, N., P. Malaise, and L. Martellini (2004). Revisiting core-satellite investing: a dynamic model of relative risk management. *Journal of Portfolio Management 31*(1), 64–75. 79

Amenc, N., L. Martellini, V. Milhau, and V. Ziemann (2009). Asset-liability management in private wealth management. *The Journal of Portfolio Management 36*(1), 100–120. 79

Ang, A. and G. Bekaert (2002). International asset allocation with regime shifts. *Review of Financial Studies 15*(4), 1137. 49

Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance 61*, 259–299. 2, 6, 30, 39

Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2009). High idiosyncratic volatility and low returns: International and further US evidence. *Journal of Financial Economics 91*(1), 1–23. 6, 30, 42

Ashley, R. (2006). Beyond optimal forecasting. *Working Papers*. 60, 63

Bai, J. and P. Perron (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 47–78. 49, 53

Bai, J. and P. Perron (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 1–22. 49

Bali, T., N. Cakici, X. Yan, and Z. Zhang (2005). Does Idiosyncratic Risk Really Matter? *The Journal of Finance 60*(2), 905–929. 2, 4, 28, 29, 30, 31, 43

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–1509. 4, 24, 25, 26

Barberis, N., M. Huang, and T. Santos (2001). Prospect Theory and Asset Prices. *Quarterly Journal of Economics 116*(1), 1–53. 6

Barry, D. and J. Hartigan (1993). A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 309–319. 50, 54, 57, 64, 170

Basak, S. (2002). A comparative study of portfolio insurance. *Journal of Economic Dynamics and Control 26*(7-8), 1217–1241. 102

Bekaert, G., R. J. Hodrick, and X. Zhang (2008). Is There a Trend in Idiosyncratic Volatility? *SSRN eLibrary*. 2, 15, 16, 19, 21, 22, 23

Bernstein, W. and D. Wilkinson (1997). Diversification, Rebalancing, and the Geometric Mean Frontier. *Unpublished manuscript*. 83, 87, 88, 96

Black, F. (1976). The pricing of commodity contracts. *Journal of financial economics 3*(1-2), 167–179. 38

Black, F. and R. Jones (1987). Simplifying portfolio insurance. *The Journal of Portfolio Management 14*(1), 48–51. 78

Black, F. and A. Perold (1989). *Theory of constant proportion portfolio insurance*. Division of Research, Harvard Business School. 78

Black, F. and A. Perold (1992). Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics and Control 16*(3-4), 403–426. 80, 81, 90, 91, 97

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of political economy 81*(3), 637. 6

Bollerslev, T., N. Sizova, and G. Tauchen (2009). Volatility in equilibrium: Asymmetries and dynamic dependencies. Technical report, Duke University Working Paper. 4, 26

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *The Review of Financial Studies 22*(11), 4463–4492. 4, 26

Booth, D. and E. Fama (1992). Diversification returns and asset contributions. *Financial Analysts Journal*, 26–32. 83

Bossaerts, P. (2002). *The paradox of asset pricing.* Citeseer. 77

Bowley, A. (1920). Elements of Statistics. *New York: Charles Scribner's Sons*. 39

Brandt, M., A. Brav, J. Graham, and A. Kumar (2009). Forthcoming. The Idiosyncratic Volatility Puzzle: Time Trend or Speculative Episodes. *Review of Financial Studies*. 2

Breiman, L. (1961). Optimal gambling systems for favorable games. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability: held at the Statistical Laboratory, University of California, June 20-July 30, 1960*, pp. 65. Univ of California Press. 76, 100

Campbell, J. (1993). Intertemporal asset pricing without consumption data. *The American Economic Review 83*(3), 487–512. 40

Campbell, J., M. Lettau, B. Malkiel, and Y. Xu (2001). Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk. *The Journal of Finance 56*(1), 1–43. 2, 6, 15, 17, 23, 43

Campbell, J. and R. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies 1*(3), 195–228. 46, 51

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average?. *Review of Financial Studies 21*(4), 1509 – 1531. 47, 68, 70

Cao, H., T. Wang, and H. Zhang (2005). Model uncertainty, limited market participation, and asset prices. *Review of Financial Studies 18*(4), 1219. 5, 37

Christensen, M. (2005). On the history of the growth optimal portfolio. *Unpublished*, 1–65. 76

Christensen, M., P. Christensen, and N. Nielsen (2005). A thesis on the growth optimal portfolio and the theory of arbitrage pricing and portfolio selection. 76, 77

Cochrane, J. (1992). Explaining the variance of price-dividend ratios. *The Review of Financial Studies 5*(2), 243–280. 46

Cochrane, J. (2008). The dog that did not bark: A defense of return predictability. *Review of Financial Studies 21*(4), 1533. 46, 47, 51

Connor, G. (1997). Sensible return forecasting for portfolio management. *Financial Analysts Journal 53*(5), 44–51. 63, 68

Cont, R. and P. Tankov (2009). Constant proportion portfolio insurance in the presence of jumps in asset prices. *Mathematical Finance 19*(3), 379–401. 133

Cuzick, J. (1995). A strong law for weighted sums of iid random variables. *Journal of Theoretical Probability 8*(3), 625–641. 122

Dangl, T. and M. Halling (2009). Predictive Regressions with Time-Varying Coefficients. *SSRN eLibrary*. 49, 59

DiBartolomeo, D. (2006). Applications of Portfolio Variety. 3

DiDonato, A. and A. Morris Jr (1992). Algorithm 708: Significant digit computation of the incomplete beta function ratios. *ACM Transactions on Mathematical Software (TOMS) 18*(3), 360–373. 57

Engle, R. F. and K. Sheppard (2001). Theoretical and Empirical properties of Dynamic Conditional Correlation Multivariate GARCH. Working Paper 8554, National Bureau of Economic Research. 107

Erdman, C. and J. Emerson (2007). bcp: an R package for performing a Bayesian analysis of change point problems. *J. Stat. Softw 23*, 1–13. 50

Erdman, C. and J. Emerson (2008). A fast Bayesian change point analysis for the segmentation of microarray data. *Bioinformatics 24*(19), 2143. 57

Ethier, S. (2004). The kelly system maximizes median wealth. *Journal of Applied Probability 41*(4), 1230–1236. 76

Fama, E. and K. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56. 14, 43

Fernholz, E. (2002). *Stochastic portfolio theory*. Springer Verlag. 80, 96

Fernholz, R. and B. Shay (1982). Stochastic portfolio theory and stock market equilibrium. *The Journal of Finance 37*(2), 615–624. 86

Ferreira, M. and P. Santa-Clara (2010). Forecasting Stock Market Returns: The Sum of the Parts is More than the Whole. *NBER Working Paper*. 47

French, K. R., G. W. Schwert, and R. F. Stambaugh (1987, September). Expected stock returns and volatility. *Journal of Financial Economics 19*(1), 3–29. 15

Fu, F. (2009). Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics 91*(1), 24–37. 6, 30, 35

Goyal, A. and P. Santa-Clara (2003). Idiosyncratic Risk Matters! *The Journal of Finance 58*(3), 975–1008. 2, 3, 4, 5, 9, 15, 17, 27, 28, 29, 31, 43, 136

Goyal, A. and I. Welch (2003). Predicting the equity premium with dividend ratios. *Management Science 49*(5), 639–654. 46, 69, 74

Goyal, A. and I. Welch (2008). A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *Review of Financial Studies 21*(4), 1455. 46, 51, 64, 67, 68, 70, 71, 74

Grauer, R. (1981). A comparison of growth optimal and mean variance investment policies. *Journal of Financial and Quantitative Analysis 16*(01), 1–21. 78

Groeneveld, R. and G. Meeden (1984). Measuring Skewness and Kurtosis. *The Statistician 33*, 391–399. 39

Grossman, S. and J. Vila (1992). Optimal dynamic trading with leverage constraints. *Journal of Financial and Quantitative Analysis 27*(02), 151–168. 102

Guo, H. and R. Savickas (2008). Average Idiosyncratic Volatility In G7 Countries. *The Review of Financial Studies 21*(3), 1259–1296. 5, 6, 35, 36

Hakansson, N. and W. Ziemba (1995). Capital growth theory. *Handbooks in Operations Research and Management Science 9*, 65–86. 76

Hamilton, J. (1989a). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, 357–384. 49

Hamilton, J. D. (1989b). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica 57*(2), 357–384. 21

Hinkley, D. (1975). On Power Transformations to Symmetry. *Biometrika 62*, 101–111. 39

Huang, W., Q. Liu, S. Rhee, and L. Zhang (2009). Return reversals, idiosyncratic risk, and expected returns. *Review of Financial Studies*. 6, 30

Inoue, A. and L. Kilian (2004). In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use? *Econometric reviews* (4), 371. 47

Kachman, S. D. (1999). Distribution of linear transformations and quadratic forms. 124

Kapadia, N. (2009). The next Microsoft ? Skewness, idiosyncratic volatility, and expected returns. *unpublished paper*. 39

Kelly, J. (1956). A new interpretation of information rate. *Information Theory, IRE Transactions on 2*(3), 185–189. 76, 100

Kim, T. and H. White (2004). On more robust estimation of skewness and kurtosis. *Finance Research Letters 1*(1), 56–73. 39

Lacerda, F. and P. Santa-Clara (2010). Forecasting Dividend Growth to Better Predict Returns. 48

Latane, H. (1959). Criteria for choice among risky ventures. *The Journal of Political Economy 67*(2), 144–155. 76

Lettau, M. and S. Ludvigson (2005). Expected returns and expected dividend growth. *Journal of Financial Economics 76*, 583–626. 47

Lettau, M. and S. van Nieuwerburgh (2008). Reconciling the return predictability evidence. *Review of Financial Studies 21*(4), 1607. 49, 50, 51, 52, 53, 54, 64, 65, 66, 67, 68, 73, 127, 156, 157

Levy, H. (1978). Equilibrium in an Imperfect Market: A Constraint on the Number of Securities in the Portfolio. *American Economic Review 68*(4), 643–658. 5

Lewellen, J. (1999). The Time Series Relations Among Expected Return, Risk, and book-to-market. *Journal of Financial Economics 54*, 5–53. 52

Lhabitant, F. (2011). Correlation vs. trends: A common misinterpretation. 80

Long, J. et al. (1990). The numeraire portfolio. *Journal of Financial Economics 26*(1), 29–69. 76, 100

Luenberger, D. (1998). *Investment science*, Volume 494. Oxford university press New York. 76

Malkiel, B. G. and Y. Xu (2002). Idiosyncratic risk and security returns. 2, 5

Markowitz, H. (1959). Portfolio Selection: Efficient Diversification of Investments, Cowles Foundation Monograph# 16. 82

Markowitz, H. (1976). Investment for the long run: New evidence for an old rule. *The Journal of Finance 31*(5), 1273–1286. 76, 100

McCracken, M. (2007). Asymptotics for out of sample tests of Granger causality. *Journal of Econometrics 140*(2), 719–752. 69

McMillan, D. (2009). Revisiting dividend yield dynamics and returns predictability: Evidence from a time-varying ESTR model. *The Quarterly Review of Economics and Finance 49*(3), 870–883. 48, 49

Merton, R. (1971). Optimal consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*, 373–413. 102

Merton, R. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of finance*, 449–470. 6

Merton, R. (1987). A simple model of capital market equilibrium with incomplete information. *Journal of Finance 42*, 483–510. 5

Merton, R. and P. Samuelson (1974). Fallacy of the log-normal approximation to optimal portfolio decision-making over many periods* 1. *Journal of Financial Economics 1*(1), 67–94. 77

Mincer, J. and V. Zarnowitz (1969). The evaluation of economic forecasts. 69

Nelson, C. C. and M. J. Kim (1993). Predictable Stock Returns: The Role of Small Sample Bias. *Journal of Finance 43*, 641–661. 52

Newey, W. and K. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica: Journal of the Econometric Society 55*(3), 703–708. 29, 42

Ophir, T. (1978). The geometric-mean principle revisited. *Journal of Banking & Finance 2*(1), 103–107. 76

Pastor, L. and R. Stambaugh (2001). The Equity Premium and Structural Breaks. *Journal of Finance 56*(4), 1207–1239. 49

Pástor, L. and R. Stambaugh (2009). Predictive systems: Living with imperfect predictors. *The Journal of Finance 64*(4), 1583–1628. 52

Paye, B. and A. Timmermann (2003). Instability of Return Predictability Models. 49

Perold, A. (1986). Constant proportion portfolio insurance. *Harvard Business School*. 78

Perold, A. and W. Sharpe (1988). Dynamic strategies for asset allocation. *Financial Analysts Journal*, 16–27. 78

Perold, A. and W. Sharpe (1995). Dynamic strategies for asset allocation. *Financial Analysts Journal*, 149–160. 90

Pestien, V. and W. Sudderth (1985). Continuous-time red and black: how to control a diffusion to a goal. *Mathematics of Operations Research*, 599–611. 76

Pettenuzzo, D. and A. Timmermann (2005). Predictability of stock returns and asset allocation under structural breaks. *Department of Economics, University of California, San Diego*. 49

Platen, E. (2005). On the role of the growth optimal portfolio in finance. *Australian Economic Papers 44*(4), 365–388. 76, 78

Platen, E. and D. Heath (2006). *A benchmark approach to quantitative finance.* Springer Verlag. 100

Pontiff, J. (2006). Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics 42*(1-2), 35–52. 32, 35

Rapach, D., J. Strauss, and G. Zhou (2009). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies*. 47, 49

Ross, S. (1976). The arbitrage theory of capital asset pricing. *Journal of economic theory 13*(3), 341–360. 7

Samuelson, P. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia 98*(4-5), 108–113. 76

Sharpe, W. (1963). A simplified model for portfolio analysis. *Management science*, 277–293. 7

Spiegel, M. (2008). Forecasting the equity premium: Where we stand today. *Review of Financial Studies 21*(4), 1453 – 1454. 46

Stambaugh, R. F. (1986). Bias in Regressions with Lagged Stochastic Regressors. *Working Paper University of Chicago*. 52

Stambaugh, R. F. (1999). Predictive Regressions. *Journal of Financial Economics 54*, 375–421. 52

Tédongap, R. (2010). Consumption Volatility and the Cross-Section of Stock Returns. *unpublished paper, Stockholm School of Economics*. 3, 24, 25

Thorp, E. (1971). Portfolio choice and the kelly criterion. *Stochastic models in finance*, 599–619. 76

Timmermann, A. (2008). Elusive return predictability. *International Journal of Forecasting 24*(1), 1–18. 47

Wei, S. and C. Zhang (2005). Idiosyncratic Risk Does Not Matter: A Reexamination of the Relationship between Average Returns and Average Volatilities. *Journal of Banking and Finance, Vol. 29, pp. 603-621, 2005*. 2, 9, 28, 29, 30, 31, 43

Wei, S. and C. Zhang (2006). Why Did Individual Stocks Become More Volatile? *The Journal of Business 79*(1), 259–292. 15

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455. 25

Wise, A. (1996). The Investment Return From A Portfolio With A Dynamic Rebalancing Policy. *British Actuarial Journal 2*(4), 975–1001. 88

Appendices

Appendix A

Appendix Chapter 1

A.1 Proof of Proposition 1

Consider the factor model decomposition

$$r_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it}\beta_{it}F_t + \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}$$

and

$$r_{it} - r_t^{(w_t)} = \left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right) F_t + \varepsilon_{it} - \sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt}$$

Under the homogeneous betas assumption, we have

$$r_{it} - r_t^{(w_t)} = \varepsilon_{it} - \sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt} \tag{A.1}$$

and therefore

$$\left[r_{it} - r_t^{(w_t)}\right]^2 = \varepsilon_{it}^2 + \left(\sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt}\right)^2 - 2\varepsilon_{it}\sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt}$$

so that

$$\begin{aligned} CSV_t^{(w_t)} &= \sum_{i=1}^{N_t} w_{it}\left(r_{it} - r_t^{(w_t)}\right)^2 \\ &= \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 + \left(\sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt}\right)^2 - 2\sum_{i=1}^{N_t}\sum_{j=1}^{N_t} w_{jt}w_{it}\varepsilon_{it}\varepsilon_{jt} \end{aligned}$$

Noting that

$$\left(\sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt}\right)^2 = \sum_{i=1}^{N_t}\sum_{j=1}^{N_t} w_{jt}w_{it}\varepsilon_{it}\varepsilon_{jt}$$

we finally have:

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 - \left(\sum_{i=1}^{N_t} w_{jt}\varepsilon_{jt}\right)^2$$

We now argue that the term $\sum_{i=1}^{N_t} w_{jt}\varepsilon_{jt}$ converges to 0 for increasingly large numbers of stocks. To show this, we need to use a recent result by Cuzick (1995) regarding the Marcinkiewcz−Zygmund strong law of large numbers for weighted sums of i.i.d. variables:

$$\frac{1}{N}\sum_{i=1}^{N} a_{Ni}X_i \longrightarrow 0 \text{ almost surely} \tag{A.2}$$

when $\{X, X_N, N \geq 1\}$ is a sequence of i.i.d. random variables with $E(X) = 0$, $E|X| < \infty$ and $\{a_{Ni}, 1 \leq i \leq N, N \geq 1\}$ is an array of constants uniformly bounded satisfying[1]

$$\sup |a_{Nit}| < \infty. \tag{A.3}$$

Here we take $a_{Nit} \equiv N_t w_{it}$ and $X_i \equiv \varepsilon_i$. For the result (A.2) to hold a_{Nit} needs to be uniformly bounded and to satisfy condition (A.3). We therefore restrict our attention to non-trivial weighting schemes, ruling out the situation such that the index is composed by a single stock. Please note that this condition together with the fact that $\sum_i w_{it} = 1$ implies $N_t > 1$ and also restrict the weights to be (strictly) positive at every given point in time. Hence, a weighting scheme (w_t), is defined as a vector process which satisfies $0 < w_{it} < 1 \; \forall \; i, t$. This condition seams reasonable since our focus is to measure idiosyncratic risk in the market.

By definition, the weighting scheme w_{it} and a_{Nit} is uniformly bounded by N_t and the following condition holds,

$$0 < w_{it} < 1 \; \forall \; i, t \tag{A.4}$$

Multiplying by N_t, we get

$$\begin{aligned} 0 &< N_t w_{it} < N_t \\ 0 &< N_t w_{it} < \infty \\ 0 &< a_{Nit} < \infty \\ |a_{Nit}| &< \infty \; \forall \; i, t \end{aligned}$$

which implies that condition (A.3) holds. Thus, for a positive weighting scheme from the strong law of large numbers for weighted sums of i.i.d. variables, it follows

[1] See Theorem 1.1, particular case of Cuzick (1995).

that:

$$\sum_{i=1}^{N_t} w_{it}\varepsilon_{it} \underset{N_t\to\infty}{\longrightarrow} 0 \text{ a.s.},$$

Using similar arguments, and the homogeneous idiosyncratic second moment assumption, $E\left[\varepsilon_{it}^2\right] \equiv \sigma_\varepsilon^2\left(t\right)$, we obtain that for a strictly positive weighting scheme, w_t, and i.i.d. ε_i,

$$\sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 \underset{N_t\to\infty}{\longrightarrow} \sigma_\varepsilon^2\left(t\right) \text{ almost surely}$$

Using these results, we finally have that:

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it}\left(r_{it} - r_t^{(w_t)}\right)^2 \underset{N_t\to\infty}{\longrightarrow} \sigma_\varepsilon^2\left(t\right) \text{ almost surely.}$$

A.2 Properties of the CSV Estimator

A.2.1 Bias of the CSV Estimator

Under the factor model decomposition (1.1) and equation (1.2) and using the homogeneous beta assumption, we have:

$$r_{it} - r_t^{(w_t)} = \left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right) F_t + \varepsilon_{it} - \sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt} = \varepsilon_{it} - \sum_{j=1}^{N_t} w_{jt}\varepsilon_{jt} \tag{A.5}$$

Replacing result (A.5) in equation (1.3) we have as before:

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 - \sum_{i=1}^{N_t}\sum_{j=1}^{N_t} w_{jt}w_{it}\varepsilon_{it}\varepsilon_{jt} \tag{A.6}$$

By definition of a strict factor model, $E\left[\varepsilon_{it}\varepsilon_{jt}\right] = 0$ for $i \neq j$, and $E(\varepsilon_{it}^2) = \sigma_{\varepsilon_i}^2$. Applying the expectation operator in equation (A.6) we get:

$$E\left[CSV_t^{(w_t)}\right] = \sum_{i=1}^{N_t} w_{it}\sigma_{\varepsilon_i}^2\left(t\right) - \sum_{i=1}^{N_t} w_{it}^2\sigma_{\varepsilon_i}^2\left(t\right) \tag{A.7}$$

The second term in (A.7) implies that the CSV would tend to underestimate the average idiosyncratic variance. Considering the equal-weighted scheme where

$w_{it} = 1/N_t\ \forall i$, (A.7) simplifies into

$$E\left[CSV_t^{EW}\right] = \left(1 - \frac{1}{N_t}\right)\frac{1}{N_t}\sum_{i=1}^{N_t}\sigma_{\varepsilon_i}^2(t)$$

and we obtain:

$$E\left[CSV_t^{EW}\right] \underset{N_t\to\infty}{\longrightarrow} \frac{1}{N_t}\sum_{i=1}^{N_t}\sigma_{\varepsilon_{it}}^2.$$

A.2.2 Variance of the CSV Estimator

Let w_t and ε_t be column vectors of the weighting scheme and residuals respectively and $\Omega_t = w_t w_t'$, $\Lambda_t = diag(w_t)$, $N_t \times N_t$ matrices, and denote Σ^ε the variance covariance matrix of the residuals, which is diagonal for a strict factor model.

For a finite number of stocks in the case where $F_t \neq r^{(w_t)}$, we have from equation (A.6):

$$CSV_t^{(w_t)} = \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 - \sum_{i=1}^{N_t}\sum_{j=1}^{N_t} w_{jt}w_{it}\varepsilon_{it}\varepsilon_{jt}$$

.

Letting $Q_t = \Lambda_t - \Omega_t$, CSV_t can be written in matrix form, as follows:

$$CSV_t^{(w_t)} = \varepsilon_t' Q_t \varepsilon_t. \tag{A.8}$$

Using the quadratic structure of the CSV and assuming normal residuals, we have (see for instance Kachman (1999))[1]:

$$Var\left(\varepsilon_t' Q_t \varepsilon_t\right) = 2tr\left(Q_t \Sigma_t^\varepsilon Q_t \Sigma_t^\varepsilon\right) \tag{A.9}$$

Under the assumption of a strict factor model, i.e. $\rho_{ijt}^\varepsilon = 0\ \forall i \neq j$, equation (A.9) simplifies to:

$$Var\left(CSV_t^{(w_t)}\right) = 2\sum_{i=1}^{N_t}\sigma_{\varepsilon_{it}}^4 w_{it}^2(1-w_{it})^2 + 2\sum_{i=1}^{N_t}\sum_{j\neq i}^{N_t} w_{it}^2 w_{jt}^2 \sigma_{\varepsilon_{it}}^2\sigma_{\varepsilon_{jt}}^2 \tag{A.10}$$

Assuming an upper bound for the individual idiosyncratic variances, denoted

[1]The operator tr stands for the trace of a matrix, which is the sum of the diagonal terms.

as $\hat{\sigma}_{\varepsilon_t}$ equation (A.10) yields to the following inequality (replacing each variance for its upper bound)

$$Var\left(CSV_t^{(w_t)}\right) < 2\hat{\sigma}_{\varepsilon_t}^4 \left(\left(\sum_{i=1}^{N_t} w_{it}^2\right)^2 + \sum_{i=1}^{N_t} w_{it}^2 - 2\sum_{i=1}^{N_t} w_{it}^3\right). \tag{A.11}$$

When $w_t = 1/N_t$, equation (A.11) simplifies to

$$Var\left(CSV_t^{(w_t)}\right) < 2\hat{\sigma}_{\varepsilon_t}^4 \left(\frac{N_t - 1}{N_t^2}\right) < 2\hat{\sigma}_{\varepsilon_t}^4 \left(\frac{1}{N_t}\right). \tag{A.12}$$

For a large number of stocks,

$$Var\left(CSV_t^{(w_t)}\right) < 2\hat{\sigma}_{\varepsilon_t}^4 \left(\frac{1}{N_t}\right) \longrightarrow 0. \tag{A.13}$$

A.2.3 Relaxing the Assumption of Homogenous Betas

The assumption that $\beta_{it} = \beta_t$ for all i is obviously a simplistic one and is done only for exposure purposes. Starting with the single factor decomposition on the definition of the CSV we have:

$$\begin{aligned}
CSV_t^{(w_t)} &= \sum_{i=1}^{N_t} w_{it}\left(r_{it} - r_t^{(w_t)}\right)^2 \\
&= \sum_{i=1}^{N_t} w_{it}\left[\left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right)F_t + \varepsilon_{it} - \sum_{i=1}^{N_t} w_{jt}\varepsilon_{jt}\right]^2 \\
&= F_t^2 \sum_{i=1}^{N_t} w_{it}\left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right)^2 + \sum_{i=1}^{N_t} w_{it}\left(\varepsilon_{it} - \sum_{i=1}^{N_t} w_{it}\varepsilon_{jt}\right)^2 + \\
&2F_t \sum_{i=1}^{N_t} w_{it}\left(\beta_i - \sum_{j=1}^{N_t} w_{jt}\beta_{jt}\right)\left(\varepsilon_{it} - \sum_{i=1}^{N_t} w_{jt}\varepsilon_{jt}\right)
\end{aligned}$$

After simple rearrangement of terms we get:

$$CSV_t^{(w_t)} = F_t^2 \sum_{i=1}^{N_t} w_{it} \left(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt} \right)^2 + \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}^2 - \sum_{i=1}^{N_t}\sum_{j=1}^{N_t} w_{jt}w_{it}\varepsilon_{it}\varepsilon_{jt}$$
$$+ 2F_t \sum_{i=1}^{N_t} w_{it}\varepsilon_{it}(\beta_{it} - \sum_{j=1}^{N_t} w_{jt}\beta_{jt})$$

Applying the expectation operator and assuming a strict factor model, the last expression simplifies so as to yield:

$$E\left[CSV_t^{(w_t)}\right] = E\left[F_t^2 CSV_t^{\beta}\right] + \sum_{i=1}^{N_t} w_{it}\sigma_{\varepsilon_i t}^2 - \sum_{i=1}^{N_t} w_{it}^2\sigma_{\varepsilon_i t}^2$$

Under an equal-weighting scheme, we finally have:

$$E\left[CSV_t^{EW}\right] = E\left[F_t^2 CSV_t^{\beta}\right] + \left(1 - \frac{1}{N_t}\right)\frac{1}{N_t}\sum_{i=1}^{N_t} \sigma_{\varepsilon_i t}^2$$

Appendix B

Appendix Chapter 2

B.1 Out of Sample Prediction using a Markov Switching Model

Similar to the BH-BCP algorithm the Markov Switching model can yield a probability of being in a particular state at each point in time. Whereas the BH-BCP algorithm produces a single posterior change-probability sequence, the Markov Model estimation procedure can produce two different sets of probability sequences - usually referred to as "filtered" (or "unsmoothed") and also a "smoothed' probabilities sequence'. Given the estimated regime means in each state and the estimated probabilities of being in each state at a point in time, we obtain an expected value for the regime mean, analogous to the posterior mean produced by the BH-BCP procedure.

In this appendix, we present results on predicting returns by using the adjusted d-p ratio series as the predictor where the adjustment level is estimated by fitting the observed series using a Markov Switching model instead of the BH-BCP procedure.

At each point in time, we estimate a model based on all data available at that time, making it a "true" out-of-sample experiment which could have been performed by an investor at each point in time.

Additionally, following Lettau and van Nieuwerburgh (2008), we also perform a "pseudo" out-of-sample experiment. Instead of relying only on data being available at each point in time, we first generate the d-p adjustment by looking ahead over the entire series in order to build the best possible estimates of the regimes and the regime means. Of course, this could not have been performed by an investor at each point in time, but it does serve to demonstrate what that investor would have been able to do if she had a more reliable estimation methodology at her disposal.

Finally, we repeat each of the true and pseudo out of sample experiments for different calibration periods and sub periods, and then for two different market series (CRSP Value Weighted Broad Market and the CRSP Equally Weighted Broad Market indices).

We first present the results for annual data then for quarterly data. In each model, we report the Mincer-Zarnowitz R^2, the Campbell R^2_{OOS} as well as the difference in utility as we did for BCP-based predictions.

B.1.1 Out of Sample Prediction using a Markov Switching Model - Annual Data

Tables E.9 and E.10 present the summary results of running a returns prediction based on a "true" out of sample experiment (denoted by "True OOS" in the tables). We examine both the CRSP Value Weighted series (denoted by $CRSP_{VW}$) and the CRSP Equally Weighted Series (denoted by $CRSP_{EW}$). For each series we examine the results both true and pseudo OOS results for different calibration periods for the unadjusted dp ratio and then using smoothed and unsmoothed probabilities by fitting a 2-state and a 3-state Markov Switching model.

We see that none of the out of sample predictions based on Markov Switching models succeeds in beating the random walk model, either for CW or EW indices. This is in contrast to the BCP procedure which produces economically meaningful predictions that outperform the random walk model.

However, if we are permitted to "peek ahead" to construct an adjustment to the d-p ratio by looking ahead in a Pseudo out of sample experiment, the Markov Switching model does outperform the Random Walk model in several periods, confirming the notion that better estimates of the regimes result in more accurate predictions (results are in Tables E.11 and E.12).

B.1.2 Out of Sample Prediction using a Markov Switching Model - Quarterly Data

Tables E.13 and E.14 present the results of using Quarterly Value Weighted data instead of Annual data, using both 2 and 3 state Markov Switching models.

As in the case of Annual data, we see that Pseudo OOS results often outperform the random walk model, but True OOS predictions using Markov Switching models are unable to do so consistently.

Appendix C

Appendix Chapter 3

C.1 Derivation of Portfolio Insurance Growth Rate

Using a Black-Scholes model for the dynamics of the Reserve (R) and performance seeking asset (S) and the fact that the CPPI can be decomposed in to a floor process, $F_t = kR_t$ and a Fixed-Mix Portfolio, we now derive an approximation of the growth rate of the CPPI, in the case with a stochastic reserve asset.

Using the interpretation of the Cushion as a Fixed-Mix portfolio with weights m and $1-m$ and result (3.11) for two assets, the value of the Portfolio Insurance strategy is given by

$$V_m^{PI}(t) = kV_0 \left(\frac{R_t}{R_0}\right) + (1-k)V_0 \left(\frac{S_t}{S_0}\right)^m \left(\frac{R_t}{R_0}\right)^{1-m} e^{\gamma_m^* t}. \tag{C.1}$$

The second term in equation (C.1) is equal to the value of a fixed mix portfolio, thus it can be expressed in terms of its growth rate, as in equation (3.12),

$$\begin{aligned} C_m(t) &= (1-k)V_0 e^{(\gamma_m^* + m\gamma_S + (1-m)\gamma_R)t + m\sigma_S W_S(t) + (1-m)\sigma_R W_R(t)} & \text{(C.2)} \\ &= (1-k)V_m^{FM}(t) & \text{(C.3)} \end{aligned}$$

Since $C_0 = (1-k)V_0$, by definition of the growth rate (eq. (3.4) or (3.5)), the Cushion process's growth rate is

$$\gamma_m^{cushion} = \gamma_m^* + m(\gamma_S - \gamma_R) + \gamma_R,$$

where $\gamma_m^* = m(1-m)\left(\sigma_S^2 + \sigma_R^2 - 2\rho\sigma_S\sigma_S\right)$. Notice that the brownian motion terms in equation (C.2) disappear after dividing by t as $t \to \infty$, because $W(t) = \sqrt{t}z(t)$ for $z \sim \mathcal{N}(0,1)$. Alternatively, replacing (C.2) in the other definition of the growth rate, i.e. equation (3.5), leads to the same result because $E[W(t)] = 0$.

By definition of the growth rate (equation (3.5)) and using result (C.1), the portfolio insurance' growth rate can be written as

$$\gamma_m^{PI} = \frac{1}{t} E\left[\ln\left(\frac{kR(t) + (1-k)V_m^{FM}(t)}{V_0}\right)\right]. \tag{C.4}$$

From Jensen's inequality and the concavity of the log it follows that

$$\gamma_m^{PI} = \frac{1}{t} E\left[\ln\left(\frac{kR(t)+(1-k)V_m^{FM}(t)}{V_0}\right)\right] \geq \frac{1}{t} E\left[k\ln\left(\frac{R(t)}{V_0}\right)+(1-k)\ln\left(\frac{V_m^{FM}(t)}{V_0}\right)\right],$$

by definition of the growth rate we have

$$\gamma_m^{PI} \geq k\gamma_R + (1-k)\gamma_m^{cushion}. \tag{C.5}$$

Inequality (C.5) implies

$$\gamma_m^{PI} = k\gamma_R + (1-k)\gamma_m^{cushion} + \phi, \text{ with } \phi \geq 0. \tag{C.6}$$

Since the growth rate of the Floor process is equal to the growth rate of the Reserve asset, γ_R, equation (C.6) implies that the growth rate of the Portfolio Insurance strategy is the weighted average of the growth rate of the Reserve asset and the growth rate of the Cushion with weights equal to k and $(1-k)$ respectively plus a positive term ϕ. The latter term can be estimated using a Taylor expansion of the log where all terms above the quadratic are omitted (see below), yielding

$$\phi = \frac{1}{2}k(1-k)\left(\gamma_R^2 + (\gamma_m^{cushion})^2 - 2\gamma_R\gamma_m^{cushion}\right).$$

The log of a weighted average of exponentials can be approximated using the following Taylor expansion where all terms above the quadratic are omitted:

$$\ln(w'\exp(x)) \approx x'w + \frac{1}{2}x'Vx \tag{C.7}$$

where $V_{i,j} = -w_i w_j \forall i \neq j$ and $V_{i,i} = w_i(1-w_i)$, and w is a vector with the same size of vector x, satisfying $\sum_i w_i = 1$. Thus, the "weights" of the weighted average of the exponentials are k and $1-k$ when applying approximation (C.7) in equation (C.4).

C.2 Portfolio Insurance Growth Rate Maximization

As shown above, the growth rate of the CPPI can be approximated as

$$\gamma_m^{PI} = k\gamma_R + (1-k)\gamma_m^{cushion} + \phi,$$

where

$$\phi = \frac{1}{2}k(1-k)\left(\gamma_R^2 + (\gamma_m^{cushion})^2 - 2\gamma_R\gamma_m^{cushion}\right). \tag{C.8}$$

In order to maximize γ_m^{PI} we assume positive values of γ_R and $\gamma_m^{cushion}$ and $\gamma_m^{cushion} \geq \gamma_R$. The latter assumption seems reasonable, as investing in a CPPI strategy in which the expected long term return (growth rate) of the Cushion is lower than the growth rate of the reserve asset would not make sense from an economic standpoint (it would better to simply invest in the reserve asset). Under these assumptions, ϕ, as given by equation (C.8) is a strictly increasing function of $\gamma_m^{cushion}$.

On the other hand, the growth rate of the Floor process is equal to γ_R, which is independent from the multiplier. Hence, under the assumptions above, maximizing the growth rate of the portfolio insurance strategy is equivalent to maximize the growth rate of the cushion (Alternatively, one could notice that for k close to 1 and typical parameter values, ϕ tends to be small compared to the weighted average of the Cushion and Floor's growth rate, thus: $\gamma_m^{PI} \approx k\ \gamma_R + (1-k)\gamma_m^{cushion}$, which implies to the same result).

As shown above, the Cushion process has a growth rate equal to

$$\gamma_m^{cushion} = \gamma_m^* + m(\gamma_S - \gamma_R) + \gamma_R, \tag{C.9}$$

taking the partial derivative of equation (C.9) with respect to m, equating to zero and solving for m yields the *growth optimal* multiplier:

$$m^* = \frac{\gamma_S - \gamma_R + \gamma^*}{2\gamma^*},$$

where, $\gamma^* = \frac{1}{2}(\sigma_S^2 + \sigma_R^2 - 2\rho\sigma_S\sigma_R)$.

C.3 Leveraged investment in Satellite

One alternative to construct a DCS with an asset with "good" properties in terms of risk and return, is to put the asset in the Core and a leveraged investment of the same asset in the Satellite. Choosing a leveraged factor equal to the multiplier of the portfolio insurance strategy, the formula of the growth rate of the Cushion has a simpler expression. The parameters of the Satellite's growth rate are in consequence given by $\mu_S = m\mu_R$, $\sigma_S = m\sigma_R$ and the rebalancing drag is minimized for a given level of volatility since $\rho = 1$. Replacing these terms in Equation (3.18) we get,

$$\gamma_m^{l=m} = (m^2 - m + 1)\mu_R + \frac{1}{2}m(m^2 - m + 1)\sigma_R^2 - \frac{1}{2}m(m^2 - 1)\sigma_R \tag{C.10}$$

For the growth rate of the cushion to be higher than the growth rate of the asset, the following condition should be satisfied:

$$(m^2 - m)\frac{\mu_R}{\sigma_R} + \frac{1}{2}(m^3 - m^2 + m + 1)\sigma_R > \frac{1}{2}(m^3 - m). \tag{C.11}$$

In particular, for $m = 2$, Equation (C.10) simplifies to,

$$\begin{aligned}
\gamma_2^{l=2} &= 3\mu_R + 3\sigma_R^2 - 3\sigma_R \\
\gamma_2^{l=2} &= 3\left(\mu_R - \sigma_R(1 - \sigma_R)\right)
\end{aligned}$$

Which is positive for an asset with a ratio of expected return to volatility satisfying $\frac{\mu}{\sigma} > 1 - \sigma$. In this particular case, for the growth rate of the cushion to be higher than the growth rate of the asset, the following condition should be satisfied:

$$3\left(\mu_R - \sigma_R(1 - \sigma_R)\right) > \mu_R - \frac{1}{2}\sigma_R^2 \Rightarrow \mu - \sigma > \frac{\mu}{3} - \frac{7}{6}\sigma^2. \tag{C.12}$$

For given values of μ_R andσ_R one may find the value of m that maximizes $\gamma_m^{l=m}$. Deriving Equation (C.10) with respect to m and equating to zero we have:

$$2m\gamma_R - \frac{3}{2}m^2\sigma_R(1 - \sigma_R) + \frac{1}{2}\sigma_R - \gamma_R = 0.$$

Which solution is given by the quadratic formula, yielding (Here the numerator and denominator of the typical presentation of the quadratic formula was multiplied by -1).

$$m^* = \frac{2\gamma_R \mp \sqrt{4\gamma_R^2 + 6\sigma_R(1 - \sigma_R)(\frac{1}{2}\sigma_R - \gamma_R)}}{3\sigma_R(1 - \sigma_R)}.$$

C.4 Alternative Time-Varying Multiplier

In Section 3.4 we claimed that a reasonable objective for an investor is to maximize the growth rate of the portfolio. An alternative and perhaps more general objective for the investor could be to maintain a target level of growth rate in the portfolio. The target level of the growth rate can be chosen, for instance, such that the probability of breaching the floor in the presence of jumps in asset prices is held constant and the level of this probability could be chosen according to the risk aversion level, as done in Cont and Tankov (2009).

Hence, instead of letting the growth rate to change as dictated by the movements in the component assets' parameters, the investor may adjust the multiplier in a way that compensates the movements in the asset's parameters.

Let the target growth rate be denoted as Γ, which depends on the assets' parameters and the multiplier as in Equation (3.18):

$$\gamma_m^{cushion}(t) = m_t(1-m_t)\gamma^*(t) + m_t(\gamma_S(t) - \gamma_R(t)) + \gamma_R(t) = \Gamma. \tag{C.13}$$

Hence we solve for m_t, leaving Γ constant as follows,

$$\begin{aligned} \Gamma &= m_t(1-m_t)\gamma^*(t) + m_t(\gamma_S(t) - \gamma_R(t)) + \gamma_R(t) \\ \Gamma - \gamma_R(t) &= m_t(\gamma_S(t) - \gamma_R(t) + \gamma^*(t)) - m_t^2\gamma^*(t). \end{aligned}$$

Which solution is given by the quadratic formula, yielding

$$m_t = \frac{b \mp \sqrt{b^2 + 4\gamma^*(t)(\gamma_R(t) - \Gamma)}}{2\gamma^*(t)} \quad \text{when} \quad b = \gamma_S(t) - \gamma_R(t) + \gamma^*(t). \tag{C.14}$$

Appendix D

Tables and Figures of Chapter 1

Table D.1: Estimates of the biases due to the cross-sectional dispersion of betas and weight concentration: This table contains a summary of the distribution of the following time series: the cross-sectional dispersion of betas CSV_t^{β}, estimated with respect to the CAPM at the end of every month using daily returns; the average idiosyncratic variance $\sigma^2_{\varepsilon_t}$ with respect to the CAPM; the product of the average return of the market portfolio squared, F_t^2, and the beta dispersion, CSV_t^{β}; the proportion of the product $F_t^2 CSV_t^{\beta}$ to $\sigma^2_{\varepsilon_t}$ and the proportion of $\sum w_{it}^2 \sigma^2_{\varepsilon_{it}}$ to $\sigma^2_{\varepsilon_t}$. The upper panel corresponds to the equal-weight scheme (CSV^{EW}) and the lower panel to the market-cap weighting (CSV^{CW}). All figures are daily. The period is July 1963 to December 2006.

Equal-Weighted	$Q_{2.5}$	Q_{25}	Q_{50}	Q_{75}	$Q_{97.5}$
CSV_t^{β}	0.282	0.970	1.563	3.022	11.437
$\sigma^2_{\varepsilon_t}(\%)$	0.043	0.065	0.103	0.241	0.485
$F_t^2 CSV_t^{\beta}(\%)$	6.57e-07	6.92e-05	3.84e-04	0.001	0.005
$\frac{F_t^2 CSV^{\beta}}{\sigma^2_{\varepsilon_t}}(\%)$	0.001	0.078	0.348	0.890	3.240
$\frac{\sum w_{it}^2 \sigma^2_{\varepsilon_{it}}}{\sigma^2_{\varepsilon_t}}(\%)$	0.014	0.020	0.030	0.054	0.154

Cap-Weighted	$Q_{2.5}$	Q_{25}	Q_{50}	Q_{75}	$Q_{97.5}$
CSV_t^{β}	0.075	0.309	0.451	0.704	3.079
$\sigma^2_{\varepsilon_t}(\%)$	0.009	0.020	0.030	0.042	0.153
$F_t^2 CSV_t^{\beta}(\%)$	1.83e-07	2.30e-05	1.09e-04	2.77e-04	0.001
$\frac{F_t^2 CSV^{\beta}}{\sigma^2_{\varepsilon_t}}(\%)$	4.85e-04	0.080	0.351	0.930	3.472
$\frac{\sum w_{it}^2 \sigma^2_{\varepsilon_{it}}}{\sigma^2_{\varepsilon_t}}(\%)$	0.173	0.281	0.426	0.637	1.463

Table D.2: Total bias associated with CSV: This table reports the output summary of the regression $CSV_t^{w_t} = bias + \psi\sigma^2_{model}(w_t) + \zeta_t$, where $\sigma^2_{model}(w_t)$ represents monthly estimates of the weighted average idiosyncratic variance estimated using the corresponding model (either CAPM or FF). The average and the CSV are computed with either the cap-weighted scheme (CW) or the equal-weighted one (EW). The period is July 1963 to December 2006.

	$CAPM^{EW}$	FF^{EW}	$CAPM^{CW}$	FF^{CW}
Bias	1.29e-05	2.23e-05	-2.09e-05	-3.74e-05
NW t-stat	1.986	2.382	-2.849	-4.767
Std. dev.	3.05e-06	5.86e-06	2.04e-06	3.50e-06
ψ	0.983	0.988	1.125	1.242
NW t-stat	153.819	100.162	39.259	39.226
Std. dev.	0.002	0.003	0.005	0.009
$\overline{R}^2(\%)$	99.866	99.503	98.946	97.117

Table D.3: Comparison of monthly measures of idiosyncratic variance: The upper panel of this table contains the annualized mean and standard deviation of the monthly time series for the CSV, the average idiosyncratic variance based on the CAPM and the Fama-French models and the average Goyal and Santa-Clara (2003) variance measure as in equations (1.4), (1.11) and (1.12) using both EW and CW weighting schemes. The lower panel presents the cross-correlation matrix among these variables. The period is January 1964 to December 2006.

	CSV^{EW}	FF^{EW}	$CAPM^{EW}$	GS^{EW}	CSV^{CW}	FF^{CW}	$CAPM^{CW}$	GS^{CW}
Mean	0.384	0.383	0.387	0.342	0.085	0.076	0.080	0.112
Std.Dev.	0.085	0.086	0.087	0.070	0.020	0.016	0.018	0.029

Correlation	CSV^{EW}	FF^{EW}	$CAPM^{EW}$	GS^{EW}	CSV^{CW}	FF^{CW}	$CAPM^{CW}$	GS^{CW}
	100.00	99.75	99.93	95.44	72.16	74.98	72.83	61.25
		100.00	99.88	93.67	68.53	72.16	69.46	56.75
			100.00	94.78	70.62	73.66	71.47	59.83
				100.00	82.38	82.88	82.36	76.60
					100.00	98.56	99.48	92.64
						100.00	99.18	88.17
							100.00	92.13
								100.00

Table D.4: Comparison of daily measures of idiosyncratic variance: The upper panel of this table contains the annualized mean and standard deviation of the daily time series for the CSV and the average idiosyncratic variance based on the Fama-French model as in equations (1.4) and (1.11) using both weighting schemes. The lower panel presents the cross-correlation matrix among these variables. The period is January 1964 to December 2006.

	CSV^{EW}	FF^{EW}	CSV^{CW}	FF^{CW}
Mean	0.384	0.383	0.085	0.078
Std.Dev.	0.021	0.019	0.005	0.004
Correlation	CSV^{EW}	FF^{EW}	CSV^{CW}	FF^{CW}
	100.00	82.63	60.33	63.96
		100.00	52.12	72.55
			100.00	73.95
				100.00

Table D.5: Regime-Switching Parameters: This table contains the parameter estimates of the Markov regime-switching model specified in equation 1.14 for the CSV and the average idiosyncratic variance based on the FF model as in equations (1.4) and (1.11) using both, equal-weighted and cap-weighted schemes. The upper panel corresponds to monthly estimates and the lower panel to daily estimates. μ_i is the average level of the variable on regime i, σ_i is the standard deviation level of the variable on regime i, ϕ is the autocorrelation coefficient, p and q are the probabilities of remaining in regimes 1 and 2 correspondingly. The period is January 1964 to December 2006.

Monthly series	CSV^{EW}	FF^{EW}	CSV^{CW}	FF^{CW}
μ_1	0.401	0.363	0.107	0.115
μ_2	0.299	0.275	0.065	0.061
σ_1	0.067	0.062	0.029	0.021
σ_2	0.010	0.009	0.004	0.003
ϕ	0.980	0.981	0.839	0.839
p	0.839	0.823	0.857	0.906
q	0.963	0.951	0.980	0.990

Daily series	CSV^{EW}	FF^{EW}	CSV^{CW}	FF^{CW}
μ_1	0.446	0.261	0.110	0.048
μ_2	0.304	0.262	0.064	0.048
σ_1	0.036	2.04e-04	0.009	4.89e-05
σ_2	0.003	0.002	0.001	3.60e-04
ϕ	0.965	1.000	0.825	1.004
p	0.695	0.962	0.778	0.870
q	0.956	0.838	0.970	0.809

Table D.6: Correlations between the monthly series of several measures of cross-sectional variance and economic variables. The sample period is January 1990 to December 2006.

	CSV^{EW}	CSV^{CW}	CSV^{EW+}	CSV^{EW-}
Consumption-Vol	0.401	0.241	0.184	0.346
Credit-Spread	0.177	0.268	0.098	0.165
Term-Spread	-0.086	-0.135	-0.107	-0.219
Inflation-Vol	-0.367	0.019	-0.137	-0.097
T-bill Rate	0.302	-0.043	0.091	0.164

Table D.7: Predictability Regression on the CRSP broad market portfolio using the CSV and CAPM-idvol equal-weighted and cap-weighted measures: The first panel of this table presents the results of a one-month ahead predictive regression of the excess cap-weighted monthly portfolio returns, denoted by r^{CW}, on the monthly lagged equal-weighted average idiosyncratic variance and CSV^{EW} for three sample periods. The second panel presents the results of a one-month ahead predictive regression of the excess equal-weighted monthly portfolio returns, denoted by r^{EW}, on the monthly lagged equal-weighted average idiosyncratic variance and CSV^{EW}. The third panel presents the results of a one-month ahead predictive regression of the excess cap-weighted monthly portfolio returns, denoted by r^{CW}, on the monthly lagged cap-weighted idiosyncratic variance and CSV^{CW}. $CAPM^{w}$ is the average idiosyncratic variance derived from CAPM, estimated using one month of daily data and CSV^{w} is the average of the daily cross-sectional variance over each month. The intercept, the regression coefficient of the corresponding lagged idiosyncratic variance, the standard errors denoted by Std, the Newey-West corrected t-stats and the adjusted coefficient of determination denoted by $\overline{R}^2$ are reported. The sample periods are 1963:08 - 1999:12, 1963:08 - 2001:12 and 1963:08 - 2006:12.

Monthly series	1963:08 - 1999:12		1963:08 - 2001:12		1963:08 - 2006:12	
Forecasting r^{CW}	$CAPM^{EW}$	CSV^{EW}	$CAPM^{EW}$	CSV^{EW}	$CAPM^{EW}$	CSV^{EW}
Intercept	-0.001	-0.002	0.001	0.001	0.002	0.002
NW t-stat	-0.424	-0.479	0.228	0.264	0.586	0.617
Std	0.003	0.003	0.003	0.003	0.003	0.003
Coefficient	0.241	0.250	0.123	0.120	0.090	0.087
NW t-stat	3.543	3.609	1.356	1.247	1.055	0.964
Std	0.092	0.094	0.082	0.083	0.077	0.078
$\overline{R}^2(\%)$	1.336	1.365	0.275	0.233	0.072	0.044
Forecasting r^{EW}	$CAPM^{EW}$	CSV^{EW}	$CAPM^{EW}$	CSV^{EW}	$CAPM^{EW}$	CSV^{EW}
Intercept	4.51e-04	3.31e-04	-1.67e-04	-1.06e-05	0.001	0.001
NW t-stat	0.086	0.063	-0.033	-0.002	0.241	0.270
Std	0.004	0.004	0.004	0.004	0.004	0.004
Coefficient	0.247	0.254	0.238	0.235	0.217	0.215
NW t-stat	2.175	2.189	2.395	2.329	2.331	2.271
Std	0.118	0.121	0.105	0.107	0.099	0.101
$\overline{R}^2(\%)$	0.774	0.773	0.885	0.824	0.726	0.678
Forecasting r^{CW}	$CAPM^{CW}$	CSV^{CW}	$CAPM^{CW}$	CSV^{CW}	$CAPM^{CW}$	CSV^{CW}
Intercept	0.001	0.001	0.007	0.007	0.007	0.008
NW t-stat	0.150	0.334	2.226	2.459	2.551	2.823
Std	0.005	0.004	0.003	0.003	0.003	0.003
Coefficient	0.856	0.688	-0.356	-0.373	-0.404	-0.421
NW t-stat	1.192	0.995	-0.871	-1.081	-1.080	-1.334
Std	0.685	0.632	0.390	0.344	0.372	0.328
$\overline{R}^2(\%)$	0.128	0.043	-0.037	0.038	0.035	0.123

Table D.8: Daily predictability Regression on CRSP broad market portfolio with average idiosyncratic variance measures: The upper panel presents the results of a one-day ahead predictive regression of the excess equal-weighted daily portfolio returns, denoted by r^{EW}, on the daily lagged equal-weighted cross-sectional variance denoted as CSV^{EW} estimated as in equation (1.4) for three sample periods. The lower panel presents the results of the predictive regression on the cap-weighted market portfolio using the cap-weighted CSV. The intercept, the regression coefficient corresponding to the CSV, the standard error of the regression coefficients denoted by std, the Newey-West corrected t-stats (30 lags) and the adjusted coefficient of determination denoted by $\overline{R}^2$ are reported. The sample periods are 1963:07 to 1999:12, 1963:07 to 2001:12 and 1963:07 to 2006:12.

Daily series	63:07-99:12	63:07-01:12	63:07-06:12
Forecasting r^{EW}	CSV^{EW}	CSV^{EW}	CSV^{EW}
Intercept	-1.58e-04	-1.40e-04	-1.29e-05
NW t-stat	-0.785	-0.714	-0.071
Std	1.09e-04	1.10e-04	1.04e-04
Coefficient	0.544	0.483	0.411
NW t-stat	4.711	4.515	4.000
Std	0.060	0.055	0.051
$\overline{R}^2(\%)$	0.883	0.788	0.573
Forecasting r^{CW}	CSV^{CW}	CSV^{CW}	CSV^{CW}
Intercept	-0.001	-1.88e-04	-1.65e-04
NW t-stat	-3.521	-0.791	-0.737
Std	1.41e-04	1.23e-04	1.19e-04
Coefficient	3.404	1.189	1.151
NW t-stat	5.919	1.948	1.966
Std	0.385	0.251	0.248
$\overline{R}^2(\%)$	0.831	0.220	0.186

Table D.9: Daily and Monthly predictability with and without market volatility: This table presents summary statistics for three predictive regressions. The first two panels correspond to daily and monthly predictions on the equal-weighted broad market index using equal-weighted CSV and the realized variance of the equal-weighted broad market index and the two lower panels their counterparts using the cap-weighted index and CSV. In each panel, the first row corresponds to the regression $r_t^w = \alpha + \beta CSV_t^w + \epsilon_t$ the second row to $r_t^w = \alpha + \vartheta Var\left(r_t^w\right) + \epsilon_t$ and the third one to $r_t^w = \alpha + \beta CSV_t^w + \vartheta Var\left(r_t^w\right) + \epsilon_t$. The sample period is July 1963 to December 2006.

Daily Estimates	Intercept	T-stat	CSV^{EW}	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	-1.29e-05	-0.071	0.411	4.000			0.573
Forecasting r^{EW}	0.001	4.928			-2.437	-0.976	0.034
Forecasting r^{EW}	1.19e-04	0.629	0.470	4.672	-5.038	-2.226	0.737
Monthly Estimates	**Intercept**	**T-stat**	CSV^{EW}	**T-stat**	$Var\left(r^{EW}\right)$	**T-stat**	$\overline{R}^2(\%)$
Forecasting r^{EW}	0.001	0.270	0.215	2.271			0.678
Forecasting r^{EW}	0.008	2.673			0.240	0.182	-0.185
Forecasting r^{EW}	0.001	0.324	0.226	2.206	-0.478	-0.394	0.515
Daily Estimates	**Intercept**	**T-stat**	CSV^{CW}	**T-stat**	$Var\left(r^{CW}\right)$	**T-stat**	$\overline{R}^2(\%)$
Forecasting r^{CW}	-1.65e-04	-0.737	1.151	1.966			0.186
Forecasting r^{CW}	2.10e-05	0.183			2.680	2.003	0.068
Forecasting r^{CW}	-1.76e-04	-0.937	1.084	1.298	0.447	0.198	0.179
Monthly Estimates	**Intercept**	**T-stat**	CSV^{CW}	**T-stat**	$Var\left(r^{CW}\right)$	**T-stat**	$\overline{R}^2(\%)$
Forecasting r^{CW}	0.008	2.823	-0.421	-1.334			0.123
Forecasting r^{CW}	0.005	2.636			-0.418	-0.569	-0.116
Forecasting r^{CW}	0.008	2.815	-0.459	-1.054	0.130	0.113	-0.065

Table D.10: Quarterly and Annual predictability with and without market variance: This table presents summary statistics for three predictive regressions. The first two panels correspond to quarterly and annual predictions on the equal-weighted broad market index using equal-weighted CSV and the realized variance of the equal-weighted broad market index and the two lower panels their counterparts with the cap-weighted index and CSV. In each panel, the first row corresponds to the regression $r_t^w = \alpha + \beta CSV_t^w + \epsilon_t$ the second row to $r_t^w = \alpha + \vartheta Var\left(r_t^w\right) + \epsilon_t$ and the third one to $r_t^w = \alpha + \beta CSV_t^w + \vartheta Var\left(r_t^w\right) + \epsilon_t$. The sample period is July 1963 to December 2006.

Quarterly Estimates	Intercept	T-stat	CSV^{EW}	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	0.009	0.567	0.191	1.761			0.682
Forecasting r^{EW}	0.008	0.730			5.662	4.027	3.124
Forecasting r^{EW}	-0.001	-0.043	0.104	0.822	5.133	3.474	2.903
Annual Estimates	**Intercept**	**T-stat**	**CSV^{EW}**	**T-stat**	**$Var\left(r^{EW}\right)$**	**T-stat**	**$\overline{R}^2(\%)$**
Forecasting r^{EW}	0.059	0.846	0.131	1.090			-0.544
Forecasting r^{EW}	0.055	1.276			4.004	2.371	-0.063
Forecasting r^{EW}	0.027	0.358	0.099	0.731	3.292	2.087	-1.548
Quarterly Estimates	**Intercept**	**T-stat**	**CSV^{CW}**	**T-stat**	**$Var\left(r^{CW}\right)$**	**T-stat**	**$\overline{R}^2(\%)$**
Forecasting r^{CW}	0.021	2.063	-0.277	-0.644			-0.288
Forecasting r^{CW}	0.006	0.869			1.800	1.749	1.195
Forecasting r^{CW}	0.021	2.467	-1.141	-2.608	3.686	2.188	3.719
Annual Estimates	**Intercept**	**T-stat**	**CSV^{CW}**	**T-stat**	**$Var\left(r^{CW}\right)$**	**T-stat**	**$\overline{R}^2(\%)$**
Forecasting r^{CW}	0.093	2.465	-0.416	-1.265			0.142
Forecasting r^{CW}	0.045	1.297			0.606	0.451	-2.112
Forecasting r^{CW}	0.079	2.125	-1.024	-2.186	3.314	1.726	3.719

Table D.11: Daily and Monthly predictability with and without market realized variance and implied variance (VIX^2): The first two panels correspond to daily and monthly predictions of the equal-weighted broad market index using equal-weighted CSV, the realized variance of the equal-weighted broad market index and the squared VIX index. The two lower panels feature similar regressions with the cap-weighted counterparts. In each panel, we report regression results for each predictor individually and for combinations of the three corresponding predictors over the period July 1990 to December 2006

Daily Estimates	Intercept	T-stat	CSV^{EW}	T-stat	VIX^2	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	-8.59e-05	-0.265	0.390	3.020					0.625
Forecasting r^{EW}	0.002	8.224			-8.400	-4.417			1.697
Forecasting r^{EW}	0.001	4.036					0.860	0.228	-0.020
Forecasting r^{EW}	0.002	7.254			-13.815	-4.890	19.928	3.675	2.866
Forecasting r^{EW}	0.001	2.440	0.788	4.503	-12.706	-6.115			3.866
Forecasting r^{EW}	-2.34e-05	-0.068	0.416	3.223			-2.692	-0.767	0.634
Forecasting r^{EW}	0.001	2.245	0.763	4.512	-17.653	-6.356	18.716	3.609	4.893
Monthly Estimates	Intercept	T-stat	CSV^{EW}	T-stat	VIX^2	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	-0.008	-0.799	0.335	2.193					1.963
Forecasting r^{EW}	0.007	1.686			0.008	0.876			-0.359
Forecasting r^{EW}	0.006	1.159					2.423	0.937	0.188
Forecasting r^{EW}	0.007	1.771			-0.005	-0.265	2.821	0.730	-0.283
Forecasting r^{EW}	-0.008	-0.751	0.392	2.082	-0.011	-0.886			1.688
Forecasting r^{EW}	-0.008	-0.787	0.325	1.819			0.295	0.104	1.478
Forecasting r^{EW}	-0.007	-0.697	0.368	1.969	-0.017	-0.903	1.416	0.368	1.329
Daily Estimates	Intercept	T-stat	CSV^{CW}	T-stat	VIX^2	T-stat	$Var\left(r^{CW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{CW}	1.79e-04	0.870	0.242	0.657					-0.011
Forecasting r^{CW}	0.001	5.852			-7.438	-4.129			0.772
Forecasting r^{CW}	-1.04e-04	-0.560					4.513	2.258	0.130
Forecasting r^{CW}	0.002	5.427			-27.744	-5.884	35.151	5.504	4.116
Forecasting r^{CW}	0.001	4.578	1.327	2.651	-9.616	-4.530			1.049
Forecasting r^{CW}	-2.04e-05	-0.102	-0.397	-0.807			5.740	2.116	0.128
Forecasting r^{CW}	0.002	5.354	-0.823	-1.690	-28.073	-5.942	38.058	5.959	4.186
Monthly Estimates	Intercept	T-stat	CSV^{CW}	T-stat	VIX^2	T-stat	$Var\left(r^{CW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{CW}	0.013	3.498	-0.640	-1.986					0.999
Forecasting r^{CW}	0.001	0.196			0.016	1.672			0.448
Forecasting r^{CW}	0.005	1.589					0.407	0.273	-0.440
Forecasting r^{CW}	-2.58e-04	-0.062			0.034	2.451	-2.230	-1.242	0.644
Forecasting r^{CW}	0.007	1.772	-1.114	-3.618	0.033	2.965			3.534
Forecasting r^{CW}	0.013	3.814	-1.470	-2.544			3.734	1.741	3.002
Forecasting r^{CW}	0.008	1.797	-1.333	-2.290	0.024	1.773	1.579	0.601	3.258

Table D.12: Quarterly and annual predictability with and without market realized variance and implied variance (VIX^2): The first two panels correspond to daily and monthly predictions of the equal-weighted broad market index using equal-weighted CSV, the realized variance of the equal-weighted broad market index and the squared VIX index. The two lower panels feature similar regressions with the cap-weighted counterparts. In each panel, we report regression results for each predictor individually and for combinations of the three corresponding predictors over the period July 1990 to December 2006

Quarterly Estimates	Intercept	T-stat	CSV^{EW}	T-stat	VIX^2	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	-0.009	-0.278	0.249	1.502					0.833
Forecasting r^{EW}	-0.007	-0.450			0.039	3.548			4.178
Forecasting r^{EW}	0.024	1.551					1.880	0.826	-1.082
Forecasting r^{EW}	-0.006	-0.345			0.057	2.964	-4.664	-1.537	4.254
Forecasting r^{EW}	-0.013	-0.426	0.051	0.211	0.036	1.979			2.751
Forecasting r^{EW}	-0.009	-0.277	0.258	1.218			-0.293	-0.098	-0.708
Forecasting r^{EW}	-0.017	-0.509	0.102	0.425	0.053	2.497	-5.049	-1.571	3.010
Annual Estimates	Intercept	T-stat	CSV^{EW}	T-stat	VIX^2	T-stat	$Var\left(r^{EW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{EW}	0.058	0.825	0.142	1.157					-5.125
Forecasting r^{EW}	0.055	1.020			0.024	1.225			-1.088
Forecasting r^{EW}	0.084	1.556					4.111	1.746	-1.528
Forecasting r^{EW}	0.054	0.988			0.015	0.427	2.167	0.422	-8.041
Forecasting r^{EW}	0.054	0.847	0.002	0.008	0.024	0.913			-8.864
Forecasting r^{EW}	0.068	0.773	0.030	0.135			3.848	1.178	-9.266
Forecasting r^{EW}	0.061	0.741	-0.017	-0.075	0.016	0.436	2.224	0.413	-17.023
Quarterly Estimates	Intercept	T-stat	CSV^{CW}	T-stat	VIX^2	T-stat	$Var\left(r^{CW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{CW}	0.047	4.082	-0.866	-2.996					4.166
Forecasting r^{CW}	-0.007	-0.431			0.027	1.865			3.574
Forecasting r^{CW}	0.018	1.528					0.281	0.136	-1.494
Forecasting r^{CW}	-0.018	-1.100			0.087	5.163	-8.025	-4.416	12.469
Forecasting r^{CW}	0.017	1.158	-1.687	-6.670	0.054	4.670			18.666
Forecasting r^{CW}	0.046	4.149	-1.999	-4.315			5.686	2.257	11.164
Forecasting r^{CW}	0.011	0.683	-1.449	-4.175	0.066	3.589	-2.109	-0.722	17.768
Annual Estimates	Intercept	T-stat	CSV^{CW}	T-stat	VIX^2	T-stat	$Var\left(r^{CW}\right)$	T-stat	$\overline{R}^2(\%)$
Forecasting r^{CW}	0.227	6.448	-1.079	-6.158					26.614
Forecasting r^{CW}	0.115	2.665			-0.006	-0.476			-6.531
Forecasting r^{CW}	0.141	3.486					-1.946	-1.197	-1.348
Forecasting r^{CW}	0.086	1.964			0.032	1.818	-5.026	-2.138	-3.052
Forecasting r^{CW}	0.171	3.068	-1.391	-7.922	0.023	2.093			28.832
Forecasting r^{CW}	0.214	4.254	-1.905	-7.880			4.579	3.342	34.155
Forecasting r^{CW}	0.224	3.222	-1.950	-5.227	-0.005	-0.186	5.193	1.406	28.789

Table D.13: Predictability Regression on CRSP broad market index with right and left CSV measures: The upper panel presents the results of a one-day, one-month, one quarter and one year ahead predictive regressions of the excess equal-weighted portfolio returns, denoted by r^{EW}, on the daily or monthly (correspondingly) lagged equal-weighted cross-sectional variance of the returns to the right (higher than) of the cross-sectional distribution mean (which is actually r_t^{EW}) denoted as CSV^+ and the cross-sectional variance of the returns to the left (lower than) the mean of the cross-sectional distribution r_t^{EW}, denoted as CSV^-. The lower panel presents the results of the predictive regressions on the cap-weighted market index using the cap-weighted CSV measures as predictors. The intercept, the regression coefficients corresponding to the CSV^+ and CSV^-, the standard error of the regression coefficients denoted by Std, the Newey-West corrected t-stats and the adjusted coefficient of determination denoted by $\overline{R}^2$ are reported. The sample period is 1963:07 to 2006:12.

Forecasting r^{EW}	$Daily^{EW}$	$Monthly^{EW}$	$Quarterly^{EW}$	$Annual^{EW}$
Intercept	0.001	0.003	0.014	0.050
NW t-stat	3.944	0.727	0.822	1.069
Std	1.10e-04	0.004	0.017	0.076
CSV^+	0.488	0.375	-0.042	-0.288
NW t-stat	3.360	2.400	-0.282	-0.845
Std	0.038	0.161	0.190	0.403
CSV^-	-1.200	-0.486	0.824	1.701
NW t-stat	-3.551	-1.129	0.842	1.346
Std	0.142	0.432	0.902	1.610
$\overline{R}^2(\%)$	1.552	1.114	-0.595	-1.862
Forecasting r^{CW}	$Daily^{CW}$	$Monthly^{CW}$	$Quarterly^{CW}$	$Annual^{CW}$
Intercept	-1.12e-04	0.008	0.020	0.107
NW t-stat	-0.588	2.753	2.656	2.147
Std	1.16e-04	0.003	0.010	0.043
CSV^+	4.942	0.071	-0.163	-2.354
NW t-stat	3.546	0.054	-0.254	-1.669
Std	0.528	1.983	0.914	1.824
CSV^-	-2.842	-1.302	-0.870	1.227
NW t-stat	-2.736	-0.879	-0.672	0.869
Std	0.669	2.233	1.477	2.527
$\overline{R}^2(\%)$	0.785	-0.069	-0.849	1.215

Table D.14: Daily and Monthly predictability with skewness for r^{EW}: This table presents the results of one-day and one-month ahead predictive regressions of the excess equal-weighted daily portfolio returns, denoted by r^{EW}. The first explanatory variable is lagged estimate of the equal-weighted CSV estimated as in equation (1.4); The second explanatory variable is the robust estimate of skewness estimated as in equations (1.20). The intercept, the corresponding regression coefficients together with their Newey-West autocorrelation corrected t-stats (with 30 lags for daily and 12 lags for monthly) and standard errors are reported. $\overline{R}^2$ denotes adjusted coefficient of determination. The regression is reported for the main sample period from 1963:07 to 2006:12.

Daily horizon	Coeff.	t-stat	Std.Dev.	$\overline{R}^2(\%)$
Intercept	-3.7e-005	-0.234	0.000	5.833
CSV^{EW}	0.402	4.013	0.053	
Skewness	0.004	20.190	0.000	

Monthly horizon	Coeff.	t-stat	Std.Dev.	$\overline{R}^2(\%)$
Intercept	0.000	0.107	0.004	4.587
CSV^{EW}	0.250	2.518	0.102	
Skewness	0.078	4.458	0.017	

Table D.15: CSV quintiles premium. The upper panel present the results for equal-weighted quintile portfolios and the lower panel on cap-weighted quintile portfolios. The first column presents the (arithmetic) average return annualized on quintiles formed at the end of every month on CSV^{EW}'s coefficient estimated with one month of daily returns. The second column presents the average return difference of the first quintile with every other quintile. The third column presents the p-value of the test of the difference to be significantly positive. The sample period is July 1963 to December 2006.

EW Quintiles	Quintile Return	$Q_1 - Q_i$	p-value(%)
Q_1	0.390	0.00e+00	
Q_2	0.083	0.307	0.00e+00
Q_3	0.044	0.346	0.00e+00
Q_4	0.050	0.340	0.00e+00
Q_5	0.237	0.154	0.130

CW Quintiles	Quintile Return	$Q_1 - Q_i$	p-value(%)
Q_1	0.383	0.00e+00	
Q_2	0.087	0.296	0.00e+00
Q_3	0.036	0.347	0.00e+00
Q_4	0.047	0.335	0.00e+00
Q_5	0.247	0.136	0.725

Table D.16: Fama-MacBeth regression. This table displays the average values, standard errors and Newey-West corrected t-stats for the coefficients in the Fama-MacBeth regression run every month in the sample, using the 3 Fama-French factors and CSV^{EW} on 100 and 25 size/book2market Fama-French equally-weighted (first two panels) and cap-weighted (last two panels) portfolios. It also displays the average R^2 across subsamples of Fama-MacBeth regressions. The sample period is July 1963 to December 2006.

100-EW Portfolios	Intercept	XMKT	SMB	HML	CSV^{EW}	$R^2(\%)$
γ	0.223	-0.067	0.029	0.048	0.005	24.657
SE	0.024	0.016	0.012	0.012	0.002	
tstat	9.222	-4.173	2.320	3.916	2.847	

25-EW Portfolios	Intercept	XMKT	SMB	HML	CSV^{EW}	$R^2(\%)$
γ	0.278	-0.133	0.042	0.064	0.009	51.962
SE	0.028	0.021	0.017	0.016	0.003	
tstat	9.969	-6.447	2.480	3.927	2.703	

100-CW Portfolios	Intercept	XMKT	SMB	HML	CSV^{EW}	$R^2(\%)$
γ	0.155	-0.007	0.004	0.037	0.003	24.262
SE	0.023	0.017	0.012	0.013	0.002	
tstat	6.896	-0.421	0.323	2.762	1.909	

25-CW Portfolios	Intercept	XMKT	SMB	HML	CSV^{EW}	$R^2(\%)$
γ	0.175	-0.034	0.010	0.048	0.004	50.815
SE	0.024	0.021	0.016	0.016	0.003	
tstat	7.417	-1.624	0.625	2.910	1.395	

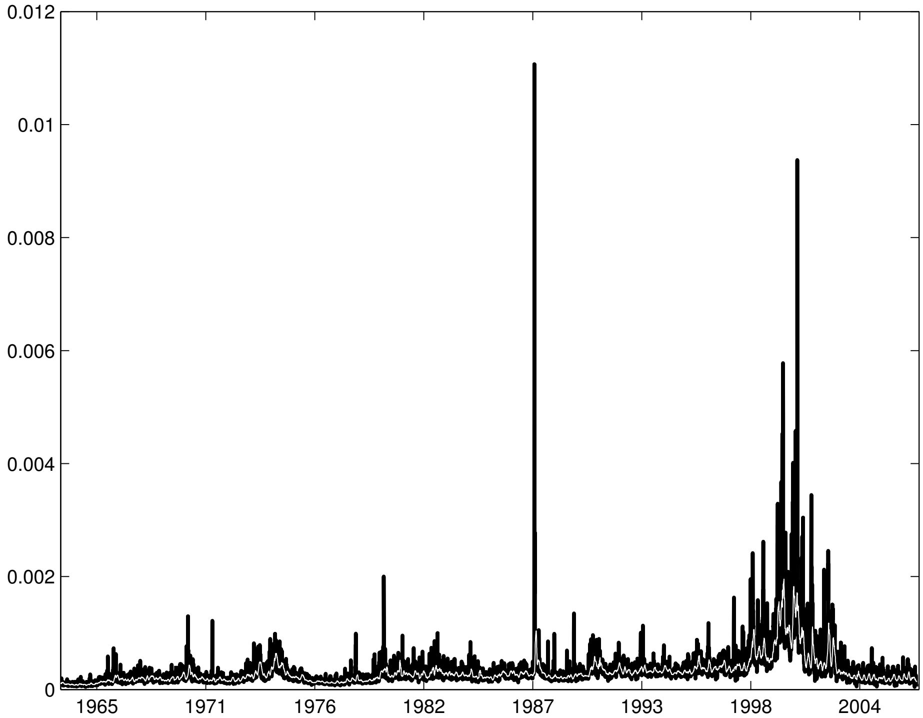

Figure D.1: Cap-weighted idiosyncratic variances, daily estimation: The white line is the time series of the cap-weighted idiosyncratic variance with respect to the FF model estimated daily as in equation 1.11. The darker line shows the time series of the cap-weighted version of CSV estimated daily as in equation 1.5. The sample period is January 1964 to December 2006.

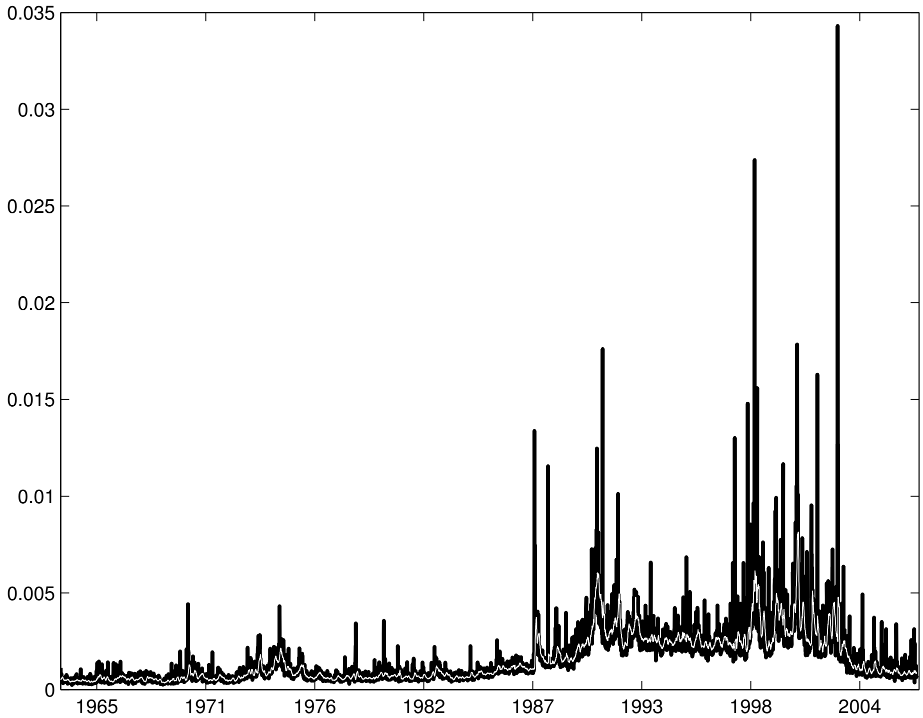

Figure D.2: Equally-weighted idiosyncratic variances, daily estimation: The white line is the time series of the equal-weighted average idiosyncratic variance with respect to the FF model estimated daily similar to equation 1.11. The darker line shows the time series of the CSV^{EW} estimated daily as in equation 1.4. The sample period is January 1964 to December 2006.

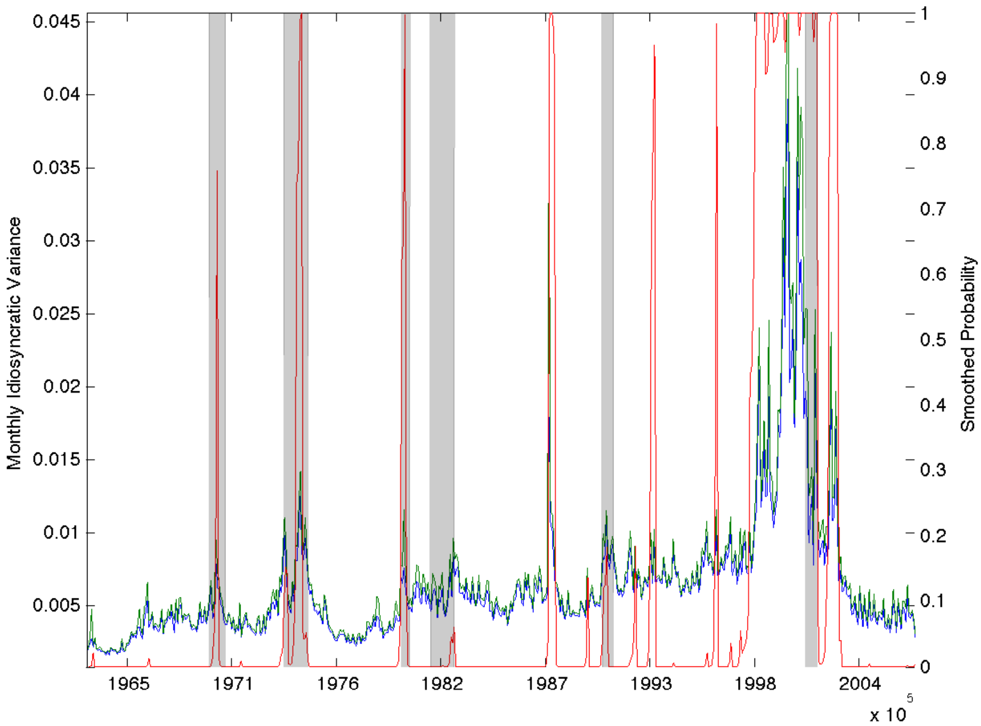

Figure D.3: Regime Switching filtered probabilities and cap-weighted CSV, monthly estimation: The red line plots the filtered probability of the CSV^{CW} being in the high-mean high-variance regime of a Markov regime-switching model specified in equation 1.14. The blue line shows the monthly time series of the CSV^{CW} estimated at the end of each month as the average of the daily estimations (as in equation 1.4) during the month. The shaded areas are the NBER recessions. The sample period is July 1963 to December 2006.

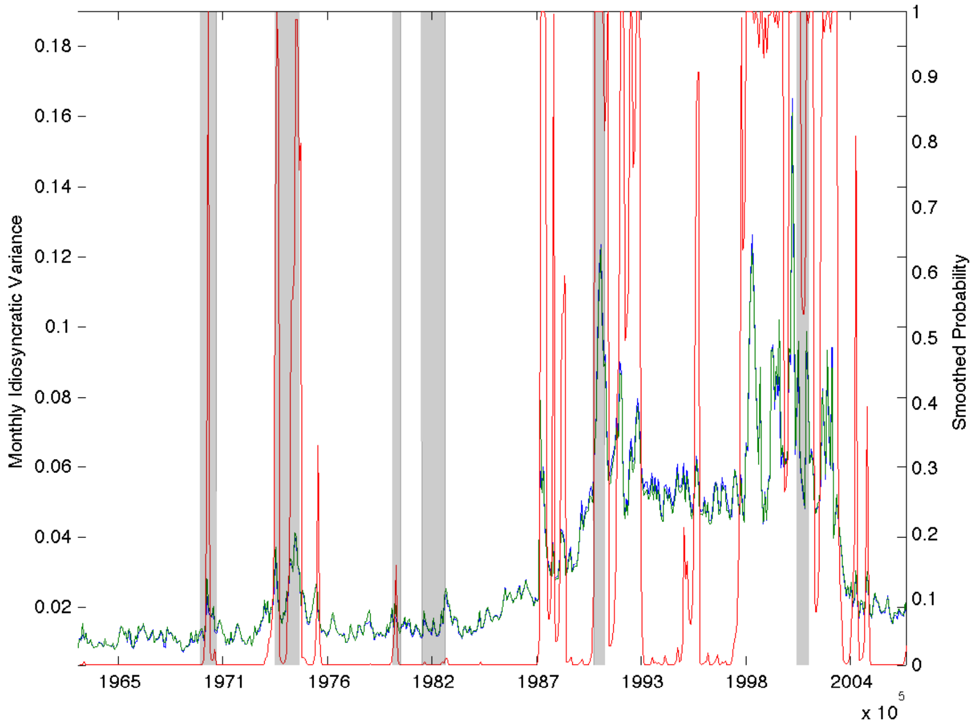

Figure D.4: Regime Switching filtered probabilities and equal-weighted CSV monthly estimation: The red line plots the filtered probability of the CSV^{EW} being in the high-mean high-variance regime of a Markov regime-switching model specified in equation 1.14. The blue line shows the monthly time series of the CSV^{EW} estimated at the end of each month as the average of the daily estimations (as in equation 1.4) during the month. The shaded areas are the NBER recessions. The sample period is July 1963 to December 2006.

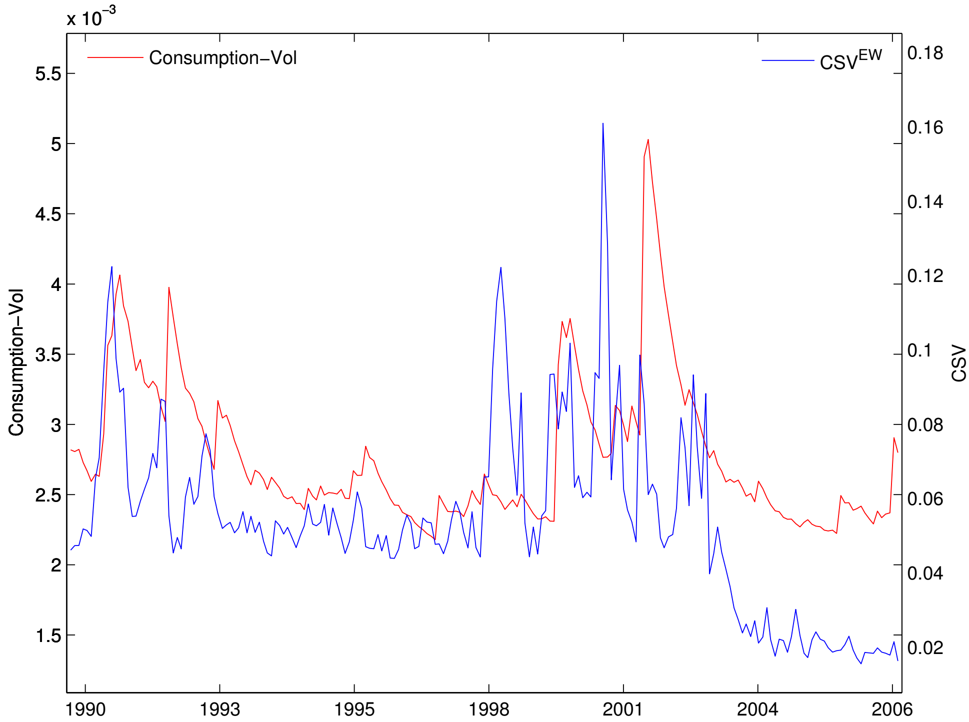

Figure D.5: CSV^{EW} **and Consumption Volatility**:Monthly time series of CSV^{EW} on the right-hand axis and Consumption Volatility on the left-hand axis. The sample period is January 1990 to December 2006.

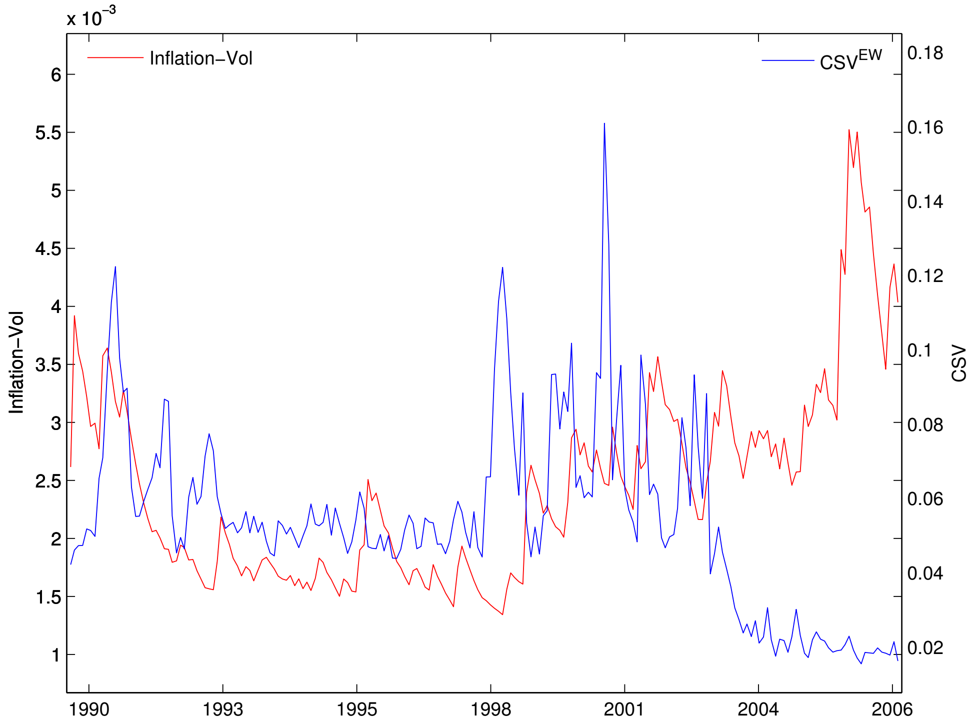

Figure D.6: CSV^{EW} **and Inflation Volatility**: Monthly time series of CSV^{EW} on the right-hand axis and Inflation Volatility on the left-hand axis. The sample period is January 1990 to December 2006.

Appendix E

Tables and Figures of Chapter 2

Table E.1: Persistence Properties of Adjusted Dividend-Price Ratio, Broad Market CRSP VW index. This table displays the autocorrelation coefficients of first and second order and the augmented Dickey-Fuller test of a unit root with the respective p-values and standard errors for the unadjusted d-p series, the adjusted d-p series using the Lettau and van Nieuwerburgh (2008) adjustment with one and two breaks and the d-p series adjusted using the BCP posterior mean. The last line of each panel present the same statistics for the BCP posterior mean of the d-p ratio. The upper panel present the figures for the annual series and the lower panel for the quarterly series. The sample period is 1927 to 2010.

Annual data	AC(1)	AC(2)	ADF test	p-value	Std.dev.
dp unadjusted	0.925	0.848	-1.508	0.123	0.429
$\tilde{dp}$, 1 break	0.775	0.555	-3.209	0.002	0.256
$\tilde{dp}$, 2 break	0.657	0.300	-4.044	0.001	0.204
$\tilde{dp}$ BCP-adjusted	0.019	-0.187	-8.973	0.001	0.069
Posterior Mean	0.958	0.898	0.587	0.841	0.407
Quarterly data	**AC(1)**	**AC(2)**	**ADF test**	**p-value**	**Std.dev.**
dp unadjusted	0.888	0.886	-4.275	0.001	0.437
$\tilde{dp}$, 1 break	0.729	0.726	-7.214	0.001	0.283
$\tilde{dp}$, 2 break	0.627	0.620	-8.702	0.001	0.241
$\tilde{dp}$ BCP-adjusted	-0.135	0.097	-20.840	0.001	0.111
Posterior Mean	0.975	0.960	0.280	0.754	0.410

Table E.2: Predictability Regression with adjusted log d-p, Broad Market CRSP VW index. This table present the summary statistics of the predictive regression using the unadjusted d-p series and the adjusted series using 1 and two breaks for the mean of the predictor, as in Lettau and van Nieuwerburgh (2008) and using the BCP adjustment.The upper panel present the figures using the annual series and the lower panel using the quarterly series to predict the return of the market one period ahead. The sample period is 1927 to 2010.

Annual data	α	t-stat	β	t-stat	$R^2(\%)$
dp unadjusted	0.089	4.295	0.095	2.281	4.033
$\tilde{dp}$, 1 break	0.090	4.321	0.241	4.489	9.473
$\tilde{dp}$, 2 break	0.090	4.653	0.387	4.579	15.270
$\tilde{dp}$ BCP-adjusted	0.089	4.651	1.284	7.770	19.178
Posterior Mean	0.317	2.116	0.068	1.545	1.891
Quarterly data	α	t-stat	β	t-stat	$R^2(\%)$
dp unadjusted	0.023	3.777	0.021	1.332	0.726
$\tilde{dp}$, 1 break	0.023	3.746	0.047	1.926	1.510
$\tilde{dp}$, 2 break	0.023	3.758	0.071	2.665	2.577
$\tilde{dp}$ BCP-adjusted	0.023	3.753	0.226	3.843	5.432
Posterior Mean	0.057	0.693	0.007	0.427	0.077

Table E.3: Out-of-Sample return prediction, Annual data, Broad Market CRSP VW index. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

OOS $CRSP_{VW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
dp unadjusted	5.54	-3.04	-0.10
dp with Shrinkage	2.97	-1.21	-0.15
$\tilde{dp}$ BCP-adjusted	7.47	8.12***	0.81
OOS $CRSP_{VW}$	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
dp unadjusted	1.32	-13.42	-0.51
dp with Shrinkage	0.24	-9.46	-0.58
$\tilde{dp}$ BCP-adjusted	3.18	2.31*	0.25
OOS $CRSP_{VW}$	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
dp unadjusted	5.80	-8.47	0.07
dp with Shrinkage	4.04	-3.96	-0.08
$\tilde{dp}$ BCP-adjusted	4.30	4.63**	0.36

Table E.4: Pseudo-Out-of-Sample return prediction, Annual data, Broad Market CRSP VW index. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

Pseudo-OOS	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	14.04	12.9***	1.16
$\tilde{dp}$, 2 breaks	15.44	14.9***	1.91
$\tilde{dp}$ BCP-adjusted	19.61	20***	3.03
Pseudo-OOS	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	9.66	9.87***	1.19
$\tilde{dp}$, 2 breaks	8.15	1.99	1.26
$\tilde{dp}$ BCP-adjusted	18.36	18.3***	2.87
Pseudo-OOS	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	8.74	3.85**	0.97
$\tilde{dp}$, 2 breaks	14.27	15.5***	1.93
$\tilde{dp}$ BCP-adjusted	22.07	22.8***	2.89

Table E.5: Pseudo-Out-of-Sample return prediction, Quarterly data, Broad Market CRSP VW index. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

Pseudo-OOS	Calibration:	1927:3-1965:3	
	OS period:	1965:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	3.70	3.24***	0.89
$\tilde{dp}$, 2 breaks	4.89	4.99***	1.52
$\tilde{dp}$ BCP-adjusted	8.98	8.4***	2.37
Pseudo-OOS	Calibration:	1927:3-1976:3	
	OS period:	1976:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	3.00	3***	0.90
$\tilde{dp}$, 2 breaks	4.01	4.57***	1.37
$\tilde{dp}$ BCP-adjusted	6.70	6.83***	1.82
Pseudo-OOS	Calibration:	1927:3-1947:3	
	OS period:	1947:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
$\tilde{dp}$, 1 break	2.35	2.46***	0.56
$\tilde{dp}$, 2 breaks	4.37	4.96***	1.37
$\tilde{dp}$ BCP-adjusted	5.42	6.01***	2.09

Table E.6: Out-of-Sample return prediction, Annual data, S&P 500 index and equal-weighted CRSP broad market index. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

OOS $SP500$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
dp unadjusted	5.19	-2.91	-0.13
dp with Shrinkage	3.12	-0.92	-0.10
$\tilde{dp}$ BCP-adjusted	4.16	5.21**	0.66
OOS $CRSP_{EW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
dp unadjusted	4.05	-0.57	-0.30
dp with Shrinkage	0.19	-0.21	-0.25
$\tilde{dp}$ BCP-adjusted	2.60	3.79**	0.20

Table E.7: Out-of-Sample return prediction, Annual data, on Fama-French Size Portfolios. The top panel presents the R^2_{MZ}, the middle panel the OOS R^2 and the bottom panel the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

R^2_{MZ}	Calibration:	1927-1965	
	OS period:	1966-2010	
	Small	Medium	Big
dp unadjusted	2.58	5.14	25.44
dp with Shrinkage	1.91	6.21	23.73
$\tilde{dp}$ BCP-adjusted	9.75	15.13	23.14
R^2_{OS}	Calibration:	1927-1965	
	OS period:	1966-2010	
	Small	Medium	Big
dp unadjusted	0.40	-1.79	-6.00
dp with Shrinkage	0.16	-0.72	-3.36
$\tilde{dp}$ BCP-adjusted	-4.15	-2.73	2.01*
Δ	Calibration:	1927-1965	
	OS period:	1966-2010	
	Small	Medium	Big
dp unadjusted	-0.21	-0.13	-0.39
dp with Shrinkage	-0.28	-0.31	-0.44
$\tilde{dp}$ BCP-adjusted	-0.72	-0.33	0.67

Table E.8: Out-of-Sample return prediction, Annual data, on Fama-French Book-to-Market Portfolios. The top panel presents the R^2_{MZ}, the middle panel the OOS R^2 and the bottom panel the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$. The prevailing historical average' coefficient of the MZ regressions were negative, hence the R^2_{MZ} for the historical average are meaningless.

R^2_{MZ}	Calibration:	1927-1965	
	OS period:	1966-2010	
	Low	Medium	High
dp unadjusted	2.59	9.49	5.48
dp with Shrinkage	7.82	6.88	4.84
$\tilde{dp}$ BCP-adjusted	12.69	7.74	14.07
R^2_{OS}	Calibration:	1927-1965	
	OS period:	1966-2010	
	Low	Medium	High
dp unadjusted	-4.66	-1.42	0.11
dp with Shrinkage	-2.69	-0.88	-0.61
$\tilde{dp}$ BCP-adjusted	4.46**	-2.60	2.94*
Δ	Calibration:	1927-1965	
	OS period:	1966-2010	
	Low	Medium	High
dp unadjusted	0.04	0.18	-0.15
dp with Shrinkage	-0.06	-0.17	-0.10
$\tilde{dp}$ BCP-adjusted	0.89	-0.58	-0.12

Table E.9: Out-of-Sample return prediction of CRSP VW Index, Annual Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

True OOS $CRSP_{VW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	5.54	-3.04	-0.10
Two States, smoothed	16.87	-6.51	-0.92
Three states, smoothed	0.00	-2.69	-0.45
Two States, unsmoothed	2.29	-5.73	-0.59
Three states, unsmoothed	0.58	-7.77	-0.63
True OOS $CRSP_{VW}$	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.32	-13.42	-0.51
Two States, smoothed	17.23	-6.82	-1.12
Three states, smoothed	3.72	-8.92	-0.97
Two States, unsmoothed	6.45	-6.05	-0.91
Three states, unsmoothed	5.62	-12.23	-1.26
True OOS $CRSP_{VW}$	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	5.80	-8.47	0.07
Two States, smoothed	0.89	-5.45	-0.39
Three states, smoothed	0.09	-8.13	-0.28
Two States, unsmoothed	0.33	-15.55	0.20
Three states, unsmoothed	0.05	-29.65	-0.34

Table E.10: Out-of-Sample return prediction of CRSP EW Index, Annual Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

True OOS $CRSP_{EW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	4.05	-0.57	-0.30
Two States, smoothed	6.50	-4.02	-0.30
Three states, smoothed	2.43	-1.46	-0.25
Two States, unsmoothed	2.37	-6.91	-0.25
Three states, unsmoothed	1.27	-4.62	-0.48
True OOS $CRSP_{EW}$	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.08	-8.30	-0.92
Two States, smoothed	13.38	-4.23	-0.22
Three states, smoothed	5.79	-4.10	-0.45
Two States, unsmoothed	7.84	-8.08	-0.46
Three states, unsmoothed	0.38	-6.65	-0.85
True OOS $CRSP_{EW}$	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.96	-0.51	-0.18
Two States, smoothed	1.23	-4.59	-0.07
Three states, smoothed	2.85	-3.47	-0.44
Two States, unsmoothed	0.46	-11.22	0.12
Three states, unsmoothed	0.05	-5.70	-0.21

Table E.11: Pseudo-OOS prediction test of CRSP VW Index, Annual Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

Pseudo OOS $CRSP_{VW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	5.54	-3.04	-0.10
Two States, smoothed	16.33	14.00	1.08
Three states, smoothed	14.37	14.05	0.86
Two States, unsmoothed	9.67	6.22	0.44
Three states, unsmoothed	4.89	2.99	1.06
Pseudo OOS $CRSP_{VW}$	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.32	-13.42	-0.51
Two States, smoothed	12.30	11.53	1.08
Three states, smoothed	13.03	10.46	0.48
Two States, unsmoothed	4.05	0.44	0.23
Three states, unsmoothed	0.79	-11.32	0.07
Pseudo OOS $CRSP_{VW}$	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	5.80	-8.47	0.07
Two States, smoothed	9.52	4.73	0.91
Three states, smoothed	15.53	15.81	1.23
Two States, unsmoothed	6.60	-4.01	0.45
Three states, unsmoothed	5.87	1.98	0.95

Table E.12: Pseudo-OOS prediction test of CRSP EW Index, Annual Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

Pseudo OOS $CRSP_{EW}$	Calibration:	1927-1965	
	OS period:	1966-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	4.05	-0.57	-0.30
Two States, smoothed	9.59	5.51	1.58
Three states, smoothed	13.81	14.59	0.86
Two States, unsmoothed	6.99	7.01	0.64
Three states, unsmoothed	0.50	2.21	0.08
Pseudo OOS $CRSP_{EW}$	Calibration:	1927-1976	
	OS period:	1977-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.08	-8.30	-0.92
Two States, smoothed	7.26	-11.82	0.41
Three states, smoothed	3.73	3.68	-0.57
Two States, unsmoothed	6.49	1.08	-0.13
Three states, unsmoothed	1.69	2.78	0.01
Pseudo OOS $CRSP_{EW}$	Calibration:	1927-1947	
	OS period:	1948-2010	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.96	-0.51	-0.18
Two States, smoothed	7.72	0.47	1.83
Three states, smoothed	11.05	12.60	0.74
Two States, unsmoothed	6.98	6.94	0.51
Three states, unsmoothed	0.19	0.03	-0.12

Table E.13: Out-of-Sample return prediction of CRSP VW Index, Quarterly Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

True OOS $CRSP_{VW}$	Calibration:	1927:3-1965:3	
	OS period:	1965:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.29	0.47	-0.03
Two States, smoothed	4.01	-0.78	-0.21
Three states, smoothed	1.37	-1.11	-0.32
Two States, unsmoothed	0.02	-0.30	-0.16
Three states, unsmoothed	1.17	0.69	-0.02
True OOS $CRSP_{VW}$	Calibration:	1927:3-1976:3	
	OS period:	1976:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	0.77	-0.98	-0.31
Two States, smoothed	3.73	-0.68	-0.19
Three states, smoothed	1.31	-1.11	-0.36
Two States, unsmoothed	0.02	-0.99	-0.20
Three states, unsmoothed	0.70	-0.18	-0.17
True OOS $CRSP_{VW}$	Calibration:	1927:3-1947:3	
	OS period:	1947:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.23	0.16	-0.09
Two States, smoothed	1.07	-1.58	-0.24
Three states, smoothed	1.63	-1.77	-0.34
Two States, unsmoothed	0.19	-0.30	0.05
Three states, unsmoothed	0.85	0.24	0.11

Table E.14: Pseudo-OOS prediction test of CRSP VW Index, Quarterly Data, using Markov Switching model with two and three regimes. The first column presents the R^2_{MZ}, the second column the OOS R^2 and the third column the difference in utility gains of using the return forecast minus the utility of using the prevailing historical average for a mean-variance investor with risk aversion parameter $\gamma = 3$ and stock index allocation limits of $[0\%, 150\%]$.

Pseudo OOS $CRSP_{VW}$	Calibration:	1927:3-1965:3	
	OS period:	1965:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.29	0.47	-0.03
Two States, smoothed	3.65	3.03	0.71
Three states, smoothed	4.16	4.25	0.95
Two States, unsmoothed	3.61	2.85	0.55
Three states, unsmoothed	0.63	0.97	0.24
Pseudo OOS $CRSP_{VW}$	Calibration:	1927:3-1976:3	
	OS period:	1976:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	0.77	-0.98	-0.31
Two States, smoothed	2.82	2.70	0.66
Three states, smoothed	3.14	2.91	0.54
Two States, unsmoothed	2.87	2.52	0.45
Three states, unsmoothed	0.50	0.82	0.16
Pseudo OOS $CRSP_{VW}$	Calibration:	1927:3-1947:3	
	OS period:	1947:6-2010:12	
	R^2_{MZ}	R^2_{OS}	Δ
Unadjusted dp	1.23	0.16	-0.09
Two States, smoothed	2.26	2.28	0.43
Three states, smoothed	3.01	2.92	0.79
Two States, unsmoothed	2.31	2.12	0.31
Three states, unsmoothed	0.10	0.68	0.14

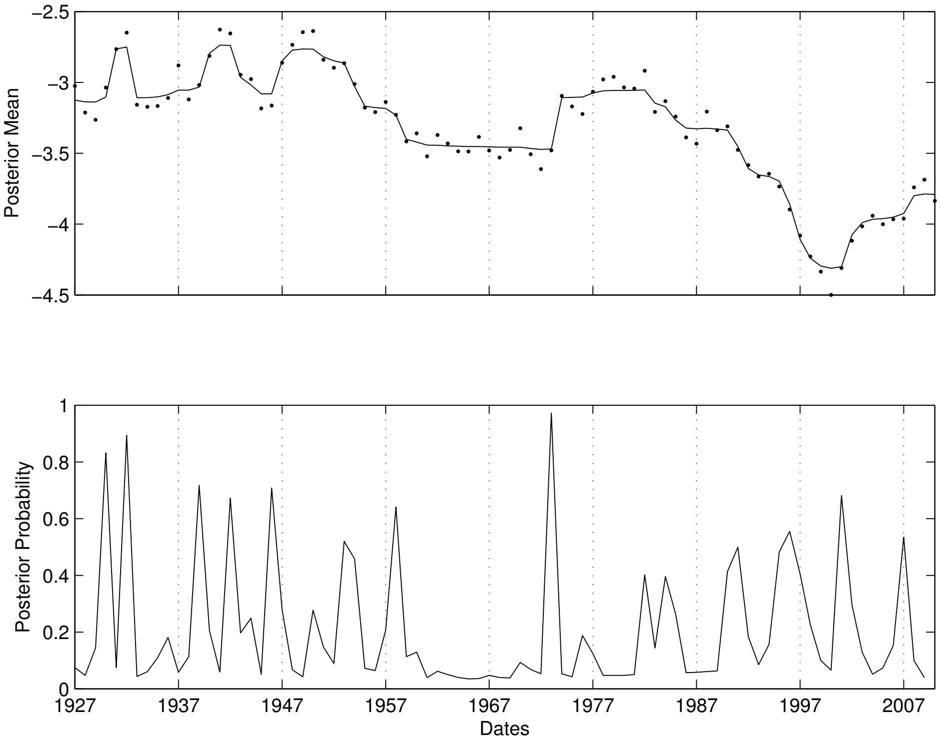

Figure E.1: Dividend-Price ratio, posterior mean and posterior change point probability series. The upper panel of the figure plots the dividend price ratio (dots) and the estimated posterior mean (straight line) using the Bayesian Change Point algorithm of Barry and Hartigan (1993). The sample period is 1927 to 2010. The persistence of the "regime" mean provide further evidence of strucutural changes in the ratio's mean.

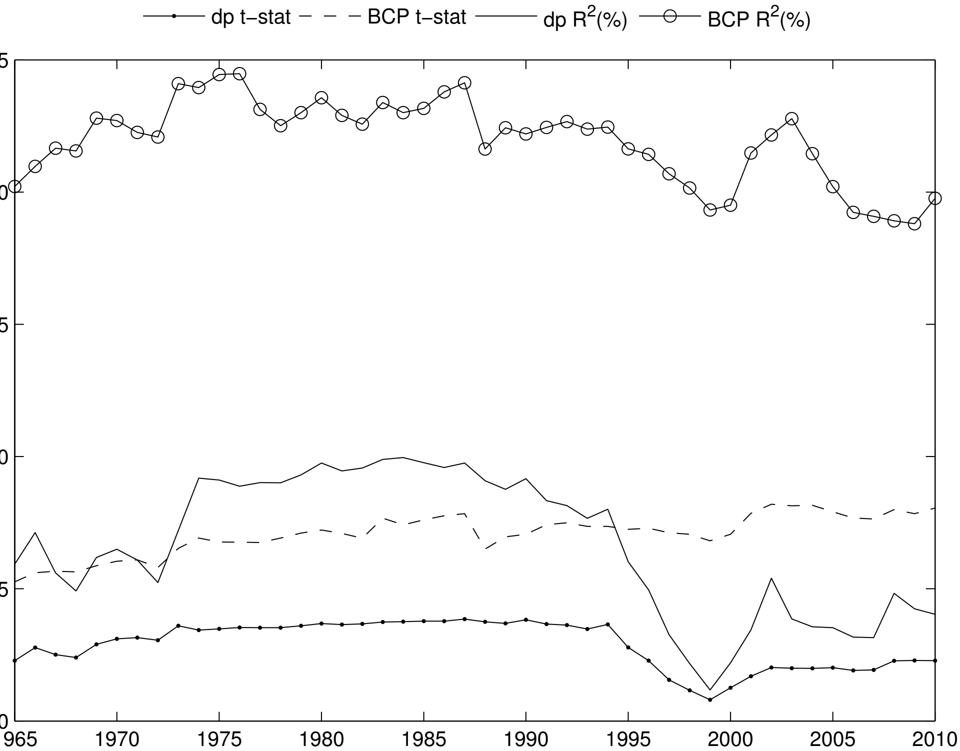

Figure E.2: Predictive regression on market return using a growing window of data with information available up to time t. The initial calibration sample is 1927 until 1965. Annual data, Broad Market CRSP index.

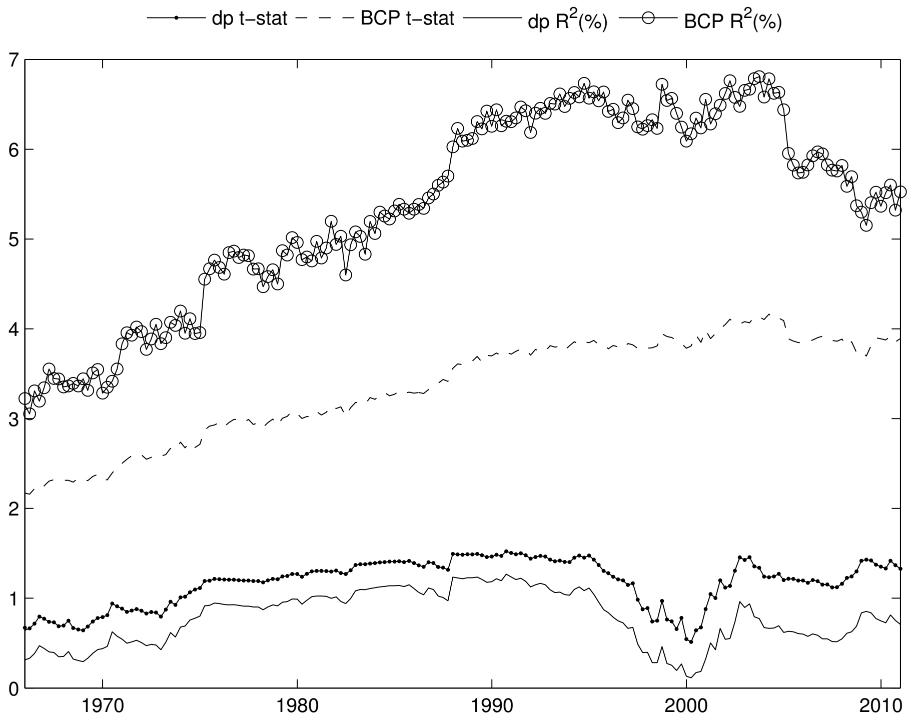

Figure E.3: Predictive regression on market return using a growing window of data with information available up to time t. The initial calibration sample is 1927 until 1965. Quarterly data, Broad Market CRSP index.

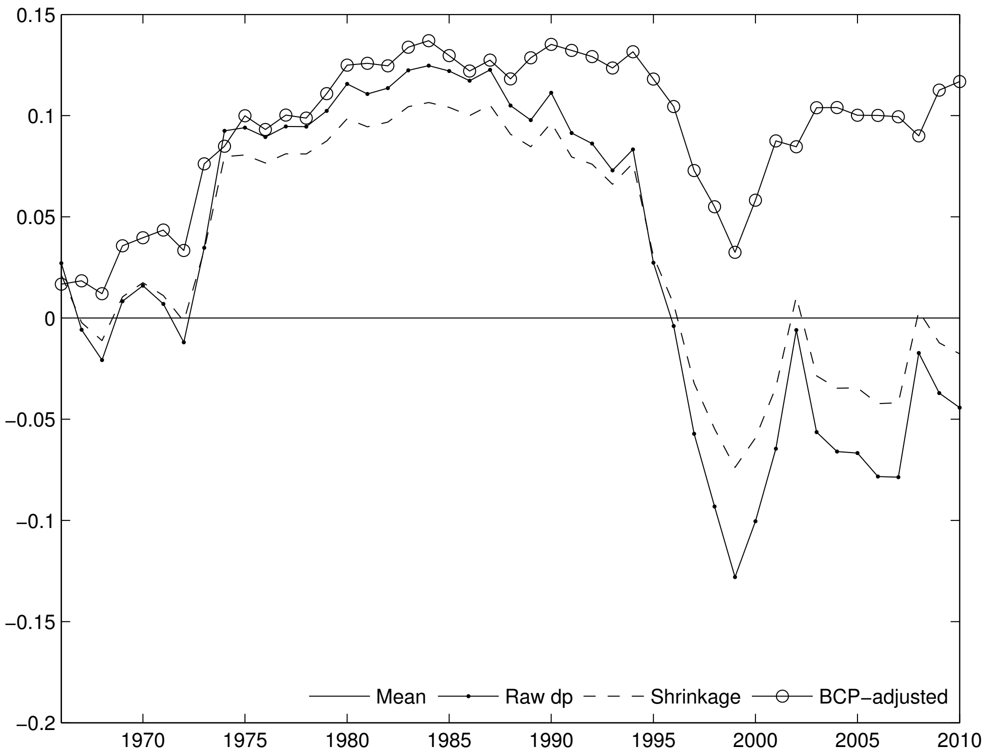

Figure E.4: Cumulative squared prediction errors of the prevailing mean minus the cumulative squared prediction error of the predictive variable from the linear historical regressions. The parameters are estimated using a growing window of data with information available up to time t to predict the market return at time t+1. The initial calibration sample has 40 data points, making the first prediction from 1965 to 1966. Annual data, Broad Market CRSP index.

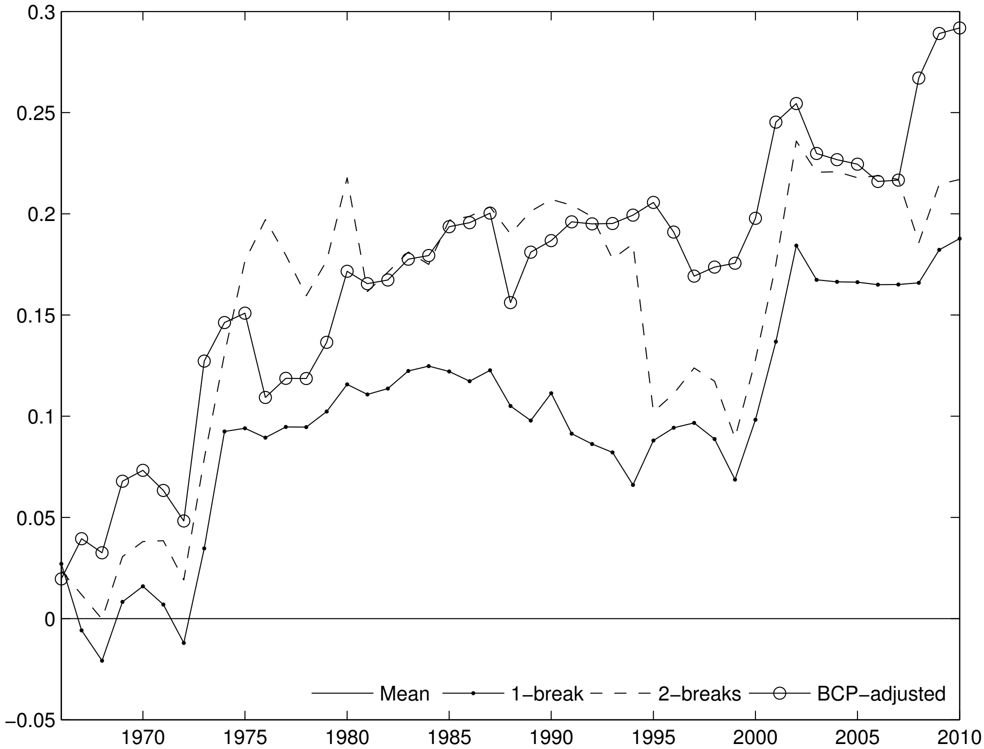

Figure E.5: Pseudo Out-of-sample test. Cumulative squared prediction errors of the prevailing mean minus the cumulative squared prediction error of the predictive variable from the linear historical regressions. Regression parameters are estimated each period using a growing window of data. However, the posterior mean of the dp ratio is estimated using the entire sample. The initial calibration sample has 40 data points, making the first prediction from 1965 to 1966. Annual data, Broad Market CRSP index.

Appendix F

Tables and Figures of Chapter 3

Table F.1: Backtests for the Growth Optimal Portfolio (GOP), the Growth Optimal Portfolio Insurance strategy (GOPI), the CPPI and the underlying assets (Cash and Stocks). The upper panel of the table presents the results for which the optimal multiplier and the allocation of the GOP were estimated as in Equation (3.23) using the sample estimate over the entire period and equal 1.6. The multiplier of the CPPI is given by Equation (3.28) and estimated using the entire sample, and equal to 3.4. The lower panel presents the results of an out-of-sample test in which the optimal multiplier and the allocation of the GOP were estimated as in Equation (3.23) using the first half of the sample (1926:01-1968:05) and the strategies performed over the second half of the available sample period (1968:06-2010:12). The optimal multiplier in this case is equal to 1.7. The multiplier of the CPPI is given by Equation (3.28) and estimated using the first half of the sample and equal to 3.4. Return stands for the annualized geometric return average, Min represents the minimum value ever attained by the portfolio for an initial value of 100 dollars of each strategy, Vol is the annualized standard deviation of returns and MDD stands for Maximum draw-down. All performance figures are presented in percentage terms.

Period: 01/1926-12/2010	GOP	GOPI	CPPI	Cash	Stocks
Return	11.27	7.88	3.58	3.71	9.62
Vol	29.50	12.73	4.60	0.88	18.86
MDD	95.56	51.45	48.46	0.00	83.72
Min	17.00	99.23	97.49	100.00	41.27
Period: 06/1968-12/2010	**GOP**	**GOPI**	**CPPI**	**Cash**	**Stocks**
Return	10.41	6.76	5.63	5.66	9.49
Vol	27.03	6.63	6.01	0.87	16.16
MDD	72.20	34.73	29.26	0.00	51.45
Min	39.17	100.00	99.53	100.00	70.28

Table F.2: Summary Statistics of Alternative Investment Indices. The upper panel presents the annualized arithmetic average and standard deviation of returns and the lower panel the cross-correlation matrix of the overall sample period.

	ConvArb	CTAs	Distress	EM	MNeutral	EventD	FixArb	GMacro	LSequity	MergArb	RelVal	ShortS	FoF
μ	0.085	0.076	0.104	0.104	0.070	0.097	0.059	0.090	0.095	0.080	0.083	0.024	0.069
σ	0.067	0.086	0.063	0.129	0.030	0.062	0.048	0.057	0.076	0.037	0.045	0.186	0.061
AR(1)	0.586	0.023	0.509	0.323	0.273	0.388	0.508	0.052	0.265	0.311	0.454	0.141	0.332

	ConvArb	CTAs	Distress	EM	MNeutral	EventD	FixArb	GMacro	LSequity	MergArb	RelVal	ShortS	FoF
ConvArb	1.00	-0.04	0.72	0.57	0.49	0.71	0.78	0.39	0.57	0.55	0.85	-0.33	0.63
CTAs	-0.04	1.00	-0.04	0.02	0.18	-0.01	-0.00	0.51	0.09	-0.01	-0.02	0.11	0.15
Distress	0.72	-0.04	1.00	0.80	0.58	0.93	0.66	0.55	0.76	0.65	0.84	-0.60	0.82
EM	0.57	0.02	0.80	1.00	0.50	0.83	0.52	0.69	0.81	0.62	0.76	-0.69	0.86
MNeutral	0.49	0.18	0.58	0.50	1.00	0.63	0.41	0.56	0.64	0.55	0.62	-0.32	0.69
EventD	0.71	-0.01	0.93	0.83	0.63	1.00	0.59	0.62	0.88	0.84	0.90	-0.68	0.89
FixArb	0.78	-0.00	0.66	0.52	0.41	0.59	1.00	0.43	0.44	0.38	0.73	-0.17	0.54
GMacro	0.39	0.51	0.55	0.69	0.56	0.62	0.43	1.00	0.72	0.49	0.56	-0.42	0.79
LSequity	0.57	0.09	0.76	0.81	0.64	0.88	0.44	0.72	1.00	0.74	0.82	-0.76	0.92
MergArb	0.55	-0.01	0.65	0.62	0.55	0.84	0.38	0.49	0.74	1.00	0.75	-0.55	0.73
RelVal	0.85	-0.02	0.84	0.76	0.62	0.90	0.73	0.56	0.82	0.75	1.00	-0.55	0.81
ShortS	-0.33	0.11	-0.60	-0.69	-0.32	-0.68	-0.17	-0.42	-0.76	-0.55	-0.55	1.00	-0.66
FoF	0.63	0.15	0.82	0.86	0.69	0.89	0.54	0.79	0.92	0.73	0.81	-0.66	1.00

Table F.3: Ranking of Satellite assets according to their geometric mean; autocorrelation coefficient; the value of a DCS portfolio rebalanced every δ/m move in its exposure for $\delta = 0.05$ and with exposure limits $\{0 \leq$e$\leq 1\}$; its theoretical value assuming continuous rebalancing and no short-settling limits and its leveraged geometric mean for three different values of $m = \{2, 3, 5\}$. The Core s the riskless asset.

Rank	Geometric Mean	AC(1)	$G^{m=2}$	$\gamma^{cushion}_{m=2}$	$V_{m=2}$	$DCS_{m=2}$
1	Distress	ConvArb	Distress	Distress	Distress	Distress
2	EM	Distress	EventD	EventD	EventD	EventD
3	EventD	FixArb	LSequity	LSequity	LSequity	LSequity
4	LSequity	RelVal	EM	GMacro	GMacro	GMacro
5	GMacro	EventD	GMacro	EM	EM	EM
6	RelVal	FoF	RelVal	RelVal	RelVal	RelVal
7	ConvArb	EM	ConvArb	ConvArb	ConvArb	ConvArb
8	MergArb	MergArb	MergArb	MergArb	MergArb	MergArb
9	CTAs	MNeutral	MNeutral	MNeutral	MNeutral	MNeutral
10	MNeutral	LSequity	CTAs	CTAs	CTAs	CTAs
11	FoF	ShortS	FoF	FoF	FoF	FoF
12	FixArb	GMacro	FixArb	FixArb	FixArb	FixArb
13	ShortS	CTAs	ShortS	ShortS	ShortS	ShortS
Rank	Geometric Mean	AC(1)	$G^{m=3}$	$\gamma^{cushion}_{m=3}$	$V_{m=3}$	$DCS_{m=3}$
1	Distress	ConvArb	Distress	Distress	Distress	Distress
2	EM	Distress	EventD	EventD	EventD	EventD
3	EventD	FixArb	LSequity	LSequity	LSequity	LSequity
4	LSequity	RelVal	GMacro	GMacro	GMacro	GMacro
5	GMacro	EventD	RelVal	RelVal	RelVal	RelVal
6	RelVal	FoF	EM	EM	EM	MergArb
7	ConvArb	EM	MergArb	MergArb	MergArb	EM
8	MergArb	MergArb	ConvArb	ConvArb	ConvArb	ConvArb
9	CTAs	MNeutral	MNeutral	MNeutral	MNeutral	MNeutral
10	MNeutral	LSequity	CTAs	CTAs	CTAs	CTAs
11	FoF	ShortS	FoF	FoF	FoF	FoF
12	FixArb	GMacro	FixArb	FixArb	FixArb	FixArb
13	ShortS	CTAs	ShortS	ShortS	ShortS	ShortS
Rank	Geometric Mean	AC(1)	$G^{m=5}$	$\gamma^{cushion}_{m=5}$	$V_{m=5}$	$DCS_{m=5}$
1	Distress	ConvArb	Distress	Distress	Distress	Distress
2	EM	Distress	EventD	EventD	EventD	EventD
3	EventD	FixArb	GMacro	GMacro	GMacro	LSequity
4	LSequity	RelVal	LSequity	LSequity	LSequity	GMacro
5	GMacro	EventD	RelVal	RelVal	RelVal	RelVal
6	RelVal	FoF	MergArb	MergArb	MergArb	ConvArb
7	ConvArb	EM	ConvArb	ConvArb	ConvArb	MergArb
8	MergArb	MergArb	MNeutral	MNeutral	MNeutral	MNeutral
9	CTAs	MNeutral	EM	EM	EM	CTAs
10	MNeutral	LSequity	FoF	FoF	FoF	FoF
11	FoF	ShortS	CTAs	CTAs	CTAs	FixArb
12	FixArb	GMacro	FixArb	FixArb	FixArb	EM
13	ShortS	CTAs	ShortS	ShortS	ShortS	ShortS

Table F.4: Ranking of Satellite portfolios according to their geometric mean, the value of a DCS portfolio rebalanced every δ/m move in its exposure with exposure limits $\{0 \leq e \leq 1\}$ and $\delta = 0.05$, its theoretical value assuming continuous rebalancing and no short-settling limits and the growth rate of its Cushion for three different values of $m = \{2, 3, 5\}$. The Core asset is the Fixed-Income Arbitrage index.

Rank	Geometric Mean	$\gamma^{cushion}_{m=2}$	$V_{m=2}$	$DCS_{m=2}$
1	Distress	Distress	Distress	Distress
2	EM	EventD	EventD	EventD
3	EventD	LSequity	LSequity	LSequity
4	LSequity	EM	EM	EM
5	GMacro	GMacro	GMacro	GMacro
6	RelVal	RelVal	RelVal	RelVal
7	ConvArb	ConvArb	ConvArb	ConvArb
8	MergArb	MergArb	MergArb	MergArb
9	CTAs	MNeutral	MNeutral	MNeutral
10	MNeutral	CTAs	CTAs	CTAs
11	FoF	FoF	FoF	FoF
12	ShortS	ShortS	ShortS	ShortS
Rank	Geometric Mean	$\gamma^{cushion}_{m=3}$	$V_{m=3}$	$DCS_{m=3}$
1	Distress	Distress	Distress	Distress
2	EM	EventD	EventD	EventD
3	EventD	LSequity	LSequity	LSequity
4	LSequity	GMacro	GMacro	GMacro
5	GMacro	EM	EM	RelVal
6	RelVal	RelVal	RelVal	EM
7	ConvArb	ConvArb	ConvArb	ConvArb
8	MergArb	MergArb	MergArb	MergArb
9	CTAs	MNeutral	MNeutral	MNeutral
10	MNeutral	FoF	FoF	FoF
11	FoF	CTAs	CTAs	CTAs
12	ShortS	ShortS	ShortS	ShortS
Rank	Geometric Mean	$\gamma^{cushion}_{m=5}$	$V_{m=5}$	$DCS_{m=5}$
1	Distress	Distress	Distress	Distress
2	EM	EventD	EventD	EventD
3	EventD	GMacro	GMacro	LSequity
4	LSequity	LSequity	LSequity	GMacro
5	GMacro	RelVal	RelVal	RelVal
6	RelVal	ConvArb	ConvArb	ConvArb
7	ConvArb	MergArb	MergArb	MergArb
8	MergArb	EM	EM	MNeutral
9	CTAs	MNeutral	MNeutral	FoF
10	MNeutral	FoF	FoF	EM
11	FoF	CTAs	CTAs	CTAs
12	ShortS	ShortS	ShortS	ShortS

Table F.5: This table presents the distribution summary of the standard multiplier (equation (3.28)) from 100 Backtest exercises for which the starting dates are selected randomly and performed over 5, 10 and 20 years horizons. The upper panel presents the result of estimating the multiplier using the same sample period of each historical scenario (perfect foresight). The lower panel presents the multiplier estimation using all data available before the starting date of each of the backtest scenarios to estimate the multiplier (sample).

Perfect	Max Multiplier Range		
	Min	Median	Max
5y	4.17	9.20	20.37
10y	4.17	7.48	16.36
20y	4.17	6.13	11.40
Sample	**Max Multiplier Range**		
	Min	Median	Max
5y	3.41	3.41	3.41
10y	3.41	3.41	3.41
20y	3.41	3.41	3.41

Table F.6: This table presents the distribution summary of the standard multiplier (equation (3.23)) from 100 Backtest exercises for which the starting dates are selected randomly and performed over 5, 10 and 20 years horizons. The upper panel presents the result of estimating the multiplier using the same sample period of each historical scenario (perfect foresight). The middle panel presents the multiplier estimation using the latest 10 years of data available at the starting date of each backtest scenario. The lower panel present the result corresponding to the dynamic estimate of the optimal multiplier. The dynamic multiplier uses all available stock returns at the starting date of each simulated scenario to fit a Dynamic Conditional Correlation model and forecast the conditional variance and covariance parameters. The expected stock returns are estimated with a 10 year moving average of past returns.

Perfect	Optimal Multiplier Range		
	Min	Median	Max
5y	1.00	2.55	17.08
10y	1.00	2.84	12.68
20y	1.00	2.83	9.21
Sample	**Optimal Multiplier Range**		
	Min	Median	Max
5y	1.00	2.68	10.34
10y	1.00	3.07	13.50
20y	1.00	1.97	13.00
Dynamic	**Optimal Multiplier Range**		
	Min	Median	Max
5y	1.00	3.04	12.14
10y	1.00	3.10	12.41
20y	1.00	2.90	12.01

Table F.7: Percentage of times that the Optimal Multiplier is lower or equal than the Standard multiplier. This table displays the percentage of times that the Optimal Multiplier is lower or equal than the Standard multiplier for each of the different estimation methodologies over 100 randomly selected Backtest exercises performed over 5, 10 and 20 years horizons. The upper panel compares the optimal multiplier estimated using the same sample period of each historical scenario (perfect foresight). The first column compares the optimal multiplier with the standard multiplier estimated using the entire sample period (denoted $m_a ll$) while the second column with the standard multiplier estimated using the same sample period of each historical scenario (perfect foresight). The middle panel compares the optimal multiplier estimated using the latest 10 years of data available at the starting date of each backtest scenario with the standard multiplier estimated using the whole sample period (first column) and with the standard multiplier estimated using all available sample before the starting date of each Backtest (second column), denoted $\hat{m}_p eriod$. The lower panel compares the optimal dynamic multiplier estimated using all available stock returns at the starting date of each simulated scenario with $m_a ll$ (first column) and $\hat{m}_p eriod$ (second column).

Perfect	Percentage of Lower	Multiplier
	$m^* <= m_{all}$	$m^* <= m_{period}$
5y	0.58	0.95
10y	0.66	0.99
20y	0.71	0.96
Sample	**Percentage of Lower**	**Multiplier**
	$\hat{m}^* <= m_{all}$	$\hat{m}^* <= \hat{m}_{period}$
5y	0.70	0.64
10y	0.68	0.61
20y	0.69	0.68
Dynamic	**Percentage of Lower**	**Multiplier**
	$\hat{m}_t^* <= m_{all}$	$\hat{m}_t^* <= \hat{m}_{period}$
5y	0.71	0.57
10y	0.66	0.55
20y	0.70	0.59

Table F.8: Optimal Multiplier's portfolio over-performance probability with Perfect Foresight on parameter values and with Sample estimates. This table displays the over-performance probability of the strategy defined by the optimal multiplier over the CPPI strategy using the standard multiplier over 100 randomly selected Backtest exercises in the available sample period and performed over 5, 10 and 20 years horizons. The upper panel compares the optimal multiplier estimated using the same sample period of each historical scenario (perfect foresight). The first column compares the optimal multiplier with the standard multiplier estimated using the entire sample period (denoted m_all) while the second column with the standard multiplier estimated using the same sample period of each historical scenario (perfect foresight). The middle panel compares the optimal multiplier estimated using the latest 10 years of data available at the starting date of each backtest scenario with the standard multiplier estimated using the whole sample period (first column) and with the standard multiplier estimated using all available sample before the starting date of each Backtest (second column), denoted $\hat{m}_period$. The lower panel compares the optimal dynamic multiplier estimated using all available stock returns at the starting date of each simulated scenario with m_all (first column) and $\hat{m}_period$ (second column).

Perfect	Over Performance -	Probability
	m_{all} =4.22	m_{period}
5y	0.850	0.670
10y	0.680	0.540
20y	0.660	0.510
Sample	Over Performance -	Probability
	m_{all} =4.22	$\hat{m}_{period}$
5y	0.480	0.470
10y	0.430	0.420
20y	0.360	0.310
Dynamic	Over-performance -	Probability
	m_{all} =4.22	$\hat{m}_{period}$
5y	0.640	0.690
10y	0.620	0.720
20y	0.670	0.720

Table F.9: This table displays the average return of the strategies defined by the standard and optimal multipliers over 100 randomly selected Backtest exercises in the available sample period and performed over 5, 10 and 20 years horizons. The first column presents the results corresponding to the optimal multiplier (either static or dynamic). The second column presents the results corresponding to the standard multiplier estimated using the entire sample period (denoted $m_a ll$) while the third column presents the results corresponding to the standard multiplier estimated using the same sample period of each historical scenario (perfect foresight) or the data available at the starting date of each historical scenario (sample estimate).

Perfect		Average Return	
	m^*	m_{all} =4.22	m_{period}
5y	0.093	0.080	0.081
10y	0.091	0.089	0.090
20y	0.091	0.090	0.090
Sample		**Average Return**	
	$\hat{m}^*$	m_{all} =4.22	$\hat{m}_{period}$
5y	0.070	0.080	0.075
10y	0.077	0.089	0.085
20y	0.076	0.090	0.093
Dynamic		**Average Return**	
	$\hat{m}_t^*$	m_{all} =4.22	$\hat{m}_{period}$
5y	0.080	0.080	0.075
10y	0.091	0.089	0.085
20y	0.096	0.090	0.093

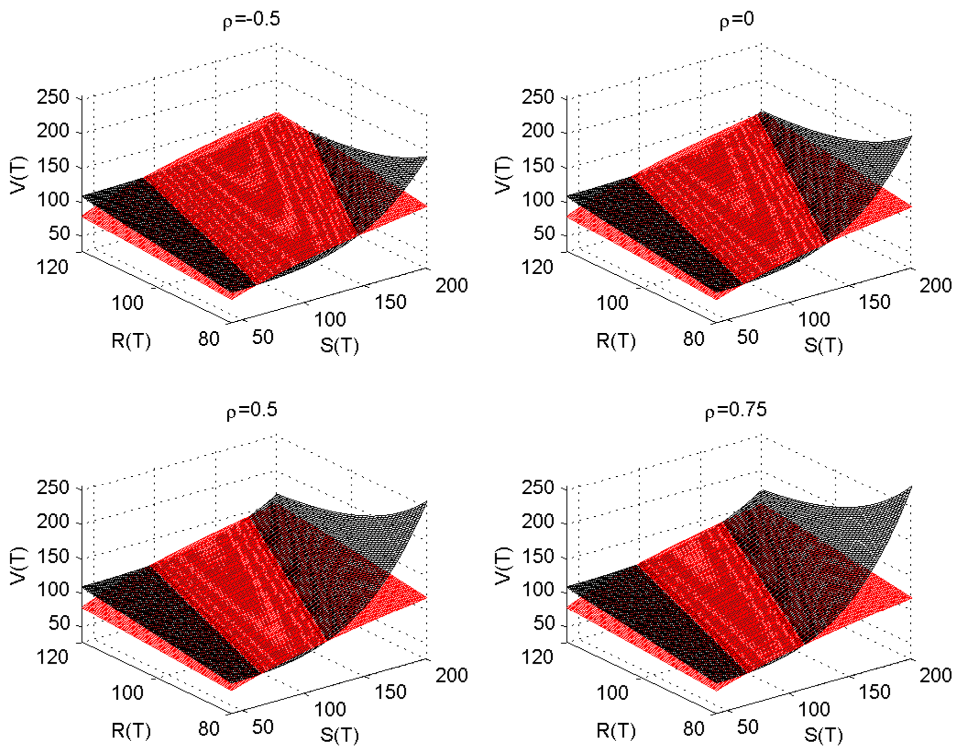

Figure F.1: Impact of assets' correlation in Portfolio's value after 5 years with $\sigma_S = 0.15$, $\sigma_R = 0.05$. The black surface draws the end of period value of a CPPI strategy with $m = 4$ and $k = 0.9$ as given by Equation (3.16). The red surface draws the end of period value of a Fixed-Mix Strategy as given by Equation (3.15) with the same initial allocation: $\pi = 0.4$. The upper panels of the figure correspond to negative and null correlation, i.e. $\rho = \{-0.5, 0\}$ and the lower panels to positive correlations, i.e. $\rho = \{0.5, 0.75\}$.

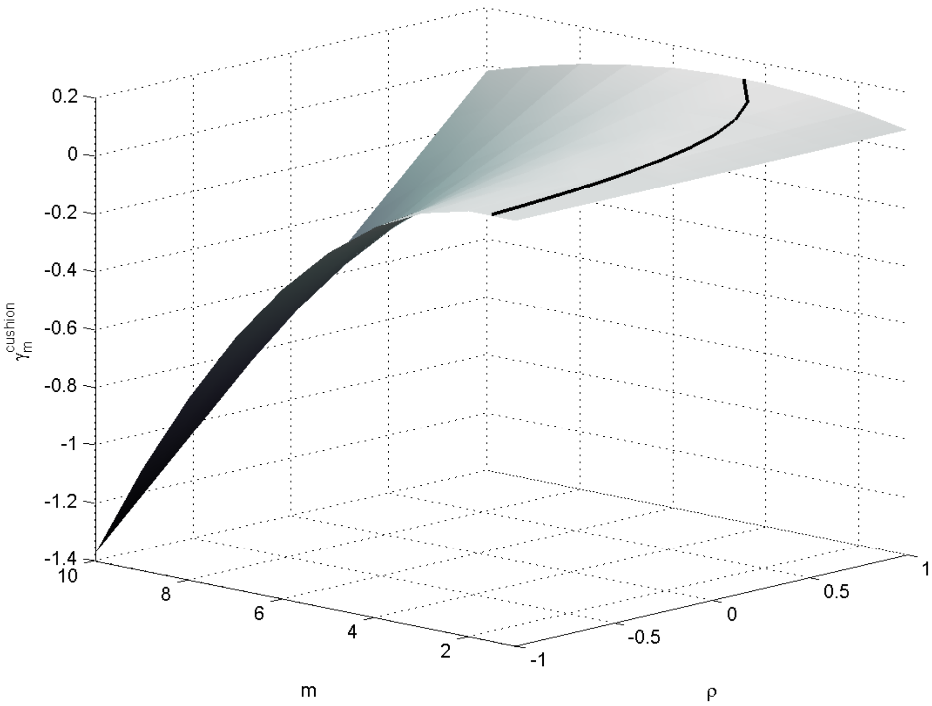

Figure F.2: Cushion Growth Rate, Correlation and Optimal Multiplier. The dark line represents the value of the Cushion's growth rate for the optimal multiplier values corresponding to the different levels of the correlation coefficient ρ. The optimal multiplier is given by Equation (3.4) and the following parameter values: $\mu_S = 0.08$, $\sigma_S = 0.15$, $\mu_R = 0.03$, $\sigma_R = 0.05$ and $\rho = [-1, 1]$. The surface represents the Cushion's growth rate for different multiplier values, i.e. $m = [1, 10]$ and correlations. This figure illustrates that for uncorrelated assets, the choice of the multiplier becomes critical (surface) and the optimal multiplier increases with correlation (dark line), everything else equal.

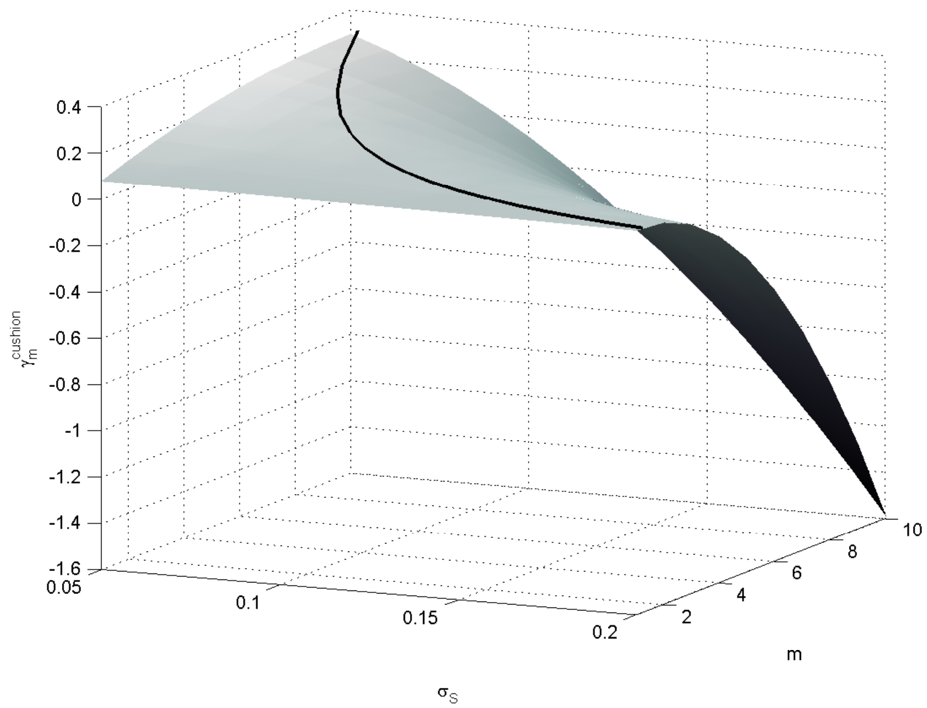

Figure F.3: Cushion Growth Rate, Volatility and Optimal Multiplier. The dark line represents the value of the Cushion's growth rate for the optimal multiplier values corresponding to the different levels of the Satellite's volatility σ_S. The optimal multiplier is given by Equation (3.4) and the following parameter values: $\mu_S = 0.08$, $\sigma_S = [0.05, 0.2]$, $\mu_R = 0.03$, $\sigma_R = 0.05$ and $\rho = -0.025$. The surface represents the Cushion's growth rate for different multiplier values, i.e. $m = [1, 10]$ and volatilities. This figure illustrates that for highly volatile assets, the choice of the multiplier becomes critical (surface) and the optimal multiplier decreases with the volatility (dark line), everything else equal.

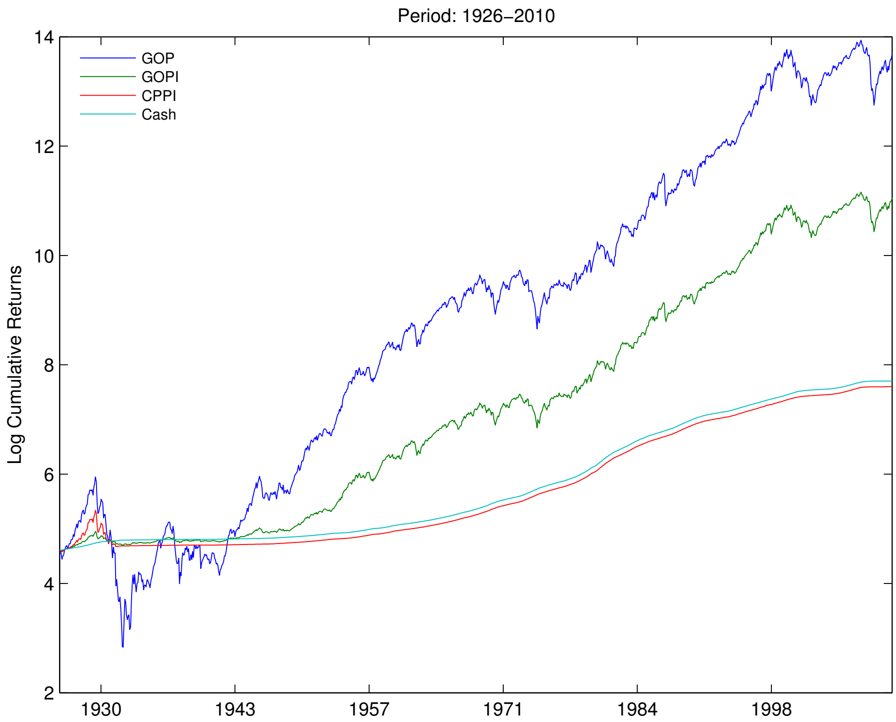

Figure F.4: The green line represents the log of the cumulative returns of the Growth Optimal Portfolio Insurance strategy (GOPI), the red line is the log of the cumulative returns of the CPPI, the dark blue line corresponds to the Growth Optimal Portfolio (GOP) and the light blue corresponds to Cash. The optimal multiplier and the allocation of the GOP were estimated as in Equation (3.23) using the sample estimates of the entire sample period. The multiplier of the CPPI is given by Equation (3.28) and estimated using the entire sample (1925-2010).

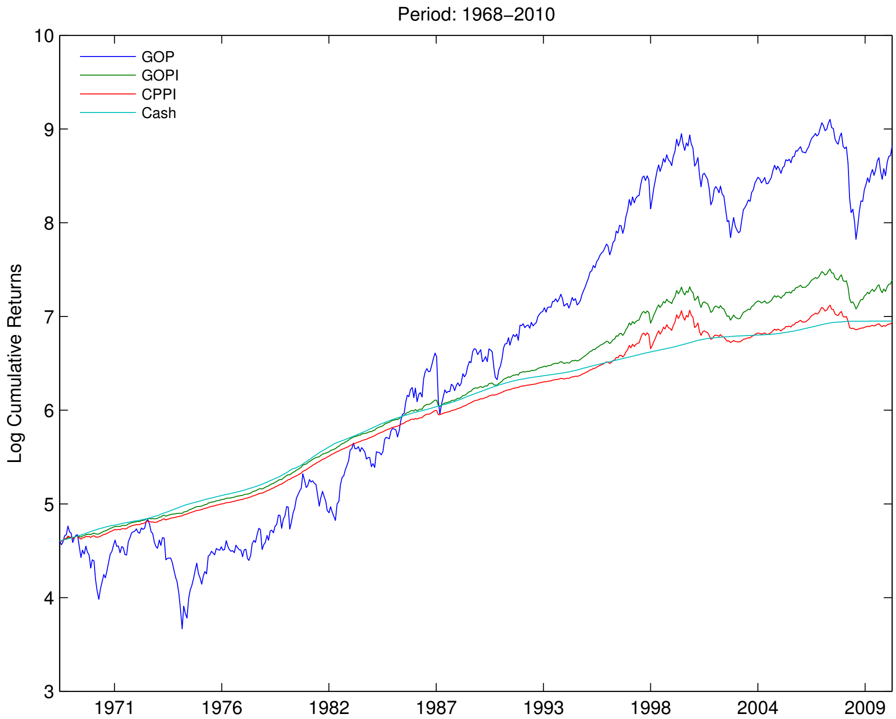

Figure F.5: The green line represents the log of the cumulative returns of the Growth Optimal Portfolio Insurance strategy (GOPI), the red line is the log of the cumulative returns of the CPPI, the dark blue line corresponds to the Growth Optimal Portfolio (GOP) and the light blue corresponds to Cash. This is an out-of-sample test in which the optimal multiplier and the allocation of the GOP were estimated as in Equation (3.23) using the first half of the sample (1926:01-1968:05) and the strategies performed over the second half of the available sample period (1968:06-2010:12). The multiplier of the CPPI is given by Equation (3.28) and estimated using the first half of the sample.