Two Essays in Empirical Asset Pricing

Flavio Nardi

October 21, 2017

A thesis submitted for the degree of Master of Philosophy
of The Australian National University

Declaration

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma in any university or equivalent institution, and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Flavio Nardi

Acknowledgements

I am grateful to the Australian National University for providing financial assistance during my post–graduate studies in finance, and to the Research School of Finance, Actuarial Studies and Statistics for the academic support and the opportunity to teach tutorial sessions. I tried my best to contribute original research to the field of finance.

I also wish to thank my panel members, Takeshi Yamada, Raymond Liu, and Phong Ngo, for their patience and guidance along the journey.

Abstract

This thesis includes two research papers in the area of empirical asset pricing. In the first research paper titled "Option implied moments and risk aversion", under reasonable assumptions, I provide empirical evidence that index options implied higher moments can predict the index returns and Sharpe ratio. Specifically, I present a method to recover option implied subjective moments of the S&P500 index under the assumption of no arbitrage and logarithmic utility. This result adds further evidence to the extensive finance literature that claims that market returns are predictable. In the second research paper titled "Expected returns: systematic risk or firm characteristics" I provide empirical evidence that expected returns can be viewed as determined by the exposure of firm returns to systematic factors that are based on firm characteristics, and not directly to the cross–sectional differences in the firm characteristics. This result addresses an ongoing debate within the empirical asset pricing literature as to whether the cross–section of expected returns is "explained" by the loadings to systematic factors or by differences in firm characteristics. The evidence I provide supports the loading to systematic factors story, consistent with the consumption asset pricing model.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis I present two research papers in the area of empirical asset pricing. The two papers address two different research questions based on different theory, data, and econometric methodologies.

In the first research paper titled "Option implied moments and risk aversion", under reasonable assumptions, I provide empirical evidence that index options implied higher moments can predict the index returns and Sharpe ratio. Specifically, I present a method to recover option implied subjective moments of the S&P500 index under the assumption of no arbitrage and logarithmic utility. Using index options prices and return data, I test the logarithmic utility assumption and obtain risk aversion estimates not statistically different from one at investment horizons of three to nine months. Under logarithmic utility, I show that the recovered subjective variance has forecasting power controlling for past realized variance. Interestingly, the risk neutral variance is larger than the subjective variance over the entire sample, an empirical fact that quantifies the implied variance premium for a log utility investor. Lastly, I also find that the forward looking Sharpe ratio implied by option prices has forecasting power; this finding can be adopted as a risk–adjusted market timing indicator to improve the return performance of either a passive indexing or a diversidied portfolio investment strategy. For example, as a long term investor would rebalance their portfolio periodically to optimize or maintain their asset allocation targets (see for example, Ang [2014]), they could use the option implied

Sharpe ratio as a "gauge" of the overall market *price level*. As such, they could take advantage of periods where there is a particurarly high expected Sharpe ratio on the market to buy more of the market index when it is at lower valuation levels. Thus, this gauge serves as a reinforcing mechanism to buy low and sell high for periodic portfolio rebabalancing.

The second research paper titled "Expected returns: systematic risk or firm characteristics" provides empirical evidence that expected returns can be viewed as determined by the exposure of firm returns to systematic factors that are based on firm characteristics, and not directly to the cross–sectional differences in the firm characteristics. This result addresses an ongoing debate within the empirical asset pricing literature as to whether the cross–section of expected returns is "explained" by the loadings to systematic factors or by differences in firm characteristics.

In this paper, I utilize cross–sectional weighted least square regressions to extract market value weighted zero–cost portfolios that are based on firm characteristics from both individual firm excess returns and risk–adjusted excess returns. The analysis shows that the value weighted size and book–to–market cross–sectional regression zero–cost portfolios have effectively the same explanatory power as the Fama French 3 factor model constructed from portfolio sorts. I compare their performance using the Fama French 25 size and book–to–market portfolios as test assets, and present evidence that the zero–cost portfolios extracted from cross–sectional regressions of risk–adjusted excess returns on firm characteristics do not have any significant explanatory power when tested on the combined Fama French 25 size and book–to–market and 30 industry portfolios. This empirical finiding implies that individual firm risk–adjusted returns do not exhibit any statistically significant cross–sectional difference related to firm characteristics.

Consistent with such interpretation, the risk–adjusted value weighted cross–sectional factors have very low correlation with the original Fama French 3 factors. Moreover, this empirical result provides suggestive evidence that risk adjustment of excess returns is effective in capturing the cross–sectional difference in firms'

expected returns. Finally, I perform a formal time series test to compare the cross-sectional regression factors based on excess returns and risk–adjusted excess returns; I conclude that there is very little correlation between the two, indicating that loadings on the factors capture the cross–section of expected returns.

Chapter 2

Option Implied Moments and Risk Aversion

2.1 Introduction

A challenging task in asset pricing is drawing inference about aggregate investors' beliefs on the ex–ante distribution of market returns. Option trading provides information on the forward looking risk neutral return distribution of the underlying security (see for example, Breeden and Litzenberger [1978], who indicate that state prices can be recovered from the price of traded options). Within the option contracts literature, studies focus on the development of parametric and/or non–parametric methods for fitting the risk neutral density from option prices (Rubinstein [1994], Jackwerth and Rubinstein [1996], Aït-Sahalia and Lo [1998], Figlewski [2010], and references therein). Specifically, some researchers test the return forecasting ability of the risk neutral density to estimate the representative agent's risk aversion implied in option prices (Jackwerth [2000], and Bliss and Panigirtzoglou [2004] (BP)). In order to recover the risk aversion coefficient implied by index option trading, one must make assumptions on the representative investor's utility functional form (investor's preferences) and on the ex–ante representative agent's subjective probability distribution of returns (investor's beliefs). Preferences are commonly specified

in the form of power or exponential utility (as in BP), and beliefs are formed based on the realized return distribution (under the additional assumption that the return distribution is stationary).

Theoretically, considering a two period finite payoff space with s states, the risk neutral probability $\pi^*(s)$ is related to the subjective probability $\pi(s)$ via the stochastic discount factor $M(s)$ and the gross risk free return R_f, as follows:

$$\pi^*(s) = R_f M(s)\pi(s). \tag{2.1}$$

This risk neutral probability, $\pi^*(s)$, can be recovered from option prices. Armed with the risk neutral density, Equation (2.1) identifies the ex–ante subjective probability (without using past realized returns). Yet, to do this, we have to pin down investor's preferences, that is, we must specify the form for the stochastic discount factor $M(s)$. Despite the difficulty in specifying a stochastic discount factor, there are a number of studies examining the information content implied in the risk neutral density. One such example, namely Conrad et al. [2013], finds that the volatility, skewness and kurtosis implied in the risk neutral density of equity options have a significant relation with future realized returns[1]. Although there is growing evidence showing that the risk neutral moments reliably predict future returns (Conrad et al. [2013]), the recovery of the option implied subjective ex-ante market return distribution remains a challenge. A recent theoretical result claiming recovery of the subjective probability distribution from index option prices without an explicit specification of the representative investor's utility function (preferences) is the recovery theorem obtained by Ross [2015]. Under a mildly restrictive assumption on the price process, Ross presents a method for the recovery of the ex–ante subjective distribution of returns, and provides some preliminary empirical evidence using S&P500 index options. Despite this, Borovička et al. [2016] present theoretical and empirical evidence that implies misspecification in the remarkable recovery result of

[1]The Federal Reserve Bank of Minneapolis (see https://www.minneapolisfed.org/banking/mpd, market based probabilities) posts weekly updates with option implied risk neutral densities of various commodities and indexes as policy and decision making guidance.

Ross [2015].

In light of the existing literature, the aim of my paper is to recover *bounds* on the moments of the subjective distribution by using index options and assuming logarithmic utility preferences. Although I am unable to claim that the logarithmic utility is the correct specification, I provide empirical evidence that such preference specification is not rejected by a formal econometric test of the combined index return and option price data. Since options have finite maturity, an investor that adds option contracts to his portfolio implicitly specifies a finite investment horizon. This is particularly true for S&P500 index option investors, as the increased demand for out–of–the money put options is driven by the need for portfolio insurance against market downside risk. The myopic investment behavior implied by the logarithmic utility specification might be a fair approximation of the representative investor in index options, and consistent with the observation that the largest volume of the CBOE traded S&P500 index option contracts have maturity less than six months.

I start with the standard assumption that asset prices are determined under the no–arbitrage condition and derive an equation that relates subjective moments with risk neutral moments. I extend the first moment bound result of Martin [2017] to higher moments of the subjective distribution implied by traded S&P500 index options. I show that all moments of the subjective distribution can be decomposed into two terms. The first term is an observable discounted risk neutral moment and the second is a covariance term that involves the stochastic discount factor. Since the logarithmic utility specification makes the covariance term zero (thus obtaining equations of subjective moments only in terms of observable discounted risk neutral quantities), I address empirically the validity of such an assumption. To this end, I present an econometric test for the logarithmic utility assumption that uses both index return and option data, and I estimate the option implied investor's risk aversion under the power utility preference specification using the Generalized Method of Moments (GMM) of Hansen [1982]. I apply an overidentified GMM specification to test whether the option implied risk aversion coeffcient is statistically different

from 1.

In this study, I implement moment conditions in a time series test and use prices from traded options on the S&P500 and Russell 2000 indexes to estimate the risk aversion coefficient under the power utility function specification. The proposed method overcomes the short horizon limitation of the estimation method in BP and, further, does not require fitting the option implied risk neutral density. I use the results of Carr and Madan [2001], Dennis and Mayhew [2002], and Bakshi et al. [2003], who provide ample evidence that option implied risk neutral moments can be recovered using a model free approach with negligible approximation error. Building on this literature, I combine the subjective moment condition with the fundamental valuation equation to estimate the risk aversion coefficient with power utility using options and realized returns time series data, and test the logarithmic utility hypothesis (risk aversion equal to 1) at investment horizons spanning between 1 and 9 months. I fail to reject that the option implied risk aversion coeffiecient is 1 at horizons between 3 and 9 months.

Given that the overidentified GMM power utility specification is not rejected with risk aversion coeffcient estimates not different from 1, I use the exact expressions for the subjective moments and test the ability of the logarithmic utility conditional variance to forecast future realized variance, and the ex–ante Sharpe Ratio on the subsequent realized Sharpe Ratio. This is of interest in the asset pricing literature because the Sharpe Ratio provides an important market timing signal with significant forecasting power. Finally, I also provide evidence of the difference between risk neutral and subjective variance and skewness at investment horizons up to 9 months, and of the forecasting power of the option implied logarithmic utility variance risk premium.

The estimation of a representative agent's risk aversion coefficient has been a central challenge that financial economist have investigated extensively over the past four decades. The range of estimates for this important parameter extends between 1 and about 50 (see, Aït-Sahalia and Lo [2000] and BP for a summary of

the values found in the literature). Although there are numerous parametric and non–parametric methods to fit the risk neutral density (for example, the methods proposed by Aït-Sahalia and Lo [1998], Jackwerth [2000], and Figlewski [2010]), there fails to be a clearly unique, objective and superior method for estimating the risk neutral density from discrete option prices. BP estimate the risk aversion coefficient implied in the index options by assuming a parametric form of the utility function and inferring the time–varying implied subjective return density function from the estimated time–varying risk neutral density. They solve for the risk aversion parameter value that improves the forecast of the future distribution of returns using the option fitted risk neutral density, and obtain values between 1 and 10. Utilizing option price data for the S&P500 and FTSE100 indexes, BP estimate the risk aversion coefficient at investment horizons between 1 and 6 weeks due to the underlying assumption of independence embedded in their empirical test. Their econometric method does not allow for the inference of option implied risk aversion at longer investment horizons of 3 to 12 months. Finally, a somewhat puzzling result from BP is the (almost monotonically) decreasing term structure of the estimated risk aversion coefficient: shorter investment horizons imply larger representative investor's risk aversion than longer investment horizons.

2.2 Modeling framework

The technical framework within this paper is motivated by the findings of Martin [2017]. Under the assumption of no–arbitrage and the (so called) negative correlation condition, Martin [2017] derives a lower bound on the expected return of S&P500 index. Utilizing an alternative mathematical derivation, I extend his result to consider higher moments of the market return subjective distribution. To begin, I introduce the following notation: time period $t < T$ where t is the present time and T is the option maturity date (investment horizon), the stochastic discount factor is given by M_T, the asset gross return, denoted R_T, $\mathbb{E}_t$ for the expectation operator under the subjective distribution conditional on time t information, and k, the mo-

ment order, an integer greater than or equal to 1. To derive the moment conditions, I assume that the conditional moments of the subjective return distribution exist. Assuming no–arbitrage, I start from the fundamental valuation equation:

$$1 = \mathbb{E}_t(M_T R_T), \tag{2.2}$$

and multiply both sides by $\mathbb{E}_t(R_T^k)$:

$$\mathbb{E}_t(R_T^k) = \mathbb{E}_t(R_T^k)\,\mathbb{E}_t(M_T R_T).$$

Using the definition of covariance between two random variables

$$\mathrm{cov}_t(M_T R_T, R_T^k) = \mathbb{E}_t(M_T R_T^{k+1}) - \mathbb{E}_t(R_T^k)\,\mathbb{E}_t(M_T R_T),$$

I rewrite the above equation as follows:

$$\mathbb{E}_t(R_T^k) = \mathbb{E}_t(M_T R_T^{k+1}) - \mathrm{cov}_t(M_T R_T, R_T^k).$$

The existence of an equivalent risk neutral representation (martingale measure) of the subjective moments allows for the substitution $\mathbb{E}_t(M_T R_T^{k+1}) = 1/R_{f,t}\,\mathbb{E}_t^*(R_T^{k+1})$, where $\mathbb{E}_t^*$ is the conditional expectation operator under the risk neutral distribution. Using this fact, I obtain a relationship between the kth moment of the subjective distribution and the difference between the $k+1$ moment of the risk neutral density and a covariance term that involves the stochastic discount factor:

$$\mathbb{E}_t(R_T^k) = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^{k+1}) - \mathrm{cov}_t(M_T R_T, R_T^k). \tag{2.3}$$

This identity is a generalized relation of the first moment condition (that is, where $k = 1$) obtained by Martin [2017]. Equation (2.3) establishes a mathematical relation between subjective and risk neutral moments: the covariance between discounted return and powers of the gross returns reflect investor's preferences. Martin

[2017] discusses the theoretical framework for which the covariance term is negative for $k = 1$: a negative covariance term allows him to establish a lower bound of the first subjective moment in terms of only the second risk neutral moment. Martin calls this condition the negative correlation condition (NCC), $\text{cov}_t(M_T R_T, R_T) \leq 0$.

A technical result on associated random variables by Esary et al. [1967] can be used to show that the NCC assumption for the first moment implies $\text{cov}_t(M_T R_T, R_T^k) < 0$ for all moments ($k > 1$). To see this, if I assume that $\text{cov}_t(-M_T R_T, R_T) \geq 0$, it follows that the two random variables $M_T R_T$ and R_T are associated. Esary et al. [1967] show that non–decreasing functions of associated random variables are also associated. Since R_T represents gross asset return and is thus a non–negative random variable, and any power of R_T (such as R_T^2) is a non–decreasing function of R_T, the NCC for the first moment implies that the covariance of $M_T R_T$ with higher powers of R_T is negative.

In this study, I provide empirical evidence that the S&P500 index options implied risk aversion coefficient with power utility is not statistically different and sufficiently close to one. This result allows me to assume that the NCC is empirically equal to zero and to justify the assumption of logarithmic utility, i.e. that the stochastic discount factor is given by $M_T = 1/R_T$, where R_T is chosen to be the S&P500 index as a proxy for the wealth portfolio.

2.2.1 Central moments with logarithmic utility

To obtain the second and third moments, I let R_T represent the market return (and its empirical proxy the S&P500 index); under the specification of logarithmic utility $M_T = 1/R_T$ and the NCC is satisfied with identity, $\text{cov}_t(M_T R_T . R_T^k) = \text{cov}_t(\frac{1}{R_T} R_T . R_T^k) = 0$ for all powers $k \geq 1$. This gives:

$$\mathbb{E}_t(R_T^2) = \frac{1}{R_{f,t}} \mathbb{E}_t^*(R_T^3),$$
$$\mathbb{E}_t(R_T^3) = \frac{1}{R_{f,t}} \mathbb{E}_t^*(R_T^4).$$

I obtain the central moments of the logarithmic investor's subjective distribution, by invoking the standard definitions of variance and skewness:

$$\mathrm{var}_t(R_T) = \mathbb{E}_t(R_T^2) - \mathbb{E}_t(R_T)^2 = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^3) - \left(\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^2)\right)^2 \tag{2.4}$$

$$\mathrm{skew}_t(R_T) = \frac{\mathbb{E}_t[(R_T - \mathbb{E}_t\,R_T)^3]}{(\mathrm{var}_t\,R_T)^{3/2}} = \frac{\mathbb{E}_t(R_T^3) - 3\,\mathbb{E}_t(R_T)\,\mathbb{E}_t(R_T^2) + 2\,\mathbb{E}_t(R_T)^3}{(\mathrm{var}_t\,R_T)^{3/2}} \tag{2.5}$$

$$= \frac{\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^4) - 3\frac{1}{R_{f,t}^2}\,\mathbb{E}_t^*(R_T^2)\,\mathbb{E}_t^*(R_T^3) + 2\frac{1}{R_{f,t}^3}\,\mathbb{E}_t^*(R_T^2)^3}{(\mathrm{var}_t\,R_T)^{3/2}}. \tag{2.6}$$

As Martin [2017] points out, the logarithmic utility function is not the only preference specification that satisfies the NCC at zero. As the empirical results I provide with this paper show, one of the advantages of the logarithmic utility specification is that it is empirically testable using option prices and return data.

Before moving on to the empirical estimation of the option implied risk aversion coefficient I show that the first two central moments derived above become bounds when the risk aversion is allowed to be greater than one as it is for the logarithmic utility specification.

2.2.2 Bounds on moments

I derive subjective moments under power utility specification and show that moments derived under the special case of logarithmic utility serve as bounds for the option implied subjective moments. As in the previous sections, I consider R_T the return on the market (or its empirical proxy the S&P500 index). Assuming power utility, the stochastic discount factor is given by $M_T = R_T^{-\gamma}$. Substitute this specification in Equation (2.3) to obtain:

$$\mathbb{E}_t\,R_T^k = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*\,R_T^{k+1} - \mathrm{cov}_t(R_T^{1-\gamma}, R_T^k). \tag{2.7}$$

Imposing the NCC assumption, I can obtain the bounds on the subjective moments. For the first central moment, as Martin [2017] shows, the NCC is sufficient to find a lower bound. The bounds on the second and higher central moments unfortu-

nately require additional assumptions. As such, in the following analysis I assume that the risk aversion coefficient $\gamma > 1$ and linearize both $R_T^{1-\gamma}$ and functions of the gross return R_T at gross return equal to 1 in the covariance term of Equation (2.7). Linearization allows me to ignore terms of the Taylor series expansion of the expressions inside the covariance term that are of order 2 or above.

Using the fact that:

$$R_T^{1-\gamma} = \exp(\log R_T^{1-\gamma}) = \exp((1-\gamma)\log R_T),$$

the first order approximation for the term $R_T^{1-\gamma}$ is given by:

$$R_T^{1-\gamma} \approx 1 + (1-\gamma)\log R_T.$$

Also, $R_T = \exp(\log R_T) \approx 1 + \log R_T$ and $R_T^2 = \exp(\log R_T^2) \approx 1 + 2\log R_T$. This approximation is known to hold only at short horizons (of less than one year, Cochrane [2005]). Consider the first moment of market return:

$$\begin{aligned}
\mathbb{E}_t R_T &\approx \frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} - \operatorname{cov}_t(1 + (1-\gamma)\log R_T, 1 + \log R_T) \\
&= \frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} - (1-\gamma)\operatorname{var}_t \log R_T.
\end{aligned}$$

The linearized power utility stochastic discount factor establishes that an investor with risk aversion coefficient greater than one expects higher return than implied by logarithmic utility. For the second moment I obtain the following relation:

$$\begin{aligned}
\mathbb{E}_t R_T^2 &\approx \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - \operatorname{cov}_t(1 + (1-\gamma)\log R_T, 1 + 2\log R_T) \\
&= \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - 2(1-\gamma)\operatorname{var}_t \log R_T.
\end{aligned}$$

With the assumption of $\gamma > 1$, the second term in both the first and second moment equations are negative. In addition, I can derive the bound on the second central moment, that is the variance implied by power utility with risk aversion greater than

1:

$$\mathbb{E}_t(R_T^2) - \mathbb{E}_t(R_T)^2 \approx \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - 2(1-\gamma)\operatorname{var}_t \log R_T - \left[\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} - (1-\gamma)\operatorname{var}_t \log R_T\right]^2.$$

Expand the square in the bracket and collect terms:

$$\begin{aligned}\mathbb{E}_t(R_T^2) - \mathbb{E}_t(R_T)^2 \approx{}& \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - 2(1-\gamma)\operatorname{var}_t \log R_T - \left(\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}}\right)^2 \\ &+ 2(1-\gamma)\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}}\operatorname{var}_t \log R_T - (1-\gamma)^2(\operatorname{var}_t \log R_T)^2 \\ &= \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - \left(\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}}\right)^2 \\ &+ (1-\gamma)\operatorname{var}_t \log R_T\left(-2 + 2\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} - (1-\gamma)\operatorname{var}_t \log R_T\right).\end{aligned}$$

The term $\left(-2 + 2\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} - (1-\gamma)\operatorname{var}_t \log R_T\right)$ is always greater than zero for $\gamma > 1$, because $\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}} \geq 1$ and $-(1-\gamma)\operatorname{var}_t \log R_T > 0$. This term multiplies $\operatorname{var}_t \log R_T > 0$ and $1-\gamma < 0$, making the overall third term in the above equation negative. Thus, I can establish the upper bound of the second central moment from the first two terms on the right hand side of the above equation, which gives:

$$0 \leq \operatorname{var}_t R_T = \mathbb{E}_t(R_T^2) - \mathbb{E}_t(R_T)^2 \leq \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - \left(\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}}\right)^2.$$

This upper bound states that an investor with risk aversion coefficient greater than 1 expects lower variance than the one implied by the logarithmic utility specification. The variance inequality relation is of the opposite sign than the one for the expected return. The linearized first and second central moment inequalities determine a lower bound on the expected market Sharpe Ratio. That is, since:

$$\begin{aligned}\mathbb{E}_t R_T - R_{f,t} &\geq \frac{\operatorname{var}_t^* R_T}{R_{f,t}}, \quad \text{and} \\ \operatorname{var}_t R_T &\leq \frac{\mathbb{E}_t^* R_T^3}{R_{f,t}} - \left(\frac{\mathbb{E}_t^* R_T^2}{R_{f,t}}\right)^2,\end{aligned}$$

a lower bound on the ex–ante Sharpe Ratio of the market follows:

$$\frac{\mathbb{E}_t\, R_T - R_{f,t}}{\sqrt{\mathrm{var}_t\, R_T}} \geq \frac{\mathrm{var}_t^*\, R_T}{R_{f,t}\sqrt{\frac{\mathbb{E}_t^*\, R_T^3}{R_{f,t}} - \left(\frac{\mathbb{E}_t^*\, R_T^2}{R_{f,t}}\right)^2}}.$$

The lower bound is binding only for the logarithmic utility investor.

2.2.3 Relation to the conditional Capital Asset Pricing Model

A well known fact is that the assumption of logarithmic utility implies the Capital Asset Pricing Model (CAPM) (Cochrane [2005], page 160). The conditional CAPM can also be derived using the linear stochastic discount factor $M_T = a_t + b_t R_T$, with R_T the return of the market portfolio. Replacing this stochastic discount factor in the fundamental valuation equation (2.2) and applying it to value R_T and the risk free rate $R_{f,t}$ allows me to solve for the two conditional parameters:

$$\begin{aligned} a_t &= \frac{1}{R_{f,t}} - b_t\, \mathbb{E}_t\, R_T, \\ b_t &= \frac{R_{f,t} - \mathbb{E}_t\, R_T}{R_{f,t}\, \mathrm{var}_t\, R_T}. \end{aligned}$$

These two parameters can be explicitly computed using the option implied log investor expected return and variance:

$$\begin{aligned} a_t &= \frac{1}{R_{f,t}} - b_t \frac{\mathbb{E}_t^*\, R_T^2}{R_{f,t}}, \\ b_t &= \frac{R_{f,t}^2 - \mathbb{E}_t^*\, R_T^2}{R_{f,t}\, \mathbb{E}_t^*\, R_T^3 - \left(\mathbb{E}_t^*\, R_T^2\right)^2}. \end{aligned}$$

The resulting explicit stochastic discount factor provides a testable implication of the ex–ante conditional CAPM.

2.3 Option implied risk aversion coefficient

Using the following version of the fundamental valuation equation (Cochrane [2005]):

$$\mathbb{E}_t(R_T) = R_{f,t} - R_{f,t}\,\mathrm{cov}_t(M_T, R_T) \tag{2.8}$$

along with Equation (2.3) I infer and test the risk aversion coefficient with the power utility specification. Both Equations (2.3) and (2.8) hold under the assumption of no–arbitrage (and the existence of the distributional moments). By equating the right hand side of Equation (2.8) and Equation (2.3) with $k = 1$, I obtain a moment condition for the stochastic discount factor M_T to price both the underlying security and its traded option:

$$0 = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^2) - \mathrm{cov}_t(M_T R_T, R_T) + R_{f,t}\,\mathrm{cov}_t(M_T, R_T) - R_{f,t}. \tag{2.9}$$

Equation (2.9) involves the security gross return R_T, the stochastic discount factor M_T, the risk free rate gross return $R_{f,t}$ from time t to T, and the second moment of the security gross return under the risk neutral distribution. The risk neutral moment can be computed from options written on the security whose return is R_T. Although I have specified R_T as the S&P500 index return, Equation (2.9) applies to any security i for which options are traded. It states that the discounted risk neutral variance must equal the difference of two covariance terms:

$$\frac{\mathrm{var}_t^*(R_{i,T})}{R_{f,t}} = \mathrm{cov}_t(M_T R_{i,T}, R_{i,T}) - R_{f,t}\,\mathrm{cov}_t(M_T, R_{i,T}).$$

As this relation holds for any security, it enables the test of a specific form of the stochastic discount factor in terms of realized security returns (other than the S&P500 index as a market portfolio proxy) from t to T and the conditional (ex–ante) risk netral variance calculated from option prices.

I focus on testing the moment condition in Equation 2.9 using realized cumulative returns in the covariance terms along with the option implied risk neutral

second moment for the S&P500 and the Russell 2000 indexes (both European index options). I follow the parametric specification of BP and implement the power utility functional form to estimate the risk aversion coefficient, effectively assuming preferences to be stationary. As mentioned above, Equation (2.9) involves the conditional second moment of the risk neutral density, itself a non–stationary density as observed, for example, by Figlewski [2010]. One advantage of Equation (2.9) is that it holds for any time t, for different return horizons T, and risk neutral moment maturity (within the limits of the traded maturities of option contracts). I test the utility specification using time series data only, thus enabling the estimation of the risk aversion coefficient without resorting to the fitting of either the risk neutral or subjective densities.

Since the identity above is in a moment condition form, the natural econometric methodology to estimate the risk aversion parameter and test the validity of the model for different investment horizons is the GMM framework of Hansen [1982]. Based on the evidence in Martin [2017], who argues that the second risk neutral moment is a strong predictor of subsequent underlying security returns; I use its lagged values as instruments in the GMM specification to enable a test of the model in over–identified form. I also consider two alternative over–identified GMM specifications: estimate the moment conditions with multiple investment horizons and with two different option indexes.

To apply the GMM framework I take the unconditional expectation of Equation (2.9):

$$\begin{aligned} 0 &= \mathbb{E}\left[\frac{1}{R_{f,t}}\,\mathbb{E}_t^*\,R_T^2 - \mathrm{cov}_t(M_T R_T, R_T) + R_{f,t}\,\mathrm{cov}_t(M_T, R_T) - R_{f,t}\right] \\ &= \mathbb{E}\left[\frac{1}{R_{f,t}}\,\mathbb{E}_t^*\,R_T^2\right] - \mathrm{cov}(M_T R_T, R_T) + R_{f,t}\,\mathrm{cov}(M_T, R_T) - R_{f,t}, \end{aligned}$$

where I used the following facts:

$$\mathbb{E}(\mathrm{cov}_t(M_T, R_T)) = \mathrm{cov}(M_T, R_T) - \mathrm{cov}(\mathbb{E}_t\,M_T, \mathbb{E}_t\,R_T) = \mathrm{cov}(M_T, R_T)$$

because, $\text{cov}(\mathbb{E}_t\, M_T, \mathbb{E}_t\, R_T) = \text{cov}(1/R_{f,t}, \mathbb{E}_t\, R_T) = 0$ (the risk free rate is a deterministic variable), and:

$$\mathbb{E}(\text{cov}_t(M_T R_T, R_T)) = \text{cov}(M_T R_T, R_T) - \text{cov}(\mathbb{E}_t(M_T R_T), \mathbb{E}_t\, R_T) = \text{cov}(M_T R_T, R_T)$$

because $\text{cov}(\mathbb{E}_t(M_T R_T), \mathbb{E}\, R_T) = 0$ since $\mathbb{E}_t(M_T R_T) = 1$. I now rewrite the covariances in terms of expectations:

$$0 = \mathbb{E}\left[\frac{1}{R_{f,t}}\, \mathbb{E}_t^*\, R_T^2\right] - \mathbb{E}(M_T R_T^2) + \mathbb{E}(M_T R_T)\, \mathbb{E}\, R_T + R_{f,t}\left[\mathbb{E}(M_T R_T) - \mathbb{E}\, M_T\, \mathbb{E}\, R_T\right] - R_{f,t}.$$

If I now assume the power utility parametric form for the stochastic discount factor, $M_T = R_T^{-\gamma}$ with γ as the risk aversion coefficient, and R_T now representing the market gross return (which will be empirically taken as the S&P500 index return):

$$0 = \mathbb{E}\left[\frac{1}{R_{f,t}}\, \mathbb{E}_t^*\, R_T^2\right] - \mathbb{E}\, R_T^{2-\gamma} + \mathbb{E}\, R_T^{1-\gamma}\, \mathbb{E}\, R_T + R_{f,t}\left[\mathbb{E}\, R_T^{1-\gamma} - \mathbb{E}\, R_T^{-\gamma}\, \mathbb{E}\, R_T\right] - R_{f,t}.$$

Let $\mathbb{E}\, R_T = \frac{1}{N}\sum_{T=1}^{N} R_T = \bar{R}_T$, and the sample equivalent expression of the moment condition above is given by:

$$0 = \frac{1}{N}\sum_{T=1}^{N}\left(\frac{1}{R_{f,t}}\, \mathbb{E}_t^*\, R_T^2 - R_T^{2-\gamma} + R_{f,t} R_T^{1-\gamma} + R_T^{1-\gamma}\bar{R}_T - R_{f,t}\bar{R}_T R_T^{-\gamma} - R_{f,t}\right),$$

where N is the size of the sample. The risk neutral expectation term can be estimated on a daily basis by using daily closing options prices. The investment horizon $T - t$ can be from 1 to 12 months. R_T is the gross return from time t to T. Since the risk neutral expected moment is measured at daily frequency, the moment condition involves overlapping gross returns. GMM estimation is executed in two steps with appropriate adjustment to account for autocorrelation and heteroskedasticity of the residuals (Cochrane [2005] is an excellent reference for such adjustments in the application of GMM).

2.4 Estimation of option implied risk neutral moments

The calculation of the powers of gross stock returns follows well know spanning results used in Carr and Madan [2001] and Bakshi et al. [2003]. To begin, I provide details of the specific computations (Bakshi et al. [2003] derive formulas for simple return or logarithmic return contracts). Let S_T be the spot price at time T, so that $R_T = S_T/S_t$. The following identities hold for any given future value S_T:

$$
\begin{aligned}
S_T^2 &= 2\int_0^\infty \max(0, S_T - K)dK, \\
S_T^3 &= 6\int_0^\infty K\max(0, S_T - K)dK, \\
S_T^4 &= 12\int_0^\infty K^2\max(0, S_T - K)dK.
\end{aligned}
$$

Since the integrands are payoffs, applying the expectation with any discount factor results in the equivalent call price Cochrane [2005]. If we apply the risk neutral expectation to the integrands above we get:

$$
\begin{aligned}
\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(S_T^2) &= 2\int_0^\infty \frac{1}{R_{f,t}}\,\mathbb{E}_t^*\left(\max(0, S_T - K)\right)dK = 2\int_0^\infty \mathrm{call}_{t,T}(K)dK \\
\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(S_T^3) &= 6\int_0^\infty \frac{K}{R_{f,t}}\,\mathbb{E}_t^*\left(\max(0, S_T - K)\right)dK = 6\int_0^\infty K\,\mathrm{call}_{t,T}(K)dK \\
\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(S_T^4) &= 12\int_0^\infty \frac{K^2}{R_{f,t}}\,\mathbb{E}_t^*\left(\max(0, S_T - K)\right)dK = 12\int_0^\infty K^2\,\mathrm{call}_{t,T}(K)dK
\end{aligned}
$$

To evaluate the integrals prices of illiquid in–the–money call options are required. A standard way around this problem is to use put–call parity and separate the interval of integration over call and (equivalent) put options:

$$
\begin{aligned}
\int_0^\infty \mathrm{call}_{t,T}(K) &= \int_0^{F_{t,T}} \mathrm{put}_{t,T}(K) + \frac{1}{R_{f,t}}(F_{t,T} - K)dK + \int_{F_{t,T}}^\infty \mathrm{call}_{t,T}(K)dK \\
&= \int_0^{F_{t,T}} K\,\mathrm{put}_{t,T}(K)dK + \frac{F_{t,T}^2}{2R_{f,t}} + \int_{F_{t,T}}^\infty K\,\mathrm{call}_{t,T}(K)dK,
\end{aligned}
$$

$$\int_0^\infty K\,\mathrm{call}_{t,T}(K) = \int_0^{F_{t,T}} K\,\mathrm{put}_{t,T}(K) + \frac{1}{R_{f,t}} K(F_{t,T} - K)dK + \int_{F_{t,T}}^\infty K\,\mathrm{call}_{t,T}(K)dK$$
$$= \int_0^{F_{t,T}} K\,\mathrm{put}_{t,T}(K)dK + \frac{F_{t,T}^3}{6R_{f,t}} + \int_{F_{t,T}}^\infty K\,\mathrm{call}_{t,T}(K)dK,$$

and:

$$\int_0^\infty K^2\,\mathrm{call}_{t,T}(K) = \int_0^{F_{t,T}} K^2\,\mathrm{put}_{t,T}(K) + \frac{1}{R_{f,t}} K^2(F_{t,T} - K)dK + \int_{F_{t,T}}^\infty K^2\,\mathrm{call}_{t,T}(K)dK$$
$$= \int_0^{F_{t,T}} K^2\,\mathrm{put}_{t,T}(K)dK + \frac{F_{t,T}^4}{12R_{f,t}} + \int_{F_{t,T}}^\infty K^2\,\mathrm{call}_{t,T}(K)dK,$$

where $F_{t,T} = R_{f,t}S_t$. Finally, plugging in the equivalent integral expressions in the risk neutral expectations above and dividing both sides by the corresponding powers of the spot price S_t we obtain:

$$\frac{1}{R_{f,t}}\,\mathbb{E}_t^*\left(\frac{S_T^2}{S_t^2}\right) = R_{f,t} + \frac{2}{S_t^2}\left[\int_0^{F_{t,T}} \mathrm{put}_{t,T}(K)dK + \int_{F_{t,T}}^\infty \mathrm{call}_{t,T}(K)dK\right] = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^2)$$
$$\frac{1}{R_{f,t}}\,\mathbb{E}_t^*\left(\frac{S_T^3}{S_t^3}\right) = R_{f,t}^2 + \frac{6}{S_t^3}\left[\int_0^{F_{t,T}} K\,\mathrm{put}_{t,T}(K)dK + \int_{F_{t,T}}^\infty K\,\mathrm{call}_{t,T}(K)dK\right] = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^3)$$
$$\frac{1}{R_{f,t}}\,\mathbb{E}_t^*\left(\frac{S_T^4}{S_t^4}\right) = R_{f,t}^3 + \frac{12}{S_t^4}\left[\int_0^{F_{t,T}} K^2\,\mathrm{put}_{t,T}(K)dK + \int_{F_{t,T}}^\infty K^2\,\mathrm{call}_{t,T}(K)dK\right] = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^4).$$

The integrals use the spot price S_t, the risk free rate $R_{f,t}$, and prices of put and calls. Strike prices K of puts and calls are available only over a finite interval, thereby forcing the truncation of the integral evaluation between the minimum and the maximum available strike prices.

2.5 Data

Daily data on S&P500 and Russell 2000 European index options from January 4, 1996 to July 31, 2014 is obtained from OptionMetrics. Risk free rate term structure, the S&P500 and Russell 2000 spot prices and dividend yields are also collected from OptionMetrics. Using index option data I estimate the first three moments of the subjective distribution by evaluating the integrals using the trapezoidal rule

suggested by Dennis and Mayhew [2002]. I also apply the following basic filters to the data following standard practice in the option literature (see for example, Dennis and Mayhew [2002], Bakshi et al. [2003], or Martin [2017]):

- Any option with a bid price less than 0.50 USD is dropped,
- Only maturities between and including 7 days and 450 days are considered (7 days is standard, whereas 450 is specific to this study – I eventually only estimate moments up to 10 months maturity due to lack of sufficiently traded option contracts),
- Maturities with at least 2 calls and 4 puts,
- Traded volume greater than 0, and,
- Any option with implied volatility greater than 1 or less than 0 is dropped.

I use the implied volatilities provided by OptionMetrics. The price used for the estimation of implied volatilities is the average of the bid and ask quotes. For options near the money, consistent with Figlewski [2010], I keep only a small overlap of option prices for the fit of the implied volatility curve at each maturity, and fit a third order spline with a knot at closing spot price S_t. It is interesting to note, from a robustness point of view, that using a second or fourth order (as Figlewski [2010] suggests) spline does not change the resulting moment estimations in a significant way. As a further robustness check, I also estimate the risk neutral moments using average bid ask quotes for all quoted options including the ones with zero traded volume. Results are unchanged.

I report the daily trading volume for the S&P500 and the Russell 2000 option indexes in Figures 2.1 and 2.2, respectively. Interestingly, two features are evident from the plots: first, the S&P500 option volume is much larger than the Russell 2000 option volume; second, both trading volumes experienced a significant increase from the end of the year 2005. Relevant volume statistics for the two index options are shown in Table 2.1. For the Russell 2000 option index there is a single day with 0 traded options, December 3, 2001. The trading volume of S&P500 options is much higher than that of the Russell 2000 options, the mean is an order of magnitude

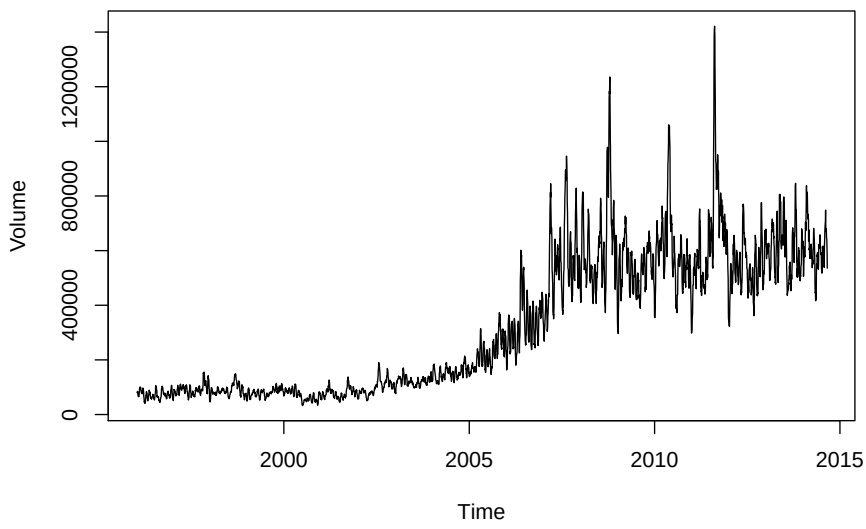

Figure 2.1: S&P500 trading volume. A 10 day moving average filter is applied to the time series.

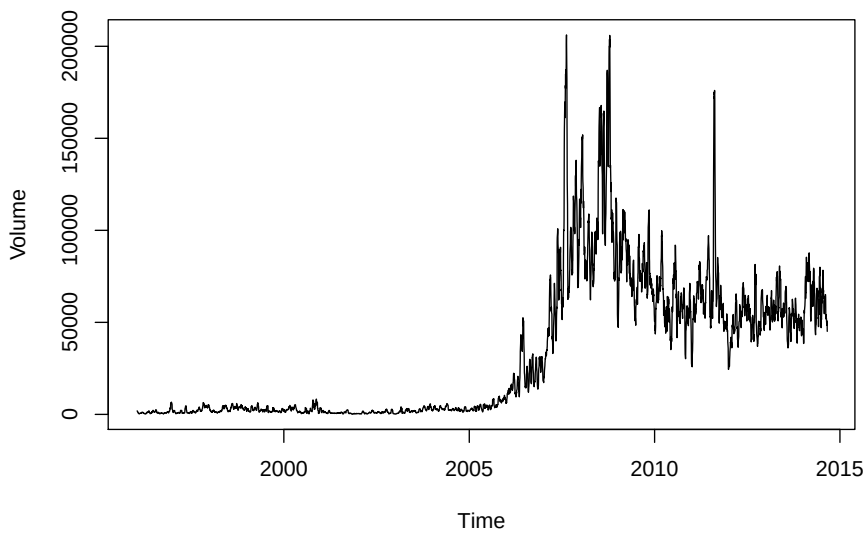

Figure 2.2: Russell 2000 trading volume. A 10 day moving average filter is applied to the time series.

larger. Over the entire sample from 1996 to 2014, the shortest maturity available on any day from the S&P500 is at least 10 months, whereas the shortest maturity contract for the Russell 2000 options is three months. When testing with both option time series I limit the inference to the three months investment horizon to

avoid the extrapolation of the estimated risk neutral moments.

Table 2.1: Option Volume statistics for S&P500 and Russell 2000 indexes. The sample data is obtained from OptionMetrics for the period 1996–2014. The statistics reported is for the data filtered according to the rules included in the data section 2.5.

	Volume (1000s)					
	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
S&P500	4.71	89.8	194	318	511	2080
Russell 2000	0	1.2	6.7	31.9	54	401

2.6 Results

2.6.1 Estimation of risk aversion coefficient

I empirically estimate the risk aversion coefficient implied in S&P500 options using the moment condition detailed in Equation (2.9). Estimating the risk aversion coefficient with the power utility function allows me to test the hypothesis of logarithmic utility (see, Brown and Gibbons [1985] for a GMM test of logarithmic utility using return data only) [2]. I test the moment condition under three different specifications for over–identification. For all three specifications I use a two stage GMM with correction for heteroskedasticity and autocorrelation of the residuals.

First, I consider the moment condition specification reported in Equation (2.9) with a lagged instrument, the predictive risk neutral variance $\frac{1}{R_{f,t}} \operatorname{var}^* R_T$,

$$0 = \frac{1}{N} \sum_{T=1}^{N} \left(\frac{\operatorname{var}_t^* R_T}{R_{f,t}} - R_T^{2-\gamma} + R_{f,t} R_T^{1-\gamma} + R_T^{1-\gamma} \bar{R}_T - R_{f,t} \bar{R}_T R_T^{-\gamma} \right),$$

$$0 = \frac{1}{N} \sum_{T=1}^{N} \left(\frac{\operatorname{var}_t^* R_T}{R_{f,t}} - R_T^{2-\gamma} + R_{f,t} R_T^{1-\gamma} + R_T^{1-\gamma} \bar{R}_T - R_{f,t} \bar{R}_T R_T^{-\gamma} \right) \frac{\operatorname{var}_{t-1}^* R_{T-1}}{R_{f,t-1}}.$$

In Table 2.2, in line with the first specification, I report the estimation of the risk

[2] Although Brown and Gibbons [1985] fail to reject the log utility model, many empirical studies of the consumption based asset pricing model of Lucas [1978] with power utility preferences have rejected it and yielded unacceptably high estimates of the risk aversion coefficient (see Chapter 21 of Cochrane [2005] for an excellent summary)

aversion coefficient for four exact and over–identified GMM tests using 1, 3, 6, and 9 months return horizon with a single instrument as the lagged expected return for the entire sample (1996–2014). The estimates of the exact identified moment condition are reported in the first two rows of the table, with standard errors below the estimates in parenthesis. Consistent with the results of BP, the estimate value of the risk aversion coefficient γ is monotonically decreasing as the investment horizon increases.

Table 2.2: GMM estimation of risk aversion coefficient γ for power utility specification. The results for horizons 1, 3, 6, and 9 months are obtained from the exact and over–identified system with 4 lags of the forecasting variable, $\frac{\text{var}^* R_T}{R_{f,t}}$, as instruments. I report statistical significance for the standard hypothesis test $H_0 : \gamma = 0$. The t-statistics associated with the hypothesis test $H_0 : \gamma = 1$ can be calculated using the ratio of the parameter estimate and the robust standard error in parenthesis.

	Return horizon			
	1 month	3 months	6 months	9 months
γ (exact)	5.714**	4.091*	2.393	1.394
	(2.273)	(2.106)	(1.815)	(2.120)
γ (over–identified)	6.310***	4.915***	3.388**	2.940***
	(2.034)	(1.869)	(1.486)	(1.132)
J stat	1.294	5.856	6.532	3.849
p value	0.524	0.119	0.088	0.278
Observations	4,422	4,422	4,422	4,422
Note:				*p<0.1; **p<0.05; ***p<0.01

Moreover, from the p–values of the J statistics I conclude that the GMM overidentified specification is not rejected for any investment horizon at the 5% significance level. Based on the exact moment estimates reported in Table 2.2, I am unable to reject the null hypothesis that risk aversion is equal to 1. In all tables I report the statistical significance associated with the standard hypothesis test $H_0 : \gamma = 0$. The t-statistics for the hypothesis test $H_0 : \gamma = 1$ can be calculated using the ratio of the parameter estimate and the robust standard error in parenthesis. The estimate and standard error for the 6 and 9 month horizons of the over–identified specification in

Table 2.2 fails to reject the null hypothesis of risk aversion equal to 1. Of course, were I to consider specific subsample, the estimation of the risk aversion coefficient would change. I choose the entire sample to increase the power of the test. An important change during the time period I consider is the significant increase in traded volume for the S&P500 index options starting around the end of 2005, the second half of the sample. Overall these findings allow me to conclude that I am unable to reject the logarithmic utility specification.

Next I test the same moment condition from Equation (2.9), but using the combination of two overlapping investment horizon time series $T_1 - t$ and $T_2 - t$:

$$0 = \frac{1}{N}\sum_{T=1}^{N}\left(\frac{\mathrm{var}_t^* R_{T_1}}{R_{f,t}} - R_{T_1}^{2-\gamma} + R_{f,t}R_{T_1}^{1-\gamma} + R_{T_1}^{1-\gamma}\bar{R}_{T_1} - R_{f,t}\bar{R}_{T_1}R_{T_1}^{-\gamma}\right),$$

and

$$0 = \frac{1}{N}\sum_{T_2}^{N}\left(\frac{\mathrm{var}_t^* R_{T_2}}{R_{f,t}} - R_{T_2}^{2-\gamma} + R_{f,t}R_{T_2}^{1-\gamma} + R_{T_2}^{1-\gamma}\bar{R}_{T_2} - R_{f,t}\bar{R}_{T_2}R_{T_2}^{-\gamma}\right).$$

I estimate the coefficient of risk aversion based on the moment specification using four investment horizon specifications: 1 and 2, 2 and 3, 3 and 6, and 9 and 10 months. These combined investment horizons are sufficiently close to each other for an over–identification model specification in the GMM test. The estimation results are reported in Table 2.3. The 2/3 and 6/9 months horizon specifications are not rejected by the GMM J statistics, whilst none of the estimate of risk aversion is statistically different from 1. The values for γ estimated from this second specification are not too different from the estimates based on the first specification in Table 2.2.

Lastly, I test the moment condition from Equation (2.9) using option data for two indexes, namely the S&P500 and the Russell 2000 index:

$$0 = \frac{1}{N}\sum_{T=1}^{N}\left(\frac{\mathrm{var}_t^* R_{SP,T}}{R_{f,t}} - R_{SP,T}^{2-\gamma} + R_{f,t}R_{SP,T}^{1-\gamma} + R_{SP,T}^{1-\gamma}\bar{R}_{SP,T} - R_{f,t}\bar{R}_{SP,T}R_{SP,T}^{-\gamma}\right),$$

$$0 = \frac{1}{N}\sum_{T=1}^{N}\left(\frac{\mathrm{var}_t^* R_{RU,T}}{R_{f,t}} + R_{SP,T}^{-\gamma}[R_{RU,T}\bar{R}_{RU,T} - R_{RU,T}^2] + R_{f,t}R_{SP,T}^{-\gamma}[R_{RU,T} - \bar{R}_{RU,T}]\right),$$

Table 2.3: GMM estimation of risk aversion coefficient γ for power utility specification. The results for horizons 1/2, 2/3, 3/6, and 6/9 months are obtained from the over–identified system that includes the moment condition in Equation (2.9) for two time series with different (overlapping) investment horizons. Standard errors are corrected for autocorrelation and heteroskedasticity. I report statistical significance for the standard hypothesis test $H_0 : \gamma = 0$. The t-statistics associated with the hypothesis test $H_0 : \gamma = 1$ can be calculated using the ratio of the parameter estimate and the robust standard error in parenthesis.

	Return horizon			
	1/2 months	2/3 months	6/9 months	9/10 months
γ	4.323**	4.360**	2.081	1.630
	(1.997)	(2.029)	(1.799)	(1.899)
J stat	5.178	0.301	0.607	4.620
p value	0.023	0.584	0.436	0.032
Observations	4,425	4,425	4,425	4,425
Note:				*p<0.1; **p<0.05; ***p<0.01

where the subscripts SP and RU are for the S&P500 index and the Russell 2000 index, respectively. I combine the Russell 2000 index option and spot price data to test the moment conditions. In this test I restricted the investment horizon to 3 months to keep option data for the entire sample: the trading volume for the Russell 2000 index options is highest within the 3 months horizon. The GMM estimation results are reported in Table 2.4. Consistent with the previous two test specifications (Tables 2.2 and 2.3), the risk aversion estimates decrease as investment horizon increases. However, the interesting difference between the results in Table 2.4 and the previous two tables is the significantly smaller values for the J statistic. In addition, the estimated risk aversion coefficient is lower for all three horizons and more precise (smaller standard errors). Logarithmic utility is not rejected at any of the 3 horizons.

Consistent with the findings of BP, the conclusion from the three test specifications above is that option implied estimated risk aversion coefficient decreases with investment horizon. The downward sloping risk aversion term structure is a somewhat puzzling result. I argue that this finding is likely due to the relatively

Table 2.4: GMM estimation of risk aversion coefficient γ for power utility specification. The results for horizons 1, 2, and 3 months are obtained from the over–identified system that includes moment conditions for both S&P500 and Russell 2000 index returns and option implied moments. I report statistical significance for the standard hypothesis test $H_0 : \gamma = 0$. The t-statistics associated with the hypothesis test $H_0 : \gamma = 1$ can be calculated using the ratio of the parameter estimate and the robust standard error in parenthesis.

	Return horizon		
	1 month	2 months	3 months
γ	2.241^{**}	1.633^{**}	1.300^{**}
	(0.913)	(0.687)	(0.555)
J stat	0.016	0.012	0.006
p value	0.900	0.914	0.937
Observations	4,424	4,424	4,424
Note:	*p<0.1; **p<0.05; ***p<0.01		

higher implied market risk premium in the short maturity versus the longer maturity option contracts. Further analysis on this issue is beyond the scope of this paper.

An important additional result from the empirical evidence reported in Table 2.4 is that the risk aversion coefficient is not statistically different from 1 at the 2 and 3 month horizons, and definitely not equal to 0 (the hypothesis $H_0 : \gamma = 0$ is rejected for both horizons). In the following sections, I use this important empirical fact to justify the assumption of logarithmic utility for the representative agent. Finally, adding more option traded securities to the tests (such as the Russell 2000 index) improves the precision of the risk aversion estimate by reducing the adjusted standard errors and the overall J statistics.

2.6.2 Logarithmic utility variance and Sharpe Ratio

I report graphical evidence of the differences between risk neutral and subjective variance implied by S&P500 index option data in Figures 2.3, 2.4, 2.5 and 2.6. Specifically, these figures show the two time series for the annualized risk neutral ($\text{var}_t^* \log R_T$, the risk neutral variance of log return, equivalent to square of VIX) and

logarithmic utility variance estimated from Equation (2.4) at the one, three, six and nine months return horizons, respectively. The risk neutral variance of log return, $\text{var}_t^* \log R_T$, is estimated using the integral expressions in Bakshi et al. [2003]. The square root of the risk neutral variance of log return corresponds to the VIX index (where VIX is calculated at the one month horizon only). I compare the log utility variance with VIX square because VIX is considered the preference free (model free) implied volatility benchmark.

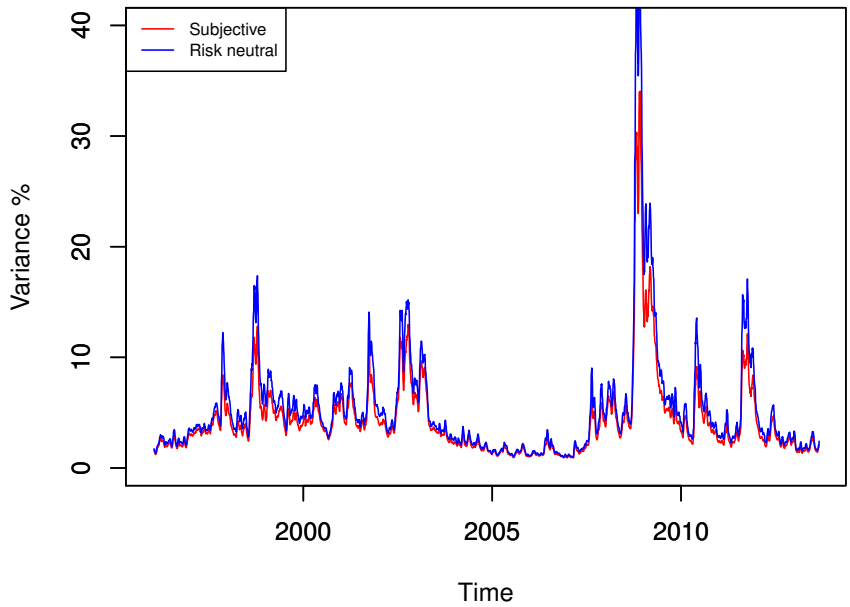

Figure 2.3: Annualized subjective and risk neutral variance, one month horizon. A 10 day moving average filter is applied to all time series.

There are two interesting points to emphasize from the graphical evidence. First, at all horizons, the logarithmic utility subjective variance is always lower than the risk neutral counterpart. This implies the presence of a variance premium between the VIX and the log utility variance. This finding may be interpreted as showing that an investor with logarithmic utility expects consistently lower future variance than implied by risk neutral pricing. If such an investor purchases options (particurarly out of the money put options on the S&P500 index), he will pay a premium because the price of such options is based on a higher expected future variance.

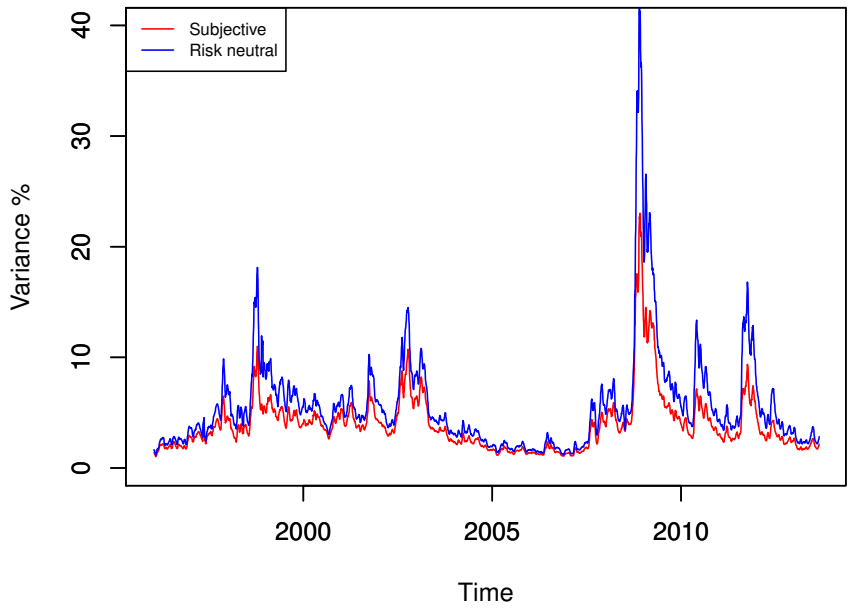

Figure 2.4: Annualized subjective and risk neutral variance, three month horizon. A 10 day moving average filter is applied to all time series.

It is also clear from these figures that the premium varies significantly over time, with the greatest deviations occurring when the change in variance is highest. The premium is lowest during "calm" market conditions such as the years that preceded the global financial crisis, namely 2005 and 2006. A second point is that the transient dynamics (the rate of change of the variance) of the two time series are also remarkably different; the risk neutral variance is more sensitive to left tail events (large negative returns) than the subjective variance, even at longer horizons. For reference, Figure 2.7 highlights the difference between the implied risk neutral and logarithmic utility subjective variance for three investment horizons. The mentioned "variance premium" increases with horizon (there is a term structure) and its mean reversion is more sluggish large upward shocks (large negative market returns).

Building on this graphical evidence, I test the forecasting power of log utility variance. Concretely, I implement the following predictive regressions with and

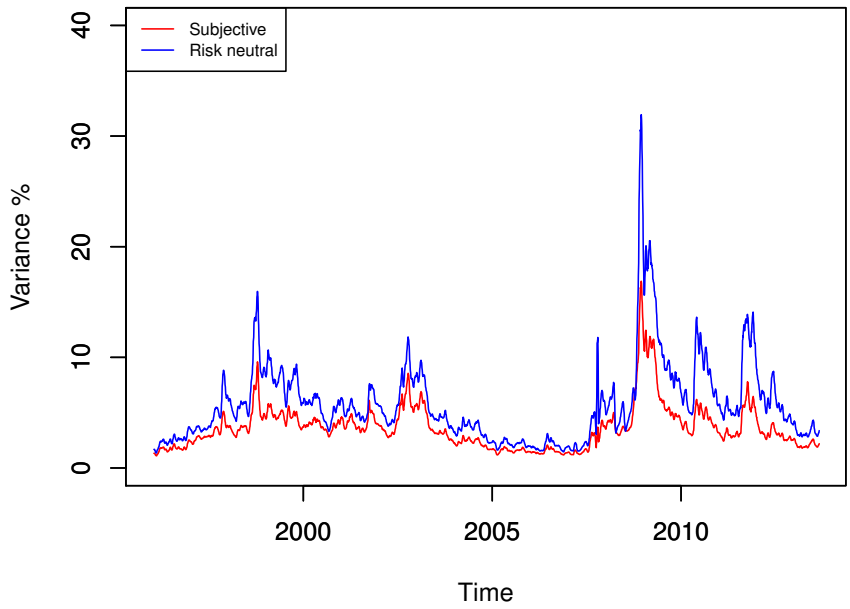

Figure 2.5: Annualized subjective and risk neutral variance, six month horizon. A 10 day moving average filter is applied to all time series.

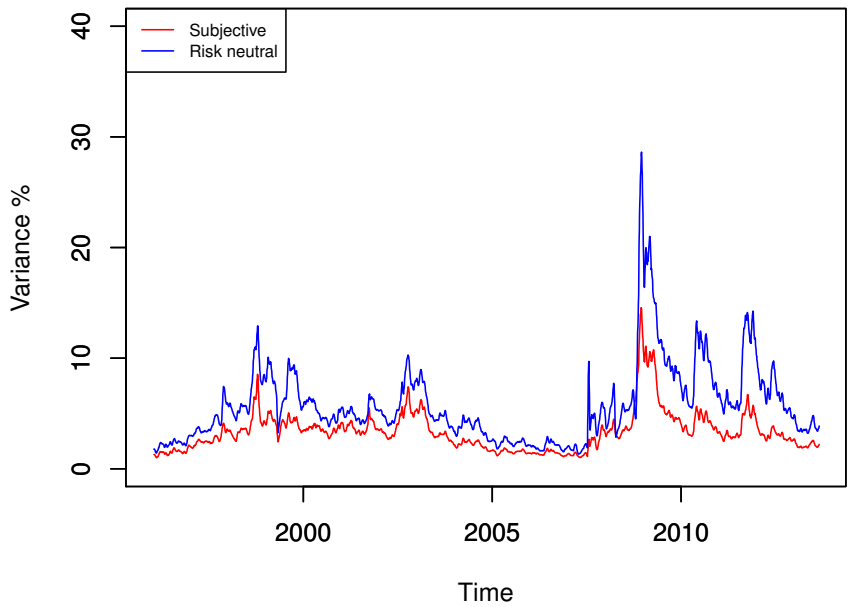

Figure 2.6: Annualized subjective and risk neutral variance, nine month horizon. A 10 day moving average filter is applied to all time series.

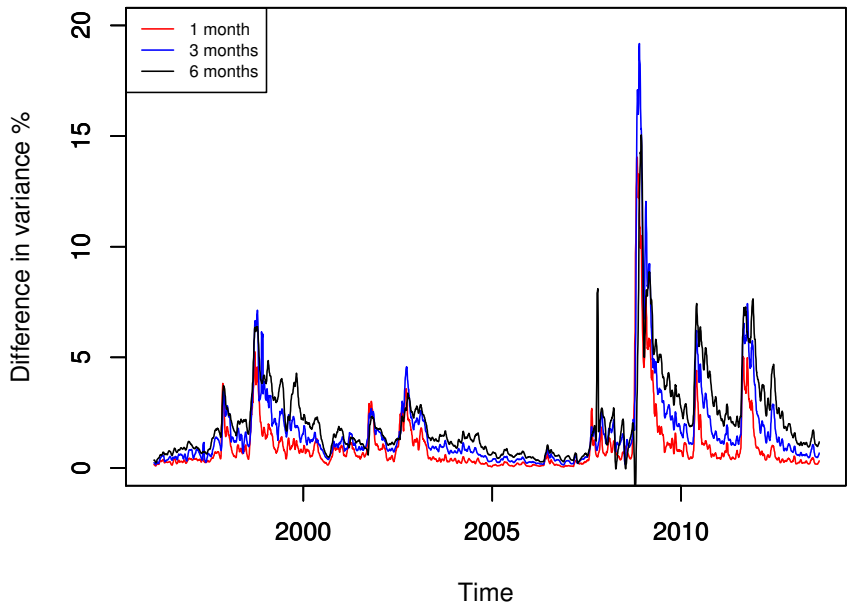

Figure 2.7: Annualized risk neutral vs. subjective variance for three different horizons. A 10 day moving average filter is applied to all time series.

without past realized variance $\text{RVAR}_{t-T,t}$:

$$\text{RVAR}_{t,t+T} = a + b\,\text{var}_t(R_T) + \epsilon_T \tag{2.10}$$

$$\text{RVAR}_{t,t+T} = a + b\,\text{var}_t(R_T) + c\,\text{RVAR}_{t-T,t} + \epsilon_T \tag{2.11}$$

where $\text{RVAR}_{t,t+T}$ is the S&P500 realized variance between time t and $t+T$, and $\text{var}_t(R_T)$ is the log utility variance at time t based on S&P500 option prices. The realized variance $\text{RVAR}_{t,t+T}$ is the sum of the square daily returns over the investment horizons (I use 22 days per month, as is convention). If the conditional variance $\text{var}_t(R_T)$ forecasts the future realized variance, the intercept term a should not be statistically different from 0, and the coefficient b should be close to 1 in both regression specifications. Since I perform these regressions at daily frequency, it involves overlapping data. To adjust for the overlapping regression specification, the standard errors are calculated using the Newey–West adjustment with $T-1$ lags in each of the regressions.

The results using 1, 3, and 6 months horizons are reported in Table 2.5. Results show that log utility conditional variance has forecasting power even when past realized variance is included in the regression specification. The values of the slope coefficient b are statistically not different from one at all three horizons (standard errors are shown in parenthesis). At the one month horizon the subjective variance and VIX^2 show comparable performance. At longer horizons past realized variance ceases to predict future realized variance, as evidenced by the values of the slope coefficient c at the three and six months horizons.

An important signal for the mean–variance optimizing investor is the Sharpe Ratio of the S&P500 index. A time varying Sharpe Ratio conveys expected risk adjusted return information and allows investors to "time" the market. I construct the Sharpe Ratio by taking the ratio of the option implied log investor expected excess return and square root of the variance:

$$\mathrm{SR}_t(R_T) = \frac{\mathbb{E}_t(R_T - R_{f,t})}{\sqrt{\mathrm{var}_t\, R_T}} = \frac{\mathrm{var}^*\, R_T}{R_{f,t}\sqrt{\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^3) - \left(\frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_T^2)\right)^2}}.$$

To determine whether the ex–ante time t market Sharpe Ratio with horizon T predicts the future realized Sharpe Ratio, I implement the following forecasting regression:

$$\mathrm{RSR}_{t,t+T} = a + b\,\mathrm{SR}_t(R_T) + \epsilon_T. \tag{2.12}$$

I expect $a = 0$ and b positive and close to one. Estimation results for the forecasting regression specification are reported in Table 2.6. Except for the 1 month horizon, the ex–ante Sharpe Ratio shows some forecasting power; the highest adjusted R^2 is equal to 1.8% at the 6 months horizon. The slope coefficient b in the regressions is not statistically different from 1, and positive. Also, the intercept a is indistinguishable from zero. The time series of the annualized conditional Sharpe Ratio is depicted in Figure 2.8. It is evident that the Sharpe Ratio is highly volatile, reaching the

Table 2.5: Realized variance forecasting regressions based on Equations (2.10) and (2.11). The dependent variable is the future realized variance, and independent variables include the square of the S&P500 volatility index (VIX^2), the subjective variance from Equation (2.4) estimated at 1 through 6 months horizons, and the past realized variance. Standard errors are corrected for autocorrelation.

	Dependent variable:							
	$\text{RVAR}_{t,t+22}$		$\text{RVAR}_{t,t+22}$		$\text{RVAR}_{t,t+65}$		$\text{RVAR}_{t,t+126}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VIX^2	0.971^{***}	0.578^{***}						
	(0.143)	(0.121)						
$\text{RVAR}_{t-22,t}$		0.321^{*}		0.332^{**}				
		(0.166)		(0.158)				
$\text{var}_t\, R_{t+22}$			1.332^{***}	0.781^{***}				
			(0.222)	(0.185)				
$\text{var}_t\, R_{t+65}$					1.088^{***}	0.923^{***}		
					(0.185)	(0.304)		
$\text{RVAR}_{t-65,t}$						0.090		
						(0.080)		
$\text{var}_t\, R_{t+126}$							0.839^{***}	0.818^{***}
							(0.194)	(0.271)
$\text{RVAR}_{t-126,t}$								0.010
								(0.117)
Const	-0.001	-0.0001	-0.001^{*}	-0.0003	-0.0002	0.001	0.006^{*}	0.006
	(0.0005)	(0.0003)	(0.001)	(0.0004)	(0.001)	(0.002)	(0.003)	(0.004)
Obs	4,425	4,425	4,425	4,425	4,425	4,425	4,425	4,425
Adj R^2	0.493	0.515	0.491	0.517	0.253	0.255	0.121	0.121
Note:								*p<0.1; **p<0.05; ***p<0.01

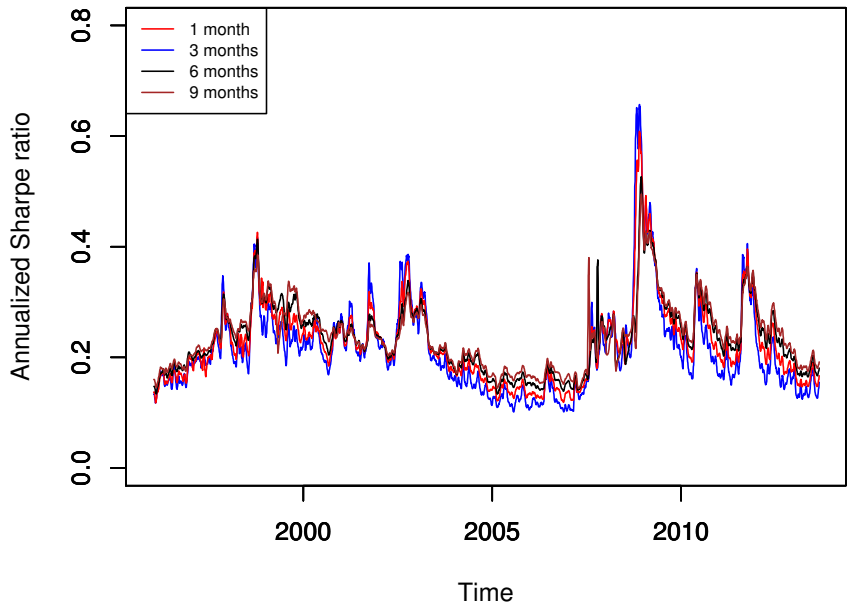

Figure 2.8: Annualized Sharpe Ratio. A 10 day moving average filter is applied to all time series.

highest peak of 0.76 (for the three month horizon) during the 2008 financial crisis. Using forward looking information from the option market can significantly improve the return of a passive cash and index holding strategy.

An interesting empirical finding is that the Sharpe Ratio of the log utility investor seems indistinguishable from the risk neutral volatility of log returns, the VIX. Figure 2.9 shows a graphical comparison of the 1 month annualized log utility investor Sharpe Ratio along with the CBOE VIX index. This observation suggests an alternative interpretation of VIX for the log utility investor. Using VIX an investor could construct a contrarian asset allocation strategy based on market timing. As a simple example, consider a market timing trading strategy that goes long (up to 100% allocation) in the S&P500 when the VIX is greater or equal to 25, and sells the market anytime the VIX is below 25. To ensure the strategy is both realistic and implementable, I assume that the investor does not leverage (borrow more than the initial wealth). The time series of the value of the resulting portfolio and the portfolio asset allocations are shown in Figures 2.10 and 2.11. The simple con-

Table 2.6: Forecasting regressions of the future realized S&P500 Sharpe Ratio on the log utility conditional Sharpe Ratio. The subscripts stand for time horizons 1, 3, 6, and 9 months, respectively.

	Dependent variable:			
	RSR_1	RSR_3	RSR_6	RSR_9
SR_1	0.548			
	(1.725)			
SR_3		1.261		
		(2.069)		
SR_6			2.655	
			(1.867)	
SR_9				1.982
				(2.028)
Constant	0.169	0.138	−0.082	0.002
	(0.122)	(0.246)	(0.339)	(0.480)
Adj R^2	−0.00001	0.003	0.018	0.012
Note:		*p<0.1; **p<0.05; ***p<0.01		

trarian strategy yields an annualized Sharpe Ratio of 0.57 (a 58% increase over the S&P500 Sharpe Ratio of 0.36 over the sample period in consideration). It is worth noting that this strategy profits by allocating 100% of the wealth to the S&P500 during market downturns, as can be observed in Figure 2.11.

2.6.3 Variance risk premium and return predictability

In the previous section I discussed the observed variance premium, the wedge between the risk neutral implied variance and the logarithmic utility investor's implied variance. This is in line with extensive research into the observed spread between the risk neutral implied and realized volatility. Recent research efforts include Bakshi and Madan [2006], who show empirically that risk aversion, left tail risk, and leptokurtic physical return distribution all contribute to the theoretically predicted positive volatility premium. Bollerslev et al. [2009] develop a general equilibrium

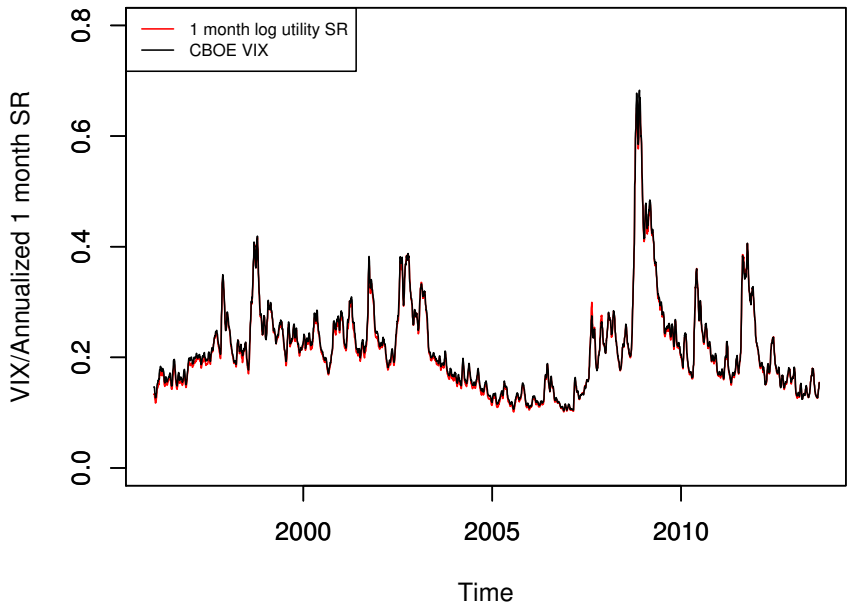

Figure 2.9: One month annualized Sharpe Ratio and VIX. A 10 day moving average filter is applied to all time series.

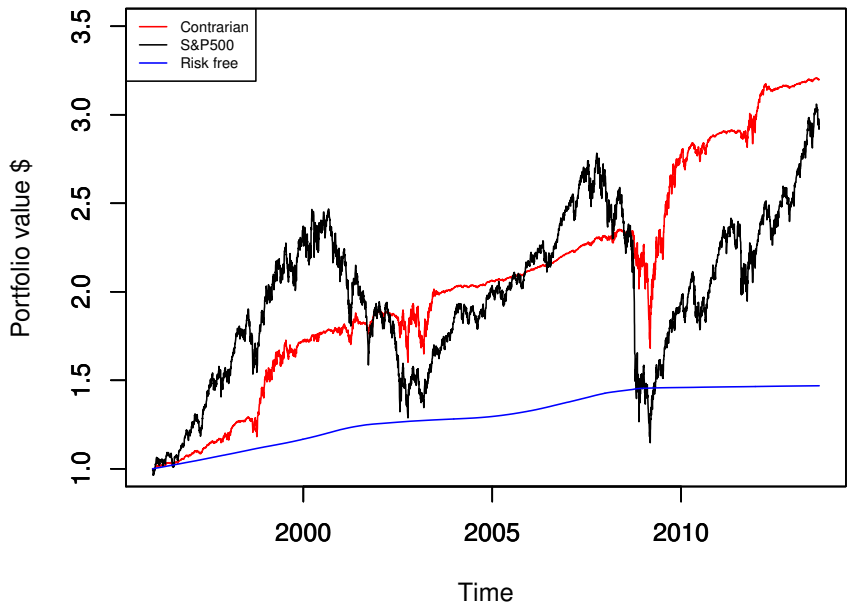

Figure 2.10: Portfolio value of simple market timing contrarian strategy. A 10 day moving average filter is applied to all time series.

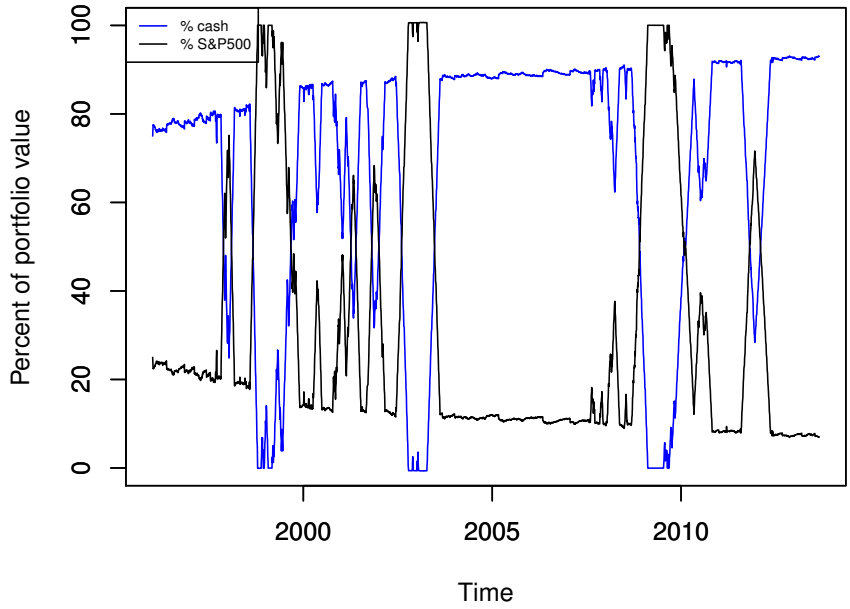

Figure 2.11: Portfolio allocation between the risk free asset and the S&P500 index of simple contrarian trading strategy. A 10 day moving average filter is applied to all time series.

model that includes a variance risk premium, showing empirically that such a premium has strong return forecasting power. Bollerslev et al. [2009] define the variance risk premium as the difference between VIX squared and the expected realized variance, using past realized variance as a proxy of future realized variance. Empirically, they find that the variance risk premium can take on negative values. This result is in contradiction with the theory which suggests that the variance risk premium should always be positive. Motivated by these findings and armed with the log utility investor's conditional variance, I test the return predictability of the variance risk premium defined as the difference between the risk neutral variance of log returns (VIX square) and the log utility conditional variance. Concretely, I estimate the following return forecasting regression:

$$ERET_{t,t+T} = R_{t,t+T} - R_{f,t} = a + b[\mathrm{var}_t^* \log R_T - \mathrm{var}_t R_T] + \epsilon_T,$$

over four different investment horizons: 1, 3, 6, and 9 months. *ERET* stands for excess market return. Figure 2.12 graphically shows the covariation of the six month variance risk premium with the subsequent 6 months return. From this plot there is

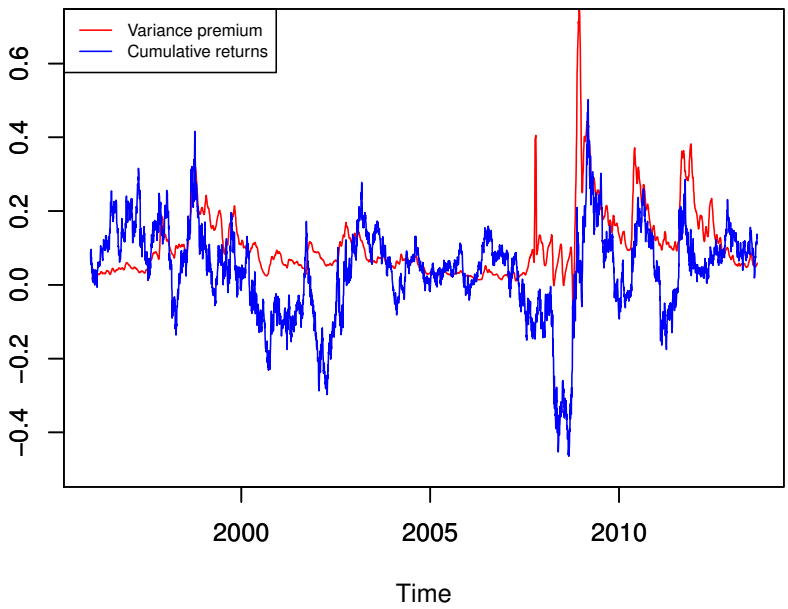

Figure 2.12: Variance risk premium and subsequent six months cumulative returns. The variance premium is multiplied by a factor of 10 plot fitting purpose. A 10 day moving average filter is applied to all time series.

an obvious covariation between the variance premium and the subsequent returns; peaks of the variance risk premium predict future positive returns. To formalize this intuition, the results of the forecasting regression are reported in Table 2.7, where I conclude that the variance risk premium forecast returns at the 3, 6 and 9 months horizons; the regression adjusted R^2 reaches its maximum of 12.2% at the 9 months horizon. Finally, as predicted by theory, the variance premium is positive over the entire sample.

2.7 Conclusions

I obtain a no–arbitrage relation for the moments of the subjective distribution of returns in terms of observable risk neutral moments and covariance terms that involve

Table 2.7: Realized return forecasting regressions based on the variance risk premium. Standard errors are corrected for autocorrelation using Newey–West with $T-1$ lags.

	$ERET_{t,t+22}$	$ERET_{t,t+65}$	$ERET_{t,t+126}$	$ERET_{t,t+189}$
$\text{var}_t^* \log R_{t+22} - \text{var}_t R_{t+22}$	0.329			
	(3.050)			
$\text{var}_t^* \log R_{t+65} - \text{var}_t R_{t+65}$		1.606		
		(2.409)		
$\text{var}_t^* \log R_{t+126} - \text{var}_t R_{t+126}$			4.472^{***}	
			(1.217)	
$\text{var}_t^* \log R_{t+189} - \text{var}_t R_{t+189}$				3.657^{***}
				(0.938)
Constant	0.003	0.003	−0.026	−0.033
	(0.003)	(0.011)	(0.023)	(0.039)
Adj R^2	0.0001	0.010	0.115	0.122
Note:				*p<0.1; **p<0.05; ***p<0.01

the stochastic discount factor. Assuming the log utility preference specification I recover the conditional variance and Sharpe Ratio of the S&P500 index. I test the log utility assumption using moment conditions in terms of both option prices and realized returns, and obtain risk aversion estimates not statistically different from 1 at all investment horizons between 3 to 9 months. I show that the recovered log utility variance has forecasting power controlling for past realized variance. I also observed that the forward looking Sharpe Ratio has forecasting power and can be used as a risk adjusted market timing indicator to improve the return performance of passive indexing.

In light of the findings, there are many possible directions for future research. These include:

1. With the market expected return and conditional variance I may be able to use Equation (B2) in the appendix of BP to estimate the risk aversion coefficient at daily frequency (their Equation assume log–normal returns and power utility; as observed from option prices, the log–normality assumption does not hold).

Applying their Equation (B2) to the S&P500 index gross return R_T:

$$\log\left(\frac{1}{R_{f,t}}\,\mathbb{E}_t\,R_T\right) = \gamma_t\,\mathrm{var}_t\,R_T.$$

2. Using the cross–section of equity and index options, if I specify the power utility I can estimate the risk aversion coefficient based on the following moment conditions:

$$\mathrm{cov}_t(R_T^{-\gamma}R_{i,T}, R_{i,T}) - R_{f,t}\,\mathrm{cov}_t(R_T^{-\gamma}, R_{i,T}) = \frac{1}{R_{f,t}}\,\mathbb{E}_t^*(R_{i,T}^2) - R_{f,t},$$

where $i = 1, \cdots, n$, where n the total number of securities. To simplify the moment condition I can linearize the stochastic discount factor:

$$R_T^{-\gamma} = \exp(-\gamma \log R_T) \approx 1 - \gamma \log R_T.$$

This approximation is only valid at short return horizons (the linear approximation of the stochastic discount factor can become negative for large horizons). Substitute in the previous equation:

$$\frac{\mathrm{var}^*(R_T)}{R_{f,t}} - \mathrm{var}_t\,R_{i,T} \approx \gamma\big[\,\mathrm{cov}_t(R_{i,T}\log R_T, R_{i,T}) - R_{f,t}\,\mathrm{cov}_t(\log R_T, R_{i,T})\big].$$

Using realized return data, the risk aversion coefficient can be estimated at daily frequency in a cross–sectional regression setting.

Chapter 3

Expected returns: systematic risk of firm characteristics?

3.1 Introduction

In this paper I provide empirical evidence that the cross–sectional differences in individual firms' expected excess returns can be attributed to the exposure to systematic factors constructed based on the firm characteristics, as opposed to direct cross–sectional difference in firm characteristics. There is still an ongoing debate within the existing literature as to whether the cross section of expected returns is "explained" by the loadings to systematic factors or to firm characteristics (see, Daniel and Titman [1997] for one of the original papers supporting the firm characteristic story, and Chordia et al. [2015] for the most recent work supporting such interpretation).

The theoretical asset pricing literature predicts that expected returns are determined by the covariance of firm returns with consumption growth (the consumption asset pricing model of Lucas [1978]). A notable exception within the literature is the recent work of Lin and Zhang [2013], who present a theoretical argument that supports the equivalence of the systematic risk and the firm characteristics stories. In a general equilibrium framework, Lin and Zhang [2013] model both the con-

sumption and production sides of the economy, and show that the covariance with characteritics based factors and the cross–sectional difference in expected returns associated with firm characteristics are a theoretically equivalent statement. The consumption side claims that the risk premia are due to covariance of excess returns with systematic factors. The production (or investment) side (for example, based on the seminal paper of Cochrane [1991]) establishes a relationship between expected returns and specific firm characteristics such as earnings and investment to capital ratio. According to their theoretical result, it is unsurprising to find strong association between cross–sectional differences in (specific) firm characteristics and expected returns. Of course, this statement does not give the researcher a license to fish for factors that appear to "determine" cross–sectional differences in expected returns in any given sample; production asset pricing theory is very specific as to which firm characteristics should be related to expected returns and why (Cochrane [1991]).

Empirically, the asset pricing literature on linear factor models has produced a disproportionate number of factors related to so called pricing "anomalies" (Harvey et al. [2016]), excess expected returns associated with specific firm characteristics (such as, gross profits, return on assets, asset growth) that are not captured by exposure to, for example, the market excess return as predicted by the Capital Asset Pricing Model (CAPM). These asset pricing anomalies are typically found by forming portfolios sorted on a selected anomaly just as Fama and French [1993] do with market value of equity and book–to–market ratio to jusitfy their market return, firm size and value factor model (FF3). The time series regression (namely, the method first proposed by Black et al. [1972] (BJS)) of a long–short factor formed from the anomaly portfolios is then used to show that prevailing models (such as the CAPM or the FF3 model) are unable to account for the cross–sectional difference in expected returns captured by the anomaly long–short portfolio.

Extended factor models that include the more "robust" (persistent over time) anomalies include the Fama French 5 factor model (FF5) of Fama and French [2015],

a model that includes two new factors based on firm's profitability and investment cross–sectional differences, and the recent factor models of Novy-Marx [2013] (which include a gross profitability factor) and Hou et al. [2015]. The model of Hou et al. [2015] (HXZ) represents an extension of the investment approach of Cochrane [1991] that leads to a 4 factor model based on the market excess return, market value of equity, investment, and return on equity, which can allegedly account for about 80 asset pricing anomalies documented in the literature. Although motivated by different theory, both the FF5 and the HXZ models are constructed using the econometric methodology of Fama and French [1992] in forming factors as long–short portfolios derived by sorting.

An important limitation of the sorting methodology is that it becomes cumbersome and challenging to perform sorts along more dimensions than 2 (as in the original FF3 model) due to the increased number of permutations in the selection of the sorting order. Fama and French [2015] are not able to summarize the FF5 model's results as they originally did for the FF3 model (see Table 1 of Fama and French [1996]). With the increased number of possible sorting order permutations, it is not obvious how (and which criterion to adopt) to select the "most meaningful" intersection of the four firm characteristics. The FF5 is ultimately constructed using the intersection of three 2×3 value-weight portfolios formed on size and book–to–market, size and operating profitability, and size and investment. Finally, as portfolios sorts are performed one variable at a time, it is virtually impossible to check ex–ante whether candidate factors are redundant in the presence of others. For example, Fama and French [2015] show that their book–to–market factor, HML, is (ex–post) redundant in "explaining" average returns, its effect being subsumed by the profitability and investment factors.

In this study, I utilize a lesser known econometric methodology that enables the extraction of firm characteristic based factors directly from cross–sectional regressions (CSRs) of individual stock excess returns. This method, based on a result from Fama [1976], identifies the slopes of predictive CSRs as being long–short (zero

cost) factor portfolios associated with the corresponding independent characteristic; Back et al. [2015] name these zero–cost portfolios "pure plays" on a specific firm characteristic. I use this technique to extract the firm size and book–to–market factors from the cross section of individual stock returns.

One of the biggest advantages of extracting firm characteristic related factors from CSRs following Fama [1976] is that of avoiding the task of sorting cross–sectional firm returns along each of the individual firm characteristics. As mentioned earlier, sorting returns along more than two dimensions is tricky due to the non–unique sorting order, which in turns yields different factors depending on the sorting order selected. With the CSR method of Fama [1976] the econometrician can obtain unique pure play characteristic related factors from a multiple linear regression. The price to pay for the characteristic based uniqueness is the imposition of a linear structural model.

Hoberg and Welch [2009] use the factors from the CSR method to reverse the inference in Daniel and Titman [1997], and Back et al. [2015] use the slopes ("pure play" factors) from cross–sectional regressions of excess returns on firm characteristics to test the FF5 and HXZ factor models. The authors compare the slopes of cross–sectional regressions on firm characteristics and loadings associated with the first stage time series regression on the respective firm characteristic factor, and conclude that both the FF5 and HXZ models fail to explain the returns to investment and momentum neutral factors.

Avramov and Chordia [2006] use risk–adjusted returns in cross–sectional regressions to show that size and book–to–market characteristics have no incremental explanatory power with respect to the CAPM. I extend their insight and construct size and book–to–market factors from risk–adjusted individual firm returns to show that firm characteristics are related to the covariance of returns but not to expected returns. Using such a test, I can also determine what proportion of the covariance of returns is captured by the factor loadings and firm characteristics. The results that I present in this paper are in contradiction with the conclusions presented by

Chordia et al. [2015], who suggest that firm characteristics take up the "lion's share" of cross–sectional expected returns coefficient of determination R^2.

I show that the factors constructed from firm risk–adjusted excess returns, the residual returns that result from a time series regression of firm excess returns on the firm characteristic based factors, have statistically insignificant alpha and low correlation with the factors used to risk–adjust the firm excess returns. At a minimum, this result shows that unconditional risk–adjustment of firm excess returns is effective at capturing the cross–sectional difference in expected returns observed across portfolios constructed based on sorts along the firm characteristics that determine the factors. Interpreted differently, the empirical results in this paper provide evidence that, even if cross–sectional differences in expected returns exist due to firm characteristics, loadings on the factors constructed from the cross section based on the firm characterstics can account for the difference in expected return. To some extent, the empirical evidence I provide in this paper supports the view that systematic risk (exposure to factors) and firm characteristics are associated.

I apply the CSR technique on a risk–adjusted sample of individual firm excess returns, and use the associated "risk–adjusted" factors as test portfolios to determine whether the such factors have explanatory power. If cross–sectional variation of expected returns is associated with firm characteristics and not with loadings on the associated factors, the risk–adjusted factors should have a statistically significant intercept (α) in a time series regression on the original (risk–unadjusted) factors. I also show that the cross–sectional explanatory power of the risk–adjusted factors is greatly reduced when compared to the power associated with the original factors. As an illustrative example, I use both the Fama French 25 size and book–to–market and 30 industry portfolios to show that risk–adjusted factors associated with size and book–to–market have almost no explanatory power compared to the original Fama French SMB and HML. This example provides compelling evidence that portfolios' expected returns are definitely accounted for by their loadings (exposure) to the underlying systematic factors.

3.2 Properties of OLS cross–sectional regressions

In this section I outline some interesting properties of ordinary least square (OLS) cross–sectional regressions of excess returns on firm characteristics. I will use OLS cross–setional regressions to construct systematic factors associated with firm characteristics, as opposed to the common method of sorting portfolios based on firm characteristics. I discuss the tradeoff in constructing factors by sorting and running cross–sectional regressions. Finally, I show how the OLS properties also hold for weighted least square (WLS) regressions.

The econometric method of extracting factors based on firm characteristics directly from cross–sectional regressions is first reported by Eugene Fama in Chapter 9 of "Foundations of Finance" (Fama [1976]). Fama effectively shows that there is an equivalence between zero–cost portfolios obtained via sorting and the slopes of the cross–sectional regression of excess returns on firm characteristics. Of course, the major differences being that sorting does not assume a structural model, and that sorting cannot really accomodate multivariable (read greater than 2) sorts (due to the increased sorting order combinations). In the context of a cross–sectional test of the CAPM, Fama shows that the least square coefficients (the slopes associated with each regressor other than the intercept) of the cross–sectional regression can be interpreted as returns of zero–cost portfolios formed on the respective regressor variable (typically a firm characteristic). To my knowledge, although apparently common knowledge in the finance community, Fama's result is first mentioned and utilized by Haugen and Baker [1996], later implemeted in Daniel and Titman [2006] to construct value weighted long–short portfolios, and more recently thoroughly described and analyzed by Hoberg and Welch [2009] in an attempt to determine whether the FF3 factors can price the respective incongruent portfolios first proposed by Daniel and Titman [1997].

I will summarize the basic result of Chapter 9 in Fama [1976]. Consider the

following ordinary least square (OLS) regression:

$$\boldsymbol{r}_{t+1} = \gamma_0 + \gamma_1 \log(ME_t) + \gamma_2 \log(BM_t) + \boldsymbol{\epsilon}_{t+1} \tag{3.1}$$

where $\boldsymbol{r}_{t+1}$ is a vector of firms' excess returns, the regressors $\log(ME_t)$ and $\log(BM_t)$ are vectors of the firms' corresponding log market capitalization and book–to–market values, and $\boldsymbol{\epsilon}_{t+1}$ is the (normally distributed) residual error. Market capitalization and book–to–market values may be demeaned at each cross section following Brennan et al. [1998] to ensure that the intercept captures the cross–sectional (equal weighted) mean of the excess returns. Back et al. [2015] demean and standardize each regressor by the respective cross–sectional standard deviation in order to facilitate comparisons between firm characteristics' contributions to the association with excess returns. Fama [1976] shows that the coefficients γ_1 and γ_2 can be interpreted as the return to zero–cost portfolios formed based on market capitalization and the book–to–market ratio, respectively. To see this, consider the equivalent matrix form representation of Equation (3.1):

$$\boldsymbol{r}_{t+1} = X\boldsymbol{\gamma} + \boldsymbol{\epsilon}_{t+1}. \tag{3.2}$$

where $X = [\ \mathbf{1} \quad \log(ME_t) \quad \log(BM_t)\ ]$. The ordinary least square solution $\boldsymbol{\gamma}_{OLS}$ is given by:

$$\boldsymbol{\gamma}_{t,OLS} = (X'X)^{-1}X'\boldsymbol{r}_{t+1}.$$

The interesting peculiarity of the OLS solution pointed out by Fama [1976] is the fact that the sum of the rows of $(X'X)^{-1}X'$ is equal to the vector $[\ 1 \quad 0 \quad 0\ ]$. This observation, along with the fact that the elements of the same matrix are weights assigned to each firms return in the cross section, implies that the first row weights represent a long only portfolio (as the row sum is equal to 1), whereas all other rows have weights for zero–cost (long–short) portfolios associated with the regressor's

characteristic (portfolios whose weights sum to 0 are zero cost long–short portfolios). The standard OLS solution results in equal weighted portfolios. If we want to control for cross–sectional differences in market capitalization we can perform a weighted least square regression; the zero cost portfolio property still holds. This fact is discussed in Hoberg and Welch [2009]. As the slopes are obtained by least square regression, the zero–cost portfolio returns have the nice property of being minimum variance returns. The weighted least square (WLS) cross–sectional regression of excess returns on the size and book–to–market firm characteristics is given by:

$$\boldsymbol{r}_{t+1} = \gamma_{VW,t}\mathbf{1} + \gamma_{ME,t}\log(ME_t) + \gamma_{BM,t}\log(BM_t) + \boldsymbol{\epsilon}_{t+1}$$

where $\boldsymbol{r}_{t+1}$ is the vector of excess returns at $t+1$, $\mathbf{1}$ is a vector of ones, $\log(ME_t)$ and $\log(BM_t)$ are vectors of the logarithm of the market value of equity and book–to–market for each firm in the cross section at time t. Both market value of equity and book–to–market values are updated yearly as in Fama and French [1993]. The matrix of weights in the cross–sectional WLS regression is strictly diagonal with lagged market capitalization values of each firm in the cross section. This matrix can be raised to a power between 0 (equal weighted market return as the intercept) and 1 (value weighted market return as the intercept). The WLS slope estimates preserve the zero–cost portfolio property of the standard OLS regression Hoberg and Welch [2009]. Its solution is of the form:

$$\boldsymbol{\gamma}_{WLS,t} = (X'\Omega^k_{ME,t}X)^{-1}X'\Omega^k_{ME,t}\boldsymbol{r}_{t+1}, \tag{3.3}$$

where k is an exponent between 0 and 1 (Hoberg and Welch [2009]). Implementing the WLS approach avoids running separate cross–sectional Fama MacBeth regressions for micro and all but micro stocks as in Fama and French [2008]; WLS enables me to obtain value weighted zero–cost portfolios that account for the cross–sectional heterogeneity in firm size (micro stocks account for 60% of the number of stocks but less than 3% of the total market capitalization). The three slopes $\gamma_{VW,t}, \gamma_{ME,t}, \gamma_{BM,t}$

are respectively the returns on a long value weighted portfolio of all stocks in the cross section, and the long short portfolios on size and book–to–market.

3.3 Data

To execute this study, I obtain the U.S. merged universe of firms in Compustat and CRSP from July 1964 till December 2013, inclusive. The merged sample is formed using the definitions and filtering criteria described in Fama and French [1996] and Davis et al. [2000]. I obtain the Compustat/CRSP merged data from the Wharton Research Data Services, and the FF3 factors, the Fama French 25 size and book–to–market (FF25) and 30 industry (FF30) portfolio returns from Prof. Kenneth French's website[3]. To remove extreme observations, I winsorize firm characteristics (the regressors, not the firm returns) at the 99.5 percentile after filtering the data sample (for example, the filtering removes firms whose book value of equity is less than or equal to 0). Winsorization is standard practice in empirical asset pricing, see for example Lewellen [2015].

In Table 3.1 I report the sample descriptive statistics for the firm characteritics commonly considered in the literature (see for example, the recent study of Lewellen [2015]). I use the following definitions to construct the variables listed in Table 3.1: BE is the book value of equity (calculated according to Davis et al. [2000], ME is the market value of equity, AT is total assets, IB is the income before extraordinary items, $REVT$ is sales revenue, $COGS$ is cost of goods sold, $XSGA$ is expense for sales and general administration, and $XINT$ is interest expense, all provided by Compustat. Table 3.1 entries are calculated as follows:

- (log book–to–market ratio) logBEME $= \log(BE_{t-1}) - \log(ME_{t-1})$,
- (log December market value of equity) logDECME $= \log(ME_{t-1})$,
- (Asset growth) logAG $= \log(AT_{t-1}) - \log(AT_{t-2})$,
- (log market leverage) logMLev $= \log(AT_{t-1}) - \log(ME_{t-1})$,

[3]http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html, accessed January 2016.

- (log book leverage) logBLev $= \log(AT_{t-1}) - \log(BE_{t-1})$,
- (Return on equity) ROE $= (IB_{t-1})/(BE_{t-1})$,
- (Return on asset) ROA $= (IB_{t-1})/(AT_{t-1})$,
- (Fama and French [2015] operating profit) OPFF $= (REVT_{t-1} - COGS_{t-1} - XSGA_{t-1} - XINT_{t-1})/(BE_{t-1})$,
- (Novy-Marx [2013] operating profit) OPNM $= (REVT_{t-1} - COGS_{t-1})/(AT_{t-1})$,
- (Fama and French [2015] investment growth) FFINV $= (AT_{t-1} - AT_{t-2})/(AT_{t-2})$.

Note that the subscripts $t-1$ and $t-2$ refer to one and two years prior respectively to the return data, a timing convention first established by the study of Fama and French [1992]. The market value of equity at $t-1$ is the December figure corresponding to the year for which the book value of equity is calculated from accounting data.

I follow the method of Brennan et al. [1998] and Avramov and Chordia [2006] in risk–adjusting the returns. I risk–adjust excess firm returns in two ways: I run unconditional risk–adjusting regressions (I use entire time series sample for each firm), and I run rolling regressions with 60 months windows for each firm. As a robustness check I also pursue the time varying risk–adjustment using instruments as of Avramov and Chordia [2006], only to conclude that the risk–adjustment on a rolling basis is sufficient to illustrate the research question.

3.4 Results

3.4.1 Extracting factors using WLS regressions

I provide a practical application of the WLS regression of excess returns on firm characteristics by utilizing the slopes from the monthly WLS regressions to estimate the FF3 size and book–to–market factors for the sample period. To ensure that the intercept represents the value weighted market return, I follow Brennan et al. [1998] and demean the market value of equity and the book–to–market ratio of all firms by the cross–sectional means for each month in the sample.

Table 3.1: Descriptive statistics for the merged Compustat and CRSP database sample for the period July 1964–December 2013, inclusive. Variable definitions are: logBEME = $\log(BE_{t-1}) - \log(ME_{t-1})$, logDECME = $\log(ME_{t-1})$, logAG = $\log(AT_{t-1}) - \log(AT_{t-2})$, logMLev = $\log(AT_{t-1}) - \log(ME_{t-1})$, logBLev = $\log(AT_{t-1}) - \log(BE_{t-1})$, ROE = $(IB_{t-1})/(BE_{t-1})$, ROA = $(IB_{t-1})/(AT_{t-1})$, OPFF = $(REVT_{t-1} - COGS_{t-1} - XSGA_{t-1} - XINT_{t-1})/(BE_{t-1})$, OPNM = $(REVT_{t-1} - COGS_{t-1})/(AT_{t-1})$, FFINV = $(AT_{t-1} - AT_{t-2})/(AT_{t-2})$. Consistent with Fama and French [1992], the subscripts $t-1$ and $t-2$ refer to one year and two years prior to the return data. The data is winsorized at the 99.5 percentile.

Variable	Mean	St. Dev.	Min	Max
logBEME	−0.48653	0.91468	−3.72404	1.75879
logDECME	4.64229	2.16047	−0.04219	10.68379
logAG	0.15036	0.35629	−0.78976	2.25479
logMlev	0.43945	1.18085	−2.61471	3.79198
logBlev	0.92470	0.74966	0.03199	3.68272
ROE	−0.03804	0.59689	−5.53508	0.77846
ROA	−0.00183	0.17840	−1.22579	0.27525
OPFF	0.14113	0.49411	−3.76508	1.81625
OPNM	0.33075	0.28718	−0.60859	1.38403
FFINV	−0.27342	0.88344	−8.53325	0.54605

Subsequently, I collect the slopes from the WLS regression for the same sample. Similar to Hoberg and Welch [2009], I use $k = 0.5$ for the weighting factor in Equation (3.3). While I perform testing with other values of k between 0 and 1, I utilize 0.5 because with this value I obtain very high correlation between the WLS factors and the respective FF3 factors. For reference, I report the correlation matrix between the factors extracted from the WLS cross–sectional predictive regressions and the corresponding Fama–French value weighted market, size, and book–to–market factors constructed via portfolio sorts in Table 3.2. All three WLS factors are highly

Table 3.2: Correlation matrix between the FF3 factors, namely $RMRF$ (market value weighted excess return), SMB (firm size), HML (book–to–market), and factors (slopes) based on cross–sectional predictive regressions, market excess return γ_{VW}, firm's market value of equity γ_{ME}, and book–to–market γ_{BM}.

	γ_{VW}	γ_{ME}	γ_{BM}	$RMRF$	SMB	HML
γ_{VW}	1					
γ_{ME}	0.639	1				
γ_{BM}	-0.383	-0.461	1			
$RMRF$	0.874	0.264	-0.393	1		
SMB	0.659	0.911	-0.371	0.311	1	
HML	-0.208	-0.254	0.908	-0.303	-0.229	1

correlated with the FF3 factors, with the correlation between γ_{VW} and $RMRF$ at 0.87, γ_{ME} and SMB at 0.91, and γ_{BM} and HML at 0.91. The unconditional means of the γ_{ME} and γ_{BM} factors differ from the unconditional means of the corresponding size and book–to–market FF3 factors, due to the fact that the FF3 factors are leveraged zero–cost portfolios: the sum of the weights on the short side is -1 and the proceeds are invested in the long side with total weight equal to 1. The WLS factors are not leveraged as discussed in Back et al. [2015], and they include all firms in the cross section. The correlation matrix in Table 3.2 is consistent with the study of Hoberg and Welch [2009].

A further test for the cross–sectional factors would be to consider the amount of variation that the WLS factors can "explain" of the Fama French 25 size and book–to–market portfolios (FF25) against the results of Fama and French [1996]. To do this, I run time series regressions of the FF25 excess returns on the three WLS factors to compare their pricing errors, t–statistics, and coefficient of determination, R^2. Concretely, for each of the 25 portfolio excess returns $r_{e,i}$ I run the following

time series regression:

$$\boldsymbol{r}_{e,i} = \alpha_i + \beta_{i,VW}\boldsymbol{\gamma}_{VW} + \beta_{i,ME}\boldsymbol{\gamma}_{ME} + \beta_{i,BM}\boldsymbol{\gamma}_{BM} + \boldsymbol{\epsilon}_i,$$

where $\boldsymbol{r}_{e,i}$ is a vector of excess returns for the ith portfolio, and $\boldsymbol{\gamma}_{VW}, \boldsymbol{\gamma}_{ME}, \boldsymbol{\gamma}_{BM}$ are the WLS factors. In Table 3.3 I summarize the results of the 25 time series regressions performed with the FF3 and WLS factors as regressors. The time series regressions of FF25 on the WLS cross–sectional factors $\boldsymbol{\gamma}_{VW}, \boldsymbol{\gamma}_{ME}, \boldsymbol{\gamma}_{BM}$ result in mispricing α's, t–statistics, and R^2, comparable to the ones obtained using the FF3 factors as regressors. Based on these results, I conclude that the factors extracted from WLS cross–sectional regressions are a valid alternative to the factors constructed via portfolio sorts. Moreover, implementing the weighted least square version of the cross–sectional regression with the firm market value of equity as weights, allows the inclusion of micro capitalization stocks from the sample, something both Fama and French [2008] and Back et al. [2015] are unable to do. WLS with demenead regressors has the additional benefit of obtaining the value weighted market return as the intercept of the cross–sectional estimation. Overall, it seems that the advantages of the linear regression structure listed above outweighs the main criticism of structural misspecification.

3.4.2 Fama MacBeth regression approach

In this section I investigate whether the Fama French size and book–to–market factors are related to systematic risk (Fama and French [1992]) or to firm characteristics (Daniel and Titman [1997]) by using the factors extracted from risk–adjusted WLS cross–sectional regressions. I implement the excess return risk–adjustment methodology adopted by Brennan et al. [1998] and Avramov and Chordia [2006] to individual firm stock returns by performing the following time series regressions for

Table 3.3: Time series regression results of the FF25 size and book–to–market portfolios on the FF3 factors and on the respective WLS cross–sectional factors. The variables reported, namely, R^2, mispricing error, and t–stats, are detailed across the FF25 portfolios, starting with the lowest book–to–market (BM) and smallest (S) market capitalization. The 25th portfolio has the highest BM and biggest (B) market capitalization. Statistics for the FF3 model's $|\alpha|$: minimum 0.006, maximum 0.515, mean 0.100, and median 0.069; for the WLS model's $|\alpha|$: minimum 0.0115, maximum 0.398, mean 0.113, and median 0.166 % monthly. The Gibbons et al. [1989] test statistics for the FF3 model is 90.536 (p value 0) and for the WLS model is 3.972 (p value 8.11×10^{-10})

.

	FF3 factors					WLS factors				
					R^2					
S	0.921	0.944	0.952	0.946	0.947	0.928	0.910	0.915	0.914	0.939
	0.952	0.944	0.939	0.940	0.948	0.917	0.909	0.894	0.896	0.894
	0.950	0.912	0.898	0.901	0.894	0.914	0.903	0.896	0.889	0.861
	0.937	0.890	0.882	0.889	0.876	0.910	0.917	0.920	0.879	0.856
B	0.941	0.902	0.856	0.892	0.803	0.886	0.886	0.848	0.859	0.734
					Mispricing α % monthly					
S	-0.515	-0.017	0.006	0.151	0.123	-0.382	0.070	0.059	0.200	0.165
	-0.178	-0.050	0.105	0.069	-0.046	-0.048	0.016	0.165	0.121	0.019
	-0.051	0.068	0.020	0.057	0.126	0.092	0.128	0.056	0.099	0.189
	0.144	-0.089	-0.037	0.050	-0.083	0.284	-0.042	-0.016	0.092	-0.050
B	0.169	0.037	-0.047	-0.104	-0.168	0.271	0.078	-0.023	-0.070	-0.116
					t–stats					
S	-5.384	-0.243	0.103	2.695	2.069	-4.180	0.790	0.789	2.813	2.573
	-2.644	-0.832	1.834	1.258	-0.795	-0.537	0.203	2.188	1.664	0.234
	-0.807	0.984	0.291	0.864	1.656	1.112	1.758	0.811	1.421	2.158
	2.280	-1.218	-0.495	0.730	-1.015	3.755	-0.652	-0.267	1.289	-0.564
B	3.494	0.619	-0.664	-1.717	-1.769	4.020	1.219	-0.313	-0.999	-1.049
	Low BM to High BM					Low BM to High BM				

each firm in the sample:

$$\boldsymbol{r}_{e,i} = a_i + b_{i,MKT}MKT + b_{i,SMB}SMB + b_{i,HML}HML + \boldsymbol{e}_i,$$

where $\boldsymbol{r}_i$ is a vector of monthly excess returns for firm i, the three regressors are the FF3 factors MKT, SMB, HML. The risk–adjusted returns for each firm i are

given by the sum of the intercept and the vector of residuals:

$$\boldsymbol{r}_{e,i,RA} = a_i + \boldsymbol{e}_i.$$

Such risk–adjustment can be performed either using the entire sample or by implementing rolling regressions of, say, 60 months at a time. Time variation of loadings can be captured by parameterizing the loadings $b_{i,MKT}, b_{i,SMB}, b_{i,HML}$ with lagged instruments, as first shown in Ferson and Schadt [1996]. Avramov and Chordia [2006] use this technique to risk–adjust individual returns over the entire sample, i.e. they implement it in an unconditional form. I report results for both unconditional and rolling regression risk–adjusted WLS factors (see for example, the results in Table 3.6).

I use the risk–adjusted returns for each firm in a WLS cross–sectional regression to extract the zero–cost portfolios corresponding to the market value of equity and book–to–market characteristics. I perform these regressions without including the additional control variables that, for example, Avramov and Chordia [2006] use in their Fama MacBeth regressions, because I am interested in determining whether the factors extracted from the cross section of individual risk–adjusted stock returns carry the same "information" as the factors extracted from non risk–adjusted returns. I compare the explanatory power of risk–adjusted and risk–unadjusted slopes associated with the FF3 factors. The risk–adjusted return WLS regression is as follows:

$$\boldsymbol{r}_{t+1,RA} = \gamma_{VW,t,RA}\mathbf{1} + \gamma_{ME,t,RA}\log(ME_t) + \gamma_{BM,t,RA}\log(BM_t) + \boldsymbol{\epsilon}_{t+1,RA}.$$

The slopes $\gamma_{VW,t,RA,}, \gamma_{ME,t,RA}, \gamma_{BM,t,RA}$ are the zero–cost portfolios (factors) associated with the value weighted risk–adjusted return and their cross–sectional variation along the market value of equity and book–to–market characteristics.

I will show that the risk–adjusted factors $\gamma_{VW,t,RA}, \gamma_{ME,t,RA}, \gamma_{BM,t,RA}$ do not exhibit the explanatory power of the original factors extracted from the excess returns

(risk–unadjusted). Using the same set of test portfolios, namely the Fama French 25 size and book–to–market portfolios (FF25), I perform time series regressions to estimate the loadings of these portfolios on the risk–adjusted factors:

$$\boldsymbol{r}_{e,i} = \alpha_{i,RA} + \beta_{i,VW,RA}\boldsymbol{\gamma}_{VW,RA} + \beta_{i,ME,RA}\boldsymbol{\gamma}_{ME,RA} + \beta_{i,BM,RA}\boldsymbol{\gamma}_{BM,RA} + \boldsymbol{\epsilon}_{i,RA},$$

and report the regression coefficient of determination R^2, the t–statistics and mispricing values ($\alpha_{i,RA}$ intercepts) thus obtained in Table 3.4. If firm characteristics

Table 3.4: Time series regression results of the FF25 size and book–to–market portfolios on the FF3 factors and on the respective risk–adjusted WLS cross–sectional factors. The variables reported, namely, R^2, mispricing error, and t–stats, are detailed across the FF25 portfolios, starting with the lowest book–to–market (BM) and smallest (S) market capitalization. The 25th portfolio has the highest BM and biggest (B) market capitalization. Statistics for the FF3 model's $|\alpha|$: minimum 0.006, maximum 0.515, mean 0.100, and median 0.069; for the risk–adjusted WLS model's $|\alpha|$: minimum 0.232, maximum 1.017, mean 0.714, and median 0.744 % monthly.

	FF3 factors					Risk–adjusted WLS factors				
					R^2					
S	0.921	0.944	0.952	0.946	0.947	0.176	0.114	0.089	0.099	0.126
	0.952	0.944	0.939	0.940	0.948	0.078	0.025	0.011	0.013	0.031
	0.950	0.912	0.898	0.901	0.894	0.071	0.009	0.001	-0.00001	0.008
	0.937	0.890	0.882	0.889	0.876	0.081	0.014	0.009	-0.001	0.010
B	0.941	0.902	0.856	0.892	0.803	0.108	0.009	0.002	0.0001	0.015
					Mispricing α % monthly					
S	-0.515	-0.017	0.006	0.151	0.123	0.232	0.730	0.744	0.896	0.957
	-0.178	-0.050	0.105	0.069	-0.046	0.593	0.722	0.894	0.867	0.885
	-0.051	0.068	0.020	0.057	0.126	0.651	0.812	0.753	0.828	1.017
	0.144	-0.089	-0.037	0.050	-0.083	0.759	0.617	0.690	0.780	0.752
B	0.169	0.037	-0.047	-0.104	-0.168	0.609	0.559	0.462	0.517	0.527
					t–stats					
S	-5.384	-0.243	0.103	2.695	2.069	0.762	2.680	3.086	3.966	4.003
	-2.644	-0.832	1.834	1.258	-0.795	2.038	2.909	3.927	3.922	3.569
	-0.807	0.984	0.291	0.864	1.656	2.418	3.531	3.565	4.006	4.422
	2.280	-1.218	-0.495	0.730	-1.015	3.187	2.850	3.257	3.867	3.273
B	3.494	0.619	-0.664	-1.717	-1.769	3.280	2.997	2.516	2.827	2.520
	Low BM to High BM					Low BM to High BM				

are associated with cross–sectional variation of returns, the R^2 of these three regressions should be very similar. From Table 3.4 it is clear that the risk–adjustment has removed a large portion of the three factor model explanatory power: the highest R^2 value is for the small ME low BM at 0.17, with the rest of the portfolios' R^2 exhibiting values at or below 0.1. The result lends further credibility to the systematic risk nature of both SMB and HML factors: characteristic based factors extracted from the cross section of individual risk–adjusted firm returns have close to no explanatory power when tested against the FF25. The time series excess return risk–adjustment regression removes the cross–sectional expected returns difference between the portfolios. This result is obtained by implementing the risk–adjusting time series regression unconditionally on the entire sample period: using conditional risk–adjustment should make the test even more powerful. A graphical representation of the significantly lower explanatory power of risk–adjusted factors is reported in Figure 3.1 for a set of test portfolios that include both the FF25 and the FF30 industry. I include the FF30 to make the R^2 comparison more challenging for the FF3 factors, following the recommendations in Lewellen et al. [2010]. Figure 3.1 provides ample evidence that the WLS cross–sectional regression based factors have equivalent explanatory power to the sorting based FF3 factors. In stark constrast, risk–adjusted excess return based factors capture significantly lower cross–sectional variation of the test portfolios, with R^2 values across all test portfolios under 0.2.

Table 3.5 provides results pertaining to the Fama MacBeth coefficients (time series average of the slopes) and the standard errors for three regression specifications: in column (1) I report the CSR of excess returns on an intercept and size and book–to–market for benchmark comparison; in columns (2) and (3), I list the results of the same regressions using market risk–adjusted returns and FF3 factor risk–adjusted returns. As a first check of the sample data, I estimate the cross–sectional statistics of Fama and MacBeth [1973] (the time series average of the slopes, and the t–statistics based on standard errors corrected for autocorrelation and heteroskedasticity) and compare the estimates from my sample with the values

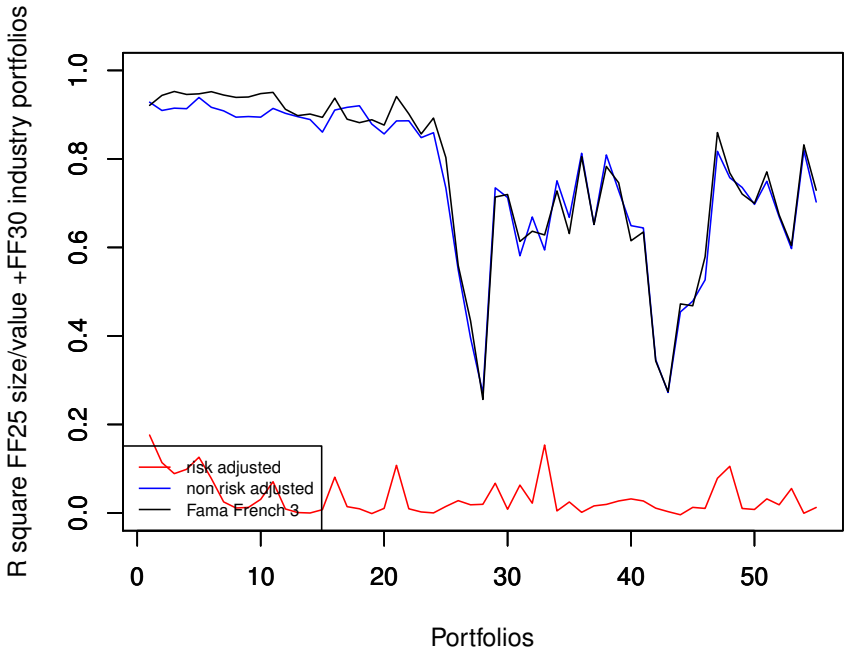

Figure 3.1: Time series of the coefficient of determination (R^2) from regressions of the combined FF25 size and book–to–market and the FF30 industry portfolios on the original FF3 factors, risk–unadjusted, and risk–adjusted WLS cross–sectional factors. On the x–axis the portfolio numbers 1–25 correspond to the FF25 portfolios, and 26–55 to FF30 industry portfolios.

reported in Fama and French [2008] and Lewellen [2015]. The slopes in column (1) for size and book–to–market are in line with those reported in both Fama and French [2008] and the recent study by Lewellen [2015]. I have slightly lower average number of firms, likely due to differences in the sequence of applying filters to the data. The average slopes are significant even after risk–adjusting the returns, but the magnitudes are significantly reduced. As found by Avramov and Chordia [2006] , albeit using more regression independent variables, risk-adjusting the individual firm returns seems to reduce the risk premia associated with the firm characteristics.

3.4.3 Time series approach of Black, Jensen, and Scholes

The classic BJS (Black et al. [1972]) time series test for α can provide more power to determine whether expected returns are related to firm characteristics: Hoberg and Welch [2009] perform such test with factors from the cross section of returns

Table 3.5: Fama MacBeth regression results for three different dependent variables, namely: (1) monthly % excess return; (2) market portfolio ($RMRF$) risk–adjusted excess returns; and, (3) FF3 factors risk–adjusted excess returns. As defined in Seciton 3.3, the independent variables used in the CSR are the log market value of equity, logDECME, and log book–to–market, logBEME.

	Dependent variable:		
	excess return	$RMRF$ risk–adjusted return	FF3 risk–adjusted returns
	(1)	(2)	(3)
Intercept	0.828^{***}	0.376^{***}	0.114^{**}
	(0.241)	(0.118)	(0.057)
logDECME	-0.104^{**}	-0.131^{***}	-0.095^{***}
	(0.043)	(0.042)	(0.027)
logBEME	0.311^{***}	0.289^{***}	0.167^{***}
	(0.070)	(0.060)	(0.038)
Obs	2,206,999	2,120,892	2,120,892
Note:			*p<0.1; **p<0.05; ***p<0.01

based on firm characteristics and associated loadings as controls. Back et al. [2015] show that the error–in–variable bias that affects the estimate of the risk premia associated with the factors present in the two pass (time series and cross–sectional) Fama MacBeth style regressions does not affect the estimate of α in BJS time series tests. This is an important advantage of the BJS time series approach, one that I will utilize when comparing the factors constructed based on CSR of excess returns and risk–adjusted excess returns on firm characteristics.

Instead of appealing to the logic of using cross–sectional regression with control variables as in Hoberg and Welch [2009] (which suffers in part to the error–in–variable problem of using estimated loadings in the second pass regression), I use the BJS test to compare characteristic based factors from the cross section of returns and risk–adjusted (using FF3 factors) returns. If the individual firm returns are explained by the FF3 factors, the CSR factors based on the risk–adjusted returns should not have any α when regressed on the original FF3. risk–adjustment should

remove the expected return through the loadings on the FF3. Moreover, I expect that CSR factors extracted from risk–adjusted returns should have a time series mean not statistically different from zero. On the other hand, the risk–adjusted factors might still have firm characteristics driven correlation with the FF3 factors; I can verify this by looking at the significance of the loadings on the FF3 factors. The regression tests are as follows:

$$\gamma_{k,t} = a + bRMRF_t + cSMB_t + dHML_t + \epsilon_{k,t}$$

versus:

$$\gamma_{k,t,RA} = a_{RA} + b_{RA}RMRF_t + c_{RA}SMB_t + d_{RA}HML_t + \epsilon_{k,t,RA}$$

where a, b, c, d are the parameters of the non risk–adjusted factor regressions, and a_{RA}, b_{RA}, c_{RA}, d_{RA} are the parameters of the risk–adjusted factor regressions, and $k = VW, ME, BM$ for the value weighted, $\log(ME)$ and $\log(BM)$ zero-cost portfolios from the WLS cross–sectional regressions, respectively. Under the null hypothesis that FF3 are risk factors that capture the cross–sectional variation of expected returns, a_{RA} should be zero. In Table 3.6, I report the results of 7 time series regressions. The first three columns (1) , (2), and (3) show the loadings and intercepts of the regressions of the risk–unadjusted factors on the FF3 factors, for reference only. The adjusted R^2 values are between 85.2% and 94.5%, indicating that the risk–unadjusted factors capture a large portion of the FF3 factor variance. The CSR factors are largely subsumed by the FF3, as the α's are statistically not different from zero and economically very small.

In column (4) and (5) I report the coefficients related to the risk–adjusted factor regressions (I only report the risk–adjusted factors associated with the firm characteristics $\log(ME_t)$ and $\log(BM_t)$); as for the first three columns, the intercepts of these two regressions are also statistically indistinguishable from zero, whereas the loadings c and d on SMB and on HML, respectively, are still significant, albeit of

Table 3.6: Time series regression of seven different dependent variables on the FF3 factors. Columns (1), (2), and (3) use risk–unadjusted WLS factors, value weighted market factor (γ_{VW}), market value of equity factor (γ_{ME}), and book–to–market factor (γ_{BM}). Columns (4) and (5) consider unconditional regressions of FF3 risk–adjusted WLS factors for market value of equity ($\gamma_{ME,RA}$) and book–to–market ($\gamma_{BM,RA}$) factors. Finally, columns (6) and (7) report unconditional regressions of FF3 rolling regression risk–adjusted WLS factors for market value of equity ($\gamma_{ME,RAR}$) and book–to–market ($\gamma_{BM,RAR}$) factors. For example, the results in column (1) are obtained from the following regression:

$$\gamma_{VW} = a + bRMRF + cSMB + dHML + \epsilon$$

where $RMRF$ is the market excess return, SMB and HML are the size and book–to–market factors.

	γ_{VW}	γ_{ME}	γ_{BM}	$\gamma_{ME,RA}$	$\gamma_{BM,RA}$	$\gamma_{ME,RAR}$	$\gamma_{BM,RAR}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
a	−0.00001	0.0001	0.0003				
	(0.001)	(0.0001)	(0.0003)				
b	1.031***	−0.007**	−0.037***				
	(0.013)	(0.003)	(0.006)				
c	0.862***	0.257***	−0.082***				
	(0.019)	(0.005)	(0.009)				
d	0.281***	−0.014***	0.530***				
	(0.020)	(0.005)	(0.010)				
a_{RA}				0.0002	0.0001	0.0002	−0.001**
				(0.0001)	(0.0003)	(0.0001)	(0.0003)
b_{RA}				0.003	−0.012	0.002	0.006
				(0.003)	(0.007)	(0.003)	(0.008)
c_{RA}				0.025***	0.033***	−0.009**	0.046***
				(0.004)	(0.010)	(0.004)	(0.011)
d_{RA}				−0.001	0.158***	0.001	0.029**
				(0.005)	(0.011)	(0.004)	(0.012)
Obs.	594	594	594	594	594	535	535
Adj. R^2	0.949	0.852	0.871	0.068	0.274	0.007	0.033
Note:							*p<0.1; **p<0.05; ***p<0.01

significantly reduced magnitude. The loading on SMB c drops by an order of magnitude between regression (2) and (4), whereas d, the loading on HML, is smaller by a factor of approximately 3 in regression (5). This result points to the fact that even after risk–adjustment, firm characteristics and corresponding risk factors still covary in a significant way. The adjusted R^2 for the risk–adjusted factor regressions are much lower, drastically lower for the SMB factor in column (4), but not so for the HML factor in column (5).

Finally, in columns (6) and (7), I report the coefficients associated with the time series regressions of risk–adjusted factors using rolling regressions with a rolling window of 60 months. The adjusted R^2 associated with these two regressions are near zero for both risk–adjusted factors regressions: the two factors extracted from the cross section do not covary at all with the original FF3 factors. Some of the loadings and one intercept α are statistically significant, but economically almost negligible (of opposite sign for the size related risk–adjusted factor, and of one full order of magnitude for the HML risk–adjusted factor if compared to column (3)). As many research studies suggest (see Lewellen and Nagel [2006]), rolling regressions should be more effective at capturing the dynamics of the loadings and the changes in expected returns over time.

Do these results support the systematic risk story for the size and book–to–market effects? In Table 3.7 I show similar BJS regressions comparing the cross–sectional factors based on risk–unadjusted and risk–adjusted using only the market excess return $RMRF$. If the Fama French SMB and HML factors are not proxies for systematic risk, risk–adjusting the excess returns based on the excess market return alone should suffice to "capture" differences in expected return in the cross section. The adjusted R^2 in columns (4) and (5) provide compelling evidence that factors based on excess market return risk–adjustment still show large common variation with Fama French's SMB and HML; moreover, the loadings c and d in columns (2) and (4) and (3) and (5) are largely unchanged. This evidence implies that $RMRF$ risk–adjustment does not capture the size and book–to–market effect.

Table 3.7: Time series regression of five different dependent variables on the FF3 factors. Columns (1), (2), and (3) use risk–unadjusted WLS factors, value weighted market factor (γ_{VW}), market value of equity factor (γ_{ME}), and book–to–market factor (γ_{BM}). Columns (4) and (5) consider unconditional regressions of $RMRF$ risk–adjusted WLS factors for market value of equity ($\gamma_{ME,RA}$) and book–to–market ($\gamma_{BM,RA}$) factors. For example, the results in column (1) are obtained from the following regression:

$$\gamma_{VW} = a + bRMRF + cSMB + dHML + \epsilon$$

where $RMRF$ is the market excess return, SMB and HML are the size and book–to–market factors.

	γ_{VW}	γ_{ME}	γ_{BM}	$\gamma_{ME,RA}$	$\gamma_{BM,RA}$
	(1)	(2)	(3)	(4)	(5)
a/a_{RA}	−0.00001	0.0001	0.0003	0.0002	0.0004
	(0.001)	(0.0001)	(0.0003)	(0.0001)	(0.0003)
b/b_{RA}	1.031***	−0.007**	−0.037***	−0.043***	0.070***
	(0.013)	(0.003)	(0.006)	(0.003)	(0.007)
c/c_{RA}	0.862***	0.257***	−0.082***	0.252***	−0.092***
	(0.019)	(0.005)	(0.009)	(0.005)	(0.009)
d/d_{RA}	0.281***	−0.014***	0.530***	−0.016***	0.479***
	(0.020)	(0.005)	(0.010)	(0.005)	(0.010)
Obs.	594	594	594	594	594
Adj. R^2	0.949	0.852	0.871	0.842	0.815
Note:					*p<0.1; **p<0.05; ***p<0.01

A quick comparison of the results in Tables 3.6 and 3.7 sheds some light on the role the two additional Fama French factors SMB and HML play in the cross section of expected returns. A direct test can be implemented by regressing the $RMRF$ risk–adjusted factors on the FF3 risk–adjusted factors, and check whether the intercepts are statistically different from zero, i.e. whether the $RMRF$ risk–adjustment does not capture parts of the cross–sectional difference in expected returns. If this were the case, it would imply that risk adjusment based only on the market excess return does not "explain" the cross–setional variation of returns. Concretely, I run

Table 3.8: Time series regression of three market excess return risk–adjusted factors on FF3 risk–adjusted factors. For example, the results in column (1) are obtained from the following regression:

$$\gamma_{VW,RA_{RMRF}} = a + b\gamma_{VW,RA_{FF3}} + c\gamma_{ME,RA_{FF3}} + d\gamma_{BM,RA_{FF3}} + \epsilon$$

where $RMRF$ is the market excess return, SMB and HML are the size and book–to–market factors.

	$\gamma_{VW,RA_{RMRF}}$ (1)	$\gamma_{ME,RA_{RMRF}}$ (2)	$\gamma_{BM,RA_{RMRF}}$ (3)
a	0.001	0.0001	0.001***
	(0.001)	(0.0003)	(0.0004)
b	0.696***	−0.109***	0.062
	(0.103)	(0.029)	(0.046)
c	3.385***	1.879***	−0.354*
	(0.408)	(0.115)	(0.183)
d	0.984***	0.105***	1.308***
	(0.116)	(0.033)	(0.052)
Obs.	594	594	594
Adj. R^2	0.356	0.378	0.558
Note:		*p<0.1; **p<0.05; ***p<0.01	

the following time series regressions (for convenience, I drop the time subscript t):

$$\gamma_{VW,RA_{RMRF}} = a + b\gamma_{VW,RA,FF3} + c\gamma_{BM,RA_{FF3}} + d\gamma_{ME,RA_{FF3}} + \epsilon_{VW,RA_{FF3}}$$

and test the significance of the a parameters for each of the three factors. The results of such regressions are reported in Table 3.8. The main takeaway from Table 3.8 is the significance of the intercept a for the book–to–market $RMRF$ risk–adjusted factor. Although economically irrelevant, the excess return a points to the fact that risk–adjustment based on market excess return alone does not fully capture the cross–sectional difference in expected returns.

3.5 Conclusions

In this paper I show that the slopes of monthly cross–sectional regressions of firm excess returns on characteristics can be used to establish whether differences in the cross-section of expected returns are determined by loadings on systematic factors or firm characteristics. To disentangle the covariance wth systematic factors from the firm characteristics, I compare the slopes obtained from the firm excess returns with those extracted from cross–sectional regressions of market or FF3 risk–adjusted excess returns. This analysis yields empirical evidence that risk–adjusting the excess returns captures cross–sectional differences in expected returns. The empirical evidence I provide does not address the theoretical aspects of the asset pricing debate on the determinants of expected returns, whether it is the covariance with systematic factors or the cross–sectional differences in firm characteristics. I also do not delve with the issue that size and book–to–market effects may not be proxies for macroeconomic risk factors (a point first raised via empirical evidence by Daniel and Titman [1997]). These topics are left for future research.

References

Yacine Aït-Sahalia and Andrew W Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2):499–547, 1998.

Yacine Aït-Sahalia and Andrew W Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1):9–51, 2000.

Andrew Ang. *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press, 2014.

Doron Avramov and Tarun Chordia. Asset pricing models and financial market anomalies. *Review of Financial Studies*, 19(3):1001–1040, 2006.

Kerry Back, Nishad Kapadia, and Barbara Ostdiek. Testing factor models on characteristic and covariance pure plays. *Available at SSRN: http://ssrn.com/abstract=2621696*, 2015.

Gurdip Bakshi and Dilip Madan. A theory of volatility spreads. *Management Science*, 52(12):1945–1956, 2006.

Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143, 2003.

Fischer Black, Michael C Jensen, and Myron S Scholes. The capital asset pricing model: Some empirical tests. *Studies in the Theory of Capital Markets, Praeger Publishers Inc*, 1972.

Robert R Bliss and Nikolaos Panigirtzoglou. Option-implied risk aversion estimates. *The Journal of Finance*, 59(1):407–446, 2004.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492, 2009.

Jaroslav Borovička, Lars P Hansen, and José A Scheinkman. Misspecified recovery. *Journal of Finance*, 71:2493–2544, 2016.

Douglas T Breeden and Robert H Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, pages 621–651, 1978.

Michael J Brennan, Tarun Chordia, and Avanidhar Subrahmanyam. Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of Financial Economics*, 49(3):345–373, 1998.

David P Brown and Michael R Gibbons. A simple econometric approach for utility-based asset pricing models. *The Journal of Finance*, 40(2):359–381, 1985.

Peter Carr and Dilip Madan. Optimal positioning in derivative securities. *Quantitative Finance*, 1:19–37, 2001.

Tarun Chordia, Amit Goyal, and Jay A Shanken. Cross-sectional asset pricing with individual stocks: betas versus characteristics. *Available at SSRN 2549578*, 2015.

John H. Cochrane. Production-based asset pricing and the link between stock returns and economic fluctuations. *The Journal of Finance*, 46(1):209–237, 1991.

John H. Cochrane. *Asset pricing*. Princeton University Press, 2005.

Jennifer Conrad, Robert F Dittmar, and Eric Ghysels. Ex ante skewness and expected stock returns. *The Journal of Finance*, 68(1):85–124, 2013.

Kent Daniel and Sheridan Titman. Evidence on the characteristics of cross sectional variation in stock returns. *The Journal of Finance*, 52(1):1–33, 1997.

Kent Daniel and Sheridan Titman. Market reactions to tangible and intangible information. *The Journal of Finance*, 61(4):1605–1643, 2006.

James L Davis, Eugene F Fama, and Kenneth R French. Characteristics, covariances, and average returns: 1929 to 1997. *The Journal of Finance*, 55(1):389–406, 2000.

Patrick Dennis and Stewart Mayhew. Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis*, 37(03):471–493, 2002.

James D Esary, Frank Proschan, and David W Walkup. Association of random variables, with applications. *The Annals of Mathematical Statistics*, 38(5):1466–1474, 1967.

Eugene F Fama. *Foundations of finance*. Basic Books, 1976.

Eugene F Fama and Kenneth R French. The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465, 1992.

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

Eugene F Fama and Kenneth R French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):55–84, 1996.

Eugene F Fama and Kenneth R French. Dissecting anomalies. *The Journal of Finance*, 63(4):1653–1678, 2008.

Eugene F Fama and Kenneth R French. A five–factor asset pricing model. *Journal of Financial Economics*, 116:1–22, 2015.

Eugene F Fama and James D MacBeth. Risk, return, and equilibrium: Empirical tests. *The Journal of Political Economy*, 1:607–636, 1973.

Wayne E Ferson and Rudi W Schadt. Measuring fund strategy and performance in changing economic conditions. *The Journal of Finance*, 51(2):425–461, 1996.

Stephen Figlewski. Estimating the implied risk neutral density for the u.s. market portfolio. *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, 2010.

Michael R. Gibbons, Stephen A. Ross, and Jay Shanken. A test of the efficiency of a given portfolio. *Econometrica*, 57(5):1121–1152, 1989.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.

Campbell R Harvey, Yan Liu, and Heqing Zhu. ... and the cross-section of expected returns. *Review of Financial Studies*, 29:5–68, 2016.

Robert A Haugen and Nardin L Baker. Commonality in the determinants of expected stock returns. *Journal of Financial Economics*, 41(3):401–439, 1996.

Gerard Hoberg and Ivo Welch. Optimized vs. sort-based portfolios. In *AFA 2010 Atlanta Meetings Paper*, 2009.

Kewei Hou, Chen Xue, and Lu Zhang. Digesting anomalies: An investment approach. *Review of Financial Studies*, 28(3), 2015.

Jens Carsten Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451, 2000.

Jens Carsten Jackwerth and Mark Rubinstein. Recovering probability distributions from option prices. *The Journal of Finance*, 51(5):1611–1631, 1996.

Jonathan Lewellen. The cross section of expected stock returns. *Critical Finance Review*, 4:1–44, 2015.

Jonathan Lewellen and Stefan Nagel. The conditional capm does not explain asset-pricing anomalies. *Journal of Financial Economics*, 82(2):289–314, 2006.

Jonathan Lewellen, Stefan Nagel, and Jay Shanken. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics*, 96(2):175–194, 2010.

Xiaoji Lin and Lu Zhang. The investment manifesto. *Journal of Monetary Economics*, 60(3):351–366, 2013.

Robert E Jr Lucas. Asset prices in an exchange economy. *Econometrica: Journal of the Econometric Society*, pages 1429–1445, 1978.

I.W. Martin. What is the expected return of the market? *Quarterly Journal of Economics*, 132:367–433, 2017.

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

Steve Ross. The recovery theorem. *The Journal of Finance*, 70(2):615–648, 2015.

Mark Rubinstein. Implied binomial trees. *The Journal of Finance*, 49(3):771–818, 1994.