Option pricing and portfolio choice

by

Sirio Aramonte

A dissertation submitted to the University of London
for the degree of Doctor of Philosophy

Department of Finance
London Business School
University of London

Declaration

The work presented in this dissertation is entirely my own.

Sirio Aramonte

a Tobias e Viviane

Abstract

The focus of this dissertation is on option pricing, in particular on the economic determinants of option risk premia, and on the interaction between investor sentiment and innovation in the context of technology diffusion. The first of the three chapters empirically investigates whether macroeconomic uncertainty is a priced risk factor in the cross-section of equity option returns. The analysis employs a factor model, estimated with the Fama-MacBeth methodology, and the macroeconomic uncertainty factor is based on options' "excess" pricing errors on days immediately before scheduled macroeconomic announcements. I find that macroeconomic uncertainty is priced in the cross-section of equity option returns, even after controlling for a large set of relevant factors. In addition, introducing the macroeconomic uncertainty factor affects the estimated risk premia on the market, the volatility of volatility and the volatility of jumps. The results are robust to measurement error in stock and option prices, to possible biases generated by the non-linearity of option returns and by non-randomly missing returns, and to several methods of measuring macroeconomic uncertainty and expected volatility. The second chapter studies how the interaction between technological innovation and investor sentiment affects firm-level investment and aggregate productivity growth. When firms face the decision to adopt a new technology with uncertain productivity, small scale experimentation is a direct way of obtaining information useful to evaluate full-scale adoption. If such information is not appropriable, a free-rider problem arises and the aggregate level of investment is sub-optimal. I hypothesize that investor sentiment mitigates the effect of informational externalities, which would make investor sentiment socially valuable in the context of technology diffusion. I find that investor sentiment increases the effect of technological innovation on investment for firms more susceptible to informational externalities, and that investor sentiment also raises the impact of technological innovation on productivity growth. The third chapter focuses on portfolio choice, investigating whether distributions implied in option prices can help reducing estimation error for expected returns, given that the cross-section of option prices contains information about the moments of expected returns that is not available in the underlying's prices. I study the performance of portfolio strategies based on option-implied distributions relative to that of the robust $\frac{1}{N}$ allocation rule, and I find that option-based models outperform the $\frac{1}{N}$ strategy, especially with (limited) short-selling. Transaction costs are relatively high for option-based strategies, and they reduce, but do not eliminate, the economic significance of the results.

Acknowledgements

I am very grateful to Raman Uppal, my advisor, Stephen Schaefer, Elroy Dimson and Tim Johnson for guidance and support during my studies. I've been fortunate to share my years at London Business School with Stefano, David, Theodosios, Ebrahim, Chao, Georg, Carlo, Oguzhan, Jungsuk and Federica.

Contents

List of Tables

List of Figures

Overview

The option pricing literature has, so far, mostly focused on explaining option prices with reduced-form models, that characterize the distribution of observed returns on the underlying and implicitly define investor preferences through the choice of the risk-neutral distribution. As Bates (2003) notes, however, it is becoming increasingly important to give clearer economic foundations to the differences between the properties of observed asset returns and of returns implied by option prices.

The first chapter of my dissertation is a contribution to the growing literature that addresses the point made by Bates. I study whether macroeconomic uncertainty is a priced risk factor in the cross-section of equity option returns, and how its inclusion changes the risk premia on other factors, in particular volatility and jumps. The analysis is based on a series of factor models, estimated with the Fama-MacBeth methodology. The macroeconomic uncertainty factor is built using three results in the literature. First, uncertainty about important economic variables is higher before the release of economic news, and is rapidly resolved upon announcement. Second, the implied volatility of equity options increases before scheduled releases and drops afterwards. Third, this change is correlated with the implied volatility of options on macroeconomic variables, which were traded between 2002 and 2007. The sensitivity of options to macroeconomic uncertainty is estimated by studying the pattern of option pricing errors around announcement days, and the macroeconomic uncertainty factor is a portfolio that buys options with the highest estimated sensitivity. The analysis includes a large set of relevant asset- and option-pricing factors, and controls for measurement error and hedging risk. The results show that macroeconomic uncertainty is a priced risk factor for equity options, and this conclusion is robust to the effects of non-linearity of option returns with respect to some of the factors, to alternative methods of estimating option sensitivity to macroeconomic uncertainty, to calculating returns by selling at the bid and buying at the ask price, to adjustments for non-randomly missing returns, and to alternative definitions of the expected objective volatility. The contribution of the macroeconomic uncertainty factor to expected excess option returns is about 70% a year for option portfolios with a beta equal to the median factor loading. This compares to 20-25% for the market factor. Including the macroeconomic uncertainty factor reduces the risk premium on the market and on the volatility of jumps, while it makes the premium on the volatility of volatility more negative and significant. This suggests that it is important to identify the economic foundations of the reduced-form sources of uncertainty that are commonly used in option pricing.

The second chapter brings together the topics of technology diffusion and behavioral finance. I study whether investor sentiment acts as a subsidy to investment in new technologies, in particular for firms that are more likely to experience informational externalities, and if such subsidy produces positive effects for aggregate productivity growth. Early adopters likely generate informational externalities about the productivity of a new technology, and this can create a free rider problem that keeps aggregate investment below the social optimum. The resulting reduction in the flow of information can delay the adoption of a valuable technology, or the rejection of a less productive one, both of which generate a loss in productivity growth. Investor sentiment can act as a subsidy to investment by reducing funding costs, and by enticing managers to undertake investments that cater to investors' optimism about the new technology. The empirical analysis investigates how the interaction between investor sentiment and technological innovation affects 1) firm-level investment, in particular for firms that are more susceptible to informational externalities, and 2) aggregate productivity. I find that the interaction has a positive effect on investment for firms subject to informational externalities, and that it raises future aggregate productivity growth. A one standard deviation increase in technological innovation adds about 7% to the difference in the investment ratio between high/low capital intensity firms if investor sentiment is higher rather than lower, as measured by changes in sentiment at the top and bottom of the interquartile range. The effect of the same increase in innovation on productivity growth, when changes in sentiment are at the third quartile, is slightly more than 0.5%.

In Chapter 3, I solve a portfolio choice problem by estimating the distributions of expected returns on cross-sections of option prices, to study whether the information about higher moments that is implied in risk-neutral distributions improves the precision of estimated expected returns. I obtain the corresponding objective distributions by applying the empirical pricing kernel of Rosenberg and Engle (2002) and, alternatively, by shifting the distribution to the right by adding the risk premium. The second method allows to evaluate the portfolio choice implications of risk-neutral skewness, which imposes a penalty on extreme weights by increasing expected losses. The marginal distributions are joined using copula theory, and the comparison of the performance of portfolios based on the normal and t copulas provides evidence on the economic value of accounting for tail dependence. The results show that option-based models outperform the $\frac{1}{N}$ rule, especially with (limited) short-selling. Time-series models, on the other hand, produce certainty equivalents that are consistently lower than the naïve strategy's. The wedge between option-based models and the $\frac{1}{N}$ rule is greater for three- rather than two-dimensional portfolios. Transaction costs are relatively high for the strategies derived from options, and reduce, but do not eliminate, the economic significance of the results.

Chapter 1

Macroeconomic uncertainty and option returns

This chapter benefited from the suggestions of Mike Chernov, Elroy Dimson, Vito Gala, Francisco Gomes, Christopher Hennessy, Oguzhan Karakas, Stephen Schaefer and Raman Uppal.

1.1 Introduction

One of the best known empirical regularities of the Black-Scholes option pricing model is the *volatility smile*, which is the expensiveness of out-of-the-money options relative to at-the-money ones.[1] A large part of the recent option pricing literature has tried to explain this and other stylised facts, like the change in the shape of the volatility smile after the October 1987 crash, by relaxing the Black-Scholes assumptions and incorporating additional sources of uncertainty, notably stochastic volatility and jumps.[2]

The most common approach to building an option pricing model involves specifying the distribution of the underlying asset's returns and identifying the sources of uncertainty that carry a risk premium. Priced sources of uncertainty have different parameters in the objective and risk-neutral distributions. If the volatility of jumps is priced, for instance, the risk-neutral jump volatility will not be equal to the objective one. This is the reason why differences between the estimated objective and risk-neutral parameters are usually interpreted as risk premia.

During estimation, risk premia effectively act as free parameters that reconcile any discrepancies between the objective and risk-neutral distributions. One important consequence is that model misspecification can appear as a risk premium, and specification tests are of primary importance (see Broadie, Chernov and Johannes (2007)).

---

1 The Black-Scholes implied volatility is the volatility that makes the Black-Scholes price equal to the market price of an option, for given stock price, interest rate, dividends, maturity and strike price. The implied volatility curve is a plot of implied volatilities across different moneyness levels, for options with the same maturity. The Black-Scholes model implies a flat curve, but observed implied volatilities for out-of-the money options are higher, hence the name volatility *smile* or *smirk*.

2 See Bakshi, Cao and Chen (1997), Bates (2003) and Benzoni, Collin-Dufresne and Goldstein (2005).

Econometric issues aside, a reduced-form approach to option pricing also poses questions in terms of economic interpretation. The difficulty in reconciling standard preferences, option returns and option holdings probably reflects the need for models that explicitly take into account frictions in financial intermediation, like agency problems in portfolio delegation (Driessen and Maenhout (2007)). The economic mechanisms that drive the sources of uncertainty in reduced-form models are also not fully specified. The volatility of expected returns, for instance, can be generated by the interaction between the asset's return volatility and uncertainty about the value of a state variable (David and Veronesi (2002)). Option prices may also depend on the microstructure of option markets (Jameson and Wilhelm (1992), Bates (2003), Gârleanu, Pedersen and Poteshman (2007)) and on market segmentation (Pan and Poteshman (2006)). Bates (2003, p.399) clearly emphasizes the need for a sharper focus on the economic fundamentals behind the differences between the objective and risk-neutral distributions:

> To blithely attribute divergences between objective and risk-neutral probability measures to the free "risk premium" parameters within an affine model is to abdicate one's responsibilities as a financial economist.

This paper examines the contribution of macroeconomic uncertainty to equity option returns. In particular, I study the effect that time-varying uncertainty about the *current* value of macroeconomic variables has on option returns.[3] The focus is not on the *time-series* relation, that is if an increase in macroeconomic uncertainty translates into higher option prices and positive returns. Instead, I investigate whether the sensitivity of option returns to one particular source of uncertainty - macroeco-

[3] Defined as holding returns: $r_{t-1,t} = \frac{opt.price_t - opt.price_{t-1}}{opt.price_{t-1}}$

nomic - explains the *cross-section* of equity option returns, which would imply that macroeconomic uncertainty is a priced risk factor.

Let us consider two examples to clarify the meaning of "macroeconomic uncertainty". The release of the official figures for August 1999's employment situation in the U.S. was scheduled on September 3, 1999. Referring to the pre-announcement day, the Financial Times wrote:[4] *"Interest rate worries returned to ambush global equity markets, as investors nervously anticipated today's US employment numbers [...]. Recent economic data have reawakened fears that the US Federal Reserve will have to move to raise interest rates again shortly"*. The following day, the same newspaper reported: *"The smaller-than-expected increase in US job creation last month [...] and the lower-than-expected increase in hourly earnings [...] was seen as reducing the likelihood of a rise in US rates [...]"*.[5]

More recently, in September 2008, the Financial Times linked heavy stock market losses to the fact that *"[...] labour market data heightened concerns that today's crucial non-farm payrolls report might be weaker than expected [...]"*.[6]

Clearly, in both cases investors were uncertain about the current state of the employment situation, and its implications for future growth prospects. Even if in one case the economy was in a phase of robust expansion, and in the other it was on the brink of a recession, uncertainty about macroeconomic fundamentals had an equally significant impact on asset prices.

Returning to the empirical analysis presented in this paper, I proxy for unobservable macroeconomic uncertainty with a factor that sorts options on the basis of their

[4] "Rate fears cast shadow ahead of jobs data", September 3, 1999.
[5] "Shares jump on bid news and US jobs report", September 4, 1999.
[6] "Worries over financial sector weigh on equities", September, 5, 2008.

pricing errors, measured on days immediately before scheduled *macroeconomic announcements*, and normalized with respect to pricing errors on non pre-announcement days (see Section 1.3.3). I find that this factor explains the cross-section of option returns, and results are robust to a large set of additional relevant factors, alternative ways of measuring macroeconomic uncertainty, and I account for potential biases due to non-randomly missing returns, measurement error, the non-linearity of option returns, and non-linearity in the relation among factors. The macroeconomic uncertainty factor is built in three steps. First, I identify days when macroeconomic uncertainty is higher. Second, I find options that have *unusual* pricing errors on such days - they likely are more sensitive to macroeconomic uncertainty. Finally, I form a factor mimicking portfolio that buys options with the highest sensitivity to macroeconomic uncertainty.

These steps deserve further discussion. The existing literature suggests that macroeconomic uncertainty is likely higher before a scheduled macroeconomic announcement. Several authors have shown that asset prices react quickly to the release of economic news (see Section 1.2.2), which is consistent with a rapid resolution of uncertainty. Beber and Brandt (2008) report that the implied volatility of options on macroeconomic variables,[7] which is likely very correlated with macroeconomic uncertainty, explains the reduction of stock implied volatilities after scheduled releases of economic news. More precisely, the higher the implied volatility of options on macroeconomic variables, the more substantial the drop in equity options' implied volatilities. While it is possible that uncertainty is partially reduced as the announcement approaches, because information may be leaked, the results in the literature

[7] Economic Derivatives, that were marketed between 2002 and 2007.

suggest that a large component is still resolved upon announcement. In addition, one of the robustness checks in the empirical analysis involves choosing different combinations of days to measure uncertainty.

As for the second step, the proxy for macroeconomic uncertainty is based on Black-Scholes pricing errors, to focus on the price component that is not explained by fundamental variables like the underlying stock price, interest rates and time to maturity. The choice of using Black-Scholes pricing errors, rather than those from a more flexible model, is to make sure that the effect of macroeconomic uncertainty is not unduly captured by one of the additional moments and risk premia. To avoid, however, that the macroeconomic uncertainty factor proxies for other variables, the empirical analysis includes a large number of relevant factors, like stochastic volatility, jumps and higher moments.

Third, the factor mimicking portfolio only buys options with the highest sensitivity to macroeconomic uncertainty, rather than also selling those with a low one, to avoid picking up the effect of inventory management by market makers. More specifically, market makers may be reluctant to fully accommodate a significant increase in the demand of options that are exposed to a specific risk factor, when such risk factor increases. Given that market makers are typically long equity options (Gârleanu, Pedersen, Poteshman (2009)), they may tend to reduce the price of an option to discourage investors from selling, which in turn may dampen the increase in price due to higher uncertainty. In fact, I provide evidence that the probability of a negative return increases when changes in trading volume are associated to a relatively large number of macroeconomic announcements. This implies that the measured sensitivity may be lower than the true one, and that the factor I use in the empirical analysis is a conservative measure of macroeconomic uncertainty.

In a recent paper, Anderson, Ghysels and Juergens (2007) find that aggregate uncertainty, proxied by the dispersion of forecasts from the Survey of Professional Forecasters, helps to explain market returns and the cross-section of expected *stock* returns. While the theoretical framework is similar, there is one important difference between their work and mine. I look for the effect of macroeconomic uncertainty on option returns *beyond* the effect it has on stocks, because this provides useful information about *why* options are non-redundant securities. If investors were able to attain payoffs across states of high/low macroeconomic uncertainty by *either* buying options *or* replicating options with stocks, then the effect of macroeconomic uncertainty on options would be routed *only* through the underlying stock, and macroeconomic uncertainty would not explain the cross-section of option returns after controlling for the exposure of the underlying stocks.

Before discussing the empirical implementation and the results in detail, I review the relevant literature in Section 1.2. Section 1.3 presents the empirical methodology, describes the data and discusses test assets and factors. Section 1.4 analyses the results, Section 1.5 focuses on the robustness checks and Section 1.6 concludes.

1.2 Related literature

1.2.1 Uncertainty about state variables and asset pricing

The macroeconomic uncertainty factor is built by analyzing option prices around announcements about the current, unobservable, value of the Consumer Price Index, the Employment Situation, Real Earnings, the Producer Price Index and Productivity and Costs. These variables are usually considered important indicators of the of the state of the economy. The often-cited[8] "misery index", for instance, is the sum of the unemployment and inflation rates, and variables related to income and productivity are also used as proxies for the business-cycle (Stock and Watson (1989)). Macroeconomic uncertainty can then be interpreted as uncertainty about the *current* value of state variables, and this subsection reviews the literature on the asset pricing implications of uncertainty about unobservable state variables.

There are several ways through which the unobservability of a state variable makes its way into asset prices, starting with heterogenous beliefs. Buraschi and Jiltsov (2006), for example, build an option pricing model where uncertainty and heterogeneous beliefs affect option trading volume and generate an asymmetric volatility smile. They test several implications of their model with a Difference in Beliefs index, which is based on survey data, and find that it helps to explain the volatility smile, future realized volatility, and violations of the Black-Scholes bounds on option deltas ($\Delta \in [0, 1]$ for calls and $\Delta \in [-1, 0]$).

Learning about the process of the unobservable state variable also affects prices and volatilities, like in David and Veronesi (2002). Investors try to learn the current value of the dividend drift, which follows a two-state regime-switching model with time-varying uncertainty about the true value, and the learning process itself generates stochastic volatility and stochastic correlation between returns and volatility. As a result, David and Veronesi (2002) can generate asymmetric smiles, whose slope is sometimes positive. Guidolin and Timmermann (2003) focus on a model where

[8] "The return of the misery index", New York Times, September 12, 2008.

dividend news evolve according to a binomial lattice with unobservable probabilities. Investors use Bayes' rule to update their estimates, and the resulting dynamics generate skewed volatility smiles and a non-constant term structure of implied volatilities. Veronesi (2000) studies a Lucas economy where investors don't observe the output growth rate, but receive a noisy signal. He finds that the risk premium can be *lower* for higher uncertainty levels. The reason is that dividend realizations affect both expectations and the hedging demand for stocks, especially if the signal is very noisy. The correlation between consumption and returns is smaller for a less precise signal, in which case the risk premium is lower. Veronesi (2000) also finds that the correlation between expected returns and volatility depends on the level of uncertainty, which suggests a reason why the empirical evidence on the time-series relation between the conditional mean and volatility of stock returns is mixed (e.g., Whitelaw (1994)).

It is also possible that *Knightian* uncertainty and model mis-specification enter directly into investors' preferences. In this case the price of risk is generally higher, equity holdings are smaller and precautionary savings increase (Hansen, Sargent and Tallarini (1999), Cagetti, Hansen, Sargent and Williams (2002), Maenhout (2004)). Investors may prefer portfolio allocations that are skewed towards assets whose return distribution is subject to relatively small ambiguity (Uppal and Wang (2003)). Liu, Pan and Wang (2005) derive option pricing implications of uncertainty aversion toward rare events, for which model specification and estimation are intrinsically difficult. They are able to reproduce the volatility smirk because options, and especially out-of-the-money puts, are very sensitive to rare events.

1.2.2 Economic news and asset prices

The contribution of my work is to examine the relation between macroeconomic uncertainty and the cross-section of option returns, and it is clearly linked to the literature on the effect that news about economic fundamentals have on asset prices. The overall conclusion is that stocks, bonds and options react quickly to macroeconomic news. McQueen and Roley (1993) focus on stock prices, for which the effect varies across the business-cycle, with "good" news increasing prices only when the economy is weak. Balduzzi, Elton and Green (2001) examine the interdealer market for Treasury bills and bonds, finding evidence of strong and rapid price effects. Andersen, Bollerslev, Diebold and Vega (2003)) study the foreign exchange market, showing that announcement surprises generate conditional mean jumps. The work of Ederington and Lee (1996) and Beber and Brandt (2006,2008) is especially relevant for my analysis, because they focus on the relation between macroeconomic news and *option prices*. Ederington and Lee (1996) examine the markets for Treasury bonds, Eurodollar and Dollar/Deutsche Mark options. The effect of macroeconomic news depends on whether the release is scheduled or not. In the first case implied volatilities drop, but rise for unscheduled releases. Beber and Brandt (2006) study bond options, too, finding that implied volatilities always decrease after announcements, irrespective of whether the content is unexpectedly positive or negative for the economy. The behavior of higher moments, however, depends on the unexpected information brought by the announcement. Beber and Brandt (2008) derive a measure of macroeconomic uncertainty by computing the implied volatility of options on macroeconomic variables, and directly test whether the drop in implied volatilities after scheduled releases is related to macroeconomic uncertainty. The results show that higher uncertainty does lead to sharper declines.

1.2.3 Microstructure and demand effects in option pricing

The empirical analysis discussed in Section 1.3 includes a series of controls to account for demand and microstructure effects on option prices. Option market-makers absorb the net demand of investors and, because imperfectly hedged imbalances expose them to the risk of losses, prices may reflect the strength of the underlying demand, especially for options that are inherently more difficult to hedge, like those with high gamma or vega.[9] The empirical literature supports this hypothesis: Figlewski and Webb (1993) find that a large short-interest in the underlying stock increases the implied volatility of puts, but not calls, which is consistent a demand effect. Jameson and Wilhelm (1992) report that option spreads are increasing in gamma. Net demand also helps to explain excess implied volatility, especially for options that are more difficult to hedge (Bollen and Whaley (2004), Gârleanu, Pedersen and Poteshman (2009)).

1.3 Empirical implementation

The analysis of the contribution of macroeconomic uncertainty to the cross-section of option returns is based on a factor model, which is estimated with the Fama-MacBeth methodology. This setting imposes little structure on the data, and reduces the potential for model misspecification problems, which can attribute risk premia to factors that are actually not priced (Broadie, Chernov and Johannes (2007)). Test assets are

[9] Gamma is the second derivative of the option price with respect to the underlying price. Vega is the first derivative of the option price with respect to volatility.

portfolios of options (as defined in Section 1.3.2), and weekly holding returns[10] are defined as:

$$r_t^i = \frac{\mathrm{OP}_t^i - \mathrm{OP}_{t-1}^i}{\mathrm{OP}_{t-1}^i} \tag{1.1}$$

where OP_t^i is the price of option i at time t. The return on a particular option portfolio, A, is the equally weighted return of all the options that belong to it at time $t-1$:

$$r_t^{\mathrm{A}} = \frac{1}{\#\mathrm{A}_{t-1}} \sum_{i \in \mathrm{A}_{t-1}} r_t^i \tag{1.2}$$

The first step of the Fama-MacBeth method consists in estimating the sensitivity of portfolio returns to a set of factors, using time-series regressions (first-stage regressions):

$$r_t^k - r_t^f = \alpha_{1,k} + \sum_{i=1}^{n} \beta_{i,k} f_{i,t} + \varepsilon_{k,t}, \forall k \tag{1.3}$$

where r_t^k is the weekly return on asset k and $f_{i,t}$ is one of the n factors.

The cross-sectional (second-stage) regressions determine the extent to which differences in the estimated factor sensitivities explain asset returns:

$$r_t^k - r_t^f = \alpha_{2,t} + \sum_{j=n+1}^{n+m} \lambda_{j,t} \mathrm{Cont}_{j,t} + \sum_{i=1}^{n} \lambda_{i,t} \hat{\beta}_{i,k} + \epsilon_k, \forall t \tag{1.4}$$

where $\mathrm{Cont}_{j,t}$ are controls for measurement error and hedging risk, to be defined later in this paragraph.

The risk premium on factor f is estimated as the time-series average of the coefficients

[10] Measured on Tuesdays, following Coval and Shumway (2001).

from the T cross-sectional regressions:

$$\hat{\lambda}_f = \frac{1}{T}\sum_{t=1}^{T}\hat{\lambda}_{f,t} \tag{1.5}$$

I use holding returns, rather than hedged returns, for two reasons. First, the computation of hedged returns requires the choice of an option pricing model to form hedging portfolios, which creates a potential source of misspecification. Second, hedging portfolios need to be rebalanced frequently in order to maintain replicating accuracy, typically every day, and with daily rebalancing measurement error and non-synchroneity have a higher impact on observed returns. The main reason why option prices are measured with noise is that Optionmetrics, the source for option data, only reports the daily bid-ask spread, and the true option price may not the bid-ask midpoint, as it is commonly assumed. Non-synchroneity arises because the database does not include information on the exact time of the trade, which may have taken place at any time during the day. The consequence is that a *daily* return may actually be a "$evening_{t-1}$ to $morning_t$" or a "$morning_{t-1}$ to $evening_t$" return. By taking weekly returns the noise generated by measurement error and non-synchroneity is a smaller proportion of the overall return variability. I also address the issue of measurement error by including controls in second stage regressions: the relative option bid-ask spread, given that measurement error is likely proportional to the spread, and the underlying stock's relative bid-ask spread, because measurement error in the stock price is by definition fed into the option's. Two additional controls are the option's log-vega and log-gamma, as proxies for the hedging risk borne by option market-makers, and the underlying stock's exposure to macroeconomic uncertainty (its beta from time-series regressions). All controls are median $t-1$ portfolio values.

Using a factor model to study option returns gives an advantage in terms of robustness, but it also requires to carefully account for one characteristic of option returns, namely their non-linearity in the returns of the underlying security and in factors that affect the underlying. As Broadie, Chernov and Johannes (2008) note, non-linearity may bias the coefficients of linear models. I take non-linearity into account in three ways. First, I use fractional polynomial regressions to study the relation between returns and factors, and among factors, and then I include non-linear terms in the Fama-MacBeth regressions when necessary. Second, I estimate cross-sectional regressions with Weighted Least Squares to account for the heteroskedasticity (returns on out-of-the-money options are more volatile than returns on in-the-money ones). Third, the analysis is based on option portfolios, rather than individual options. One of the characteristics used to sort options into portfolios is moneyness, which means that the sensitivity of option returns to returns in the underlying tends to be quite homogeneous, minimizing the impact of non-linearity (see Section 1.3.2 for more details).

Finally, a simulation confirms that, when accounting properly for non-linearity by including squared terms and moneyness dummies, the Fama-MacBeth procedure accurately estimates factor risk premia.

1.3.1 Data description

Daily option bid-ask prices, together with additional variables like option volume, stock prices and interest rates, are from OptionMetrics, which covers all U.S.

exchange-traded index- and equity-options. The Fama-French, momentum and Pastor-Stambaugh liquidity factors, and stock bid-ask spreads are from CRSP. Intradaily returns on the Dow Jones Industrial are from Global Financial Data, while the VIX series is from the Chicago Board of Options Exchange. VIX is a daily index of implied volatility, based on S&P500 option prices, and is often used as a proxy for expected market volatility. Days of scheduled releases for the Consumer price index, the Employment situation, Real earnings, the Producer price index and Productivity and costs are from the Bureau of Labor Statistics. The sample includes 11.5 years, from January 1996 to June 2007.

As is customary with option data, I apply a series of filters to eliminate illiquid prices and recording errors. Table 1.1 shows the average number of monthly observations, by year, before and after the filters. The database grows substantially between 1996 and 2007, with a temporary reduction after 2000, and it includes almost 390 million entries. OptionMetrics provides an option identifier (*optionid*) which should be uniquely assigned to a single option contract. This is not always the case, so I define a new identifier that merges *optionid* and the suffix of the option symbol. This reduces but does not eliminate identifiers with more than one entry on a given date, and I drop all the observations that can not be uniquely identified. The requirement of non-zero trading volume eliminates many observations, but it is important to have option prices that reflect information from a recent transaction, which is not the case without trades on the day. I also drop observations where either the ask or the bid price is zero, and those with a bid greater than the ask. I only keep options on common stocks and with standard settlement (e.g., investors are sometimes able to deliver securities other than the underlying). All equity options have American exercise, but

I delete those with a missing exercise style flag because additional data fields may be incorrectly recorded. The remaining sample size is about 16% of the total, with more than 60 million observations.

1.3.2 Option portfolios

In this section I explain how option portfolios are formed, then present the summary statistics of their holding returns, and finally discuss the effect of non-linearity on the coefficients of time-series OLS regressions.

Definition of option portfolios

While many option pricing studies have focused on index options, mainly because of data availability and market liquidity, I have decided to use equity options. The reason is that stock characteristics are useful to create dispersion in options' sensitivities to macroeconomic uncertainty. In addition, returns on equity options do not embed a substantial compensation for jump risk, which minimizes the possibility that the macroeconomic uncertainty risk premium proxies for jump risk premium. Size and book-to-market are a common way of sorting stocks in asset pricing studies, because they are well-known sources of empirical regularities. Such sorting may be helpful in my case, too, because Barinov (2007) shows that the value effect is related to the real option nature of growth stocks, which makes them particularly sensitive to uncertainty. OptionMetrics, however, is not representative of CRSP, especially in terms of firm size, so I sort options on the basis of the underlying's sensitivity to the three Fama-French factors, given that Hml and Smb are strictly related to the book-to-market and size effects. Sensitivities are computed by regressing daily stock excess returns on the Fama-French factors for each year from 1995 to 2006, and stocks are then sorted into beta tertiles for each factor (Pmkt,Psmb,Phml=1,2,3 - 3 being the highest-sensitivity tertile). The sorting based on year t's betas applies to year $t+1$.

Options are also sorted into portfolios on the basis of the following contract characteristics: moneyness, maturity and put/call type. Moneyness should be especially useful to create dispersion in the sensitivity to macroeconomic uncertainty, because several authors have linked the effect of uncertainty about state variables to the volatility smile (Buraschi and Jiltsov (2006), David and Veronesi (2002) and Liu, Pan and Wang (2005)). In particular, letting K be the strike price and S the stock price, I sort options into five moneyness categories:

$$\text{MonDummy}_t = \begin{cases} \text{Put} & \text{Call} & \\ 1 & 5 & -0.200 < ln(\text{K}/\text{S}_{t-1}) \leq -0.100 \\ 2 & 4 & -0.100 < ln(\text{K}/\text{S}_{t-1}) \leq -0.025 \\ 3 & 3 & -0.025 < ln(\text{K}/\text{S}_{t-1}) \leq 0.025 \\ 4 & 2 & 0.025 < ln(\text{K}/\text{S}_{t-1}) \leq 0.100 \\ 5 & 1 & 0.100 < ln(\text{K}/\text{S}_{t-1}) \leq 0.200 \end{cases}$$

In addition, options with time to maturity between 15 and 90 calendar days at time $t-1$ are assigned to MatDummy=1, while those with 90 to 360 days to maturity belong to MatDummy=2. Options that do not fall into a moneyness and maturity category are discarded.[11] Excluding deep out-of-the-money/in-the-money and very long/short maturity options may create concerns of sample selection. Generally speaking, it is quite common to apply criteria aimed at reducing econometric problems due to outliers and at eliminating thinly-traded options. Studies that use tick data sometimes focus on options that trade in a specific minute of the day (Constantinides, Jackwerth and

[11] I also drop options if the bid-ask spread is smaller than the tick size (see Goyal and Saretto (2007)), if the underlying pays dividends during the week, and if the stock price is below 1 dollar. I censor the sample of individual option returns at 0.5% and the sample of option portfolios returns at 1% on both tails. I have tried several other symmetric and non-symmetric cutoffs, without altering the conclusions.

Perrakis (2009)). The simulation results I present in Section 1.5.6 are based on the sample selection criteria described above, and show that my variables of interest are accurately estimated.

To summarize, individual equity options are sorted into portfolios using six characteristics: three related to the underlying (sensitivity to the market, Smb and Hml factors) and three related to the option contract (maturity, moneyness, put/call type), and the sorting is based on the value of the characteristic in year $t-1$ (for the Fama-French sensitivities) or in week $t-1$ (for moneyness and maturity).

Return statistics

Table 1.2 shows summary statistics for option portfolios' weekly holding returns, averaged across different portfolio characteristics. Sorting by moneyness produces the most evident pattern, with out-of-the-money options earning more than 7% a week, and in-the-money ones losing almost 10%. Short maturity options and puts also lose about 1.2% and 2.5%, while long maturity options earn 0.79%. Sorting across stock characteristics does not produce a clear-cut pattern. Options on stocks with low sensitivity to the market and Hml lose slightly more than 0.5% a week, and those with medium sensitivity to Hml gain about 0.80%. Interestingly, the mean return for portfolios of options with average factor-sensitivity is always greater than for portfolios with a high/low sensitivity, although the statistical significance is weak, with the exception of Hml, as noted above. The distributions of returns are mainly positively skewed, and their kurtosis is only slightly above 3, although it is about 7 for portfolios on stocks with high sensitivity to the market and to Hml.

As discussed in Section 1.3, option returns are non-linear in returns on the underlying and in factors that affect returns on the underlying, like the market. To gauge the effect of such non-linearity on the coefficients of a time-series OLS regression, Figures 1 and 2 plot returns on a call and a put option portfolio against market returns. Both portfolios have high sensitivity to the market, low sensitivity to Hml and Smb, short maturity, and include at-the-money options, for which non-linearity is more pronounced (at-the-money options have a large gamma).

Figure 1.1 shows returns on the portfolio of call options, together with two regression lines. One is calculated from standard OLS, while the other is computed by weighing down observations that generate non-linearity. More specifically, observations corresponding to the bottom 25% of market returns receive a weight of 0.25 in the Weighted Least Squares regression, rather than 1. The reason for doing so is apparent in the leftmost portion of Figure 1. When call options are at-the-money, high market returns (routed through returns on the underlying stock) translate into higher option returns. On the other hand, low market returns do not generate correspondingly low option returns, because the value of a call option rapidly goes to zero. Indeed, Figure 1 shows that the relation between option returns and market returns is almost a flat line for low market returns, while it is positively sloped for high returns. By weighing down the bottom 25% of market returns, I can evaluate the impact of non-linearity on the estimated market beta. The plot suggests there is a slight difference in the estimated slopes, with the market beta being larger when non-linearity is accounted for. More in detail, the standard OLS regression gives a market beta equal to 9.52 and an intercept of -0.0036, while Weighted Least Squares estimates the slope at 9.97 and the intercept at -0.0077. The percentage difference between the betas is less than 5.0%.

Figure 1.2 repeats the analysis described above for the put option portfolio. In this case non-linearity is generated by high market returns, because the price of a put option quickly declines toward zero, so the weight of 0.25 applies to observations with market returns in the last quartile. Standard OLS estimates are equal to -10.15 for the market beta, and -0.031 for the intercept. In the case of WLS, beta is equal to -10.79 and the intercept to -0.037, giving a percentage difference of less than 6.0% between the betas.

1.3.3 The macroeconomic uncertainty factor

The proxy for unobservable macroeconomic uncertainty is built by exploiting the pattern of option prices around days of macroeconomic announcements.

The starting point is that, before scheduled releases, uncertainty about the value of macroeconomic variables is relatively high. There is concurring evidence about this, as discussed in Section 1.2.2. Stock, bond and option prices react quickly to macroeconomic announcements, which suggests a resolution of uncertainty immediately after the release of economic news. Furthermore, Beber and Brandt (2008) and Savor and Wilson (2008) analyse, respectively, the behavior of equity implied volatilities and interest rates around macroeconomic announcements, and infer that investors increase their hedging activity before scheduled announcements, which is also consistent with higher uncertainty.

Secondly, options that are more sensitive to macroeconomic uncertainty should be relatively more expensive before an economic release is made, consistently with pricing the higher macroeconomic uncertainty. Beber and Brandt (2008) provide evidence of this price effect, studying the relation between the implied volatility of options on macroeconomic variables, which is likely correlated with unobservable macroeconomic uncertainty, and the implied volatility of equity options around scheduled releases of economic news. They find that higher implied volatilities from macroeconomic options are associated with larger reductions in equity options' implied volatilities after the release of news.

Lastly, I identify options that have *unusual* prices on days before macroeconomic announcements, and I form a portfolio that goes long such options. To define which options have *unusual* prices, I first of all normalize option prices by taking out the effect of fundamental variables like the underlying price, interest rates, time to maturity and strike price. In other words, I compute pricing errors to isolate the component that is not explained by these variables. I do so by using the Black-Scholes formulas, because more complex models may proxy for macroeconomic uncertainty with one of the additional moments and risk premia. Of course, there is the possibility that the macroeconomic uncertainty factor proxies for variables omitted in the Black-Scholes model, and this is the reason why the Fama-MacBeth regressions include a large set of relevant asset- and option-pricing factors, like stochastic volatility, jumps and volatility of jumps. After calculating the Black-Scholes residuals, I compute the difference between pricing errors on the day preceding a macroeconomic announcement and the median pricing errors on all other days. Options that are in the highest quintile of this difference are at the basis of the macroeconomic uncertainty factor (EU).

Detailed definition of the macroeconomic uncertainty factor

First, I compute pricing errors (PE) for all options, on each day:

$$\text{PE}_{i,t} = \frac{\text{SBS}_{i,t} - \text{BS}_{i,t}}{\text{OP}_{i,t}} \tag{1.6}$$

where $\text{BS}_{i,t}$ is the option's Black-Scholes price computed with the expected objective volatility, $\text{SBS}_{i,t}$ is the Black-Scholes price computed with OptionMetrics' implied volatility[12] and $\text{OP}_{i,t}$ is the option's bid-ask midpoint. I then calculate the daily PE for each option portfolio, $\text{PE}_{p,t}$, as the median $\text{PE}_{i,t}$ of all the options belonging to p. Next, I collect scheduled announcement dates for the Consumer price index, the Employment situation, Real earnings, Producer price index and Productivity and costs. For each option portfolio, in each quarter, I define the Excess PE (EPE) as the difference between the median PE on pre-news days and the median PE on normal days. More precisely, the EPE for portfolio p in quarter q is:

$$\text{EPE}_{p,q} = \bar{\text{PE}}^{\text{pre-news}}_{p,q} - \bar{\text{PE}}^{\text{normal}}_{p,q} \tag{1.7}$$

where $\bar{\text{PE}}^{\text{pre-news}}_{p,q}$ is the median PE on days immediately before (i.e., $t-1$) scheduled news releases, and $\bar{\text{PE}}^{\text{normal}}_{p,q}$ is the median PE on non pre-announcement days.

I then sort option portfolios into quintiles on the basis of their EPE,[13] and the EU factor is the median return of the high EPE quintile. The reason why I do not build the factor mimicking portfolio as a long/short position is that inventory management by option market makers may partially offset the price effect of macroeconomic uncer-

[12] This implied volatility is calculated using a binomial tree based on the Cox-Ross-Rubinstein model. See the OptionMetrics Data Reference Manual for more details.

[13] The sorting in quarter q applies to quarter $q+1$.

tainty. Specifically, when a risk factor increases, market makers may not wish to fully accommodate the incremental demand for options that expose them to such factor, because imperfect hedging increases the risk of unexpected losses. On average, market makers have positive holdings of equity options (Gârleanu, Pedersen, Poteshman (2009)), so they may reduce the price of an option to discourage investors from selling, which in turn may dampen the increase in price due to higher uncertainty. The implication is that the lowest EPE quintile may contain options whose true sensitivity to macroeconomic uncertainty would place them in higher quintiles, and this would make EU a noisier proxy for macroeconomic uncertainty. Of course, this issue holds true even when the EU factor is defined as the return on the highest EPE quintile, but it works against finding the results that I present in Section 1.4. Consistently with an inventory management effect, Table 1.3 shows that the probability of a negative return increases when changes in trading volume are associated with a relatively high number of macroeconomic announcements. More in detail, the table shows odds ratios from probit models in which the dependent variable is a dummy equal to 1 if the residual of a regression of option returns on the following factors is negative: Mkt, Mkt^2, Smb, Hml, Hml^2, Vix, Vix^2, Vix_v, Skew, Put, Put_v and EU (see Section 1.3.4 for definitions). The independent variables are the weekly percentage change in the average trading volume of all options, the number of macroeconomic announcements in the week, and their interaction. The odds ratio of interest is for the interaction, and to be consistent with an inventory management effect it must be greater than one. The reason is that market makers, on average, have a positive holding of equity options, so a string of events that generate uncertainty (announcements) coupled with an increase in demand from investors (proxied for by the change in trading volume) should increase the probability of price reductions (i.e., negative returns) to keep investors from selling. Interestingly, market makers have larger holdings of calls (Gârleanu, Pedersen, Poteshman (2009)), and the odds ratio is greater for calls than for puts, while the odds ratio is also large and significant for at-the-money options, which are relatively difficult to hedge (they have a high gamma).

Alternative definitions of the macroeconomic uncertainty factor

The definition of the macroeconomic uncertainty factor deserves further discussion. First of all, there is the question of how to determine the volatility used to compute the Black-Scholes price. The main results are based on a trailing 30 days volatility, while the robustness checks section shows results based on alternative definitions: 1) trailing 30 days volatility, similar to the base case but the current day is excluded to make sure that volatility does not contain more recent information than the option price (which can be the case if the option stops trading before the stock); 2) trailing 30 days volatility, where zero stock returns are dropped, to avoid that pricing errors reflect trading liquidity in the underlying stock;[14] 3) for options belonging to MatDummy=1 (=2), the volatility realized in the two (seven) months *following* the day when the pricing error is calculated, to make sure that the difference in pricing errors across maturity does not reflect expectations about the term structure of volatility. Figure 1.3 shows the term structure of implied volatility for call and put options belonging to different moneyness categories. The implied volatility of equity options is not very different from the corresponding objective volatility, unlike for index options, so the term structure of the implied volatility should be a good proxy for the term structure

[14] The calculation includes the latest 30 non-zero returns from the previous 50 days.

of the expected objective volatility. The term structure is calculated by normalizing the implied volatility of an option with the implied volatility of the longest maturity option on the same underlying, with the same strike price and of the same put/call type, as long the longest maturity is between 200 and 250 days, and then taking the median for each maturity across moneyness categories and put/call type. The figure shows that the volatility term structure flattens substantially after about one month but, for in- and out-of-the-money options, the short term implied volatility is noticeably higher.

Second, the assumption held so far is that macroeconomic uncertainty is relatively higher on the day before a macroeconomic announcement is made. In addition, the pricing error on a pre-announcement day is normalized by subtracting the median on all other days. It is possible, however, that uncertainty is partially resolved as the day of the release approaches, and that the normalizing pricing errors incorporate a significant amount of macroeconomic uncertainty, given that there is more than one announcement per quarter. To address these points the EU factor is also built by computing pricing errors 1) two days before scheduled announcement and normalizing them with the median pricing errors on non pre-announcement days; 2) one day before announcements and normalizing them with the pricing errors on the announcement day. In additional robustness checks, meant to isolate the uncertainty generated by individual announcements, the factor is built by only considering announcements that are three and seven days apart.

Properties of the macroeconomic uncertainty factor

The composition of the high EPE quintile, in terms of option portfolio characteristics, is examined in Figure 1.4. For each characteristic, the table shows the proportion of portfolios that belong to the high EPE quintile, across the values of the characteristic. The clearest pattern is across moneyness, with out-of-the-money options belonging more often to the high EPE quintile, and the proportion declining monotonically as moneyness increases. There is no difference across maturity, while put options are only slightly more likely than calls to be in the high EPE quintile. The sensitivities of the underlying stocks to the Fama-French factors show an interesting pattern, with options on stocks with medium sensitivities belonging to the top quintile.

To better understand its properties, it can be useful to compare the EU factor to other macroeconomic variables and to alternative measures of aggregate uncertainty. Figure 1.5 shows the quarterly average of the residuals from regressing EU on a set of other factors, namely Mkt, Smb, Hml, Vix, Vix_v, Skew and Put (see Section 1.3.4), plotted together with the end-of-quarter level of the S&P 500, of the Industrial capacity utilization index, and with an index of dispersion of analyst forecasts. The figure shows residuals, rather than the EU factor itself, because the way it is computed means that EU contains the effect of variables omitted by the Black-Scholes model, in addition to macroeconomic uncertainty. The factor spikes between 2001 and 2002, and is relatively high in 1997 and early 2000. It is interesting to note that in early 2000 the robust economic growth of the late 1990s was about to come to an end, while in 2002 the economy started to pick up again (the NBER identifies the through in the last quarter of 2001). This suggests that macroeconomic uncertainty is higher when the growth trend of the economy turns from positive to negative, but also when it turns from negative to positive. The first panel of Figure 1.5 shows the EU residuals together with the level of the S&P 500 index, and it is apparent that, with the exception of the very beginning of the sample, macroeconomic uncertainty does rise when the S&P 500 is entering phases of growth or decline, and stabilizes when the trend is set. The second panel of the figure plots the EU residuals together with the level of industrial capacity utilization. The relation between the two series once again suggests that macroeconomic uncertainty increases at the inflection points of economic activity. This is also clear in the first part of the sample, where the relation between EU and the S&P 500 is less defined. The final panel shows the macroeconomic uncertainty residuals and a measure of analyst forecast dispersion. The FD index is built by normalizing a firm's one-year-ahead forecasts of earnings-per-share by the monthly average of the absolute value of the firm's forecasts, and then computing the quarterly volatility of all standardized forecasts. The two measures of uncertainty peak at the same time in 2002, and share a common pattern in the first part of the sample. The correlation over the 1996-2007 period is 0.21, which is relatively low, but it increase to 0.31, 0.43 and 0.73 when EU is greater than the 50^{th}, 60^{th} and 75^{th} percentile. This suggests that the FD index normally measures a different type of uncertainty, although it reflects macroeconomic uncertainty to a greater extent when this is higher.

1.3.4 Additional factors

As mentioned above, the EU factor is based on Black-Scholes pricing errors and, being residuals, pricing errors contain the effect of omitted factors. To control for this, the regressions include 13 additional factors that are likely to be relevant when explaining option returns (Bates (2003), Broadie, Chernov and Johannes (2007)).

Asset pricing factors. *Mkt, Smb, Hml, Umd and Liquidity.*

Weekly Fama-French factors are computed by compounding daily returns. The Pastor-Stambaugh liquidity factor is only available at monthly frequency, so I construct a factor mimicking portfolio. I sort CRSP stocks on the basis of their liquidity beta, computed by regressing monthly excess returns on the market, Smb, Hml, Umd and on innovations to the Pastor-Stambaugh liquidity factor (Pastor and Stambaugh (2003)). The regressions cover the 1996-2004 period. Stocks are sorted into quintiles on the basis of their liquidity beta, and the factor is the difference between the equally weighted returns on stocks in the fifth and in the first quintile.

Option pricing factors. *Volatility, Volatility of volatility, Skewness of volatility, S&P 500 put returns, Volatility of S&P 500 put returns, Skewness of S&P 500 put returns, Market skewness, Market kurtosis.*

The *Volatility* factor (Vix) is the series of weekly changes in the VIX index. *Volatility of volatility* (Vix_v) and *Skewness of volatility* (Vix_s) are changes in the weekly volatility and skewness of daily VIX changes. *S&P 500 put returns* (Put) is the weekly average of daily returns on S&P 500 put options with $-0.2<\ln(K/S)<-0.1$. *Volatility and Skewness of S&P 500 put returns* (Put_v and Put_s) are changes in the weekly volatility and skewness of daily returns on out-of-the-money S&P 500 options, as defined above. *Market skewness* (Skew) and *Market kurtosis* (Kurt) are weekly changes in the skewness and kurtosis of intradaily Dow Jones Industrial returns. Driessen, Maenhout and Vilkov (2009) find evidence of a priced correlation risk factor, which contributes to explain the expensiveness of index options with respect to individual ones. It is not included in the list of factors because I focus on equity options, whose returns are "likely to be much less dependent on correlation shocks".[15] In addition, the Put factor would account for the correlation movements that are usually associated with periods of market stress (Roll (1988)).

Table 1.4 shows time series summary statistics for the 14 factors used. Only three have statistically significant average returns, namely Hml, Umd and Put. Kurtosis is mostly slightly above three, with the exception of Smb, Hml, Umd, Vix and Vix_v, for which it lies in the range 6-10. Table 1.5 reports factor correlations. As expected, the market factor is highly correlated with Vix and Put (-76% and -56%), but also moderately correlated (37%) with the macroeconomic uncertainty factor.

1.4 Results

Factor risk premia ($\hat{\lambda}_f$) are estimated with the Fama-MacBeth methodology, using weekly excess returns on the option portfolios defined in Section 1.3.2. The sample starts in the second quarter of 1996 - the first is used to sort options into Excess Pricing Error (EPE) quintiles - and ends in June 2007. All t-stats are Newey-West- (4 lags) and Shanken-adjusted. Unless specified otherwise, the EU factor is built using trailing 30 days volatility as a proxy for expected objective volatility.

Table 1.6 reports the first set of estimates, without the controls for measurement error and hedging risk. The base-case specification (1) includes nine factors: Mkt, Smb, Hml, Vix, Vix_v, Put, Put_v and EU. Five of them are statistically significant: the market, Hml, Vix, Put and EU. The macroeconomic uncertainty factor has a large risk premium, about 16%, which is substantially greater than other factors' risk premia.

[15] Driessen, Maenhout and Vilkov (2009), page 28.

The estimated $\hat{\lambda}_{\text{EU}}$ has the same magnitude and statistical significance throughout the other specifications of Table 1.6, while the market and Hml are rendered insignificant by the inclusion of the whole set of factors in specification (4). There are two reasons why the risk premium on EU is so large. First, the model appears to be misspecified, as demonstrated by the large and very significant intercept. Second, factor exposures have to be taken into account when assessing the contribution of a factor to expected excess returns. As discussed in detail in Section 1.5, EU increases the expected excess return on options with a EU beta equal to the median by about 50% a year. While large, this number has to be compared with an average annualized return upwards of 400% for out-of-the-money option portfolios (Table 1.2). Both the first- and second-stage average adjusted R^2 are relatively high, at 38.68% and 48.35% in specification (4).

Measurement error, hedging risk and the underlying stock sensitivity to EU are taken into account in Table 1.7. The coefficient on the option spread is positive and significant, as one would expect given that measurement error positively biases observed returns (Blume and Stambaugh (1983)). If hedging risk exerted an effect on option returns, higher gamma and vega would have a positive coefficient. This is the case for $\hat{\lambda}_{\nu}$, which is strongly significant, while $\hat{\lambda}_{\gamma}$ is negative, much smaller in absolute value, and statistically insignificant in the full specifications (3) and (4). Introducing the hedging risk and measurement error controls reduces the risk premium on EU by about half, and the statistical significance of the intercept is substantially smaller, although it is now positive and larger in absolute value. A comparison of specifications (3) and (4) gives three interesting results. First of all, EU increases the cross-sectional adjusted R^2 by 0.70% in absolute value. Second, $\hat{\lambda}_{\text{EU}}$ takes explanatory power away from $\hat{\lambda}_{\text{Mkt}}$. Third, comparing $\hat{\lambda}_{\text{Vix}_v}$ and $\hat{\lambda}_{\text{Put}_v}$ across the two specifications shows the consequences of explicitly considering macroeconomic uncertainty as a separate source of uncertainty. The risk premium on the volatility of volatility becomes more negative, and the risk premium on the volatility of jumps becomes non-significant. Put differently, the results suggest that, in a reduced form specification, Vix_v and Put_v capture the effect of macroeconomic uncertainty, and because EU and Vix_v have opposite signs, it is difficult to identify $\hat{\lambda}_{\text{Vix}_v}$. While the model is still misspecified, as suggested by the significant intercept, the effect of EU on $\hat{\lambda}_{\text{Vix}_v}$ and $\hat{\lambda}_{\text{Put}_v}$ is found in most of the specifications presented in the rest of the paper.

The most likely source of misspecification in the models discussed in Tables 1.6 and 1.7 is the non-linearity of option returns with respect to returns in the underlying (and hence with respect to some factors). As discussed by Broadie, Chernov and Johannes (2008), applying linear models to option returns may result in biased regression coefficients. To evaluate which factors should enter non-linearly in the specification, and if the relation among any of the factors is also non-linear, I use fractional polynomial regressions to compare specifications in which the factors enter with different powers. The first step is to run time-series regression in which the factors can enter with powers 1 or 1 and 2. If, for a given factor, the percentage of option portfolios in which the best-fitting regression includes a squared term exceeds 5%, the factor is included in the model. The first column of Table 1.8 shows that this is the case only for Mkt, Hml and Umd. Interestingly, only "asset pricing" factors, that is those that have been shown to be important for stocks, enter non-linearly in the time-series regressions. The reason is that the main effect of these factors on options is routed through returns on the underlying stocks.

The second column of Table 1.8 shows whether the relation between the selected factor and other factors includes non-linear terms. Each factor is regressed on all the others using a fractional polynomial model, in which factors can enter with any combination of the following powers: -2, -1, -0.5, 0, 0.5, 1, 2, 3, where 0 means that the log of the factor is included. The same power can also be repeated twice, in which case factor x would enter as $\beta_{x,1}x^p + \beta_{x,2}x^p \ln x$. The model is estimated over five time periods - the full sample and four subperiods - to evaluate the robustness of the non-linearities. Column (2) of Table 1.8 shows whether, for each factor, any of the other factors enters the best-fitting regression non-linearly in at least three of the five time periods. On the basis of the regression of Mkt on the other factors, only Vix squared will be included as a factor in the Fama-MacBeth procedure. The difference between Vix^2 and Vix_v is that the first is the square of the weekly changes in VIX, while the second is the weekly volatility of daily VIX changes. In both columns (1) and (2), the best fitting model is determined as the one with the lowest $D = n(1 + ln\frac{2\pi RSS}{n})$, where n is the number of observations and RSS the residual sum of squares.

Table 1.9 shows the estimated risk premia when including the additional squared factors selected above. To account for heteroskedasticity in cross-sectional regressions, I use Weighted Least Squares, where the weight is given by the inverse of the squared residual predicted by the log relative option spread (os).[16] Unreported Breusch-Pagan tests show that heteroskedasticity is related to os, and indeed this variable is monotonically decreasing in moneyness. The risk premium on EU is still positive and strongly significant, slightly below 10%, and introducing EU reduces $\hat{\lambda}_{\text{Mkt}}$, and still makes $\hat{\lambda}_{\text{Vix}_v}$ more negative and $\hat{\lambda}_{\text{Put}_v}$ insignificant. The risk premia on Mkt^2,

[16] Unless stated otherwise, the results presented from now on include squared terms and the cross-sectional regressions are based on WLS.

Hml^2 and Umd^2 are all positive, quite small but strongly statistically significant. The intercept is now smaller and statistically equal to zero. The model in Table 1.10 deals with heteroskedasticity in a different way, by adding moneyness dummies in the cross-sectional regressions. More specifically, it includes six moneyness dummies, for calls and puts with MonDummy from 1 to 3. The macroeconomic uncertainty risk premium $\hat{\lambda}_{\text{EU}}$ is still positive and strongly significant, although it is now less than half the corresponding values in Table 1.9. The coefficient on *os* is much smaller and less statistically significant, and introducing EU does not produce discernible results on $\hat{\lambda}_{\text{Vix}_v}$ and $\hat{\lambda}_{\text{Put}_v}$. The reason is that the volatility of jumps plays an important role in the cross-section of moneyness (see also Liu, Pan and Wang (2005)), and the dummies absorb the effect. Similarly, as noted above, *os* is monotonic in moneyness, and the dummies take much of its explanatory power away.

An important source of revenue for option market-makers is the bid-ask spread. Assuming that the option price is the mid-point of the bid-ask spread, every trade gives a revenue equal to half the spread. The revenue is used to pay for the running costs of the business, the costs of inventory management, and part of it may go as compensation for the risks associated with trading on a particular source of uncertainty. Table 1.11 shows the time-series average of the coefficients from cross-sectional regressions of a measure of bid-ask revenue on the factor loadings of option returns. The dependent variable in the first two specifications is $\text{BArev}_{t,1} = \frac{1}{2}\frac{0.5 \cdot (\text{ask}_t - \text{bid}_t) \cdot \text{volume}_t}{\text{oi}_t \cdot \text{oprice}_t}$, where the total theoretical revenue is at the numerator, and it is scaled by a proxy for the value of the net holdings of market makers. This measure is then multiplied by $\frac{1}{2}$ to account for fixed and inventory management costs. In specifications (3) and (4) the dependent variable is $\text{BArev}_{t,2} = \frac{1}{2}\frac{0.5 \cdot (\text{ask}_t - \text{bid}_t) \cdot \text{volume}_t}{\text{oi}_{t-1} \cdot \text{oprice}_t}$, that is the proxy for net holdings is the lagged value of the open interest. In the last two specifications the theoretical revenue is multiplied by $\frac{1}{4}$ as an additional robustness. The results show that the option exposure to EU is an important determinant of the bid-ask revenue, and including EU takes explanatory power away from Mkt, Put_v and Skew. Interestingly, higher exposure to Hml reduces BArev.

Before proceeding to the robustness checks section, it is important to discuss the economic interpretation of the risk premium on macroeconomic uncertainty. As pointed out in Section 1.3.3, EU is higher at economic inflection points, when the trend of growth turns from positive to negative and viceversa. This means that, when it comes to the relation with the marginal utility of investors, macroeconomic uncertainty is not entirely interpretable, on average, as a "bad" state of the world, which would explain why $\hat{\lambda}_{\text{EU}}$ is non-negative. An implication is that the interaction of EU with a factor that increases when investors expect a deterioration of the economic conditions should have a negative risk premium, because macroeconomic uncertainty would signal an increased chance of a downturn, rather than of an upturn. In unreported results, this is indeed the case. The risk premium on the interaction between EU and Put is slightly less than -1% a week, with a t-stat of 2.08.

1.5 Robustness checks

The macroeconomic uncertainty factor is built by aggregating announcements about five macroeconomic variables. The first robustness check presents estimated risk premia when the factor is based on one variable at the time, and the main conclusions are not affected. Table 1.12 shows that $\hat{\lambda}_{\text{EU}}$ ranges from 6.93% for the Producer Price

Index to 11.56% for Real Earnings. The intercept is insignificant in every specification, and $\hat{\lambda}_{\text{Put}_v}$ is always lower when EU is included, although it remains statistically significant and the difference is marginal. The risk premium on the market is also smaller, and marginally significant at 10%, while $\hat{\lambda}_{\text{Vix}_v}$ is practically unchanged.

1.5.1 Alternative definitions of returns

Option returns are calculated from prices defined as the mid-point of the bid-ask spread, but investors are unable to trade on the basis of this theoretical price. Rather, they need to buy and sell options at the quoted ask and bid prices. Table 1.13 shows the estimated risk premia when returns are computed by selling at the bid and buying at the ask. The intercept is naturally negative and significant, because the relatively large bid-ask spread generates a substantial net loss for every weekly trade, but $\hat{\lambda}_{\text{EU}}$ is still positive and significant, and both $\hat{\lambda}_{\text{Vix}_v}$ and $\hat{\lambda}_{\text{Put}_v}$ are smaller when EU is introduced, although $\hat{\lambda}_{\text{Put}_v}$ is never significant. The coefficient on *os* is now negative and significant, but this follows the fact that the higher *os*, the higher the spread loss that investors have to incur. Interestingly, the risk premium on Vix is now negative and significant, which confirms the findings of Table 1.11 and suggests that, when aggregate volatility increases, market makers widen the bid-ask spread.

In the results presented so far, if an option has a valid price in week $t-1$, but not in week t, the return is not calculated. This, however, implies that the trading strategy is not replicable, because it is based on information unavailable at time $t-1$. For this reason, the first two specifications of Table 1.14 assume a return of -50% and 0% if the price at t is not available. The 0% return is essentially equivalent to the riskless rate, given the short horizon. In specification (3) the EU factor is computed as the *average* rather than *median* return of the option portfolios in the high Excess Pricing Error quintile, which implies that, trading costs aside, the factor can be interpreted as the return on an equally weighted portfolio. Specification (4) shows risk premia when first stage regressions are based on the Cochrane-Orcutt procedure, to account for any autocorrelation in option returns. The results discussed in Section 1.4 prove to be robust to the four specifications, with $\hat{\lambda}_{\mathrm{EU}}$ equal to about 7% and the intercept remaining statistically insignificant. The adjusted R^2s are noticeably lower if missing returns are substituted with 0%, and especially if with -50%. Specifications (5) to (7) present a sub-period analysis, for years 1996-1999, 2000-2003, and 2004-2007. The risk premium on macroeconomic uncertainty is still positive and significant, equal to 5.2%, 6.5% and 3.5% in the three subperiods. The intercept is positive and significant for 2000-2003, while $\hat{\lambda}_{\mathrm{Vix}_v}$ is positive in 1996-1999 and negative in the second part of the sample.

1.5.2 Alternative ways to measure macro uncertainty

As discussed in Section 1.3.3, the rationale behind the way the EU factor is constructed is that macroeconomic uncertainty is higher before scheduled announcements, that it is resolved afterwards, and that this change is measurable in option prices. In this subsection I investigate how four alternative definitions of the factor impact the results. First of all, the day of higher macroeconomic uncertainty is assumed to be $t-2$ rather than $t-1$, and pricing errors on such days are still normalized by subtracting the median on all other days. Second, the day of higher macroeconomic uncertainty is assumed to be $t-1$, and pricing errors are normalized by subtracting those on day t. In addition, to isolate the contribution of individual announcements, I repeat the second robustness check above by only considering announcements that are set apart by at least 3 and 7 days. Specifications (1) to (6) in Table 1.15 show that $\hat{\lambda}_{\text{EU}}$ is unchanged, but it is interesting to compare the effect of the different definitions of EU on the magnitude and statistical significance of the risk premia on Mkt, Vix_v and Put_v. All of them are larger than in the equivalent specifications in Table 1.9, and the pattern of statistical significance is reversed: $\hat{\lambda}_{\text{Vix}_v}$ is insignificant while $\hat{\lambda}_{\text{Put}_v}$ and $\hat{\lambda}_{\text{Mkt}}$ are significant. Interestingly, $\hat{\lambda}_{\text{Put}_v}$ is larger when the day of higher uncertainty is assumed to be $t-2$, with the exception of specifications (5)-(6). The overall conclusion is that the risk premium on EU is robust to measuring uncertainty in different ways around announcements, although the more stable pattern in how $\hat{\lambda}_{\text{EU}}$ takes explanatory power away from other risk premia suggests that the most accurate measurement is the one presented in Section 1.3.3.

The last two specifications of Table 1.15 exclude options on stocks that pay dividends within expiration, to rule out that pricing errors may be related to unobserved differences in the valuation of options on dividend paying stocks. The macroeconomic uncertainty risk premium, in this case, is more than 10%.

1.5.3 Alternative definitions of expected volatility

The calculation of the macroeconomic uncertainty factor is based on pricing errors, defined with respect to the Black-Scholes model. The formula includes the Black-Scholes price computed with the implied volatility provided by OptionMetrics ($\text{SBS}_{i,t}$, see Section 1.3.3), which is a synthetic Black-Scholes price that can be compared with the one calculated using the expected objective volatility ($\text{BS}_{i,t}$). The implied volatility is not available, however, when the mid-point of the bid-ask spread is lower than the intrinsic value, when vega is below 0.5 or when the optimization fails to converge.[17] Duarte and Jones (2007) note that this sample selection is related to censoring on the basis of measurement error, because large measurement errors are more likely to result in a violation of arbitrage bounds. If the censoring is not randomly distributed across the days used to measure macroeconomic uncertainty, the EU factor may proxy for the causes of non-random measurement error. To address this problem I define an augmented implied volatility, which is equal to the implied volatility of the same option contract on the previous day, or two days before, depending on availability. Specifications (1) and (2) of Table 1.16 show the estimated risk premia when the augmented implied volatility is used instead of the implied volatility provided by OptionMetrics. It is important to note that the returns of options with missing implied volatility are still excluded, because the controls ν and γ are not available without implied volatility. Table 1.6 shows results when ν and γ are not included and there is no sample selection in option returns due to missing implied volatility.

The remaining specifications in Table 1.16 are based on different definitions of the expected objective volatility that enters the calculation of pricing errors. In (3) and (4) I use a trailing 30 days volatility, with the exclusion of day t, to make sure that the volatility does not contain more information about future returns than the option implied volatility. (5) and (6) are based on a trailing 30 days volatility, excluding zero returns, so that there is no bias due to the liquidity of the underlying stock. Finally, specifications (7) and (8) use realized volatility in the subsequent 2 (7) months for options with MatDummy=1 (=2), to make sure that differences in the pricing errors between long/short maturity options are not due to the term structure of expected

[17] IvyDB reference manual v.2.6, page 33.

volatility.

The results in Table 1.16 show that the risk premium on EU remains positive and significant throughout the specifications and, with the exception of (1) and (2), $\hat{\lambda}_{\text{EU}}$ makes $\hat{\lambda}_{\text{Vix}_v}$, $\hat{\lambda}_{\text{Put}_v}$ and $\hat{\lambda}_{\text{Mkt}}$ smaller.

1.5.4 Non-randomly missing option returns

The filters applied to the dataset (see Table 1.1) mean that option portfolio returns may be missing on certain dates. I need to evaluate whether this is a random occurrence, and if there is a pattern I have to account for it to avoid possible biases in the estimated risk premia. The first step is to establish a relation between the probability of missing returns, factors and other relevant variables. I do so by estimating a logit for each option portfolio, in which the dependent variable is a dummy equal to one if the return is missing, and the independent variables are a set of factors plus a time trend, and the natural logarithm of the trading volume, open interest and option spread. The last three variables are defined as either the median across put/call type on a given date (specification (1) of Table 1.17) or the median for each portfolio (specification (2)). Table 1.17 shows the median odds ratios across portfolios, with the relative 90% confidence intervals. The results suggest that the time trend reduces the probability of missing returns, as the option market grows and becomes more liquid, while an increase in the open interest makes it more likely to have missing returns. Of the factors considered, the odds ratios for the market, Smb, Hml, and Vix are less than one, while the macroeconomic uncertainty factor has an odds ratio equal to about two in both specifications. Trading volume and open interest are two measures of activity on the option market, and it may be natural to expect that their odds ratio are less than one, especially in specification (2), where open interest is measured at the portfolio level. While volume is a direct gauge of trading activity - and the odds ratio is indeed less than one - open interest is the number of outstanding contracts. The consequence is that changes in the open interest not only reflect the number of contracts added on a given date, but also the number of contracts that are delivered, and this can be especially significant for equity options, which have American exercise. It is then possible that the odds ratio for volume reflects the negative relation between trading activity and missing returns, while the odds ratio on the open interest measures the effect a reduction in the number of contracts that are delivered.

Table 1.18 shows the betas from a regression where the dependent variable is the change in the number of equity options that satisfy the filters, and the independent variables are the same as in the logit described above. The log-volume increases the number of observations, as does the Smb factor, while the log-open interest and the EU factor decrease them. The sign of these coefficients is consistent with the corresponding odds ratios in Table 1.17, as a positive/negative effect on the number of observations is associated with a lower/higher probability of observing a missing return. The coefficient on Vix, however, is negative, while the odds ratio is much smaller than one. This suggests that, when volatility increases, trading becomes concentrated in a smaller set of options and the likelihood of a missing return is lower (on a given week, option portfolios need to include at least two options for the return to be valid).

After having studied the relation between the probability of missing returns and a set of factors and measures of market activity, I re-estimate the option risk premia with a Heckman model in place of the standard time-series regressions. The selection equation includes the variables that have a significant odds ratio in the logit described above, namely Mkt, Smb, Hml, Vix, EU, a time trend, log-volume, log-open interest and log-option spread. Table 1.19 shows the results, which confirm that $\hat{\lambda}_{\mathrm{EU}}$ is positive and statistically significant, and that it decreases the risk premia on Mkt, and Put_v. It also makes the risk premium on Vix_v more negative in specifications (3) and (4), where the log-volume, log-open interest and log-option spread are median portfolio values.

1.5.5 Factor contributions to expected excess returns

The results discussed so far suggest that the risk-premium on the macroeconomic uncertainty factor is between 5% and 7% per week. This is quite a large number, but, to evaluate the contribution of a factor to expected excess returns, it is important to take the distribution of factor betas into consideration. Table 1.20 shows, for different specifications, the annualized product between factor risk premia and 25^{th}, 50^{th} and 75^{th} percentile betas. With the exception of specifications (1) and (2), which include moneyness dummies in second-stage regressions, all cross-sectional regressions are estimated with WLS. Dots mean that the factor is not included in the specification, and zeros that the risk premium, or the product, is not significant. There is a marked difference in the contribution of macroeconomic uncertainty, depending on the estimation method for second-stage regressions. When moneyness dummies are included, the median effect is about 30% a year, while it increases up to about 75% when the estimation is based on WLS. The magnitude is relatively insensitive to what other factors are included and to alternative factor definitions. The market, Mkt^2 and Hml^2

are also significant across most specifications, with a median contribution of, respectively, about 20-25%, 10% and 10-13%, although it falls to about 2-4% for Hml^2 for alternative definitions of EU. Vix_v has a zero median effect, but it is about 10% and -10% for 25^{th} and 75^{th} percentile betas. Lastly, the risk premium on Put_v is significant in 5 of the 11 specifications, with a median contribution to expected returns of about 4% a year.

1.5.6 Simulation

Option returns are non-linear in the underlying's returns, and this can bias the estimates of linear factor models (see Broadie, Chernov and Johannes (2008)). In order to account for non-linearity, I have included selected squared factors in first and second-stage regressions, and used moneyness dummies or WLS in cross-sectional regressions. I now rely on a simulation to evaluate whether this approach generates accurate estimates of factor risk premia.

I simulate returns on 50 assets using a single factor model, and then compute Black-Scholes returns on several options on each asset. Following the empirical analysis discussed in the previous sections, I form option portfolios on the basis of moneyness, maturity, put/call contract type and stock sensitivity to the single factor (or *market*). While the Black-Scholes model is relatively simple and omits factors like volatility and jumps, the simulation is meant to analyse the effect of option returns' non-linearity, and this can be accomplished with a single factor model. Tellingly, Table 1.8 shows that, in the actual data, Mkt is the only factor for which a squared term clearly improves the fit of the regressions.

The simulated option returns have weekly frequency, for a total of 11.5 years. The simulation is repeated with the market risk premium equal to 0% and 10%. The riskless rate and the volatility of the market are constant, at 3% and 18% year. Asset betas and volatilities match the range of the corresponding variables in the sample, with market betas starting at 0.4 and up to 1.9, while annualized volatilities are between 20% and 95%. Every week I compute Black-Scholes prices for call and put options with maturity of 2 and 7 months, and moneyness ln(K/S) equal to -15, -6.25, 0, 6.25, 15%. These values are the mid-points of the moneyness and maturity intervals defined in section 1.3.2. I then form option portfolios by intersecting the following characteristics: underlying stock's beta with respect to the market (10 deciles), moneyness (5 categories), maturity (2) and call/put type (2). Weekly portfolio returns are the equally weighted average of the constituent options' returns. To reduce the effect of outliers on the results, I censor option returns at both the option and portfolio level with cut-offs of, respectively, 0.175% and 0.35%. These are a fraction ($\frac{1}{3}$) of the cutoffs used in the empirical analysis, because simulated returns are not affected by measurement error and non-synchroneity.

I test whether the estimated market risk premium, $\hat{\lambda}_{\text{Mkt}}$, is statistically different from the risk premium used to generate stock returns, and if the intercept $\hat{\alpha}_2$ is different from zero. First and second stage regressions include the market squared, and moneyness dummies (category 1 to 3, for calls and puts separately) are added to cross-sectional regressions.

Table 1.21 reports the average $\hat{\lambda}_{\text{Mkt}}$ and $\hat{\alpha}_2$ across the 1,000 replications, along with 90% confidence intervals based on bootstrapped standard errors. The results suggest that, after accounting for option returns' non-linearity, the Fama-MacBeth method accurately estimates option risk premia. More precisely, when the risk premium is equal to 10%, both $\hat{\alpha}_2$ and $\hat{\lambda}_{\text{Mkt}}$ are not statistically different from their true values (0% and 10%). For a zero risk premium, on the other hand, $\hat{\lambda}_{\text{Mkt}}$ is negatively biased by about 36 basis points a year. The 95% confidence interval, however, does include zero.

1.6 Conclusion

I study whether macroeconomic uncertainty is a priced risk factor in the cross-section of equity option returns, and how its inclusion changes the risk premia on other volatility and jump factors. The analysis is based on a series of factor models, estimated with the Fama-MacBeth methodology. The macroeconomic uncertainty factor is built using three results in the literature. First, uncertainty about important economic variables is higher before the release of economic news, and is rapidly resolved upon announcement. Second, the implied volatility of equity options increases before scheduled releases and drops afterwards. Third, this change is correlated with the implied volatility of options on macroeconomic variables, which were traded between 2002 and 2007. The sensitivity of options to macroeconomic uncertainty is estimated by studying the pattern of option pricing errors around announcement days, and the macroeconomic uncertainty factor is a portfolio that buys options with the highest estimated sensitivity.

The analysis includes a large set of relevant asset- and option-pricing factors, and controls for measurement error and hedging risk. The results show that macroeconomic uncertainty is a priced risk factor for equity options, and this conclusion is robust to the effects of non-linearity of option returns with respect to some of the factors, to alternative methods of estimating option sensitivity to macroeconomic uncertainty, to calculating returns by selling at the bid and buying at the ask price, to adjustments for non-randomly missing returns, and to alternative definitions of the expected objective volatility.

The contribution of the macroeconomic uncertainty factor to expected excess option returns is about 70% a year for option portfolios with a beta equal to the median factor loading. This compares to 20-25% for the market factor. Including the macroeconomic uncertainty factor reduces the risk premium on the market and on the volatility of jumps, while it makes the premium on the volatility of volatility more negative and significant. This suggests that it is important, also in terms of future research, to identify the economic foundations of the reduced-form sources of uncertainty that are commonly used in option pricing.

Figure 1.1: Non-linearity of option returns and OLS. Call options.

Market returns, and returns on the at-the-money call option portfolio with the highest sensitivity to Mkt (third tertile), short maturity and low sensitivity to Smb and Hml (first tertile). One regression line is from standard OLS of option returns on market returns ($\beta = 9.52$, $\alpha = -0.0036$), the other is calculated by assigning a weight of 0.25 to observations with market returns below -0.028, which is the 25% percentile ($\beta = 9.97$, $\alpha = -0.0077$).

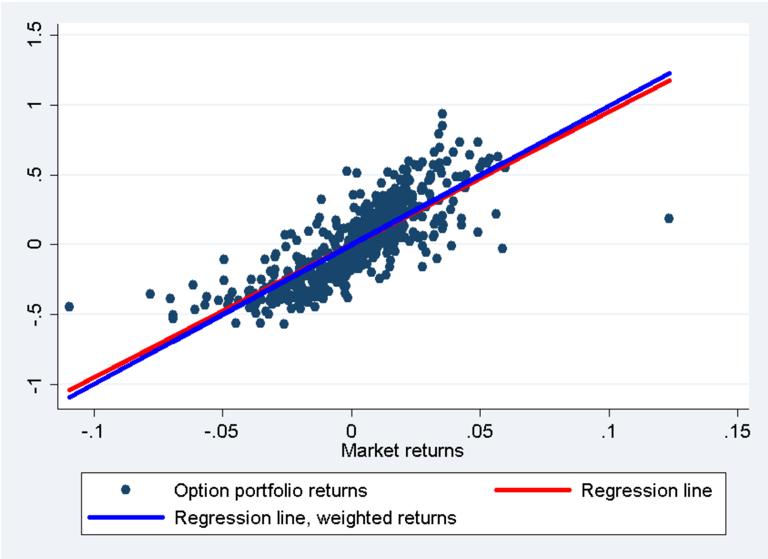

Figure 1.2: Non-linearity of option returns and OLS. Put options.

Market returns, and returns on the at-the-money put option portfolio with the highest sensitivity to Mkt (third tertile), short maturity and low sensitivity to Smb and Hml (first tertile). One regression line is from standard OLS of option returns on market returns ($\beta = -10.15$, $\alpha = -0.031$), the other is calculated by assigning a weight of 0.25 to observations with market returns above 0.027, which is the 75% percentile ($\beta = -10.79$, $\alpha = -0.037$).

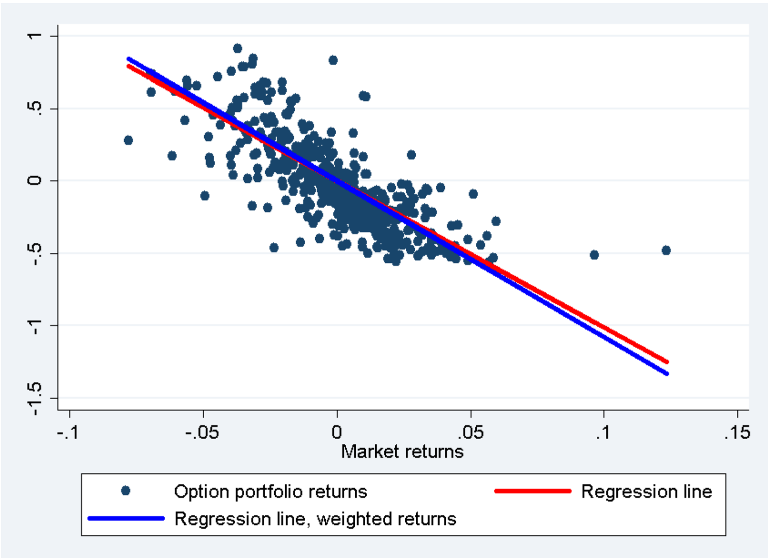

Figure 1.3: Term structure of implied volatility.

The implied volatility of individual options is standardized using the implied volatility of the longest maturity option with the same strike, put/call type and on the same stock, provided that its days to expiration are between 200 and 250. The figure shows median implied volatilities, by time to maturity, over the whole sample, across moneyness categories and put/call type.

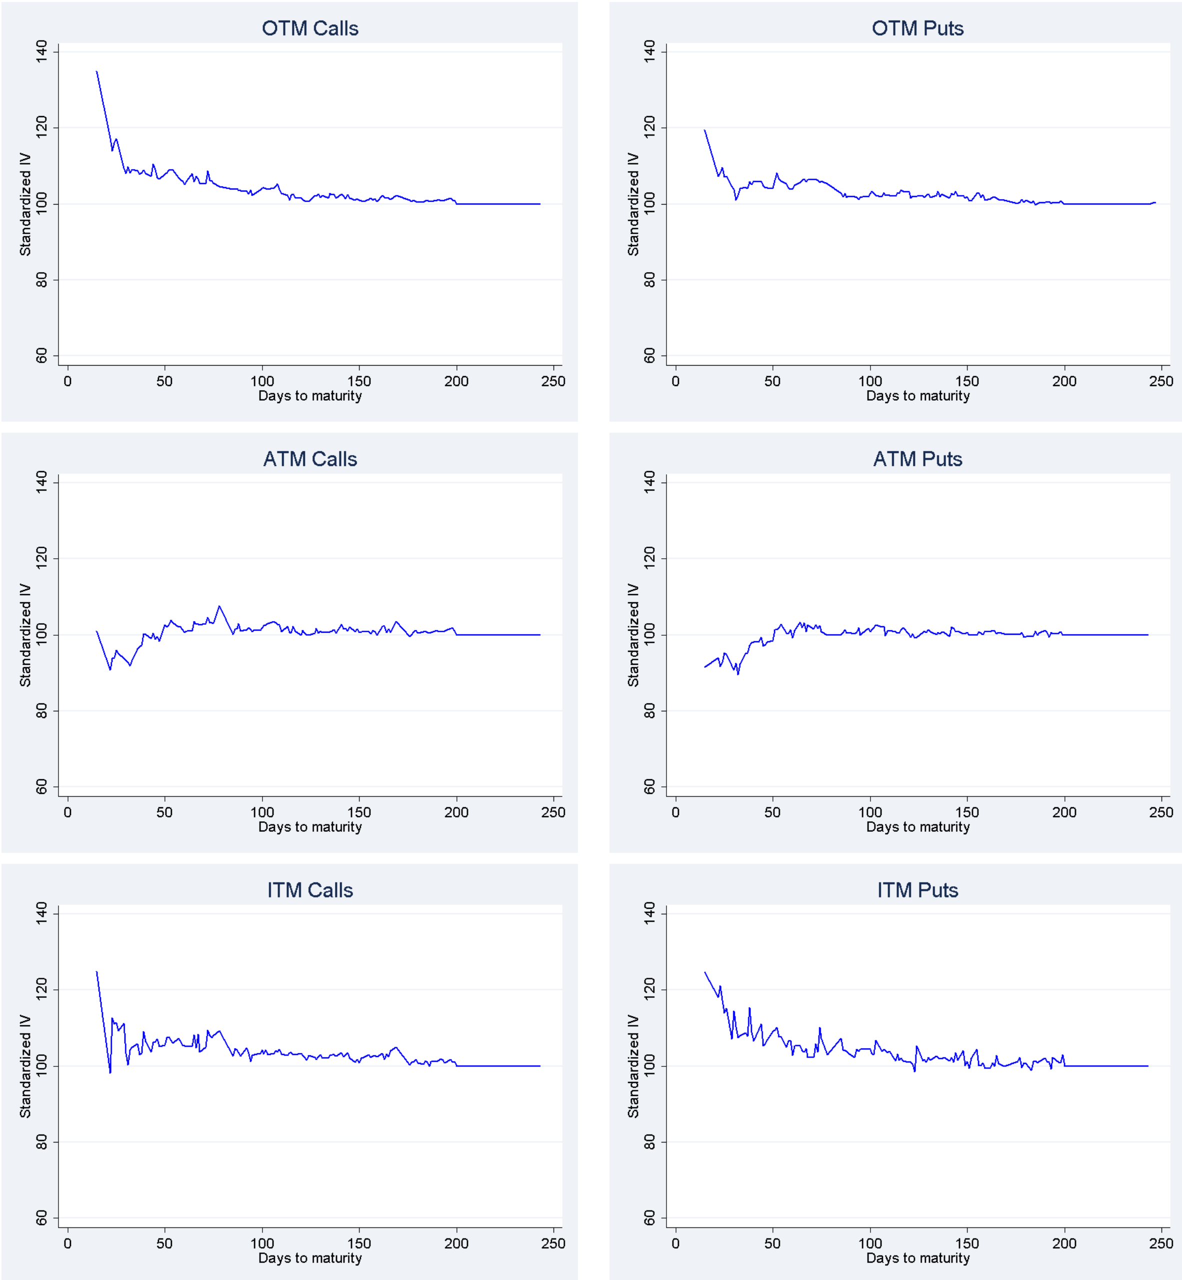

Figure 1.4: Composition of high EPE quintile.

For a given characteristic, the figure shows the fraction of option portfolios that belong to the high Excess Pricing Error (EPE) quintile, across different values of the characteristic. Pmkt, Psmb and Phml are the sensitivity of the underlying stock to the Fama-French factors (1=Low, 3=High). MatDummy and MonDummy are the maturity and moneyness categories (MatDummy=1,2 for short and long maturity options. MonDummy=1,...,5 for out-of-the-money to in-the-money options). Pricing errors are defined using realized volatility over the previous 30 days.

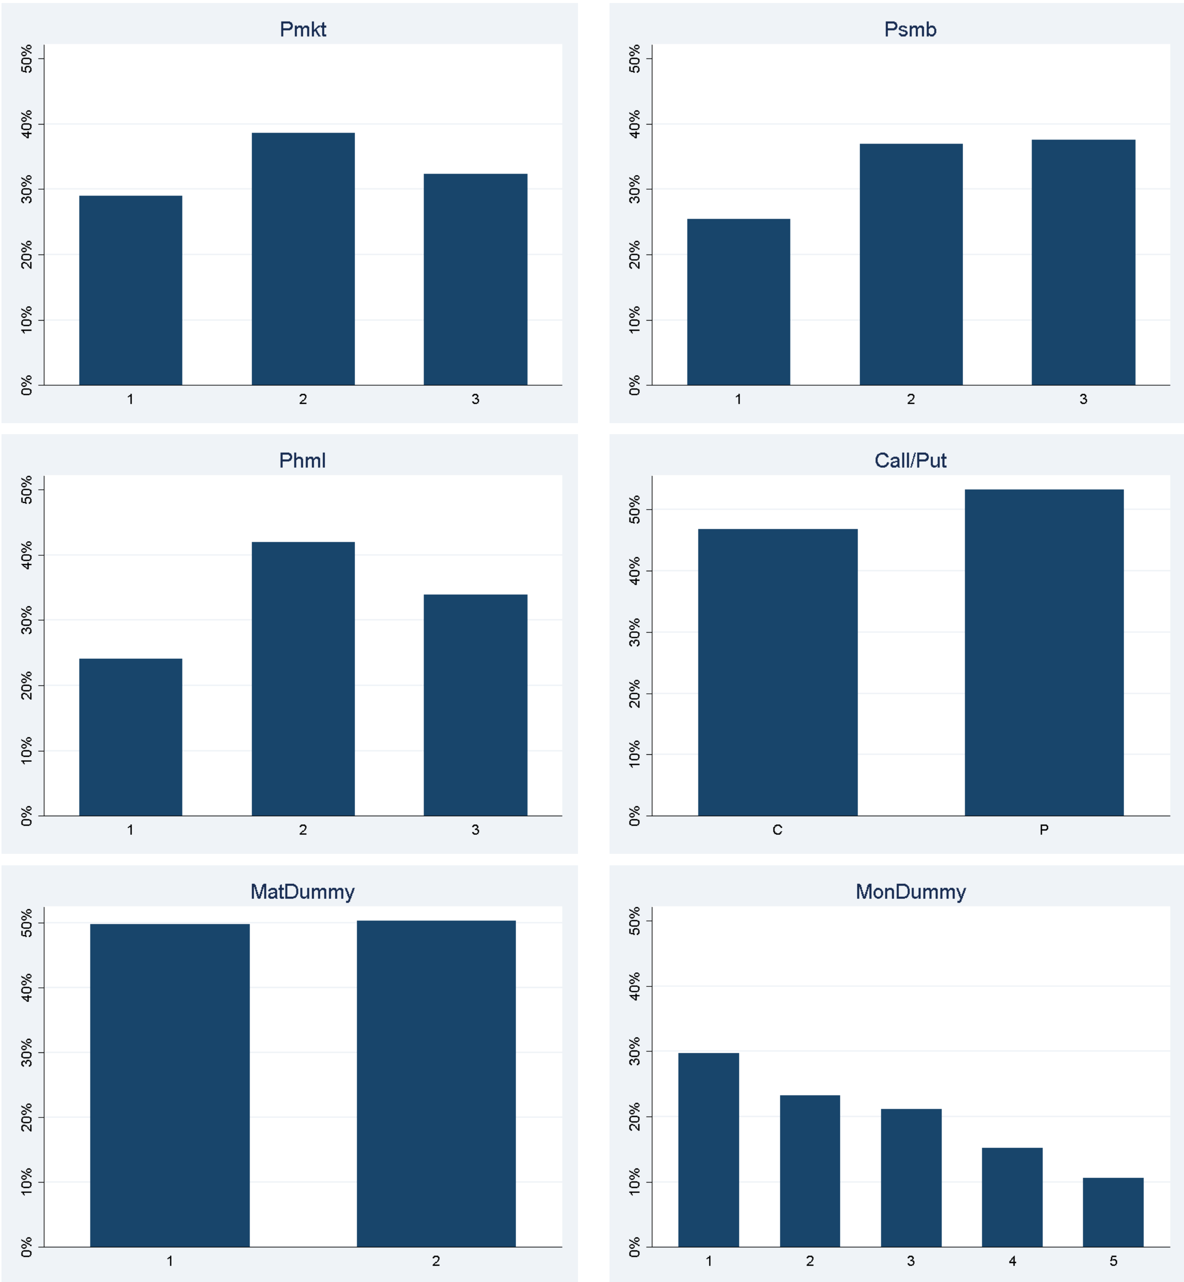

Figure 1.5: Macroeconomic uncertainty, S&P 500 level, capacity utilization and forecast dispersion.

Quarterly average of the residuals from regressing EU on Mkt, Smb, Hml, Vix, Vix_v, Skew and Put, and (top) end-of-quarter level of the S&P500, (center) end-of-quarter level of the Capacity utilization index and (bottom) the FD index. FD is a forecast dispersion index, computed as the quarterly volatility of analyst forecasts for one-year-ahead earnings-per-share. The correlation between EU and FD over the whole sample is 0.21. It is equal to 0.31, 0.43 and 0.73 when EU is greater than the 50^{th}, 60^{th} and 75^{th} percentile. The EU factor is defined using realized volatility over the previous 30 days and average option portfolio returns.

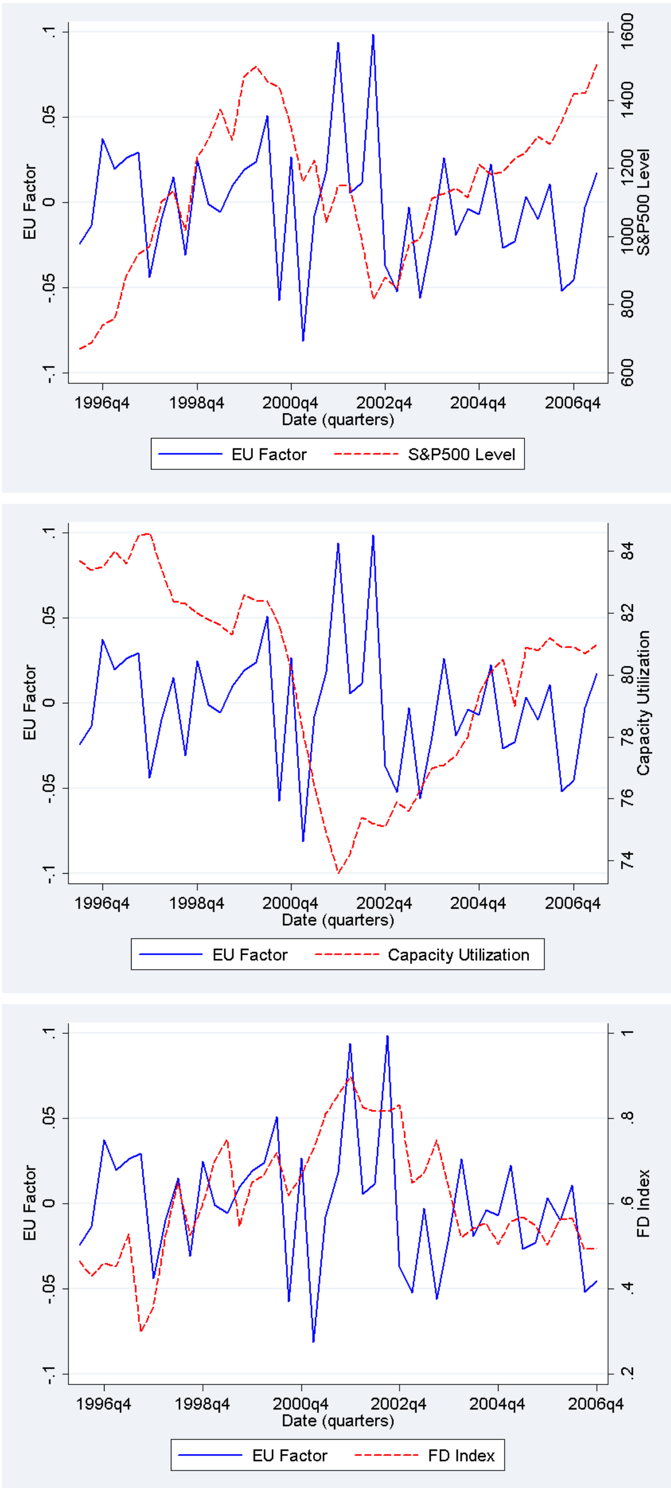

Table 1.1: Average number of monthly observations, by year.

Average number of monthly observations before and after a series of filters. Filters drop observations as follows. (1) An option identifier refers to more than one option series. (2) Volume is equal to zero. (3) Bid or ask price is equal to zero. (4) Ask price less than bid. (5) Contract is non-standard. (6) Underlying is not a stock. (7) Exercise style is European or unspecified. Figures and percentages are rounded to the nearest integer.

Year	Total	ID unique (1)	Non-zero volume (2)	Non-zero ask/bid (3)	Ask> bid (4)	No special settlement (5)	Stock options (6)	American exercise (7)	Used as % of total
1996	1,218,015	1,218,015	219,275	213,996	213,580	211,878	206,488	206,488	17%
1997	1,777,007	1,777,007	276,877	270,572	270,190	266,594	259,326	259,326	15%
1998	2,197,941	2,197,941	331,838	323,959	322,766	315,602	305,768	305,768	14%
1999	2,489,147	2,489,147	420,261	411,388	409,160	401,198	387,729	387,685	16%
2000	3,156,323	3,155,881	603,607	589,670	579,924	568,110	545,233	545,218	17%
2001	2,688,616	2,686,705	447,577	432,050	430,091	423,774	403,207	403,193	15%
2002	2,570,823	2,569,443	406,548	390,237	389,081	384,719	361,707	361,707	14%
2003	2,378,004	2,375,249	415,635	401,125	400,380	396,690	368,123	368,123	15%
2004	2,857,701	2,854,478	512,339	493,711	493,056	488,309	449,097	449,097	16%
2005	3,226,817	3,219,752	597,811	574,714	574,194	569,441	518,830	518,830	16%
2006	3,723,522	3,713,102	752,253	725,609	725,194	718,762	645,143	645,143	17%
2007	3,972,290	3,960,454	815,671	787,682	787,490	780,847	682,246	682,246	17%

Table 1.2: Return statistics.

Time-series statistics for option returns, averaged across the values of the characteristic used to form option portfolios. μ is the mean, σ_t the volatility of the mean, μ^3 and μ^4 the skewness and kurtosis, and the remaining columns show selected percentiles. All figures, with the exception of skewness and kurtosis, are percentages. Means with a superscripted asterisk are significant at 10% or less. 1996(q2) to 2007(q2).

		μ	σ_t	μ^3	μ^4	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Pmkt	**1**	-0.54*	0.23	0.51	3.37	-6.95	-4.36	-1.16	2.99	7.05
	2	-0.08	0.27	0.28	3.08	-8.29	-4.46	-0.62	3.88	8.61
	3	-0.40	0.30	0.94	6.66	-8.72	-5.15	-1.42	3.62	9.05
Psmb	**1**	-0.25	0.20	0.32	3.07	-5.80	-3.46	-0.72	2.82	6.33
	2	-0.05	0.43	0.48	4.88	-11.73	-5.89	-1.22	4.75	12.90
	3	-0.70	0.51	0.44	4.70	-14.29	-7.68	-1.65	5.88	14.08
Phml	**1**	-0.78*	0.20	0.25	3.32	-6.68	-3.76	-0.74	2.05	5.16
	2	0.86*	0.46	0.42	3.91	-11.19	-6.27	-0.09	7.37	15.52
	3	1.10	0.73	1.22	7.13	-17.39	-8.62	-0.41	8.06	20.43
Moneyness	**1**	7.41*	0.43	0.64	4.07	-3.79	0.55	5.55	13.99	20.76
	2	1.09*	0.31	0.69	4.16	-7.29	-3.81	0.49	5.17	10.21
	3	-1.02*	0.28	0.23	3.24	-9.40	-5.32	-1.14	2.99	7.49
	4	-3.52*	0.25	-0.15	3.69	-11.19	-7.27	-3.48	0.41	3.97
	5	-9.49*	0.33	-0.79	3.93	-19.63	-13.76	-8.40	-4.09	-1.10
Maturity	**1**	-1.22*	0.27	0.34	3.03	-9.01	-5.62	-1.77	2.74	7.56
	2	0.79*	0.22	0.26	3.15	-5.71	-2.68	0.58	3.92	7.41
Style	**Call**	0.89	0.67	0.10	2.76	-20.02	-11.11	1.04	12.32	20.61
	Put	-2.48*	0.77	0.65	3.13	-22.92	-15.59	-6.31	9.44	23.40

Table 1.3: Sign of option return residuals, trading volume and macroeconomic announcements.

For each option portfolio, returns are regressed on Mkt, Mkt^2, Smb, Hml, Hml^2, Vix, Vix^2, Vix_v, Skew, Put, Put_v and EU. Mkt, Smb and Hml are the market, Small-minus-Big and High-minus-Low Fama-French factors. Vix and Vix_v are weekly changes in VIX, and changes in the weekly volatility of daily VIX changes. Skew is the change in the weekly skewness of Dow Jones Industrial intradaily returns. Put and Put_v are the weekly average of daily returns on out-of-the-money S&P 500 put options, and changes in the weekly volatility of daily returns on out-of-the-money S&P 500 put options. EU is the macroeconomic uncertainty factor. A dummy variable is equal to zero if the regression residuals are positive, and one if they are negative. The dummy is the dependent variable in a probit model, in which the independent variables are weekly percentage changes in the average trading volume of all options, the number of macroeconomic announcements in the week, and their interaction. The table reports the median odds-ratio across the indicated characteristics, and the 90% confidence interval, calculated using the binomial method of Mood and Graybill (1963).

		Δ Volume			# Weekly Ann.ts			Interaction		
		Med.	90% C.I.		Med.	90% C.I.		Med.	90% C.I.	
Pmkt	**1**	1.25	1.01	1.33	1.00	0.99	1.03	1.13	1.05	1.37
	2	1.20	1.11	1.41	0.98	0.95	1.01	1.08	1.01	1.16
	3	1.16	1.01	1.30	0.96	0.94	0.97	1.06	1.00	1.14
Psmb	**1**	1.21	1.12	1.31	0.99	0.97	1.01	1.07	1.01	1.12
	2	1.08	0.95	1.28	0.94	0.93	0.99	1.22	1.04	1.55
	3	1.24	1.01	1.48	0.94	0.91	0.97	1.13	0.97	1.29
Phml	**1**	1.24	1.16	1.32	0.97	0.95	0.98	1.10	1.05	1.14
	2	1.01	0.78	1.26	1.02	0.96	1.08	1.02	0.89	1.28
	3	1.07	0.92	1.48	0.97	0.93	1.01	1.14	0.93	1.37
Moneyness	**1**	1.10	0.99	1.24	0.98	0.96	1.01	1.04	0.97	1.18
	2	1.32	1.11	1.44	0.97	0.95	1.00	1.05	0.98	1.14
	3	1.16	0.97	1.30	0.99	0.96	1.03	1.15	1.06	1.29
	4	1.19	0.92	1.30	0.94	0.89	0.97	1.10	1.01	1.23
	5	1.28	0.87	1.48	0.94	0.88	1.00	1.11	0.94	1.28
Maturity	**1**	1.16	1.07	1.24	0.95	0.94	0.97	1.10	1.04	1.16
	2	1.27	1.11	1.40	0.98	0.97	1.00	1.09	1.01	1.15
Style	**Call**	1.10	1.01	1.18	0.97	0.96	0.99	1.10	1.04	1.16
	Put	1.43	1.24	1.59	0.96	0.93	0.98	1.05	0.97	1.14

Table 1.4: Factor statistics.

The table reports factor statistics. μ is the mean, σ_t the volatility of the mean, μ^3 and μ^4 the skewness and kurtosis, and the remaining columns show selected percentiles. Mkt and Liq are the market and liquidity factors. Hml, Smb and Umd are the High-minus-Low, Small-minus-Big and Momentum factors. Vix, Vix_v, Vix_s are weekly changes in VIX, and changes in the weekly volatility and skewness of daily VIX changes. Skew and Kurt are changes in the weekly skewness and kurtosis of Dow Jones Industrial intradaily returns. Put, Put_v and Put_s are the weekly average of daily returns on out-of-the-money S&P 500 put options and changes in the weekly volatility and skewness of daily returns on out-of-the-money S&P 500 put options. EU is the macroeconomic uncertainty factor. All figures, with the exception of Kurt and Skew, are percentages. Means with a superscripted asterisk are significant at 10% or less. 1996(q2) to 2007(q2).

	μ	σ_t	μ^3	μ^4	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Mkt	0.14	0.10	-0.20	5.45	-2.88	-1.29	0.34	1.54	2.81
Smb	0.04	0.06	-0.88	10.97	-1.57	-0.80	0.11	0.92	1.66
Hml	0.12*	0.06	0.69	6.73	-1.50	-0.66	0.08	0.77	1.68
Umd	0.23*	0.09	-0.63	7.56	-2.08	-0.72	0.28	1.35	2.64
Liq	0.00	0.02	0.06	3.69	-0.46	-0.22	0.01	0.25	0.47
Vix	0.00	0.10	0.18	6.86	-2.57	-1.13	-0.04	1.12	2.81
Vix_v	0.00	0.03	0.59	7.15	-0.78	-0.38	-0.02	0.33	0.86
Vix_s	0.02	0.08	-0.26	4.02	-2.36	-1.09	0.12	1.04	2.36
Skew	0.00	0.05	0.02	2.99	-1.66	-0.86	0.02	0.87	1.57
Kurt	0.00	0.09	0.02	4.33	-2.60	-1.19	-0.01	1.15	2.42
Put	-5.04*	0.54	0.92	4.79	-19.48	-13.69	-6.67	0.91	11.87
Put_v	0.11	0.53	0.11	3.22	-15.33	-7.16	-0.14	7.50	15.58
Put_s	0.19	3.57	0.18	3.51	-108.41	-54.47	0.78	51.8	103.31
EU	0.34	0.62	0.54	3.86	-16.15	-8.63	-1.64	8.72	20.82

Table 1.5: Factor correlations.

The table shows factor correlations. Mkt and Liq are the market and liquidity factors. Hml, Smb and Umd are the High-minus-Low, Small-minus-Big and Momentum factors. Vix, Vix_v, Vix_s are weekly changes in VIX, and changes in the weekly volatility and skewness of daily VIX changes. Skew and Kurt are changes in the weekly skewness and kurtosis of Dow Jones Industrial intradaily returns. Put, Put_v and Put_s are the weekly average of daily returns on out-of-the-money S&P 500 put options and changes in the weekly volatility and skewness of daily returns on out-of-the-money S&P 500 put options. EU is the macroeconomic uncertainty factor. 1996(q2) to 2007(q2).

	Mkt	Smb	Hml	Umd	Liq	Vix	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU
Mkt	1													
Smb	0.01	1												
Hml	-0.56	-0.29	1											
Umd	-0.15	0.22	0.06	1										
Liq	0.13	-0.01	-0.08	-0.06	1									
Vix	-0.76	0.10	0.34	0.11	-0.07	1								
Vix_v	-0.29	-0.05	0.15	0.01	-0.07	0.28	1							
Vix_s	-0.44	0.17	0.24	0.05	0.01	0.65	0.16	1						
Skew	0.08	-0.07	-0.08	-0.01	0.01	-0.12	0.09	-0.16	1					
Kurt	0.08	0.03	-0.10	-0.01	0.02	0.00	0.04	-0.08	0.20	1				
Put	-0.56	0.00	0.27	-0.02	-0.10	0.51	0.12	0.19	-0.01	-0.02	1			
Put_v	-0.20	-0.02	0.17	-0.08	-0.06	0.16	0.41	0.03	0.10	0.06	0.24	1		
Put_s	0.03	0.01	0.05	0.00	0.03	0.00	0.07	0.07	-0.04	-0.05	-0.16	0.23	1	
EU	0.37	0.05	-0.10	0.00	0.04	-0.21	-0.01	-0.12	0.00	0.04	-0.25	-0.01	0.02	1

Table 1.6: Risk-premia estimates.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). *os* and *ss* are the median relative option and stock bid-ask spreads. ν and γ are the median log vega and log gamma. β_s is the mean stock beta with respect to the EU factor. See Table 1.4 for factor definitions. $\bar{R}_1^2$ and $\bar{R}_1^2$ are the average adjusted R-squared from first and second-stage regressions.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix$_v$	Vix$_s$	Skew	Kurt	Put	Put$_v$	Put$_s$	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
						-0.0321	0.0028		-0.0015	-0.0029					0.0063	
						-10.43	2.42		-0.98	-2.04					2.22	
(1)																
	-0.0005		0.4061		0.0349	0.0190		0.1673		0.3802	0.4511					
	-0.48		1.54		1.65	1.11		8.72								
						-0.0291	0.0014		-0.0031	-0.0013						
						-11.25	1.39		-2.22	-0.97						
(2)																
								0.1697		0.3793	0.4213					
								9.25								
						-0.0305									0.0017	
						-10.73									0.87	
(3)																
	0.0010		0.8784		0.0491	0.0227		0.1559		0.2536	0.3815					
	0.96		3.44		2.65	1.06		8.87								
						-0.0297	0.0016		-0.0025	-0.0005		-0.0135		0.0002	0.0034	
						-10.38	1.46		-1.64	-0.34		-4.19		0.31	1.26	
(4)																
	-0.0018	0.0009	0.1754	0.9613	0.0060	0.0096	0.0519	0.1444		0.3868	0.4835					
	-1.62	0.37	0.70	2.59	0.29	0.54	0.31	8.18								

Table 1.7: Risk-premia estimates, with controls for measurement error and hedging risk.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). *os* and *ss* are the median relative option and stock bid-ask spreads. ν and γ are the median log vega and log gamma. β_s is the mean stock beta with respect to the EU factor. See Table 1.4 for factor definitions. $\bar{R}_1^2$ and $\bar{R}_1^2$ are the average adjusted R-squared from first and second-stage regressions.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix$_v$	Vix$_s$	Skew	Kurt	Put	Put$_v$	Put$_s$	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
	0.0543	0.0038	0.0367	-0.0068	0.0174	0.0490	0.0024		-0.0006	-0.0020					-0.0003	
	13.00	0.98	13.53	-2.47	0.19	2.87	2.24		-0.47	-1.81					-0.12	
(1)																
	-0.0015		0.1286		0.0121	0.0052		0.0906		0.3802	0.4757					
	-1.89		0.71		0.80	0.42		6.49								
	0.0661	0.0066	0.0375	-0.0087	0.0210	0.0893	0.0033		-0.0010	-0.0009					-0.0008	
	17.09	2.00	15.44	-3.71	0.28	6.39	3.29		-0.87	-0.96					-0.47	
(2)																
	-0.0006		0.1018		0.0111	0.0265				0.3726	0.4684					
	-0.83		0.66		0.89	2.42										
	0.0511	0.0009	0.0398	-0.0023	-0.0075	0.0416	0.0016		-0.0009	-0.0006		-0.0094		0.0001	-0.0018	
	11.53	0.22	12.63	-0.70	-0.07	2.35	1.46		-0.71	-0.53		-3.46		0.20	-0.80	
(3)																
	-0.0019	0.0019	-0.0304	0.4230	-0.0088	0.0032	-0.0380	0.0831		0.3868	0.5064					
	-2.27	0.95	-0.16	1.46	-0.55	0.23	-0.29	5.72								
	0.0612	0.0040	0.0404	-0.0037	-0.0009	0.0804	0.0024		-0.0017	0.0001		-0.0114		0.0007	-0.0028	
	14.51	1.10	13.73	-1.28	-0.01	5.08	2.43		-1.40	0.07		-4.60		1.43	-1.51	
(4)																
	-0.0013	-0.0004	-0.0525	0.2392	-0.0031	0.0246	-0.1198			0.3792	0.4994					
	-1.60	-0.24	-0.31	0.92	-0.22	1.82	-1.04									

Table 1.8: Analysis of non-linearity in the relation between returns and factors, and among factors.

Column (1) reports the percentage of option portfolios in which the best-fitting time-series regression includes a squared term for the selected factor, among regressions in which each factor enters with powers 1 or 1 and 2. Column (2) shows whether the relation between the selected factor and other factors includes non-linear terms. Each factor is regressed on the others using a fractional polynomial model, in which factors can enter with any combination of the following powers: -2, -1, -0.5, 0, 0.5, 1, 2, 3, where 0 means that the log of the factor is included. The model is estimated over five time periods: the full sample and four subperiods (1996-98, 1999-2001, 2002-04, 2005-07). Column (2) reports whether any factor enters the best-fitting regression non-linearly in at least three of the five time periods. Only factors included in the *base case* specification are analyzed (see Table 1.6). In both columns (1) and (2), the best fitting model is determined as the one with the lowest $D = n(1 + ln\frac{2\pi RSS}{n})$, where n is the number of observations and RSS the residual sum of squares.

Factor	**(1)**	**(2)**
Mkt	28.43	Vix, with power=2
Smb	1.52	no
Hml	7.61	no
Umd	10.66	.
Liq	2.54	.
Vix	0.51	no
Vix_v	1.02	no
Vix_s	0.51	.
Skew	0.51	no
Kurt	0.51	.
Put	0	no
Put_v	1.52	no
Put_s	0	.
EU	2.03	no

Table 1.9: Risk-premia estimates, with controls for measurement error and hedging risk, selected squared factors. Second stage regressions estimated with WLS.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4 and 1.7 for factor and control definitions. Mkt^2, Hml^2, Umd^2 and Vix^2 are squared Mkt, Hml, Umd and Vix factors.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}^2_1$	$\bar{R}^2_2$					
(1)	0.0324	-0.0013	0.0276	-0.0076	0.0024	-0.0165	0.0016	0.0006	-0.0005	-0.0002	0.0005				-0.0025	0.0000
	6.72	-0.26	8.71	-2.36	0.02	-0.79	1.44	5.13	-0.34	-0.12	5.84				-1.03	0.15
	-0.0019		-0.0763		-0.0022	0.0266		0.0967		0.3862	0.5039					
	-1.89		-0.33		-0.12	1.51		5.67								
(2)	0.0420	0.0016	0.0270	-0.0092	0.0106	0.0215	0.0024	0.0007	-0.0010	0.0008	0.0005				-0.0032	0.0001
	9.40	0.35	8.64	-3.00	0.10	1.13	2.34	5.53	-0.69	0.67	6.16				-1.40	0.55
	-0.0016		-0.1503		-0.0076	0.0429				0.3788	0.4966					
	-1.60		-0.67		-0.46	2.57										
(3)	0.0351	-0.0028	0.0308	-0.0065	0.0278	-0.0158	0.0022	0.0006	0.0007	-0.0012	0.0004	-0.0033	0.0009	0.0012	-0.0018	-0.0001
	6.92	-0.61	9.58	-1.80	0.25	-0.80	1.93	5.19	0.45	-0.94	5.08	-1.13	5.74	1.80	-0.75	-0.50
	-0.0018	0.0015	-0.0687	-0.2596	-0.0066	0.0261	0.0339	0.0826		0.3945	0.5337					
	-1.82	0.67	-0.31	-0.79	-0.37	1.39	0.25	5.24								
(4)	0.0420	-0.0005	0.0301	-0.0082	0.0241	0.0115	0.0029	0.0006	0.0002	-0.0006	0.0004	-0.0040	0.0010	0.0016	-0.0027	-0.0001
	8.04	-0.12	8.97	-2.28	0.22	0.58	2.80	5.27	0.16	-0.48	5.13	-1.29	6.02	2.38	-1.18	-0.34
	-0.0015	0.0000	-0.1632	-0.5916	-0.0043	0.0451	-0.0272			0.3871	0.5297					
	-1.42	0.02	-0.72	-1.66	-0.24	2.31	-0.19									

Table 1.10: Risk-premia estimates, with controls for measurement error and hedging risk, selected squared factors. Second stage regressions include moneyness dummies.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. Second stage regressions include dummies for moneyness categories 1 to 3, interacted with put/call type (dummy_{ij}, i=MonDummy(1,2,3),j=put,call).

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix$_v$	Vix$_s$	Skew	Kurt	Put	Put$_v$	Put$_s$	EU		$\bar{R}^2_1$	$\bar{R}^2_2$					
(1)	0.0081	0.0009	0.0333	0.0223	0.0360	-0.0069	0.0023	0.0004	-0.0007	-0.0002	0.0002				-0.0032	0.0001
	1.92	0.24	10.90	5.66	0.42	-0.40	2.12	3.89	-0.62	-0.18	3.80				-1.68	0.69
	-0.0024		-0.0491		-0.0165	-0.0123		0.0380		0.3862	0.5046					
	-3.13		-0.27		-1.11	-0.85		2.98								
(2)	0.0089	0.0020	0.0324	0.0227	0.0293	0.0027	0.0026	0.0004	-0.0010	0.0002	0.0002				-0.0035	0.0001
	2.20	0.55	10.76	5.99	0.35	0.16	2.48	4.18	-0.85	0.22	3.71				-1.88	0.54
	-0.0024		-0.0966		-0.0186	-0.0084				0.3788	0.4995					
	-3.01		-0.52		-1.33	-0.59										
(3)	0.0109	0.0000	0.0360	0.0232	0.0934	-0.0042	0.0028	0.0004	0.0001	-0.0008	0.0002	0.0015	0.0004	0.0002	-0.0023	0.0001
	2.47	0.01	11.77	5.82	1.07	-0.26	2.49	3.80	0.08	-0.72	3.17	0.63	3.45	0.32	-1.27	0.84
	-0.0020	0.0010	0.0127	-0.2744	-0.0180	-0.0119	-0.0110	0.0306		0.3945	0.5384					
	-2.69	0.61	0.07	-1.04	-1.20	-0.77	-0.10	2.55								
(4)	0.0120	0.0012	0.0358	0.0234	0.0708	0.0046	0.0030	0.0004	-0.0001	-0.0006	0.0002	0.0012	0.0004	0.0003	-0.0025	0.0001
	2.63	0.33	11.29	5.87	0.81	0.28	2.81	3.93	-0.05	-0.54	3.10	0.47	3.56	0.52	-1.37	0.33
	-0.0022	0.0008	-0.0358	-0.2562	-0.0178	-0.0070	0.0150			0.3871	0.5348					
	-2.81	0.43	-0.20	-0.93	-1.22	-0.45	0.14									

Table 1.11: Risk factor loadings and bid-ask spread revenue.

The table shows the results of regressing a measure of bid-ask revenue on the factor loadings of option portfolio returns. The coefficients are time-series averages of cross-sectional regressions. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor definitions. The dependent variable of the cross-sectional regressions, which is calculated every Tuesday, is $\text{BArev}_{t,1}=\frac{1}{2}\frac{0.5\cdot(\text{ask}_t-\text{bid}_t)\cdot\text{volume}_t}{\text{oi}_t\cdot\text{oprice}_t}$ in specifications (1)-(2) and $\text{BArev}_{t,2}=\frac{1}{2}\frac{0.5\cdot(\text{ask}_t-\text{bid}_t)\cdot\text{volume}_t}{\text{oi}_{t-1}\cdot\text{oprice}_t}$ in (3)-(4). The numerator is half the bid-ask spread, which is the theoretical revenue on each trade if the true option price is the bid-ask mid-point, times trading volume. The denominator is a proxy for the value of net holdings. The theoretical revenue is multiplied by $\frac{1}{2}$ to account for fixed and inventory management costs. In (5) and (6) ($\text{BArev}_{t,1}$ and $\text{BArev}_{t,2}$, respectively) the theoretical revenue is multiplied by $\frac{1}{4}$. Coefficients are expressed as percentages.

	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Vix	Vix^2	Vix$_v$	Skew	Put	Put$_v$	EU
(1)	0.0924	-0.0006	0.0008	0.0078	-0.0148	-0.0001	0.0504	-0.0003	0.0106	-0.674	-0.0004	0.2099	0.5507
	24.41	-0.24	6.30	5.94	-6.98	-0.67	11.86	-1.41	4.43	-1.72	-0.01	5.13	18.59
(2)	0.1360	0.0075	0.0009	0.0038	-0.0073	0.0006	0.0328	0.0000	0.0129	-1.3197	0.1313	0.3814	
	33.77	2.37	7.23	2.29	-2.59	4.99	9.06	-0.09	5.24	-3.43	1.38	8.15	
(3)	0.1223	-0.0023	0.0012	0.0101	-0.0173	0.0000	0.0700	-0.0002	0.0310	0.3734	-0.1040	0.4261	0.7842
	20.29	-0.64	4.41	5.67	-5.07	0.04	9.40	-0.54	6.18	0.72	-0.85	7.20	17.98
(4)	0.1836	0.0151	0.0007	0.0035	-0.0154	0.0004	0.0545	0.0000	0.0174	-1.3798	0.3038	0.6493	
	31.34	4.08	4.12	2.07	-3.85	2.58	8.56	0.16	4.74	-2.55	2.88	13.36	
(5)	0.0005	0.0000	0.0000	0.0000	-0.0001	0.0000	0.0003	0.0000	0.0001	-0.0034	0.0000	0.0010	0.0028
	24.41	-0.24	6.30	5.94	-6.98	-0.67	11.86	-1.41	4.43	-1.72	0.00	5.13	18.59
(6)	0.0006	0.0000	0.0000	0.0001	-0.0001	0.0000	0.0003	0.0000	0.0002	0.0019	-0.0005	0.0021	0.0039
	20.29	-0.64	4.41	5.67	-5.07	0.04	9.40	-0.54	6.18	0.72	-0.85	7.20	17.98

Table 1.12: Risk-premia estimates, marginal contribution of the five macroeconomic announcements.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. The coefficients of the controls for measurement error and hedging risk are omitted. Second stage regressions are estimated using WLS. In (CPI) the macroeconomic uncertainty factor is defined using announcements about the Consumer Price Index, in (PPI) about the Producer Price Index, in (EMPL) about the Employment situation, in (PRCO) about Productivity and Costs, in (RE) about Real Earnings.

	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Vix	Vix^2	Vix_v	Skew	Put	Put_v	EU	$\bar{R}_1^2$	$\bar{R}_2^2$
	0.0215	0.0024	0.0007	-0.0010	0.0008	0.0005	-0.0032	0.0001	-0.0016	-0.1503	-0.0076	0.0429		0.3788	0.4966
	1.13	2.34	5.53	-0.69	0.67	6.16	-1.40	0.55	-1.60	-0.67	-0.46	2.57			
(CPI)	-0.0121	0.0017	0.0006	-0.0005	-0.0008	0.0004	-0.0004	-0.0001	-0.0014	-0.0562	0.0022	0.0336	0.1087	0.3847	0.5027
	-0.59	1.70	5.02	-0.33	-0.58	5.37	-0.16	-0.31	-1.32	-0.22	0.11	1.90	6.40		
(EMPL)	-0.0053	0.0016	0.0007	-0.0008	0.0003	0.0005	-0.0033	0.0002	-0.0016	-0.0323	0.0059	0.0316	0.0758	0.3884	0.5023
	-0.26	1.59	5.46	-0.56	0.27	6.01	-1.45	0.93	-1.60	-0.15	0.33	1.93	5.56		
(PPI)	0.0054	0.0019	0.0006	-0.0017	0.0003	0.0005	-0.0035	0.0000	-0.0018	-0.1643	-0.0016	0.0331	0.0693	0.3869	0.5040
	0.26	1.77	4.98	-1.14	0.26	6.12	-1.47	0.05	-1.82	-0.73	-0.09	1.92	4.97		
(PRCO)	0.0116	0.0020	0.0005	-0.0020	-0.0010	0.0004	-0.0034	0.0000	-0.0011	-0.0553	0.0103	0.0410	0.0698	0.3801	0.5020
	0.60	1.90	4.64	-1.39	-0.79	5.89	-1.42	0.28	-1.12	-0.25	0.58	2.57	4.84		
(RE)	-0.0148	0.0017	0.0006	-0.0005	-0.0008	0.0004	-0.0003	-0.0001	-0.0014	-0.0714	0.0049	0.0335	0.1156	0.3844	0.5030
	-0.70	1.72	4.99	-0.30	-0.57	5.31	-0.11	-0.38	-1.34	-0.28	0.25	1.86	6.45		

Table 1.13: Risk-premia estimates. Option returns computed using bid/ask prices.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. Second stage regressions are estimated using WLS. Option returns are computed by selling at the bid and buying at the ask.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
(1)	-0.0415	0.0066	0.0433	0.0118	0.0685	-0.2366	0.0018	0.0005	0.0004	0.0006	0.0004				-0.0048	0.0001
	-10.12	1.64	16.58	4.25	0.76	-13.64	1.75	4.21	0.34	0.55	5.74				-2.32	0.47
	-0.0024		-0.1394		0.0025	0.0102		0.0511		0.3839	0.5053					
	-2.84		-0.74		0.17	0.68		3.35								
(2)	-0.0372	0.0087	0.0432	0.0114	0.0677	-0.2157	0.0021	0.0005	0.0001	0.0012	0.0004				-0.0049	0.0001
	-9.64	2.21	16.45	4.22	0.78	-13.36	2.11	4.59	0.05	1.09	6.04				-2.45	0.48
	-0.0020		-0.1675		-0.0036	0.0194				0.3770	0.4996					
	-2.40		-0.87		-0.26	1.30										
(3)	-0.0382	0.0052	0.0441	0.0110	0.0948	-0.2355	0.0027	0.0005	0.0012	-0.0003	0.0003	0.0003	0.0008	0.0007	-0.0047	0.0001
	-8.45	1.24	14.80	3.75	0.97	-13.44	2.61	4.49	0.85	-0.23	4.86	0.12	6.34	1.17	-2.08	0.60
	-0.0021	0.0002	-0.1170	-0.7936	0.0102	0.0184	-0.1568	0.0341		0.3922	0.5360					
	-2.25	0.11	-0.58	-2.81	0.68	1.04	-1.32	2.29								
(4)	-0.0366	0.0065	0.0440	0.0108	0.0796	-0.2257	0.0028	0.0005	0.0010	-0.0001	0.0003	0.0000	0.0008	0.0009	-0.0047	0.0001
	-7.70	1.58	14.37	3.65	0.83	-12.81	2.75	4.55	0.72	-0.08	5.00	-0.02	6.16	1.38	-2.12	0.43
	-0.0017	-0.0004	-0.1605	-0.9549	0.0068	0.0254	-0.1294			0.3855	0.5317					
	-1.89	-0.20	-0.81	-3.19	0.46	1.41	-1.05									

Table 1.14: Risk-premia estimates. Robustness checks I.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. Second stage regressions are estimated using WLS. (1) and (2): returns for options with a valid price at date t-1 but not at date t are set equal to -50% and 0%. (3): the EU factor is computed by taking the mean rather than the median of option portfolio returns. (4): first-stage regressions are based on the Cochrane-Orcutt procedure. (5) to (7): sub-period analysis, including years 1996 to 1999, 2000 to 2003 and 2004 to 2007.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix$_v$	Vix$_s$	Skew	Kurt	Put	Put$_v$	Put$_s$	EU		$\bar{R}^2_1$	$\bar{R}^2_2$					
(1)	0.0473	-0.0093	-0.0085	0.0087	0.0650	-0.0190	0.0036	0.0009	-0.0028	0.0001	0.0003				0.0045	0.0000
	12.45	-2.06	-2.07	2.15	0.50	-0.92	3.06	6.36	-1.68	0.09	4.41				1.83	-0.04
	-0.0011		-0.5563		-0.0121	-0.0021		0.0711		0.2265	0.3975					
	-1.22		-3.28		-0.72	-0.14		4.42								
(2)	0.0247	0.0014	0.0256	-0.0055	0.0045	-0.0105	0.0019	0.0007	-0.0006	0.0000	0.0005				-0.0022	0.0001
	8.53	0.44	11.33	-1.98	0.05	-0.73	1.55	5.54	-0.43	0.02	5.55				-0.92	0.70
	-0.0024		0.0993		-0.0089	0.0041		0.0758		0.3084	0.4644					
	-2.56		0.54		-0.48	0.25		3.94								
(3)	0.0332	-0.0008	0.0278	-0.0069	0.0097	-0.0099	0.0017	0.0006	-0.0008	-0.0004	0.0005				-0.0037	0.0000
	6.80	-0.16	8.76	-2.10	0.11	-0.49	1.51	5.04	-0.52	-0.31	5.95				-1.51	0.03
	-0.0018		-0.1660		-0.0058	0.0274		0.0663		0.3895	0.5036					
	-1.89		-0.73		-0.31	1.61		5.31								
(4)	0.0544	0.0004	0.0438	0.0011	0.0622	0.0537	0.0026	0.0002	-0.0007	-0.0002	0.0000				-0.0035	-0.0002
	11.85	0.10	12.83	0.34	0.83	3.15	3.15	2.18	-0.70	-0.20	1.44				-2.28	-2.41
	0.0004		0.1097		-0.0292	0.0347		0.0819			0.4957					
	0.52		0.63		-2.34	2.50		4.99								

Continued

Table 1.14, continued.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
	0.0612	-0.0053	0.0395	-0.0057	0.0510	0.0189	0.0058	0.0002	-0.0021	-0.0038	0.0002				-0.0042	0.0003
	8.75	-0.60	7.22	-0.97	0.33	0.54	3.27	2.29	-1.09	-2.04	3.91				-1.88	2.23
(5)																
	0.0017		-0.1404		-0.0128	0.0161		0.0520		0.3189	0.5098					
	3.74		-0.81		-0.84	1.76		2.55								
	0.0551	0.0041	0.0378	-0.0085	0.0674	0.0483	0.0002	0.0007	0.0014	-0.0006	0.0005				0.0016	0.0006
	8.76	0.60	6.16	-1.32	0.42	1.93	0.10	3.19	0.64	-0.29	4.11				0.89	1.88
(6)																
	-0.0014		0.4147		0.0374	0.0109		0.0646		0.4242	0.5676					
	-2.72		6.26		2.11	1.19		3.74								
	0.0280	0.0002	0.0290	-0.0058	0.0473	-0.0099	0.0018	0.0000	0.0007	0.0009	0.0000				-0.0022	-0.0001
	6.93	0.06	12.64	-2.61	0.42	-0.57	1.74	0.56	0.71	1.08	2.52				-2.05	-0.59
(7)																
	-0.0008		-0.0237		-0.0151	0.0056		0.0351		0.4624	0.5578					
	-1.11		-0.15		-1.45	0.66		3.32								

Table 1.15: Risk-premia estimates. Robustness checks II.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. Second stage regressions are estimated using WLS. (1) and (2): the EU factor is computed by comparing pricing errors two days before announcements and on all other days, and one day before announcements and on announcement days. (3)-(4) and (5)-(6): as in (2), but announcements are considered only if set at least 3 and 7 days apart. (7)-(8): options on stocks that pay dividends before the expiration date are excluded.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}^2_1$	$\bar{R}^2_2$					
(1)	0.0354	-0.0003	0.0289	-0.0070	0.0325	-0.0050	0.0021	0.0006	-0.0014	0.0002	0.0005				-0.0032	0.0000
	7.88	-0.06	9.27	-2.19	0.38	-0.27	1.94	5.06	-0.97	0.15	6.04				-1.34	0.08
	-0.0014		-0.0597		-0.0016	0.0383		0.0893		0.3867	0.5039					
	-1.47		-0.26		-0.09	2.32		5.77								
(2)	0.0328	0.0007	0.0291	-0.0047	-0.0587	-0.0003	0.0022	0.0006	-0.0009	-0.0005	0.0005				-0.0030	0.0000
	6.88	0.16	9.40	-1.48	-0.59	-0.02	2.03	4.56	-0.63	-0.41	5.97				-1.29	-0.06
	-0.0014		-0.0129		0.0087	0.0289		0.0933		0.3863	0.5013					
	-1.46		-0.06		0.49	1.81		6.07								
(3)	0.0340	-0.0011	0.0284	-0.0059	-0.0046	-0.0063	0.0022	0.0007	-0.0011	-0.0008	0.0005				-0.0014	0.0001
	7.43	-0.26	9.52	-1.91	-0.05	-0.33	2.01	5.32	-0.77	-0.59	5.96				-0.65	0.82
	-0.0011		-0.1457		0.0020	0.0345		0.0852		0.3866	0.5035					
	-1.15		-0.63		0.12	2.20		5.79								
(4)	0.0356	-0.0032	0.0315	-0.0046	0.0129	-0.0102	0.0027	0.0006	0.0001	-0.0017	0.0004	-0.0035	0.0009	0.0013	-0.0008	0.0000
	7.38	-0.75	10.35	-1.34	0.13	-0.57	2.57	5.16	0.06	-1.33	5.19	-1.23	6.18	2.05	-0.36	0.01
	-0.0012	0.0004	-0.0979	-0.2638	-0.0035	0.0320	0.0645	0.0710		0.3951	0.5350					
	-1.29	0.18	-0.45	-0.84	-0.21	1.87	0.47	4.52								

Continued

Table 1.15, continued.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
(5)	0.0360	-0.0005	0.0290	-0.0064	0.0042	-0.0008	0.0021	0.0007	-0.0012	-0.0004	0.0005				-0.0027	0.0002
	7.81	-0.12	9.59	-2.01	0.04	-0.04	2.05	5.33	-0.84	-0.32	6.05				-1.22	0.91
	-0.0013		-0.1454		0.0022	0.0317		0.0692		0.3884	0.5030					
	-1.33		-0.65		0.13	1.93		5.10								
(6)	0.0379	-0.0026	0.0321	-0.0062	0.0354	-0.0075	0.0026	0.0006	0.0002	-0.0014	0.0004	-0.0032	0.0010	0.0014	-0.0018	0.0000
	7.43	-0.61	9.95	-1.74	0.33	-0.42	2.61	4.97	0.14	-1.10	5.19	-1.07	6.10	2.22	-0.83	-0.09
	-0.0011	0.0002	-0.1058	-0.4401	-0.0026	0.0352	0.0133	0.0565		0.3967	0.5343					
	-1.09	0.08	-0.47	-1.33	-0.15	1.91	0.10	3.87								
(7)	0.0351	0.0057	0.0362	-0.0025	0.0323	0.0140	0.0013	0.0007	-0.0002	-0.0033	0.0004				0.0025	-0.0002
	5.79	0.80	6.14	-0.41	0.25	0.48	1.02	4.18	-0.09	-1.79	4.39				0.83	-0.95
	-0.0022		-0.1579		-0.0062	0.0332		0.1157		0.3553	0.4656					
	-1.54		-0.58		-0.29	1.52		3.81								
(8)	0.0387	0.0062	0.0357	-0.0028	0.0597	0.0312	0.0014	0.0009	0.0011	-0.0025	0.0006				0.0012	-0.0001
	6.04	0.84	5.51	-0.42	0.43	1.06	1.11	4.78	0.59	-1.23	4.91				0.37	-0.50
	-0.0021		-0.4287		-0.0175	0.0732				0.3465	0.4613					
	-1.37		-1.42		-0.79	3.10										

Table 1.16: Risk-premia estimates. Robustness checks III.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. Second stage regressions are estimated using WLS. (1)-(2): if an option's implied volatility is missing on day t, it is set equal to the implied volatility on day t-1 or t-2. The remaining specifications report estimated risk-premia for alternative definitions of the expected objective volatility used in computing EPEs (excess pricing errors). (3)-(4): previous month's realized volatility, excluding the current day. (5)-(6): previous month's realized volatility, excluding days with returns equal to zero. (7)-(8): for options with MatDummy=1 (=2) expected volatility is the realized volatility in the subsequent 2 (7) months.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix$_v$	Vix$_s$	Skew	Kurt	Put	Put$_v$	Put$_s$	EU		$\bar{R}^2_1$	$\bar{R}^2_2$					
	0.0311	-0.0045	0.0240	-0.0076	0.0070	-0.0226	0.0013	0.0007	-0.0007	0.0003	0.0005				-0.0031	0.0001
	6.84	-0.96	7.95	-2.46	0.06	-1.17	1.15	5.51	-0.47	0.25	6.03				-1.25	0.85
(1)																
	-0.0011		-0.0012		-0.0143	0.0299		0.0928		0.3868	0.5038					
	-1.07		0.00		-0.82	1.74		5.77								
	0.0396	-0.0020	0.0238	-0.0086	0.0276	0.0125	0.0019	0.0007	-0.0012	0.0013	0.0005				-0.0036	0.0002
	8.93	-0.46	7.89	-2.83	0.26	0.70	1.88	5.82	-0.85	1.07	6.41				-1.51	0.98
(2)																
	-0.0008		-0.1299		-0.0240	0.0444				0.3798	0.4970					
	-0.81		-0.55		-1.50	2.64										
	0.0303	-0.0001	0.0271	-0.0065	-0.0122	-0.0126	0.0016	0.0007	-0.0001	0.0000	0.0005				-0.0026	0.0000
	6.18	-0.03	8.35	-1.90	-0.13	-0.62	1.50	5.26	-0.09	0.04	5.89				-1.04	-0.07
(3)																
	-0.0018		0.0475		0.0003	0.0332		0.0903		0.3856	0.5045					
	-1.77		0.21		0.01	1.96		5.85								
	0.0421	0.0024	0.0272	-0.0091	-0.0288	0.0248	0.0024	0.0007	-0.0010	0.0007	0.0005				-0.0031	0.0001
	9.53	0.52	8.66	-2.90	-0.32	1.29	2.35	5.45	-0.70	0.61	6.18				-1.36	0.57
(4)																
	-0.0016		-0.1878		-0.0039	0.0406				0.3786	0.4954					
	-1.57		-0.80		-0.23	2.52										

Continued

Table 1.16, continued.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
	0.0339	-0.0004	0.0280	-0.0080	0.0052	-0.0114	0.0018	0.0007	-0.0004	-0.0010	0.0005				-0.0025	0.0000
	6.50	-0.08	8.12	-2.12	0.05	-0.51	1.83	5.02	-0.26	-0.74	5.44				-1.04	-0.18
(5)																
	-0.0023		0.0940		-0.0147	0.0257		0.0891		0.3913	0.5073					
	-2.44		0.38		-0.82	1.36		5.28								
	0.0434	0.0022	0.0276	-0.0097	0.0107	0.0227	0.0027	0.0007	-0.0008	0.0000	0.0005				-0.0026	0.0001
	9.19	0.48	8.57	-2.81	0.10	1.17	2.65	5.13	-0.55	-0.04	5.74				-1.18	0.44
(6)																
	-0.0015		0.1078		-0.0089	0.0433				0.3814	0.4989					
	-1.67		0.43		-0.54	2.50										
	0.0321	-0.0008	0.0317	-0.0034	-0.0534	-0.0097	0.0023	0.0005	-0.0005	-0.0008	0.0004				-0.0032	0.0000
	6.65	-0.18	10.67	-1.06	-0.58	-0.48	2.37	4.43	-0.36	-0.59	5.21				-1.37	-0.05
(7)																
	-0.0018		-0.0396		0.0308	0.0176		0.0972		0.3845	0.5039					
	-1.59		-0.17		1.74	1.14		6.02								
	0.0423	0.0019	0.0291	-0.0071	-0.0406	0.0257	0.0032	0.0006	-0.0011	0.0001	0.0005				-0.0019	0.0001
	9.38	0.41	9.41	-2.22	-0.44	1.32	3.28	4.84	-0.76	0.08	6.11				-0.79	0.66
(8)																
	-0.0017		-0.2182		0.0192	0.0385				0.3751	0.4946					
	-1.56		-0.93		1.11	2.52										

Table 1.17: Determinants of the probability of missing option returns.

A dummy variable is set equal to one if option portfolio returns are missing, and to zero otherwise. The dummy is the dependent variable of a logit model, in which the independent variables are a year dummy, the natural logarithm of trading volume, open interest and relative option bid-ask spread, and the following factors: Mkt, Smb, Hml, Vix, Vix_v, Skew, Put, Put_v and EU. Trading volume, open interest and option spread are the median across call/put type (1) and for each portfolio (2). A probit model gives similar odds-ratios and an identical pattern of statistical significance. The table reports the median odds-ratio, across all option portfolios, and the 90% confidence interval, calculated using the binomial method of Mood and Graybill (1963).

	(1)			(2)		
Factor	**Median**	90% C.I.		**Median**	90% C.I.	
year	0.745*	0.720	0.764	0.764*	0.737	0.789
log-volume	0.861*	0.808	0.918	1.021	0.979	1.063
log-op.interest	1.736*	1.341	2.204	1.175*	1.129	1.247
log-spread	0.949	0.683	1.297	0.799*	0.658	0.905
Mkt	0.015*	0.001	0.157	0.015*	0.001	0.072
Smb	0.001*	0.000	0.013	0.002*	0.000	0.011
Hml	0.000*	0.000	0.000	0.000*	0.000	0.001
Vix	0.074*	0.014	0.608	0.038*	0.009	0.306
Vix_v	1.993	0.207	44.208	0.590	0.089	7.244
Skew	1.006	0.990	1.021	1.005	0.992	1.014
Put	1.052	0.861	1.306	0.881	0.740	1.217
Put_v	1.128	0.882	1.354	1.054	0.814	1.321
EU	2.046*	1.749	2.562	2.111*	1.815	2.785

Table 1.18: Determinants of changes in the number of available observations.

The weekly change in the number of options with the indicated characteristic is regressed on selected factors and variables. ln-o.s., ln-v. and ln-o.i. are the natural logarithm of the median (over the whole cross-section of options, on each Tuesday) relative option spread, trading volume and open interest. The number of observation is calculated on the basis of options that satisfy the filters to be included in the computation of returns. The subscripts indicate statistical significance at the 10% (1), 5% (2) and 1% (3) level. See Table 1.4 for factor definitions.

		ln-o.s.	ln-v.	ln-o.i.	Mkt	Smb	Hml	Vix	Vix_v	Skew	Put	Put_v	EU	α
Pmkt	**1**	0.015	0.428^{3}	-0.242^{3}	-0.536	1.281^{1}	0.818	-2.356^{3}	0.853	-0.011	0.121	-0.121	-0.204^{3}	0.296
	2	-0.019	0.454^{3}	-0.254^{3}	-0.239	1.181^{1}	0.711	-2.426^{3}	2.012	-0.009	0.159^{1}	-0.137	-0.158^{3}	0.210
	3	-0.033	0.439^{3}	-0.253^{3}	0.115	1.692^{3}	0.741	-2.376^{3}	1.467	-0.010	0.156^{1}	-0.125	-0.145^{2}	0.216
Psmb	**1**	-0.015	0.422^{3}	-0.238^{3}	-0.348	1.071^{1}	0.923	-2.203^{3}	1.428	-0.008	0.153^{1}	-0.143^{1}	-0.165^{2}	0.212
	2	-0.028	0.453^{3}	-0.262^{3}	0.162	2.065^{3}	0.903	-2.721^{3}	1.701	-0.013	0.144	-0.129	-0.201^{3}	0.246
	3	-0.021	0.497^{3}	-0.281^{3}	0.732	2.139^{3}	0.143	-2.771^{3}	1.921	-0.016^{1}	0.138	-0.072	-0.134	0.257
Phml	**1**	-0.011	0.436^{3}	-0.249^{3}	-0.147	1.406^{2}	0.597	-2.342^{3}	1.741	-0.008	0.151^{1}	-0.135^{1}	-0.181^{3}	0.253
	2	-0.048	0.433^{3}	-0.249^{3}	-0.033	1.719^{2}	1.584^{1}	-2.208^{3}	0.535	-0.01	0.125	-0.095	-0.152^{2}	0.171
	3	-0.056	0.463^{3}	-0.248^{3}	-0.555	1.704^{2}	1.320	-2.651^{3}	1.460	-0.019^{2}	0.177^{1}	-0.192^{1}	-0.012	0.043
Moneyness	**1**	-0.048	0.449^{3}	-0.253^{3}	-0.738	0.974	0.204	-1.816^{3}	0.059	-0.006	0.137	-0.071	-0.170^{2}	0.140
	2	-0.020	0.451^{3}	-0.244^{3}	0.092	1.349^{2}	0.743	-2.853^{3}	0.775	-0.014^{1}	0.190^{1}	-0.109	-0.217^{3}	0.142
	3	-0.014	0.465^{3}	-0.263^{3}	0.378	1.568^{2}	1.592^{1}	-2.963^{3}	1.216	-0.007	0.133	-0.141	-0.168^{2}	0.247
	4	0.000	0.394^{3}	-0.244^{3}	-0.189	1.812^{3}	0.736	-2.227^{3}	3.102^{2}	-0.012	0.154^{1}	-0.159^{1}	-0.153^{2}	0.399^{1}
	5	0.006	0.407^{3}	-0.246^{3}	-0.611	1.387^{2}	0.407	-1.801^{3}	3.866^{3}	-0.010	0.100	-0.218^{3}	-0.063	0.382
Maturity	**1**	-0.029	0.735^{3}	-0.383^{3}	-0.978	1.860	1.012	-4.381^{3}	-0.056	-0.024^{1}	0.278^{1}	-0.117	-0.355^{3}	0.145
	2	-0.018	0.112^{3}	-0.079^{3}	0.703	0.625	-0.118	0.132	2.652^{3}	0.007	-0.040	-0.094	0.076	0.143
Style	Call	-0.024	0.430^{3}	-0.239^{3}	0.739	1.334^{2}	0.806	-2.733^{3}	0.985	-0.004	0.187^{2}	-0.111	-0.128^{2}	0.169
	Put	-0.012	0.447^{3}	-0.264^{3}	-1.673^{2}	1.547^{2}	0.700	-1.877^{3}	2.308	-0.019^{2}	0.091	-0.166^{2}	-0.227^{3}	0.329

Table 1.19: Risk-premia estimates, accounting for non-randomly missing returns.

The table shows second-stage coefficients ($\hat{\lambda}$) and t-stats. Standard errors are adjusted according to Shanken and Newey-West (4 lags). See Tables 1.4, 1.7 and 1.9 for factor and control definitions. First stage estimates are obtained through the Heckman model, in which the selection equation includes the following factors and variables: Mkt, Smb, Hml, Vix, EU, a year dummy, and median log-trading volume, log-open interest and log-option spread. The median is calculated across call/put type [(1)-(2)] and for each portfolio [(3)-(4)]. The choice of variables and factors to include in the selection equation is based on the results reported in Table 1.17. Second stage regressions are estimated using WLS.

	os	ss	ν	γ	β_s	Cons	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2
	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU		$\bar{R}_1^2$	$\bar{R}_2^2$					
(1)	0.0364	-0.0010	0.0288	-0.0086	0.0191	-0.0084	0.0020	0.0006	-0.0009	-0.0002	0.0005				-0.0023	-0.0001
	7.41	-0.20	9.30	-2.71	0.18	-0.40	1.86	4.77	-0.63	-0.17	5.75				-1.07	-0.37
	-0.0017		-0.0522		0.0095	0.0325		0.0817			0.5027					
	-1.85		-0.25		0.52	1.99		5.11								
(2)	0.0432	0.0020	0.0294	-0.0076	0.0180	0.0277	0.0030	0.0006	-0.0014	0.0006	0.0005				-0.0027	0.0000
	9.89	0.46	9.23	-2.44	0.17	1.50	2.78	4.82	-0.97	0.52	5.85				-1.28	0.19
	-0.0019		-0.1381		0.0124	0.0463					0.4950					
	-2.00		-0.66		0.73	2.91										
(3)	0.0352	0.0009	0.0297	-0.0063	0.0007	-0.0005	0.0020	0.0006	-0.0010	-0.0003	0.0005				-0.0025	0.0000
	7.27	0.19	9.28	-1.88	0.01	-0.03	1.87	4.78	-0.68	-0.26	5.68				-1.13	0.01
	-0.0021		-0.0534		0.0078	0.0321		0.0836			0.5030					
	-2.16		-0.25		0.41	1.86		5.26								
(4)	0.0453	0.0027	0.0277	-0.0102	0.0220	0.0312	0.0028	0.0006	-0.0012	0.0005	0.0005				-0.0030	0.0001
	10.51	0.62	8.90	-3.33	0.21	1.70	2.68	4.91	-0.85	0.42	5.86				-1.47	0.68
	-0.0017		-0.1335		0.0045	0.0495					0.4942					
	-1.78		-0.62		0.27	3.05										

Table 1.20: Factor contributions to expected excess returns, across beta percentiles.

The table shows factor risk premia multiplied by the 25^{th}, 50^{th} and 75^{th} percentile beta. Dots indicate that the factor is not included in the reported specifications. Zero entries mean that either the corresponding risk premium is non-statistically significant, or that the risk premium times the beta is statistically equal to zero. Second stage regressions are estimated with WLS, with the exception of specifications (1) and (2), which include dummies for moneyness categories 1 to 3, interacted with put/call type (dummy_{ij}, i=MonDummy(1,2,3),j=put,call). (1) to (5): expected volatility is realized volatility over the previous 30 days. (5): first stage regression based on the Heckman model. (6): if an option's implied volatility is missing on day t, it is set equal to the implied volatility on day t-1 or t-2. (7): previous month's realized volatility, excluding the current day. (8): previous month's realized volatility, excluding days with returns equal to zero. (9): for options with MatDummy=1 (=2), expected volatility is the realized volatility in the subsequent 2 (7) months. (10) and (11): the EU factor is computed by comparing pricing errors two days before announcements and on all other days, and one day before announcements and on announcement days.

	p(β)	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU
(1)	0.25	-56.10	-21.34	0	0	-30.94	.	.	.	4.54	0	11.64	.	0	.	0	0	.	2.38
	0.50	24.85	6.41	0	0	6.25	.	.	.	-6.39	0	0.07	.	0	.	0	0	.	33.06
	0.75	75.40	33.55	0	0	34.64	.	.	.	-17.25	0	-13.73	.	0	.	0	0	.	54.23
(2)	0.25	-66.93	-15.74	0	0	-18.08	0	-26.32	0	0	0	10.63	0	0	0	0	0	0	1.78
	0.50	27.84	6.55	0	0	5.75	0	1.99	0	0	0	-0.12	0	0	0	0	0	0	26.58
	0.75	95.91	36.64	0	0	26.66	0	19.35	0	0	0	-13.14	0	0	0	0	0	0	44.15
(3)	0.25	0	-35.57	0	0	-66.94	.	.	.	0	0	9.22	.	0	.	0	0	.	6.06
	0.50	0	10.69	0	0	13.53	.	.	.	0	0	0	.	0	.	0	0	.	84.23
	0.75	0	55.94	0	0	74.94	.	.	.	0	0	-10.88	.	0	.	0	0	.	138.16
(4)	0.25	-51.78	-25.79	0	0	-41.14	0	-53.41	-10.34	0	0	9.46	0	0	0	0	0	0	4.79
	0.50	21.54	10.74	0	0	13.09	0	4.05	2.15	0	0	0	0	0	0	0	0	0	71.68
	0.75	74.2	60.03	0	0	60.68	0	39.27	15.19	0	0	-11.7	0	0	0	0	0	0	119.09

Continued

Table 1.20, continued.

	$p(\beta)$	Mkt	Mkt^2	Smb	Hml	Hml^2	Umd	Umd^2	Liq	Vix	Vix^2	Vix_v	Vix_s	Skew	Kurt	Put	Put_v	Put_s	EU
(5)	0.25	-51.51	-30.83	0	0	-55.48	.	.	.	0	0	12.29	.	0	.	0	-8.01	.	7.83
	0.50	19.78	10.26	0	0	13.71	.	.	.	0	0	0.98	.	0	.	0	3.39	.	70.16
	0.75	69.18	47.75	0	0	76.38	.	.	.	0	0	-13.38	.	0	.	0	15.45	.	118.35
(6)	0.25	0	-34.93	0	0	-65.91	.	.	.	0	0	0	.	0	.	0	-8.09	.	8.23
	0.50	0	10.97	0	0	8.88	.	.	.	0	0	0	.	0	.	0	3.10	.	79.17
	0.75	0	58.99	0	0	72.11	.	.	.	0	0	0	.	0	.	0	13.78	.	126.2
(7)	0.25	0	-32.74	0	0	-75.53	.	.	.	0	0	9.80	.	0	.	0	-9.36	.	2.99
	0.50	0	10.77	0	0	10.08	.	.	.	0	0	0.27	.	0	.	0	3.70	.	76.21
	0.75	0	54.83	0	0	75.81	.	.	.	0	0	-10.25	.	0	.	0	17.79	.	140.18
(8)	0.25	-44.57	-46.59	0	0	-71.94	.	.	.	0	0	15.53	.	0	.	0	0	.	10.76
	0.50	14.14	12.47	0	0	11.82	.	.	.	0	0	-0.50	.	0	.	0	0	.	80.55
	0.75	59.22	60.24	0	0	77.40	.	.	.	0	0	-14.74	.	0	.	0	0	.	140.52
(9)	0.25	-56.27	-25.38	0	0	-61.40	.	.	.	0	0	0	.	0	.	-17.08	0	.	17.60
	0.50	26.05	12.27	0	0	1.86	.	.	.	0	0	0	.	0	.	0	0	.	96.65
	0.75	73.88	47.18	0	0	53.83	.	.	.	0	0	0	.	0	.	12.05	0	.	165.98
(10)	0.25	-50.00	-34.58	0	0	-74.28	.	.	.	0	0	0	.	0	.	0	-10.80	.	18.14
	0.50	23.11	12.73	0	0	1.29	.	.	.	0	0	0	.	0	.	0	4.40	.	76.47
	0.75	70.31	58.04	0	0	62.65	.	.	.	0	0	0	.	0	.	0	19.69	.	112.00
(11)	0.25	-52.73	-31.03	0	0	-71.32	.	.	.	0	0	0	.	0	.	0	-7.19	.	-5.71
	0.50	19.64	9.98	0	0	4.21	.	.	.	0	0	0	.	0	.	0	4.33	.	81.88
	0.75	69.97	46.70	0	0	61.06	.	.	.	0	0	0	.	0	.	0	14.95	.	138.15

Table 1.21: Market risk premium estimated on simulated option returns.

Fama-MacBeth estimated market risk premium, and second stage intercept. Confidence intervals are based on bootstrap standard errors (200 replications). Option returns are censored at 0.175% and option portfolio returns at 0.35%. 1,000 replications. The 95% confidence interval for $\hat{\lambda}_{\text{Mkt}}$ when Rp=0 includes 0.

Rp		**Average**	**90% C.I.**	
0.10	$\hat{\alpha}_2$	-.0003748	-.0024221	.0016726
	$\hat{\lambda}_{mkt}$	.0976612	.0944565	.1008659
0.00	$\hat{\alpha}_2$	.0005656	-.0015481	.0026793
	$\hat{\lambda}_{mkt}$	-.0036066	-.0067519	-.0004612

Chapter 2

Technology diffusion and the social value of investor sentiment

This chapter benefited from the suggestions of Francesca Cornelli, Theodosios Dimopoulos, Elroy Dimson, Carlo Fezzi, Tim Johnson, Chris Malloy, Stefano Sacchetto and Raman Uppal.

2.1 Introduction

When firms face the decision of whether to invest in a new technological vintage with uncertain productivity, they need to account for different types of costs. Besides the direct capital expenditure, and the need to retrain the workforce (Chari and Hopenhayn (1991)), early adopters likely generate informational externalities that are difficult to appropriate, and benefit competitors instead (Bolton and Harris (1999)). The difference between costs like retraining the workforce and unpriced externalities is one of substance. The diffusion of a new capital vintage may be slowed by technology-specific human capital (Jovanovic and Nyarko (1996)), but this simply reflects the presence of adjustment costs in the learning process. Unpriced externalities, on the other hand, are the result of a market failure, namely the inability of early adopters to appropriate the information that their initial investment reveals about the new technology's productivity. The consequence is a free rider problem, which normally implies that the aggregate level of investment falls short of the social optimum, and the resulting reduction in the flow of information can either delay the adoption of a valuable technology, delay the rejection of a bad one[1], or lead to rejecting a good or adopting a bad technology (Bolton and Harris (1999)).

The question at the heart of this paper is whether higher investor sentiment acts as a subsidy to investment in new technologies, in particular for firms that are more likely to experience informational externalities, and if such subsidy produces positive effects for the economy's productivity growth. Sentiment can be interpreted as the tendency of investors' decision making to deviate from rationality, in particular when forming beliefs. Some market participants may trade on noise rather than

[1] A "good"/"bad" new technology has a higher/lower productivity than the one currently in use.

fundamental news (Black (1986)), and rational investors may be unable to eliminate the resulting mispricing (De Long, Shleifer, Summers and Waldmann (1990)). Baker and Wurgler (2006) define sentiment as the "propensity to speculate", while in Dumas, Kurshev and Uppal (2008) sentiment is reflected in the overconfidence of some investors about a noisy public signal.

The sign of investor sentiment's contribution to aggregate productivity is key to evaluate whether the firm-level effect actually brings investment closer to the social optimum. If investor sentiment is to have a positive social value in the context of technology diffusion, it is essential that it offsets costs generated by a market failure, rather than costs reflecting relative scarcity. As an example, suppose that there are no informational externalities, and that investor sentiment increases the rate of adoption of a new, better technology because firms use the subsidy to offset the cost of lost productivity while retraining the workforce. At the aggregate level the additional investment is wasteful, because too much of the production is moved to a technology that, for the time being, is less productive. In other words, unless there is a slack generated by suboptimal investment, the subsidy provided by sentiment is unlikely to yield a positive aggregate benefit, because it simply mis-allocates resources. There are several ways in which investor sentiment can act as a subsidy to investment. First, it can reduce funding costs (Baker, Stein and Wurgler (2003), Farhi and Panageas (2004)). Second, managers may be willing to invest in projects that cater to investors' optimism about the new technology, in order to maximize stock prices in the short term (Polk and Sapienza (2009), Jensen (2005)). The firm-level empirical evidence is that stock mispricing explains both investment and future returns (Polk and Sapienza (2009)). At the aggregate level, Chirinko and Schaller (2001) support the hypothesis that sentiment affects investment. They analyse the Japanese stock market bubble of the late 1980s, and find that it increased business fixed investment by about 5%-10%. It is important to clarify that the focus of this paper is not on informational externalities at the *invention* stage, that is when a new technology is developed by a relatively small group of firms. I am rather interested in the externalities produced during the *adoption* stage, when the larger population of firms decides whether to invest in the new technology. The distinction is important because, as Johnson (2007) notes, patent protection and first-mover advantages can reduce the impact of informational externalities on technological innovation. This is probably true for technology *providers* during the *invention* process but, from the point of view of technology diffusion, the focus is on technology users, and externalities are generated while *learning* about uncertain productivity, in which case information is likely more difficult to appropriate.

The empirical analysis I present in section 2.3 investigates how the interaction between investor sentiment and technological innovation affects 1) firm-level investment, in particular for firms that are more susceptible to informational externalities, and 2) aggregate productivity, in the form of changes in the multi-factor productivity index, provided by the Bureau of Labor Statistics. I find that, when investor sentiment is higher, the effect of technological innovation on investment is greater for firms subject to informational externalities (in years t and $t+1$), and a positive effect on productivity growth is detectable after three years.

The paper is organized as follows. Section 2.2 is a literature review. Section 2.3 discusses the empirical implementation. Section presents 2.4 the results and Section 2.5 concludes.

2.2 Related literature

The hypothesis at basis of this paper is that investor sentiment has a positive social value in the context of technology diffusion, because it offsets the costs generated by a free-rider problem that arises when firms experiment with the new technology in order to learn its productivity. This brings together three strands of research, on learning by doing and informational externalities, on investor sentiment, and on the real effects of mispricing. The next three sections briefly review the relevant literature.

2.2.1 Uncertain productivity and learning

Learning by doing is an important ingredient of many vintage-capital investment models (Chari and Hopenhayn (1991), Jovanovic and Nyarko (1996), Cooley, Greenwood and Yorukoglu (1997)). It is typically used to explain the decrease in productivity that follows investment in a new technological vintage (Hugget and Ospina (2001)). As such, it is a variable that enters the problem of whether to invest or not. The decision maker has to factor in the upfront cost of investment, the expected period of learning, and the gains in productivity that the investment will likely generate. As long as the relevant variables are known, the decision maker needs to look no further than her own firm. However, it is reasonable to assume that the productivity of a new capital vintage is not precisely known. From this point of view, learning by doing is not only a cost to factor in, but also an opportunity to refine initial beliefs about the productivity. As long as the results of a firm's investment are known to outsiders, firms can learn about the productivity of a new technology by waiting and observing how well early adopters fare, rather than investing themselves. It is well known that the presence of unpriced externalities can reduce the production of a public good below the optimal level. When the productivity of a new technology is uncertain, strategic interaction among firms is then a key element of the diffusion process (Bhattacharya, Chatterjee and Samuelson (1986), Jovanovic and MacDonalds (1994), Reinganum (1983)). Bolton and Harris (1999) solve a continuous time two-armed bandit problem where firms decide, at each point in time, how much to produce with the old and with the new technology. They show that, in general, equilibrium experimentation is less than the socially optimal level. The literature on learning by doing and informational externalities is mostly theoretical, but there are a few empirical contributions that provide evidence of substantial learning effects. Zimmerman (1982) focuses on the construction of nuclear power plants, while Thornton and Thompson (2001) analyse wartime ship-building in the US.

2.2.2 Investor sentiment

Investor sentiment has increasingly been a subject of interest in behavioral finance, a research field that relaxes the paradigm of fully rational investors to explain anomalies that the standard asset pricing theory struggles to reconcile (see Barberis and Thaler (2003)). In a recent contribution, Baker and Wurgler (2006) build a proxy for investor sentiment from variables that the literature has linked to investors' irrational behavior, like the closed-end fund discount (Lee, Shleifer and Thaler (1991)), and find that stocks more attractive to speculators, like those of small, young firms, earn low future returns when the sentiment proxy is high, and viceversa, which is consistent with an over-/under-valuation generated by high/low sentiment. The efficient market hypothesis implies that any mispricing is quickly eliminated by arbitrageurs, but there are several reasons why this may not always be the case in practice. First, arbitrage is risky if the persistence of mispricing is uncertain (De Long, Shleifer, Summers and Waldmann (1990), Dumas, Kurshev and Uppal (2009)). Even if rational investors collectively have the resources to immediately eliminate deviations from fundamental values, coordination problems can delay arbitrage, and may actually provide incentives to trade on a continuation of mispricing (Abreu and Brunnermeier (2003)). Brunnermeier and Nagel (2004) and Lakonishok, Lee and Poteshman (2004) provide evidence that, during the late 1990s, sophisticated investors actually tried to accommodate what they perceived as mispricing, in order to realize short-term profits. Even if they had tried to arbitrage it out, it might have been difficult to do so. In IPOs, for instance, a large part of the floated capital is usually held in lock-up agreements for at least 6 months (Hong, Scheinkman and Xiong (2006)). In addition, it can be harder to take a stand against aggregate rather than single-firm mispricing, because of the large funds this requires. Arbitrageurs usually raise capital with open-ended structures, and need to take the performance-flow relation (Shleifer and Vishny (1997)) into account. They may find it difficult to raise capital through closed-end funds, even if it would be more efficient, because asymmetric information and agency problems usually make the early liquidation option valuable to investors.

2.2.3 The real effects of mispricing

The effect of mispricing on economic growth has been studied, from a theoretical point of view, by Olivier (2000), who suggests that bubbles can enhance growth because high stock valuations increase the creation of firms and investment. This conclusion, however, is unambiguous only for small open economies, where the interest rate is loosely dependent on domestic investment. For large economies, the increase in interest rates potentially offsets the effect of additional investment on growth. Farhi and Panageas (2004) suggest that mispricing can relax financing constraints, and its aggregate effect depends on the proportion of newly available funds that finance efficient versus wasteful investments. Their empirical analysis shows that the net macroeconomic effect is negative. More precisely, they estimate a vector auto-regression for NYSE aggregate profits and turnover (turnover is a function of mispricing, as in Scheinkman and Xiong (2003)), and the impulse response function of shocks to turnover on profits is negative. Polk and Sapienza (2009) focus on the cross-sectional relation between discretionary accruals and investment, in particular for firms where capital misallocation is more difficult to pin down (high R&D expenses) and whose shareholders more likely have a short investment horizon (high share turnover). Discretionary accruals are used as a proxy for firm-level mispricing because they measure the extent of possible earnings manipulation by the management (Sloan (1996)). Consistent with the hypothesis that mispricing affects both the quantity and quality of investment, their results show that investment by firms with either high R&D or high turnover is more sensitive to discretionary accruals, and both classes of firms have low returns following abnormally high investment. Baker, Stein and Wurgler (2003) also study the relation between non-fundamental stock valuation and investment, finding that equity dependent firms (according to Kaplan and Zingales (1997)'s index) issue more equity and invest more when their stock is likely overvalued, as measured by a high Tobin's Q and low future stock returns.

2.3 Empirical analysis

The hypothesized effect of the interaction between investor sentiment and technological innovation implies three testable restriction. First, from a microeconomic point of view, the difference in the marginal effect of technological innovation on investment, across firms more/less susceptible to informational externalities, should be larger when sentiment is higher. Second, from a macroeconomic perspective, the marginal effect of technological innovation on aggregate productivity should be larger for higher sentiment. Third, the effect on firm-level investment has to take place before the one on productivity.

The empirical implementation relies on different econometric methods. The microeconomic test is based on a panel regression, estimated with fixed effects, while the macroeconomic analysis employs regressions, vector auto-regressions and the method proposed by Den Haan (2000) to estimate the comovement of macroeconomic variables.

Before presenting the results, the next two sections discuss the data and the econometric approach.

2.3.1 Data description and summary statistics

The proxy for investor sentiment is the index of Baker and Wurgler (2006), which is based on the principal component analysis of a set of variables that have been linked to sentiment by previous research. More specifically, the variables included in the "short" series, which runs from 1962 to 2005, are the closed-end fund discount, NYSE share turnover, the number of IPOs, the average first day return on IPOs, the equity share in new issues, and the dividend premium. Baker and Wurgler (2006) also calculate a "long" series, from 1934 to 2005, which is based on three variables: the closed-end fund discount, the equity share in new issues and NYSE share turnover. To make sure that the index is not proxying for a business-cycle component, Baker and Wurgler (2006) orthogonalize the sentiment variables with respect to industrial production growth, growth in consumer durables, nondurables and services, and a dummy for NBER recessions. The results I present in Section 2.4 are mainly based on the "long" sentiment index, in order to increase the available time-series, although some specifications also use the "short" series for robustness.

To further reduce the risk that the sentiment index contains a business-cycle component, the empirical analysis includes several controls. Two variables are related to the financial market: $\Delta \mathrm{SP}_{\mathrm{disc}}$, which is the return on the S&P consumer discretionary index, and lnP/E, the log P/E ratio of the S&P 500. Two are proxies for real economic activity: ΔIP, the log-change in industrial production, and ΔRS, the log change in retail sales. Another two controls are monetary variables: D.spread, the yield on BBB bonds minus the yield on the CBOE 10 year government bond index, and T.spread, the yield on the CBOE 10 year government bond index minus the yield on 6 months T-bills. All these series are from Global Financial Data. The last control for aggregate investment opportunities is a proxy for investor expectations derived from the Livingston Survey. The survey is maintained by the Federal Reserve Bank of Philadelphia, and collects academic and professional economists' forecasts for a set of macroeconomic variables. I define the proxy, LS, as the first principal component of changes in the median forecasts for real GDP, Weekly wages for the non-farm sector and Industrial production, which are the only variables with forecasts available through the whole post-war period. Changes are defined as the difference between the 12- and the 6-months-ahead median forecasts, in order to avoid measurement error in the current value of the variable.

The aggregate productivity measure is the Multifactor Productivity Index for the private non-farm business sector (excluding government enterprises), provided by the Bureau of Labor Statistics. Following Griliches (1990), the proxy for technological innovation is ΔP, the log change in the number of patents granted by the United States Patent and Trademark Office (USPTO), which I obtain from the USPTO's website. Granted patents provide a first screening of the novelty of the claim - the rejection rate is about 30% - and on average two years pass from application to grant. One may be tempted to lag granted patents accordingly, but strategic interaction among firms means that patent applications usually *lead* the actual availability of a technology. Early patents, for instance, are valuable in case of disputes about the attribution of later patents on the more mature technology (Barzel (1968), Griliches (1990)). I do not use citation-weighted patents (Trajtenberg (1990)) because future citations are not part of the information set available to investors when patents are granted. As Griliches (1990) notes, the variation in granted patents at the end of the 1970s is related to staff shortages at the USPTO, rather than to actual innovation, and 1979 and 1980 do show extreme changes. To make sure that this is not driving the results, I multiply the change in patents for the years 1979-1981 by 30% in the macroeconomic test, and I drop 1979-1981 in the microeconomic analysis. The conclusions, however, are similar when using the original ΔP series. There is one last important point to discuss, namely the potential endogeneity between innovation and economic growth. Technological innovation is a key contributor to economic growth, but it's not unreasonable to expect that business cycle conditions also affect the resources spent on R&D and, eventually, innovation itself. Indeed, while in "supply push" models innovation is driven by exogenous and unpredictable advances in scientific knowledge, in a "demand pull" framework it is economic conditions that affect future innovation. Geroski and Walters (1995) find that demand plays a modest role relative to supply side factors.

The microeconomic test is based on firm-level investment decisions, where the dependent variable is the ratio of investment to installed capital, and the independent variables are controls for firm profitability and aggregate investment opportunities, plus changes in sentiment, ΔS, and changes in patents. End-of-year balance sheet items are from Compustat, share prices and the number of shares outstanding from CRSP, while analyst forecasts are obtained from IBES. A detailed definition of firm-level variables is available in Appendix A.

Tables 2.1 and 2.2 report summary statistics for the macroeconomic and firm-level variables. Changes in multifactor productivity, industrial production, retail sales and returns on the S&P consumer discretionary index are all positive and significant, reflecting the strong economic growth of the post-war period. The mean of the business cycle control built from the Livingston Survey, LS, is not significant because the principal component analysis is run on demeaned variables. For firm-level variables, I use the definitions of Polk and Sapienza (2009), and the statistics in Table 2.2 are very similar to what they report, even if the sample is slightly different. The average investment ratio is about 32%, the mean Tobin's Q is 1.7, and the expected analyst forecast for ROA is just above 5.1%. Table 2.3 reports the correlations among the macroeconomic variables. Changes in the multifactor productivity index are strongly correlated with industrial production and retail sales growth, while ΔS and ΔP are essentially uncorrelated with the other variables.

2.3.2 Econometric methods

This section describes the implementation of the tests on firm-level investment and multifactor productivity. In both cases the objective is to identify the marginal effect of innovation *conditional* on higher or lower investor sentiment. In the case of firm investment, the analysis is based on a panel regression. For the test on multifactor productivity, I use standard regressions, vector auto-regressions and the method of Den Haan (2000) based on forecast errors.

Microeconomic test

The hypothesis to test is that, for firms that are more susceptible to informational externalities, the marginal effect of ΔP is larger for higher ΔS. The empirical literature on informational externalities in a learning by doing framework is relatively small, and it has focused on firms with high capital intensity (Zimmerman (1982), and Thornton and Thompson (2001)). Indeed, it is reasonable to expect that, in industries characterized by a high ratio of capital to total assets, firms can more easily observe the level of investment in a new technology, and measure the effect on their competitors' profitability. Capital intensity, then, is the variable I use to proxy for the relevance of informational externalities. I compute the median ratio of book capital to book assets within industries, defined by the first two digits of the SIC code, and sort industries in three groups, according to whether the median capital intensity is in the bottom 30%, middle 40% or top 30%. The marginal effect of ΔP, conditional on ΔS, is then compared across firms that belong to industries in the bottom 30% and top 30% of the capital intensity ratio. To do so, I estimate a model that explains the investment ratio with controls for firm-level profitability, aggregate growth expectations, and with ΔS, ΔP and their interaction. I then define a dummy, equal to one if the firm belongs to an industry in the top 30% of the capital intensity ratio, and this dummy is interacted with all the variables in the model. In other words, the estimation includes "marginal" effects for high capital intensity firms. More in detail, the model is as follows (the loadings are omitted):

$$
\begin{aligned}
\frac{\mathrm{I}_t}{\mathrm{K}_{t-1}} &= \alpha + \mathrm{Q}_{t-1} + \mathrm{CF}_{t-1} + \mathrm{dBE}_{t-1} + \mathrm{DD}_{t-1} + \{\mathrm{ctrl}_{t-s}\}_{s=0}^{2} + \\
&+ \{\mathrm{LS}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\cdot\Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + \\
&+ d_{>70\%}\cdot\big(1 + \mathrm{Q}_{t-1} + \mathrm{CF}_{t-1} + \mathrm{dBE}_{t-1} + \mathrm{DD}_{t-1} + \{\mathrm{ctrl}_{t-s}\}_{s=0}^{2} + \\
&+ \{\mathrm{LS}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\cdot\Delta\mathrm{S}_{t-s}\}_{s=0}^{2}\big)
\end{aligned}
$$

where $\frac{\mathrm{I}_t}{\mathrm{K}_{t-1}}$ is the investment to capital ratio, Q_{t-1} Tobin's q, CF_{t-1} the cash-flow to capital ratio, dBE_{t-1} the total amount of dividends paid in the previous year, as a fraction of book equity, DD_{t-1} a dummy equal to one if the firm has paid dividends in the previous year, ctrl one of the macroeconomic controls, LS the expectation index built from the Livingston Survey, and ΔP and ΔS changes in patents and investor sentiment. Finally, the dummy $\mathrm{d}_{>70\%}$ is equal to one if the firm belongs to an industry in the top 30% of the capital intensity ratio. When testing the main hypothesis, I test that $\Delta(\beta_t^d) = \beta^d_{\Delta\mathrm{P}\cdot\Delta\mathrm{S}} \cdot (\Delta\mathrm{S}^{75} - \Delta\mathrm{S}^{25})$ is different from zero, where $\beta^d_{\Delta\mathrm{P}\cdot\Delta\mathrm{S}}$ is the coefficient on the interaction between ΔP, ΔS and the dummy $\mathrm{d}_{>70\%}$. $\Delta(\beta_t^d)$ quantifies how much the difference in the marginal effect of innovation, between high/low levels of capital intensity, changes when ΔS is at its 75^{th} rather than the 25^{th} historical percentile. The model is estimated with a fixed-effect panel regression, and the standard errors for $\Delta(\beta_t^d)$ are computed with the delta method.

Macroeconomic test

The joint effect of investor sentiment and technological innovation on multifactor productivity is assessed using three methodologies. First of all, I estimate regressions[2] of ΔMFP on controls, ΔP, ΔS and ΔP·ΔS, and then compute the marginal effect of ΔP conditional on different levels of ΔS. The specification is:

$$\Delta\text{MFP}_t = \alpha+\sum_{s=1}^{3}\beta_{k,t-s}\text{ctrl}_{k,t-s}+\sum_{s=1}^{4}\gamma_{t-s}\Delta\text{P}_{t-s}+\sum_{s=1}^{4}\delta_{t-s}\Delta\text{S}_{t-s}+\sum_{s=1}^{4}\lambda_{t-s}\Delta\text{S}_{t-s}\Delta\text{P}_{t-s} \tag{2.1}$$

and marginal effects are calculated as:

$$\Delta\text{P}_{t-s}^{m} = \gamma_{t-s} + \lambda_{t-s}\cdot\Delta\text{S}_{t-s}^{q} \tag{2.2}$$

where $\{ctrl_k\}_{k=1}^{3}$ are investment opportunities controls, and ΔS_{t-s}^{q} is the q^{th} percentile of ΔS_{t-s}. Given that there is no a-priori restriction on the exact lag of ΔS·ΔP that may have an effect on ΔMFP, the significance will be judged on the basis of a Bonferroni adjustment.

The second method of testing the relation between investor sentiment, aggregate productivity and innovation is a vector auto-regression. I estimate several specifi-

[2] I have also estimated an autoregressive distributed lags (ARDL) model, with unchanged results. ARDL models allow to analyse the effect of a longer lag structure without increasing the number of regressors. By adding a lagged dependent variable in the right hand side of the regression described above, coefficients on lagged interaction terms, for instance, can be written as: $\Delta\text{MFP}_t = \ldots + \frac{\lambda_1 L+\ldots+\lambda_4 L^4}{(1-\Theta L)}$, where Θ is the coefficient on the lagged dependent variable and L is the lag operator.

cations, that will be discussed in detail in Section 2.4, which include at least one business-cycle control, in addition to ΔMFP, ΔS, and ΔP. The number of lags is 3, selected on the basis of the Akaike Information Criterion, but the results are also robust to choosing 2 or 4 lags. To identify structural shocks I impose that B is a lower triangular matrix:

$$(I - A(L))\, y_t = B\epsilon_t \tag{2.3}$$

where y_t is the vector of variables included in the VAR. The economic interpretation of this restriction is that the first variable is contemporaneously affected only by its own shocks, the second by it own and the first variable's shocks, and so on. The ordering of the variables in this case is important, and I will discuss the economic intuition of the various specifications while presenting the results.

The third econometric approach is based on the method proposed by Den Haan (2000), which has the advantage of computing conditional correlations without requiring identifying assumptions. After estimating a VAR, I use the parameters[3] to generate a simulated time series for the variables in the model, and then I calculate $\rho\left(\Delta \mathrm{P}\,(p = 40\%, 60\%)_t^{\mathrm{FE}}, \Delta \mathrm{MFP}_{t+s}^{\mathrm{FE}}\right)$, which is the correlation between time t forecast errors for $\Delta \mathrm{P}\,(p = 40\%, 60\%)$ and time $t + s$ forecast errors for ΔMFP, with $s \in [1, 5]$. $\Delta \mathrm{P}\,(40\%)$ and $\Delta \mathrm{P}\,(60\%)$ are variables equal to ΔP if ΔS is smaller than the 40^{th} percentile or greater than the 60^{th}, and zero otherwise. Each simulation consists of 250 observations, and I keep the last 55 to eliminate the effect of the starting conditions (the unconditional mean of each variable), and to have the same sample size as in the actual dataset. I repeat the simulation 1,000 times, and the correlation $\rho\left(\Delta \mathrm{P}\,(p = 40\%, 60\%)_t, \Delta \mathrm{MFP}_{t+s}\right)$ is estimated as the aver-

[3] The parameters are also randomly drawn, using the standard errors provided by the VAR.

age $\rho\left(\Delta\mathrm{P}\,(p=40\%,60\%)_t^{\mathrm{FE}},\Delta\mathrm{MFP}_{t+s}^{\mathrm{FE}}\right)$ across the 1,000 repetitions, with bootstrap standard errors.

2.4 Results

As I discuss in the next two sections, the results provide strong evidence that the impact of technological innovation on firm-level investment ratios and on productivity growth depends on investor sentiment. The timing is also as expected, with higher $\Delta\mathrm{S}\cdot\Delta\mathrm{P}$ in year t increasing investment in t and $t+1$, and raising productivity growth after three years.

2.4.1 Microeconomic analysis

The first step in understanding the contribution of $\Delta\mathrm{S}$ to firm investment is to estimate the model in equation (2.1) on the whole sample, which runs from 1969 to 2005, setting $\mathrm{d}_{>70\%}$ to zero. Specification (3) in Table 2.4 shows that, of the controls for firm investment opportunities, Tobin's Q and cash-flows increase the investment ratio, and the coefficients are similar to those reported by Polk and Sapienza (2009). Retail sales have a positive effect when lagged by one year, unlike LS, which has a statistically significant negative impact. The sign of $\Delta\mathrm{P}$ and $\Delta\mathrm{S}$ is also negative, with high t-stats for lagged effects, while the interaction between $\Delta\mathrm{P}$ and $\Delta\mathrm{S}$ is not statistically significant at any lag. In specification (4) the model is estimated on the sample of firms that belong to industries in the top/bottom 30% of the distribution of capital intensity, and the dummy $\mathrm{d}_{>70\%}$ is set to 1 for firms in the top 30%. Some of the parameters are substantially different when interacted. The coefficients on cash flows and retail sales are smaller for firms with high capital intensity, while those on ΔP and ΔP·ΔS are larger. The coefficient on the interaction, in particular, is statistically significant only for high capital-intensity firms, which suggests that the effect of innovation, made conditional on ΔS, is indeed larger for such firms.

To directly test whether this is the case, Table 2.5 shows the values and t-stats of $\Delta(\beta_t^d) = \beta_{\Delta P \cdot \Delta S}^d \cdot (\Delta S^{75} - \Delta S^{25})$ for different specifications. As discussed in Section 2.3.2, $\Delta(\beta_t^d)$ quantifies how much the difference in the marginal effect of innovation, between high/low levels of capital intensity, changes when ΔS is at its 75^{th} rather than the 25^{th} percentile. In Table 2.5, model (2.1) is estimated with one business-cycle control at the time, and specification (1) excludes 1979-1981 because, as reported by Griliches (1990), staff shortages at the USPTO created spurious changes in the number of granted patents. The second specification covers the whole 1969-2005 sample, and there is little difference in the conclusions. The contemporaneous and lagged $\Delta(\beta_t^d)$ is positive across the whole set of controls, with t-stats equal about 2.5, and usually the lagged $\Delta(\beta_t^d)$ is smaller than the contemporaneous one. In economic terms, ΔP has a volatility of about 13% a year, so a one standard-deviation increase in ΔP adds 7% to the difference in the investment ratio, between high/low capital intensity firms, when ΔP is at its 75^{th} rather than the 25^{th} percentile.

Table 2.6 reports the first set of robustness checks, estimating the model with 1 and 3 lags rather than 2, clustering standard errors by firm to account for autocorrelation, and using the shorter proxy of investor sentiment, which covers 1962-2005 (see Section 2.3.1). The results prove to be robust, with the contemporaneous $\Delta(\beta_t^d)$ being higher when the panel regression includes three lags, and staying statistically significant in 4 out of 7 cases in the specification based on the shorter sentiment proxy. Table 2.7 shows additional specifications, that add discretionary accruals (following Polk and Sapienza (2009)) and include two business-cycle controls. Including accruals (specification (1)) makes every $\Delta(\beta_t^d)$ statistically insignificant, but this is a sample size issue, because accruals are available for only less than half the firm-year observations. Specification (2) shows the results from models that do not include discretionary accruals, but are estimated on observations that have a valid accruals entry. Almost all coefficients are statistically insignificant, confirming that the results of specification (1) are due to a substantial reduction in sample size. The remaining panels of Table 2.7 show $\Delta(\beta_t^d)$ when the regression includes two business-cycle controls. The coefficients and t-stats are slightly smaller when the controls are the lnP/E and ΔSP_{disc}, but the overall magnitude and pattern of significance carry through this last robustness check.

2.4.2 Macroeconomic analysis

The first evidence on the effect of investor sentiment and innovation on aggregate productivity is based on a set of regressions, estimated both with OLS and with the method of Rousseeuw and Leroy (1987), which improves the robustness to outliers. The dependent variable is ΔMFP, while the right-hand side includes several sets of controls, ΔS, ΔP and ΔS·ΔP. The regression estimates may be biased by endogeneity, although the coefficients are remarkably stable through alternative specifications, but the conclusions are confirmed by a VAR analysis that I discuss later in this section. In Table 2.8 the specifications include four lags of ΔS, ΔP and of their interaction, because it is difficult to formulate restrictions on the exact lag that should affect multifactor productivity, and four years is a reasonably long period of time. This obviously changes the significance level of the t-tests but, as I discuss below, the only significant interaction between ΔP and ΔS has large t-stats in most specifications, which make it significant at 10% even after a Bonferroni adjustment that accounts for all the combinations of the 4 lags. The results show that high ΔIP predicts lower ΔMFP, while ΔRS does not enter significantly. LS, on the other hand, is positive and significant. ΔS and ΔP are mostly statistically insignificant when taken alone, but the third lag of the interaction is highly significant and with the expected positive sign. It is the only significant lag in all the specifications, so it is the lag I focus on when computing the marginal effects, that are reported at the bottom of the table for the three full specifications. In every case the marginal effect of ΔP is positive and significant when ΔS is at its top quartile, while it is usually non significant for ΔS equal to the median or the first quartile. The results suggest that a one standard deviation increase in ΔP would raise ΔMFP in three year's time by about 0.70% for high values of ΔS. This finding is confirmed by Table 2.9, which provides initial robustness checks with specifications that include different lags for the controls, and OLS estimation rather than robust regressions.

The second set of evidence comes from a series of vector auto-regressions, which provide an insight on the dynamic relations between the variables. As explained in Section 2.3.2, I impose the restriction that the matrix B in equation (2.3) is lower triangular, so the shocks can be given a structural interpretation. In this setting, deciding the order of the variables in the VAR is important to give economic content to the restriction on the matrix B. Most of the impulse response functions that I present are based on models where ΔS is the first variable, then ΔP, LS, one of the business cycle controls and finally ΔMFP. The rationale for this ordering is that contemporaneous innovations to variable n should not be relevant for variable $n-1$. More in detail, investor sentiment should behave like noise, and demand factors have a modest impact on innovation (Geroski and Walters (1995)). In addition, expectations (LS) should react more quickly than business-cycle controls to new information, and productivity growth should be associated with changes in business-cycle variables. The first row of Figure 2.1 shows the cumulative orthogonalized impulse response function (COIRF) of ΔS to ΔMFP in the VAR specifications ΔS, ΔRS, ΔMFP and ΔS, LS, ΔMFP. The functions suggest that ΔS has an initially negative impact, which is consistent with investor sentiment misallocating resources and reducing productivity, when considered in isolation. The effect of ΔP on ΔMFP, on the other hand, is not significant, as suggested by the sub-figures in the second row, which show COIRFs of ΔP on ΔMFP in the VAR that includes ΔS, ΔP, ΔRS (LS), ΔMFP. In the third row the specifications are ΔS, ΔP·ΔS, ΔRS (LS), ΔMFP, and in both cases there is a significant effect of the interaction after three years, consistent with the results presented in Tables 2.8 and 2.9. The same conclusion can be drawn from the last row of Figure 2.1, in which the VARs include the variable ΔP interacted with a dummy equal to one if ΔS is above (below) its 60^{th} (40^{th}) percentile. The implication is that ΔP produces a delayed positive effect on ΔMFP only if the contemporaneous ΔS is relatively high. These findings are confirmed in Figure 2.2, where the VARs include the controls LS and ΔRS at the same time. The last three panels show COIRFs from the the vector auto-regressions LS, ΔRS, ΔMFP, ΔS and LS, ΔRS, ΔP, ΔMFP, ΔS which are useful to evaluate the assumption that ΔS is only affected by its own innovations. Indeed, the impulse response functions of the controls do not show a significant effect, while ΔP does have a slightly negative impact on ΔS after one year, although the effect is smaller and only marginally significant when ΔP is multiplied by 0.3 in the years 1979-1981 (unreported result), as discussed in Section 2.3.1. The final set of vector auto-regressions changes the order of the variables, so that ΔRS is now affected by all the structural innovations, and ΔMFP by all with the exception of those for ΔRS. Indeed, it can be reasonable to assume that contemporaneous shocks to productivity affect sales, rather than viceversa. The impulse response functions confirm the findings of the previous specifications, with both ΔP(60) and ΔP$\cdot\Delta$S having a positive effect on productivity growth after three years.

The last set of results is based on the method of Den Haan (2000), which is particularly useful for computing multi-horizon correlations between macroeconomic variables because it does not require identifying assumptions. The procedure for estimating the correlations (see Section 2.3.2) is based on a VAR that includes ΔMFP, ΔS, ΔP, a business-cycle control and either ΔP(60) or ΔP(40). Table 2.10 reports correlations up to five years ahead, across business-cycle controls and for both the "long" and "short" sentiment index. Interestingly, the pattern of correlations is different across the two sentiment proxies, but not across controls. When the "long" index is used, the correlation between ΔP(60) and ΔMFP is small for the first two years, but then reaches 30% in year $t + 3$, in line with the results presented so far. It is then negative in year $t + 4$, meaning that the effect on productivity *growth* is not permanent, which is quite reasonable from an economic point of view, and consistent with the VAR analysis presented above, where the *cumulative* impulse response function is positive after three years but zero in the long run. The correlations between ΔP(40) and ΔMFP, on the other hand, are negative and usually small, with a few exceptions in year $t + 4$. The second panel of Table 2.10 presents results when the "short" sentiment index is used. The correlations between ΔP(60) and ΔMFP are positive and gradually increasing from year $t+1$ to $t+3$, rather than high in just $t+3$, and still substantially negative in $t+4$.

2.5 Conclusion

I study whether investor sentiment acts as a subsidy to investment in new technologies, in particular for firms that are more likely to experience informational externalities, and if such subsidy produces positive effects for aggregate productivity growth. Early adopters likely generate informational externalities about the productivity of a new technology, and this can create a free rider problem that keeps aggregate investment below the social optimum. The resulting reduction in the flow of information can delay the adoption of a valuable technology, or the rejection of a less productive one, both of which generate a loss in productivity growth. Investor sentiment can act as a subsidy to investment by reducing funding costs, and by enticing managers to undertake investments that cater to investors' optimism about the new technology, in order to maximize stock prices in the short term.

The empirical analysis investigates how the interaction between investor sentiment and technological innovation affects 1) firm-level investment, in particular for firms that are more susceptible to informational externalities, and 2) aggregate productivity. I find that the interaction has a positive effect on investment for firms subject to informational externalities, and that it raises future aggregate productivity growth. A one standard deviation increase in technological innovation adds about 7% to the difference in the investment ratio between high/low capital intensity firms if investor sentiment is higher rather than lower, as measured by changes in sentiment at the top and bottom of the interquartile range. The effect of the same increase in innovation on productivity growth, when changes in sentiment are at the third quartile, is slightly more than 0.5%.

Figure 2.1: Cumulative orthogonalized impulse response functions, I/III

The VAR variables are (in the same order as in the Cholesky decomposition): (row 1, col. 1) ΔS, ΔRS, ΔMFP. (2,1) ΔS, ΔP, ΔRS, ΔMFP. (3,1) ΔS, ΔP·ΔS, ΔRS, ΔMFP. (4,1) ΔS, ΔP(40), ΔRS, ΔMFP. (1,2) ΔS, LS, ΔMFP. (2,2) ΔS, ΔP, LS, ΔMFP. (3,2) ΔS, ΔP·ΔS, LS, ΔMFP. (4,2) ΔS, ΔP(60), ΔRS, ΔMFP. ΔS is the change in the sentiment index, ΔRS the log change in retail sales, ΔP the log change in patents, LS the investor expectation index, ΔMFP the log change in multifactor productivity. ΔP(60) and ΔP(40) are ΔP multiplied by a dummy equal to one if ΔS is greater/smaller than the $60^{th}/40^{th}$ percentile. 90% confidence intervals shown.

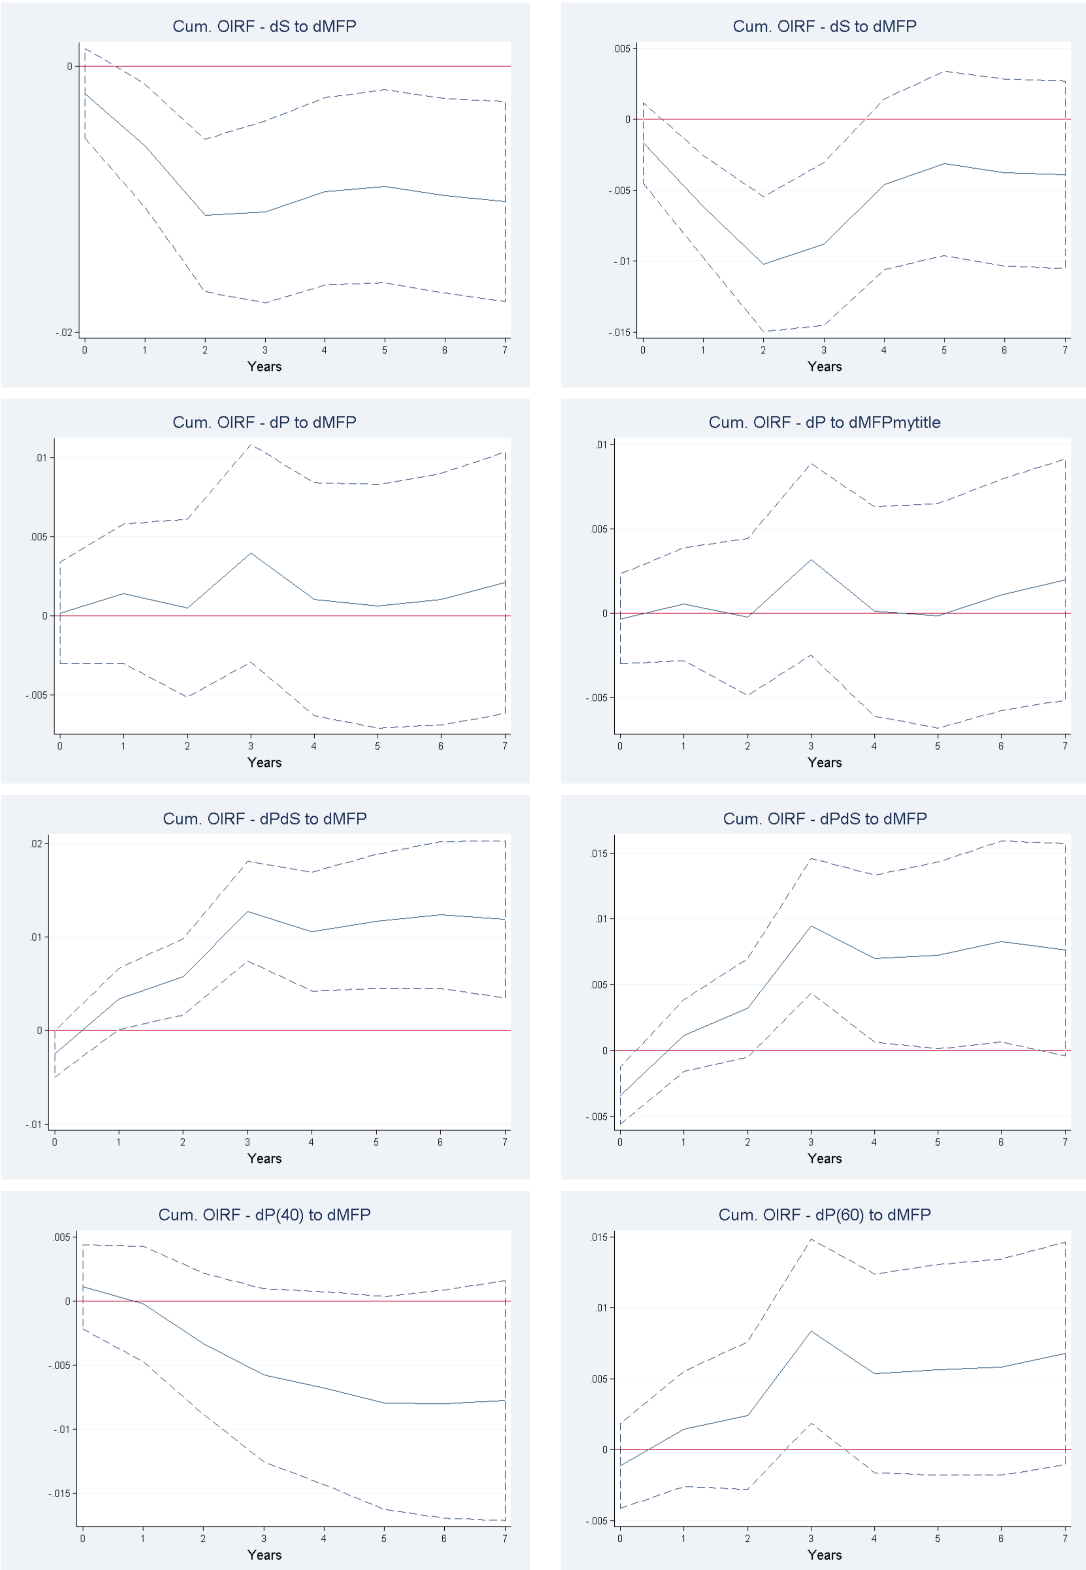

Figure 2.2: Cumulative orthogonalized impulse response functions, II/III.

The VAR variables are (in the same order as in the Cholesky decomposition): (1,1) ΔS, LS, ΔRS, ΔMFP. (2,1) ΔS, ΔP·ΔS, LS, ΔRS, ΔMFP. (3,1) ΔS, ΔP(60), LS, ΔRS, ΔMFP. (4,1) LS, ΔRS, ΔMFP, ΔS. (1,2) ΔS, ΔP, LS, ΔRS, ΔMFP. (2,2) ΔS, ΔP(40), LS, ΔRS, ΔMFP. (3,2) LS, ΔRS, ΔMFP, ΔS. (4,2) LS, ΔRS, ΔP, ΔMFP, ΔS. See Figure 2.1 for variable definitions. In the subtitles, ctrl1 is LS and ctrl2 is ΔRS. 90% confidence intervals shown.

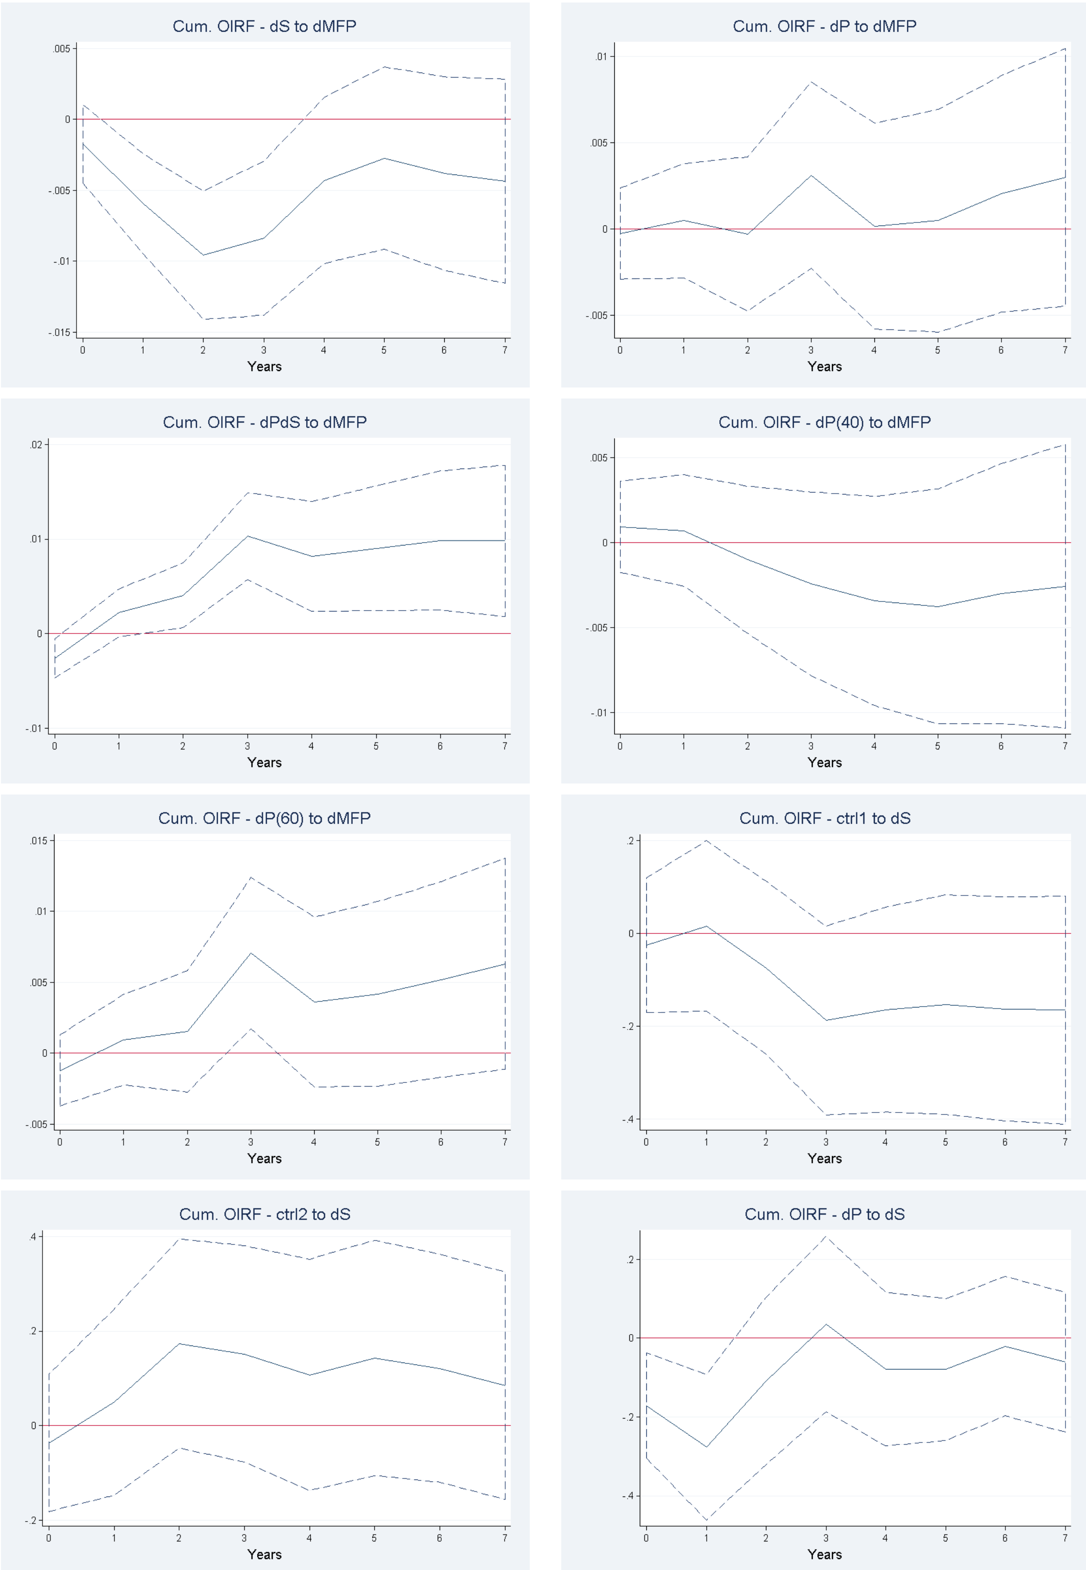

Figure 2.3: Cumulative orthogonalized impulse response functions, II/III.

The VAR variables are (in the same order as in the Cholesky decomposition): (1,1) ΔS, LS, ΔMFP, ΔRS. (2,1) ΔS, ΔP·ΔS, LS, ΔMFP, ΔRS. (1,2) ΔS, ΔP, LS, ΔMFP, ΔRS. (2,2) ΔS, ΔP(40), LS, ΔMFP, ΔRS. (3,2) ΔS, ΔP(60), LS, ΔMFP, ΔRS. See Figure 2.1 for variable definitions. 90% confidence intervals shown.

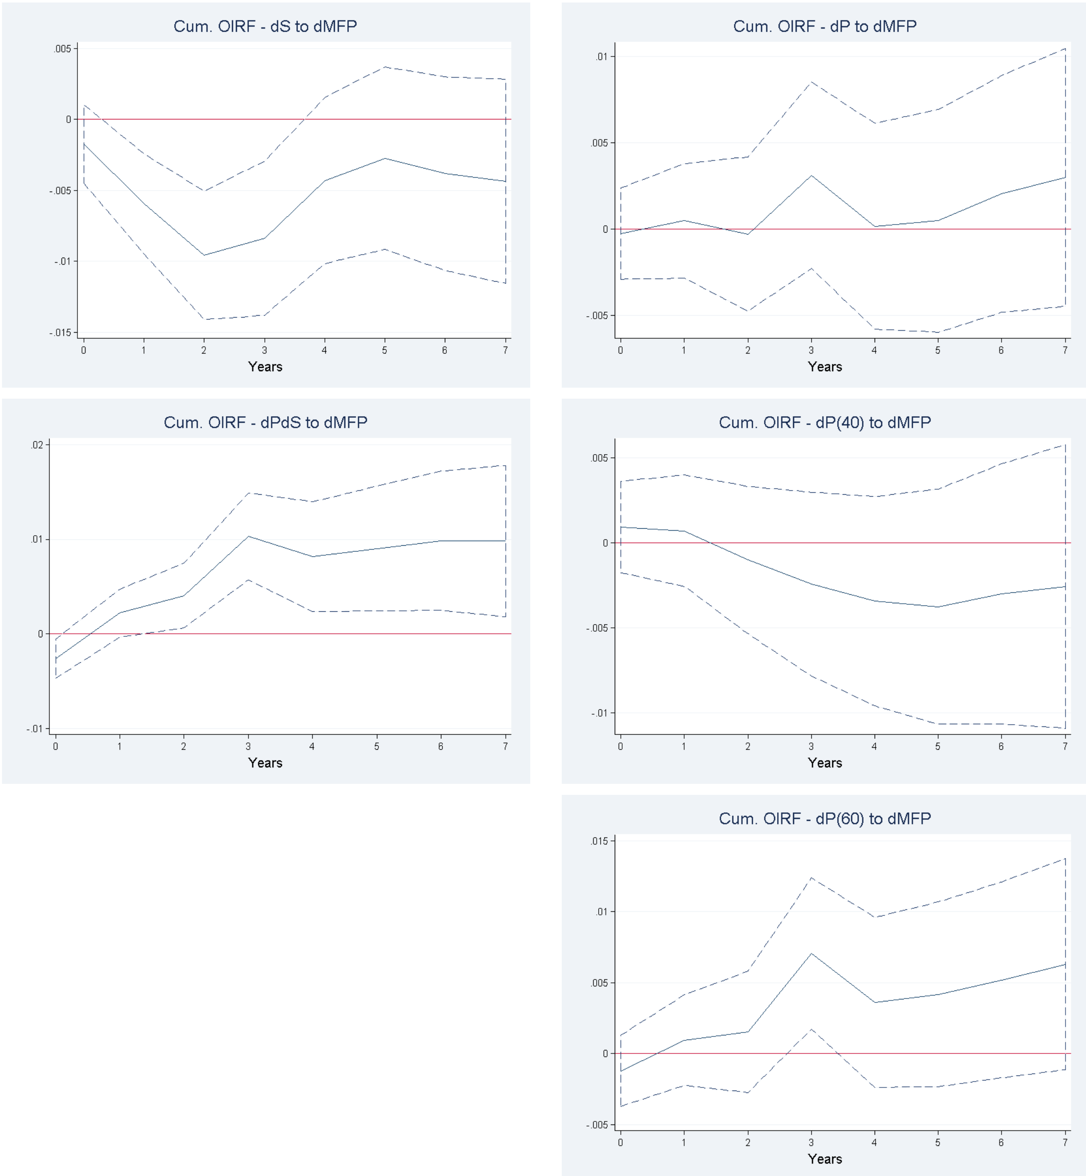

Table 2.1: Summary statistics.

Means, standard deviations and percentiles for the macroeconomic variables. ΔMFP is log-changes in the multifactor productivity index. ΔIP and ΔRS are log-changes in industrial production and retail sales. lnP/E is the log P/E ratio of the S&P 500. ΔSP_{disc} is the return on the S&P consumer discretionary index. D.spr. and T.spr. are the default and term spread (BBB - 10 year T.bonds and 10 year T.bonds - 6 months T.bills, respectively). LS is the principal component of Livingston Survey's forecasted 6 to 12 months ahead changes of real GDP, weekly wages for the non-farm sector and industrial production. ΔP is the log-change in granted patents and ΔS is the change in Baker and Wurgler's "long" sentiment index. 1951 to 2005.

	Mean	σ_t	**25**th	**50**th	**75**th
ΔMFP	0.01157	0.00235	0.00075	0.01205	0.02528
ΔIP	0.03247	0.00600	0.01017	0.03656	0.0656
ΔRS	0.03397	0.00276	0.01886	0.03314	0.05271
lnP/E	2.75013	0.05366	2.45873	2.81661	2.94969
ΔSP_{disc}	0.07784	0.02392	-0.05636	0.10368	0.19399
D.spr.	0.01640	0.00114	0.00900	0.01584	0.02270
T.spr.	0.01116	0.00160	0.00290	0.00945	0.01920
LS	0.00115	0.00168	-0.00487	0.00087	0.00881
ΔP	0.02524	0.01749	-0.06228	0.01770	0.08505
ΔS	0.01546	0.10757	-0.46737	0.08194	0.46822

Table 2.2: Summary statistics of firm-level variables.

I is investment, K capital, Q Tobin's q, CF cash flows, dBE the ratio of dividends through the year on book equity, $E_{t-1}[ROA_t]$ the ratio of the median analyst forecast of year t earnings during year $t-1$ divided by book assets, and daccr discretionary accruals. See Appendix A for a detailed definition of the variables.

	Mean	**Median**	σ	**Min.**	**Max.**	**Obs.**
$\frac{I_t}{K_{t-1}}$	0.323	0.214	0.392	0.000	6.652	34740
Q	1.702	1.255	1.353	0.453	13.631	37478
CF	0.287	0.292	1.621	-11.489	6.857	37388
dBE	0.018	0.000	0.029	0.000	0.150	38275
$E_{t-1}[ROA_t]$	0.051	0.047	0.077	-0.744	0.361	13508
daccr	0.022	0.025	0.130	-0.818	0.530	20671

Table 2.3: Correlations.

Correlations for the macroeconomic variables. ΔMFP is the log-change in the multifactor productivity index. ΔIP and ΔRS are log-changes in industrial production and retail sales. lnP/E is the log P/E ratio of the S&P 500. ΔSP_{disc} is the return on the S&P consumer discretionary index. D.spr. and T.spr. are the default and term spread (BBB - 10 year T.bonds and 10 year T.bonds - 6 months T.bills, respectively). LS is the principal component of Livingston Survey's forecasted 6 to 12 months ahead changes of real GDP, weekly wages for the non-farm sector and industrial production. ΔP is the log-change in granted patents and ΔS is the change in Baker and Wurgler's "long" sentiment index. 1951 to 2005.

	ΔMFP	ΔIP	ΔRS	lnP/E	ΔSP_{disc}	D.spr.	T.spr.	LS	ΔP	ΔS
ΔMFP	1									
ΔIP	0.63	1								
ΔRS	0.75	0.84	1							
lnP/E	0.18	-0.09	0.00	1						
ΔSP_{disc}	-0.03	-0.31	-0.11	0.23	1					
D.spr.	-0.32	-0.43	-0.36	0.36	-0.04	1				
T.spr.	0.03	-0.32	-0.18	0.25	0.24	0.49	1			
LS	0.05	-0.36	-0.31	0.35	0.17	0.10	0.30	1		
ΔP	-0.05	0.10	0.04	0.19	0.02	-0.02	0.01	0.06	1	
ΔS	-0.01	0.09	-0.01	0.01	0.08	-0.14	-0.25	0.01	-0.06	1

Table 2.4: Investment, investor sentiment and technological innovation.

The following model is estimated on a panel of Compustat firms: $\frac{I_t}{K_{t-1}} = \alpha + Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^2 + \{LS_{t-s}\}_{s=0}^2 + \{\Delta P_{t-s}\}_{s=0}^2 + \{\Delta S_{t-s}\}_{s=0}^2 + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^2 + d_{>70\%} \cdot (Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^2 + \{LS_{t-s}\}_{s=0}^2 + \{\Delta P_{t-s}\}_{s=0}^2 + \{\Delta S_{t-s}\}_{s=0}^2 + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^2)$. $d_{>70\%}$ is zero in specifications (1) to (3). Specification (4) focuses on firms in the top/bottom 30% of the distribution of capital intensity, and $d_{>70\%}$ is equal to 1 for firms in the top 30%. Q is Tobin's q, CF cash-flow/capital ratio, dBE the ratio of dividends through the year on book equity, DD a dummy equal to 1 if dividends have been paid in the year. See Table 2.1 and Appendix A for detailed variable definitions. Reported t-stats are based on standard errors computed with the heteroskedasticity-consistent sandwich estimator, and are clustered by year. 1969 to 2005, but 1979-1981 are excluded, as staff shortages at the USPTO spuriously affected ΔP.

	(1)	(2)	(3)	(4)	
					$d_{ci>70\%}$
Q_{t-1}			0.083	0.081	0.020
			18.86	12.39	1.30
CF_{t-1}			0.029	0.037	-0.057
			5.65	6.21	-2.27
dBE_{t-1}			-0.115	-0.262	0.218
			-1.16	-1.10	0.79
dd_{t-1}			-0.008	-0.027	0.008
			-0.98	-0.95	0.26
ΔRS_t		0.029	0.041	0.072	-0.067
		0.73	1.42	2.46	-3.35
ΔRS_{t-1}		0.089	0.030	0.105	-0.110
		2.54	0.96	2.62	-3.34
ΔRS_{t-2}		0.088	0.057	0.132	-0.094
		2.31	1.67	2.89	-2.61
LS_t	-2.661	-1.663	-1.963	-2.483	0.863
	-3.57	-2.65	-3.90	-2.30	0.75
LS_{t-1}	-0.499	-1.101	-1.262	-0.715	-0.018
	-0.54	-0.98	-1.40	-0.50	-0.02
LS_{t-2}	-0.192	-0.172	-0.274	-1.346	2.029
	-0.24	-0.19	-0.37	-1.30	2.50
ΔP_t	0.088	0.097	0.036	0.074	0.071
	1.05	1.14	0.46	0.60	0.59
ΔP_{t-1}	-0.197	-0.276	-0.213	-0.299	0.317
	-1.75	-2.56	-2.26	-2.36	2.99
ΔP_{t-2}	-0.315	-0.376	-0.275	-0.402	0.324
	-2.72	-4.44	-3.38	-2.03	1.76
ΔS_t	0.015	0.005	0.001	0.010	-0.023
	1.00	0.28	0.06	0.36	-1.08
ΔS_{t-1}	-0.022	-0.018	-0.024	-0.026	0.008
	-1.69	-1.75	-2.35	-1.40	0.52
ΔS_{t-2}	-0.006	-0.008	-0.012	-0.014	0.004
	-0.37	-0.63	-1.13	-0.84	0.29
$(\Delta P \cdot \Delta S)_t$	-0.169	0.010	0.003	-0.144	0.495
	-1.12	0.06	0.02	-0.63	2.15
$(\Delta P \cdot \Delta S)_{t-1}$	-0.083	-0.064	0.017	0.002	0.280
	-0.55	-0.53	0.19	0.02	2.26
$(\Delta P \cdot \Delta S)_{t-2}$	-0.248	-0.250	-0.101	-0.235	0.351
	-2.41	-2.56	-1.21	-0.96	1.39
intercept	0.002	0.003	0.001	0.001	-0.001
	0.36	0.48	0.33	0.29	-0.27
R^2	1.10	1.38	7.91	8.98	
Obs	33621	33621	33531	17136	

Table 2.5: Investment, investor sentiment and technological innovation: test of the microeconomic channel.

The following model is fitted on a panel of Compustat firms with a fixed-effects regression: $\frac{I_t}{K_{t-1}} = \alpha + Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^{2} + \{LS_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s}\}_{s=0}^{2} + \{\Delta S_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^{2} + d_{>70\%} \cdot (Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^{2} + \{LS_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s}\}_{s=0}^{2} + \{\Delta S_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^{2})$. ctrl is a business-cycle control (shown in the columns of the table). The table reports a test of whether the sensitivity of investment to the interaction between sentiment and innovation is higher for firms belonging to industries with high capital intensity (capital to book value of assets). The panel includes firms in the top/bottom 30% of the distribution of capital intensity, and the dummy $d_{>70\%}$ is equal to 1 when firms fall in the top 30%. See Table 2.1, 2.4 and Appendix A for detailed variable definitions. The table reports values and t-stats of the following difference: $\Delta(\beta_t^d) = \beta_{\Delta P \cdot \Delta S}^d \cdot (\Delta S^{75} - \Delta S^{25})$, where $\beta_{\Delta P \cdot \Delta S}^d$ is the coefficient on the interaction between innovation, changes in sentiment and the dummy $d_{>70\%}$. $\Delta(\beta_t^d)$ quantifies how much the difference in the marginal effect of innovation, between high/low levels of capital intensity, changes when ΔS is at its 75^{th} rather than the 25^{th} historical percentile. Reported t-stats are based on standard errors computed with the heteroskedasticity-consistent sandwich estimator, and are clustered by year. 1969 to 2005, but 1979-1981 are excluded in specification (1) because staff shortages at the USPTO created spurious changes in the number of granted patents.

		lnP/E	ΔSP_{cd}	ΔIP	ΔRS	D.spr.	T.spr	ΔSP
	$\Delta(\beta_t^d)$	0.518	0.549	0.448	0.478	0.686	0.645	0.755
		2.04	2.15	1.95	2.24	2.69	2.34	2.71
(1)	$\Delta(\beta_{t-1}^d)$	0.138	0.311	0.506	0.35	0.425	0.302	0.446
		0.61	2.26	3.84	2.74	2.50	1.61	3.31
	$\Delta(\beta_{t-2}^d)$	0.193	0.39	0.428	0.363	0.29	0.42	0.44
		0.67	1.39	1.64	2.14	1.06	1.29	1.33
	$\Delta(\beta_t^d)$	0.512	0.508	0.469	0.496	0.48	0.505	0.489
		2.73	2.52	2.50	2.56	2.47	2.47	2.42
(2)	$\Delta(\beta_{t-1}^d)$	0.49	0.475	0.49	0.487	0.499	0.504	0.478
		2.72	2.57	2.69	2.59	2.73	2.65	2.58
	$\Delta(\beta_{t-2}^d)$	0.276	0.287	0.265	0.26	0.289	0.289	0.294
		1.20	1.24	1.17	1.18	1.21	1.21	1.24

Table 2.6: Investment, investor sentiment and technological innovation: test of the microeconomic channel. Robustness checks, I/II.

The following model is fitted on a panel of Compustat firms with a fixed-effects regression: $\frac{I_t}{K_{t-1}} = \alpha + Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^{2} + \{LS_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s}\}_{s=0}^{2} + \{\Delta S_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^{2} + d_{>70\%} \cdot (Q_{t-1} + CF_{t-1} + dBE_{t-1} + DD_{t-1} + \{ctrl_{t-s}\}_{s=0}^{2} + \{LS_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s}\}_{s=0}^{2} + \{\Delta S_{t-s}\}_{s=0}^{2} + \{\Delta P_{t-s} \cdot \Delta S_{t-s}\}_{s=0}^{2})$. ctrl is a business-cycle control (shown in the columns of the table). $\Delta(\beta_t^d) = \beta_{\Delta P \cdot \Delta S}^d \cdot (\Delta S^{75} - \Delta S^{25})$, where $\beta_{\Delta P \cdot \Delta S}^d$ is the coefficient on the interaction between innovation, changes in sentiment and the dummy $d_{>70\%}$. $\Delta(\beta_t^d)$ quantifies how much the difference in the marginal effect of innovation, between high/low levels of capital intensity, changes when ΔS is at its 75^{th} rather than the 25^{th} historical percentile. (1) and (2): the macroeconomic controls, ΔP, ΔS and (ΔP·ΔS) are lagged by 3 and 1 years. (3): standard errors are clustered by firm rather than year. (4): ΔS is the the "short" Baker and Wurgler's sentiment index. Standard errors are computed with the heteroskedasticity-consistent sandwich estimator, and are clustered by year. See Table 2.1, 2.4 and Appendix A for detailed variable definitions. 1969 to 2005, but 1979-1981 are excluded because staff shortages at the USPTO created spurious changes in the number of granted patents.

		lnP/E	ΔSP_{cd}	ΔIP	ΔRS	D.spr.	T.spr	ΔSP
(1)	$\Delta(\beta_t^d)$	0.624	0.990	0.628	0.492	0.642	1.090	1.055
		1.80	3.57	2.08	2.22	2.41	2.95	3.38
	$\Delta(\beta_{t-1}^d)$	0.229	0.338	0.458	0.473	0.542	0.363	0.385
		1.03	2.56	3.62	3.58	3.84	2.50	2.58
	$\Delta(\beta_{t-2}^d)$	0.249	0.552	0.388	0.325	0.131	0.325	0.558
		0.99	2.26	1.51	1.76	0.75	1.23	2.00
	$\Delta(\beta_{t-3}^d)$	0.241	-0.118	0.021	0.167	0.023	0.082	-0.083
		1.04	-0.68	0.13	0.81	0.13	0.39	-0.40
(2)	$\Delta(\beta_t^d)$	0.437	0.418	0.108	0.240	0.407	0.350	0.436
		2.41	2.22	0.47	1.06	2.17	1.71	2.14
	$\Delta(\beta_{t-1}^d)$	0.067	0.276	0.391	0.287	0.333	0.213	0.348
		0.29	1.30	2.03	1.43	1.71	0.85	1.72
(3)	$\Delta(\beta_t^d)$	0.518	0.549	0.448	0.478	0.686	0.645	0.755
		2.01	2.17	1.66	1.90	2.62	2.50	2.85
	$\Delta(\beta_{t-1}^d)$	0.138	0.311	0.506	0.350	0.425	0.302	0.446
		0.52	1.27	1.91	1.43	1.69	1.22	1.79
	$\Delta(\beta_{t-2}^d)$	0.193	0.390	0.428	0.363	0.290	0.420	0.440
		0.83	1.70	1.80	1.59	1.32	1.82	1.81
(4)	$\Delta(\beta_t^d)$	-0.053	0.327	0.641	0.554	0.159	0.644	0.585
		-0.31	1.17	3.39	2.98	0.54	2.49	2.13
	$\Delta(\beta_{t-1}^d)$	0.113	0.037	0.228	0.369	-0.021	0.082	0.049
		0.84	0.21	1.57	3.06	-0.10	0.54	0.27
	$\Delta(\beta_{t-2}^d)$	-0.417	0.048	0.129	0.331	-0.171	0.041	-0.013
		-2.78	0.31	0.94	2.30	-0.97	0.22	-0.07

Table 2.7: Investment, investor sentiment and technological innovation: test of the microeconomic channel. Robustness checks, II/II.

The following model is fitted on a panel of Compustat firms with a fixed-effects regression: $\frac{\mathrm{I}_t}{\mathrm{K}_{t-1}} = \alpha + \mathrm{Q}_{t-1} + \mathrm{CF}_{t-1} + \mathrm{dBE}_{t-1} + \mathrm{DD}_{t-1} + \{\mathrm{ctrl}_{t-s}\}_{s=0}^{2} + \{\mathrm{LS}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s} \cdot \Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + d_{>70\%} \cdot (\mathrm{Q}_{t-1} + \mathrm{CF}_{t-1} + \mathrm{dBE}_{t-1} + \mathrm{DD}_{t-1} + \{\mathrm{ctrl}_{t-s}\}_{s=0}^{2} + \{\mathrm{LS}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{S}_{t-s}\}_{s=0}^{2} + \{\Delta\mathrm{P}_{t-s} \cdot \Delta\mathrm{S}_{t-s}\}_{s=0}^{2})$. ctrl is a business-cycle control (shown in the columns of the table). $\Delta(\beta_t^d) = \beta_{\Delta\mathrm{P}\cdot\Delta\mathrm{S}}^d \cdot (\Delta\mathrm{S}^{75} - \Delta\mathrm{S}^{25})$, where $\beta_{\Delta\mathrm{P}\cdot\Delta\mathrm{S}}^d$ is the coefficient on the interaction between innovation, changes in sentiment and the dummy $d_{>70\%}$. $\Delta(\beta_t^d)$ quantifies how much the difference in the marginal effect of innovation, between high/low levels of capital intensity, changes when $\Delta\mathrm{S}$ is at its 75^{th} rather than the 25^{th} historical percentile. (1): the panel regression includes contemporaneous discretionary accruals. (2): the panel regression does not include observations with missing discretionary accruals. (3) and (4): the regression includes two macroeconomic controls, in addition to LS, and the macroeconomic controls, ΔP, ΔS and (ΔP·ΔS) are lagged by 2 and 3 years. Standard errors are computed with the heteroskedasticity-consistent sandwich estimator, and are clustered by year. See Table 2.1, 2.4 and Appendix A for detailed variable definitions. 1969 to 2005, but 1979-1981 are excluded because staff shortages at the USPTO created spurious changes in the number of granted patents.

		lnP/E	$\Delta\mathrm{SP}_{cd}$	ΔIP	ΔRS	D.spr.	T.spr	ΔSP
	$\Delta(\beta_t^d)$	0.415	0.403	0.164	0.290	0.337	0.464	0.614
		1.34	1.35	0.52	0.99	0.99	1.44	1.84
(1)	$\Delta(\beta_{t-1}^d)$	-0.244	-0.030	0.137	0.118	0.192	0.016	0.052
		-1.62	-0.22	0.94	0.87	1.10	0.11	0.38
	$\Delta(\beta_{t-2}^d)$	0.277	0.409	0.359	0.314	0.189	0.310	0.461
		0.76	1.16	1.00	1.21	0.61	0.78	1.10
	$\Delta(\beta_t^d)$	0.407	0.399	0.157	0.288	0.320	0.456	0.609
		1.30	1.32	0.49	0.98	0.94	1.39	1.81
(2)	$\Delta(\beta_{t-1}^d)$	-0.257	-0.036	0.134	0.114	0.193	0.007	0.045
		-1.70	-0.27	0.92	0.83	1.09	0.05	0.32
	$\Delta(\beta_{t-2}^d)$	0.266	0.395	0.346	0.302	0.167	0.299	0.442
		0.72	1.10	0.94	1.15	0.54	0.74	1.03

		lnP/E & $\Delta\mathrm{SP}_{cd}$	ΔIP & ΔRS	D.spr. & T.spr
	$\Delta(\beta_t^d)$	0.309	0.393	0.590
		1.46	1.77	2.33
(3)	$\Delta(\beta_{t-1}^d)$	0.081	0.436	0.690
		0.44	3.98	3.25
	$\Delta(\beta_{t-2}^d)$	0.232	0.275	0.069
		0.94	1.72	0.26
	$\Delta(\beta_t^d)$	0.588	0.540	0.942
		1.66	2.08	2.75
	$\Delta(\beta_{t-1}^d)$	0.213	0.622	0.711
(4)		1.13	4.56	3.89
	$\Delta(\beta_{t-2}^d)$	0.290	0.365	-0.095
		1.26	2.16	-0.47
	$\Delta(\beta_{t-3}^d)$	0.063	0.278	-0.171
		0.25	1.44	-1.02

Table 2.8: Multifactor productivity, sentiment and technological innovation. Regressions I/II.

Time-series robust regressions of the log-change in the multifactor productivity index on controls (1 lag), log-change in the number of granted patents, change in the "long" Baker-Wurgler sentiment index and interaction between log-changes in patents and changes in sentiment. The marginal effects panel shows the contribution of log-changes in the number of patents when changes in sentiment are equal to the 25^{th}, 50^{th}, 75^{th} percentile. Control C^3, LS, is the same through all specifications. C^1 and C^2 change through specifications. (1), (2), (3): ΔIP and ΔRS. (4): lnP/E and ΔSP_{disc}. (5): D.spread and T.spread. I_t is $(\Delta S \cdot \Delta P)_t$. See Tab. 2.1 for definitions. $\bar{R}^2$ from OLS regressions. The marginal effects of the third lag of ΔP are defined as $\gamma_{-3} + \lambda_{-3} \cdot \Delta S_q$. T-stats in parentheses. 1951 to 2005.

	Cons	C^1_{-1}	C^2_{-1}	C^3_{-1}	ΔP_{-1}	ΔP_{-2}	ΔP_{-3}	ΔP_{-4}	ΔS_{-1}	ΔS_{-2}	ΔS_{-3}	ΔS_{-4}	I_{-1}	I_{-2}	I_{-3}	I_{-4}	$\bar{R}^2$ (%)
(1)	0.0128												0.0576	0.0518	0.1004	-0.0143	41
	6.38												2.94	2.68	5.21	-0.74	
(2)	0.0087	-0.1326	0.1562	0.5988	0.0132	-0.0074	0.0160	0.0072	-0.0053	-0.0027	-0.0031	0.0042					45
	1.89	-1.39	0.79	3.16	0.85	-0.44	0.95	0.45	-1.85	-0.93	-1.09	1.43					
(3)	0.0122	-0.1800	0.1075	0.2521	0.0245	-0.0044	-0.0004	0.0050	-0.0020	-0.0028	-0.0044	0.0042	0.0285	0.0340	0.0848	0.0113	67
	3.00	-2.46	0.66	1.59	1.98	-0.33	-0.03	0.40	-0.90	-1.25	-2.01	1.80	1.59	2.12	5.20	0.62	
(4)	0.0006	0.0035	0.0200	0.3868	-0.0039	-0.0118	0.0056	0.0030	-0.0055	-0.0033	-0.0008	0.0036	0.0456	0.0170	0.0637	0.0014	67
	0.04	0.65	1.81	2.08	-0.27	-0.75	0.34	0.19	-2.17	-1.24	-0.31	1.36	2.32	0.94	3.52	0.08	
(5)	0.0100	-0.0393	0.2219	0.3700	0.0018	-0.0096	0.0127	-0.0001	-0.0043	-0.0038	0.0002	0.0045	0.0435	0.0132	0.0764	0.0085	63
	2.12	-0.11	1.12	1.95	0.12	-0.57	0.74	-0.01	-1.60	-1.36	0.08	1.74	2.07	0.64	3.72	0.38	

Marginal effects	**(3)**		**(4)**		**(5)**	
$q=25^{th}$	-0.0414	(-2.51)	-0.0252	(-1.29)	-0.0242	(-1.07)
$q=50^{th}$	0.0065	(0.50)	0.0108	(0.66)	0.0190	(1.14)
$q=75^{th}$	0.0528	(3.45)	0.0456	(2.42)	0.0607	(3.45)

Table 2.9: Multifactor productivity, sentiment and technological innovation. Regressions II/II.

Time-series robust regressions of the log-change in the multifactor productivity index on controls (1 lag), log-change in the number of granted patents, change in the "long" Baker-Wurgler sentiment index and interaction between log-changes in patents and changes in sentiment. The marginal effects panel shows the contribution of log-changes in the number of patents when changes in sentiment are equal to the 25^{th}, 50^{th}, 75^{th} percentile. Controls C^1 and C^2 are ΔIP and ΔRS in all specifications. Control C^3, is LS in specifications (1) to (3). In regressions (3) and (4) it is the ΔSP. Controls are lagged 2 years in specification (1), 3 years in (2) and (3) and 1 year in specification (4). Specification (4) reports OLS estimates. I_t is $(\Delta S \cdot \Delta P)_t$. See Tab. 2.1 for definitions. $\bar{R}^2$ from OLS regressions. The marginal effects of the third lag of ΔP are defined as $\gamma_{-3} + \lambda_{-3} \cdot \Delta S_q$. T-stats in parentheses. 1951 to 2005.

	Cons	C^1	C^2	C^3	ΔP_{-1}	ΔP_{-2}	ΔP_{-3}	ΔP_{-4}	ΔS_{-1}	ΔS_{-2}	ΔS_{-3}	ΔP_{-4}	I_{-1}	I_{-2}	I_{-3}	I_{-4}	$\bar{R}^2$ (%)
(1)	0.0120	0.1170	-0.1210	0.2282	0.0038	-0.0132	0.0053	0.0118	-0.0048	-0.0059	0.0013	0.0084	0.0524	0.0187	0.0685	-0.0134	61
	2.40	1.22	-0.58	1.06	0.23	-0.75	0.30	0.67	-1.64	-2.00	0.42	2.94	2.16	0.85	3.21	-0.61	
(2)	0.0059	-0.1343	0.3195	0.0606	0.0080	-0.0108	0.0039	0.0040	-0.0033	-0.0056	0.0024	0.0087	0.0680	0.0344	0.0853	-0.0119	73
	1.31	-1.67	1.80	0.32	0.52	-0.66	0.23	0.25	-1.14	-1.98	0.90	3.02	3.27	1.64	4.28	-0.59	
(3)	0.0063	-0.1369	0.3109	0.0001	0.0078	-0.0109	0.0058	0.0045	-0.0036	-0.0056	0.0025	0.0082	0.0669	0.0350	0.0859	-0.0102	61
	1.41	-1.61	1.65	0.01	0.50	-0.65	0.36	0.28	-1.29	-2.01	0.92	3.09	3.25	1.76	4.38	-0.54	
(4)	0.0069	-0.1193	0.2841	-0.0013	0.0068	-0.0055	0.0114	0.0053	-0.0039	-0.0049	0.0018	0.0081	0.0662	0.0318	0.0819	-0.0107	61
	1.74	-1.59	1.72	-0.10	0.50	-0.37	0.81	0.37	-1.57	-2.00	0.75	3.47	3.66	1.82	4.74	-0.64	

Marginal effects	**(1)**		**(2)**		**(3)**		**(4)**	
$q=25^{th}$	-0.0277	(-1.27)	-0.0373	(-1.80)	-0.0356	(-1.76)	-0.0281	(-1.58)
$q=50^{th}$	0.0109	(0.62)	0.0108	(0.64)	0.0128	(0.81)	0.0181	(1.30)
$q=75^{th}$	0.0483	(2.32)	0.0574	(2.94)	0.0597	(3.29)	0.0629	(3.93)

Table 2.10: Correlation between forecast errors of ΔMFP and of ΔP when contemporaneous ΔS is high or low.

The correlations are calculated by estimating a VAR that includes ΔMFP, the indicated business-cycle control, ΔS, ΔP and either 1) ΔP(60), a variable equal to ΔP if ΔS is greater than its 60^{th} historical percentile, and zero otherwise or 2) ΔP(40), a variable equal to ΔP if ΔS is smaller than its 40^{th} historical percentile, and zero otherwise. Using the VAR parameters I simulate a series of length 250, discard the first 195 observations to eliminate the effect of starting conditions (each series' unconditional mean) and to have a simulated sample with the same number of observations as the actual sample. I then compute the simulated forecast errors at horizons of 1 to 5 years, and calculate the correlation between the simulated forecast errors of ΔMFP and of ΔP(60)) or ΔP(40). The simulation is repeated 1,000 times, and the table shows the average across simulations. Entries with a "0" ("1","2") superscript are not statistically significant (significant at 10%, 5%), and those without a superscript are significant at 1%. Bootstrap standard errors are calculated with 50 replications.

	LS	ΔIP	ΔRS	lnP/E	ΔSP_{cd}	D.spr.	T.spr.	ΔSP
	Investor sentiment: "long" Baker and Wurgler's index							
ΔP(ΔS>p(60))								
1yr	0.14	0.15	0.14	0.15	0.15	0.10	0.18	0.17
2yr	0.04	0.06	0.04	0.01^{2}	0.05	0.09	0.04	0.08
3yr	0.28	0.32	0.30	0.29	0.31	0.26	0.28	0.28
4yr	-0.21	-0.14	-0.14	-0.15	-0.15	-0.15	-0.20	-0.14
5yr	0.01^{0}	0.03	0.01	0.01	-0.01^{2}	-0.05	0.00^{0}	0.01^{0}
ΔP(ΔS<p(40))								
1yr	0.06	-0.01	-0.01^{0}	-0.01	0.06	-0.02	0.00^{0}	0.04
2yr	-0.04	-0.01	-0.06	-0.08	-0.06	-0.06	-0.04	-0.06
3yr	-0.07	-0.12	-0.11	-0.14	-0.12	-0.09	-0.07	-0.11
4yr	-0.18	-0.13	-0.15	-0.11	-0.16	-0.14	-0.08^{0}	-0.20^{0}
5yr	-0.06	-0.05	-0.07	-0.03	-0.01	-0.05	-0.01	0.00
	Investor sentiment: "short" Baker and Wurgler's index							
ΔP(ΔS>p(60))								
1yr	0.08	0.19	0.16	0.15	0.15	0.18	0.19	0.16
2yr	0.13	0.14	0.11	0.12	0.20	0.20	0.17	0.22
3yr	0.12	0.18	0.13	0.17	0.21	0.20	0.20	0.20
4yr	-0.28	-0.27	-0.25	-0.27	-0.28	-0.27	-0.32	-0.26
5yr	0.00^{0}	-0.04	-0.01	-0.04	-0.05	-0.10	-0.04	-0.05
ΔP(ΔS<p(40))								
1yr	0.03	-0.04	-0.08	-0.03	0.02	0.00^{0}	-0.01^{0}	0.04
2yr	0.04	0.05	0.03	0.02	0.05	0.04	0.04	0.06
3yr	-0.02	-0.04	-0.07	-0.09	-0.10	-0.01	-0.05	-0.08
4yr	-0.03	-0.07	-0.02	0.01	-0.10	-0.01^{2}	-0.01^{0}	-0.13
5yr	0.01	0.03	0.06	0.17	0.12	0.07	0.05	0.13

Appendix A

The definition of firm-level variables follows Polk and Sapienza (2009), and the first part of this Appendix is adapted from Appendix A in their paper. Investment (I) is capital expenditure (Compustat Item128). Capital (K) is net property, plant, and equipment (Item8). Tobin's q (Q) equals the market value of assets divided by the book value of assets (Item6). The market value of assets equals the book value of assets plus the market value of common stock less the sum of book value of common stock (Item60) and balance sheet deferred taxes (Item74) in year $t-1$. Cash flow (CF) equals the sum of earnings before extraordinary items (Item18) and depreciation (Item14) over beginning-of-year capital. Sales (Item12) is net sales. Accruals (accr) equal the change in accounts receivable (Item2) plus the change in inventories (Item3) plus the change in other current assets (Item68) minus the change in accounts payable (Item70) minus the change in other current liabilities (Item72) minus depreciation (Item178). Accruals are scaled by total assets (the average of Item6 at the beginning and end of the fiscal year). Discretionary accruals (daccr) are:

$$\text{daccr}_t = accr_t - normal\_accr_t \tag{2.4}$$

$$normal\_accr_t = \frac{\Sigma_{k=1}^{5} accr_{i,t-k}}{\Sigma_{k=1}^{5} sales_{i,t-k}} \cdot sales_{i,t} \tag{2.5}$$

BE is stockholders' equity, plus balance sheet deferred taxes (Item74) and investment tax credit (Item208, set to zero if unavailable), plus post-retirement benefit liabilities (Item330, set to zero if unavailable), minus the book value of preferred stock. Depending on availability of preferred stock data, I use redemption (Item56), liquidation (Item10), or par value (Item130), in that order, for the book value of preferred stock.

Stockholders' equity is as follows. If Item216 is not available, stockholders' equity is the book value of common equity (Item60), plus the book value of preferred stock. If common equity is not available, stockholders' equity is the book value of assets (Item6) minus total liabilities (Item181). The sample only includes firms with fiscal years ending in December.

The analysis focuses on common stocks (CRSP share code 10 and 11) traded on NYSE, AMEX and NASDAQ (exchange code 1, 2 and 3). Only ordinary cash dividends are considered (distribution code 12). The following variables are censored at the 0.5 and 99.5 percentiles: investment rate ($\mathrm{I}_t/\mathrm{K}_{t-1}$), Tobins's Q, discretionary accruals and expected ROA. The cash-flow ratio ($\mathrm{CF}_t/\mathrm{K}_{t-1}$) is censored at the 1.5 and 98.5 percentiles, because it exhibits a larger proportion of extreme values. Censoring at 3% rather than 1% of the cash-flow ratio's has a noticeable effect on the statistical and economic significance of the coefficient. The ratio of dividends to book equity (dBE) is censored at 0 and 0.15 to limit the effect of small BE values.

Chapter 3

Portfolio choice with distributions implied in option prices

This chapter benefited from the suggestions of Victor De Miguel, Theodosios Dimopoulos, Carlo Fezzi, Tim Johnson, Andrew Patton, Franco Peracchi, Stefano Sacchetto and Raman Uppal. I would also like to credit Andrew Patton and Michael Rockinger the Matlab code they made available on their web sites.

3.1 Introduction

In a recent contribution on the implications of estimation error for portfolio choice, De Miguel, Garlappi and Uppal (2009) compare the out-of-sample performance of several portfolio choice models with that of a naïve strategy that allocates wealth in equal proportions to N securities. They show that, far from being badly inefficient, the $\frac{1}{N}$ strategy often yields the best out-of-sample results. The literature has indeed recognized that model misspecification and estimation error can have important consequences for portfolio choice problems. One solution is to acknowledge that portfolio weights may be driven by noise rather than fundamentals, and to prevent the model from taking extreme positions. This can be done by imposing constraints on the domain of the distributions' parameters (Ledoit and Wolf (2003)) or on the portfolio weights (Brandt, Santa-Clara and Valkanov (2009), Jagannathan and Ma (2003)). Investors can use the Bayesian framework to formalize uncertainty about the return generating function (Jorion (1986), Pastor (2000), Pastor and Stambaugh (2000)), or they may adopt robust decision rules. In Garlappi, Uppal and Wang (2007), for instance, it is possible to account for model uncertainty by minimizing the objective function over admissible parameter values, which are represented by confidence intervals. On the other hand, the precision of estimated returns can be improved by exploiting the statistical properties of returns themselves (as in Pastor and Stambaugh (2009)), or by recognizing the implications that omitted factors have on the relation between expected returns and the covariance structure (MacKinlay and Pastor (2000)).

My contribution is to suggest a method that improves the precision of estimated expected returns. I solve a portfolio choice problem by deriving the distribution of expected returns from the cross-section of option prices, under the assumption that the information about higher moments implied in the moneyness structure can be helpful to reduce uncertainty about estimated expected returns. There is substantial evidence, for instance, that implied volatility has significant forecasting power for future realized volatility (Christensen and Prabhala (1997), Jiang and Tian (2005)). Distributions derived from option prices are risk-neutral, which means that the probabilities incorporate a correction for investors' risk aversion, and I rely on the empirical pricing kernel of Rosenberg and Engle (2002) to obtain estimates of the expected objective distributions. Alternatively, I simply shift the risk-neutral distribution to the right by adding the risk premium, which allows to evaluate the portfolio choice implications of risk-neutral skewness, that imposes a penalty on extreme weights in terms of magnified expected losses.

The empirical analysis solves a portfolio choice problem for an investor with power utility, and compares the certainty equivalent from the $\frac{1}{N}$ rule with those from strategies based on distributions of expected returns estimated with option prices, and with times-series models that allow an autoregressive structure for the first four moments.[1] The marginal distributions for expected returns are joined into a multivariate one by using the normal and t copulas. Comparing the results across the two specifications allows to evaluate the economic importance of accounting for tail dependence with the t copula. Other families of copulas can also model asymmetric dependence, in particular the Archimedean copulas used by Patton (2004), but they can seldom be extended beyond the bivariate case, and parameter instability may increase estimation error (Nelsen (1999)). The focus is on how the "active" portfolio choice models com-

[1] The relatively large number of parameters in the time-series models may generate estimation error, so in unreported robustness checks I have used a simpler GARCH-in-mean, without altering the conclusions.

pare with the $\frac{1}{N}$ strategy. If distributions implied in option prices actually increase the precision of estimated expected returns, their certainty equivalents should be closer to $\frac{1}{N}$'s than those obtainable from the time-series models. The portfolio problem is solved for four international equity indices, and this, as noted by De Miguel, Garlappi and Uppal (2009), increases the chance that the naïve strategy performs well, because portfolios have lower idiosyncratic risk than individual stocks. The results show that option-based models generate higher certainty equivalents than the $\frac{1}{N}$ rule, especially when limited short-selling is allowed, while time-series models have consistently lower performance than the naïve strategy. The wedge between the $\frac{1}{N}$ allocation and option-based models is largest for portfolios consisting of three assets. Transaction costs are very low for the $\frac{1}{N}$ rule, which requires little rebalancing, while they are relatively high for option-based strategies. As a consequence, they reduce, but do not eliminate, the economic significance of the results.

Section 3.2 reviews the relevant literature, Section 3.3 describes the methodology, Section 3.4 presents the data and the results, and Section 3.5 concludes.

3.2 Literature review

I now discuss the literature on option-implied distributions, models for dynamic skewness and kurtosis, and copula theory, which are at the basis of the methodology used to solve the portfolio problem.

3.2.1 Risk neutral distributions of expected returns

There are different ways of obtaining distributions of expected returns from option prices. Some exploit the fact that the distribution is equal to the second derivative of the call price function with respect to the strike price (Breeden and Litzenberger (1978)). To compensate for the limited number of observable strike prices, the volatility smile is usually fitted with non-parametric techniques, like parabolic functions (Shimko (1993)) or cubic and fourth order splines (Andersen and Wagener (2002)). A major problem is that liquid options aren't usually far in- or out-of-the-money, so the shape of the tails of the implied distributions depends substantially on how the volatility smile is extrapolated. Other approaches require distributional assumptions on the process of the underlying. The reduced flexibility is compensated by robustness to outliers, which is quite valuable when there is only a small number of liquid strike prices available. Jondeau and Rockinger (2000) compare several methods and conclude that a diffusion process with jumps is well suited for long maturity options, while a mixture of lognormals performs better for short maturities. Other studies have used binomial trees (Rubinstein (1994), Jackwerth and Rubinstein (1996)) and kernel regressions (Aït-Sahalia and Lo (1998)).

3.2.2 Models for dynamic skewness and kurtosis

Fama (1965) provided the first evidence that stock returns are not normally distributed, and since then the literature has shown that excess kurtosis and skewness are a common characteristic of the unconditional distribution of daily equity returns (Ané and Geman (2000)). The ARCH family provides a dynamic model that accounts for volatility clustering and produces excess kurtosis even with normal innovations. Departure from normality has been first studied by Bollerslev (1987) and Nelson (1991). Hansen (1994) proposed a class of autoregressive density functions (ARD), which build on the GARCH structure and make the parameters controlling the third and fourth moments time varying. Previous research had already introduced distributions with dynamic moments beyond the second (Gallant, Hsieh, and Tauchen (1997)) but Hansen (1994)'s formulation is more parsimonious and is based on the familiar Student t. Rockinger and Jondeau (2002) estimate a conditional density from moment conditions through the entropy principle: this imposes a low amount of a-priori information, but requires a rather complex estimation procedure. They find that the model is flexible enough to span a greater set of skewness and kurtosis values than Hansen (1994)'s skew t.

In continuous time, normal diffusion processes have been augmented with stochastic volatility and jump components (Aït-Sahalia (2002), Benzoni (2002), Eraker, Johannes and Polson (2003)). Recent contributions also added jumps in volatility (Chernov, Gallant, Ghysels and Tauchen (2003), Eraker (2004)). Most of these models can explain the cross section of option prices quite well but, from a term structure perspective, Das and Sundaram (1999) show that jump-diffusions and stochastic volatility models cannot account for both long and short term anomalies.

While the behavior of daily returns can't be easily reconciled with a normal distribution, longer horizons generate weaker rejections of the normality assumption (Upton and Shannon (1979)). A natural question is whether models for dynamic skewness and kurtosis are justified for monthly data. Even if monthly returns tend to be more "well behaved" than the daily counterpart, there is evidence that the return series used in this paper have significant skewness and excess kurtosis (Harris and Küçüközmen (2001)).

3.2.3 Analysis of dependence

The multivariate normal distribution has been under intense scrutiny for quite some time in the equity returns literature. Three of its characteristics - absence of tail dependence, symmetry and linear dependence - are at odds with the evidence accumulated over the years (Ang and Bekaert (2002), Ang and Chen (2002), Erb, Harvey and Viskanta (1994)). Alternative models either focus on a particular aspect of the distribution, as in the case of extreme value analysis (Longin and Solnik (2001)), or introduce a level of complexity that creates estimation problems (Bauwens and Laurent (2005)). The main advantage of using copula functions to generate multivariate models is flexibility, both from the computational and descriptive point of view. Parameters can be estimated separately for each marginal and for the copula, and the loss of efficiency does not compromise consistency and asymptotic normality (Patton (2005)). At the same time, features like asymmetric dependence structures or tail dependence can be easily described by Student t or Archimedean copulas (see Gagliardini and Gourieroux (2004), Nelsen (1999) for references). Archimedean copulas, in particular, are one of the most useful families. They can model a wide range of dependence structures and can be extended to generate stochastic dependence (Dias and Embrechts (2004)). Patton (2004), for instance, shows that time-varying asymmetry in the dependence structure is important for portfolio choice. A major shortcoming of Archimedean copulas is that they can be extended beyond the bivariate case only under restrictive conditions. This helps explaining why the great variety of families used in bivariate applications turns into a near-monopoly of the normal and Student t copulas for higher dimensions.

Recent applications of copula theory include Jondeau and Rockinger (2006), who estimate dynamic normal and Student t copulas on stock-returns series. Cherubini and Luciano (2002) study the option pricing implications of tail dependence. Li (2000) explores the effects of different dependence structures on default correlation models.

3.3 Methodology

This section describes the set-up of the portfolio choice problem and the details of the optimization. It first defines the utility function and then explains how the distribution of joint returns is built. The marginal distributions and the dependence structure are described separately.

3.3.1 The investor problem

Each month, from January 2000 to November 2004, the investor allocates her wealth to four equity indices: S&P 500, Ftse 100, Nasdaq 100, Nikkei 225.

She solves:

$$\mathbf{w}_t^* = \arg\max_{\mathbf{w}_t \in W} E_t[U(\mathbf{w}_t \cdot \mathbf{r}'_{t+1})]$$

where

$$\mathbf{w}_t = [w_{1,t}, ..., w_{N,t}] \, , \, \{w_{i,t}\}_{i=1}^N \ \textit{are portfolio weights}$$

$$\mathbf{r}_t = [r_{1,t}, ..., r_{N,t}] \, , \, \{r_{i,t}\}_{i=1}^N \ \textit{are monthly gross returns}$$

$$U\left(\mathbf{w}_t \cdot \mathbf{r}'_{t+1}\right) = \begin{cases} \frac{(\mathbf{w}_t \cdot \mathbf{r}_{t+1})^{1-\gamma}}{1-\gamma} & \text{for } \gamma \neq 1 \\ \log\left(\mathbf{w}_t \cdot \mathbf{r}_{t+1}\right) & \text{for } \gamma = 1 \end{cases}$$

$$E_t[U(\mathbf{w}_t \cdot \mathbf{r}'_{t+1})] = \int \cdots \int U(\mathbf{w}_t \cdot \mathbf{r}'_{t+1}) f(\ln(\mathbf{r}_{t+1})) d\mathbf{r}_{t+1}$$

The multivariate density function $f(\ln(\mathbf{r}_{t+1}))$ is equal to the product of the marginal densities $m(\ln(r_{i,t+1}))$ and the copula density $c(M_1(\ln(r_{1,t+1})), ..., M_N(\ln(r_{N,t+1})))$, where $M_i(\ln(r_{i,t+1}))$ is the marginal cumulative probability function of index i's returns:

$$f(\ln(\mathbf{r}_{t+1})) = c(M_1(\ln(r_{1,t+1})), ..., M_N(\ln(r_{N,t+1}))) \prod_i m_i(\ln(r_{i,t+1}))$$

The optimization is performed under two sets of shortselling constraints. In the first case, shortselling is not allowed:

$$W_C = \{[w_{1,t}, ..., w_{N,t} \in [0,1]^N : \sum_i w_{i,t} = 1, \forall t\}$$

In the second case a limited amount of shortselling is possible (this is referred to as the *unconstrained* optimization throughout the paper):

$$W_U = \{[w_{1,t}, ..., w_{N,t}] \in [-2, 2(N-1)+1]^N : \sum_i w_{i,t} = 1, \forall t\}$$

Shortselling is limited because the comparison among different models should not be based on returns driven by extreme positions. Results of completely unconstrained optimizations (not reported) show that option-based models typically produce much larger returns. For the Ftse 100 - S&P500 portfolio, for instance, the gross return over the five years can be as high as 1.70 (0.95 with limited shortselling), but is mainly driven by three months with shortselling in excess of 5,000%.

The portfolio problem does not allow to invest in the riskless asset. This is inefficient (Ang and Bekaert (2002)), but preventing the investor from taking a safe position is a way of "stress testing" the models, because the certainty equivalent depends only on the model's ability to describe expected returns.

The choice of the utility function warrants further discussion. The optimization problem is solved as if the investor were myopic, even with $\gamma > 1$. The reason is that liquid options have short maturity, and a feasible multi-period model would not be very different from a myopic one. Different values of γ can still give useful insights in the asset allocation choices.

The details of the optimization procedure are as follows. At time t, after having estimated $f(\ln(\mathbf{r}_{t+1}))$ as described in the following sections, the expected value of the utility function is approximated with a Montecarlo simulation:

1. a matrix S with N columns and T rows ($T = 400,000$ for $N = 3,4$; $T = 10,000$ for $N = 2$) is generated according to a standardized N-variate normal (or Student t).

2. the standardized univariate normal Φ (or Student t, T_v) cumulative function is applied to each column of S to obtain a $T \times N$ matrix U.

3. the inverse integral transformation of each estimated marginal distribution is applied to the corresponding column of U to obtain a $T \times N$ matrix R, whose generic element (i, j) is then defined as:

$$R_{[i,j]} = M_j^{-1}(F(S_{[i,j]})), \qquad F = \Phi \text{ or } T_v.$$

4. the expected utility is approximated by:

$$\hat{E}_t[U(\mathbf{w}_t \cdot \mathbf{r}'_{t+1})] = \frac{1}{T}\sum_i U(\exp(R_{[i,1]}w_{1,t} + ... + R_{[i,N]}w_{N,t}))$$

Simulated returns outside the 0.1% and 99.9% percentiles (0.2% and 99.8% for Nasdaq 100) are winsorized. The number T of simulated returns has been chosen by trading off accuracy (as defined by the average variance of gross returns on the portfolios over 20 simulated optimizations) and computational speed. Average variance appears to be a hyperbolic function of T: for $N = 4$, at $T = 400,000$ it is half than at $T = 100,000$, and it is further halved only for $T = 2,500,000$.

3.3.2 Distributions estimated on historical returns

The time-series process presented below is an extension of the AR(1)-GARCH(1,1) model. It accounts for dynamic skewness and kurtosis, where the dynamics are obtained through an autoregressive structure for the skewness (λ_t) and the degrees-of-freedom (η_t) parameters of Hansen (1994)'s skew t. The logistic transformation is used to make sure that they remain within their bounds [$\lambda_t \in (-1,1), \eta_t \in (2,\infty)$]. Letting y_t be returns for month t, the model is defined as follows:

$$y_t = k + \alpha_m y_{t-1} + \varepsilon_t$$

$$z_t = \frac{1}{\sqrt{h_t}}\varepsilon_t$$

$$g\left(z_t|\eta_t,\lambda_t\right) = \begin{cases} d_t e_t \left(1 + \frac{1}{\eta_t - 2}\left(\frac{d_t z_t + f_t}{1-\lambda_t}\right)^2\right)^{-(\eta_t+1)/2} & z_t < -f_t/d \\ d_t e_t \left(1 + \frac{1}{\eta_t - 2}\left(\frac{d_t z_t + f_t}{1+\lambda_t}\right)^2\right)^{-(\eta_t+1)/2} & z_t \geq -f_t/d_t \end{cases}$$

$$d_t = \sqrt{1 + 3\lambda_t^2 - f_t^2}$$

$$e_t = \frac{\Gamma(\frac{\eta_t+1}{2})}{\sqrt{\pi(\eta_t-2)}\Gamma(\frac{\eta_t}{2})}$$

$$f_t = 4\lambda_t e_t(\frac{\eta_t - 2}{\eta_t - 1})$$

where the laws of motion for the parameters are:

$$h_t = k_v + \alpha_v \varepsilon_{t-1}^2 + \beta_v h_{t-1}$$

$$\eta_t = 2.1 + \frac{47.9}{1+\exp(\hat{\eta}_t)} \quad \text{with} \quad \hat{\eta}_t = k_d + \alpha_d (\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}})^2 + \beta_d \hat{\eta}_{t-1}$$

$$\lambda_t = \frac{1.98}{1+\exp(-\hat{\lambda}_t)} - .99 \quad \text{with} \quad \hat{\lambda}_t = k_s + \alpha_s \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \beta_s \hat{\lambda}_{t-1}$$

Models with time varying degrees of freedom are typically difficult to estimate, so I trade off efficiency for parameters' stability with a two-steps estimation. First, the AR(1) process is fitted through quasi-maximum likelihood. Then, the dynamics of variance, skewness and degrees of freedom are estimated on the residuals via exact maximum likelihood. The results are remarkably robust to different initial values.

The AR(1) describes the series of returns reasonably well and captures the low autocorrelation. To quantify the effects of misspecified conditional mean on portfolio optimization, an ARMA(1,1) is also fitted. Given the low autocorrelation of the returns series, the ARMA(1,1) is likely to yield estimates for the autoregressive and moving average parameters that are very similar in value but with opposite sign. This induces spurious autocorrelation (often alternating in sign for odd/even lags) and increases the volatility of the conditional mean.

The time series of Ftse 100 and Nasdaq 100 are quite short (about 250 monthly observations), and estimating the full model may be problematic. To avoid unstable estimates, I use a traditional AR(1)-GARCH(1,1) with normal innovations for the Ftse 100 and Nasdaq 100.

3.3.3 Distributions implied in option prices

Jondeau and Rockinger (2000) compare several methods of extracting risk neutral distributions from option prices. For long maturity options, they find supportive evidence for a jump-diffusion, while a mixture of log-normals is more suited for shorter maturities. The estimation procedure follows Jondeau and Rockinger (2000), minimizing the sum of squared differences between observed and theoretical option prices. A skew t is then fitted to returns generated by a random draw of 1,000 observations from the mixture of normals. The next step is to transform risk neutral into objective distributions. Rubinstein (1994) suggests that simply adding the risk premium to the risk neutral distribution is a good approximation. As a first estimate of the objective distribution, I shift the risk neutral one to the right by adding an arbitrary but reasonable 7% to risk neutral expected returns.

I also rely on the power utility specification of Rosenberg and Engle (2002)'s empirical pricing kernel, which is estimated on S&P 500 returns:

$$M(r_{t+1}; \delta_0, \delta_1) = \delta_0 (r_{t+1})^{-\delta_1}$$

where $\delta_0 = 1.0051$, $\delta_1 = 7.56$ and r_{t+1} is the gross return on the asset.

The five-years gap between the estimation of the pricing kernel and the beginning of the portfolio choice analysis may create spurious differences between the certainty equivalents of the models based on historical returns and option prices. The pricing kernel, however, can not be estimated on the same dataset used to solve the portfolio choice problem, because there would not be any difference between the objective distributions obtained from option prices and those estimated from the data.

3.3.4 The dependence structure

I use normal and Student t copulas to model dependence among the series of stock indices returns. The comparison between certainty equivalents from the two copulas quantifies the contribution of accounting for tail dependence on investor's utility. In light of the relatively small number of available observations, I use static, elliptic copulas in the interest of parameter stability.

When coupled with skew marginal distributions, like in this case, elliptical copulas don't generate elliptical multivariate distributions. This means that correlations must be estimated by maximum likelihood rather than through sample moments. Copulas' parameters are estimated only once, with kernel densities for the marginals, from Feb. 1983 to Dec. 2004. Estimating the normal copula has been quite straightforward[2], while the presence of degrees of freedom makes the Student t copula less tractable. I then use the methodology of Mashal and Zeevi (2002), which yields estimates very similar to full maximum likelihood. Kendall τ (ρ^τ) is a suitable measure of dependence for non elliptical copulas, and it is related to the linear correlation ρ^s by: $\rho^s = \sin\left(\frac{\pi}{2}\rho^\tau\right)$. The correlation matrix can be estimated as:

$$\rho^s = [\rho^s_{ij}]_{i,j=1,\ldots,N} = [\sin(\tfrac{\pi}{2}\rho^\tau)]_{i,j=1,\ldots,N}$$

Degrees of freedom are then estimated by maximum likelihood. The correlation matrix is not guaranteed to be positive definite, but this has never been a problem in practice.

The normal copula is:

[2] The positive definiteness of the covariance matrix has been enforced by setting the likelihood function to -∞ (more precisely, to the lower overflow threshold in Matlab, which is still a number) if the determinant is less than 1e-8. This is correct only if the the likelihood is decreasing for the determinant going between 1e-8 and zero, however the estimation has proven robust across on several asset combinations and subsamples.

$$C(u_1, ..., u_N; \rho) = \Phi_\rho(\Phi^{-1}(u_1), ..., \Phi^{-1}(u_N))$$

with density:

$$c(u_1, ..., u_N; \rho) = \frac{\phi_\rho(\Phi^{-1}(u_1), ..., \Phi^{-1}(u_N))}{\prod_{i=1}^{N} \phi(\Phi^{-1}(u_i))}$$

The Student t copula is:

$$C(u_1, ..., u_N; \rho, \upsilon) = T_{\rho,\upsilon}(T_\upsilon^{-1}(u_1), ..., T_\upsilon^{-1}(u_N))$$

with density:

$$c(u_1, ..., u_N; \rho, \upsilon) = \frac{t_{\rho,\upsilon}(T_\upsilon^{-1}(u_1), ..., T_\upsilon^{-1}(u_N))}{\prod_{i=1}^{N} t_\upsilon(T_\upsilon^{-1}(u_i))}$$

Where Φ_ρ, $T_{\rho,\upsilon}$, ϕ_ρ, $t_{\rho,\upsilon}$ are the cumulative and density functions of the multivariate standardized normal and Student t with correlation matrix ρ and υ degrees of freedom, and ϕ, t_υ, Φ^{-1}, T_υ^{-1} are the probability distributions and inverse functions of the univariate standardized normal and Student t with υ degrees of freedom.

3.4 Data and results

Monthly returns for S&P 500, Nikkei 225 (from May 1950), Nasdaq 100 (from Feb. 1983) and Ftse100 (from Feb. 1978) are obtained from Datastream and Global Financial Data. Daily closing prices and trading volume for call and put options are provided by the Chicago Board of Options Exchange (CBOE), the London International Financial Futures Exchange (LIFFE) and the Osaka Securities Exchange (OSE). Riskless rates are one-month LIBORs.

The last trading day is the third Thursday of each expiration month[3] and the number of strike prices varies for each expiration month, being typically greater for options expiring on the quarterly cycle. I filter out illiquid contracts (less than 5 trades a day), those that violate basic and vertical spread arbitrage. The number of options used to estimate each implied distribution is typically 14-20.

3.4.1 Conditional and unconditional returns statistics

Table 3.1 shows the first four unconditional moments of the indices. Excess kurtosis is very high only for Ftse 100 but is noticeable in all of the series. Jarque-Bera tests easily reject the assumption of normality, but unconditional distributions can be non-normal even if the conditional ones are. Ljung-Box tests (not reported) provide evidence in favor of models with time varying third and fourth moments. They show little autocorrelation in returns but support time-dependence in higher moments, in particular for Nasdaq 100 and S&P500. Ftse 100 has the highest unconditional kurtosis and skewness, and both the Ljung-Box and Kolmogorov-Smirnov tests (Table 3.4) suggest that conditional volatility models do not provide a good fit (although the AR-GARCH cannot be rejected at $\alpha = 5\%$). This is probably due to the very low return on November 1987, which increases unconditional third and fourth moments but is an isolated event. A model with both stochastic volatility and jumps would probably increase the goodness of fit and reduce the standard errors reported in Table 3.4.

[3] Nikkei 225 options expire one week before the others. The implied distribution is always estimated 30 days before the last trading day, to avoid liquidity concerns. but the optimization is carried out as if they expired on the same date as the others. In other words, the risk neutral distribution for Nikkei 225 is assumed to be the same on the second and third Thursdays of each month.

The conditional means of the four series are modeled with both an AR(1) and ARMA(1,1). ARMA models are known to produce spurious autocorrelation when the series is weakly autocorrelated. Detecting the misspecification is not always easy. The ARMA(1,1)-GARCH(1,1) fitted to the Ftse 100 has an Akaike criterion equal to −3.262369, while the same statistic is −3.253499 for the AR(1)-GARCH(1,1). A Kolmogorov-Smirnov test on the ARMA specification of the S&P 500 has quite a high p-value, too. The correlogram, however, is quite indicative of the misspecification.

3.4.2 Evaluation of portfolio choice models

Figure 3.1 plots certainty equivalents (CE) against risk aversion for the five models defined in Section 3.3, estimated on three portfolios with unconstrained optimization. They represent (to my judgement, admittedly) the best, worst and *typical* performance of the option-based models with respect to the naïve strategy. Historical models are always outperformed, which confirms the findings of De Miguel, Garlappi and Uppal (2009) that historical models often produce worse results than the simple $\frac{1}{N}$ strategy. Interestingly, the “worst scenario” refers to the portfolio Nasdaq 100 - S&P500, which has a very high correlation. Figure 3.2 shows CE for the same models and portfolios as Figure 3.1, this time for constrained optimization. While the differences are smaller, the pattern is unchanged.

Table 3.5 reports the mean CE for the normal and Student t copulas, averaged across all portfolios and models. The differences are quite small in economic terms and not statistically significant, both in the constrained and unconstrained optimization case. This suggests that the choice of dependence structure is a rather secondary issue, at least in the framework I consider. Table 3.6 shows CE averaged across risk aversion, constrained/uncostrained optimization, portfolio choice models and copulas for all two-assets portfolios, with the exception of S&P500-Nasdaq100. The reason for excluding this pair is that its high correlation generates extreme portfolio weights, which would significantly affect the results. Tables 3.7 and 3.8 have the same structure as Table 3.6. The first includes three-asset portfolios, again excluding those with the highly correlated S&P500-Nasdaq100 pair, while the second refers to the only four-asset portfolio.

Tables 3.6-3.8 contain several interesting results. First of all, option-based models tend to perform better than the others when optimization is constrained, and much better with unconstrained optimization. Furthermore, CE generated by option-based models are remarkably stable across two- and three-assets portfolios. They are lower in Table 3.8, which shows the results for the only four-assets portfolios, that includes the highly correlated S&P500-Nasdaq100 pair. Option-based models still outperform those estimated on historical returns.

Finally, a note on transaction costs. With a trading fee of 0.1%, the performance of the naïve strategy is hardly affected, because it requires little rebalancing. For unconstrained optimization, annualized transaction costs are between 0.5 and 2% (depending on risk aversion) for ARMA and option-based models, and about 0.5% for AR models. Costs for constrained models are about 60% smaller. The effect of transaction costs reduces the advantage of using option-implied distributions with respect to the naïve strategy, but does not eliminate it.

3.5 Conclusion

I solve a portfolio choice problem by estimating the distributions of expected returns on cross-sections of option prices, to study whether the information about higher moments that is implied in risk-neutral distributions improves the precision of estimated expected returns. I obtain the corresponding objective distributions by applying the empirical pricing kernel of Rosenberg and Engle (2002) and, alternatively, by shifting the distribution to the right by adding the risk premium. The second method allows to evaluate the portfolio choice implications of risk-neutral skewness, which imposes a penalty on extreme weights by increasing expected losses. The marginal distributions are joined using copula theory, and the comparison of the performance of portfolios based on the normal and t copulas provides evidence on the economic value of accounting for tail dependence. The results show that option-based models outperform the $\frac{1}{N}$ rule, especially with (limited) short-selling. Time-series models, on the other hand, produce certainty equivalents that are consistently lower than the naïve strategy's. The wedge between option-based models and the $\frac{1}{N}$ rule is greater for three- rather than two-dimensional portfolios. Transaction costs are relatively high for the strategies derived from options, and reduce, but don't eliminate, the economic significance of the results.

Figure 3.1: Certainty equivalents, unconstrained optimization.

The vertical axes report certainty equivalents for varying risk aversion coefficients (x-axes). On the left column, the four models are: options-based correcting risk neutrality by adding the risk premium (dashed line) and with the pricing kernel (dotted line), historical-AR(1) (solid line) and $\frac{1}{N}$ (dash-dot line). In the right column the solid line represent the historical-ARMA(1,1). Optimization is unconstrained. The certainty equivalent is computed over the realized returns from Feb. 2000 to Dec. 2004. The three portfolios are (across rows): Nikkei 225 and Nasdaq 100, Nasdaq 100 and S&P500, and Nikkei 225, Ftse 100 and S&P500. They represent, respectively, the best, worst and "typical" relative performance of the options-based models under unconstrained optimization.

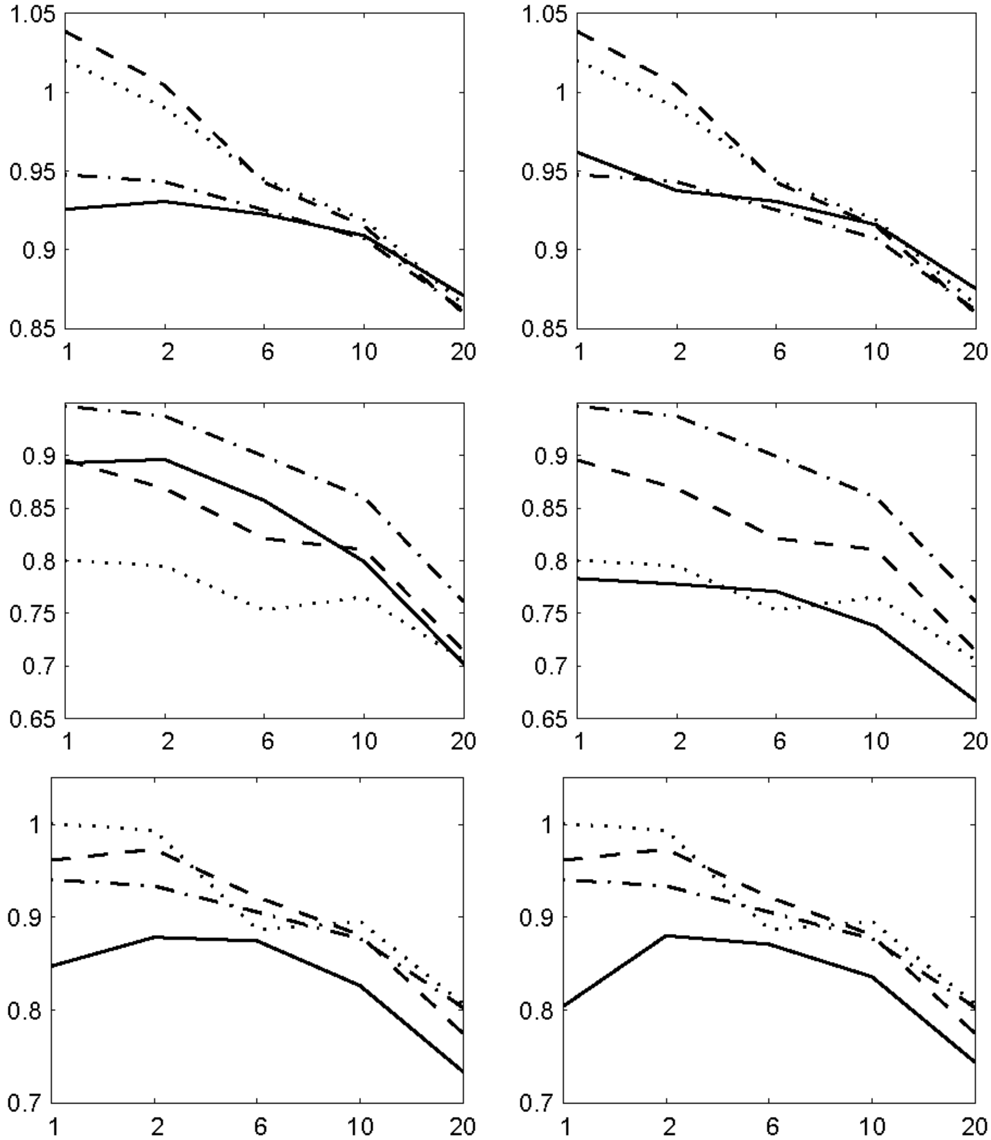

Figure 3.2: Certainty equivalents, constrained optimization.

The vertical axes report certainty equivalents for varying risk aversion coefficients (x-axes). On the left column, the four models are: options-based correcting risk neutrality by adding the risk premium (dashed line) and with the pricing kernel (dotted line), historical-AR(1) (solid line) and $\frac{1}{N}$ (dash-dot line). In the right column the solid line represent the historical-ARMA(1,1). Optimization is constrained. The certainty equivalent is computed over the realized returns from Feb. 2000 to Dec. 2004. The three portfolios are (across rows): Nikkei 225 and Nasdaq 100, Nasdaq 100 and S&P500, and Nikkei 225, Ftse 100 and S&P500. They represent, respectively, the best, worst and "typical" relative performance of the options-based models under constrained optimization.

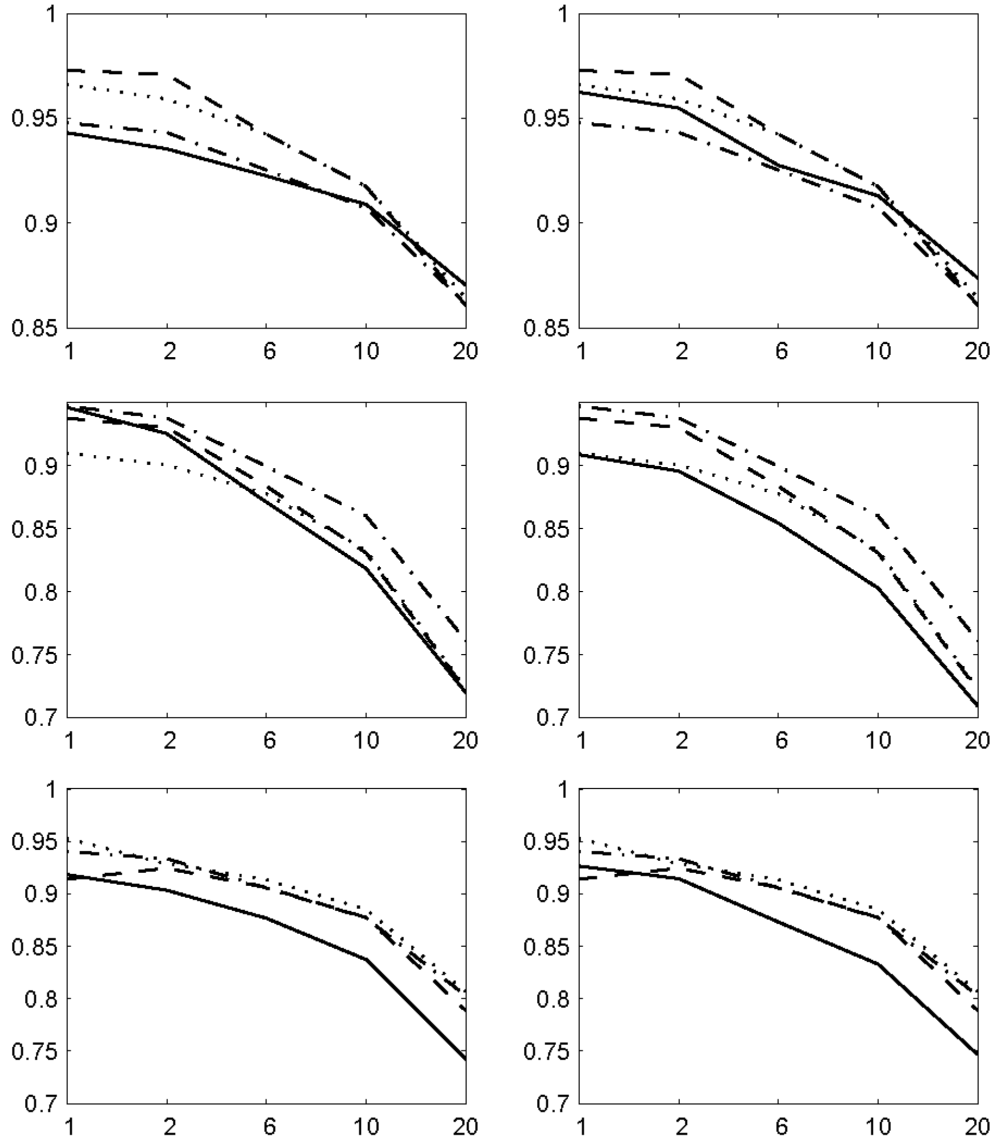

Table 3.1: Returns statistics, 2000-2004

	Nikkei 225	Nasdaq 100	Ftse 100	S&P500
Mean	0.00731	0.00972	0.00738	0.00638
St. Dev.	0.05409	0.07429	0.04718	0.04055
Skewness	-0.31	-0.59	-1.46	-0.47
Kurtosis	4.44	4.51	10.80	3.87
Jarque-Bera	67.63	40.64	929.06	45.43
p-value	0	0	0	0

Table 3.2: Returns statistics, 2000-2004

Estimated correlations for the normal (upper half matrix) and Student-t (lower half matrix) copula in the four-dimensional case. Estimated degrees of freedom for the Student-t are 16.594.

	Nikkei 225	Nasdaq 100	Ftse 100	S&P500
Nikkei 225	1	0.356	0.3929	0.4076
Nasdaq 100	0.2859	1	0.5327	0.8339
Ftse 100	0.3482	0.4847	1	0.6349
S&P500	0.3699	0.8206	0.5833	1

Table 3.3: Predicted one-period-ahead moments across different models. Summary statistics.

Minimum, median and maximum [Min Median Max] one-period-ahead moments, over the 59 months from Jan. 2000 to Nov. 2004, for the processes fitted on historical data and for the risk neutral distributions extracted from option prices. Note that, in the latter case, the mean is not corrected for risk neutrality (add 0.0057). For each index, the first row refers to the historical (AR), the second to the historical (ARMA) and the third to the options based one. When the skewness parameter is equal to .9 or -.9 it hits the upper and lower bound (higher absolute values do not change the shape much but can give numerical problems during the estimation). D.o.f stands for degrees of freedom.

	Mean			**Variance**			**Skewness**			**D.o.f.**		
	Min	Median	Max	Min	Median	Max	Min	Median	Max	Min	Median	Max
Nikkei 225												
a)	0.0088	0.0093	0.0106	0.0019	0.0043	0.0106	-0.1437	-0.0001	0.1232	2.9	7.6	12.8
b)	-0.0006	0.001	0.0096	0.0019	0.0043	0.0097	-0.3326	-0.1083	0.1113	3.2	7.6	13.7
c)	-0.0627	-0.0013	0.0558	0.0023	0.0077	0.0916	-0.9	-0.0086	0.9	2.1	4.1	341.9
Nasdaq 100												
a)	0.0147	0.0155	0.0198	0.0035	0.006	0.0126						
b)	-0.0423	0.0085	0.0332	0.0042	0.0062	0.0134						
c)	-0.0633	0.0027	0.0466	1.30E-05	0.0162	0.2118	-0.9	-0.2701	0.9	2.1	4.8	341.9
Ftse 100												
a)	0.0106	0.0117	0.0137	0.0019	0.0024	0.0026						
b)	0.0006	0.0082	0.0242	0.0015	0.0023	0.0033						
c)	-0.0222	-0.0001	0.0045	0.0009	0.0029	0.0197	-0.5346	-0.3286	-0.0467	3.6	8.9	341.9
S&P500												
a)	0.0073	0.0077	0.0087	0.0012	0.0023	0.0047	-0.2933	-0.2215	-0.1519	2.2	12.4	16.4
b)	0.0016	0.0063	0.012	0.0012	0.0023	0.0047	-0.2773	-0.2297	-0.141	2.2	13.6	17.4
c)	-0.0112	0.0009	0.0154	1.03E-05	0.0054	0.088	-0.8022	-0.2241	0.7344	2.1	3.2	347.9

Table 3.4: Parameters for the one-period-ahead distributions. Summary statistics.

Coefficient estimates for the processes fitted to the four index returns series from the beginning of the series to November 2004. Models with conditional means following an AR(1) are indicated with "I". Models with (mis-specified) ARMA(1,1) conditional mean are indicated with "II". Panels contain the estimates for Nasdaq 100 (a), Ftse 100 (b), Nikkei 225 (c) and S&P500 (d). The p-value for the two-tailed Kolmogorov-Smirnov test over the estimation period is also reported (H_0: the cdf of the hypothesized distribution is equal to the one generating the data).

a)		**Mean**			**Variance**		KS-test
	κ_m	α_m	β_m	κ_v	α_v	β_v	p-value
I	0.0115	0.1029		0.0005	0.0672	0.8396	0.88
	0.0051	0.0745		0.0004	0.05	0.1286	
II	0.0114	0.3089	-0.2077	0.0005	0.0662	0.841	0.84
	0.0052	0.5503	0.5848	0.0004	0.0499	0.1282	
b)		**Mean**			**Variance**		KS-test
	κ_m	α_m	β_m	κ_v	α_v	β_v	p-value
I	0.0091	-0.0436		0.0009	-0.0206	0.6413	0.0641
	0.0029	0.0363		0.0011	0.0181	0.5086	
II	0.0072	-0.883	0.9376	0.0001	0.0929	0.8454	3.00E-05
	0.0028	0.0597	0.0423	0.0001	0.0381	0.0641	
c)		**Mean**			**Variance**		KS test
	κ_m	α_m	β_m	κ_v	α_v	β_v	p-value
I	0.0073	0.0671		7.00E-05	0.132	0.8576	0.9811
	0.0022	0.0391		4.00E-05	0.0483	0.0461	
II	0.0073	-0.7936	0.8323	7.00E-05	0.1331	0.8592	0.9155
	0.0021	0.0999	0.0928	4.00E-05	0.0457	0.0428	
		Skewness			**D.o.f**		
	κ_s	α_s	β_s	κ_d	α_d	β_d	
I	-0.0215	-0.0781	0.9194	-0.1076	0.2546	0.8705	
	0.0198	0.0439	0.0571	0.2029	0.1322	0.0712	
II	-0.0237	-0.0775	0.9188	-0.1198	0.2637	0.872	
	0.0199	0.044	0.0576	0.2383	0.1246	0.0806	
d)		**Mean**			**Variance**		KS test
	κ_m	α_m	β_m	κ_v	α_v	β_v	p-value
I	0.0063	0.0368		1.00E-04	0.1305	0.8111	0.9256
	0.0016	0.0392		2.00E-05	0.042	0.0361	
II	0.0064	-0.5336	0.574	1.00E-04	0.1335	0.8107	0.9533
	0.0016	0.3779	0.3665	6.00E-05	0.0439	0.0498	
		Skewness			**D.o.f**		
	κ_s	α_s	β_s	κ_d	α_d	β_d	
I	-0.2408	0.0592	0.521	0.2424	0.4588	1.00E-08	
	0.3192	0.1088	0.612	1.0293	0.246	0.129	
II	-0.2578	-0.04146	0.5073	0.0613	0.5094	1.00E-08	
	0.4374	0.1111	0.8083	1.03	0.3616	0.1314	

Table 3.5: Model performance, across copula specification.

Average out-of-sample certainty equivalent (in %) across the type of copula used for the dependence structure. N and T denote the normal and Student-t copula. The average is calculated across all portfolios. Jackknife standard errors are reported below the average CE. γ is risk aversion and *constr.*, *unc.* indicate constrained and unconstrained optimization. Feb. 2000 to Dec. 2004.

	$\gamma = 1$		$\gamma = 2$		$\gamma = 6$		$\gamma = 10$		$\gamma = 20$	
	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.
N)	93.347	92.381	92.603	92.293	89.778	88.6	86.77	86.187	78.417	77.147
	0.195	0.312	0.195	0.266	0.201	0.236	0.214	0.235	0.267	0.297
T)	93.343	92.285	92.631	92.207	89.75	88.576	86.685	86.123	78.565	77.334
	0.195	0.323	0.195	0.276	0.201	0.242	0.214	0.244	0.262	0.311

Table 3.6: Model performance, two-dimensional case.

Average out-of-sample certainty equivalent (in %) across different risk aversions and models. Panels A and B refer to the normal and Student-t copula. Models are as follows: a) historical-ARMA; b) historical-AR; c) options plus equity premium; d) options and pricing kernel e) $\frac{1}{N}$. γ is the risk aversion and *constr.*, *unc.* indicate constrained and unconstrained optimization. Only two-dimensional portfolios are included. The portfolio Nasdaq100-S&P500 is excluded because of the high correlation between the two series. Jackknife standard errors are reported below the average CE. Feb. 2000 to Dec. 2004.

	$\gamma = 1$		$\gamma = 2$		$\gamma = 6$		$\gamma = 10$		$\gamma = 20$	
	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.
	Panel A									
a)	93.342	90.316	92.314	91.282	88.416	88.471	85.068	84.891	75.951	75.23
	0.073	0.132	0.074	0.101	0.078	0.083	0.083	0.086	0.105	0.108
b)	92.669	90.871	91.724	91.241	88.748	88.82	85.352	85.41	76.343	76.176
	0.07	0.088	0.071	0.073	0.074	0.074	0.08	0.08	0.102	0.102
c)	93.787	96.184	93.376	95.214	90.544	90.967	87.282	87.275	78.876	78.262
	0.068	0.097	0.067	0.086	0.07	0.074	0.074	0.076	0.092	0.093
d)	94.507	96.035	93.887	94.503	90.965	91.392	88.023	88.17	80.19	79.28
	0.071	0.106	0.07	0.092	0.071	0.076	0.074	0.076	0.087	0.089
e)	94.012	94.012	93.326	93.326	90.524	90.524	87.645	87.645	80.29	80.29
	0.062	0.062	0.063	0.063	0.067	0.067	0.071	0.071	0.087	0.087
	Panel B									
a)	93.098	90.145	92.212	90.772	88.392	88.271	85.137	84.932	76.692	76.238
	0.074	0.134	0.075	0.104	0.078	0.084	0.083	0.087	0.101	0.106
b)	92.616	90.318	91.694	90.907	88.552	88.479	85.258	85.24	76.872	76.935
	0.069	0.091	0.07	0.074	0.074	0.074	0.08	0.08	0.1	0.1
c)	93.736	96.107	93.385	95.267	90.489	90.974	87.258	87.672	78.972	74.966
	0.068	0.101	0.067	0.088	0.07	0.074	0.074	0.075	0.089	0.15
d)	94.435	96.177	93.667	94.674	90.914	91.388	87.765	88.186	79.997	80.152
	0.071	0.109	0.07	0.094	0.071	0.076	0.075	0.076	0.088	0.088
e)	94.012	94.012	93.326	93.326	90.524	90.524	87.645	87.645	80.29	80.29
	0.062	0.062	0.063	0.063	0.067	0.067	0.071	0.071	0.087	0.087

Table 3.7: Model performance, three-dimensional case

Average out-of-sample certainty equivalent (in %) across different risk aversions and models. Panels A and B refer to the normal and Student-t copula. Models are as follows: a) historical-ARMA; b) historical-AR; c) options plus equity premium; d) options and pricing kernel e) $\frac{1}{N}$. γ is the risk aversion and *constr.*, *unc.* indicate constrained and unconstrained optimization. Only three-dimensional portfolios are included. The portfolio Nasdaq100-S&P500 is excluded because of the high correlation between the two series. Jackknife standard errors are reported below the average CE. Feb. 2000 to Dec. 2004.

	$\gamma = 1$		$\gamma = 2$		$\gamma = 6$		$\gamma = 10$		$\gamma = 20$	
	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.
	Panel A									
a)	91.825	87.818	90.957	89.748	87.384	87.226	84.35	81.456	74.826	65.739
	0.046	0.11	0.048	0.084	0.051	0.058	0.054	0.063	0.072	0.113
b)	91.997	88.336	90.658	88.736	88.119	87.602	85.248	85.364	75.762	75.077
	0.046	0.063	0.046	0.047	0.046	0.047	0.05	0.051	0.066	0.068
c)	92.41	94.05	92.505	95.794	90.732	91.988	87.761	88.362	78.256	74.947
	0.041	0.068	0.039	0.058	0.039	0.043	0.042	0.043	0.053	0.08
d)	95.086	97.19	93.613	95.901	91.704	84.085	89.305	89.99	81.662	82.082
	0.045	0.077	0.045	0.065	0.04	0.049	0.041	0.043	0.048	0.048
e)	93.987	93.987	93.396	93.396	90.987	90.987	88.509	88.509	82.123	82.123
	0.037	0.037	0.037	0.037	0.039	0.039	0.041	0.041	0.048	0.048
	Panel B									
a)	92.423	84.57	91.661	88.294	87.712	86.973	84.118	83.839	75.483	74.226
	0.046	0.114	0.048	0.086	0.051	0.059	0.055	0.061	0.068	0.075
b)	91.869	87.543	90.86	89.311	87.894	87.773	84.476	83.723	75.389	71.15
	0.046	0.067	0.046	0.05	0.047	0.047	0.05	0.05	0.065	0.08
c)	92.428	95.329	92.423	96.595	90.684	91.921	87.776	88.259	78.961	77.405
	0.041	0.071	0.039	0.061	0.039	0.044	0.041	0.044	0.051	0.058
d)	95.18	98.376	93.586	95.353	91.742	86.743	89.346	89.574	81.721	82.131
	0.045	0.08	0.045	0.069	0.04	0.059	0.041	0.043	0.048	0.048
e)	93.987	93.987	93.396	93.396	90.987	90.987	88.509	88.509	82.123	82.123
	0.037	0.037	0.037	0.037	0.039	0.039	0.041	0.041	0.048	0.048

Table 3.8: Model performance, four-dimensional case.

Out-of-sample certainty equivalent (in %) across different risk aversions and models. Panels A and B refer to the normal and Student-t copula. Models are as follows: a) historical-ARMA; b) historical-AR; c) options plus equity premium; d) options and pricing kernel e) $\frac{1}{N}$. γ is the risk aversion and *constr.*, *unc.* indicate constrained and unconstrained optimization. The figures refer to the only four-dimensional portfolio. Jackknife standard errors are reported below the CE. Feb. 2000 to Dec. 2004.

	$\gamma = 1$		$\gamma = 2$		$\gamma = 6$		$\gamma = 10$		$\gamma = 20$	
	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.	constr.	unc.
	Panel A									
a)	90.013	72.539	89.371	75.927	86.299	78.023	83.429	75.732	74.436	66.277
	0.032	0.092	0.033	0.081	0.037	0.061	0.039	0.061	0.05	0.073
b)	92.155	87.129	90.908	87.249	88.427	86.609	84.75	82.264	76	73.669
	0.03	0.043	0.031	0.036	0.031	0.031	0.036	0.037	0.049	0.058
c)	92.386	91.518	92.845	93.081	91.696	85.789	89.192	85.444	79.748	79.098
	0.032	0.066	0.03	0.061	0.028	0.061	0.03	0.049	0.039	0.041
d)	92.92	89.988	92.036	86.906	89.906	67.988	87.953	81.438	81.647	79.389
	0.034	0.067	0.035	0.067	0.034	0.085	0.033	0.067	0.038	0.046
e)	94.23	94.23	93.588	93.588	90.966	90.966	88.268	88.268	81.323	81.323
	0.027	0.027	0.027	0.027	0.029	0.029	0.03	0.03	0.036	0.036
	Panel B									
a)	90.899	71.709	90.243	73.637	86.671	77.424	83.481	74.195	74.274	62.817
	0.032	0.095	0.033	0.087	0.037	0.065	0.04	0.068	0.05	0.092
b)	91.905	86.915	91.186	86.414	88.233	86.663	84.584	83.18	75.24	71.633
	0.03	0.045	0.032	0.038	0.032	0.033	0.037	0.039	0.049	0.061
c)	92.389	93.423	92.841	94.665	91.702	83.876	89.217	84.352	78.761	81.535
	0.032	0.068	0.03	0.061	0.028	0.074	0.03	0.055	0.04	0.041
d)	93.032	91.71	92.123	89.355	90.022	67.128	88.07	78.748	81.527	79.064
	0.034	0.074	0.035	0.075	0.034	0.077	0.033	0.085	0.038	0.05
e)	94.23	94.23	93.588	93.588	90.966	90.966	88.268	88.268	81.323	81.323
	0.027	0.027	0.027	0.027	0.029	0.029	0.03	0.03	0.036	0.036

Chapter 4

Bibliography

Abreu D.; Brunnermeier, M.K. (2003), *Bubbles and crashes*, Econometrica, 71(1), 173-204

Aït-Sahalia, Y. (2002), *Telling from discrete data whether the underlying continuous-time model is a diffusion*, Journal of Finance, 57(5), 2075-2112

Aït-Sahalia, Y.; Lo, A.W. (1998), *Nonparametric estimation of state price densities implicit in financial asset prices*, Journal of Finance, 53(2), 499-547

Andersen, T.G.; Bollerslev, T.; Diebold, F.X.; Vega, C. (2003), *Micro effects of macro announcements: real time price discovery in foreign exchange*, American Economic Review, 93(1), 38-62

Andersen, A.B.; Wagener, T. (2002), *Extracting risk neutral probability densities by fitting implied volatility smiles: some methodological points and an application to the 3M Euribor futures option prices*, Working paper n.198, European Central Bank

Anderson, E.W.; Ghysels, E.; Juergens, J.L. (2005), *Do heterogeneous beliefs matter for asset pricing?*, Review of Financial Studies, 18(3), 875-924

Anderson, E.W.; Ghysels, E.; Juergens, J.L. (2007), *The impact of risk and uncertainty on expected returns*, Working paper

Ané, T.; Geman, H. (2000), *Order flow, transaction clock, and normality of asset returns*, Journal of Finance, 55(5), 2259-2284

Ang, A.; Bekaert,G. (2002), *International asset allocation with regime shifts*, Review of Financial Studies, 15(4), 1137-1187

Ang, A.; Chen, J. (2002), *Asymmetric correlations of equity portfolios*, Journal of Financial Economics, 63(3), 443-494

Baker, M.; Stein, J.C.; Wurgler, J.C. (2003), *When does the market matter? Stock prices and the investment of equity-dependent firms*, Quarterly Journal of Economics, 118, 969-1005

Baker, M.; Wurgler, J.C. (2006), *Investor sentiment and the cross section of stock returns*, Journal of Finance, 61(4), 1645-1680

Bakshi, G.; Cao, C.; Chen, Z. (1997), *Empirical performance of alternative option pricing models*, Journal of Finance, 52(5), 2003-2049

Balduzzi, P.; Elton E.J.; Green, T.C. (2001), *Economic news and bond prices: evidence from the U.S. Treasury market*, Journal of Financial and Quantitative Analysis, 36(4), 523-543

Barberis, N.; Thaler, R. (2003), *A survey of behavioral finance*, in Handbook of the Economics of Finance, Elsevier

Barinov, A. (2007), *Idiosyncratic volatility, growth options, and the cross-section of returns*, Working paper

Barzel, Y. (1968), *Optimal timing of innovations*, Review of Economics and Statistics, 50(3), 348-355

Bates, D.S. (2003), *Empirical option pricing: a retrospection*, Journal of Econometrics, 116, 387-404

Bauwens, L.; Laurent, S. (2005), *A new class of multivariate skew densities, with applications to GARCH models*, Journal of Business and Economic Statistics, 23(3), 346-354

Beber, A.; Brandt, M.W. (2006), *The effect of macroeconomic news on beliefs and preferences: evidence from the options market*, Journal of Monetary Economics, 53, 1997-2039

Beber, A.; Brandt, M.W. (2008), *Resolving macroeconomic uncertainty in stock and bond markets*, Review of Finance, forthcoming

Benzoni, L. (2002), *Pricing options under stochastic volatility: an empirical investigation*, Working paper

Benzoni, L.; Collin-Dufresne, P.; Goldstein, R.S. (2005), *Can standard preferences explain the prices of out-of-the-money S&P 500 put options?*, Working paper

Bhattacharya, S.; Chatterjee, K.; Samuelson, L. (1986), *Sequential research and the adoption of innovations*, Oxford Economic Papers, 38, 219-243

Black, F. (1986), *Noise*, Journal of Finance, 41(3), 529-543

Blume, M.; Stambaugh, R. (1983), *Bias in computed returns: an application to the size effect*, Journal of Financial Economics, 12, 387-553

Bollen, N. P.B.; Whaley, R.E. (2004), *Does net buying pressure affect the shape of implied volatility functions?*, Journal of Finance, 59(2), 711-753

Bollerslev, T. (1987), *A conditionally heteroskedastic time series model for speculative prices and rates of return*, Review of Economics and Statistics, 69(3), 542-547

Bolton, P.; Harris, C. (1999), *Strategic experimentation*, Econometrica, 67(2), 349-374

Brandt, M.W.; Santa-Clara, P.; Valkanov, R. (2009), *Parametric portfolio policies: exploiting characteristics in the cross section of equity returns*, Review of Financial Studies, forthcoming

Breeden, D.T.; Litzenberger, R.H. (1978), *Prices of state-contingent claims implicit in option prices*, Journal of Business, 51(4), 621-651

Broadie, M.; Chernov, M.; Johannes, M. (2007), *Model specification and risk premia: evidence from futures options*, Journal of Finance 62(3), 1453-1490

Broadie, M.; Chernov, M.; Johannes, M. (2008), *Understanding index options returns*, Review of Financial Studies, forthcoming

Brunnermeier, M.K.; Nagel, S. (2004), *Hedge funds and the technology bubble*, Journal of Finance, 59(5), 2013-2040

Buraschi, A.; Jiltsov, A. (2006), *Model uncertainty and options markets with heterogeneous beliefs*, Journal of Finance, 61(6), 2841-2897

Cagetti, M.; Hansen, L.P.; Sargent, T.J.; Williams, N.(2002), *Robustness and pricing with uncertain growth*, Review of Financial Studies, 15(2), 363-404

Chari, V.V.; Hopenhayn, H. (1991), *Vintage human capital, growth and the diffusion of new technology*, Journal of Political Economy, 99(6), 1142-1165

Chernov, M.; Gallant, A.R.; Ghysels, E.; Tauchen, G. (2003), *Alternative models for stock price dynamics*, Journal of Econometrics, 116(1-2), 225-257

Cherubini, U.; Luciano, E. (2002), *Bivariate option pricing with copulas*, Applied Mathematical Finance, 9(2), 69-85

Chirinko, R.S.; Schaller, H. (2001), *Business fixed investment and "bubbles": the Japanese case*, American Economic Review 91, 663-680

Christensen, B.J.; Prabhala, N.R. (1998), *The relation between implied and realized volatiltiy*, Journal of Financial Economics, 50, 125-150

Constantinides, G.M.; Jackwerth, J.C.; Perrakis, S. (2009), *Mispricing of S&P 500 index options*, Review of Financial Studies, 22(3), 1247-1277

Cooley, T.F.; Greenwood, J.; Yorukoglu, M. (1997), *The replacement problem*, Journal of Monetary Economics, 40, 457-499

Coval, J.D. (2001); Shumway, T. (2001), *Expected option returns*, Journal of Finance, 56(3), 983-1009

Cremers, M.; Driessen, J.; Maenhout, P.; Weinbaum, D. (2008), *Individual stock-option prices and credit spreads*, Journal of Banking and Finance, forthcoming

Das, S.R.; Sundaram, R.K. (1999), *Of smiles and smirks: a term structure perspective*, Journal of Financial and Quantitative Analysis, 34(2), 211-239

David, A.; Veronesi, P. (2002), *Option prices with uncertain fundamentals: theory and evidence on the dynamics of implied volatilities*, Working paper

De Long, J.B.; Shleifer, A.; Summers, L.H.; Waldmann, R.J. (1990), *Noise trader risk in financial markets*, Journal of Political Economy, 98(4), 703-738

De Miguel, A.V.; Garlappi, L.; Uppal, R. (2009), *Optimal versus naïve diversification: how inefficient is the 1/N portfolio strategy?*, Review of Financial Studies, 22(5), 1915-1953

Den Haan, W.J. (2000), *The comovement between output and prices*, Journal of Monetary Economics, 46, 3-30

Dias, A.; Embrechts, P. (2004), *Dynamic copula models for multivariate high-frequency data in finance*, Working paper

Driessen, J.; Maenhout, P. (2007), *An empirical portfolio perspective on option pricing anomalies*, Review of Finance, 11(4), 561-603

Driessen, J.; Maenhout, P.; Vilkov, G. (2009), *The price of correlation risk. Evidence from equity options*, Journal of Finance, forthcoming

Duarte, J.; Jones, C.S. (2007), *The price of market volatility risk*, Working paper

Dubinsky, A.; Johannes, M. (2005), *Earnings announcements and equity options*, Working paper

Dumas, B.; Kurshev, A.; Uppal, R. (2008), *Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility*, Journal of Finance, forthcoming

Ederington, L.H.; Lee, J.H. (1996), *The creation and resolution of market uncertainty: the impact of information releases on implied volatility*, Journal of Financial and Quantitative Analysis, 31(4), 513-539

Eraker, B. (2004), *Do stock prices and volatility jump? Reconciling evidence from spot and option prices*, Journal of Finance, 59(3), 1367-1403

Eraker, B., Johannes, M.S.; Polson, N.G. (2003), *The impact of jumps in returns and volatility*, Journal of Finance, 58(3), 1269-1300

Erb, C.; Harvey, C.; Viskanta, T. (1994), *Forecasting international equity correlation*, Financial Analysts Journal, 50, 32-45

Fama, E. (1965), *The behavior of stock market prices*, Journal of Business, 38(1), 34-105

Fama, E. (1998), *Market efficiency, long term returns, and behavioral finance*, Journal of Financial Economics, 49, 283-306

Farhi, E.; Panageas, S. (2004), *The real effect of stock market mispricing at the aggregate: theory and empirical evidence*, Working paper

Figlewski, S.; Webb, G.P. (1993), *Options, short sales, and market completeness*, Journal of Finance, 48(2), 761-777

Gagliardini, P.; Gourieroux, C. (2004), *Constrained nonparametric dependence with applications in finance and insurance*, Working paper

Gallant, A.R.; Hsieh, D.; Tauchen, G. (1997), *Estimation of stochastic volatility models with diagnostics*, Journal of Econometrics, 81, 159-192

Gârleanu, N.; Pedersen, L.H.; Poteshman, A.M. (2009), *Demand based option pricing*, Review of Financial Studies, forthcoming

Geroski, P.A.; Walters, C.F. (1995), *Innovative activity over the business cycle*, Economic Journal, 105(431), 916-928

Goyal, A.; Saretto, A. (2007), *Option returns and volatility mispricing*, Working paper

Griliches, Z. (1990), *Patent statistics as economic indicators: a survey*, Journal of Economic Literature, 28(4), 1661-1707

Guidolin, M.; Timmermann, A. (2003), *Option prices under Bayesian learning: implied volatility dynamics and predictive densities*, Journal of Economic Dynamics & Control, 27(5), 717-769

Hansen, B.E. (1994), *Autoregressive conditional density estimation*, International Economic Review, 35(3), 705-730

Hansen, L.P.; Sargent, T.J.; Tallarini, T.D.Jr. (1999), *Robust permanent income and pricing*, Review of Economic Studies, 66(4), 873-907

Harris, R.D.F.; Küçüközmen, C.C (2001), *The empirical distribution of UK and US stock* *returns*, Journal of Business, Finance & Accounting, 28(5-6), 715-740

Hong, H.; Scheinkman, J.; Xiong, W. (2005), *Asset float and speculative bubbles*, Journal of Finance, 61(3), 1073-1179

Hugget, M.; Ospina, S. (2001), *Does productivity growth fall after the adoption of new technology?*, Journal of Monetary Economics, 48, 173-195

Jackwerth, J.C.; Rubinstein, M. (1996), *Recovering probability distributions from option prices*, Journal of Finance, 51(5), 1611-1631

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

Jameson, M.; Wilhelm, W. (1992), *Market making in the options markets and the cost of discrete hedge rebalancing*, Journal of Finance, 48(2), 765-779

Jensen, M.C. (2005), *Agency costs of overvalued equity*, Financial Management, 34(1), 5-19

Jiang, G.J.; Tian, Y.S. (2005), *The model-free implied volatility and its information content*, Review of Financial Studies, 18(4), 1305-1342

Johnson, T. (2007), *Optimal learning and new technology bubbles*, Journal of Monetary Economics, 54(8), 2486-2511

Jondeau, E.; Rockinger, M. (2000), *Reading the smile: the message conveyed by methods which infer risk neutral densities*, Journal of International Money and Finance, 19(6), 885-915

Jondeau, E.; Rockinger, M. (2006), *The copula-GARCH model of conditional dependencies: an international stock-market application*, Journal of International Money and Finance, 25(5), 827-853

Jorion, P. (1986), *Bayes-Stein estimation for portfolio analysis*, Journal of Financial and Quantitative Analysis, 21(3), 279-292

Jovanovic, B.; MacDonald, G.M. (1994), *Competitive diffusion*, Journal of Political Economy, 102(1), 24-52

Jovanovic, B.; Nyarko, Y. (1996), *Learning by doing and the choice of technology*, Econometrica, 64(6), 1299-1310

Kaplan, S.N.; Zingales, L. (1997), *Do investment-cash flow sensitivities provide useful measures of financing constraints?*, Quarterly Journal of Economics, 112, 169-215

Lakonishok, J.; Lee, I.; Poteshman, A.M. (2004), *Investor behavior in the option market*, Working paper

Ledoit, O.; Wolf, M. (2003), *Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection*, Journal of Empirical Finance, 10(5), 603-621

Lee, C.; Shleifer, A.; Thaler, R. (1991), *Investor sentiment and the closed-end fund puzzle*, Journal of Finance, 46(1), 75-109

Li, D.X. (2000), *On default correlation: a copula function approach*, Journal of Fixed Income, 9(4), 43-54

Liu, J.; Pan, J.; Wang, T. (2005), *An equilibrium model of rare-event premia and its implication for options smirks*, Review of Financial Studies, 18(1), 131-164

Longin, F.; Solnik, B. (2001), *Extreme correlation of international equity markets*, Journal of Finance, 56(2), 649-676

MacKinley, A.C.; Pastor, L. (2000), *Asset pricing models: implications for expected returns and portfolio selection*, Review of Financial Studies, 13, 883-916

Maenhout, P.J. (2004), *Robust portfolio rules and asset pricing*, Review of Financial Studies, 17(4), 951-983

Mashal, R.; Zeevi, A. (2002), *Beyond correlation: extreme-comovements between financial assets*, Working paper

McQueen, G.; Roley, V.V. (1993), *Stock prices, news, and business conditions*, Review of Financial Studies, 6(3), 683-707

Mood, A.M.; Graybill, F.A. (1963), *Introduction to the theory of statistics*, McGraw-Hill, New York

Nelsen, R.B. (1999), *An introduction to copulas*, Springer-Verlag, London

Nelson, D.B. (1991), *Conditional heteroskedasticity in asset returns: a new approach*, Econometrica, 59(2), 347-370

Olivier, J. (2000), *Growth-enhancing bubbles*, International Economic Review, Vol. 41(1), 133-151

Pastor, L. (2000), *Portfolio selection and asset pricing models*, Journal of Finance, 55(1), 179-223

Pastor, L.; Stambaugh, R.F. (2000), *Comparing asset pricing models: an investment perspective*, Journal of Financial Economics, 56(3), 335-517

Pastor, L.; Stambaugh, R.F. (2009), *Predictive systems: living with imperfect predictors*, Journal of Finance, forthcoming

Patton, A.J. (2004), *On the out-of-sample importance of skewness and asymmetric dependence for asset allocation*, Journal of Financial Econometrics, 2(1), 130-168

Patton, A.J. (2005), *Modelling asymmetric exchange rate dependence*, International Economic Review, 47(2), 527-556

Polk, C.; Sapienza, P. (2009), *The stock market and corporate investment: a test of catering theory*, Review of Financial Studies, 22(1), 187-217

Reinganum, J.F. (1983), *Technology adoption under imperfect information*, Bell Journal of Economics, 14(1), 57-69

Rockinger, M.; Jondeau, E. (2002), *Entropy densities with an application to autoregressive conditional skewness and kurtosis*, Journal of Econometrics, 106(1), 119-142

Roll, R. (1988), *The international crash of October 1987*, Financial Analysts Journal, 44, 19-35

Rosenberg, J.V.; Engle, R. (2002), *Empirical pricing kernels*, Journal of Financial Economics, 64(3), 341-372

Rousseeuw, P.J.; Leroy, A.M. (1987), *Robust regression and outlier detection*, John Wiley & Sons

Rubinstein, M. (1994), *Implied binomial trees*, Journal of Finance, 49(3), 771-818

Savor, P.; Wilson, M. (2008), *Asset returns and scheduled macroeconomic news announcements*, Working paper

Scheinkman, J.A.; Xiong, W. (2003), *Overconfidence and speculative bubbles*, Journal of Political Economy, 111(6), 1183-1219

Shimko, D. (1993), *Bounds of probability*, Risk, 6

Shleifer, A.; Vishny, R.W. (1997), *The limits of arbitrage*, Journal of Finance, 52(1), 35-55

Sloan, R. (1996), *Do stock prices fully reflect information in accruals and cash flows about future earnings?*, Accounting Review, 71, 289-315

Stock, J.H.; Watson, M.W. (1989), *New indexes of coincident and leading economic indicators*, NBER Macroeconomics Annual, 4, 351-394

Thornton, R.A.; Thompson, P. (2001), *Learning from experience and learning from others: an exploration of learning and spillovers in wartime shipbuilding*, American Economic Review, 91(5), 1350-1368

Trajtenberg, M. (1990), *A penny for your quotes: patent citations and the value of innovations*, RAND Journal of Economics, 21(1), 172-187

Uppal, R.; Wang, T. (2003), *Model misspecification and underdiversification*, Journal of Finance, 58(6), 2465-2486

Upton, D.E.; Shannon, D.S. (1979), *The stable Paretian distribution, subordinated stochastic processes and asymptotic lognormality: an empirical investigation*, Journal of Finance, 34(4), 1031-1039

Veronesi, P. (2000), *How does information quality affect stock returns?*, Journal of Finance, 55(2), 807-837

Whitelaw, R.F. (1994), *Time variations and covariations in the expectation and volatility of stock market returns*, Journal of Finance, 49,515-541

Zimmerman, M.B. (1982), *Learning effects and the commercialization of new energy technologies: the case of nuclear power*, Bell Journal of Economics, 13(2), 297-310