Three Essays in Finance

Benjamin Golež

TESI DOCTORAL UPF / 2011

DIRECTOR DE LA TESI

Prof. José M. Marín

TUTOR DE LA TESI

Prof. Xavier Freixas

UNIVERSITAT POMPEU FABRA

Dipòsit legal:

ISBN:

to my parents

Acknowledgement

I owe my gratitude to several individuals who in one way or another contributed in the preparation of this thesis. First and foremost, I would like to thank my supervisor José M. Marín for his guidance and inspiration as well as collaboration on the third essay of this thesis. I also owe much gratitude to Xavier Freixas for his advices and help during my graduate studies. I am indebted to Jens C. Jackwerth for his co-authorship on the second essay of this thesis and his invaluable support in my academic development. I received many insightful comments from Geert Bekaert, Miloš Božović, Paolo Colla, Francesco Corielli, Martijn Cremers, Peter Hoffmann, Gueorgui I. Kolev, Peter Koudijs, Francisco Peñaranda, Christopher Polk, Jesper Rangvid, Bernhard Silli, Roland Umlauft, Joachim Voth, and Robert Zymek. I am grateful to Marta Araque and Laura Agusti for their administrative help. Finally, I would like to thank my friends and my family, whose continuous encouragement and support was truly indispensable on my journey towards the PhD.

Benjamin Golež, Barcelona, March 2011

Abstract

This thesis consists of three essays. In the first essay, I show that information about dividends implied in derivative markets predicts future dividend growth and thereby improves the forecasts of short-run returns on the aggregate market. In the second essay, we analyze the impact of options trading on the price distribution of the underlying asset. Specifically, we show that S&P 500 futures finish in the proximity of the closest strike price on days when options on S&P 500 futures expire. We document that this effect is mainly driven by the rebalancing of delta hedges of the market maker. In the third essay, we develop a theory of price support in security markets that arises from conflict of interests, and we test our hypothesis in the context of the Spanish mutual fund industry. In particular, we analyze how bank-affiliated mutual funds trade in the stock of the parent bank and show that, consistently with the price support hypothesis, affiliated mutual funds tend to increase their holdings of the parent bank's stock following a large drop in its price.

Keywords: *Return predictability, implied dividend growth, options, futures, pinning, hedging, agency problem, price support, mutual funds*

Resumen

Esta tesis consta de tres capítulos. En el primer capítulo, muestro que la información sobre dividendos implícita en los mercados de derivados predice el crecimiento futuro de los dividendos, mejorando así las predicciones de los rendimientos a corto plazo en el mercado agregado. En el segundo capítulo, analizamos el impacto de la compraventa de opciones en la distribución del precio del activo subyacente. En concreto, mostramos que los futuros del S&P 500 terminan en el entorno del precio de ejercicio más próximo en los días en que las opciones sobre los futuros del S&P 500 expiran. Documentamos que este efecto está principalmente motivado por el reajuste de la cobertura delta de los intermediarios. En el tercer capítulo, desarrollamos una teoría de sostenimiento de precios en los mercados de valores motivado por un conflicto de intereses y testamos nuestra hipótesis en el contexto de la industria española de fondos de inversión. En concreto, analizamos cómo los fondos de inversión afiliados a un banco operan las acciones del banco matriz y mostramos que, consecuentemente con la hipótesis del sostenimiento de precios, los fondos de inversión filiales tienden a incrementar sus posiciones en las acciones del banco matriz después de una caída importante de su cotización.

Palabras clave: *Predicción de rendimientos, crecimiento de dividendos implícito, opciones, futuros, pinning, cobertura, problema de agencia, sostenimiento de precios, fondos de inversión*

Table of Contents

Page

Page

Page

Preface

This thesis is a collection of three self-contained essays. In the first essay, I reexamine the role of the dividend–price ratio (DP) for predicting market returns. I argue that the DP's documented poor performance in predicting returns is related to the time-varying expected dividend growth. I introduce a novel proxy for expected dividend growth, which is extracted from index options and futures, and derive a simple present value model to guide the empirical analysis. I find that the DP, when corrected for the implied dividend growth, is much better than the uncorrected DP at predicting short-term S&P 500 returns in the period 1994–2009. This improved performance in predicting returns is especially pronounced at the monthly horizon and is robust to various statistical treatments. Since the improvement yields a sizable gain in the Sharpe ratio, the results have a bearing on portfolio decisions. Because these results are driven by the time-varying expected cash flows, the study has important implications for the risk–return relationship in equilibrium models.

The second essay is a joint work with Jens C. Jackwerth in which we analyze the impact of options trading on the price distribution of the underlying asset. Specifically, we extend the study of "pinning"—that is, the tendency of stocks to finish in the proximity of the closest strike price on the option's expiration day. We find that pinning is also evident in the large and liquid market for S&P 500 futures on days when options on S&P 500 futures expire and the underlying (first-to-maturity) future continues to trade. In exploring the possible explanations for this phenomenon, we provide evidence that pinning is driven by the interplay of market makers' rebalancing of delta hedges due to the time-decay of those hedges as well as in response to reselling (and early exercise) of in-the-money options by individual investors. We find that the effect is asymmetric and stronger above the strike price. Owing to increased options activity, pinning has become more pronounced in recent years. The associated shift in notional futures value is at least \$1.6 bn per expiration day and we suggest a trading strategy to exploit the pinning. To corroborate that the documented pinning is indeed related to options expirations, we show that there is no pinning in second to maturity futures, on which there exist no expiring options. Also, there is no pinning in the S&P 500 index itself or in the exchanged traded fund on the S&P 500 (SPDR) as they are much harder to move through trading than the future.

The third essay reports on work conducted jointly with José M. Marín. In this essay we develop a theory of price support in security markets that arises from conflict of interests, and we test our hypothesis in the context of the Spanish mutual fund industry. In particular, we analyze how bank-affiliated mutual funds (i.e., funds managed by asset management firms that are controlled by banks) trade in the stock of the parent bank. We show that bank-affiliated funds systematically increase, relative to nonaffiliated funds, their holdings of the controlling bank stock when it suffers a large price drop (large negative return). Similar purchases are not made of competing banks' stock when they suffer a similar shock. Neither are these purchases associated with an overall increase in the allocation to banking stocks. Finally, we show that purchases following such large drops do not outperform a portfolio that includes all the other banks, which implies that the documented trading patterns are not driven by private information, either. The patterns are consistent, however, with our price support hypothesis, according to which affiliated funds increase their holdings of the parent bank's stock in times of turmoil in order to limit—in the interest of the bank's shareholders and management—the downside potential of the stock price.

Chapter 1

EXPECTED RETURNS AND DIVIDEND GROWTH RATES IMPLIED IN DERIVATIVE MARKETS

1.1 Introduction

The predictability of market returns is of great interest to market practitioners and has important implications for asset pricing. However, there is still no consensus on whether returns are predictable. Although many studies argue that returns can be predicted by price multiples such as the dividend–price ratio (Fama and French, 1988; Lewellen, 2004; Cochrane, 2008a), others document that predictability is subject to statistical biases and is difficult to exploit for purposes of portfolio allocation (Stambaugh, 1999; Goyal and Welch, 2008).

In this paper I reexamine the role of dividend ratios for predicting market returns. I argue that the poor performance of the dividend-price ratio (DP) in predicting returns is largely due to the time-varying nature of the expected dividend growth. I introduce a novel proxy for expected dividend growth, which is extracted from index options and futures, and derive a simple present value model to guide the empirical analysis. Using the dividend growth implied in derivative markets to correct the DP for variation in expected dividend growth, I find that short-term market returns are strongly predictable. Indeed, the corrected DP predicts monthly market returns both in-sample and out-of-sample, and it is also robust to the statistical biases that have been shown to hinder the predictive ability of the uncorrected DP.

The insight that the time-varying expected dividend growth can reduce the ability of the DP to predict returns has long been part of the predictability literature (Campbell and Shiller, 1988; Fama and French, 1988). According to the textbook treatment, the DP may vary over time not only because of changes in expected returns but also because of changes in expected dividend growth. Therefore, as pointed out by Fama and French (1988), the DP is only a noisy proxy for expected returns in the presence of time-varying expected dividend growth (see also Cochrane, 2008a; Rytchkov, 2008; Binsbergen and Koijen, 2010). Moreover, since the DP increases with expected returns and decreases with expected dividend growth, the problems caused by time-varying expected dividend growth are pronounced when expected returns and expected dividend growth are positively correlated (Menzly et al., 2004; Lettau and Ludvigson, 2005).[1] This positive correlation offsets the changes in expected returns and those in expected dividend growth, which further reduces the DP's ability to predict returns.

Thus, if our task is to predict returns, then the DP is insufficient: We must also account for the time-varying value of expected dividend growth. Yet this value is difficult to estimate because it aggregates investors' expectations about future growth opportunities. Recent studies on return predictability typically assume that the future will be similar to the past and then go on to extract expected dividend growth from historical data. For example, Binsbergen and Koijen (2010) take a latent variable approach within the present value model to filter out both expected returns and expected dividend growth from the history of dividends and prices (see also Rytchkov, 2008). Lacerda and Santa-Clara (2010) use a simple average of historical dividend growth as a proxy for expected dividend growth. These authors all conclude that improved prediction of dividend growth will, in turn, improve the predictability of longer-term (i.e., annual) returns. Nevertheless, their methods exploit only the information that can be derived from past dividends and prices. In contrast, investors base expectations about future cash flows on a much richer—and forward-looking—information set.

[1] Menzly et al. (2004) show that a positive correlation between expected returns and expected dividend growth arises (in a general equilibrium model) as a natural consequence of dividend growth predictability.

This paper takes a different approach to estimating expected dividend growth. Instead of relying on historical data, I extract a proxy for investors' expected dividend growth from derivative markets (index options and index futures). Prices of options and futures depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts. Therefore, derivative markets provide us with a unique laboratory for estimating the dividends that investors expect to realize in the near future. Because index derivatives are highly liquid, new information about future cash flows is rapidly incorporated into the estimated implied dividends. For this reason, implied dividends are particularly well suited for revealing expectations over short horizons, where the constant flow of information causes rapid changes in investors' expectations regarding future dividends and returns.

To provide an analytical framework for the empirical analysis, I first derive a simple present value model. Like Binsbergen and Koijen (2010), I combine the Campbell and Shiller (1988) present value identity with a simple, first-order autoregressive process for the expected return and the expected dividend growth. In this environment, the future return is a function of the DP and the expected dividend growth, where both terms enter linearly. We can therefore consider predicting returns through a multivariate regression of returns on the DP and an estimate for the expected dividend growth, or we can combine them in a single predictor—the so-called corrected DP. The corrected DP can be interpreted as the dividend–price ratio adjusted for variation in expected dividend growth.

Following the implications of the present value model, I proceed with estimating the proposed proxy for the expected dividend growth. To extract the dividend growth implied in index options and index futures, I first estimate an implied dividend yield. By combining the no-arbitrage, cost-of-carry formula for index futures and the put–call parity condition for index options, I derive an expression that enables estimation of the implied dividend yield in a model-free way, and solely in terms of the observed prices of derivatives and their underlying asset. Once estimated, I combine the implied dividend yield with the realized DP to calculate the implied dividend growth and the corrected DP.

I apply the empirical analysis to the S&P 500 index. Given the requirement for data on both options and futures, the analysis is restricted to the period from January 1994 through December 2009.[2] The main results can be summarized as follows. Consistent with previous studies, I find that the standard DP is a rather poor predictor of both future returns and dividend growth. The predictive coefficients on the DP are insignificant in all the forecasting regressions for horizons ranging from one to six months. In contrast, the implied dividend growth reliably predicts dividend growth for all the considered horizons. In line with this observation, the ability to predict market returns improves considerably when implied dividend growth is included as an additional regressor in the standard DP regression for predicting returns. Furthermore, the results confirm that the DP and the implied dividend growth can be replaced by a single predictor: the corrected DP. The predictive coefficient on the corrected DP is statistically significant for all the considered return horizons. The improvement in the predictability is especially strong for short time horizons. In the predictive regressions with monthly returns, the corrected DP exhibits an in-sample adjusted R^2 of 4.61% and an out-of-sample R^2_{OS} of 6.06%, as compared with 0.33% and -0.15% (respectively) for the uncorrected DP. For a mean-variance investor, the documented improvement in predicting returns translates into a gain of 0.32 in terms of the Sharpe ratio. Since the corrected DP is less persistent than the uncorrected DP and since innovations to the corrected DP are only weakly related to returns, the corrected DP has the additional advantage of being robust to small sample bias that that has been shown to hinder the predictive ability of the uncorrected DP. Furthermore, the documented improvement in predictive accuracy is not due to duplication by implied dividend growth of information embedded within other options-implied predictors such as variance risk premia (Bollerslev et al., 2009) and cannot be

[2] Notice that the post 1994 period is not affected by the breaks in the mean of the DP, which have been shown to affect the forecasting relationship of returns and the DP over longer periods of time (Lettau and Nieuwerburgh, 2008; Favero et al., 2010).

replicated by using historical dividend growth in place of implied dividend growth.

Consistent with the empirical results and on contrary to the standard view, a variance decomposition of the DP reveals that there is considerable variation not only in expected returns but also in expected dividend growth. However, like Lettau and Ludvigson (2005), I find that expected returns and expected dividend growth are highly correlated (0.88). This high correlation means that movements in expected returns and expected dividend growth offset each other's effect in the DP, which renders the DP relatively smooth. Correcting the DP for the implied dividend growth restores the variation that is offset by this strong comovement, and thus implies that expected returns vary significantly more than is suggested by variation in the uncorrected dividend–price ratio.

The paper draws upon a large number of studies in the predictability literature and is also related to other papers using implied dividends. Dividends implied in derivative markets have been used as an input in the calculation of risk-neutral densities (Ait-Sahalia and Lo, 1998), and to study empirical properties of dividend strips (Binsbergen, Brandt and Koijen, 2010). However, this paper is the first to employ implied dividends for the purpose of predicting market returns. I also use a new techinique which enables me to extract dividends from derivative prices without resorting to the use of proxies for the implied interest rate. This is important as interest rates implied in derivative markets may differ from observable interest rates (Naranjo, 2009).

The rest of the paper is organized as follows. Section 1.2 derives the present value model. Section 1.3 details the technique proposed to extract the dividend growth that is implied in the market for derivatives. Section 1.4 presents the data, and Section 1.5 reports on the results of predictive regressions involving dividend growth and market returns. Section 1.6 considers additional statistical tests and compares the documented predictability with alternative predictors. Section 1.7 presents a variance decomposition of the dividend–

price ratio, and Section 1.8 is devoted to robustness checks. Section 1.9 concludes the paper.

1.2 Present value model

To provide an analytical framework for the empirical analysis, this section derives a simple log-linear present value model. The model combines the Campbell and Shiller (1988) present value identity with AR(1) processes for expected returns and expected dividend growth rates. A similar approach is used in Binsbergen and Koijen (2010) and Rytchkov (2008).[3,4] The main innovation of this study lies in the empirical estimation of this setup. I use the present value model mainly to motivate the return predictive regressions.

Define log return r_{t+1}, log dividend growth Δd_{t+1}, and log dividend-price ratio dp_t as:

$$r_{t+1} = \log\left[\frac{P_{t+1} + D_{t+1}}{P_t}\right], \quad \Delta d_{t+1} = \log\left[\frac{D_{t+1}}{D_t}\right], \quad dp_t = \log\left[\frac{D_t}{P_t}\right] \tag{1}$$

Rewrite returns as in Campbell and Shiller (1988):

$$r_{t+1} \simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} \tag{2}$$

where $\rho = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})}$ and $\kappa = \log\left[1 + \exp(-\overline{dp})\right] + \rho\overline{dp}$ are constants related to the long-run average of the dividend-price ratio, $\overline{dp}$. Iterate (2) forward to obtain the Campbell and Shiller (1988) present value identity:

$$dp_t \simeq -\frac{\kappa}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j (r_{t+1+j}) - E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j}) \tag{3}$$

[3] The AR(1) structure is motivated by growing evidence that both expected returns and expected dividend growth rates are time-varying and persistent (Menzly et al., 2004; Lettau and Ludvigson, 2005; Bansal and Yaron, 2004).

[4] Present value models with different processes for expected returns and expected dividend growth are extensively analyzed in Cochrane (2008b).

Let $\mu_t = E_t(r_{t+1})$ be the conditional expected return and let $g_t = E_t(\Delta d_{t+1})$ be the conditional expected dividend growth. Suppose that μ_t and g_t follow $AR(1)$ processes:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t) + \varepsilon^{\mu}_{t+1} \quad (4)$$

$$g_{t+1} = \gamma_0 + \gamma_1(g_t) + \varepsilon^{g}_{t+1} \quad (5)$$

$$\Delta d_{t+1} = g_t + \varepsilon^{d}_{t+1} \quad (6)$$

where ε^{μ}_{t+1}, ε^{g}_{t+1} and ε^{d}_{t+1} are zero mean errors. Combine the present value identity in (3) with the $AR(1)$ assumptions to find the dividend-price ratio:

$$dp_t \simeq \varphi + \left(\frac{1}{1-\rho\delta_1}\right)\mu_t - \left(\frac{1}{1-\rho\gamma_1}\right)g_t \quad (7)$$

where φ is a constant related to $\kappa, \rho, \delta_0, \delta_1, \gamma_0, \gamma_1$ (details are provided in Appendix).

Equation (7) states that the log dividend-price ratio is related to expected returns and is therefore a good candidate for predicting future returns. However, according to (7), dp_t also contains information about expected dividend growth. Hence, if expected dividend growth varies over time, the dp_t is only a noisy proxy for expected returns and an imperfect predictor for future returns (Fama and French, 1988; Binsbergen and Koijen, 2010; Rytchkov, 2008; Lacerda and Santa-Clara, 2010). Since the dp_t increases with expected returns and decreases with expected dividend growth, the problem is pronounced when expected returns and expected dividend growth are positively correlated (Menzly et al., 2004; Lettau and Ludvigson, 2005). This positive correlation offsets the changes in expected returns and those in expected dividend growth, which further reduces the ability of the dp_t to predict returns.

Thus, if our task is to predict returns, then the dp_t is insufficient: We must also account for the time-varying value of expected dividend growth. To see this formally, combine (2), (6) and (7) to obtain a return forecasting equation:

$$r_{t+1} \simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} \quad (8)$$

$$\simeq \psi + (1-\rho\delta_1)dp_t + \left(\frac{1-\rho\delta_1}{1-\rho\gamma_1}\right) g_t + v^r_{t+1} \quad (9)$$

where $v^r_{t+1} = \varepsilon^d_{t+1} - \rho\left(\frac{\varepsilon^\mu_{t+1}1}{1-\rho\delta_1} - \frac{\varepsilon^g_{t+1}}{1-\rho\gamma_1}\right)$ and ψ is a constant related to $\kappa, \rho, \delta_0, \delta_1, \gamma_0$, and γ_1.

In line with the above argument, equation (8) reveals that, if our task is to predict returns, we need both dp_t and an estimate for expected dividend growth.

Since dp_t and the expected dividend growth are linearly related to future returns, we can also replace them by a single predictor:

$$r_{t+1} \simeq \psi + (1-\rho\delta_1)dp_t + \left(\frac{1-\rho\delta_1}{1-\rho\gamma_1}\right) g_t + v^r_{t+1} \quad (10)$$

$$\simeq \psi + (1-\rho\delta_1)\left[dp_t + g_t\left(\frac{1}{1-\rho\gamma_1}\right)\right] + v^r_{t+1} \quad (11)$$

$$\simeq \psi + (1-\rho\delta_1)dp^{Corr}_t + v^r_{t+1} \quad (12)$$

where $dp^{Corr}_t = dp_t + g_t\left(\frac{1}{1-\rho\gamma_1}\right)$ is the corrected dividend-price ratio and can be interpreted as the dividend-price ratio that is adjusted for variation in the expected dividend growth. The corrected dividend-price ratio depends on the dp_t, the expected dividend growth, the linearization constant and the persistence of the expected dividend growth.[5,6]

[5] The fact that correction depends on the persistence of the expected dividend growth is an interesting insight since persistence of the expected dividend growth is one of the driving forces of the return predictability in the long-run risk models pioneered by Bansal and Yaron (2004).

[6] Lacerda and Santa-Clara (2010) derive a similar correction for the adjusted dividend-price ratio:
$dp^{Adj.}_t = dp_t + \overline{g}_t\left(\frac{1}{1-\rho_t}\right)$

1.3 Estimating implied dividend growth

The present value model outlined in the previous section implies that the dividend-price ratio is not enough to capture variation in expected returns. Additionally, we need an estimate for the expected dividend growth.

In this study, I propose to extract a proxy for expected dividend growth from derivative markets (index options and index futures). Prices of options and futures depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts. Therefore, we can invert the pricing relations to extract a proxy for expected dividend growth from the observable prices of derivatives.

I employ a two step approach to estimating the implied dividend growth. In the first step, I extract an implied dividend yield embedded in derivative markets. In the second step, I combine the estimated implied dividend yield with the realized dividend-price ratio to calculate the implied dividend growth.

Below, I describe the proposed method for the estimation of the implied dividend yield. Transition from the implied dividend yield to the implied dividend growth is presented along with the estimation of the realized dividend-price ratio in the next section.

Implied dividend yield. To express the implied dividend yield in terms of the observable prices of derivatives, I combine two well-known no-arbitrage conditions, the cost-of-carry formula for index futures and the put-call parity condition for index options.

Under a standard assumption that the index pays a continuously compounded dividend yield (λ), the cost-of-carry formula for the future price is:

In their version, the adjusted dividend-price ratio ($dp_t^{Adj.}$) does not depend on the persistence of the expected dividend growth because they assume that expected dividend growth is equal to the average historical dividend growth ($\overline{g}_t$).

$$F_t(\tau) = S_t \exp\left[(r_t(\tau) - \lambda_t(\tau))\tau\right] \tag{13}$$

where F_t is the future's price, S_t is the price of the underlying, $\lambda_t(\tau)$ is the annualized continuously compounded dividend yield between t and $t+\tau$ and $r_t(\tau)$ is the annualized continuously compounded interest rate from t to $t+\tau$.

Similarly, by no-arbitrage, the difference between a European call and a European put written on the index can be expressed as:

$$C_t(K,\tau) - P_t(K,\tau) = S_t \exp\left[-\lambda_t(\tau)\tau\right] - K \exp\left[-r_t(\tau)\tau\right] \tag{14}$$

where $C_t(K,\tau)$ and $P_t(K,\tau)$ are the prices of a European call and a European put option with the same maturity τ and the same strike price K.

Both no-arbitrage conditions relate prices of derivatives to the future dividend yield and the risk-free rate. Hence, we can combine them to first solve for the interest rate implied in the derivative markets:

$$r_t(\tau) = \frac{1}{\tau} \log\left[\frac{F_t(\tau) - K}{C_t(K,\tau) - P_t(K,\tau)}\right] \tag{15}$$

Once we have an expression for the implied interest rate, we can plug it back in (14) to obtain an expression that enables us to estimate the implied dividend yield:

$$\lambda_t(\tau) = -\frac{1}{\tau} \log\left[\left(\frac{C_t(K,\tau) - P_t(K,\tau)}{S_t}\right) + \frac{K}{S_t}\left(\frac{C_t(K,\tau) - P_t(K,\tau)}{F_t(\tau) - K}\right)\right] \tag{16}$$

Equation (16) has several appealing features. First, it enables us to estimate the implied dividend yield using only information that is available at time t. All we need is a European call option and a European put option with the same strike and the same maturity, the future price with the same expiration date as the options, and the price of the underlying. Second, as the expression is based on a combination of two no-arbitrage conditions, it allows us to substitute out the interest rate and estimate the implied dividend yield without resorting to the use of proxies for the implied interest rate. This is important because the implied interest rate may deviate from the observable proxies for the interest rate (Naranjo, 2009).

Note, however, that the estimated implied dividend yield contains dividend yield risk premia because I do not consider an adjustment for the stochastic dividend yield. Although such adjustment is possible by assuming an exogenous process for the implied dividend yield,[7] the dividend yield risk premia is believed to play a secondary role given the low volatility of the realized dividend yield and the fact that dividends are partially known in advance. Empirical results support this notion by showing that the information about future returns steaming from implied dividends is orthogonal to the information about expected returns contained in other proxies for risk and risk aversion (e.g. variance risk premia).

1.4 Data

I use the S&P 500 index as a proxy for the aggregate market. The S&P 500 price index and total return index (dividends reinvested) are downloaded from Datastream. The S&P 500 futures data comes from Chicago Merchandile Exchange and the S&P 500 options data is obtained from Market Data Express.

Futures on S&P 500 have been traded since April 1982 and European options on S&P 500 have existed since April 1986. However, Market Data Express options data only goes back to January 1990. Also, until 1994, the settlement procedure for S&P 500 options and futures differed. While futures are settled

[7] See Lioui (2006) for the derivation of the put-call parity under the stochastic dividend yield.

in the opening value of the index since June 1987, the most liquid S&P 500 options expired in the closing value of the index until December 1993.[8] Since liquid options and futures with matching expiration times are needed to estimate the implied dividend growth, I further restrict the analysis to the period from January 1994 through December 2009. The analysis is based on end-of-month observations.

In some parts of the paper I also make use of other variables. In particular, I download constant maturity 3-month and 6-month Treasury yield from the Federal Reserve Bank of St. Louis and I obtain the S&P 500 earnings-price ratio and the 6-month LIBOR rate from Datastream. Additionally, I obtain the implied variance index (VIX) and the variance risk premia from Hao Zhou's homepage. Finally, I download the consumption-to-wealth ratio from Sydney C. Ludvigson's website.

1.4.1 Empirical estimation

Implied dividend yield. I estimate the implied dividend yield at the end of each month according to (16). I use daily settlement prices for futures, mid-point between the last bid and the last ask price for options and closing values for the S&P 500 price index.[9]

It is well-known that no-arbitrage conditions hold well for the S&P 500 index (Kamara and Miller, 1995). Still, due to market frictions (transaction costs

[8] When S&P 500 futures and S&P 500 options were introduced, they initially expired in the closing value of the index (P.M. settlement). In 1987, the Chicago Merchandile Exchange (CME) changed the expiration procedure of S&P 500 futures from the P.M. settlement to the A.M. settlement (A.M. settlement value is based on the opening prices of the index constituents on the expiration date). As a response, the Chicago Board of Options Exchange (CBOE) introduced a new version of its S&P 500 options that also settle A.M. However, the P.M. settled options remained the most liquid and the A.M. settled options were initially hardly traded. In 1992, CBOE decided that all the S&P 500 options should expire A.M. Since long dated P.M. settled options were already traded on the market, it took until December 1993 before all the traded S&P 500 options became A.M. settled.

[9] Market Data Express end-of-day data covers all the options written on the S&P 500 index, including mini options, quarterlies, weeklies and long-dated options. With the kind help of Market Data Express support team, I first eliminated all but standard S&P 500 options. Additionally, I imposed the standard filters to eliminate missing observations and options that violate the basic no-arbitrage bounds.

and demand imbalances), particular pairs of options and futures may violate no-arbitrage conditions. To take this into account, I calculate the implied dividend yield by aggregating information from a wide set of options and futures.

For each end of the month, I use 10 days of backward-looking data and I construct option pairs (put-call pairs with the same strike and the same maturity) from all the reliable options (options with positive volume or open interest greater than 200 contracts).[10] Then I combine option pairs with the futures of matching maturity and the current value of the underlying index. To eliminate some extreme observations, I discard observations where $\frac{C_t(K,\tau)-P_t(K,\tau)}{F_t(\tau)-K}$ is smaller than 0.5 or greater than 1.5 (and where $F_t(\tau) = K$).[11] Using this data, I obtain several estimates for the implied dividend yields at the end of each month, which I aggregate into a single market's implied dividend yield by taking the median across all the implied dividend yields with the same maturity.

Since within year dividends exhibit seasonality, the common approach in the predictability literature is to calculate the dividend-price ratio by aggregating dividends over one year. In line with this literature, the implied dividend yield should ideally be estimated using options and futures with one year to expiration. However, long maturity derivatives are illiquid. As illustrated in Figure 1, open interest concentrates strongly on near to maturity options and futures. The tilt towards short maturities is especially pronounced for futures, for which there is almost no open interest for maturities above 9 months. For this reason, we cannot reliably estimate the implied dividend yield with the maturity of one year and we have to resort to the use of options and futures with shorter expiration dates. This may, nevertheless, introduce some seasonality into the estimated implied dividend yield.

[10] Note: the formula for the implied dividend yield holds for all the moneyness levels. Unreported results show that there is no strike price effect, i.e. the implied dividend yield does not depend on the moneyness level.

[11] This filter eliminates a bit less than 2% of observations.

To examine the effect of the seasonality in dividend payments on the implied dividend yield, I first estimate the whole term structure of the implied dividend yields. Since there are only four dates per year when options and futures expire simultaneously (third Friday in March, June, September and December),[12] I proceed as follows. In January, April, July and October, I extract the implied dividend yield for the maturities of 2, 5, or 8 months. In February, May, August and November, I extract the implied dividend yield for the maturities of 1, 4 or 7 months. Finally, in March, June, September and December, I estimate the implied dividend yield for the maturities of 3, 6 or 9 months. Then I linearly interpolate the estimated yields to obtain the term structure of the implied dividend yields with constant maturities (between 3 and 7 months).

Table 1 presents the summary statistics for the implied dividend yields with different maturities. All the yields have approximately the same mean, but differ with respect to their volatility. As expected, due to the seasonality in dividend payments, implied dividend yields with short maturities (3 and 4 months) are the most volatile. With the increase of the maturity, the volatility of the implied dividend yields first decays and then stabilizes, so that implied dividend yields with 6 and 7 months to maturity exhibit approximately the same volatility (see also Figure 2). This suggests that the problem of seasonality in dividend payments is largely diminished for the implied dividend yield with a maturity of at least 6 months. Given these results, I choose to conduct the main analysis using the implied dividend yield with maturity of 6 months.

By construction, the estimated implied dividend yield is continuously compounded. To make it comparable with the realized dividend-price ratio, I transform it into a raw (effective) implied dividend yield, $IDY_t = \exp(\widehat{\lambda}_t) - 1$. The log implied dividend yield is simply $idy_t = log(IDY_t)$.

[12] Options expire on a monthly cycle (third Friday in a month) and futures expire on a quarterly cycle (third Friday in March, June, September and December).

Market returns, dividend growth and dividend-price ratio. I follow the standard definitions for the realized variables. Monthly returns are defined as:

$$r_t^M = \log\left[\frac{P_t + D_t}{P_{t-1}}\right] \tag{17}$$

where P_t and D_t denote the price and dividends in month t. The dividend-price ratio is calculated by aggregating dividends over one year:

$$dp_t = \log\left[DP_t\right] = \log\left[\frac{D_t^{12}}{P_t}\right] \tag{18}$$

where D_t^{12} is the sum of dividends over the last 12 months. Monthly dividend growth is defined as in Ang and Bekaert (2007):

$$\Delta d_t^M = \log\left[\frac{D_t^{12}}{D_{t-1}^{12}}\right] \tag{19}$$

All the ratios are calculated from the S&P 500 price index and the total return index downloaded from Datastream. Since Datastream calculates the total return index by reinvesting dividends daily, I first extract the daily amount of dividends. Then I calculate D_t and D_t^{12} by summing dividends over the past month and year, respectively.

Implied dividend growth and the corrected dividend-price ratio. Based on the implied dividend yield and the dividend-price ratio, I calculate the implied dividend growth (idg) and the corrected dividend-price ratio (dp_t^{Corr}) as:

$$idg_t = \log\left[\frac{IDY_t}{DP_t}\right] = idy_t - dp_t \tag{20}$$

$$dp_t^{Corr} = dp_t + \left(\frac{1}{1-\widehat{\rho\gamma_1}}\right) idg_t \tag{21}$$

where $\widehat{\rho}$ is the estimated linearization constant and $\widehat{\gamma_1}$ is the $AR(1)$ coefficient of the implied dividend growth.

1.4.2 Data description

Table 2 reports the summary statistics for the variables sampled monthly. All the variables are annualized and expressed in logs. Returns and dividend growth rates are on average 7.33% and 3.61%, respectively.

The proxy for the expected dividend growth (implied dividend growth) is on average somewhat higher than the realized dividend growth rate (6.02%) and it nicely reflects market conditions. As shown in Figure 3, the implied dividend growth is positive during the market booms (1994-97 and 2002-2007), when investors were optimistic about future growth opportunities, and it is negative in times of stock market busts, such as in 1998 (Asian-Russian-LTCM crisis), in 2001 (dot.com bubble burst), and in 2008/2009 (the recent financial crisis), when investors were rather pessimistic about growth opportunities. The implied dividend growth is also relatively persistent. It exhibits a first order autocorrelation coefficient of 0.53 and it thereby justifies modeling expected dividend growth rate as a persistent process.

The corrected dividend-price ratio is calculated as:

$$dp_t^{Corr} = dp_t + \left(\frac{1}{1 - (0.98 * 0.53)} \right) idg_t = dp_t + 2.08 * idg_t. \qquad (22)$$

where $\widehat{\rho} = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})} = \frac{\exp(-4.03)}{1+\exp(-4.03)} = 0.98.$[13]

Figure 4 plots dp_t^{Corr} along with the dp_t. Both dividend ratios exhibit strong comovement (pairwise correlation coefficient of 0.72), but they differ in three important aspects.

First, in line with the patterns revealed by the expected dividend growth, the dp_t^{Corr} is on average higher than the dp_t in the boom periods and it is lower than the dp_t in the bust periods. This means that the dp_t tends to

[13] Note: the construction of the corrected dividend-price ratio introduces a look-ahead bias because $\widehat{\rho}$ and $\widehat{\gamma_1}$ are estimated using the data of the whole sample and are therefore not available at time t. However, the out-of-sample predictability results in the Section 6 show that the look-ahead bias plays only a minor role when predicting returns with the corrected dividend-price ratio.

predict returns that are too low to be justified with the market's optimism about growth opportunities during the boom periods. Simultaneously, the dp_t tends to forecast returns that are too high during the crisis periods. This is especially apparent at the end of the sample when the market experienced one of the largest drops in the history of the U.S. market, but the uncorrected dividend-price ratio rose and therefore implied unrealistically high returns.

Second, the corrected dividend-price ratio is notably more volatile than the uncorrected dividend-price ratio. The standard deviation is 0.27 for the dp_t and 0.54 for the dp_t^{Corr}. In the context of the present value model, this increase in volatility implies that expected returns and expected dividend growth are highly correlated. To see this formally, expand the variance of the corrected dividend-price ratio as:

$$var(dp_t^{Corr}) = var(dp_t) + 2\left(\frac{1}{1-\rho\gamma_1}\right)\left(\frac{1}{1-\rho\delta_1}\right)cov(\mu_t, g_t) - \left(\frac{1}{1-\rho\gamma_1}\right)^2 var(g_t) \quad (23)$$

Equation (23) says that the variance of the dp_t^{Corr} can be higher than the variance of the dp_t only if expected returns and expected dividend growth rates covary and the covariation is big enough:

$$\left(2\left(\frac{1}{1-\rho\delta_1}\right)cov(\mu_t, g_t) > \left(\frac{1}{1-\rho\gamma_1}\right)var(g_t)\right) \quad (24)$$

Furthermore, since dp_t increases with expected returns and decreases with expected dividend growth, this positive covariation also affects the uncorrected dividend-price ratio. It offsets shocks to expected returns and expected dividend growth and reduces the volatility of the dp_t (Lettau and Ludvigson, 2005, Rytchkov, 2008; Van Binsbergen and Koijen, 2010). Thus, correcting the dp_t for the implied dividend growth restores the variation, which is otherwise offset by the comovement of the expected return and the expected dividend growth (see also Lacerda and Santa-Clara, 2010).

Last, consistent with the increase in volatility of the dp_t^{Corr}, the dp_t^{Corr} is also less persistent than the dp_t. While dp_t exhibits first order autocorrelation coefficient of 0.98, the $AR(1)$ for dp_t^{Corr} is notably lower and amounts to 0.74. This decrease in persistence is important because highly autocorrelated predictors are typically subject to small sample bias (Stambaugh, 1999) and produce inaccurate inference results in the case of overlapping observations (Boudoukh et al., 2008). Given its lower persistence, the dp_t^{Corr} is therefore largely free of the common concern related to the use of highly persistent variables for predicting returns.

By applying equation (15) and following the same estimation procedure as for the implied dividend yield, I additionally estimate the implied interest rate (IIR_t). Although IIR_t is not of special interest for this study, it is important to note that the IIR_t behaves as we would expect. As shown in Figure 5, IIR_t strongly covaries with the T-bill rate and the LIBOR rate and it is on average closer to the LIBOR rate (see also Naranjo, 2009). Still, IIR_t is more volatile than the T-bill rate and the LIBOR rate at the beginning of the analyzed period and it deviates from both proxies for the interest rate during the recent financial crisis, when it is notably lower than the LIBOR rate. This shows that the implied interest rate may deviate from the observable proxies for the interest rate and it therefore points at the importance of isolating the effect of the interest rate when estimating the implied dividend yield.

1.5 Empirical results

This section presents dividend growth and market return predictability results. Since derivative markets subsume market expectations about the near future, the implied dividend ratios should be especially suitable for tracking short term variations in future dividends and returns as opposed to long term tendencies in asset markets. To investigate this, I consider predicting dividend growth rates and market returns at the horizons ranging from one to six months.

I use standard predictive regressions, in which returns or dividend growth rates are regressed on the lagged predictors. I report OLS t-statistics for the case of non-overlapping monthly observations and Hodrick (1992) t-statistics for the case of longer horizon regressions with overlapping observations.[14] Additionally, I report the adjusted R^2. Note however that the R^2 in the context of overlapping observations needs to be interpreted with caution because it tends to increase with the length of the overlap even in the absence of true predictability (Valkanov, 2003; Boudoukh et al., 2008).

1.5.1 Predicting dividend growth

Figure 3 shows that the implied dividend growth tracks general market conditions and it therefore seems to be a good proxy for the expected dividend growth. In this subsection, I complement this argument by showing that the implied dividend growth also uncovers part of the variation in the future dividend growth.

For a comparison with the implied dividend growth, I consider whether the dividend-price ratio predicts future dividend growth. I use dp_t as a competing predictor for two reasons. Firstly, dp_t is itself a function of the expected dividend growth and could therefore predict future dividend growth as opposed to future returns. Secondly, implied dividend growth is defined as the difference between the implied dividend yield and the dividend-price ratio. Therefore, it is necessary to show that the implied dividend growth does not predict future dividend growth simply because it is duplicating information contained in the dp_t.

The main regression takes the following form:

$$\Delta d_{t+h} = a_0 + a_1(X_t) + \varepsilon_{t+1} \tag{25}$$

[14] Ang and Bekaert (2007) show that the performance of Hodrick (1992) standard errors, which are based on summing the predictors in the past, is superior to other standards errors that are frequently employed in the literature, such as the Newey-West (1987) standard errors, or the Hansen and Hodrick (1980) standard errors.

where $\Delta d_{t+h} = (12/h)\sum_{i=1}^{h} \Delta d_{t+i}^{M}$ is the annualized dividend growth with $h = 1, 2, 3$ or 6 months and X_t is either idg_t, or dp_t, or both. For $h = 1$, t-statistics are based on the simple OLS. For $h = 2, 3$ or 6, t-statistics are computed according to Hodrick (1992).

Table 3 presents results. I start by analyzing regression results with the dividend-price ratio. The estimated parameter on the dp_t is negative, just as the theory suggests, but the associated t-statistics are insignificant at the conventional 5% level and range between 1.41 and 1.96. Also, the *adj.* R^2 is low and ranges from 1.47% for monthly dividend growth to 3.28% for half-annual dividend growth. In comparison, the implied dividend growth is positively related to future dividend growth and explains 4.79% of the variation in the monthly dividend growth and 18.42% of the variation in the half-annual dividend growth. Furthermore, all the estimated coefficients on the implied dividend growth are statistically significant and range between 3.25 and 4.76.

As reported in the last panel of Table 3, adding dividend-price ratio as an additional predictor to the implied dividend growth boosts statistical significance of the implied dividend growth and leads to further increase in the *adj.* R^2. The *adj.* R^2 in a bivariate predictive regression amounts to 9.31% for monthly dividend growth and to 30.81% for half annual dividend growth. Since this is more than the sum of the *adj.* $R^{2\prime}s$ in the univariate regressions, it clearly indicates that the implied dividend growth is not duplicating information about future returns that is already captured in the dividend-price ratio.

1.5.2 Predicting market returns

I employ three specifications for the return predictive regressions. The first is the standard predictive regression, in which returns are regressed on the lagged dividend-price ratio:

$$r_{t+h} = b_0 + b_1(dp_t) + \varepsilon_{t+1} \tag{26}$$

The second regression augments the first by using the proxy for the expected dividend growth (implied dividend growth):

$$r_{t+h} = c_0 + c_1(dp_t) + c_2(idg_t) + \varepsilon_{t+1} \tag{27}$$

The last return regression replaces the dividend-price ratio and the implied dividend growth by the corrected dividend-price ratio:

$$r_{t+h} = d_0 + d_1(dp_t^{Corr}) + \varepsilon_{t+1} \tag{28}$$

In all the regressions, $r_{t+h} = (12/h)\sum_{i=1}^{h} r_{t+i}^{M}$ is the annualized market return with $h = 1, 2, 3$ or 6 months. For $h = 1$, t-statistics are based on the simple OLS. For $h = 2, 3$ or 6, t-statistics are computed according to Hodrick (1992).

Table 4 presents the regression results. I start by analyzing univariate regression results of returns on the lagged dp_t. The estimated coefficient on the dp_t is positive, as suggested by the theory, but the t-statistics are insignificant at the 5% level of statistical significance and range between 1.28 and 1.67. Also, the associated *adj.* R^2 is relatively low and ranges from 0.33% for monthly returns to 7.02% for half-annual returns.

When implied dividend growth is added as an additional regressor to the dividend-price ratio, the return predictability improves for all the considered horizons. The *adj.* R^2 increases from 0.33% to 5.20% in the regression with monthly returns and from 7.02% to 8.71% in the regression with half-annual returns. This result is directly in line with the observation that the implied dividend growth predicts future dividend growth and thereby implies that variation in the expected dividend growth plays an important role for uncovering variation in the future returns.

As suggested by the present value model and confirmed by the last regression, the dp_t and the implied dividend growth can also be replaced by a single predictor, the corrected dividend-price ratio. The corrected dividend-price ratio predicts returns approximately as well as the dividend-price ratio and the implied dividend growth together. The *adj.* R^2 amounts to 4.61% at the monthly horizon and to 8.56% at the half-annual horizon. Also, the estimated parameter on the dp_t^{Corr} is always statistically significant with the t-statistics ranging from 3.19 at the monthly horizon to 2.33 at the half-annual horizon.

1.6 Additional tests

The results imply that the corrected dividend-price ratio predicts returns significantly better than the realized dividend-price ratio, and that the improvement in the predictability is especially pronounced over the monthly horizon. However, all the results so far are based on the in-sample predictive regressions, which have been criticized on the grounds that they may be subject to the small sample bias (Stambaugh, 1999), and they may not necessarily imply that the documented predictability can be exploited in real time (Goyal and Welch, 2008).

To address these issues, this section considers small sample bias correction, out-of-sample predictability and a simple out-of-sample trading strategy. Additionally, I compare the return predictive ability of the corrected dividend-price ratio to the alternative corrections for the dividend-price ratio and to other popular predictors. To avoid the statistical problems inherent in the use of overlapping observations (Boudoukh et al., 2008), the analysis is restricted to predicting non-overlapping monthly returns.

1.6.1 Is there a small sample bias?

Dividend ratios are very persistent and an extensive literature argues that the standard OLS predictive regressions applied to highly persistent variables may lead to severe biases in small samples (Stambaugh, 1999; Amihud and Hurvich, 2004).

To analyze the source of the bias, consider a model where returns are predicted by a variable (X_t) that follows first-order autoregressive process:

$$r_{t+1} = \alpha + \beta X_t + u_{t+1} \tag{29}$$

$$X_{t+1} = \delta + \rho X_t + v_{t+1} \tag{30}$$

where $|\rho| < 1$ and the errors (u_{t+1}, v_{t+1}) are distributed as:

$$\begin{pmatrix} u_{t+1} \\ v_{t+1} \end{pmatrix} \sim_{iid} N(0, \Sigma), \Sigma = \begin{pmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{pmatrix} \tag{31}$$

If errors are correlated ($\sigma_{uv} \neq 0$), OLS produces a biased estimate of β in small samples (Stambaugh, 1999). The larger the ρ, i.e. the persistence of shocks to the predictor variable, the larger the bias. For dividend-price ratios, σ_{uv} is negative and ρ is close to one. This results in upward biased estimates of β and the corresponding t-statistics.

To correct for the small sample bias, I follow the correction methodology proposed by Amihud and Hurvich (2004) and employed in several recent studies (Boudoukh et al. 2007; Kolev, 2008; Lioui and Rangvid, 2009). First I estimate (29) to obtain an OLS estimate $\widehat{\rho}$. Then, I calculate the bias corrected estimator for $\widehat{\rho}$:

$$\widehat{\rho}^c = \widehat{\rho} + (1 + 3\widehat{\rho})/n + 3(1 + 3\widehat{\rho})/n^2 \tag{32}$$

where n is the length of the time series. The estimator $\widehat{\rho}^c$ is then used to calculate the bias corrected errors:

$$v^c_{t+1} = X_{t+1} - [(1 - \widehat{\rho}^c)\Sigma^n_{t=1}(X_{t+1}/n) + \widehat{\rho}^c X_t] \tag{33}$$

Finally, I run an OLS regression of returns on the predictor variable X_t and the v^c_{t+1} :

$$r_{t+1} = \alpha + \beta^c X_t + \phi^c v^c_{t+1} + \varepsilon_{t+1} \tag{34}$$

The estimate of β^c gives us the bias corrected estimator of β. The corresponding bias corrected t-statistic is calculated as:

$$t^c = \widehat{\beta^c} / \sqrt{\left(\widehat{\phi^c}\right)^2 \left(\widehat{SE}(\widehat{\rho})\right)^2 (1 + 3/n + 9/n^2)^2 + \left(\widehat{SE}(\widehat{\beta^c})\right)^2} \tag{35}$$

I apply the bias correction to the realized dividend-price ratio and to the corrected dividend-price ratio. Table 5 compares and contrasts the slope estimates and the t-statistics based on the standard OLS with those obtained after correcting for the small sample bias.

The realized dividend-price ratio is an insignificant predictor for monthly returns even before correcting for the small sample bias. After correction, the estimated predictive coefficient even changes its sign and becomes negatively related to future returns. Unlike the realized dividend-price ratio, the corrected dividend-price ratio is largely unaffected by the small sample bias correction. The adjusted slope coefficient is almost identical to the OLS slope coefficient (0.22 in comparison to 0.23) and the adjusted t-statistic is only marginally smaller than the OLS t-statistic (3.10 in comparison to 3.19).

A rather small effect of the small sample bias correction on the inference of the dp^{Corr}_t is due to a combination of two effects. First, the corrected dividend-price ratio is less persistent than the realized dividend-price ratio (0.74, in comparison to 0.98). Second, the innovations to the predictor variable and to the returns are only weakly correlated (-0.24 for the dp^{Corr}_t in comparison to -0.97 for the realized dividend-price ratio). The combination of both effects enables the dp_t^{Corr} to remain statistically significant predictor for monthly returns and hence, implies that the corrected dividend-price ratio is by and large robust to small sample bias.

1.6.2 Out-of-sample predictability

Goyal and Welch (2008) demonstrate that variables with in-sample predictive power may not necessarily predict returns out-of-sample. I follow their approach to test whether the corrected dividend-price ratio predicts returns out-of-sample better than the realized dividend-price ratio.

I calculate the out-of-sample R^2 as in Campbell and Thompson (2008) and Goyal and Welch (2008):

$$R_{OS}^2 = 1 - \frac{\sum_{t=1}^{T} (r_{t+1} - \widehat{\mu}_t)^2}{\sum_{t=1}^{T} (r_{t+1} - \overline{r_t})^2} \tag{36}$$

where $\widehat{\mu}_t$ is the fitted value from a predictive regression estimated through period t and $\overline{r_t}$ is the historical average return estimated through period t. A positive out-of-sample R^2 indicates that the predictive regression has a lower mean-squared prediction error than the historical average return.

To make out-of-sample forecasts, I split the sample in two subperiods. I use the period from January 1994 through December 1999 for the estimation of the initial parameters and the period from January 2000 through December 2009 for the calculation of the R_{OS}^2. All out-of-sample forecasts are based on a recursive scheme using all the available information up to time t. I calculate R_{OS}^2 for the realized and the corrected dividend-price ratio.

Recall that the corrected dividend-price ratio is defined as:

$$dp_t^{Corr} = dp_t + \left(\frac{1}{1 - \widehat{\rho\gamma_1}} \right) idg_t \tag{37}$$

where ρ (linearization constant) and γ_1 (AR(1) coefficient of the implied dividend growth) are estimated using the whole sample period and therefore introduce a slight look-ahead bias in the construction of the corrected dividend-price ratio. To alleviate the concern that the look-ahead bias may be influencing the results, I additionally estimate the so called No-Look-Ahead-Bias corrected dividend-price ratio:

$$dp_t^{NLAB\_Corr} = dp_t + \left(\frac{1}{1 - \widehat{\rho}_t \widehat{\gamma}_t} \right) idg_t \tag{38}$$

where $\widehat{\rho}_t$ and $\widehat{\gamma}_t$ are time-varying and estimated using the same recursive scheme as in the calculation of the out-of-sample R^2_{OS}.

Table 6 reports results. The dp_t that exhibits poor ability to predict returns in-sample also fails to predict returns out-of-sample. The R^2_{OS} for the dp_t is -0.15%. In comparison, the out-of-sample R^2 for the dp_t^{Corr} is as high as 6.06%. Thus, the dp_t^{Corr} does not predict returns only in-sample, but it also delivers superior out-of-sample forecasts of the monthly returns relative to the forecasts based on the historical average. Furthermore, approximately the same R^2_{OS}, if not even slightly higher, is also obtained with the corrected dividend-price ratio that is adjusted for the look-ahead bias (R^2_{OS} 6.09%). Hence, the look-ahead bias is not a concern and the dp_t^{Corr} can be effectively used in real time for the portfolio allocation decisions.[15]

To illustrate the relative success of the dp_t^{Corr} in predicting returns out-of-sample, Figure 6 plots out-of-sample forecasts along with the realized returns. Although realized returns are significantly more volatile than any of the forecasted returns, there are considerable differences between the forecasts.

[15] The rather small difference in the R^2_{OS} between the dp_t^{Corr} and the $dp_t^{NLAB\_Corr}$ is driven by the fact that the persistence of the implied dividend growth $\widehat{\gamma}_t$ and the linearization constant $\widehat{\rho}_t$ are very stable ($\widehat{\gamma}_t$ ranges from 0.47 to 0.61 and $\widehat{\rho}_t$ is always between 0.98 and 0.99). This makes dp_t^{Corr} and $dp_t^{NLAB\_Corr}$ highly correlated (0.99) and almost indistinguishable from each other.

The forecasts based on the realized dividend-price ratio and the forecasts based on the historical average return are both very smooth and almost indistinguishable from each other. In comparison, the forecasts based on the corrected dividend-price ratio vary significantly more and the changes of the forecasts are typically of the same sign as the changes of the realized returns.

1.6.3 Economic value of the corrected DP

To assess the economic value of the documented improvement in predicting returns, I run a simple out-of-sample trading strategy. I consider a mean-variance investor who invests in the stock market and the risk-free rate. Each period the investor uses different predictor variables to estimate one period ahead expected return $\widehat{\mu}_t$. Based on these estimates, the investor's portfolio weight on the stock market at time t is given by:

$$w_t = \frac{\widehat{\mu}_t - rf_{t+1}}{\gamma \widehat{\sigma}^2} \tag{39}$$

where rf_{t+1} is the one period ahead risk-free rate, γ is the risk-aversion coefficient and $\widehat{\sigma}^2$ is the variance of the stock market. I set γ equal to 3 and I proxy the variance of the market by the variance as implied in the options on the S&P 500 (VIX). The time-series of portfolio returns is then given by:

$$Rp_{t+1} = w_t r_{r+1} + (1 - w_t) rf_{t+1} \tag{40}$$

I assess economic value of predictors by the certainty equivalent return CE and the Sharpe ratio SR :

$$CE = \overline{Rp} - \frac{\gamma}{2}\widehat{\sigma}^2(Rp) \tag{41}$$

$$SR = \frac{\overline{Rp^e}}{\widehat{\sigma}(Rp^e)} \tag{42}$$

where $\overline{Rp}$ and $\widehat{\sigma}^2(Rp)$ are the mean and the variance of the portfolio return, and the superscript e stands for returns in excess of the risk-free rate. As in the calculation of the out-of-sample R^2_{OS}, I use the period from January 1994 through December 1999 for the estimation of the initial parameters and the period from January 2000 through 2009 for the calculation of the $CE's$ and the $SR's$. All out-of-sample forecasts are based on a recursive scheme using all the available information up to time t.

Table 7 reports results. The first column reports certainty equivalents and the second column reports Sharpe ratios. All the values are annualized. Note that the average excess return on the S&P 500 in the period from 2000 to 2009 is negative, which points at the difficulty of building trading strategies with positive Sharpe ratios. Indeed, a trading strategy based on the historical average return delivers a CE of 1.75%[16] and a negative Sharpe ratio (−0.10). Using dividend-price ratio to time the market yields slightly better results. The CE amounts to 2.78% and the Sharpe ratio becomes positive, but remains at the low level of 0.09.

In comparison, the corrected dividend-price ratio yields a CE as high as 5.07% and a Sharpe ratio of 0.41. This is a 0.32 gain in terms of the Sharpe ratio and a 2.29% gain in terms of the CE. In other words, an investor who is timing the market with the dp_t would be willing to pay as much as 2.29% of the invested wealth to get the access to the dp_t^{Corr}.

1.6.4 Alternative predictors

To further assess the return-predictive ability of the corrected dividend-price ratio, I consider a set of alternative return predictors.

Lacerda and Santa-Clara (2010) show that correcting the dividend-price ratio for the 10 year moving average of dividend growth improves predictability of longer horizon (i.e. annual) returns. Following their approach, I construct

[16] A comparable number for the period from 1947 to 2007 is 7.4% (Ferreira and Santa-Clara, 2010).

the dividend-price ratio corrected for the changes in the average historical dividend growth. Since the focus of this study lies on the short horizon predictability and the 10 year moving average of the dividend growth rate is rather slowly evolving, I calculate average historical dividend growth as a moving average of one year of annualized monthly dividend growth rates. Furthermore, to foster comparability with the dividend-price ratio corrected for the implied dividend growth, I assume that the persistence of the historical dividend growth is the same as the persistence of the implied dividend growth. The dividend-price ratio adjusted for the historical dividend growth is then defined as dp_t^{HIST} :

$$dp_t^{HIST} = dp_t + 2.08 * dg_t^{\overline{M}} \quad (43)$$

where $dg_t^{\overline{M}}$ is the moving average of annualized monthly dividend growth rates over the past year.

In addition, I use the variance risk premia (vrp_t) as implied in the S&P 500 (Bollerslev et al., 2009). The variance risk premia is arguably one of the strongest predictors for short horizon returns and is also estimated from the S&P 500 derivatives. Therefore, it is instructive to compare its return predictive ability to the one afforded by the corrected DP. This comparison also enables us to address the concern that the corrected dividend-price ratio is influenced by the time-varying risk premia.

Furthermore, I employ two other standard predictors, the earnings-price ratio (ep_t) and the consumption-to-wealth ratio (cay_t) proposed by Lettau and Ludvigson (2001).

Summary statistics. Table 8 reports the basic summary statistics and the unconditional correlation structure for the predictors sampled monthly.[17] The numbers are in line with previous studies. Except for the variance risk premia (vrp_t) and the corrected dividend-price ratio (dp_t^{Corr}), all the

[17] Since cay_t is available only at the quarterly frequency, monthly observations of cay_t are defined by the most recently available quarterly observation.

predictors are highly persistent with a first-order autocorrelation of more than 0.9. The persistence of the dividend-price ratio corrected for the average historical dividend growth (dp_t^{HIST}) is slightly lower than the persistence of the realized dividend-price ratio, but it is still of the similar magnitude (0.97).

The realized dividend-price ratio (dp_t), dividend-price ratio corrected for the historical dividend growth (dp_t^{HIST}) and the dividend-price ratio corrected for the implied dividend growth (dp_t^{Corr}) are all highly correlated and they exhibit similar relationships with respect to alternative predictors. They are all positively correlated with the earnings-price ratio and the consumption-to-wealth ratio and they are negatively related to the variance risk premia.

Predicting market returns. Table 9 reports results for predicting monthly returns. Since in-sample and out-of-sample results are largely consistent, I evaluate predictors mainly on the in-sample evidence. The traditional predictors based on the realized data explain only a small part of the variation in the future monthly returns. The earnings-price ratio exhibits a slightly negative *adj.* R^2. The dividend-price ratio, as already documented, explains 0.33% of the variation in the future monthly returns. The consumption-to-wealth ratio exhibits *adj.* R^2 of around one percent. Furthermore, correcting dividend-price ratio for the variation in the average historical dividend growth does not seem to improve predictability of monthly market returns in the analyzed period. The *adj.* R^2 in a univariate regression with the dp_t^{HIST} is approximately the same, if not slightly lower, as the *adj.* R^2 in the regression with the realized dividend-price ratio.[18]

[18] Unreported results show that this result is robust to using alternative proxies for the historical dividend growth, such as lagged monthly dividend growth, lagged annual dividend growth or the moving average of 10 years of monthly dividend growth rates.

Furthermore, approximately the same results are also obtained using the correction for the dividend-price ratio proposed by Lacera and Santa-Clara (2010):

$$dp_t^{LSC} = dp_t + dg_t^{\overline{M}} \left(\frac{1}{1-\rho_t}\right)$$

where ρ_t is a time-varying linearization constant and the correction does not depend on the persistence of the dividend growth. The only difference between the dp_t^{LSC} and the dp_t^{HIST} is that the dp_t^{LSC} is more volatile and less persistent because ρ_t is close to 0.98, implying that the typical correction for the dividend-price ratio is around $50 * dg_t^{\overline{M}}$ as opposed to the $2.08 * dg_t^{\overline{M}}$ used in this paper.

In comparison, the dividend-price ratio corrected for the implied dividend growth (dp_t^{Corr}) and the variance risk premia (vrp_t) explain a significantly higher portion of the variation in the future monthly returns. The vrp_t exhibits an *adj.* R^2 of 4.06%. The dp_t^{Corr}, as already documented, exhibits an *adj.* R^2 of 4.61%. The variance risk premia and the corrected dividend-price ratio are also the only predictors that are significant at the conventional levels of statistical significance.[19]

Since the corrected dividend-price ratio and the variance risk premia are both based on variables that are extracted from derivative markets, the relative success of the dp_t^{Corr} in predicting future returns could be driven by the fact that the dp_t^{Corr} is simply duplicating information contained in the vrp_t. To address this concern, I additionally consider a bivariate regression with the dp_t^{Corr} and the vrp_t. Quite interestingly, adding vrp_t as an additional predictor to the dp_t^{Corr} boosts statistical significance of both predictors and the *adj.* R^2 increases to as much as 9.99%. This result is even more remarkable because both predictors also predict returns out-of-sample with the R_{OS}^2 of 11.55%. This suggests that the corrected dividend-price ratio is not duplicating information about future returns that is already captured in other options implied variables. Also, it implies that the corrected dividend-price ratio is not capturing variation in expected returns that steams from changes in risk and risk aversion. Additional support for this interpretation is provided by the fact that a bivariate regression with the variance risk premia and, either the realized dividend-price ratio (dp_t), or the dividend-price ratio adjusted for the historical dividend growth (dp_t^{HIST}) results in a considerably smaller *adj.* R^2 (approximately 5%).

[19] Note that the results with the vrp_t are sensitive to the definition of the variance risk premia. A predictive regression with the variance risk premia defined as the difference between the VIX and the objective expectations of the realized variance (as opposed to the difference between the VIX and the actual realized variance) exhibits a slightly negative *adj.* R^2.

1.7 Variance decomposition of the DP

Until now I used the present value model merely to motivate the predictive regressions. In this Section, I employ the model to decompose the variance of the dividend-price ratio and to provide further insights for the interpretation of the results. As before, I treat the implied dividend growth as a true proxy for the expected dividend growth (no measurement error) and I use annualized variables taken at the monthly frequency.[20]

Within the framework of the present value model, the variance of the dividend-price ratio can be decomposed as:

$$\begin{aligned} var(dp_t) \simeq & \left(\frac{1}{1-\rho\delta_1}\right)^2 var(\mu_t) + \left(\frac{1}{1-\rho\gamma_1}\right)^2 var(g_t) \\ & -2 * \left(\frac{1}{1-\rho\delta_1}\right)\left(\frac{1}{1-\rho\gamma_1}\right) cov(\mu_t, g_t) \end{aligned} \quad (44)$$

where the first term on the right hand-side presents the contribution of the expected return, the second term presents the contribution of the expected dividend growth and the last term presents the contribution of the covariation between the expected return and the expected dividend growth.

I set $(1 - \widehat{\rho}_t\widehat{\delta_1})$ equal to the estimated parameter on the corrected dividend-price ratio (see equation (12)) and I calculate the variance of the expected return by inverting equation (7). I standardize all terms on the right-hand side of (44) by the left-hand side, so that the terms sum up to 100%.

The results imply that 400% of the variance of the dp_t is driven by the variation in the expected returns and 210% of the variance of the dp_t is driven by the time-varying expected dividend growth rate. This means that the covariance term account for as much as 510% of the variation in the

[20] Unreported results show that repeating the exercise with non-overlapping annual data does not change results qualitatively.

dividend-price ratio, which further implies that the correlation between the expected return and the expected dividend growth is as high as 0.88.

Thus, contrary to the standard result that virtually all the variation in the dividend-price ratio is driven by the time-varying expected returns (Campbell 1991; Cochrane 2005), the results show that there is a lot of variation in both expected returns and expected dividend growth. However, the positive correlation between them offsets each other within the dividend-price ratio. This dampens the volatility of the dividend-price ratio and explains why the dp_t fails to predict returns (and dividend growth rates) (see also Menzly et al., 2004 and Lettau and Ludvigson, 2005).

Furthermore, as discussed in Section 3, correcting the dividend-price ratio for changes in the expected dividend growth restores the variation that is offset by the positive correlation between the expected returns and the expected dividend growth. This makes the corrected dividend-price ratio more volatile than the uncorrected dividend-price ratio, and thus implies that expected returns vary significantly more than is suggested by the uncorrected dividend-price ratio.

1.8 Robustness checks

To validate the documented improvement in predicting market returns, I show that results are robust to several methodological changes in the calculation of the corrected dividend-price ratio.

1.8.1 Maturity of the implied dividend yield

The corrected dividend-price ratio analyzed throughout the paper is defined as:

$$dp_t^{Corr} = dp_t + 2.08 * idg_t = dp_t + 2.08 * (idy_t - dp_t). \qquad (45)$$

where idg_t is the log implied dividend growth rate calculated as the difference between the log implied dividend yield (idy_t) and the log realized dividend-price ratio (dp_t). Given the trade-off between the seasonality in dividend payments and the liquidity of the derivatives, the results in the main analysis are based on the annualized implied dividend yield with 6 months to maturity and the realized dividend-price ratio estimated in a standard way by summing dividends over the past 12 months. To address the concern that the maturity mismatch between the implied dividend yield and the realized dividend-price ratio could be a source of seasonality driving the documented improvement in predicting returns, I consider two robustness checks.

In the first robustness check, I re-estimate the corrected dividend-price ratio using implied dividend yields with maturities between 3 and 7 months. In the second robustness check, I repeat the same exercise, but instead of the standard dividend-price ratio with dividends summed over the past 12 months, I use a dividend-price ratio based on dividends summed over the past 6 months, $dp_t^{6m} = \log\left[\frac{D_t^6}{P_t}\right]$,where D_t^6 is the annualized sum of dividends over the past 6 months.

To foster comparability between the different versions of the corrected dividend-price ratios, I impose that the persistence of all the implied dividend growth rates is the same and equals the persistence of the implied dividend growth rate used in the main analysis. In other words, the corrected dividend-price ratio is always calculated as the dividend-price ratio plus 2.08 times the implied dividend growth.

Table 10 and Table 11 report results for the first and the second robustness check, respectively.[21] I start by analyzing results in Table X. The results seem to offer two general conclusions. First, irrespective of the maturity of the implied dividend yield, the corrected dividend-price ratio exhibits statistically significant predictive coefficients. Second, the maturity of the implied

[21] Due to the high volatility of the implied dividend yields with short maturity (see Table I), the implied dividend yield with maturity of 3 months takes a negative value on three occasions (December 1999, May 2000 and July 2000). I replace these observations with the implied dividend yield with maturity of 4 months.

dividend yield seems to matter. The corrected dividend-price ratio based on the implied dividend yields with longer maturities (5, 6 and 7 months) predicts returns better than the dp_t^{Corr} based on the implied dividend yields with maturities of 3 or 4 months. Since implied dividend yields with longer maturities are less prone to the seasonality in dividend payments, the results suggest that the documented improvement in the predictability of monthly returns is unlikely to be driven by the seasonality in dividend payments.

Results reported in Table 11 further reveal that the documented predictability is robust to the alternative way of constructing the realized dividend-price ratio. Specifically, the comparison of the results reported in Table 10 and Table 11 shows that the dividend-price ratio based on either half-annual or annual dividends exhibit identical *adj.* R^2 and the predictive coefficients differ only marginally. Also, the ability of the corrected dividend-price ratio to predict returns is largely unaffected by the alternative way of constructing the realized dividend-price ratio. All in all, results show that the documented improvement in predicting market returns cannot be explained by the maturity mismatch in the calculation of the implied dividend growth.

1.8.2 Moneyness and backward-looking data

The implied dividend yield used in the calculation of the corrected dividend-price ratio is estimated from no-arbitrage relations spanning the prices of index derivatives. Since no-arbitrage relations can be violated for particular pairs of options and futures, but hold well in general (Kamara and Miller, 1995), I calculate implied dividend yield by aggregating information from a wide set of options and futures. Each end of month, I use 10 days of backward-looking data and options across all the moneyness levels.

The use of such a wide set of data is necessary to smooth dividend yield estimates, but it may lead to inclusion of unreliable data. For example, the wider the moneyness level, the more observations we have for calculation of the implied dividend yield. Nevertheless, deep out-of-the money options are less liquid and therefore deemed unreliable. In the main run, I use all the options across all the moneyness levels. Now, I consider filtering out options with moneyness levels below 0.8 (0.9) and above 1.2 (1.1), respectively. Similar argument applies to the use of backward-looking data. We should expect that the most recent data is most important for forecasting purposes. However, more data may be needed to smooth the implied dividend yield estimates. In the main run, I use 10 days of backward-looking data. Now, I consider using either 5 or 15 days of backward-looking data. As before, I always calculate corrected dividend-price ratio as the dividend-price ratio plus 2.08 times the implied dividend growth.

Results are reported in Table 12 and confirm the above conjectures. Filtering out unreliable deep out-of-the money options improves the ability of the corrected dividend-price ratio to predict monthly returns. For example, the corrected dividend-price ratio based on options with moneyness levels between 0.9 and 1.1 explains as much as 5.82 percent of the variation in the future monthly market returns. Imposing even tighter restrictions on the moneyness levels should lead to even better results, but is unfortunately limited by the relatively low level of options liquidity at the beginning of the sample period. The liquidity issues are even more pronounced in the second exercise. Using 15 days as opposed to 10 days of backward-looking data does not seem to influence considerably either volatility of the corrected dividend-price ratio or its ability to predict future returns. In comparison, using 5 days of backward-looking data makes the corrected dividend-price ratio more prone to violations of no-arbitrage relations and significantly more volatile. This also reduces its ability to predict returns. Nevertheless, the estimated parameter on the dp_t^{Corr} remains significant and it still predicts returns significantly better than the uncorrected dividend-price ratio.

1.9 Conclusions

That variation in the expected dividend growth reduces the ability of the dividend-price ratio to predict returns is a long-standing notion in the predictability literature (Fama and French, 1988). However, empirical analysis of this issue is complicated because the expected dividend growth is an aggregate of investors' expectations about future growth opportunities and is therefore difficult to estimate.

In this paper I propose extracting the expected dividend growth from derivative markets (index options and futures). Because prices of derivatives depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts, they provide a unique laboratory for estimating the dividend growth that investors expect to realize in the near future. Indeed, I find that the implied dividend growth uncovers variation in future dividend growth and thereby allows for improvements in predicting market returns. Using implied dividend growth as an additional regressor in the standard dividend-price ratio return predictive regression—or correcting the dividend-price ratio for variation in the implied dividend growth—significantly improves the predictability of short-run S&P 500 returns over the past 16 years.

This predictive improvement is especially strong over a short horizon (i.e. monthly returns), holds both in-sample and out-of-sample, yields a sizable gain in the Sharpe ratio, and is robust to small sample bias. Furthermore, these results are not driven by the fact that implied dividend growth duplicates information in other, well-known options-implied predictors (e.g., the variance risk premia), and neither can they be replicated using historical rather than implied dividend growth.

Importantly, the results show that the expected return and expected dividend growth are highly correlated. This high correlation means that movements in expected returns and expected dividend growth offset each other's effect in the dividend-price ratio, which renders the dividend-price ratio relatively smooth. Correcting the dividend-price ratio for the implied dividend growth restores the variation which is otherwise obscured by this strong comovement, and hence implies that expected returns vary significantly more than is apparent from the observed variation in the uncorrected dividend-price ratio.

1.10 APPENDIX

Derivation of the present value model

Define log return r_{t+1}, log dividend growth Δd_{t+1}, and log dividend-price ratio dp_t as:

$$r_{t+1} = \log\left[\frac{P_{t+1} + D_{t+1}}{P_t}\right], \quad \Delta d_{t+1} = \log\left[\frac{D_{t+1}}{D_t}\right], \quad dp_t = \log\left[\frac{D_t}{P_t}\right] \quad (46)$$

Let $\mu_t = E_t(r_{t+1})$ be conditional expected return and let $g_t = E_t(\Delta d_{t+1})$ be conditional expected dividend growth. Assume that μ_t and g_t follow $AR(1)$ processes:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t) + \varepsilon^{\mu}_{t+1} \quad (47)$$
$$g_{t+1} = \gamma_0 + \gamma_1(g_t) + \varepsilon^{g}_{t+1} \quad (48)$$
$$\Delta d_{t+1} = g_t + \varepsilon^{d}_{t+1} \quad (49)$$

where ε^{μ}_{t+1}, ε^{g}_{t+1} and ε^{d}_{t+1} are zero mean errors.

Derive Campbell and Shiller (1988) present value identity:

Rewrite log returns as:

$$r_{t+1} = dp_t + \Delta d_{t+1} + \log\left[1 + \exp(-dp_{t+1})\right] \quad (50)$$

Use first order Taylor expansion to linearize $\log\left[1 + \exp(-dp_{t+1})\right]$ around $\overline{dp} = E(dp_t)$:

$$\log\left[1 + \exp(-dp_{t+1})\right] \simeq \log\left[1 + \exp(-\overline{dp})\right] + \frac{\exp(-\overline{dp})}{1 + \exp(-\overline{dp})}\left[-dp_{t+1} + \overline{dp}\right] \quad (51)$$

Define $\rho = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})}$ and $\kappa = \log\left[1 + \exp(-\overline{dp})\right] + \rho\overline{dp}$, such that:

$$\log\left[1+\exp(-dp_{t+1})\right] \simeq \kappa - \rho dp_{t+1} \tag{52}$$

Plug (7) into (5) to get an expression for the one-period return:

$$r_{t+1} \simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} \tag{53}$$

Iterate equation (8) forward:

$$dp_t \simeq -\kappa + \rho dp_{t+1} + r_{t+1} - \Delta d_{t+1} \tag{54}$$

$$dp_t \simeq -\kappa + \rho\left(-\kappa + \rho dp_{t+2} + r_{t+2} - \Delta d_{t+2}\right) + r_{t+1} - \Delta d_{t+1} \tag{55}$$

$$dp_t \simeq -\kappa - \kappa\rho + \rho^2 dp_{t+2} + r_{t+1} + \rho r_{t+2} - \Delta d_{t+1} - \rho\Delta d_{t+2} \tag{56}$$

$$dp_t \simeq -\frac{\kappa}{1-\rho} + \rho^{\infty} dp_{t+\infty} + \sum_{j=0}^{\infty} \rho^j \left(r_{t+1+j} - \Delta d_{t+1+j}\right) \tag{57}$$

Assume that $\lim_{j\to\infty}\rho^j dp_{t+j} = 0$ to obtain the Campbell and Shiller (1988) approximation for the log dividend-price ratio (since the relationship holds ex-ante and ex-post, an expectation operator can be added to the right hand sight):

$$dp_t \simeq -\frac{\kappa}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j \left(r_{t+1+j}\right) - E_t \sum_{j=0}^{\infty} \rho^j \left(\Delta d_{t+1+j}\right) \tag{58}$$

Combine Campbell and Shiller (1988) present value identity with the AR(1) processes for the expected return and the expected dividend growth to solve for the dividend-price ratio:

Iterate equations (2) and (3) forward to obtain:

$$E_t(r_{t+1+j}) = \delta_0 \frac{1-\delta_1^j}{1-\delta_1} + \delta_1^j \mu_t \tag{59}$$

$$E_t(\Delta d_{t+1+j}) = \gamma_0 \frac{1-\gamma_1^j}{1-\gamma_1} + \gamma_1^j g_t \tag{60}$$

Work out the expectations of $E_t \sum_{j=0}^{\infty} \rho^j \left(r_{t+1+j}\right)$ and the $E_t \sum_{j=0}^{\infty} \rho^j \left(\Delta d_{t+1+j}\right)$:

$$E_t \sum_{j=0}^{\infty} \rho^j (r_{t+1+j}) = \sum_{j=0}^{\infty} \rho^j (\delta_0 \frac{1-\delta_1^j}{1-\delta_1} + \delta_1^j \mu_t) \tag{61}$$

$$= \frac{\delta_0}{1-\delta_1} \sum_{j=0}^{\infty} \rho^j - \frac{\delta_0}{1-\delta_1} \sum_{j=0}^{\infty} \rho^j \delta_1^j + \mu_t \sum_{j=0}^{\infty} \rho^j \delta_1^j \tag{62}$$

$$= \frac{\delta_0}{(1-\delta_1)(1-\rho)} - \frac{\delta_0}{(1-\delta_1)(1-\rho\delta_1)} + \mu_t \left(\frac{1}{1-\rho\delta_1} \right) \tag{63}$$

$$E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j}) = \sum_{j=0}^{\infty} \rho^j (\gamma_0 \frac{1-\gamma_1^j}{1-\gamma_1} + \gamma_1^j g_t) \tag{64}$$

$$= \frac{\gamma_0}{1-\gamma_1} \sum_{j=0}^{\infty} \rho^j - \frac{\gamma_0}{1-\gamma_1} \sum_{j=0}^{\infty} \rho^j \delta_1^j + g_t \sum_{j=0}^{\infty} \rho_1^j \gamma_1^j \tag{65}$$

$$= \frac{\gamma_0}{(1-\gamma_1)(1-\rho)} - \frac{\gamma_0}{(1-\gamma_1)(1-\rho\gamma_1)} + g_t \left(\frac{1}{1-\rho\gamma_1} \right) \tag{66}$$

Finally, insert (18) and (21) in the Campbell and Shiller (1988) present value identity to find the dividend-price ratio:

$$dp_t \simeq \varphi + \mu_t \left(\frac{1}{1-\rho\delta_1} \right) - g_t \left(\frac{1}{1-\rho\gamma_1} \right) \tag{67}$$

where $\varphi = -\frac{\kappa}{1-\rho} + \frac{\delta_0}{(1-\delta_1)(1-\rho)} - \frac{\delta_0}{(1-\delta_1)(1-\rho\delta_1)} - \frac{\gamma_0}{(1-\gamma_1)(1-\rho)} + \frac{\gamma_0}{(1-\gamma_1)(1-\rho\gamma_1)}$.

1.11 TABLES AND FIGURES

Table 1: Implied dividend yields with different maturities

This table reports the summary statistics (Panel A) and the unconditional correlations (Panel B) for the implied dividend yields (IDY) with maturities between 3 months (3m) and 7 months (7m). The period is from January 1994 through December 2009.

	IDY $(3m)$	IDY $(4m)$	IDY $(5m)$	IDY $(6m)$	IDY $(7m)$
Panel A: Summary statistics					
Mean	0.0207	0.0206	0.0207	0.0206	0.0203
Std. Dev.	0.0116	0.0086	0.0077	0.0074	0.0073
Panel B: Unconditional correlations					
IDY $(3m)$	1.0000	0.8702	0.7932	0.7776	0.7541
IDY $(4m)$	.	1.0000	0.9565	0.9260	0.8769
IDY $(5m)$	.	.	1.0000	0.9781	0.9252
IDY $(6m)$	.	.	.	1.0000	0.9789
IDY $(7m)$	.	.	.	.	1.0000

Table 2: Summary statistics

This table reports the summary statistics (Panel A) and the unconditional correlations (Panel B) for annualized S&P 500 log monthly returns (r_t^M), annualized log monthly dividend growth rates (Δd_t^M), log dividend-price ratio (dp_t), annualized log implied dividend growth rates (idg_t), and log corrected dividend-price ratio (dp_t^{Corr}). The period is from January 1994 through December 2009.

	r_t^M	Δd_t^M	dp_t	idg_t	dp_t^{Corr}
Panel A: Summary statistics					
Mean	0.0733	0.0361	-4.0198	0.0602	-3.8946
Std. Dev.	0.5443	0.1463	0.2687	0.1868	0.5364
Skewness	-0.9465	-0.1728	0.3758	-1.0235	-0.9809
Kurtosis	4.6857	4.2485	2.6228	5.0496	4.6206
AR (1)	0.1292	0.0871	0.9786	0.5286	0.7359
Panel B: Unconditional correlations					
r_t^M	1.0000	-0.0328	-0.0735	0.0725	0.0157
Δd_t^M	.	1.0000	-0.0854	0.1641	0.0761
dp_t	.	.	1.0000	0.3091	0.7248
idg_t	.	.	.	1.0000	0.8793
$dp_t^{Corr.}$	.	.	.	.	1.0000

Table 3: Dividend growth regressions

This table reports in-sample regression results for predicting annualized S&P 500 log dividend growth. The predictor variables include log dividend-price ratio (dp_t) and log implied dividend growth (idg_t). All of the regressions are based on monthly observations. For regressions with non-overlapping observations ($h = 1$), t-statistics are calculated according to OLS and are reported in parentheses. For regressions with overlapping observations ($h = 2, 3$ or 6), t-statistics are computed according to Hodrick (1992) and are reported in brackets. The period is from January 1994 through December 2009.

Dividend growth horizon (h)	1	2	3	6
Const.	-0.2728	-0.2644	-0.2423	-0.2013
	(-1.728)	[-1.506]	[-1.344]	[-1.164]
dp_t	-0.0767	-0.0749	-0.0696	-0.0599
	(-1.958)	[-1.735]	[-1.571]	[-1.405]
adj. R^2	0.0147	0.0304	0.0373	0.0328
Const.	0.0248	0.0248	0.0261	0.0283
	(2.284)	[2.206]	[2.264]	[2.517]
idg_t	0.1800	0.1983	0.1906	0.1902
	(3.249)	[4.401]	[4.262]	[4.761]
adj. R^2	0.0479	0.1160	0.1501	0.1842
Const.	-0.4919	-0.5019	-0.4713	-0.4415
	(-3.067)	[-2.463]	[-2.239]	[-2.187]
dp_t	-0.1276	-0.1301	-0.1228	-0.1158
	(-3.229)	[-2.626]	[-2.402]	[-2.355]
idg_t	0.2368	0.2562	0.2454	0.2440
	(4.165)	[4.630]	[4.393]	[4.826]
adj. R^2	0.0931	0.2093	0.2661	0.3081

Table 4: Return regressions

This table reports in-sample regression results for predicting annualized log S&P 500 returns. The predictor variables include log dividend-price ratio (dp_t), log implied dividend growth (idg_t), and log corrected dividend-price ratio (dp_t^{Corr}). All of the regressions are based on monthly observations. For regressions with non-overlapping observations ($h = 1$), t-statistics are calculated according to OLS and are reported in parentheses. For regressions with overlapping observations ($h = 2, 3$ or 6), t-statistics are computed according to Hodrick (1992) and are reported in brackets. The period is from January 1994 through December 2009.

Return horizon (h)	1	2	3	6
Const.	0.8234	0.9316	0.9726	1.1540
	(1.393)	[1.339]	[1.418]	[1.784]
dp_t	0.1870	0.2137	0.2238	0.2686
	(1.275)	[1.241]	[1.317]	[1.672]
adj. R^2	0.0033	0.0145	0.0264	0.0702
Const.	0.1689	0.5018	0.6385	0.9402
	(0.276)	[0.634]	[0.832]	[1.317]
dp_t	0.0348	0.1137	0.1461	0.2189
	(0.231)	[0.589]	[0.779]	[1.250]
idg_t	0.7075	0.4637	0.3579	0.2173
	(3.271)	[2.066]	[1.712]	[1.201]
adj. R^2	0.0520	0.0503	0.0571	0.0871
Const.	0.9647	0.7863	0.7061	0.6404
	(3.415)	[3.097]	[2.802]	[2.715]
dp_t^{Corr}	0.2292	0.1833	0.1626	0.1458
	(3.192)	[2.743]	[2.445]	[2.333]
adj. R^2	0.0461	0.0529	0.0619	0.0856

Table 5: In-sample bias correction

This table reports the effect of small sample bias on the statistical significance of predictor variables in the regressions for predicting annualized log monthly S&P 500 returns (r_{t+1}^{M}). β and $t-stat.$ are the slope estimate and its corresponding t-statistic according to OLS. β^{c} and $t^{c}-stat.$ are the slope estimate and its corresponding t-statistic according to Amihud and Hurvich (2004) bias correction methodology (see regression (33) and equation (34) in the main text). The predictor variables include log dividend-price ratio (dp_t) and log corrected dividend-price ratio (dp_t^{Corr}). The correlation between the innovations to the predictor variable and the errors of the predictive regression is denoted by ρ. The period is from January 1994 through December 2009.

Dependent variable: r_{t+1}^{M}

	OLS	Bias correction
	$\beta_x, (t_x)$	$\beta_x^c, (t_x^c), [\rho]$
dp_t	0.1870	-0.0407
	(1.275)	(-0.273)
		[-0.966]
$dp_t^{Corr.}$	0.2292	0.2234
	(3.192)	(3.098)
		[-0.236]

Table 6: Out-of-sample predictability

This table reports out-of-sample R^2_{OS} for predicting annualized log monthly S&P 500 returns $\left(r^M_{t+1}\right)$ The predictor variables include log dividend-price ratio $\left(dp_t\right)$, log corrected dividend-price ratio $\left(dp^{Corr}_t\right)$, and log corrected dividend-price ratio adjusted for the look-ahead bias $\left(dp^{Corr-NLAB}_t\right)$. The R^2_{OS} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast).

Dependent variable: r^M_{t+1}	
	R^2_{OS}
dp_t	-0.0015
dp^{Corr}_t	0.0606
$dp^{Corr-NLAB}_t$	0.0609

Correlation$(dp^{Corr}_t, dp^{Corr-NLB}_t)$=0.9985

Table 7: Trading strategies

This table reports certainty equivalent (CE) and Sharpe ratio (SR) of a trading strategy based on timing log monthly S&P 500 returns with different predictor variables. The predictor variables include historical average return estimated through period t $(\overline{r_t})$, log dividend-price ratio (dp_t), log corrected dividend-price ratio $\left(dp_t^{Corr}\right)$, and log corrected dividend-price ratio adjusted for the look-ahead bias $\left(dp_t^{Corr-NLAB}\right)$. The CE and the SR are calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). All the values are annualized.

	CE	SR
$\overline{r_t}$	0.0175	-0.1035
dp_t	0.0278	0.0888
dp_t^{Corr}	0.0507	0.4113
$dp_t^{Corr-NLAB}$	0.0512	0.4170

Table 8: Alternative predictors: Summary statistics

This table presents the summary statistics (Panel A) and the unconditional correlations (Panel B) for the log dividend-price ratio (dp_t), log dividend-price ratio corrected for the implied dividend growth $\left(dp_t^{Corr}\right)$, log dividend-price ratio corrected for the 12 month average monthly dividend growth $\left(dp_t^{HIST}\right)$, variance risk premia (vrp_t), log earnings-price ratio (ep_t), and consumption-to-wealth ratio (cay_t) (monthly observations of cay_t are defined by the most recently available quarterly observation). The period is from January 1994 through December 2009.

	dp_t	dp_t^{Corr}	dp_t^{HIST}	vrp_t	ep_t	cay_t
Panel A: Summary statistics						
Mean	-4.0198	-3.8946	-3.9282	18.2236	-3.2247	-0.0016
Std. Dev.	0.2687	0.5364	0.3007	22.2653	0.4179	0.0204
Skewness	0.3758	-0.9809	-0.2751	-2.8320	-2.2244	0.1480
Kurtosis	2.6228	4.6206	2.1317	36.9296	9.0900	1.8789
AR (1)	0.9786	0.7359	0.9687	0.3014	0.9419	0.9669
Panel B: Unconditional correlations						
dp_t	1.0000	0.7248	0.8858	-0.0979	0.1497	0.4858
$dp_t^{Corr.}$	.	1.0000	0.7416	-0.1146	0.3286	0.3402
dp_t^{HIST}	.	.	1.0000	-0.2295	0.4906	0.2622
vrp_t	.	.	.	1.0000	-0.3095	0.1349
ep_t	.	.	.	.	1.0000	0.0224
cay_t	.	.	.	.	.	1.0000

Table 9: Alternative predictors: Monthly return regressions

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t), log dividend-price ratio corrected for the implied dividend growth $\left(dp_t^{Corr}\right)$, log dividend-price ratio corrected for the 12 month average monthly dividend growth $\left(dp_t^{HIST}\right)$, variance risk premia (vrp_t), log earnings-price ratio (ep_t), and consumption-to-wealth ratio (cay_t).

Dependent variable: r_{t+1}^{M}

Const.	0.8234 (1.393)	0.9647 (3.415)	0.6552 (1.261)	-0.0233 (-0.467)	-0.0695 (-0.221)	0.0765 (1.943)	0.9045 (1.565)	0.9671 (3.525)	0.9454 (1.842)
dp_t	0.1870 (1.275)						0.2320 (1.611)		
dp_t^{Corr}		0.2292 (3.192)						0.2576 (3.667)	
dp_t^{HIST}			0.1486 (1.126)						0.2503 (1.897)
vrp_t				0.0052 (3.007)			0.0054 (3.165)	0.0059 (3.506)	(0.0059) (3.381)
ep_t					-0.0438 (-0.452)				
cay_t						3.1200 (1.621)			
adj. R^2	0.0033	0.0461	0.0014	0.0406	-0.0042	0.0085	0.0487	0.0999	0.0537
R_{OS}^{2}	-0.0015	0.0606	-0.0033	0.0450	-0.0157	0.0083	0.0447	0.1155	0.0496

Table 10: Robustness check: Maturity of implied dividend yield I

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t), and log corrected dividend-price ratio $\left(dp_t^{Corr}(T)\right)$ defined as $dp_t^{Corr.}(T) = dp_t + 2.08 * (idy_t^T - dp_t)$, where idy_t^T is log implied dividend yield with maturities between 3 months $(3m)$ and 7 months $(7m)$.

Dependent variable: r_{t+1}^{M}

Const.	0.4054 (2.905)	0.4154 (2.595)	0.8284 (3.283)	0.9647 (3.415)	1.0117 (3.458)	0.8234 (1.393)
$dp_t^{Corr}(3m)$	0.0849 (2.491)					
$dp_t^{Corr}(4m)$		0.0871 (2.215)				
$dp_t^{Corr}(5m)$			0.1939 (3.035)			
$dp_t^{Corr}(6m)$				0.2292 (3.192)		
$dp_t^{Corr}(7m)$					0.2408 (3.241)	
dp_t						0.1870 (1.275)
adj. R^2	0.0267	0.0201	0.0414	0.0461	0.0477	0.0033
R_{OS}^2	0.0365	0.0268	0.0560	0.0606	0.0608	-0.0015

Table 11: Robustness check: Maturity of implied dividend yield II

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include annualized log dividend-price ratio based on dividends summed over the past 6 months $\left(dp_t^{6m}\right)$ and log corrected dividend-price ratio $\left(dp_t^{Corr}(T)\right)$ defined as $dp_t^{Corr}(T) = dp_t^{6m} + 2.08 * (idy_t^T - dp_t^{6m})$, where idy_t^T is log implied dividend yield with maturities between 3 months $(3m)$ and 7 months $(7m)$.

Dependent variable: r_{t+1}^{M}

Const.	0.4066 (2.907)	0.4166 (2.597)	0.8327 (3.290)	0.9701 (3.423)	1.0154 (3.462)	0.8303 (1.396)
$dp_t^{Corr}(3m)$	0.0850 (2.494)					
$dp_t^{Corr}(4m)$		0.0872 (2.217)				
$dp_t^{Corr}(5m)$			0.1946 (3.042)			
$dp_t^{Corr}(6m)$				0.2301 (3.200)		
$dp_t^{Corr}(7m)$					0.2412 (3.246)	
dp_t^{6m}						0.1890 (1.278)
adj. R^2	0.0268	0.0202	0.0417	0.0464	0.0478	0.0033
R_{OS}^2	0.0366	0.0269	0.0563	0.0609	0.0610	-0.0000

Table 12: Robustness check: Options moneyness and backward-looking data

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t) and log corrected dividend-price ratio $\left(dp_t^{Corr}\right)$ defined as $dp_t^{Corr} = dp_t + 2.08 * (idy_t - dp_t)$, where idy_t is log implied dividend yield calculated from different number of days of backward-looking data (d), or based on restricted moneyness levels of options (M).

Dependent variable: r_{t+1}^{M}

Const.	0.9647 (3.415)	0.9743 (3.537)	0.8224 (3.847)	0.6675 (2.639)	0.9849 (3.309)	0.8234 (1.393)
dp_t^{Corr}	0.2292 (3.192)					
$dp_t^{Corr}(0.8 \leq M \leq 1.2)$		0.2315 (3.309)				
$dp_t^{Corr}(0.9 \leq M \leq 1.1)$			0.1900 (3.570)			
$dp_t^{Corr}(-5 \leq d \leq 0)$				0.1531 (2.385)		
$dp_t^{Corr}(-15 \leq d \leq 0)$					0.2341 (3.095)	
dp_t						0.1870 (1.275)
adj. R^2	0.0461	0.0498	0.0582	0.0241	0.0432	0.0033
R_{OS}^{2}	0.0606	0.0632	0.0594	0.0235	0.0599	-0.0015

Figure 1: Open interest by maturity

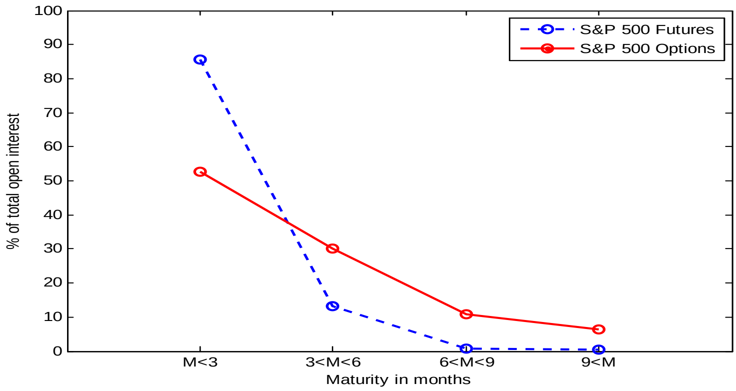

This figure plots the percentage of open interest by maturity for S&P 500 options and S&P 500 futures. The percentage of open interest is calculated as the total open interest for a given maturity over the total open interest for all the maturities. The period is from January 1994 through December 2009.

Figure 2: Term structure for the implied dividend yields

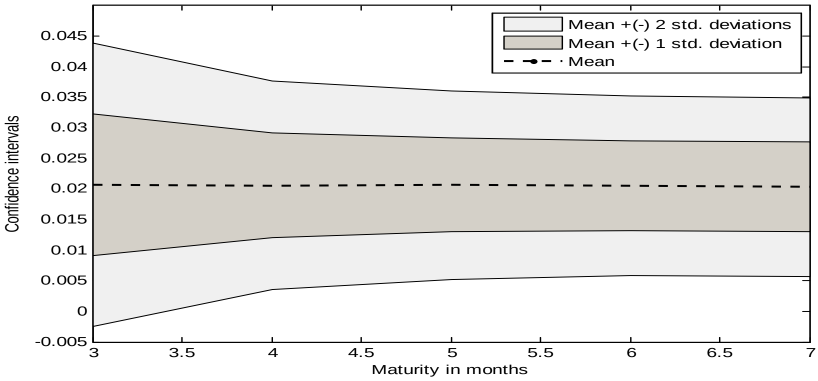

This figure plots the term structure for the implied dividend yields. The figure is based on the summary statistics for the implied dividend yields reported in Table I. Dashed line denotes the mean of the implied dividend yields with different maturities. Dark grey color denotes the area that is one standard deviation away from the mean of the implied dividend yields. Bright grey color denotes the area that is two standard deviations away from the mean of the implied dividend yields. The period is from January 1994 through December 2009.

Figure 3: Implied dividend growth

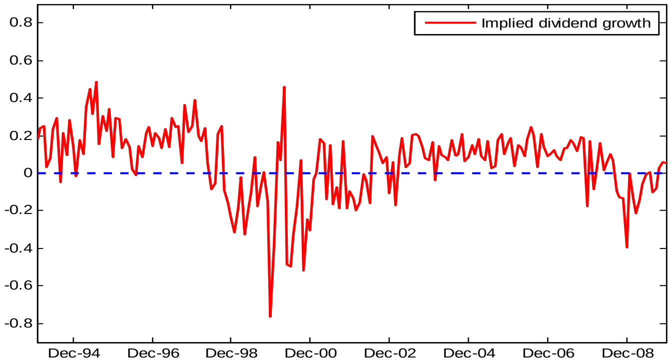

This figure plots log implied dividend growth for the S&P 500. The period is from January 1994 through December 2009.

Figure 4: Dividend-price ratio and corrected dividend-price ratio

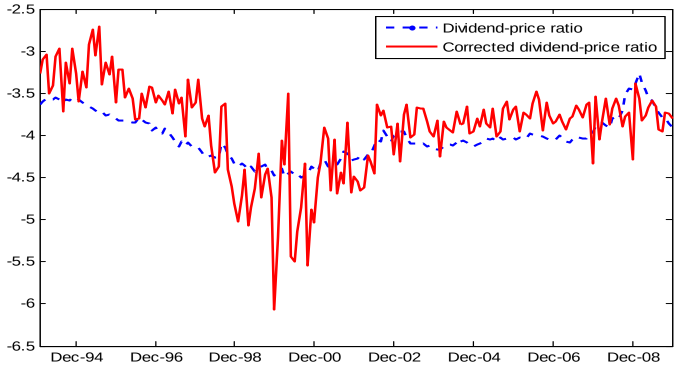

This figure plots log dividend-price ratio and log corrected dividend-price ratio for the S&P 500. The period is from January 1994 through December 2009.

Figure 5: Implied interest rate, T-bill rate and LIBOR rate

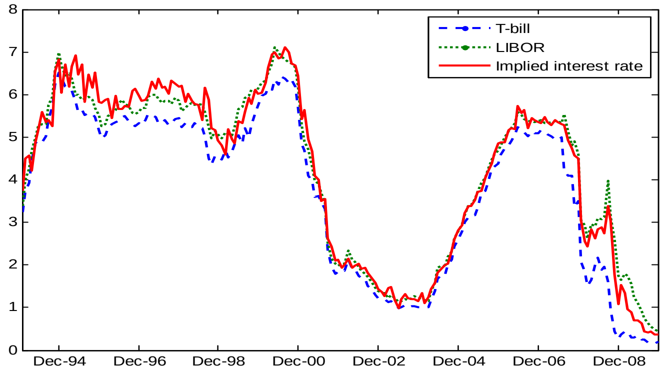

This figure plots the 6-month implied interest rate along with the 6-month T-bill rate and the 6-month LIBOR rate. The period is from January 1994 through December 2009.

Figure 6: Realized vs. forecasted returns

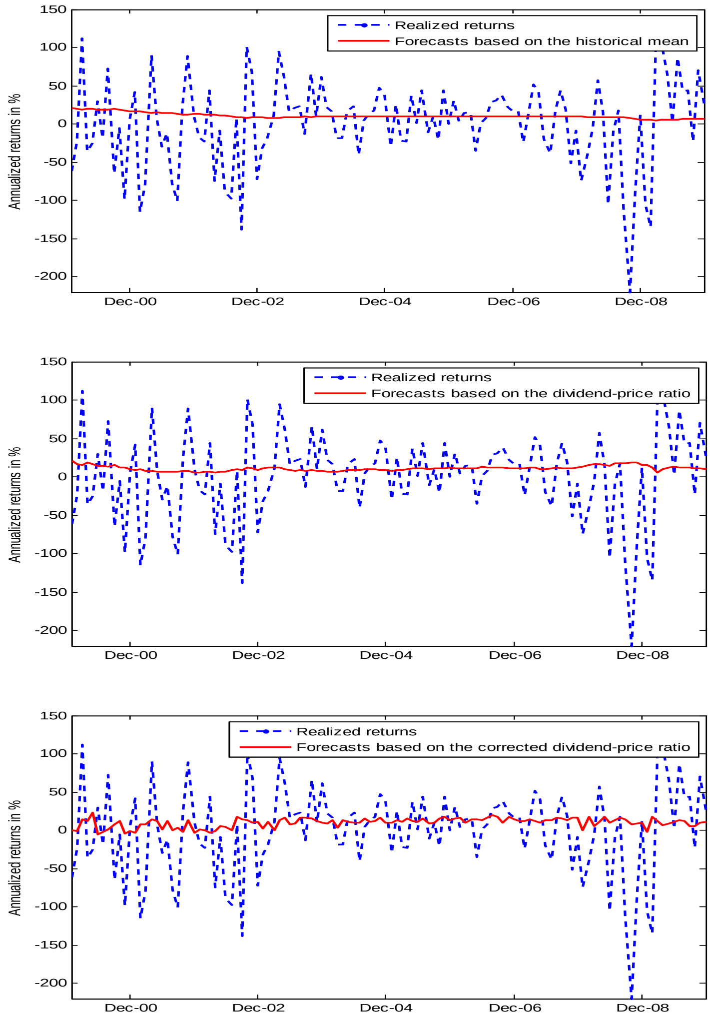

The figures plot annualized monthly S&P 500 log returns along with the return forecasts. The forecasts are based either on historical average return (upper figure), on the dividend-price ratio (middle figure), or on the corrected dividend-price ratio (lower figure). The period is from January 1994 through December 2009.

Chapter 2

PINNING IN THE S&P 500 FUTURES (WITH JENS C. JACKWERTH)

2.1 Introduction

From first principles, we would expect stock prices to be uniformly distributed on any small interval - there should not be any attraction to one particular stock price or another. However, pinning exactly describes such tendency of stock prices to finish at the expiration date of an option more frequently near a strike price.[3] This is a fascinating feature as it involves effects across two markets: the options market and the market for the underlying asset.

Pinning has been documented for individual stocks, see the instances described in Anders (1982), Krishnan and Nelken (2001), or Augen (2009, pp. 26). Ni, Pearson, and Poteshman (2005) study stock option pinning and provide statistical evidence of its existence. In their paper, the main driving force for pinning is the market maker's adjustment of the delta hedge due to the time-decay of the hedges, according to Avellaneda and Lipkin (2003), and stock price manipulation of proprietary traders.

In this paper, we take the analysis to the aggregate level and analyze the behavior of S&P 500 futures (henceforth futures) on expiration days of options on S&P 500 futures (henceforth SP options). Since SP options expire on a monthly cycle and futures expire on a quarterly cycle, we primarily focus on serial expiration months (all months excluding the quarterly cycle) as those days provide us with a unique laboratory of cases when SP options expire and the underlying future continues to trade for an additional month or two.

[3] The fact that pinning occurs only on expiration dates is different from clustering which is the tendency of prices to be quoted on particular round values. Such clustering is independent of a day being an expiration day or not. See Schwarz, Van Ness, and Van Ness (2004) for a recent account of clustering in S&P 500 futures trade prices.

It is exactly this feature that enables the future to pin to the nearest strike price.

As futures are highly liquid (the typical expiration day notional open interest from August 1987 until November 2009 was some 90 billion dollars and each expiration day some 15% thereof were traded), it is hard to imagine that futures could be subject to manipulation and we will provide evidence to this effect. Further, as opposed to individual stocks, where likely delta hedgers (market makers) tend to hold long option positions, market makers are typically short the options on the S&P 500 index (henceforth SPX options) (Garleanu, Pedersen, and Poteshman, 2009). We will argue that this fact extends to the very similar market for SP options. Given such short position of the market maker in the SP options, the time decay of the delta hedge should then lead, according to the model of Avellaneda and Lipkin (2003), to anti-pinning in the S&P 500 futures and not to pinning.

Surprisingly however, we find evidence of pinning in the serial expirations of closest to maturity S&P 500 futures and not of the predicted anti-pinning. We document this behavior in Figure 1 where we depict the percentage of future settlement prices finishing within \$0.25 of the closest strike price. Assuming that future prices are uniformly distributed between the option strike prices – which are spaced \$5 apart – we would expect 10% of the futures prices to finish within \$0.25 of the closest strike price.[2]

Nevertheless, from August 1987 until November 2009 we see in Figure 1, Panel A that the expected frequency of 10% is elevated on expiration days as opposed to the five preceding or following days and amounts to 13.56%. This implies that on each expiration date 3.56% of all futures prices moved by at least 50 cents in absolute terms. The associated average price move is then at least 1.6 billion dollars in notional terms on each expiration date. Panel B documents that this effect is even stronger from October 1998 until November

[2] Clearly, over longer horizons the futures prices are distributed according to some more complicated distribution which is often assumed to be the lognormal. However, locally over any small interval, any continuous distribution will be rather well approximated through a uniform distribution.

2009. In this period, 9.10% of all futures prices have their prices altered on the options expiration dates. This effect is significant in statistical terms and translates into a move in notional terms of at least 7.0 billion dollars on each expiration day. The effect could be large enough to enable profitable trading as the bid/ask spread in the futures is most of the time only 10 cents, thus much smaller than the 50 cent shift which we document. Below we suggest a simple trading strategy where we hedge our exposure with the second to maturity future. We find this trading strategy to be statistically significantly more profitable on expirations where there exists pinning than on expirations where there is no pinning. The annualized returns of the weekend trading strategy amount to a surprisingly large 24%. Our findings might thus also be of interest to the exchange and regulators due to the large size of the distortions and the importance of this market.

In additional tests, we find that pinning is especially pronounced from above the strike price. Intuitively, this asymmetry is due to limits to arbitrage: pinning from above is more difficult to arbitrage away as it involves buying the depressed future. In turn, one would need to sell (short) the S&P 500 basket which is difficult or resort to the much less liquid future with the next longer maturity. Thus, the transaction costs on the related arbitrage are large and we can have periods of time where the future price is pinned from above.

Given that the observed pinning is seemingly at odds with the main story for pinning due to Avellaneda and Lipkin (2003) we explore in detail other potential explanations for pinning. Anders (1982) suggests that last minute sales of in-the-money option by individual investors lead to pinning as the market maker needs to adjust the hedge afterwards. By the same token, pinning can also arise because of early exercise of options. We test the competing three mechanisms via logistic regressions which explain pinning and anti-pinning based on option volume, open interest, and early option exercise. To the best of our knowledge, no other study used early options exercise data in our context.

Our regressions confirm that the time decay of the delta hedge of Avellaneda and Lipkin (2003) indeed leads to anti-pinning, but the effects of the other two mechanisms are overcompensating. Further, the results confirm that manipulation is an unlikely explanation for the documented pinning. Robustness checks find these results stable with regard to changes in the methodology.

To corroborate that the documented pinning is indeed related to options expirations, we show that there is no pinning in second to maturity futures, on which there exist no expiring SP options. Also, there is no pinning in the first to maturity futures on the quarterly expirations when SP options and futures expire simultaneously in the value of the S&P 500 index and cash-settlement is used. This is understandable since the S&P 500 basket is much harder to move through trading than the future. For the same reason, there is neither pinning in the S&P 500 index itself due to expiration of SPX options on the S&P 500 index nor in the exchanged traded fund on the S&P 500 (SPDR) due to expiration of its SPY options.

Based on a literature review, the paper develops the hypotheses in Section 2.2. Section 2.3 introduces the econometric methodology of testing for pinning and documenting the driving mechanisms. All data are presented in Section 2.4. Results for different option classes follow in Section 2.5 while robustness checks are presented in Section 2.6. Section 2.7 concludes.

2.2 Hypotheses and literature

We motivate our study by evidence that pinning exists in the near to maturity futures on serial expiration days of SP options. Next, we turn to possible reasons as to why such pinning might occur. Since many arguments relate to the delta hedging of the market maker, we argue that the market maker tends to be short at-the-money straddles (positions of sold calls and puts in roughly equal number) and that only the market maker delta hedges.

We know that the market maker in the S&P 500 index options is short gamma, thus selling mainly straddles (see Table 1 in Garleanu, Pedersen, and Poteshman 2009). Unfortunately, such proprietary data is unobtainable for the futures options market. However, we argue that market maker positions are likely to be rather similar since the two markets are closely related. The correlation between the S&P 500 index and the shortest to maturity future in the period from 1983 to 2009 is 0.9999. Further, trading activity in the SPX options market and the SP options market are highly related. Correlations of near-the-money open interest and volume between the two markets during the last 5 days leading up to expiration Friday are 0.86 and 0.79, respectively.[3] This leads us to assume that the market maker in the SP option market holds similar positions as in the SPX options market, i.e. the market maker is typically short at-the-money straddles in SP options.

With respect to delta hedging, we argue that only market makers delta hedge as they are faced with large aggregate positions which they take on from trading with many (small) individual investors. Individual investors do not normally hedge their smaller positions since they would often be constrained by transaction costs and financial know-how in hedging. Furthermore, Ni, Pearson, and Poteshman (2005) report that institutional trading in the index options market amounts to a rather small fraction of total volume. Consistent with this finding, Savickas and Wilson (2003) report that approximately 70% of all option trades (equity and index) in 1995 are due to trades between public customers and market makers. We argue that we can follow their example by ignoring the effect of institutional traders in the similar market for futures options. Next, we detail all the explored pinning mechanisms.

The simplest mechanism is the change of delta hedging sold straddle (i.e. short gamma) positions as the underlying future moves. In Figure 2 we can see that just before expiration a sold straddle with the future being above the strike price has a negative delta of almost -1, which is hedged

[3] We define near-the-money as options with moneyness between 0.95 and 1.05. We calculate correlations for all the expirations in the period from January 1990 to December 2009 where we have data for SPX options from Market Data Express.

with almost a long future. As the future falls, the straddle gains delta, thus the hedge needs to lose delta by selling futures in the falling market. The reverse mechanism operates in increasing markets. This effect amplifies the movement of the underlying in the presence of movement in the underlying future which leads to higher volatility; see Pearson, Poteshman, and White (2007). However, this mechanism does not lead to pinning as postulated by Krishnan and Nelken (2001) because the hedging pressure does not revert at the strike price of the straddle; it merely amplifies the movement of the underlying future.

2.2.1 Delta hedging and time decay effect (Avellaneda and Lipkin, 2003)

Avellaneda and Lipkin (2003) argue that the time decay of delta-hedges of long option positions leads to pinning.[4] Alas, given that the market maker typically holds a sold straddle position, their mechanism leads in that case to anti-pinning. As we can see in Figure 3, initially at t_0 the hedge around the strike price of the straddle is zero or almost zero. Now imagine that the future goes slightly above the strike price and the delta of the position is about -0.5.[5] However, as expiration comes very close at time t_1, the delta of the position moves from about -0.5 to almost -1, thus, the hedge involves buying the future as expiration nears and the future is above the strike. As a result, the future is being pushed upwards and away from the strike. A similar mechanism establishes the predicted anti-pinning for the case where the future is below the strike price.

The main hypothesis related to Avellaneda and Lipkin (2003) is (we express all hypotheses in terms of pinning; anti-pinning being then a lessening of pinning):

[4] See also Jeannin, Lori, and Samuel (2007). The implications of Avellaneda and Lipkin (2003) hold also in the equilibrium model with feedback effects of Nayak (2007).

[5] The delta values are not exact but simply meant to be suggestive of possible values and magnitudes.

- **Hypothesis AL-1:** *At-the-money (ATM) open interest decreases pinning.* The larger the ATM open interest on expiration Friday,[6] the larger is the sold straddle position that the market maker needs to hedge. Thus, the higher the open interest, the weaker the pinning, or the stronger the anti-pinning effect.

As we will see later, we use end-of-day options data in this study. As end-of-day open interest on expiration Friday is theoretically zero, we use open interest on the Thursday before expiration Friday. We follow here Ni, Pearson, and Poteshman (2005). However, as options are actively traded during expiration Friday, Thursday open interest does not reflect exactly Friday open interest.

Thus, we complement the main hypothesis with two additional hypotheses related to option trading activity on expiration Friday: option volume and option early exercise.

- **Hypothesis AL-2:** *ATM option volume increases pinning.* ATM option volume on expiration Friday is partly related to the closing of open positions which will expire at day end. Thus, it is reasonable to presume that while some option volume will open new positions, the net effect is to close positions. Assuming reasonably stable proportions, larger volume should then lead to more closures of positions, thus reducing open interest and the hedging need of the market maker. As a result, there should be less anti-pinning and more pinning.
- **Hypothesis AL-3:** *ATM early option exercise increases pinning.* Early exercise of individual investors long positions would lead to reduced short positions of the market maker and thus to reduced hedging needs. This would weaken the anti-pinning and thus strengthen pinning.

[6] Options on S&P 500 futures typically expire on the third Friday of the month. Occasionally, due to holidays, the expiration falls on the Thursday before the third Friday of the month. In our sample, this happens three times. As reported in the robustness section, omitting these days does not change the results.

2.2.2 Reselling of slightly in-the-money options

Anders (1982) argues that individual investors dislike long in-the-money (ITM) option positions at expiration because they expose individual investors to price risk over the weekend. This concern is relevant for the SP options since they settle with physical delivery and the investor is then stuck with the unhedged futures position until after the weekend. Thus, investors sell their ITM positions to the market maker who will then need to adjust the hedge on the typical market maker short straddle position. In Figure 4 we can appreciate what ensues as the future starts at the strike price at some time t_0 before expiration and increases above the strike price. The short straddle changes its delta from about zero to about -0.5 and the hedge requires buying half a future. Now, the call is ITM and has a delta of about 0.75. Next, the investor sells the ITM call to the market maker and the market maker's reduced short straddle position (= one short out-of-the-money (OTM) put) requires 0.75 futures less in the hedge, for a net effect of -0.25 futures after the adjustment at time t_1. Thus, the market maker sells a quarter of the future when the futures price goes above the strike price. The opposite story unfolds below the strike price. Note that these effects are asymmetric as downward pressure from above the strike price is due to calls being ITM, upward pressure from below the strike price is due to ITM puts.

Note that the described effects are applicable only to ATM options. As deeply ITM options are unlikely to finish OTM, investors re-sell them already several days before the expiration. It is only for close to ATM options that investors are uncertain whether their options will expire in- or out-of-the-money. Hence, investors wait until right before expiration and then sell their options if they go in-the-money.

Our hypotheses related to Anders (1982) are twofold:

- **Hypothesis AN-1:** *ATM call volume increases pinning from above the strike price.* As we argued above, volume is related to the closing of positions. Thus, ATM call volume measures investor activities as calls go ITM and will lead to directional pinning, namely, to increased pinning from above the strike price as the future is being pushed downward closer towards the strike price.

- **Hypothesis AN-2:** *ATM put volume increases pinning from below the strike price.* The mechanism is exactly the opposite of hypothesis AN-1.

2.2.3 Early exercise of slightly in-the-money options

The next potential explanation of pinning is due to early exercise of ITM call options and simultaneous selling of the delivered underlying future. This puts downward pressure on the price of the future and as the effect reverses for ITM put options, options exercise can explain pinning.[7]

The mechanism is very similar to Anders (1982) but based on individual investors exercising their American ITM options instead of selling them as in Anders (1982). Again, this is a realistic concern as the SP options are American. Individual investors will then buy the necessary future for delivery (in case of a put) or sell the delivered future (in case of a call) right away in the market. However, the results are just the same in terms of hedging and pinning as in Anders (1982). In detail, we start again at some time t_0 before expiration with the future starting near the strike price and increasing above the strike price. The short straddle position of the market maker changes from a delta of zero to a delta of about -0.5 and the hedge requires a purchase of 0.5 futures. Furthermore, when the investor exercises the ITM call, the market maker needs to buy additional 0.5 futures for the delivery of one future, which is then sold on the market by the investor. The net effect is thus the purchase of one future (0.5 + 0.5 futures) by the market maker and the selling of one future by the investor. However, the market maker still needs to hedge the remaining OTM put leg of the original sold straddle. As the OTM put option has a delta of about 0.25, the market maker needs to

[7] This behavior has been documented in Chiang (2010), albeit with a focus on the underlying stock returns and without referring to pinning.

sell 0.25 of the future. The total market effect is hence 0.25 sold futures (= $0.5 + 0.5 - 1 - 0.25$ futures) at time t_1 and the mechanism creates exactly the same downward pressure as in Anders (1982). Again, the effect is strongest ATM and weakens at a distance from ATM.

The hypotheses related to early option exercise are twofold again:

- **Hypothesis EARLY-1:** *ATM call early option exercise increases pinning from above the strike price.* ATM call option exercise measures investors winding down positions as the calls go ITM and will lead to directional pinning, namely, to increased pinning from above the strike price as the future is being pushed downward closer towards the strike price.
- **Hypothesis EARLY-2:** *ATM put early option exercise increases pinning from below the strike price.* The mechanism is exactly the opposite of the hypothesis EARLY-1.

2.2.4 Manipulation of the underlying

Observationally equivalent to the pinning mechanisms of Anders (1982) and the early exercise explanation is the market manipulation mechanism of Ni, Pearson, and Poteshman (2005). Here, sophisticated market participants with short positions (i.e. the typical market maker in the SP options) could gain from manipulating the future. Namely, pushing the future downward from above the strike price would reduce payments to individual investors with long call option positions while pushing the future upward from below the strike price would reduce payments to individual investors with long put positions. We investigate to what extent pinning can be explained by the hedging mechanisms of Anders (1982) and early exercise mechanism. Only residual pinning should then be attributable to market manipulation and would show up as additional explanatory power of volume of future trading which we use to measure manipulation.

However, as the futures market is very large and liquid, any manipulation should be rather difficult as it would involve large unhedged trades in order to move the future sufficiently for the purpose of manipulation. Such trades would leave the market maker vulnerable to price risk over the weekend which is undesirable for the market maker. Further, the risk of detection of the manipulation will also diminish the interest of the market maker in such activities. Moreover, pinning itself is risky for the market maker (so-called pin risk) and manipulation would increase this risk. Pin risk arises because, due to transaction costs, the option writer (i.e. the market maker) cannot predict with certainty whether the marginally ITM options will be exercised at expiration. Hence, pinning aggravates the risk of ending with a naked position in the future over the weekend. Finally, small movements of the future through the strike price will lead to dramatic adjustments in the hedge (for a vanilla short call the delta of the hedge goes from 0 to unity as the future moves through the strike price from below). As a result, the market maker should be wary to increase pinning through manipulation and needs to carefully balance benefits and costs.

- **Hypothesis NI-1:** *Futures volume is insignificantly related to pinning after accounting for delta-hedging.* Once we account for the delta hedging based explanations of pinning, we do not expect manipulation to play a large role anymore. Thus, adding futures volume as a variable should only contribute insignificantly to explaining pinning.

2.2.5 Volatility and pinning

Pinning may also be related to general conditions in the futures market. In times of high volatility when the futures price crosses several strikes in a single day, we may expect that future volatility obscures pinning effects, a point also made by Avellaneda and Lipkin (2003). In their model the "strength" of the anti-pinning force is inversely related to the volatility of the underlying. The same logic of volatility weakening pinning effects applies also to other explanations of pinning.

- **Hypothesis NI-2:** *Futures volatility decreases pinning.* Future volatility makes delta hedging of the market maker more difficult and is thus negatively related to pinning.

2.3 Methodology

We are interested in testing for pinning in different option classes associated with the S&P 500 and, given that we find such pinning, in explaining which mechanisms drive this pinning. For the purpose of testing for pinning, we employ logistic regressions and additionally use a binomial test based on the uniform distribution of futures prices. The first test is a logistic regression with fixed effects; see Ni, Pearson, and Poteshman (2005). We use 5 days before and after each expiration day.[8]

$$Pinn\_sym_t = \alpha + \beta Dumm_t + \varepsilon_t \tag{68}$$

$Pinn\_sym_t$ is taken to be a zero/one variable which is 1 if the future price at settlement is within \$0.25 below or above the ATM strike price.[9] We always take the ATM strike price to be the strike price closest to the future settlement price.

$$Pinn\_sym_t = \begin{cases} 1, & if \left|Fut_t - K_t^{ATM}\right| \leq 0.25; \\ 0, & otherwise. \end{cases} \tag{69}$$

We define as 1 for expiration days and 0 otherwise. Similar to $Pinn\_sym_t$, we define $Pinn\_above_t$ and $Pinn\_below_t$ to be \$0.25 half-intervals above and below the ATM strike price, respectively, and use them as alternative dependent variables in equation (68):

[8] We use the above simplified notation merely to ease exposition for our logistic regressions. Formally, regression model (1) should be written as: $Pinn\_sym_t = 1/\left(1 + \exp\left[-\left(\alpha + \beta Dumm_t\right)\right]\right)$.

[9] We vary the size of the interval in the robustness section to \$0.125 and \$0.5, respectively.

$$Pinn\_above_t = \begin{cases} 1, & if\ 0 \leq Fut_t - K_t^{ATM} \leq 0.25; \\ 0, & otherwise. \end{cases} \quad (70)$$

$$Pinn\_below_t = \begin{cases} 1, & if\ -0.25 \leq Fut_t - K_t^{ATM} \leq 0; \\ 0, & otherwise. \end{cases} \quad (71)$$

The above logistic regression tests if pinning on expiration days is significantly higher than on non-expiration days. However, we are also interested in testing if pinning on expiration days is significantly higher compared to pinning due to independent draws from a uniform distribution of futures prices. We can then compute the p-value based on the following approximation which can be used as long as the number of observations n exceeds 90, based on Johnson, Kotz, and Kemp (1992, p. 114, equation 3.27):

$$\Pr\left[X \geq x\right] \approx 1 - \Phi\left(\{x - 0.5 - nq\} / \{nq\,(1-q)\}^{0.5}\right) \quad (72)$$

where Φ is the cumulative normal density and where the probability q of symmetric pinning is 10%. In our further work, we simply report this binomial p-value alongside the p-value of the logistic regression.

Once we establish pinning, we explore which mechanisms can explain the pinning. We use logistic regressions where we drop the expiration dummy and focus only on expiration Fridays. In accordance with our hypotheses in form of equation (73), we use additional right-hand-side variables such as option open interest, option volume, option early exercise, and others. We detail our independent variables below.

$$Pinn\_sym_t = \alpha + \beta\ right - hand - side\ variables_t + \varepsilon_t \quad (73)$$

2.4 Data

We now turn to the description of the data sources and presentation of descriptive statistics of all variables.

2.4.1 Data sources

We obtain the whole history of daily data for S&P 500 futures and SP options on S&P 500 futures directly from the Chicago Mercantile Exchange (CME).

The futures data provides daily open, low, high, close, and settlement prices along with the daily open interest and volume for all maturities of futures from their introduction on April 21^{st}, 1982 to December 31^{st}, 2009. Similarly, the SP options data provides daily open, high, low, and close prices along with the daily open interest, volume, and early exercise for all individual options from their introduction on January 28^{th}, 1983 to December 31^{st}, 2009.

To test for pinning in the S&P 500 cash index, we additionally obtain the special A.M. exercise-settlement values (SET) of the S&P 500 from Market Data Express. Quarterly SET values run from June 1991 to December 2009 and serial SET values run from November 1992 to November 2009.

We also employ daily data for SPX options on the S&P 500 index for the period January 2^{nd}, 1990 to December 31^{st}, 2009, which we obtain from Market Data Express. The SPX options data comes along with daily open, high, low, close prices, open interest, and volume for all individual SPX options and the value of the underlying S&P 500 cash index.

Finally, in tests for pinning in the SPDR exchange traded fund on the S&P 500, we employ daily prices of SPDR for the period from January 29^{th}, 1993 to December 31^{st}, 2009, which we obtain from Datastream.

In our main tests, we focus on settlement prices of nearest to maturity futures on serial expiration dates (usually the third Friday of the month) which are available in the data from August 1987 to November 2009.

We use settlement prices as those determine the value of the expiring SP options. Further, since SP options always expire in nearest to maturity futures, we abstract from longer dated futures. We corroborate this by showing below that there is no pinning in the second to maturity futures.

Since SP options trade on a monthly cycle and futures trade on a quarterly cycle (March, June, September, and December), we primarily focus on serial expiration months (all months excluding the quarterly cycle: January, February, April, May, July, August, October, and November). These serial expiration days provide us with a unique laboratory of cases when SP options expire and the underlying future continues to trade for an additional month or two. It is exactly this feature that enables the future to finish in the proximity of the strike price. As opposed to serial expirations, on quarterly expiration days futures and SP options expire simultaneously in the cash value of the S&P 500 index. As the whole basket of S&P 500 stocks is difficult to move, we do not expect to find pinning in quarterly expirations. Again, the results below confirm our conjecture.

Finally, since SP options were first traded on a quarterly cycle, just like futures, and serial expirations for SP options were introduced only in June 1987 (and there is no data for SP options expiring in July 1987), we restrict the analysis to the period from August 1987 to November 2009. Also, as it is standard in derivatives research, we regularly eliminate two crash months, October 1987 and October 2008. However, adding them back to the analysis has virtually no effect on the results. Appendix A elaborates further on the main characteristics of the S&P 500 derivatives and the changes in the settlement procedures of these derivatives. Appendix B details the raw data processing.

2.4.2 Variable definition

Having defined our dependent variables already above, we now turn to defining our independent variables. First, ATM open interest is measured on the Thursday before expiration with respect to the ATM strike price on the expiration Friday. We add 10 to each variable and take logarithms where the addition of 10 serves to avoid taking logarithms of zero. The transformed variable is labeled OI.

Second, ATM volume is measured with respect to the ATM strike price on the serial expiration Friday. Again, it is composed of ATM put and ATM call volume, respectively. We add 10 to each variable and take logarithms. The transformed variables are labeled VOL, $Call\_VOL$, and Put_VOL.

Third, ATM early option exercise is measured with respect to the ATM strike price on the serial expiration Friday. And again, it is composed of ATM put and ATM call early option exercise, respectively. We add 10 to each variable and take logarithms. The transformed variables are labeled OE, $Call\_OE$, and Put_OE.

Fourth, Fut_vol measures the logarithm of 10 plus the volume of futures contracts traded on the serial expiration Friday.

Last, Fut_sigma measures the volatility of futures one day before the expiration Friday. We use the Thursday before the expiration Friday to avoid endogeneity problems arising from the fact that pinning itself could lower the volatility of the future on expiration Friday. We approximate volatility by the Parkinson (1980) scaled daily realized range (see Martens and van Dijk, 2007):

$$Pinn\_sigma = \frac{\left(\log\left(Fut\_high\right) - \log\left(Fut\_low\right)\right)^2}{4\log\left(2\right)} \tag{74}$$

where Fut_high is the intra-daily futures high price and Fut_low is the intra-daily futures low price.

Some independent variables have missing values. Since there is no generally accepted treatment for missing variables, we replace missing observations by the sample mean of the untransformed missing variable. If needed, we then transform all variables. Results are robust to alternative treatments of missing observations and are not affected if we use zero instead, nor if we eliminate the missing observations altogether. We provide details in the robustness section below.

More problematic is the early option exercise variable where the first half of the variable is missing. We are uncomfortable with imputing the missing values through some statistical procedure since too many observations are missing. Instead, we argue that we can start with the long sample of the standard model, which performs much like the short sample standard model where the option exercise is observable (October 1998 until November 2009). Then, we can analyze the effect of options early exercise on the short sample. We detail the implementation of this approach in the results below.

2.4.3 Descriptive statistics

We first look at the time pattern of SP option trading activity on serial expiration dates as depicted in Figure 5. The data is used as reported in the original data (without taking logarithms and without the addition of 10) with missing values replaced by zeros. Panel A depicts ATM open interest, Panel B depicts ATM volume, and Panel C depicts ATM option exercise.

Figure 5 demonstrates that the SP option activity on the serial expiration dates rose over the years. Open interest and volume steadily increased from 1987 to approximately 1997, then they decreased somewhat and ramped again from 2004 to 2009. Hence, if pinning is related to option activity, we should observe more pinning in the more recent period which is confirmed in the results.

Table 1 reports the summary statistics of the transformed data for the full period August 1987 to November 2009. Summary statistics for early options exercise is based on the short period from October 1998 to November 2009.

Table 1 already reveals some interesting phenomena. First, from August 1987 to November 2009, 13.56% of futures prices settle within 0.25$ of the strike price on the serial expiration Fridays; on average much higher than the 10% expected under a uniform distribution. The result is even stronger if we focus on a more recent period which is characterized by increased options trading activity. Indeed, from October 1998 to November 2009, as much as 19.10% of futures prices settle within 0.25$ of the strike price.

Second, we notice that pinning from above the strike price is especially pronounced. For the full sample, pinning from above the strike price amounts to 8.47% and pinning from below the strike price amounts to 7.34%. In the short sample, the values are 12.36% for pinning from above and 10.11% for pinning from below.[10]

Table 2 reports the unconditional correlation structure between the main variables. Note that among the set of considered variables, option volume, and early option exercise exhibit the highest unconditional correlations with symmetric pinning. The correlations are 0.14 and 0.11, respectively. Although the correlation between early option exercise and symmetric pinning is not directly comparable to the correlation between open interest and symmetric pinning because the latter is based on the short period, these results already suggest that option volume as well as option exercise play an important role for the documented pinning. Further, note that open interest is positively related to symmetric pinning. Although the correlation is very weak (0.01), the positive sign is somewhat surprising because, according to Avellaneda and Lipkin's (2003) anti-pinning argument, open interest should be negatively related to pinning. However, this unconditional correlation could be positive simply because open interest is highly correlated with op-

[10] Note that pinning from above and pinning from below do not sum up to symmetric pinning as both include observations with zero deviation from the strike price.

tion volume (pairwise correlation of 0.81). Indeed, we show in Section 2.5 that, conditional on option volume, open interest is always negatively related to symmetric pinning. Finally, as expected, futures volatility is negatively related to pinning. Slightly surprising is the negative correlation between futures volume and pinning.

Table 3 complements Table 2 by providing unconditional correlations for the subvariables; symmetric pinning is broken down into pinning from above the strike price and pinning from below the strike price. Option volume and early option exercise are reported separately for calls and puts. Although the correlations are not entirely conclusive, Table 3 demonstrates that volume and early option exercise for calls are specially related to pinning from above the strike price. For puts these quantities exhibit stronger correlations with pinning from below the strike price.

In unreported results, we find that recomputing Tables 1, 2, and 3 for the short sample changes the point estimates but generally confirms the above descriptive statistics.

2.5 Results

In our results section, we normally use the full sample from August 1987 to November 2009. As detailed in Section III, there are some missing values and we replace those with the sample mean of the respective variable. The variable early option exercise (OE) misses the first half of its values and we therefore analyze the effect of early option exercise only on the short sample from October 1998 to November 2009. The dependent variable, $Pinn\_sym$, will be labeled $Pinn\_sym^L$ in the long sample and $Pinn\_sym^S$ in the short sample.

2.5.1 Pinning does exist in the near maturity future due to serial SP options

Using the long and short samples respectively, we analyze in equations (75) and (76) if expiration Friday pinning (within $0.25 below and above the ATM strike price) is stronger than pinning on the 5 days before and after expiration Friday. The p-values in parentheses are based on the logistic regression. The p-values in brackets are based on the binomial distribution of comparing Friday pinning against a uniform distribution without pinning.

$$\begin{aligned} Pinn\_sym_t^L &= -2.32 + 0.46 Dumm_t + \varepsilon_t \\ p - value \quad & (0.00) \qquad (0.05)[0.07] \end{aligned} \tag{75}$$

$$\begin{aligned} Pinn\_sym_t^S &= -2.40 + 0.96 Dumm_t + \varepsilon_t \\ p - value \quad & (0.00) \qquad (0.00)[0.00] \end{aligned} \tag{76}$$

For the long sample containing 177 expirations, we find supporting evidence for symmetric pinning, which affects the near maturity futures with p-values of just slightly below 0.05 when compared to other days and of 0.07 when compared to the uniform distribution. Note however that in the period from June 1987 to December 1993 there exist liquid SPX options that also expire P.M. and could therefore disrupt pinning in the futures. Excluding this period, we find that both p-values decrease to 0.02. As reported in equation (76), the evidence for pinning is even stronger, with p-values of 0.00 and 0.00 respectively, in the short sample from October 1998 to November 2009 despite the reduced sample size (89 expirations), indicating that pinning increased substantially during the last decade.

We next investigate asymmetric pinning in the long sample (177 expirations). In equations (77) and (78) we analyze pinning from above and below, respectively.

$$\begin{aligned} Pinn\_above_t^L &= -2.92 + 0.54 Dumm_t + \varepsilon_t \\ p-value &\quad (0.00) \qquad (0.07)[0.03] \end{aligned} \tag{77}$$

$$\begin{aligned} Pinn\_below_t^L &= -2.96 + 0.43 Dumm_t + \varepsilon_t \\ p-value &\quad (0.00) \qquad (0.17)[0.10] \end{aligned} \tag{78}$$

While we find supporting evidence for pinning from above (p-values of 0.07 against other days and 0.03 against the uniform distribution[11]), the evidence is weaker for pinning from below with p-values of 0.17 and 0.10. These findings are consistent with the fact that arbitraging away pinning from below is much harder than pinning from above. When pinning from below occurs, then the arbitrage trade involves selling the overpriced future and buying the index basket which is expensive but feasible. However, arbitraging pinning from above involves buying the underpriced future and (short) selling the index basket which is much harder to do. Thus, pinning from above is likely to exist more often than pinning from below as the latter is easier to arbitrage away.

Same as for symmetric pinning, asymmetric pinning is stronger in the short sample (89 expirations) with all p-values being lower. For pinning from above the p-values are now 0.01 and 0.00, respectively. For pinning from below, p-

[11] Here we use the probability of 0.05 that the future lies within \$0.25 above the ATM strike and the same for the interval below the ATM strike as opposed to a probability of 0.1 for the symmetric interval.

values are 0.04 and 0.02, and thus provide stronger evidence for asymmetric pinning during the last decade.

2.5.2 Mechanisms of pinning in the near maturity future

Since we established pinning in the near maturity futures, we now embark on analyzing the mechanisms which drive this pinning. Our first set of hypotheses is based on the time-decay in the delta hedge as modeled by Avellaneda and Lipkin (2003).

- **Hypothesis AL-1:** *ATM open interest decreases pinning* and
- **Hypothesis AL-2:** *ATM option volume increases pinning.*

$$\begin{aligned} Pinn\_sym_t^L &= -4.68 - 1.28OI_t + 1.65VOL_t + \varepsilon_t \\ p - value &\quad (0.02) \quad (0.00) \qquad (0.00) \end{aligned} \tag{79}$$

In the long sample (177 expirations), we find according to equation (79) that both variables are strongly significant and have indeed the expected signs. ATM open interest reduces pinning and ATM volume increases pinning.[12]

In testing the remaining hypothesis of Avellaneda and Lipkin, we would like to investigate if it is true that

- **Hypothesis AL-3:** *ATM early option exercise increases pinning.*

The variable early option exercise (OE), which we would like to use here is problematic as we do not trust the first part of the sample where there

[12] There is little autocorrelation in the residuals as the AR(1) of the residuals is statistically insignificant at 0.05.

are mostly zero values recorded and some few extremely low values. However, starting October 1998, the values are much more realistic. Rather than imputing the first half of the sample, we suggest the following method. We first reestablish the above results for hypotheses AL-1 and AL-2 on the short sample from October 1998 to November 2009 (89 expirations). This is demonstrated in equation (80) and while the point estimates vary somewhat when compared to equation (79), the signs are stable and all coefficients are significant. Then we use the short sample while including early option exercise (OE) in our model and report the results in equation (81). While ATM early option exercise exhibits the correct sign, it is insignificant. We are afraid that this could be due to the reduced power as we are only using 89 observations in the short sample as opposed to 177 observations in the long sample.

$$\begin{aligned} Pinn\_sym_t^S &= -4.12 - 1.06OI_t + 1.42VOL_t + \varepsilon_t \\ p - value & \quad (0.06) \quad (0.04) \qquad (0.02) \end{aligned} \tag{80}$$

$$\begin{aligned} Pinn\_sym_t^S &= -4.16 - 1.09OI_t + 1.32VOL_t + 0.15OE_t + \varepsilon_t \\ p - value & \quad (0.05) \quad (0.03) \qquad (0.04) \qquad (0.62) \end{aligned} \tag{81}$$

Next, we turn our attention to the asymmetric pinning effects of Anders (1982). We conceptually revert to the setting of equation (79) where we investigated the effects of open interest and volume in the long sample (177 expirations). However, we now separate the volume effects into ATM call and put volume. The hypotheses are:

- **Hypothesis AN-1:** *ATM call volume increases pinning from above the strike price* and

- **Hypothesis AN-2:** *ATM put volume increases pinning from below the strike price.*

We use ATM call volume and ATM put volume in addition to ATM open interest in order to explain pinning from above the strike price in equation (82) and pinning from below the strike price in equation (83).

$$\begin{array}{lccccc} Pinn\_above_t^L = & -3.00 & -\ 1.28 OI_t & +\ 0.97 Call\_VOL_t & +\ 0.55 Put\_VOL_t & + \varepsilon_t \\ p - value & (0.15) & (0.01) & (0.02) & (0.13) & (82) \end{array}$$

$$\begin{array}{lccccc} Pinn\_below_t^L = & -0.98 & -\ 0.82 OI_t & +\ 0.27 Call\_VOL_t & +\ 0.43 Put\_VOL_t & + \varepsilon_t \\ p - value & (0.62) & (0.07) & (0.47) & (0.23) & (83) \end{array}$$

Pinning from above the strike price is supported by equation (82) since ATM call volume increases the propensity of pinning from above the strike price and ATM put volume is insignificant, as expected. The evidence in favor of pinning from below the strike price is somewhat weaker. ATM call volume in equation (83) is insignificant, as expected, but ATM put volume has the right sign but is insignificant with a p-value of 0.23. This asymmetry is not entirely surprising as we argued above in the light of arbitrage trades: since pinning from below involves buying the index, it is easier to arbitrage away than pinning from above where the arbitrage trade involves (short) selling the index. Using the short sample, the signs remain as in equations (82) and (83) while the p-values decrease.

Now we investigate the closely related mechanism of early option exercise which leads to following two hypotheses:

- **Hypothesis EARLY-1:** *ATM call early option exercise increases pinning from above the strike price* and

- **Hypothesis EARLY-2:** *ATM put early option exercise increases pinning from below the strike price.*

We are using again the standard model of equation (79) applied to the short sample (89 expirations) due to the missing data in the early option exercise variable. Equation (84) shows the results for pinning from above the strike price and equation (85) for pinning from below the strike price. Pinning from above the strike price indeed seems to be influenced by call early option exercise (although the p-value is insignificant at 0.15) and put early option exercise does not have an impact on pinning from above, as expected. Pinning from below the strike price is also supported with a significant contribution on the put early option exercise and an insignificant coefficient of the call early option exercise. All early option exercise coefficients are positive.

$$\begin{aligned} Pinn\_above_t^S &= -2.15 - 0.58OI_t + 0.32VOL_t \\ p-value \quad & \quad\ (0.36) \quad (0.27) \qquad (0.64) \end{aligned}$$

$$\begin{aligned} &+0.42Call\_OE_t + 0.00Put\_OE_t + \varepsilon_t \\ &\quad (0.15) \qquad\qquad (0.99) \end{aligned} \tag{84}$$

$$\begin{aligned} Pinn\_below_t^S &= -1.31 - 1.40OI_t + 0.43VOL_t \\ p-value \quad & \quad\ (0.58) \quad (0.03) \qquad (0.59) \end{aligned}$$

$$+0.44Call\_OE_t + 0.82Put\_OE_t + \varepsilon_t$$
$$(0.13) \qquad (0.02) \qquad (85)$$

We next turn to potential market manipulation and investigate:

- **Hypothesis NI-1:** *Futures volume is insignificantly related to pinning after accounting for delta-hedging.*

We use as a point of departure the model in equation (81) which includes ATM open interest, ATM volume, and early option exercise. As we use early option exercise, we can only use the short sample. We then add the variable future volume and report the result in equation (86). The result repeats much of equation (81) in that ATM open interest and volume are significant and while all variables have the right sign, early option exercise is insignificant. The addition of future volume leads, as expected, to insignificant coefficients. We conclude that market manipulation does not seem to explain pinning.

$$Pinn\_sym_t^S = -2.56 - 1.11OI_t + 1.33VOL_t + 0.14OE_t - 0.14Fut\_vol_t + \varepsilon_t$$
$$p - value \quad (0.74) \quad (0.04) \quad (0.04) \quad (0.65) \quad (0.83) \quad (86)$$

Finally, we analyze the influence of volatility on pinning and test

- **Hypothesis NI-2:** *Futures volatility decreases pinning.*

As in the case of Hypothesis NI-1, we use as a point of departure the short sample and model (81), which includes ATM open interest, ATM volume, and early option exercise. We then add future volatility and report the results in equation (87). In line with our hypothesis, futures volatility is negatively related to pinning but the p-value is insignificant at 0.12. Furthermore, adding futures volatility to model (81) slightly increases the significance of other variables, such as open interest and option volume. This suggests that volatility indeed weakens the pinning forces of the market maker's delta hedging activity.

$$
\begin{aligned}
Pinn\_sym_t^S &= -3.36 - 1.35 OI_t + 1.62 VOL_t \\
p-value \quad & \quad (0.15) \quad (0.02) \quad\quad (0.02)
\end{aligned}
$$

$$
\begin{aligned}
&+0.11 OE_t - 6.54 Fut\_sigma_t + \varepsilon_t \\
&\quad (0.02) \quad\quad\quad (0.12)
\end{aligned} \tag{87}
$$

In summary, regarding the serial SP options we find evidence that pinning is explained by the interplay of time-decay of the delta hedge (anti-pinning due to Avellaneda and Lipkin 2003) and pinning due to the hedging effects of Anders (1982). Pinning due to the hedging effect caused by early option exercise is insignificant, possibly due to the shorter sample over which the data is available. Market manipulation does not seem to contribute to the explanation. Volatility of the underlying seems to have little impact on the pinning effects of delta hedging.

2.5.3 No pinning in second to maturity futures

SP options expire in the nearest (first) to maturity futures. Hence, if pinning is related to option expiration, it should be present in the first to maturity futures, as documented in Subsection 2.5.1, and it should be absent for longer maturity futures. To investigate whether there is any evidence for pinning in longer maturity futures, we next measure symmetric pinning in the second to maturity futures on serial expiration dates.[13] We repeat the model of equation (68):

$$
\begin{aligned}
Pinn\_sym_t^L &= -1.95 - 0.36 Dumm_t + \varepsilon_t \\
p - value \quad & \quad (0.00) \qquad (0.19)[0.71]
\end{aligned}
\tag{88}
$$

In the long sample, based on the insignificant p-values we conclude that there is no evidence for pinning in the serial expiration dates for second to maturity futures. This finding continues to hold in the short sample.

2.5.4 No pinning in quarterly expirations

We next investigate whether there is any evidence for pinning in the future settlement price on quarterly expiration days. As opposed to the above analyzed serial expirations, on quarterly expiration days futures and SP options expire simultaneously in the value of the underlying S&P 500 index. Thus, pinning should be much harder in the quarterly expirations as the future needs to finish in the value of the underlying and any pinning in the settlement prices of futures would imply that there is pinning in the S&P 500 index. As the whole basket of S&P 500 stocks is difficult to move, we do not expect to find pinning in quarterly expirations for futures.

Also, on quarterly expirations, SP options expire into the cash-settled value of the underlying whereas the above serial SP options expire in the physically delivered future. Therefore, the fear of ending up with a naked position in the underlying does not apply to SP quarterly options and Anders (1982) story of reselling options on the expiration date does not work. For the same reason, the early exercise story does not apply. The only remaining explanation for pinning (ignoring manipulation) is the anti-pinning story of

[13] We do not use third to maturity futures since their strike price intervals are either $10 or $25 instead of always $5 for the first and second to maturity futures.

Avellaneda and Lipkin (2003). Thus, if there is any pinning in the futures settlement prices on quarterly expirations, it should be anti-pinning.

Before measuring pinning on quarterly expirations, one final remark is in order. Note that we have the data for quarterly expirations for the whole period from March 1983 to December 2009. But the settlement procedures for quarterly expirations underwent two important changes (see Appendix A for details). To take into account these changes, we measure pinning on quarterly expirations using the P.M. settlement futures prices on the third Thursday for the period from March 1983 to March 1984, the P.M. settlement futures price on the third Friday from June 1984 to March 1987, and the A.M. settlement price (which is determined by the first opening price of all the constituents of the S&P 500 index) for the period from June 1987 to December 2009. Altogether we have 108 quarterly expirations.

We cannot use our usual logistic regressions of equations (75) and (76) as we do not have A.M. settlement prices for the days surrounding the expiration Friday. However, we can test against the uniform distribution. With a p-value of 0.77 in the long sample, we do not find evidence in favor of pinning in the quarterly expirations. The same picture emerges in the short sample with a p-value of 0.69 which due to a small number of observations (46), we calculate directly from the binomial distribution.

To further complicate matters, the quarterly SPX options on the S&P 500 index also expire at the same time into the same value. However, from June 1987 to December 1993 two types of SPX options co-existed, one that expired P.M. and another that expired A.M. Since this can potentially disrupt pinning in the quarterly expirations, we rerun the long sample without this period (the short sample remains unaffected by this period). The p-value goes up to 0.91. Realizing that 1 minus the p-value is the one-sided test for documenting anti-pinning, this, if anything, is slight evidence for anti-pinning in line with Avellaneda and Lipkin (2003). However, extending our research into this direction and running the mechanism regression of equation (79), we only obtain insignificant coefficients on open interest and option volume.

2.5.5 No pinning in the S&P 500 index due to serial SPX options

If pinning is related to option expirations, we could potentially also observe pinning in the S&P 500 index itself as there are very liquid SPX options written on the index. However, as we argued in Section 2.5.4, it is hard to imagine that market makers' hedging needs would be strong enough to move the whole basket of 500 stocks.

Also, the market maker does not necessarily hedge SPX options by trading the underlying basket of S&P 500 stocks. As an alternative to trading the basket, the market maker could hedge SPX options by trading the SPDR (an exchange traded fund replicating the S&P 500 index), by trading S&P 500 futures, or by trading options on the SPDR (SPY options). It is beyond the scope of the current paper to analyze the cross effects between these markets, but it is on our research agenda. Any such hedging would weaken potential pinning in the index itself.

Furthermore, from April 1986 to May 1987, all SPX options expired in the P.M. value of the underlying S&P 500 index. Some SPX options continued to do so until December 1993. However, starting in June 1987 some SPX options settle in the so-called A.M. exercise-settlement value which became the standard settlement for all SPX options starting January 1994. Since the main reason behind the introduction of the so-called A.M. settlement was to prevent manipulation, the likelihood of detecting pinning should be even smaller.

Last, contrary to serial SP options, SPX options are cash settled (just like quarterly SP options). Therefore, the fear of ending with a naked position in the underlying does not apply to SPX options and Anders (1982) story of reselling options on the expiration date does not work. Further, the early exercise story is ruled out as SPX options are European. The only remaining explanation for pinning (ignoring manipulation) is the negative pinning story of Avellaneda and Lipkin (2003). Thus, if there is any pinning in the S&P 500 itself, it should be anti-pinning.

We test for pinning in the S&P 500 index on serial expirations in the long sample (we exclude quarterly observations as they are already analyzed above). The European serial SPX options exist since April 1986, but the settlement procedures for SPX options underwent some changes (see Appendix A for details). In line with these changes, we measure pinning using P.M. values of S&P 500 from April 1986 until October 1992 and A.M. exercise-settlement values of the S&P 500 from November 1992 to November 2009. As usual, we eliminate October 1987 and October 2008. We are left with 188 expirations.

We cannot use our usual logistic regressions of equations (75) and (76) as we do not have A.M. settlement prices for the days surrounding the expiration Friday. However, we can test against the uniform distribution. With a p-value of 0.62 in the long sample, we do not find evidence in favor of pinning. The same picture emerges in the short sample with a p-value of 0.80.

Again, we are concerned about the period from July 1987 until November 1993 where SPX options existed expiring both A.M. and P.M. Eliminating this period from the sample yields a p-value of 0.73. We conclude that there is no evidence of pinning due to the SPX options in the index.

2.5.6 No pinning in the SPDR due to SPY options

The last market for which we investigate pinning is the SPDR exchange traded fund on the S&P 500 index. American SPY options on the SPDR exist since January 2005 (with data available until December 2009), trade on a monthly cycle, have physical delivery, and expire on the third Friday of the expiration month. Thus, in principle all three theoretical mechanisms could lead to pinning in this market. However, shorting the SPDR is fraught with the typical difficulties of shorting any equity security. It is thus easier than shorting the basket of 500 securities for the index but harder than selling the future. We repeat the model of equation (76) where the p-value for the uniform distribution is calculated from the binomial distribution directly as the number of observations is only 59:

$$\begin{aligned} Pinn\_sym_t^S &= -2.20 + 0.19 Dumm_t + \varepsilon_t \\ p - value &\quad (0.00) \qquad (0.65)[0.38] \end{aligned} \tag{89}$$

Based on the insignificant p-values we conclude that there is no evidence for pinning in the SPDR. Also, the short time series makes statistical inference difficult.

We also checked if there is pinning in the future on the SPDR which would be an alternative hedging instrument to the SPDR itself. However, we could not find any pinning in the future, either.

2.6 Robustness

We show that our results are robust to a number of methodological changes. Changes to the methodology such as the treatment of missing variables, the treatment of holidays occurring on the third Friday of a month, the inclusion of two crash months, or the elimination of two outlying observations do not affect the results at all with details relegated to Appendix C. Our biggest concern is the correct choice of the pinning interval which we so far set to \$0.25 above and below the ATM strike price. Here, the significance of the results can be affected.

Our definition of pinning as cases when the futures settlement price is within \$0.25 of the nearest strike price is somewhat arbitrary, even though Ni, Pedersen, and Poteshman (2005) also use this value as well as \$0.125. Theory does not provide a clear suggestion for the size of the pinning interval. Choosing the interval too small results in very few instances of pinning and the associated test statistics will be very noisy. Choosing the interval too large and beyond the region where hedging pressure is influencing futures prices will again lead to insignificant results.

We recall the standard model in equations (75) and (76) where we test for pinning on expiration Fridays compared to other days and compared to the uniform distribution. The p-values for the interval of \$0.25 below or above the strike price were 0.05 and 0.07, respectively, in the long sample (177 expirations), and even stronger at 0.00 and 0.00 in the short sample (89 expirations).

We now repeat the regressions with intervals restricted to \$0.125 below or above the strike price:

$$\begin{aligned} Pinn\_sym_t^L &= -3.05 + 0.24 Dumm_t + \varepsilon_t \\ p-value \quad & \quad (0.00) \qquad (0.50)[0.41] \end{aligned} \tag{90}$$

$$\begin{aligned} Pinn\_sym_t^S &= -3.23 + 0.77 Dumm_t + \varepsilon_t \\ p-value \quad & \quad (0.00) \qquad (0.08)[0.16] \end{aligned} \tag{91}$$

As expected, all the p-values increase. In the long sample both p-values are insignificant, 0.50 and 0.41. In the short sample, the p-value based on the logistic regression is significant at the 10% level, but the p-value based on the binomial distribution is insignificant. The sign of the dummy variable is always positive.

We finally repeat the regressions with intervals restricted to \$0.5 below or above the strike price:

$$Pinn\_sym_t^L = -1.28 + 0.26 Dumm_t + \varepsilon_t$$
$$p - value \quad (0.00) \quad (0.14)[0.02] \tag{92}$$

$$Pinn\_sym_t^S = -1.26 + 0.54 Dumm_t + \varepsilon_t$$
$$p - value \quad (0.00) \quad (0.03)[0.00] \tag{93}$$

The p-values increase somewhat compared to the interval based on \$0.25. For the long sample, the p-values are 0.14 and 0.02. For the short sample, the p-values are 0.03 and 0.00. Thus, only the logistic regression is insignificant at the 10% level in the long sample. All coefficients are positive.

We interpret these findings to imply that the interval over which pinning effects occur is indeed restricted and is approximately \$0.25 below and above the ATM strike price. This relatively small interval further suggests that ITM or OTM variables from outside that interval will not lead to pinning. For this reason, in the explanations for pinning, we focused exclusively on trading activity of ATM options.

With regard to the mechanisms explaining pinning, we recomputed all regression models in equations (79) through (87) while varying the size of the interval. Results in the long sample (177 expirations) are generally robust to the choice of the size of the interval, but results weaken somewhat in the short sample (89 expirations). Also, the anti-pinning effect of Avellaneda and Lipkin (2003) seems to be somewhat stronger for smaller intervals.

2.7 Conclusions

We investigate SP option induced pinning in the market for futures on the S&P 500 index. Pinning describes the tendency of the underlying future to be attracted to strike prices on expiration Friday of the option. Such behavior is surprising in light of our typical understanding of finance which suggests that any closing price of the underlying is reached with equal probability.

Ni, Pearson, and Poteshman (2005) documented such behavior for stock options and practitioners believe strongly that stock pinning exists even though the statistical verification is at times replaced by verbal assertion, e.g. Augen (2009, pp. 26). Here, we document pinning in the much larger and more liquid futures market on the S&P 500 index. We show that S&P 500 futures finish in the proximity of the strike price more often on days when SP options on S&P 500 futures expire and the underlying future continues to trade than on other days. Interestingly, there is no pinning in the S&P 500 index itself nor in the SPDR exchange traded fund on the S&P 500 index as both underlying securities (the basket of 500 securities and the SPDR, respectively) are harder to trade and short than the future on the S&P 500 index.

In analyzing the economic mechanisms which drive index futures pinning, we find that they differ considerably from the mechanisms driving stock pinning. Concerning stock pinning, Ni, Pearson, and Poteshman (2005) suggest that the effect is largely driven by the time-decay of the delta hedge of market makers who are typically long the stock options, see Avellaneda and Lipkin (2003) for the model. Also, Ni, Pearson, and Poteshman (2005) argue that manipulation plays a role. For index futures pinning, neither of these two effects is wholly convincing. For one, Garleanu, Pedersen, and Poteshman (2009) report that the market maker is typically short index options as opposed to long stock options which suggests anti-pinning and not pinning in the closely related market for SP options on the index futures. Second, manipulation seems much harder in the index futures and is thus less likely to serve as an explanation.

We resolve the puzzle by introducing two additional effects which lead to pinning, namely hedging pressure resulting from individual investors selling their in-the-money options (Anders 1982) and a related mechanism of individual investors early exercising their in-the-money options. We document that the time-decay of the delta hedge does indeed lead to anti-pinning but is overcompensated by the two additional mechanisms. We do not find evidence of manipulation.

We document extensive movements of 1.6 billion dollars worth of futures value per expiration day that is temporarily being moved closer to the strike nearest the settlement price of the future. This number actually increases to 7.0 billion dollars in the more recent period from October 1998 until November 2009. These large shifts of notional value in the futures market might also be of concern to the regulators and the exchanges as they constitute quite some fraction of the typical expiration day trading volume of approximately 14 billion dollars. The magnitude of the shift per future of more than 50 cents is large in relation to the typical bid/ask spread of only 10 cents.

We were curious to find out if profitable trading opportunities arise. For that purpose, we calculate on each expiration day the futures implied interest rate using the cost-of-carry formula for the price of the future and the realized daily dividends discounted at the one-month LIBOR rate. If this implied interest rate is lower than the median futures implied interest rate over the past 30 days minus one-half the standard deviation of the implied rates, we consider the future to be cheap. We thus buy one future and hedge it with one short second maturity future adjusted by the ratio of futures values. For implied rates above the median implied rate plus one-half the standard deviation, we reverse the trade accordingly. We enter all trades on Friday 3:15 P.M. at closing prices and close them on Monday 8:30 A.M. where we skip 5 days when either expiration falls on a Thursday or the futures opening price is missing. On non-pinning expirations (56 observations), these trade generate average losses of 13 cents. On expiration days with pinning (8 observations), these trades generate average profits of \$1.66. A two-sided t-test for the difference in means is significant at the 10% level. Removing the very volatile year 2009 decreases the p-value to 0.01. The profit of \$1.66 can be related to an average futures level of \$1037 on expirations with pinning. The associated annualized return over the weekend is an astonishing 24%. Given the large trading volume in the S&P 500 futures of some \$14bn on an average expiration day, we thus document sizeable distortions in futures prices due to hedging of market makers.

An exciting field of study beyond the scope of the current paper is the interaction across markets. So could pinning in the future also be driven by the SP options on the underlying S&P 500 index while pinning from the options on the index future should be much less likely to lead to pinning in the (hard to move) index itself.

2.8 APPENDIX

Appendix A: The main characteristics of S&P 500 derivatives

S&P 500 futures and SP options on S&P 500 futures. Futures on S&P 500 and options on S&P 500 futures are traded on the Chicago Merchandise Exchange (CME). Futures were introduced on April 21^{st}, 1982 and SP options were introduced approximately one year later, on January 28^{th}, 1983. SP options are American. First and second closest to maturity options have strike price intervals of $5 and options for deferred months trade with strike price intervals of either $10 or $25.

When futures and SP options were first introduced, they initially expired in the P.M. cash value of the S&P 500 index on the third Thursday in a quarterly cycle (March, June, September, and December). This settlement procedure however underwent three important changes. In June 1984, CME decided to shift expiration dates from the third Thursday to the third Friday of the month.

In June 1987, two additional changes were introduced. First, quarterly futures and SP options no longer expired in the P.M. value of the index, but in the special opening value of the index on the third Friday of the month, the so called A.M. expiration. The special opening value of the index is determined by the first opening prices of all the constituents of the index. It is also called special opening quotation (SOQ) or exercise-settlement value of the index (SET). Second, CME introduced serial SP options that expire in the closest to maturity futures on the third Friday of the serial months (January, February, April, May, July, August, October, and November).

This last introduction of the serial options is of crucial importance for our study. We state three main differences between quarterly and serial SP options:

- First, while quarterly SP options expire simultaneously with the underlying future, serial SP options expire while the underlying future continues to trade for an additional month or two.
- Second, while quarterly SP options expire in the A.M. value of the S&P 500 index (like futures), serial SP options expire in the P.M. value of the underlying future.
- Third, while quarterly SP options are cash settled, serial SP options lead to physical delivery of the underlying future.

SPX options on the S&P 500 index. European SPX options on S&P 500 index are traded on the Chicago Board of Options Exchange (CBOE) since

April 2^{nd}, 1986. All SPX options are cash-settled and trade on a monthly cycle (serial expirations plus quarterly expirations). Nearest to maturity options have strike price intervals of $5 and options for deferred months have strike price intervals of $25.

Initially, SPX options expired in the P.M. value of the S&P 500 on the third Friday in a month.[14] With the introduction of the A.M. settlement for futures and quarterly SP options by the CME in June 1987, CBOE decided to introduce another set of options that also expire A.M. For a while both sets of options coexisted, until in June 1992, CBOE decided that all SPX options should expire A.M. Ever since, all SPX options (serial expirations and quarterly expirations) expire in the special A.M. opening value. Thus, on quarterly expirations, SPX options, SP options, and futures expire in the same special opening value of the S&P 500 index.

SPY options on the SPDR exchange Traded Funds on the S&P 500 Index. SPY options on SPDR exist since January 2005. Like SPX options, SPY options trade on a monthly cycle and expire on the third Friday of the expiration month. However, unlike SPX options, SPY options are American and are settled by delivery of the underlying.

The above information is also summarized in Table A.1.

[14] Technically, it is the Saturday following the third Friday of the month. However, since the settlement value is being determined on Friday already, we will keep on referring to all expiration dates as (third) Fridays.

Appendix B: Treatment (filtering) of the main data

We obtain the data from two sources: the Chicago Mercantile Exchange (CME) and Market Data Express, the official provider of Chicago Board of Options Exchange (CBOE) data. From CME, we obtain the whole history of daily data for S&P 500 futures and SP options on S&P 500 futures. From Market Data Express, we obtain daily data for SPX options on the S&P 500. From Market Data Express, we also obtain a separate file with exercise-settlement values (SET) for SPX options.

All the data spans until December 2009. The futures data starts in April 1982, SP options data starts in January 1983, and SPX options data starts in January 1990. The A.M. exercise settlement values (SET) start either in June 1991 (quarterly expirations) or November 1992 (serial expirations).

In all the datasets, we first filter out observations with missing values for any of the key variables. Further, we eliminate duplicate entries (that is, we keep one of the duplicate entries). Below, we describe specific adjustments to each dataset.

S&P 500 futures. CME does not provide exact expiration dates. The data only contains the expiration year and the expiration month. Therefore, we manually complement the data with the exact expiration dates (usually the third Friday in the quarterly cycle) and eliminate futures with negative time to maturity.

SP options on S&P 500 futures. Similarly to the futures data, the SP options data only contains the expiration year and the expiration month. Therefore, we manually complement the data with the exact expiration dates (usually the third Friday of the month) and eliminate options with negative time to maturity.

SPX options on the S&P 500 index. Market Data Express is a comprehensive source for SPX options, covering not only standard SPX options, but also LEAPS (long dated options), quarterlies, weeklies, and mini options.

First, we eliminate all the options with last bid equal to 998 and last ask equal to 999 as those values stand for erratum in the data. Further, we eliminate LEAPS, non-index options, weeklies, quarterlies, and mini options. Finally, as all the expiration dates in the Market Data Express are set to Saturday following the third Friday in a month, we move the expiration dates one day to the third Friday.

Appendix C: Details on robustness

Missing values. Results seem robust to alternative treatments of missing values. Our main runs are based on filling in missing values with the sample mean of the variable (the first half of option early exercise is never filled in but treated separately by running short sample regressions). We compare two different treatments of missing values by rerunning equation (79). However, eliminating observations with missing values from the sample (8 missing observations) does not change the qualitative results. As reported in model (C.1), the estimated parameters are only marginally different from model (79) and they remain highly statistically significant.

$$\begin{aligned} Pinn\_sym_t^L &= -4.65 - 1.23 OI_t + 1.60 VOL_t + \varepsilon_t \\ p - value &\;\; (0.02) \quad (0.00) \qquad (0.00) \end{aligned} \tag{C.1}$$

Results are also robust to replacing missing observations with zeros which we test in equation (C.2). In this case, the estimated parameters on open interest and option volume decrease by about half a standard deviation and open interest has a slightly higher p-value but remains significant at the 5% level.

$$\begin{aligned} Pinn\_sym_t^L &= -5.19 - 0.98 OI_t + 1.43 VOL_t + \varepsilon_t \\ p - value &\;\; (0.01) \quad (0.02) \qquad (0.00) \end{aligned} \tag{C.2}$$

Holidays on the third Friday. Expiration days usually fall on the third Friday of the month. If the third Friday is a holiday then the Thursday before the third Friday is used. In our sample, it occurs only three times that the serial expiration does not fall on a Friday. Since these are unusual expiration days, we next estimate our main expiration dummy models (75) and (76) using only serial expirations that occur on Fridays. We are left with 174 observations. The sign of the coefficient stays positive and all p-values stay the same or even decrease. Finally, we rerun equation (79) and, as reported in equation (C.3), all results survive and even exhibit larger coefficients.

$$\begin{aligned} Pinn\_sym_t^L &= -4.12 - 1.44 OI_t + 1.74 VOL_t + \varepsilon_t \\ p - value &\;\; (0.04) \quad (0.00) \qquad (0.00) \end{aligned} \tag{C.3}$$

Crash months. In our main runs we always exclude October 1987 and October 2008. We rerun our main expiration dummy models (75) and (76) using all the serial expiration days which gives us 179 observations. All results come through with minimally smaller coefficients and minimally larger p-values. Rerunning equation (79), we confirm in equation (C.4) that the inclusion of the crash months does not have an impact on our conclusions as the results remain virtually unchanged.

$$\begin{aligned} Pinn\_sym_t^L &= -4.71 - 1.29 OI_t + 1.67 VOL_t + \varepsilon_t \\ p-value &\quad (0.02) \quad (0.00) \qquad (0.00) \end{aligned} \tag{C.4}$$

Outlying observations. Visual inspection of Figures 5 let us wonder about two outlying observations in August 2007 and August 2008. Elimination of those two months does not change any of our results. Point estimates tend to be affected by less than 10% and p-values by less than 1%.

2.9 TABLES AND FIGURES

Table 1: Summary statistics

This table collects the summary statistics for the serial expiration dates in the period August 1987 to November 2009 (excluding October 1987 and October 2008). Symmetric pinning $Pinn\_sym$ is a zero/one variable, which is 1 if the future settlement price is within \$0.25 to the left or right of the ATM strike price. Similarly, $Pinn\_above$ and $Pinn\_below$ are zero/one variables that take a value of one if the future settlement price is within \$0.25 above or below of the ATM strike price, respectively. ATM open interest OI is measured one day before the serial expiration day with respect to ATM strike price on the serial expiration day. ATM volume VOL and ATM options exercise OE are both measured on the serial expiration date with respect to the ATM strike price. ATM volume is the sum of ATM call volume $Call\_VOL$ and ATM put volume Put_VOL. Similarly, ATM option exercise OE is the sum of ATM call option exercise $Call\_OE$ and ATM put option exercise Put_OE. Futures volume Fut_vol measures the number of contracts traded on the serial expiration Friday. Futures volatility Fut_sigma is a scaled realized daily range measured one day before the expiration date. We replace missing observations by the mean of non-missing observations. Numbers in brackets next to number of observations denote number of non-missing observations for each variable. Summary statistics for OE are based on the period October 1998 to November 2009. We add 10 to all open interest, option and futures volume, and early option exercise values and take logarithms.

Variable	Subvariable	NObs	Mean	StdDev	Min	Max
$Pinn\_sym$		177 (177)	0.14	0.34	0.00	1.00
	$Pinn\_above$	177 (177)	0.08	0.28	0.00	1.00
	$Pinn\_below$	177 (177)	0.07	0.26	0.00	1.00
OI		177 (169)	7.43	0.98	3.53	9.69
VOL		177 (175)	7.27	1.04	3.09	9.72
	$Call\_VOL$	177 (175)	6.56	1.17	2.94	9.38
	Put_VOL	177 (177)	6.21	1.34	2.30	8.68
OE		89(88)	6.01	1.30	2.30	9.18
	$Call\_OE$	89(88)	4.43	2.00	2.30	8.16
	Put_OE	89(89)	4.44	1.96	2.30	8.89
Fut_vol		177 (177)	10.79	0.43	9.62	11.99
Fut_sigma		177 (177)	0.13	0.10	0.03	0.91

Table 2: Unconditional correlation structure for the main variables

This table collects the unconditional correlations for the main variables for the serial expiration dates in the period August 1987 to November 2009 (excluding October 1987 and October 2008). Symmetric pinning $Pinn\_sym$ is a zero/one variable, which is 1 if the future settlement price is within \$0.25 to the left or right of the ATM strike price. ATM open interest OI is measured one day before the serial expiration day with respect to ATM strike price on the serial expiration day. ATM volume VOL and ATM options exercise OE are both measured on the serial expiration date with respect to the ATM strike price. Futures volume Fut_vol measures the number of contracts traded on the serial expiration Friday. Futures volatility Fut_sigma is a scaled realized daily range measured one day before the expiration date. We replace missing observations by the mean of non-missing observations. Correlations for OE are based on the period October 1998 to November 2009. We add 10 to open interest, option and futures volume, and early option exercise values and take logarithms.

	$Pinn\_sym$	OI	VOL	OE	Fut_vol	Fut_sigma
$Pinn\_sym$	1.00	0.01	0.14	0.11	-0.10	-0.04
OI	.	1.00	0.81	0.55	-0.07	-0.18
VOL	.	.	1.00	0.60	0.00	-0.13
OE	.	.	.	1.00	-0.35	-0.11
Fut_vol	.	.	.	.	1.00	0.18
Fut_sigma	.	.	.	.	.	1.00

Table 3: Unconditional correlation structure for the subvariables

This table collects unconditional correlations for the subvariables for the serial expiration dates in the period August 1987 to November 2009 (excluding October 1987 and October 2008). Symmetric pinning $Pinn\_sym$ is a zero/one variable, which is 1 if the future settlement price is within $0.25 to the left or right of the ATM strike price. Similarly, $Pinn\_above$ and $Pinn\_below$ are zero/one variables that take a value of one if the future settlement price is within $0.25 above or below of the ATM strike price, respectively. ATM open interest OI is measured one day before the serial expiration day with respect to ATM strike price on the serial expiration day. It is a sum of ATM call open interest and ATM put open interest. ATM volume VOL and ATM options exercise OE are both measured on the serial expiration date with respect to the ATM strike price. ATM volume is the sum of ATM call volume $Call\_VOL$ and ATM put volume Put_VOL. Similarly, ATM option exercise OE is the sum of ATM call option exercise $Call\_OE$ and ATM put option exercise Put_OE. We replace missing observations by the mean of non-missing observations. Correlations for OE are based on the period October 1998 to November 2009. We add 10 to all open interest, option volume, and early option exercise values and take logarithms.

		$Pinn\_sym$	$Pinn\_above$	$Pinn\_below$
$Pinn\_sym$		1.00	0.77	0.71
	$Pinn\_above$	.	1.00	0.23
	$Pinn\_below$	.	.	1.00
OI		0.01	-0.05	-0.09
VOL		0.14	0.06	0.01
	$Call\_VOL$	0.12	0.09	-0.02
	Put_VOL	0.12	0.06	0.02
OE		0.11	-0.01	-0.02
	$Call\_OE$	0.25	0.22	0.00
	Put_OE	0.06	-0.14	0.18

Figure 1: Percentage of S&P 500 futures finishing within +/-0.25$ of the strike price

Panel A: August 1987 – November 2009

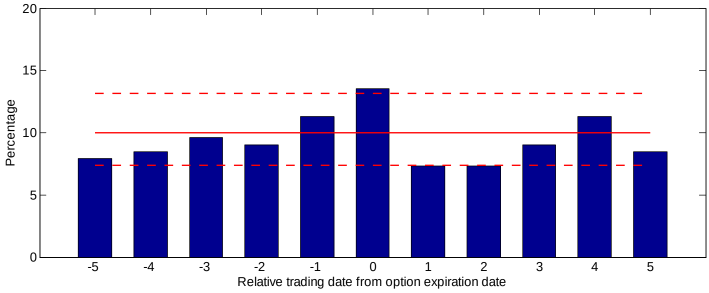

Panel B: October 1998 – November 2009

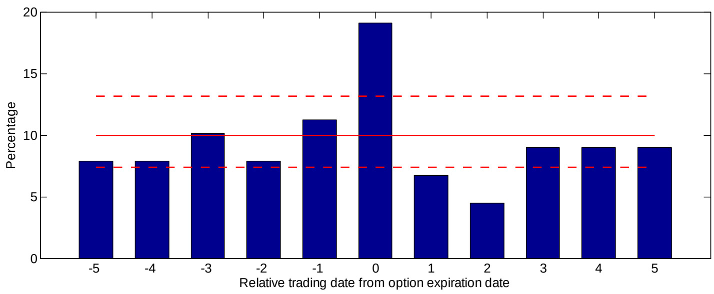

This figure depicts the percentage of S&P 500 futures settlement prices which are within the +/-\$0.25 range around the ATM strike price. This proportion should be 10% if prices are uniformly distributed around the strike price (horizontal line plus bounds for the 10^{th} and 90^{th} percentile based on equation (72)). The figure presents results for the 5 days before and after the serial expiration dates and for the serial expirations themselves. Panel A depicts results for the period from August 1987 to November 2009 and Panel B depicts results for the period from October 1998 to November 2009. Both panels exclude October 1987 and October 2008.

Figure 2: Delta-hedging of short straddles (sold gamma positions)

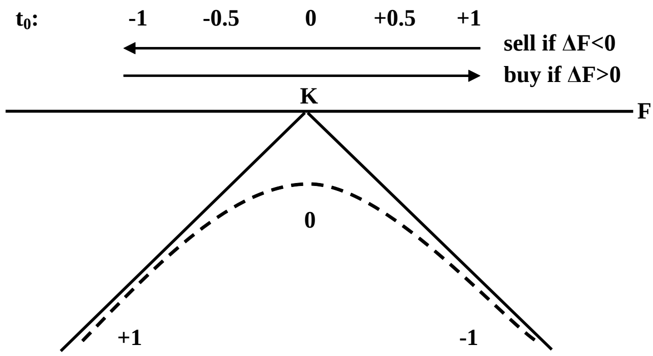

This figure depicts the hedging of an ATM sold straddle where the delta of the hedge is noted at the top of the figure and the delta of the straddle at the bottom.

Figure 3: Pinning mechanism of Avellaneda and Lipkin (2003)

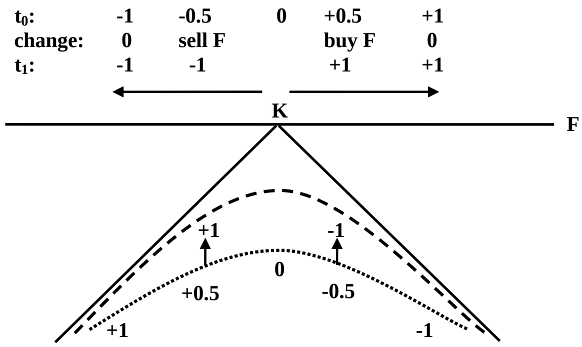

This figure depicts the hedging of an ATM sold straddle where the delta of the hedge is noted at the top of the figure and the delta of the straddle at the bottom. The figure demonstrates how the hedge for different levels of the future changes as time passes from t_0 (dotted line) to t_1 (dashed line). The resulting adjustment trades to the hedge cause anti-pinning as indicated by the arrows.

Figure 4: Pinning mechanism of Anders (1982)

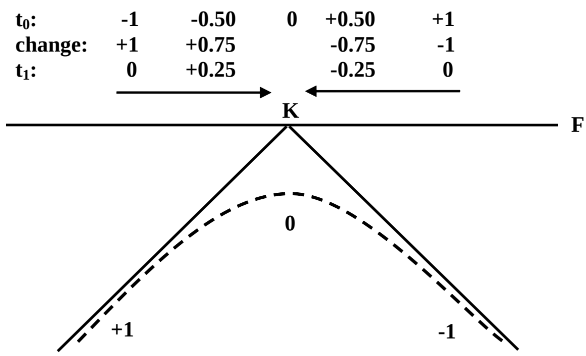

This figure depicts the hedging of an ATM sold straddle where the delta of the hedge is noted at the top of the figure and the delta of the straddle at the bottom. The figure demonstrates how the hedge for different levels of the future changes as the unhedged investor sells the ITM option to the market maker (t_0 is before the sale and t_1 thereafter). The resulting adjustment trades to the hedge cause pinning as indicated by the arrows.

Figure 5: Time pattern of open interest, volume, and options exercise of ATM options

Panel A: ATM Open Interest

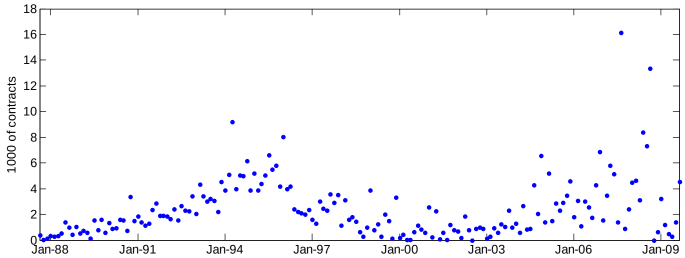

Panel B: ATM Option Volume

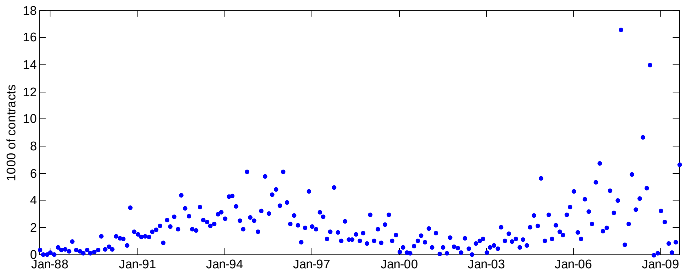

Panel C: ATM Early Option Exercise

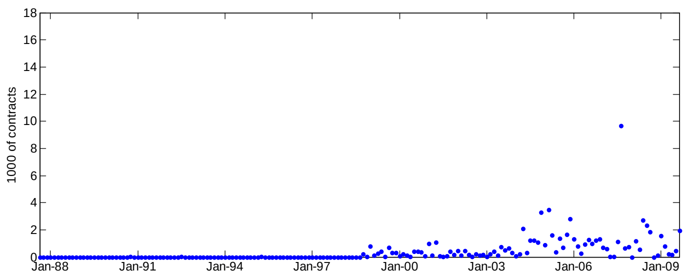

This figure depicts SP option trading activity for serial expiration dates between August 1987 and November 2009. Panel A depicts SP option open interest measured one day before the serial expiration Friday with respect to ATM strike price on the serial expiration date. Panel B depicts ATM option volume measured on the serial expiration Friday. Panel C depicts ATM early option exercise measured on the serial expiration Friday. Missing values are replaced by zeros. Quantities have not been transformed by logarithms or the addition of 10.

Table A.1: The main characteristics of the S&P 500 derivatives

This table collects the main characteristics of the S&P 500 index related derivatives. Q stands for quarterly expiration cycle (March, June, September and December). S stands for serial expiration months (January, February, April, May, July, August, October and November). A.M. settlement price is based on the special opening value of the underlying. P.M. settlement price is based on last prices of the underlying on an expiration day. The table is based on the information obtained from CME webpage (http://www.cmegroup.com), CBOE webpage (www.cboe.com), Stoll and Whaley (1991), and the CBOE Regulatory Circular Number RG92-46.

	SPX options	SP futures	SP options	SPY options
Underlying	S&P 500 index	S&P 500 index	SP futures	SPDR (ETF, 1/10th of S&P 500 index)
Opening date	7/1/1983	4/21/1982	1/28/1983	January 2005
Strike price interval	5 points (first nearest month) 25points (deferred months)	-	5 points (first and second nearest month) 10 and 25 points (deferred months)	1 point
Type	European (American before April 1986)	-	American	American
Trading hours (Central Time)	8:30 A.M. - 3:15 P.M.	8:30 A.M. - 3:15 P.M.	8:30 A.M. - 3:15 P.M.	8:30 A.M. - 3:15 P.M.
Expiration months	3 serial months 3 months in the quarterly cycle	8 months in the quarterly cycle	3 serial months 8 months in the quarterly cycle	3 serial months 3 months in the quarterly cycle
Settlement at expiration	Q+S: Cash-settlement	Q: Cash-settlement	Q: Cash-settlement S: Physical delivery	Q+S: Physical delivery
		Settlement value		
Until June 1984	Q: P.M. settlement (3rd Friday)	Q: P.M. settlement (3rd Thursday)	Q: P.M. settlement (3rd Thursday)	-
June 1984 to June 1987	Q+S*: P.M. settlement (3rd Friday)	Q: P. M. settlement (3rd Friday)	Q: P.M. settlement (3rd Friday)	-
June 1987 to December 1993	Q+S: Co-existence of A.M. and P..M.. settled options (3rd Friday)	Q: A.M. settlement (3rd Friday)	Q: A.M. settlement (3rd Friday) S: P.M. settlement (3rd Friday)	-
Dec. 1993 onwards	Q+S: A.M. settlement (3rd Friday)	Q: A.M. settlement (3rd Friday)	Q: A.M. settlement (3rd Friday) S: P.M. settlement (3rd Friday)	Q+S: P.M. price (3rd Friday)

*Serial expirations for SPX options were introduced in 1986 (they exist for sure since April 2nd 1986).

Chapter 3

PRICE SUPPORT IN THE STOCK MARKET (WITH JOSÉ M. MARÍN)

3.1 Introduction

The term "price support" traditionally refers to public interventions in security markets that aim to maintain asset prices above their market values. These public interventions are mainly policy driven and are especially important in times of turmoil. The intervention of a central bank defending a country's currency is probably the most widespread price support activity in financial markets, but such activity is not the only type. For instance, the U.S. Troubled Asset Relief Program (TARP) of 2008 is similar in nature, but targets a different type of asset (mortgage-related securities) and has a different goal—namely, to provide financial intermediaries with liquidity and solvency. In this paper we provide evidence of similar "interventions" that are implemented by traders, rather than public entities, in the stock market. These interventions also tend to occur during turbulent times; however, unlike public interventions—which are motivated by policy reasons and, perhaps, are for the greater good of the country—private interventions arise as a natural response to conflicts of interest and agency problems that prevail in security markets.

Delegated portfolio management is the norm rather than the exception in the asset management industry: most individuals do not invest in security markets directly but rather via mutual, pension, and hedge funds.[5] These

[5] In April 2008 the assets under management, as reported by the 11,030 U.S. SEC-registered investment advisers, totaled $42.3 trillion (Investment Adviser Association, 2008). Moreover, according to the 2007 Survey of Consumer Finance, the direct stock holdings of U.S. households amount to 17.9% of their total financial assets (Bucks et al., 2009); indirect household stock holdings (defined as the sum of retirement accounts and other managed assets) are more than twice as large (accounting for 41.1% of assets).

funds are managed by asset management firms, which are legal entities different from the funds they manage. Consequently, even though fund managers have the obligation to manage fund assets in the interest of the fund's investors, they are actually employees of the asset management firm. Fund managers are thus "double agents" in that they have two principals: the fund's investors and the management firm's owners. The interests of these two principals are not necessarily aligned. By law the interest of the fund investors must prevail, but in practice we may observe cases where fund managers act more on behalf of the management firm's shareholders. This means that, in the presence of delegated portfolio management, the ownership of asset management firms matters because it may give rise to agency problems that affect asset trading. To the best of our knowledge, the literature has addressed neither this economic problem nor its clear implications for asset management, asset pricing, performance evaluation, and regulation.

This "generic" conflict of interest gains ground when we observe that the main purpose of funds is to hold a portfolio of securities over time and that some management firms are fully owned and/or controlled by publicly traded companies. In this case, the stock of the controlling publicly traded shareholder qualifies as an investment in the fund. A conflict of interest then arises because the fund manager may make decisions involving such stock that are in the interest of the controlling shareholder and not of the fund's investors. As a result, the double agency problem analyzed in this paper bears strongly on portfolio theory and performance evaluation because it generates a new motive for trading. Double agency problem also affects asset pricing in that the new trading may have a direct impact on asset prices. Finally, it opens a policy debate concerning the status of the controlling shareholder's stock in the portfolio of affiliated funds.

In this paper we focus on a particular type of controlling shareholders of asset management firms: financial conglomerates or banks. We provide evidence that bank-affiliated funds (i.e., funds managed by asset management firms that are controlled by banks) systematically increase, relative to nonaffiliated funds, their holdings of the controlling bank stock when it suffers a large price drop (large negative return). Similar purchases are not made of competing banks' stock when they suffer a similar shock, which indicates that extra purchases of the controlling bank's stock do not stem from a "contrarian" style of investment. Neither are these purchases the result of an overall increase in the allocation to banking stocks. We also provide evidence that such purchases are not part of a "timing" strategy associated with anticipation of a large drop in the parent bank's stock price, and we find that purchases following such large drops do not outperform a portfolio that includes all the other banks. These last two results lead us to conclude that the identified patterns of trading by an affiliated fund in its controlling bank's stock are not driven by private information, either. However, the patterns are consistent with our price support hypothesis, according to which affiliated funds increase their holdings of the parent bank's stock in times of turmoil in order to limit—in the interest of the bank's shareholders and management—the downside potential of the stock price.[2]

Our study is undertaken in the context of the Spanish mutual fund industry. The choice of Spain is not capricious; in fact, given the nature of the problem we address, it is fully justified. First, notice that a necessary condition for testing our hypothesis is the existence of bank-affiliated funds with considerable assets under management. If the amount of assets managed by bank-affiliated funds is negligible, then we should not expect price support activities to succeed—that is, to have an impact on the bank's stock price. Hence, in that case we would not expect fund managers to pursue price support activities. This means that our price support hypothesis (a) is not testable in countries without bank-affiliated funds and (b) will likely be rejected in countries where banks have relatively little involvement in fund

[2] There are many (obvious) reasons why bank shareholders and managers benefit from price support. First, if price support is effective (i.e., if it limits the downside potential of stock returns) then the stock becomes a less risky investment; this directly benefits current bank shareholders. It also makes the bank's securities a more attractive investment for new investors as compared with the issues of banks having no (or unsuccessful) price support. Second, to the extent that management is remunerated via performance-based compensation packages, it also benefits directly from successful price support activities. In any event, successful price support is beneficial in terms of career concerns and peer competition.

asset management. Second, prosecution of crimes is a key issue. Strictly speaking, price support activities by mutual funds are illegal because the trades may not be in the interest of the fund's investors. We should therefore expect price support to be more likely in countries where such activity is neither closely monitored nor prosecuted and severely punished by the authorities. Spain clearly qualifies on both counts and so is a clear candidate for our study.

One of the distinctive characteristics of the Spanish financial sector is its historically strong reliance on banks, rather than markets. In their classical distinction between intermediated financial systems and those driven by stock markets, Allen and Gale (1994) include Spain among the countries where intermediation dominates. This dominance is also reflected in delegated portfolio management. In 2009, funds affiliated with Spanish banks represented nearly half of the assets under management by the Spanish mutual fund industry. If we add the funds affiliated with Spanish savings and loan institutions, then this share of total assets under management increases to nearly 80%.[3] Hence Spain features an asset management industry in which the presence of financial intermediaries in delegated portfolio management is significant enough to test our hypothesis.

With regard to crime prosecution, the history of the Spanish "SEC"[4] shows that Spain is an ideal candidate for testing our hypothesis. Table 1 compares the U.S. and Spanish securities commissions in terms of their track records on crime investigation and prosecution during the past six years. Panel A shows U.S. SEC data as given in its Performance and Accounting Report and by SEC Statistics; panel B shows equivalent Spanish data as reported in its commission's Annual Reports.[5] Large differences are evident in the

[3] In the United States, about 40% of mutual funds belong to financial conglomerates involved in banking and insurance (Massa and Rehman (2008)). Note that, unlike the Spanish data, this figure includes funds affiliated with insurance institutions.

[4] Comision Nacional del Mercado de Valores (CNMV).

[5] The figures for civil cases are not entirely comparable. The Spanish SEC does not have the authority to initiate civil or penal proceedings, so all cases are of an administrative nature. The CNMV merely forwards to the General Attorney those cases that could probably be taken to trial. In contrast, the U.S. figures reflect actual civil proceedings.

number of cases initiated, cases closed, and enforcement. Perhaps much of these differences is explained by the relative size of the two countries' financial sectors. However, the differences in enforcement related to asset management cannot be explained on this basis: Table 1 makes it clear that, unless we assume Spanish investors and financial institutions to be extremely law-abiding, crime investigation and prosecution in Spain is much weaker than in the United States.

The previous two paragraphs justify setting our study in the context of the Spanish mutual fund industry. Of course, testing our price support hypothesis in other countries will be essential for drawing inferences on the actual level of investor protection offered by different countries. There is a fairly recent literature, led by La Porta et al. (2000), showing that the legal system's level of investor protection is key to understanding the patterns of some financial variables (e.g., the cost of capital, capital flows, corporate ownership structures) across countries. Most countries have legislation in place stipulating that delegated portfolio management must be in the interest of investors only, which means that price support activities are unlawful virtually everywhere. However, the level of effort put into investigating and prosecuting transgressions—and the efficacy of prosecution—varies widely from country to country. The same type of law can therefore result in different levels of investor protection, and the extent of variation in these levels can be clarified by testing for the existence of price support activities.

The purpose of price support activities is to affect the return distributions of the supported stocks. There is ample evidence that trading by mutual funds has a price impact. For instance, Sias et al. (2001) document a positive correlation between changes in institutional ownership and returns. Further, the authors suggest that the price impact results from the information extracted from the institutional trades. In principle, we suspect that a mutual fund's trades to support its parent bank's stock can have an even larger price effect than documented by Sias et al. (2001) because price support occurs in turbulent times, when inferring information is especially critical. Unfortunately, it is beyond the scope of this paper to analyze the impact of price support activities on return distributions. Note that it would not be straightforward to test for this impact. One possibility would be to identify the (partially) unobservable left tail of the return distribution by comparing the return distribution of price-supported stocks with that of comparable stocks for which there is no price support. In fact, this is the approach followed by Ruud (1993) when testing the underwriters' price support hypothesis in IPOs. Because all Spanish banks have asset management arms that could practice price support, there are no comparable yet nonsupported banking stocks that would make this approach a viable one for this study.

That being said, price support activities by affiliated funds must affect the fund's overall performance. However, we do not address this important issue, either. In principle we expect price support to affect performance negatively, but identifying this trend in the data is difficult. The main reason is that the relationship between the bank and its affiliated funds has many dimensions (beyond price support) that affect performance. Masa and Rehman (2008) analyze the flow of information within financial conglomerates and offer evidence of information "leakages" from parent banks to affiliated funds on the banks' lending activities—information that results in affiliated funds increasing their holdings in companies whose loans have been renewed. The authors show that returns on such transactions outperform portfolios of similar control companies. Hence price support will tend to worsen performance whereas information leakages will tend to improve it,[6] and the outcome of these two competing effects is an empirical question. Analyzing the impact of price support on both price and performance are top priorities on our research agenda.

Our study is related to several areas of research in finance. For example, price support activities have been analyzed in the context of the IPO underwriting

[6] Affiliated mutual funds also enjoy other benefits that may improve their performance. Mola and Guidolin (2009) document that affiliated analysts are likely to assign frequent and favorable ratings to stocks that are part of the affiliated mutual funds or that have large portfolio weights in affiliated fund families. Ritter and Zhang (2007) show that affiliated investment banks allocated "hot" IPOs during the 1999–2000 Internet bubble to affiliated mutual funds.

business. This literature argues that (part) of the underpricing of IPOs is explained by the trading of IPO underwriters. This literature can be classified into two broad categories. On the one hand are papers, such as those by Ruud (1993) and Prabhala and Puri (1998), in which price support–based trading by underwriters is not directly measured but rather inferred from documented differences in the distribution of the returns of IPO versus non-IPO stocks. On the other hand are papers in which underwriter activity is either directly analyzed (Ellis et al. (2000), Schulz and Zaman (1994)) or proxied by the inventory of market makers (Lewelen (2006)). Our analysis is closer to the second approach in that we analyze the actual trading activity of price supporters. However, we must emphasize that the economic forces behind price support of IPOs versus banking stocks are very different. First, price support in the IPO business does not result from a double agency problem like the one analyzed here but rather from the underwriter's interest in improving its reputation or compensating IPO participants for providing relevant information in the pre-IPO stage. Second, IPO price support by underwriters is—unlike the price support analyzed in this paper—entirely legal.[7] Finally, price support by underwriters is confined a single brief period in the security's history, whereas the price support that we document may occur during many different periods of the stock's life span.

Our research is also related to a large and growing literature on the role played by incentives and agency conflicts in funds trading. For instance, previous research has identified clear patterns of calendar-driven risk taking by asset management firms (Chevalier and Ellison (1997)) and of so-called window dressing (Lakonishok et al. (1991)). Closely related to our research is the previously mentioned paper by Massa and Rehman (2008), who analyze the trading patterns of bank-affiliated mutual funds that arise from conflicts of interest. However, Massa and Rehman analyze information leakages from the bank to the affiliated funds resulting in trades of potential *benefit* to fund investors, whereas we analyze activity of affiliated funds on the parent bank's behalf resulting in trades that could *harm* fund investors. Finally,

[7] For instance, see Jenkinson and Ljungqvist (2001).

our analysis also contributes to the literature on the role played by financial institutions in general—and by the legal system in particular—in the efficient allocation of capital.[8]

In this paper we focus on trading patterns in the mutual fund industry that are consistent with the existence of price support. However, we believe this is just the tip of the iceberg of price support activities in security markets arising from agency conflicts. It is typical for financial conglomerates to own not only management firms but also a portfolio of firms that they control. The conflict of interest that the managers of affiliated funds face may result in trading to support not only the price of the bank stock, but also the prices of these affiliated firms' stock. Therefore, the analysis of price support for stock of the bank's corporate affiliates seems like a promising topic for future research. Other examples of possible price support lie outside the asset management industry. For instance, it is typical for firms or wealthy investors to place some of their assets in foundations. Although such a foundation generally has a charter that spells out specific investment goals, it seems likely that the foundation's assets are sometimes used to support the prices of stocks related to the individuals or corporations that funded the foundation. In general, we can suspect a degree of price support activity in any area of the financial industry where the double agency problem is present.

The balance of the paper is organized as follows. In Section 3.2 we formulate the price support hypothesis and address some methodological issues involved in its testing. In Section 3.3 we describe the data and define the main variables used in the study. In Sections 3.4 and 3.5, respectively, we test our price support hypothesis at the level of individual funds and the level of fund families and banking groups. In Section 3.6 we rule out private information as an explanation for the patterns of trading observed in affiliated funds. Section 3.7 is dedicated to several robustness checks, and we offer some concluding remarks in Section 3.8.

[8]See, for example, La Porta et al. (2000), Wurgler (2000), Bhattacharya and Daouk (2002), and Durnev and Kim (2005).

3.2 Hypotheses

The intervention of central banks in currency markets does not occur in a continuous fashion but rather when the currency is under attack. Likewise, according to our price support hypothesis, bank-affiliated funds increase their holdings of their parent bank's stock only when it is "under attack"—in other words, when it seems necessary to limit the stock price's downside potential. We can imagine many alternative definitions of attack episodes. One would involve identifying missing parts in the tail distribution of returns on the bank's stock and comparing them to the return distribution of a comparable bank that does not enjoy price support. Unfortunately, as we mentioned before, all Spanish banks have asset management arms which, ex ante, could be price supporters. Another possibility is to identify attacks exogenously, or independently of the returns on the bank's stock. For instance, we could identify attack periods with episodes of bad public news about the stock. However, this approach would require ad hoc assumptions regarding the severity of the bad news, and the best available metric would likely be the price impact of the bad news.[9] In the absence of a better definition, in this paper we use the realized return on the bank's stock as our metric. In particular, we define an attack episode as a period during which the bank's stock suffers a large negative return (large price drop).

We must emphasize that this definition of attacks suffers an endogeneity problem that actually favors *rejection* of the price support hypothesis. To understand this, consider an economy where all banks are identical and suffer only systemic shocks. Let's analyze first the extreme case in which all the affiliated funds of all banks trade frequently and successfully to support prices. In this case we would not observe stock price corrections (since price support is fully effective) and so would not be able to test the hypothesis. Next assume that only the affiliated funds of some (but not all) banks en-

[9] Another possibility would be to proxy attacks by episodes of abnormal short selling interest in the stock. Unfortunately, Spanish data on short selling interest is partial and available only for the last two years.

gage in price support and that this activity is only moderately successful. Ceteris paribus, the stock price of non–price-supported banks will fall more than that of banks whose stock price is being supported. Hence, any definition of a price drop—whether in isolation or in excess of some benchmark index—will result in drops most likely associated to banks whose stock price is not being supported. Thus, any empirical test will associate events (price drops) to the lack of price support activity. This discussion illustrates how the endogeneity problem generates a bias toward rejecting the price support hypothesis. This is important as in this paper we find strong evidence of price support even in the presence of this endogeneity issue.

Another problem with testing the price support hypothesis concerns the choice of a time interval in which to test for it. In this study, we use quarterly holdings of mutual funds. A large price drop can occur arbitrarily during the quarter. Similarly, price support–based trading by mutual funds can take place at any time within the quarter (or in future quarters). Under our price support hypothesis, then, the quarter in which the drop occurs can exhibit arbitrary amounts of price support–based trading relative to the expected (average) amount—depending on the timing of the price drop within the quarter. If the drop happens at the very start of a quarter, then the price support–based trading will be in line with the expected quarterly amount of such trading; but if the drop takes place near the end of the quarter, then the former will be much smaller. For this reason, in our analysis we exclude the "drop quarter" when computing changes in holdings and instead test the price support hypothesis with reference to trading in the quarter *after* a large drop in the bank's stock price.[10]

We can now formally state our price support hypothesis as follows.

[10] In the robustness checks described in Section 7, we document that the results of our study survive (at even higher levels of statistical significance) inclusion of the drop quarter in the analysis.

- **Hypothesis 1:** *Price support activity.* After large drops in the parent bank's stock price, affiliated funds increase their holdings of the parent bank's stock more than do nonaffiliated funds.

However, the increase in holdings after a large drop is consistent not only with the price support hypothesis but also with many other hypotheses. First, it may be that affiliated fund managers view a large drop in the parent bank's stock price as a signal to increase their exposure to the whole banking sector. Under this "sector rebalancing" hypothesis, affiliated funds should also increase their holdings of other bank stocks when the parent's stock price suffers a large drop. Second, the increase in holdings following negative returns is also consistent with a "contrarian" style of trading. Note that this hypothesis implies that if affiliated funds increase their holdings of the parent bank's stock after a large negative return then they should also increase their holdings of other banking stocks that have suffered a similar shock. Third, affiliated funds may anticipate large drops in their parent bank's stock price and therefore implement a "timing" strategy whereby the stock is sold just before the drop and then repurchased immediately thereafter. According to this hypothesis, affiliated banks should reduce holdings of the parent bank stock *before* large negative returns. Finally, an increase in holdings following a large drop could be motivated by private information regarding the parent bank's future stock performance. Under this hypothesis, the parent bank's stock should outperform a portfolio that invests in alternative banking stocks.

These considerations motivate an explication of Hypothesis 1 in terms of the four competing hypotheses as follows.

- **Hypothesis 2:** *No sector rebalancing explanation.* After large drops in the parent bank's stock price, affiliated funds increase their holdings of that stock relative to other bank stocks more than do nonaffiliated funds.

- **Hypothesis 3:** *No contrarian explanation.* After large drops in the parent bank's stock price, affiliated funds increase their holdings of the parent bank's stock, relative to other bank stocks that also suffer large price drops, more than do nonaffiliated funds.

- **Hypothesis 4:** *No anticipating explanation.* Before large drops in the parent bank's stock price, affiliated funds do not reduce their holdings of the parent bank's stock more than do nonaffiliated funds.

- **Hypothesis 5:** *No private information explanation.* After a large drop in the parent bank's stock price, a portfolio consisting of that stock does not outperform a portfolio consisting of other banking stocks.

Hence, our theory of price support is formulated in terms of Hypotheses 1 to 5. These five hypotheses can be tested at three different levels of aggregation of mutual funds holdings: individual funds, fund families, and banking groups. In principle, supporting evidence at one level need not imply supporting evidence at other levels. It might be the case, for example, that price support is implemented via many affiliated funds and so requires only small increases in each of their holdings; in this case, our hypotheses would likely be rejected if the analysis were performed at the level of individual funds. On the other hand, if each affiliated fund family holds a large fraction of the parent bank then price support activities would probably require only relatively small purchases by each of the fund families or by the group as a whole (when aggregating holdings of all fund families controlled by the same bank); in this case, the price support hypothesis would likely be rejected at the levels of fund families and banking groups.

3.3 Data and variable definitions

3.3.1 Data sources

We merge data from two main sources: the Spanish SEC and Datastream. From the Spanish SEC we obtain all the mutual fund related data used in this study.[11] The data set is very comprehensive and specially suitable for the purposes of the present paper. The data set:

- Covers the *whole universe* of mutual funds registered with the Spanish SEC.
- Provides the quarterly portfolios of all mutual funds for the period 1995Q1 to 2009Q3.
- It also includes many other fund related variables such as the fund's Net Asset Value (NAV), Assets Under Management (AUM), fund family affiliation, number of investors in the fund, fund's inception day, fees, inflows/outflows and investment style.
- Identifies the asset management (Fund Family) managing each of the funds and the controlling group (Ownership) of each of these asset management firms.

Further, doing direct searches in the Spanish SEC registry we obtain the funds merger dates.[12] Figure 1 provides the evolution of the assets under management and number of funds in the Spanish mutual fund industry. The figure illustrates that assets under management peaked by the end of year 2007, reaching a total close to 280 billion euros. The number of funds steadily increases during the sample period until reaching a total of almost 3000 funds in the second quarter of 2008. The current crisis has reduced dramatically the number of funds and the assets under management in this industry. As of September 2009, there were 164 billion euros under management in the Spanish mutual fund industry divided among 2633 mutual funds.

From Datastream we obtain stock data for banks and the Madrid SE General Index for the period January 1990 to September 2009. In particular,

[11] During the period 1995Q3 to 2008Q3, all funds in the data set are uniquely identified by their ISIN code. Starting 2008Q4, funds were allowed to create sub-funds. In these cases, the ISIN code was assigned to one of the subfunds. Only 10 funds in the sample opted for this split. We manually aggregated these subfunds into the fund that could be tracked by its ISIN code.

[12] First mergers between funds in Spain took place in 2001.

from Datastream we obtain the stock split adjusted price, the market capitalization and the turnover (or volume). The merge of portfolio holdings information (Spanish SEC dataset) and banks stock data (Datastream) is done using a unique identifier for each security in the portfolio, the ISIN code. ISIN codes, however, change over time due to stock splits. Datastream only provides the current codes. For this reason we got directly from the Madrid Stock Exchange the ISIN history for all banks. Using this information we match 100% of the mutual fund bank holdings to their corresponding banking stocks.

In some parts of the study we also make use of other data sources. In particular, in the analysis of the performance of strategies based on price support activities performed in Section 6 we make use of the Spanish 4 Fama-French factors computed by Professor Rafael Santamaria, the leading expert in the subject in Spain, for the period July 1990 to December 2009.

3.3.2 Controlling banks and affiliated funds

We are interested in studying the different trading patterns of funds affiliated to banks. In order for a bank to be included in the treatment group in our study, it must satisfy the following three criteria:

- The bank has to control at least one asset management firm.
- The bank is incorporated in Spain and it is traded on the Madrid SIBE (the Spanish computerized trading platform) for at least two consecutive years during our sample period.
- The bank is not controlled by another corporation. If the bank is controlled by another corporation, the Spanish SEC data assigns the asset management company of the controlled bank to its controlling shareholder.[13]

[13] For example, the asset management firm of Banesto is assigned to Banco Santander Central Hispano who owns more than 90% of Banesto.

To implement these criteria we manually match the controlling groups of asset management firms reported in the Spanish SEC dataset with the Datastream data on active and dead banks listed on the Madrid SIBE for at least two consecutive years during our sample period. This process results in a total of 12 banking groups.

The sample of banks is quite heterogeneous including two of the largest banks in the world, Banco Santander and Banco Bilbao Vizcaya Argentaria (BBVA), as well as medium and small size banks. Some small banks have very low free float and are very illiquid. The stock of these banks rarely enters the portfolios of institutional investors. Since the inclusion of illiquid banks in the analysis would only add noise, we additionally impose that the stock has to be liquid enough. Formally, we impose the criteria that the bank has to have an average annualized standardized trading volume (annualized ratio of volume to market capitalization) above 20%.[14] Table 2 reports this measure for all the banks that meet the previous three criteria. Consequently, the four bottom banks in the table are excluded from the filtered sample of banks in our analysis.[15]

The final sample consists of 8 banks: Banco Bilbao Vizcaya Argentaria (BBVA), Argentaria (dead), Banco Santander, Banco Central Hispano (dead), Bankinter, Banco Popular, Banco Pastor and Banco Sabadell. In the rest of the paper we will refer to funds managed by asset management firms controlled by the 8 banks included in the treatment group as "funds affiliated to banks" or just "affiliated funds". The rest of the funds will be included in the control group and correspond to funds managed by asset management firms controlled by non-financial intermediaries or financial intermediaries that do not meet our criteria, such as Saving and Loans institutions (not traded), foreign banks (not incorporated in Spain) and illiquid or rarely traded banks.

[14] Due to lack of data, we could not compute the free float for all the banks for the whole sample period. For the period April 2002 to September 2009, for which the free-float data on Datastream is available, we verified that the banks with the smallest standardized volume also consistently have the smallest free-float.

[15] In any case, we have verified that all the main results in this paper survive the inclusion of these 4 banks. Explicit computations are available upon request

We will often refer to these funds as "funds non-affiliated to banks" or just "non-affiliated funds".

Notice that the universe of funds includes money market funds, fixed income funds, etc. In the spirit of Masa and Rehman (2008), we filter this universe by eliminating all funds that never held shares of any banking stocks during the whole sample period.[16] After this filter, the initial universe of 4254 unique funds that exist in the whole database is reduced to 1236 funds, of which 418 funds clasify as affiliated funds.[17,18]

Table 3 presents the distribution of affiliated funds and fund families among the 8 banks, and Figure 2 plots the time series of the number of funds and the AUM of affiliated vs. non-affiliated funds. As we can see in Table 3, approximately one third (149 funds) of the affiliated funds are affiliated to Banco Santander. This compares to 87 funds for BBVA and 76 funds for Banco Sabadell. The rest of the banks have between 8 to 31 funds. The same picture is also obtained by looking at the number of fund families belonging to each bank. Banco Santander controls 5 distinct asset management firms, BBVA and Banco Sabadell own 3 asset management firms and the rest of the banks only have one or two fund families.

Figure 2 provides the time series of the number of funds and the AUM of affiliated versus non-affiliated funds in our filtered sample. The 8 banks in the sample manage approximately the same amount of assets as all the other non-affiliated mutual funds together. At the same time, the number of affiliated mutual funds is approximately 3 times lower than the number of non-affiliated mutual funds, which implies that affiliated funds tend to be larger funds.

[16] We also eliminate 29 funds that have incomplete time series.

[17] To allieviate most obvious errors present in the holdings data as well as incubated funds, it is standard in the literature to impose funds to be older than one year to be included in the analysis. We do not impose this requirement because all our exercises involve the use of the fund's lagged annual return, which means that the portfolio holdings of the first year of existence of a fund are never used in the study.

[18] Notice that some funds can be clasified as non-affiliated funds in some parts of the sample and as affiliated funds in other parts of the sample. For example, Banco Sabadell funds are classified as non-affiliated funds before 2003Q3 and as affiliated funds after 2003Q3.

3.3.3 Definition of large price "drops"

To test our price support hypothesis we need to formally define large negative returns events that could trigger price support activities. We suspect that price support is more likely not only when the bank stock price falls significantly in historical terms, but also when it falls relatively to other stocks, including other banking stocks. When these two conditions are met, the bank is singled out in the market and the pressure to revert the situation is maximal. This suggests defining price drop events as events that meet the following two conditions simultaneously:

- **Condition 1:** The bank stock price suffers a large drop in historical terms. This criteria only relies on the times series of each bank return. In the spirit of Marín and Olivier (2008), we formally define this condition as follows. For each bank i, we first compute the monthly average return, $\bar{r}^M_{i,t}$, and the monthly standard deviation of returns, $\sigma^M_{i,t}$, at the end of each quarter t over a rolling window of the past 5 years, or 60 months, of data[19]. Then we transform the monthly average and standard deviation of returns to quarterly frequency as follows:

$$\bar{r}_{i,t} = (1 + \bar{r}^M_{i,t})^3 - 1$$
$$\sigma_{i,t} = \sqrt{3}\sigma^M_{i,t}$$

  Finally, for any bank i, Condition 1 is met in any quarter t in which bank $i's$ return in quarter t, $r_{i,t}$, satisfies

$$r_{i,t} - \bar{r}_{i,t-1} \leq -a\sigma_{i,t-1}$$

  where a is a constant. For instance, when $a = 1$, a drop occurs when the quarterly return is one standard deviation below its historical mean.

[19]In the cases in which the stock has not been publicly traded for 5 years, we require at least 2 years, or 24 months, of data.

- **Condition 2:** The previously defined drop also results in a drop large enough relatively to some market index. This extra criteria relies on the time series of the excess return of the bank stock over the index return. Again, following Marín and Olivier (2008), we formally define this second condition as follows. First we compute the series of monthly *excess returns* for each stock, $r_{i,t}^{ER,M} \equiv r_{i,t}^{M} - r_{INX,t}{}^{M}$, where $r_{INX,t}^{M}$ is the monthly return of the Madrid general SE Index. Then, we proceed as before, computing first the historical monthly average and standard deviation of the excess return, and then converting monthly moments into quarterly moments. Finally, for any bank i, Condition 2 is met in any quarter t in which bank $i's$ excess return in quarter t, $r_{i,t}^{ER}$, satisfies

$$r_{i,t}^{ER} - \bar{r}_{i,t-1}^{ER} \leq -a\sigma_{i,t-1}^{ER}$$

Now, we formally define the $DROP_{i,t}$ variable as:

$$DROP_{i,t} = \begin{cases} 1, & \text{if } r_{i,t} - \bar{r}_{i,t-1} \leq -a\sigma_{i,t-1} \text{ AND } r_{i,t}^{ER} - \bar{r}_{i,t-1}^{ER} \leq -a\sigma_{i,t-1}^{ER}\ ; \\ 0, & \text{otherwise.} \end{cases} \tag{94}$$

In the main body of our analysis we set $a = 0.5$. The choice of 0.5 standard deviations away below the historical mean in the definition of a raw drop is arbitrary. Theory does not provide a threshold level of price drop triggering price support activities. On the one hand, we know the drop has to be large enough for price support activities to take place. But, according to our discussion in section 3.2, the drop should not be "too" large, as in this case, under the hypothesis of price support being effective, we would end up associating price drops to bank stocks where there is no price support. In our study we corroborate this point by doing sensitivity analysis on a and report the results also for the cases of $a = 0$ and $a = 1$.

In order to test the five hypothesis previously defined we need analyze the trading behavior of the affiliated mutual funds both before and after big drops. To avoid contaminating the trading patterns by the simultaneous existence of another drop, we exclude all overlapping drops in a one quarter window. That is, we only include the drops for which the previous and following quarters are not drop events.[20] To avoid contaminating the analysis by the process of banking mergers, we also exclude drops in a one-quarter window (quarter of the merger and preceding and following quarter) around bank mergers or acquisitions. These criteria eliminate many interesting drop events, but is recommended given the nature of our analysis. In any case, the main results of the paper hold true when overlapping drops and drops around merger dates are not excluded.

Using the above definition of the DROP variable, the leading case of $a = 0.5$ clasifies approximately 10% of all the observations as drops. More specifically, altogether we have 323 quarter-bank observations, of which 31 classify as drops. The mean return for all quarterly returns is 4.02% and ranges from -42.96% to 64.93%. The mean return for drops is -14.45% and ranges from -42.96% to -0.09%.

One final remark is in order. Note that we analyze trading patterns of affiliated and non affiliated funds around drop events. In order to test the price support hypothesis against competing explanations it is important to keep a fixed panel of funds around each event. To achieve this we impose the standard minimum of 1 million of AUM filter for the quarter of the drop and the preceding and following quarters. More explicitly, if a drop occurs in quarter t, and the fund has AUM below 1 million EUR in any of the quarters from $t-2$ to $t+1$, the fund is not included in the analysis for that particular drop event. This criteria yields a fixed panel of funds for all fund level regressions and also maximizes on the number of funds included in the analysis as funds excluded in a given drop may be present in other drop events if the minimum 1 million requirement is met.

[20] This implies that drops cannot take place in the first quarter nor in the last quarter of the sample. By the same token, drops cannot take place in the first quarter in which the bank is included in the analysis, nor in the last quarter before the bank ceased to exist.

3.3.4 Definition of the rest of variables

Our basic goal is to test for different patterns of trading of affiliated funds versus non-affiliated funds on the different banking stocks. To achieve this we must first define a measure of trading activity. For each fund, f, the Spanish SEC dataset provides the *quarterly holdings* of each banking stock i at the end of quarter t, $A_{f,i,t}$, in euro terms. These holdings change mechanically when the price of the security changes and, perhaps also when the fund receives inflows/outflows. To alleviate the effect of price changes, we normalize holdings by the market capitalization of the security. More specifically, we define $H_{f,i,t}$ as the percentage of bank i held by fund f at the end of quarter t:

$$H_{f,i,t} \equiv \frac{A_{f,i,t}}{MC_i}$$

In our study we look at changes in holdings around large price drops. Abusing in the use of notation and hoping that no confusion arises, with t we index both a generic quarter and the quarter in which the price drop occurs. Further, to make results comparable across exercises, all changes in holdings are expressed in per quarter terms (average of all quarters involved). More specifically, we use the following general expression for the change in holdings:

$$\Delta H_{f,i,t+j,t+k} \equiv \frac{H_{f,i,t+k} - H_{f,i,t+j}}{k - j} \tag{95}$$

where $k > j$, and which reads as the average change in holdings between quarters $t + j$ and $t + k$ when a price drop occurs in quarter t. Some leading examples:

- $k = 1$ and $j = 0$. In this case we have $\Delta H_{f,i,t,t+1}$, which denotes the change in holdings the quarter after a price drop in quarter t. Note that in this case we do not include the change in holdings that took place in the quarter of the drop, as holdings are computed at quarter ends.

- $k = 0$ and $j = -2$. In this case we have $\Delta H_{f,i,t-2,t}$, which denotes the average change in holdings in the quarter of the crash (t) and the preceding one ($t-1$).

To take into account the effect of the net flows on holdings, we use the fund's net flow as a control variable in all our empirical exercises. Given that the Spanish SEC only provides data on inflows and outflows from April 1999 onwards, we approximate the net flows of fund f in quarter t as:

$$NETFLOW_{f,t} = AUM_{f,t} - AUM_{f,t-1} - AUM_{f,t-1} r_{f,t}$$

where $r_{f,t}$ is the net return of the fund during quarter t. $NETFLOW$ is in millions euros. As a robustness check, in section 3.7 we replicate all main exercises for the period 1999Q2- 2009Q3 using the actual inflows and outflows and show that all the results hold true in this subsample.

In our analysis we use many other control variables. We closely follow the choice of controls in Massa and Rehman (2008).[21] In particular, we control for:

- $AUM(fund)_{f,t}$: Assets under management of fund f in quarter t. They are reported monthly in the data set. We take the last month of the quarter figure and express it in millions euros.
- $AUM(fund\ family)_{F,t}$: Assets under management of the fund family F in quarter t. It is obtained by adding up all the AUM of funds that belong to family F.
- $LAGRET_{f,t}$: Fund f's lagged annual return in quarter t. It is computed using the change in the fund's NAV during the previous year. Formally,
$$LAGRET_{f,t} \equiv \frac{NAV_{f,t-1}}{NAV_{f,t-5}} - 1$$

[21] The only notorious difference is the inclusion of the fund's net flows as a control variable in our analysis.

- $FEES_{f,t}$: Fund f fees in quarter t. The data set provides management fees, performance fees, custodian fees, up-front fees, redemption fees and rebate fees. All the fees are reported monthly in annual percentages. In our construction of the fee variable we follow Sirri and Tufano (1998). Unfortunately, the data set does not report the fund's expense ratio. For this reason we express fees in terms of the percentage derived from the previous fees (rather than in euro terms) taking the last month of the quarter figure. Performance fees are only added up when when fund has a positive return for the year. More specifically, we compute the fee variable according to the following formula:[22]

$$FEES_{f,t} = management_{f,t} + effective\ performance_{f,t}$$
$$+custodian_{f,t} + (1/7) * (up - front_{f,t} + redemption_{f,t} - rebate_{f,t})$$

- $STYLE_{f,t}$: Fund f style of investment in quarter t. Unfortunately, the fund style is only reported from 1999Q2 onwards. For this reason, we do not include it in our leading regressions as a control variable. However, in the robustness section we show that all the results of the paper survive in the sub-sample starting in 1999Q2 where the fund's style is included as a control variable. We follow the general criteria of selecting the fund style as reported in the last month of each quarter.
- Time dummies. All the regressions include yearly dummies.

3.3.5 Descriptive statistics

The following table collects the main characteristics of the variables used in the present paper.

Table 4 already reveals some interesting phenomena. First notice that the (unconditional) average change in holdings of banking stocks is always negative (for the all funds, affiliated funds and non-affiliated funds samples) when

[22] A missing fee in the Spanish SEC data set is reported as the value 99.99. In this cases, we use the average of the minimum and maximum fees reported

we compute it using the whole sample. This means that in historical terms, Spanish funds have reduced the percentage of Spanish banks outstanding shares in their portfolios. However, when we compute the changes in holdings conditional on a drop occurring in the previous quarter, the average remains negative for all funds and non-affiliated funds, but it becomes positive for the sub-sample of affiliated funds. This means that following large price drops affiliated funds reverse the historical decreasing trend and increase their holdings of banking stocks. The reverse in the historical trend is also remarkably strong in economic terms. After a large drop in the price of the bank, affiliated funds increase their holdings in the parent bank by 0.000700%. This compares to a *decrease* of 0.000240% for non-affiliated funds and hence, implies that the change in holdings for affiliated funds is not only of the opposite sign but also almost 3 times bigger.

This evidence already points at the price support hypothesis. It is not conclusive however, as the patterns could be explained by the controls reported in the table. Further, as explained in section 3.2, the evidence could be consistent with competing hypotheses of trading after large negative returns. In the next section we address all these issues formally. The table also confirms some characteristics of affiliated funds when compared to non-affiliated funds: the typical affiliated fund is larger, less expensive, performs better in terms of (not risk-adjusted) net returns, and belongs to a larger fund family.

3.4 Price support at the individual funds level

In this section we test hypotheses 1 to 4 at the individual funds level. We start by comparing the trading of affiliated funds in their parent bank stock with the trading in that stock done by non affiliated funds. To achieve this, we estimate the following model:

$$\Delta H_{f,i,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \quad (96)$$

where $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$. It is defined as the difference between the percentage stock holding in quarter $t+1$ and the percentage stock holding in quarter t. $AFFDUMMY$ is a dummy equal to one if fund f is affiliated to bank i in quarter $t+1$. $BANKDUMMY$ is a dummy equal to 1 if fund f in quarter $t+1$ belongs to any of the banks in our sample. $CONTROLS$ is a vector of control variables including: $NETFLOW$, $LAGRET$, $FEES$, $AUM(fund)$ and $AUM(fund\ family)$. The choice of controls is motivated by Massa and Rehman (2008).

Model (3) is estimated in two alternative scenarios: the complete sample and after price drop events occurring in quarter t according to our DROP definition (equation (94)). Table 5 reports the estimation results. More specifically, column (1) reports the regression results for changes in holdings of banking stocks for all the analyzed banks and all quarters. Columns (2), (3) and (4) report the same regression results, but only for DROPS associated to $a = 0$, $a = 0.5$ and $a = 1$, respectively. The key variable for our analysis is $AFFDUMMY$. Under the price support hypothesis the coefficient of this variable should be positive ($\beta > 0$). As we can observe in Table 5, beta is not statistically significant in the whole sample (column 1), but consistent with our price support hypothesis, it becomes positive and statistically significant in all the regressions conditional on price drops. Further, our sensitivity analysis on a shows that, as expected, both β and its significance level increase when we increase a from a low level (as we capture more severe price drops), and then decrease as a becomes too large (as both the number of drops decreases and these are more likely to be associated to no price supported banks). From now on we fix $a = 0.5$ as our leading price support scenario.

Notice that the differential trading of affiliated funds is not only statistically significant, but it also bears economic significance. The coefficient on the $AFFDUMMY$ in the leading case $a = 0.5$ amounts to 0.001155%. Since the average change for the control funds is -0.000240%, this means that the change of holdings for affiliated funds is almost 5 times stronger than for the control funds.

Our previous results are consistent with the price support hypothesis, but could also be consistent with alternative hypotheses of trading after large price drops. In particular, the results could be consistent with affiliated funds increasing the holdings of all banking stocks and not just the stock of their parent bank, when this suffers a negative shock. In this case, the drop in the price of the parent bank can be interpreted as an entry signal into the whole banking sector. To test this alternative we estimate the following model of relative changes in portfolio holdings:

$$\Delta H_{f,i,t,t+1} - \Delta H_{f,ob,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \quad (97)$$

where $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$, $\Delta H_{f,ob,t,t+1}$ is the change in the holdings of fund f of other banking stocks in the same quarter $t+1$.[23] According to the hypothesis 2 stated in section 3.2, the coefficient associated to $AFFDUMMY$ must be positive and significant.

The increase in holdings reported in Table 5, can also be consistent with trading generated by a contrarian style of investment of affiliated funds. To test this alternative hypothesis we estimate the following model of relative

[23] $\Delta H_{f,ob,t,t+1}$ is calculated as if all the other banking stocks would comprise a single stock. Specifically, we define $\Delta H_{f,ob,t,t+1}$ as:

$$\Delta H_{f,ob,t,t+1} = \frac{\sum_{j \neq i}^{N} A_{f,j,t+1}}{\sum_{j \neq i}^{N} MC_{j,t+1}} - \frac{\sum_{j \neq i}^{N} A_{f,j,t}}{\sum_{j \neq i}^{N} MC_{j,t}}$$

where $A_{f,j,t}$ is fund f's holdings (in euros) of bank j stock in quarter t, $MC_{j,t}$ is Market Capitalization of j in quarter t and N stands for the number of banks in the treatment group.

changes in portfolio holdings:

$$\Delta H_{f,i,t+1} - \Delta H_{f,od,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} \\ + \gamma BANKDUMMY_{f,i,t,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \quad (98)$$

which is identical to the previous one with the only exception that we now subtract from $\Delta H_{f,i,t,t+1}$ (the change in the holdings of fund f in bank i in quarter $t+1$ after a price drop in bank i in quarter t), the change in holdings in other banking stocks *that also suffered a large price drop* in quarter t, $\Delta H_{f,od,t,t+1}$. According to hypothesis 3 in section 3.2, the $AFFDUMMY$ must be positive and statistically significant.

Finally, we analyze the anticipation hypothesis, according to which the increase in holdings after drops is part of a dynamic strategy where funds also buy in anticipation of the drop, by estimating the model:

$$\Delta H_{f,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{f,i,t-2,t-1} \\ + \gamma BANKDUMMY_{f,i,t-2,t-1} + \delta CONTROLS_{f,t-2,t-1} + \varepsilon_{f,i,t-2,t-1} \quad (99)$$

which is the same as model (3), but considering the change in holdings in the quarter preceding the quarter in which the drop takes place. According to hypothesis 4 in section 3.2, the $AFFDUMMY$ must be statistically non-significant.

The results of the estimation of models (4), (5) and (6) for $DROP_{i,t}$ when $a = 0.5$ are reported in Table 6. Columns (1), (2) and (3) collect the estimation results for the tests of the price support hypothesis against the portfolio rebalancing, the contrarian and the anticipation hypotheses, respectively. The results of the estimation confirm our hypotheses. The estimated parameter on the $AFFDUMMY$ is positive and highly significant in the test for the rebalancing hypothesis and the contrarian hypothesis and is insignificant in the test for the aniticpation hypothesis.

The analysis so far provides supportive evidence of the price support hypothesis, as outlined in hypotheses 1 to 4, at the individual funds level. Before testing the price support hypothesis against the private information hypothesis (hypothesis 5), in the next section we provide evidence in favor of the price support hypothesis at the fund families and banking groups levels.

3.5 Price support at the fund families and banking groups levels

In this section we investigate whether there is also evidence of price support at the level of fund families affiliated to banks and at the banking groups as a whole level. The methodology is similar to the one developed in the previous section with variables aggregated to the corresponding level of analysis.

3.5.1 Price support at the fund families level

All variables at the families level are computed by adding up the corresponding variables from the funds in the corresponding family. The only exceptions are the $LAGRET$ and $FEES$, which are calculated as an assets-under-management-weighted average of the fund-level variables. All fund families controlled by banks in our sample are assigned to the treatment group and the rest to the control group. We then estimate the "fund families" model equivalent to (3.4) defined as:

$$\Delta H_{F,i,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1} + \gamma BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1} + \varepsilon_{F,i,t,t+1} \quad (100)$$

where relative to a price drop in bank i in quarter t, $\Delta H_{F,i,t,t+1}$ is the change in the holdings of fund family F in the stock of bank i in quarter $t+1$. $AFFDUMMY$ is a dummy equal to one if fund family F belongs to bank i in quarter $t+1$. $BANKDUMMY$ is a dummy equal to 1 if fund family F in quarter $t+1$ belongs to any of the analyzed banks. $CONTROLS$ is a vector of control variables including the fund families: $NETFLOW$, $LAGRET$, $FEES$ and AUM ($fund\ family$).

We have also estimated the models equivalent to (5), (6) and (7) at the fund family level in a similar fashion. In the present environment, these are defined as:

$$\Delta H_{F,i,t,t+1} - \Delta H_{F,ob,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1} + \gamma\ BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1} + \varepsilon_{F,i,t,t+1} \quad (101)$$

$$\Delta H_{F,i,t+1} - \Delta H_{F,od,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1} + \gamma BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1} + \varepsilon_{F,i,t,t+1} \quad (102)$$

$$\Delta H_{F,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{F,i,t-2,t-1} + \gamma BANKDUMMY_{F,i,t-2,t-1} + \delta CONTROLS_{F,t-2,t-1} + \varepsilon_{F,i,t-2,t-1} \quad (103)$$

The results of the estimation of the four models at the fund families levels are reported in Table 7. As we can observe, all coefficients have the expected sign and significance according to the price support hypothesis. The parameter on the $AFFDUMMY$ is positive and significant in the test for the leading price support hypothesis, the rebalancing hypothesis and the contrarian hypothesis and is insignificant in the test for the anticipation hypothesis. Hence, the price support story survives when the analysis is performed at the fund families level.

3.5.2 Price support at the banking groups level

The last layer of aggregation consists on testing the price support hypothesis at the banking groups level. Bank level variables are computed as the sum of the fund family variables for the fund families controlled by any of our banks.

The only exception is the $LAGRET$ and $FEES$, which are calculated as an assets-under-management-weighted average of the fund-level variables. Fund families that do not belong to any of the banks in the treatment group are aggregated in a single non-bank or control group.

We then estimate the "banking group" model equivalent to (3.4) defined as:

$$\Delta H_{BG,i,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1} + \gamma \ BANKDUMMY_{BG,i,t,t+1} + \delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \quad (104)$$

where relative to a price drop in bank i in quarter t, $\Delta H_{BG,i,t,t+1}$ is the change in the holdings of banking group BG in the stock of bank i in quarter $t+1$. It is defined as the difference between the percentage stock holding in quarter $t+1$ and the percentage stock holding in quarter t. $AFFDUMMY$ is a dummy equal to one if banking group BG belongs to bank i in quarter $t+1$. $CONTROLS$ is a vector of control variables including the banking groups': $NETFLOW$, $LAGRET$, $FEES$ and AUM (*banking group*).

We have also estimated the models equivalent to (5), (6) and (7) at the banking groups level, which are defined in a similar fashion:

$$\Delta H_{BG,i,t,t+1} - \Delta H_{BG,ob,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1} + \gamma \ BANKDUMMY_{BG,i,t,t+1} + \delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \quad (105)$$

$$\Delta H_{BG,i,t,t+1} - \Delta H_{BG,od,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1} + \gamma BANKDUMMY_{BG,i,t,t+1} + \delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \quad (106)$$

$$\Delta H_{BG,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{BG,i,t-2,t-1} + \gamma BANKDUMMY_{BG,i,t-2,t-1} + \delta CONTROLS_{BG,t-2,t-1} + \varepsilon_{BG,i,t-2,t-1} \quad (107)$$

The results of the estimation of the four models at the banking groups level are reported in Table 8. Generally, the price support hypothesis survives this level of aggregation. The parameter on the $AFFDUMMY$ is positive and significant in the test for the leading price support hypothesis and the rebalancing hypothesis and is insignificant in the test for the anticipation hypothesis. The only news worthwhile mentioning is that at the banking groups level, the coefficient of the $AFFDUMMY$ associated to the contrarian regression (column (3)) looses significance at conventional significant levels, but still keeps the right sign.

3.6 Price support or informed trading?

Affiliated funds may increase the holdings of their parent bank stock because they have private information on the future performance of the stock. To test this alternative explanation we adopt a portfolio approach. If affiliated funds increase holdings due to private information, a trading strategy that tracks these transaction must offer outstanding performance. In particular, a trading strategy long the banking stocks after large drops must beat a strategy that buys all other banks in the economy. We analyze this strategy in the context of the most widely accepted methodology for performance evaluation, namely, a 4-factor model including the three Fama and French factor plus a fourth momentum factor.

First we analyze the case of a trading strategy that can be implemented by an outside investor who observes the portfolios of affiliated funds, which are reported quarterly. Being conservative, we assume that mutual funds portfolios become public right at the end of the quarter, rather than several months later as it occurs in practice. In this case, the outside investor can take a position in the bank stock at the end of the quarter after the price drop, when he observes that affiliated funds have increased their holdings of the parent bank stock. The long and short legs of the strategy are constructed as follows. At every quarter following the drop in the price of the bank, in which the investor observes an increase in holdings of a banking group in its parent bank stock, the investor takes a long position in a portfolio of all bank stocks for which the previous condition is met. Simultaneously the investor takes a short position in a portfolio that includes all banking stocks except those included in the long portfolio. We then compute the returns of the long, sort and long-short portfolios for several months after observing the increase in holdings in the banking groups. We analyze the cases of 3 and 6 months holding periods, creating both equally and value weighted portfolios. The process is repeated for all quarters in which there is an increase in holdings in any of the banking groups in the parent bank stock after a drop. We then estimate the 4-factor model for the periods in which a return exist for our long-short portfolio and test for a positive and significant intercept. This strategy should exhibit outstanding performance if affiliated funds increase holdings because of private good news about the bank's future performance.

The results of the estimation of the 4-factor model for both the 3 and the 6 months holding periods are reported in Table 9. While neither the long nor the short legs of the strategy in isolation exhibit any abnormal performance, the intercept for the long-short portfolio is always negative and even becomes significantly negative at conventional statistical levels in the case of value-weigthed portfolios. As this is exactly the opposite to what is expected according to the private information hypothesis, our hypothesis 5 of price support is also validated.

The previous exercise can be considered a weak test of private information as it does not include the quarter in which affiliated funds increase holdings. Strictly speaking, this test would be the right one only if most of the increase in holding by affiliated funds takes place at the end of the quarter after the drop. But it may be the case that affiliated funds implement purchases earlier in the quarter. For this reason we now repeat the previous exercise including the quarter in which the increase in holding of affiliated funds takes place. The construction of the short and long legs of the strategy are identical as before but putting the trades at the beginning of the quarter in which affiliated funds increase holdings. A very important remark is in order here. The inclusion of the quarter in which the increase in holdings takes place may capture returns in the long leg of the strategy generated by the same trades of the affiliated funds. As mentioned in the introduction, Sias et al. (2001) document a significant price impact of institutional trades. This means that the purchases of affiliated funds during the period may generate both price support (prevent the price from falling any further) and return reversals (a positive return). The latter may be captured as an abnormal return of the long leg of the previous strategy. Further, notice that there is ample evidence of short run reversals in asset returns (Fama (1965)) and that reversals seem to be especially strong after large drops (e.g. Bremer and Sweeney (1991)). This means that the long leg in the previous strategy may also capture positive abnormal returns arising from well documented return reversals not related to private information, but driven by other factors that have been found to contribute to price reversal, such as overreaction (e.g. Jegadeesh and Titman (1995a)) and market illiquidity (e.g. Jegadeesh and Titman (1995b); Cox and Peterson (1994)).

In Table 10 we report the equivalent of Table 9, but for a strategy in which the portfolios are created at the beginning of the quarter after the drops. Now, the long leg of the strategy exhibits a positive and statistically significant intercept. However, the fact that the long leg exhibit outstanding performance is indicative of the two forces mentioned above being in place. In particular the long portfolio abnormal return in the 3 month holding period case can be just a byproduct of the affiliated funds trading in the stock and the mechanical reversals. In any case, the abnormal performance of the long leg decays over time. Further, these abnormal returns wanish once we consider a long-short return. The long-short strategy does not produce abnormal returns neither in the case of 3-month nor 6-month holding periods. Hence, the price support hypothesis survives this strong test possibly contaminated by forces not related to private information.

To clear any remaining suspicions on the abnormal performance of the long portfolio in the very short run being consistent with affiliated fund mangers in possession of private information, we perform a final test. The two previous tables make clear that the abnormal returns of the long portfolio gradually die out. If affiliated fund managers were privately informed on the future of their parent bank stock they should reverse their trades in the quarter after they increase it. Not doing this implies holding the parent bank stock for a period in which the manager would have done better holding alternative banking stocks. Hence, the private information hypothesis would be corroborated if affiliated fund managers significantly decrease holding two quarters after the price drop. To test for this we estimate the following model:

$$\begin{aligned}\Delta H_{f,i,t+1,t+2} &= \alpha + \beta AFFDUMMY_{f,i,t+1,t+2} \\ &+ \gamma BANKDUMMY_{f,i,t+1,t+2} + \delta CONTROLS_{f,t+1,t+2} + \varepsilon_{f,i,t+1,t+2}\end{aligned} \tag{108}$$

Unreported results show that there is no sign that affiliated funds decrease holdings of their parent bank two quarters after the price drop. The estimated parameter on the $AFFDUMMY$ is even slightly positive 0.0000795, and is in any case insignificant with a t-statistic of 0.46. Hence, this reinforces the notion that the post drop trading by the affiliated funds is consistent with the price support hypothesis.

3.7 Robustness checks

In this section we practice several robustness checks on the previous results. All the exercises are practiced at the main level of the analysis, the individual funds level, and for the main price support model (hypothesis 1):

$$\begin{aligned}\Delta H_{f,i,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1} \\ + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1}\end{aligned} \tag{109}$$

Funds Mergers. The results of the present study could be influenced not only by mergers between banks, but also by mergers between individual funds. In the main analysis, we alleviate the effect of banking mergers by eliminating all drops around bank mergers for all the banks involved in the mergers. However, we do not control for the mergers between the funds. For this reason, we next use our hand-collected data on the funds mergers and re-estimate the main regression model while controlling for the fund mergers. Specifically, we exclude all funds from the nalysis in quarter t if they absorbed another fund in this same quarter. Altogether we identify 1952 funds mergers in the perid 2001 to 2009. However, only a few mergers overlap with our post drop quarters, so that the number of fund-quarter observations in the new regression model is reduced by less than 0.5%. Further, as reported in Column (1) in Table 11, the estimated coefficient on $AFFDUMMY$ is only marginally different from the original model and the t-statistic remains identical.

Inclusion of the drop quarter. So far, we have only analyzed the trading patterns in the quarter after the large price drop in the banking stocks. Nevertheless, looking at trading in the quarter after the drop may miss some of the purchases that take place at the time of the drop. Therefore, we now re-estimate the main model by also including the drop quarter. Specifically, we estimate the following model:

$$\Delta H_{f,i,t-1,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t-1,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t-1,t+1} \quad (110)$$

where $\Delta H_{f,i,t,t+1}$ is the average change in the holdings of fund f in the stock of bank i between the quarter $t-1$ and the quarter $t+1$. All the rest of the variables are the same as in the main model. As reported in Column (2) in Table 11, the estimated parameter on the $AFFDUMMY$ variable is approximately half as big as in the after drop regression. This implies that most of the differential trading indeed takes place after the drop quarter. However, the statistical significance of the estimated parameter on the $AFFDUMMY$ is higer when we include the drop quarter. Hence, exclusion of the drop quarter indeed misses some of the trading that takes place at the time of the drop.

Subsample analysis. The Spanish SEC provides data on some of the funds related variables, such as investment style and flows, only from 1999Q2 onwards. For this reason, we re-estimate our model for the subperiod 1999Q2 to 2009Q3 and check for sub-sample robustness as well as for robustness to the incorporation of the newly reported fund characteristics. Column (3) in Table 11 reports the estimation results with our initial set of controls. Column (4) reports the estimation results when using the officially reported net flows instead of our approximated net flows. Finally, Column (5) reports the results of the estimation of our original model (Column (3)), but now including the investment style dummies as a control variable. The reported results show that price support based trading is statistically even stronger in the more recent subsample than in the whole sample. Further, replacing approximated flows by the officially reported flows increases the statistical significance of the estimated coefficient on the $AFFDUMMY$ variable even further as the estimated parameter on the officially reported net flows is insignificant. Finally, the inclusion of the funds investment style dummies only has a marginal effect on the results.

3.8 Conclusions

We provide evidence that is consistent with trading by private entities to support the prices of traded securities. Our study is conducted in the context of the asset management industry, where a conflict of interest naturally arises because fund managers have two principals: the fund's investors and the owners of the asset management firm for whom the managers work. The interests of these two principals are not necessarily aligned, and the interest of the latter may prevail in some circumstances. The evidence indicates that, in times of turmoil, bank-affiliated funds trade in the parent bank's stock in order to support its price.

We also show that bank-affiliated funds increase their holdings of the parent bank's stock—after large drops in its price—to a significantly greater extent than do nonaffiliated funds. These trading patterns are not consistent with alternative hypotheses, such as trading to rebalance the fund's portfolio of banking stocks, contrarian trading, or trading in anticipation of a price drop. Furthermore, evidence suggests that the observed trading patterns are not driven by private information. However, the patterns are definitely consistent with trading to support the parent bank's stock price.

This paper's approach is minimalist in the sense that we analyze only one conflict of interest that can result in trading to support prices. There are many other areas of finance where similar conflicts of interest exist and hence where price support activities could naturally arise. The characteristics of banks that are relevant to our study are that they are publicly traded and own asset management firms. Funds affiliated with any other entity that features these two characteristics could be trading the entity's stock in a similar fashion. In addition, trading to support prices may go beyond these entities (banks or otherwise) and target other companies in which they have an interest—for example, companies in their stock portfolio. Outside the realm of asset management, we suggest that foundations might undertake a similar trading strategy in the interest of their funding parties. We believe that these and all similar venues are worthy of being explored by future research.

The approach of this paper is limited also because we confine ourselves to documenting evidence that is consistent with price support activities. In other words, we do not explore the impact of these activities on the returns of the supported stock, and neither do we explore the impact of price support on the performance of the funds that practice it. These issues are key areas of investigation on our research agenda.

3.9 TABLES AND FIGURES

Table 1: Law enforcement by US versus Spanish SEC

This table reports the track record on crime investigation and prosecution during the last five years of the US SEC versus the CNMV (Spanish SEC). *Panel A* collects US SEC data reported in the *SEC Performance and Accounting Report and SEC Statistics*, and *Panel B* Spanish SEC data reported in *CNMV Annual Reports*. The figures for civil cases are not 100% comparable. The Spanish SEC does not have the authority to initiate civil or penal proceedings. Hence, all cases are of an administrative nature. It just communicates the General Attorney cases that could probably be taken to trial. The US figures reflect actual civil proceedings.

	2004	*2005*	*2006*	*2007*	*2008*	*2009*
Panel A: USA						
Investigations opened	973	947	914	776	890	944
Investigations closed	639	625	868	374	1355	716
Enforcement						
Administrative proceedings	375	294	356	394	386	352
Civil Actions	264	335	218	262	285	312
Total	639	629	574	656	671	664
Cases related to asset management	90	97	95	79	88	81
Source: SEC Performance and Accounting Report and SEC Statistics						
Panel B: Spain						
Investigations opened	27	14	10	4	13	21
Investigations closed	37	21	14	9	7	9
Enforcement						
Administrative proceedings	48	36	22	15	8	14
Civil Actions	0	0	0	0	1	3
Total	48	36	22	15	9	17
Cases related to asset management	4	0	0	0	2	0
Source: CNMV (Spanish SEC) Annual Reports						

Table 2: Liquidity filter for banks

This table reports the annualized average of volume to market value ratio for all the banks traded on the Madrid SIBE that own at least one asset management company in the period 1995Q1 to 2009Q3. Formally, the standardized volume of bank i, $StdVol_i$, is defined as:

$$StdVol_i = 252 * \frac{\sum_{i=1}^{N}\left(\frac{P_{i,t}Vol_{i,t}}{MC_{i,t}}\right)}{N_i}$$

where $Vol_{i,t}$ is the turnover (in number of shares) of bank i on day t, $P_{i,t}$ is bank i stock price on day t, $MC_{i,t}$ is Market Capitalization of bank i on day t and N_i is the number of trading days for the period in which bank i was traded on the Madrid SIBE.

	StdVol (%)	*Period*
Banco Santander	140.09	1995Q1-2009Q3
Banco Bilbao Vizcaya Argentaria	136.06	1995Q1-2009Q3
Banco Popular	106.97	1995Q1-2009Q3
Argentaria	100.79	1995Q1-1999Q4
Bankinter	84.96	1995Q1-2009Q3
Banco Central Hispano	62.35	1995Q1-1999Q1
Banco Sabadell	52.81	2001Q2-2009Q3
Banco Pastor	24.51	1995Q1-2009Q3
Banco Zaragozano	16.36	1995Q1-2003Q3
Banco Guipuzcoano	8.41	1995Q1-2009Q3
Banco Atlantico	1.30	1995Q1-2004Q1
Banco Herrero	1.27	1995Q1-2000Q3
AVERAGE	61.32	

Table 3: Banks in our study

This table collects the main characteristics of the banks in the treatment group and the affiliated funds. Only funds that traded with any of the analyzed banks at least once in the period $1995Q1$ to $2009Q3$ are included. *Period in the sample* denotes the period for which the bank is included in the analysis. $Distinct\ funds$ is the number of unique funds that belong to each bank during any time in the analyzed period. $Distinct\ fund\ families$ stand for the number of unique asset management firms controlled by each bank at any point during the analyzed period. AUM is the total of assets under management for all the mutual funds that belong to each bank at the end of the sample period (or when the bank ceased to exist). AUM is expressed in billions euros.

	Period in the sample	*Funds*	*Fund families*	AUM
Banco Santander	1995Q1-2009Q3	149	5	4.9250
Banco Central Hispano*	1995Q1-1999Q1	31	2	5.0430
Banco Bilbao Vizcaya Argentaria	1995Q1-2009Q3	87	3	2.1441
Argentaria**	1995Q1-1999Q4	22	2	5.4102
Bankinter	1995Q1-2009Q3	26	1	0.7173
Banco Popular	1995Q1-2009Q3	19	2	0.6183
Banco Pastor	1995Q1-2009Q3	8	1	0.3527
Banco Sabadell***	2003Q3-2009Q3	76	3	0.6336
TOTAL		418	19	17.5223

*In 1999Q2, Banco Central Hispano and Banco Santander merged into Banco Santander Central Hispano, which was later renamed as Banco Santander.

**In 2000Q1, Argentaria and Banco Bilbao Vizcaya merged into Banco Bilbao Vizaya Argentaria.

***Banco Sabadell became publicly traded in 2001Q2. Due to a requirement of at least two years of data for the calculation of the drop variable Sabadell enters our analysis in 2003Q3.

Table 4: Descriptive statistics of fund related variables

This table collects the summary statistics for affiliated funds, non-affiliated funds and all funds together. Affiliated funds are funds that are owned by a bank in our treatment group. Non-affilliated funds are all the other funds. The period is $1995Q1$ to $2009Q3$. The observations are based on a fund-bank-quarter level. $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$. AUM $(fund)$ are total assets under management of fund f in quarter $t+1$ in millions euros. AUM $(fund\ family)$ is the sum of AUM $(fund)$ for all the funds in the fund family. $NETFLOW$ of fund f in quarter $t+1$ is calculated as:$NETFLOW_{f,t+1} \equiv AUM_{f,t+1} - AUM_{f,t} - AUM_{f,t} r_{f,t+1}$, where $r_{f,t+1}$ is net return of fund f during quarter $t+1$. $NETFLOW$ is in millions euros. $LAGRET$ is fund f's lagged annual return in quarter $t+1$. $FEES$ are total fees in quarter $t+1$ computed as: $FEES \equiv management\ fees + effective\ performance\ fees + custodian\ fees + (1/7) * (up - front\ fees + redemption\ fees - rebate\ fees)$. The statistics are based on either all bank-quarter observations, or only on the quarters proceeded by a drop in the price of the bank. Drop in the price of bank i in quarter t is defined as:

$$DROP_{i,t} = \begin{cases} 1, & \text{if } r_{i,t} - \bar{r}_{i,t-1} \leq -a\sigma_{i,t-1} \text{ AND } r^{ER}_{i,t} - \bar{r}^{ER}_{i,t-1} \leq -a\sigma^{ER}_{i,t-1} \ ; \\ 0, & \text{otherwise.} \end{cases}$$

where $a = 0.5$, $r_{i,t}$ is the return of bank i in quarter t, $\bar{r}_{i,t-1}$ and $\sigma_{i,t-1}$ are the average and the standard deviation of the quarterly return of bank i calculated over the 5 years rolling window of monthly observations. $r^{ER}_{i,t}$, $\bar{r}^{ER}_{i,t-1}$, $\sigma^{ER}_{i,t-1}$ are calculated in the same way, but in excess of market index (Madrid General SE). Only non-overlapping drops with a one quarter window are considered.

	All funds		Affiliated funds		Non-affiliated funds	
	ALL (N=323)	*DROP* (N=31)	*ALL* (N=323)	*DROP* (N=31)	*ALL* (N=323)	*DROP* (N=31)
$\Delta H_{f,i,t,t+1}$	-0.000026	-0.000199	-0.000021	0.000700	-0.000026	-0.000240
$NETFLOW$	-0.000	-0.728	-0.001	3.008	-0.000	-0.898
$LAGRET$	3.33%	-3.05%	3.89%	-5.13%	3.31%	-2.96%
$FEES$	1.80%	1.79%	1.77%	1.80%	1.80%	1.79%
AUM $(fund)$	63.660	64.347	125.853	140.476	61.205	60.890
AUM $(family)$	1866.313	1713.166	6043.858	7955.808	1701.433	1429.720
No. obs.	218,070	20,998	8,280	912	209,790	20,086

Table 5 : Price support at the individual funds level

This table reports results of the regression model:

$$\Delta H_{f,i,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1}$$
$$+ \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \qquad (Price\ support)$$

where $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$. $AFFDUMMY$ equals one if fund f belongs to bank i in quarter $t+1$, and zero otherwise. $BANKDUMMY$ equals 1 if fund f belongs to any of the analyzed banks in quarter $t+1$, and zero otherwise. The set of controls, defined in Table 3, includes net flow of the fund, $NETFLOW$, lag annual return of the fund, $LAGRET$, total fees of the fund, $FEES$, assets under management of the fund, $AUM(fund)$, and assets under management of the fund family, $AUM(fund\ family)$. The period is $1995Q1$ to $2009Q3$. Column (1) reports regression results for the complete sample (all banks and all quarters). The rest of the columns report regression results only for quarters preceded by a drop in the price of the bank. Drops are defined as in Table 4. Only non-overlapping drops with a one quarter window are considered. Columns (2), (3) and (4) report results for $a = 0$, $a = 0.5$ and $a = 1$, respectively. All regressions include year dummies. Errors are clustered at the individual funds level and $t - statistics$ are reported in brackets below the estimated coefficients. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	(1) (ALL) (N=323)	(2) (DROP) (a=0, N=45)	(3) (DROP) (a=0.5, N=31)	(4) (DROP) (a=1,N=16)
CONST.	0.000299***	0.000239**	0.000292*	0.000436**
	(4.80)	(2.22)	(1.68)	(2.35)
AFFDUMMY	0.000000	0.000565**	0.001155***	0.001098***
	(0.00)	(2.41)	(3.55)	(2.73)
BANKDUMMY	0.000146**	-0.000160	-0.000570**	-0.000640
	(2.49)	(-0.97)	(-2.13)	(-1.59)
NETFLOW	0.015388**	0.000008	0.000010*	0.000010**
	(2.25)	(1.37)	(1.94)	(2.39)
LAGRET	0.001447***	0.000621*	-0.000592	-0.001120
	(5.47)	(1.86)	(-0.68)	(-1.25)
FEES	-0.000005	-0.000056	-0.000214***	-0.000360***
	(-0.32)	(-1.20)	(-2.70)	(-3.00)
AUM (fund)	-0.000001**	-0.000001	-0.000002	-0.000007**
	(-2.23)	(-1.02)	(-1.47)	(-2.52)
AUM (family)	0.000000	0.000000	0.000000*	0.000000
	(-1.36)	(1.02)	(1.70)	(1.26)
Clustering	Fund	Fund	Fund	Fund
Year dummies	Yes	Yes	Yes	Yes
R^2	0.008	0.003	0.007	0.033
No. obs.	218,070	31,304	20,998	11,820

Table 6: Alternative hypotheses at the individual funds level

This table reports results of the regression models:

$$\Delta H_{f,i,t,t+1} - \Delta H_{f,ob,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \qquad (Rebalancing)$$

$$\Delta H_{f,i,t+1} - \Delta H_{f,od,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \qquad (Contrarian)$$

$$\Delta H_{f,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{f,i,t-2,t-1} + \gamma BANKDUMMY_{f,i,t-2,t-1} + \delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \qquad (Anticipation)$$

where relative to a price drop in bank i in quarter t, $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$, $\Delta H_{f,ob,t,t+1}$ is the change in the holdings of fund f in all other banking stocks in quarter $t+1$, $\Delta H_{f,od,t,t+1}$ is the change in the holdings of fund f in quarter $t+1$ in all other banking stocks that also suffered a price drop in quarter t and $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t-1$. $AFFDUMMY$ equals one if fund f belongs to bank i, and zero otherwise. $BANKDUMMY$ equals 1 if fund f belongs to any of the analyzed banks, and zero otherwise. The set of controls, defined in Table 3, includes net flow of the fund, $NETFLOW$, lag annual return of the fund, $LAGRET$, total fees of the fund, $FEES$, assets under management of the fund, $AUM(fund)$, and assets under management of the fund family, $AUM(fund\ family)$. The period is $1995Q1$ to $2009Q3$. Drops are defined as in Table 4. Only non-overlapping drops with a one quarter window are considered. All regressions include year dummies. Errors are clustered at the individual funds level and $t-statistics$ are reported in brackets below the estimated coefficients. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	(1) *(Rebalancing)* *(a=0.5, N=31)*	(2) *(Contrarian)* *(a=0.5, N=31)*	(3) *(Anticipation)* *(a=0.5, N=31)*
CONST.	0.000070	0.000580**	0.000271
	(0.41)	(2.43)	(1.26)
AFFDUMMY	0.001732***	0.001830***	-0.000129
	(5.14)	(3.78)	(-0.54)
BANKDUMMY	-0.000742***	-0.000873***	-0.000534*
	(-2.67)	(-3.11)	(-1.94)
NETFLOW	0.000002	-0.000002*	0.000014***
	(0.99)	(-1.77)	(2.92)
LAGRET	-0.001229	0.001317	0.000070
	(-1.37)	(1.37)	(0.15)
FEES	-0.000142*	0.000001	-0.000079
	(-1.70)	(0.01)	(-0.73)
AUM (fund)	0.000000	0.000003	-0.000005***
	(0.12)	(1.60)	(-4.47)
AUM (family)	0.000000	0.000000	0.000000***
	(0.59)	(-0.21)	(3.11)
Clustering	Fund	Fund	Fund
Year dummies	Yes	Yes	Yes
R^2	0.006	0.004	0.015
No. obs.	20,998	20,998	20,998

Table 7: Price support at the fund families level

This table reports results of the regression models:

$$\Delta H_{F,i,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1}$$
$$+\gamma BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1}$$
$$+\varepsilon_{F,i,t,t+1} \qquad (Price\ support)$$

$$\Delta H_{F,i,t,t+1} - \Delta H_{F,ob,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1}$$
$$+\gamma BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1}$$
$$+\varepsilon_{F,i,t,t+1} \qquad (Rebalancing)$$

$$\Delta H_{F,i,t+1} - \Delta H_{F,od,t,t+1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1}$$
$$+\gamma BANKDUMMY_{F,i,t,t+1} + \delta CONTROLS_{F,t,t+1}$$
$$+\varepsilon_{F,i,t,t+1} \qquad (Contrarian)$$

$$\Delta H_{F,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{F,i,t,t+1} + \gamma BANKDUMMY_{F,i,t,t+1}$$
$$+\delta CONTROLS_{F,t,t+1} + \varepsilon_{F,i,t,t+1} \qquad (Anticipation)$$

where relative to a price drop in bank i in quarter t, $\Delta H_{F,i,t,t+1}$ is the change in the holdings of fund family F in the stock of bank i in quarter $t+1$, $\Delta H_{F,ob,t,t+1}$ is the change in the holdings of fund family F in all other banking stocks in quarter $t+1$, $\Delta H_{F,od,t,t+1}$ is the change in the holdings of fund family F in quarter $t+1$ in all other banking stocks that also suffered a price drop in quarter t and $\Delta H_{F,i,t-2,t-1}$ is the change in the holdings of fund family F in the stock of bank i in quarter $t-1$. $AFFDUMMY$ equals one if fund family F belongs to bank i, and zero otherwise. $BANKDUMMY$ equals 1 if fund family F belongs to any of the analyzed banks, and zero otherwise. The set of controls includes net flow of the fund family, $NETFLOW$, lag annual return of the fund family, $LAGRET$, total fees of the fund family, $FEES$, and assets under management of the fund family AUM $(fund\ family)$. Variables at the fund families level are computed by adding up the corresponding variables from the funds in the corresponding family. The only exceptions are the $LAGRET$ and the $FEES$, which are calculated as AUM-weighted average of the fund-level variables. The period is $1995Q1$ to $2009Q3$. Drops are defined as in Table 4. Only non-overlapping drops with a one quarter window are considered. All regressions include year dummies. Errors are clustered at the fund families level and $t-statistics$ are reported in brackets below the estimated coefficients. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	(1) (Price support) (a=0.5, N=31)	(2) (Rebalancing) (a=0.5, N=31)	(3) (Contrarian) (a=0.5, N=31)	(4) (Anticipation) (a=0.5, N=31)
CONST.	0.001937	0.001153	0.004754*	0.000955
	(1.08)	(0.61)	(1.88)	(0.66)
AFFDUMMY	0.024014**	0.034467***	0.031682*	-0.000428
	(2.47)	(2.77)	(1.86)	(-0.08)
BANKDUMMY	-0.009766*	-0.011862**	-0.010241*	-0.006813
	(-1.89)	(-1.99)	(-1.76)	(-1.45)
NETFLOW	0.000012***	0.000004	-0.000002	0.000020***
	(2.90)	(1.24)	(-0.27)	(3.86)
LAGRET	0.002028	0.000063	0.009654	0.001279
	(0.30)	(0.01)	(0.89)	(0.19)
FEES	-0.000847	-0.000338	-0.000741	0.000006
	(-1.04)	(-0.35)	(-1.00)	(0.00)
AUM (family)	0.000000	0.000000	0.000002	-0.000003**
	(-0.23)	(0.32)	(1.51)	(-2.45)
Clustering	Fund family	Fund family	Fund family	Fund family
Year dummies	Yes	Yes	Yes	Yes
R^2	0.018	0.025	0.011	0.034
No. obs.	3,134	3,134	3,134	3,170

Table 8: Price support at the banking groups level

This table reports results of the regression models:

$$\Delta H_{BG,i,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1}$$

$$+\delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \qquad (\textit{Price support})$$

$$\Delta H_{BG,i,t,t+1} - \Delta H_{BG,ob,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1}$$

$$+\delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \qquad (\textit{Rebalancing})$$

$$\Delta H_{BG,i,t+1} - \Delta H_{BG,od,t,t+1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1}$$

$$+\delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \qquad (\textit{Contrarian})$$

$$\Delta H_{BG,i,t-2,t-1} = \alpha + \beta AFFDUMMY_{BG,i,t,t+1}$$

$$+\delta CONTROLS_{BG,t,t+1} + \varepsilon_{BG,i,t,t+1} \qquad (\textit{Anticipation})$$

where relative to a price drop in bank i in quarter t, $\Delta H_{BG,i,t,t+1}$ is the change in the holdings of banking group BG in the stock of bank i in quarter $t+1$, $\Delta H_{BG,ob,t,t+1}$ is the change in the holdings of banking group BG in all other banking stocks in quarter $t+1$, $\Delta H_{BG,od,t,t+1}$ is the change in the holdings of banking group BG in quarter $t+1$ in all other banking stocks that also suffered a price drop in quarter t. $\Delta H_{BG,i,t-2,t-1}$ is the change in the holdings of banking group BG in the stock of bank i in quarter $t-1$. $AFFDUMMY$ equals one if banking group BG belongs to bank i, and zero otherwise. The set of controls includes net flow of the banking group, $NETFLOW$, lag annual return of the banking group, $LAGRET$, total fees of the banking group, $FEES$, and assets under management of the banking group AUM (*banking group*). Variables at the banking groups level are computed by adding up the corresponding variables from the fund family in the corresponding banking group. The only exceptions are the $LAGRET$ and the $FEES$, which are calculated as AUM-weighted average of the fund-level variables. Fund families that do not belong to any of the banks in the treatment group are aggregated in a single non-bank or control group. The period is $1995Q1$ to $2009Q3$. Drops are defined as in Table 4. Only non-overlapping drops with a one quarter window are considered. All regressions include year dummies. Errors are clustered at the banking groups level and $t-statistics$ are reported in brackets below the estimated coefficients. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	(1) *(Price support)* *(a=0.5, N=31)*	(2) *(Rebalancing)* *(a=0.5, N=31)*	(3) *(Contrarian)* *(a=0.5, N=31)*	(4) *(Anticipation)* *(a=0.5, N=31)*
CONST.	0.007454	0.004526	0.082877	0.018192
	(0.18)	(0.13)	(1.52)	(0.41)
AFFDUMMY	0.035391**	0.048546***	0.034146	-0.000593
	(2.51)	(3.33)	(1.31)	(-0.06)
NETFLOW	0.000017	0.000006	0.000044	0.000007
	(1.50)	(0.82)	(1.13)	(1.46)
LAGRET	-0.003977	-0.100835	-0.007198	0.132945
	(-0.02)	(-0.89)	(-0.02)	(0.61)
FEES	-0.005643	-0.016957	-0.052050***	-0.003269
	(-0.33)	(-1.03)	(-3.94)	(-0.20)
AUM (ba. group)	0.000003***	0.000000	0.000008***	0.000005***
	(-5.85)	(-0.51)	(5.56)	(-10.37)
Clustering	Bank	Bank	Bank	Bank
Year dummies	Yes	Yes	Yes	Yes
R^2	0.088	0.078	0.056	0.188
No. obs.	223	223	223	225

Table 9: Price support versus private information

This table reports performance evaluation results for portfolios of banking stocks. At every quarter following the drop in the price of the bank and simultaneous increase in holdings of an affiliated banking group in its parent bank stock we form a $Long$ portfolio of all bank stocks for which the previous condition is met. Simultaneously we form a $Short$ portfolio that takes a long position in a portfolio that includes all banking stocks except those included in the $Long$ portfolio. We then compute the returns of both portfolios for either 3 months or 6 months after observing the increase in holdings in the banking groups. The process is repeated for all quarters preceded by a drop in the price of the bank and a simultaneous increase in holdings of the banking group in the parent bank stock. We identify 16 quarter-bank observations that satisfy this criteria. Depending on the holding period of the portfolios, this construction gives rise to a time series of either 36 or 72 monthly returns. We then estimate the 4-factor model for the periods in which a return exists for the portfolios:

$$r_t = \alpha + \beta Mkt_t + \gamma SMB_t + \delta HML_t + \theta MOM_t + \varepsilon_t$$

where r_t is either excess return of the $Long$ portfolio, excess return of the $Short$ portfolio, or the difference between the return of the $Long$ and of the $Short$ portfolio ($L - S$), MKT_t is the market return (Madrid SIBE) in excess of the risk-free rate, SMB_t is the return on the small minus big factor, HML_t is the return on the high minus low factor and MOM_t is the return on the Momentum factor. Portfolio returns are either equally- or market capitalization-weighted. The period is $1995Q1$ to $2009Q4$. $OLS\ t-statistics$ are reported in brackets. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	3 Months					
	Equally-weighted			*Value-weighted*		
	Long	*Short*	*L-S*	*Long*	*Short*	*L-S*
CONST.	0.001509	0.007959	-0.006450	0.002498	0.016381	-0.013883*
	(0.13)	(0.71)	(-0.79)	(0.21)	(1.21)	(-1.72)
MKT	0.624388***	0.599185***	0.025203	0.587313***	0.553698**	0.033615
	(3.24)	(3.30)	(0.19)	(3.10)	(2.52)	(0.26)
SMB	-0.582815*	-0.183517	-0.399297*	-0.587813*	-0.555771	-0.032042
	(-1.94)	(-0.64)	(-1.95)	(-1.98)	(-1.62)	(-0.15)
HML	-1.259515***	-0.835688**	-0.423826*	-1.137103***	-1.004345**	-0.132758
	(-3.61)	(-2.54)	(-1.78)	(-3.31)	(-2.52)	(-0.55)
MOM.	-0.091698	0.044423	-0.136121	-0.207260	0.059276	-0.266536
	(-0.37)	(0.20)	(-0,82)	(-0.86)	(0.22)	(-1.62)
R^2	0.453	0.378	0.138	0.432	0.312	0.116
No. obs.	36	36	36	36	36	36

	6 Months					
	Equally-weighted			*Value-weighted*		
	Long	*Short*	*L-S*	*Long*	*Short*	*L-S*
CONST.	-0.003164	0.008768	-0.011932	-0.001677	0.013225	-0.014901**
	(-0.33)	(1.23)	(-1.65)	(-0.18)	(1.66)	(-2.31)
MKT	0.591556***	0.418905***	0.172651	0.607627***	0.401234***	0.206393*
	(3.63)	(3.33)	(1.36)	(3.92)	(2.86)	(1.82)
SMB	-0.399976*	-0.235476	-0.164500	-0.399879*	-0.437978**	0.038099
	(-1.76)	(-1.34)	(-0.92)	(-1.85)	(-2.24)	(0.24)
HML	-0.669278***	-0.575670***	-0.093608	-0.608040**	-0.691770***	0.083730
	(-2.67)	(-2.97)	(-0.47)	(-2.55)	(-3.21)	(0.48)
MOM.	-0.186666	-0.152465	-0.034201	-0.251407	-0.119259	-0.132148
	(-0.86)	(-0.91)	(-0.20)	(-1.22)	(-0.64)	(-0.88)
R^2	0.311	0.302	0.053	0.334	0.292	0.077
No. obs.	72	72	72	72	72	72

Table 10: Price support versus private information: Including contemporaneous quarter

This table reports performance evaluation results for portfolios of banking stocks. At every quarter following the drop in the price of the bank and simultaneous increase in holdings of an affiliated banking group in its parent bank stock we form a *Long* portfolio of all bank stocks for which the previous condition is met. Simultaneously we form a *Short* portfolio that takes a long position in a portfolio that includes all banking stocks except those included in the *Long* portfolio. We then compute the returns of both portfolios either for the quarter when we observe the increase in holdings in the banking groups (contemporaneous quarter) or for the contemporaneous quarter plus 3 months. The process is repeated for all quarters preceded by a drop in the price of the bank and a simultaneous increase in holdings of the banking group in the parent bank stock. We identify 16 quarter-bank observations that satisfy this criteria. Depending on the holding period of the portfolios, this construction gives rise to a time series of either 36 or 72 monthly returns. We then estimate the 4-factor model for the periods in which a return exists for the portfolios:

$$r_t = \alpha + \beta Mkt_t + \gamma SMB_t + \delta HML_t + \theta MOM_t + \varepsilon_t$$

where r_t is either excess return of the *Long* portfolio, excess return of the *Short* portfolio, or the difference between the return of the *Long* and of the *Short* portfolio ($L - S$), MKT_t is the market return (Madrid SIBE) in excess of the risk-free rate, SMB_t is the return on the small minus big factor, HML_t is the return on the high minus low factor and MOM_t is the return on the Momentum factor. Portfolio returns are either equally- or market capitalization-weighted. The period is 1995Q1 to 2009Q4. $OLS\ t - statistics$ are reported in brackets. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	3 Months					
	Equally-weighted			*Value-weighted*		
	Long	*Short*	*L-S*	*Long*	*Short*	*L-S*
CONST.	0.022014***	0.008286	0.013728	0.020530**	0.010752	0.009778
	(2.67)	(0.81)	(1.27)	(2.67)	(0.86)	(0.89)
MKT	0.625625***	0.667135***	-0.041510	0.700392***	0.649943***	0.050450
	(4.48)	(3.87)	(-0.22)	(5.39)	(3.08)	(0.27)
SMB	-0.424386	-0.618364*	0.193978	-0.408909*	-1.014105**	0.605196*
	(-1.67)	(-1.97)	(0.59)	(-1.73)	(-2.65)	(1.81)
HML	-0.642561**	-0.574477*	-0.068084	-0.581991**	-0.925641**	0.343650
	(-2.47)	(-1.79)	(-0.19)	(-2.40)	(-2.36)	(1.00)
MOM.	-0.816165***	-0.857073***	0.040908	-0.749181***	-1.010036***	0.260855
	(-4.91)	(-4.18)	(0.19)	(-4.84)	(-4.03)	(1.19)
R^2	0.669	0.595	0.025	0.7097	0.5721	0.1198
No. obs.	36	36	36	36	36	36

	6 Months					
	Equally-weighted			*Value-weighted*		
	Long	*Short*	*L-S*	*Long*	*Short*	*L-S*
CONST.	0.014969**	0.013151*	0.001819	0.013957**	0.019360**	-0.005402
	(2.02)	(1.71)	(0.28)	(2.00)	(2.09)	(-0.78)
MKT	0.496644***	0.537022***	-0.040378	0.546473***	0.512580***	0.033894
	(4.14)	(4.30)	(-0.37)	(4.83)	(3.42)	(0.31)
SMB	-0.680982***	-0.483666**	-0.197316	-0.648867***	-0.872389***	0.223522
	(-3.37)	(-2.30)	(-1.09)	(-3.41)	(-3.45)	(1.20)
HML	-0.904368***	-0.658362***	-0.246006	-0.842884***	-0.922435***	0.079551
	(-4.02)	(-2.81)	(-1.22)	(-3.97)	(-3.28)	(0.38)
MOM.	-0.538867***	-0.476701***	-0.062166	-0.539096***	-0.524732***	-0.014364
	(-3.64)	(-3.09)	(-0.46)	(-3.86)	(-2.83)	(-0.10)
R^2	0.464	0.403	0.026	0.504	0.377	0.026
No. obs.	72	72	72	72	72	72

Table 11 : Robustness checks

This table reports variations to the regression model:

$$\Delta H_{f,i,t,t+1} = \alpha + \beta AFFDUMMY_{f,i,t,t+1} + \gamma BANKDUMMY_{f,i,t,t+1}$$
$$+\delta CONTROLS_{f,t,t+1} + \varepsilon_{f,i,t,t+1} \qquad (Price\ support)$$

where relative to a price drop in bank i in quarter t, $\Delta H_{f,i,t,t+1}$ is the change in the holdings of fund f in the stock of bank i in quarter $t+1$. $AFFDUMMY$ equals one if fund f belongs to bank i in quarter $t+1$, and zero otherwise. $BANKDUMMY$ equals 1 if fund f belongs to any of the analyzed banks in quarter $t+1$, and zero otherwise. The set of controls, defined in Table 3, includes net flow of the fund, $NETFLOW$, lag annual return of the fund, $LAGRET$, total fees of the fund, $FEES$, assets under management of the fund, $AUM(fund)$, and assets under management of the fund family, $AUM(fund\ family)$. Drops are defined as in Table 4. Only non-overlapping drops with a one quarter window are considered. Column (1) reports regression results for the period 1995Q1 to 2009Q3 excluding funds that absorbed any other fund (in a one quarter window). Column (2) reports regression results for the period 1995Q1 to 2009Q3 using as a dependent variable $\Delta H_{f,i,t-1,t+1}$ instead of $\Delta H_{f,i,t,t+1}$, where $\Delta H_{f,i,t-1,t+1}$ is the average change in the holdings of fund f in the stock of bank i between quarter $t-1$ and quarter $t+1$. Column (3) reports regression results with the original set of controls for the sub-period 1999Q2 to 2009Q3. Column (4) reports the same results as column (3), but using officially reported flows, $NETFLOW\ (Off.)$, instead of approximated flows. Finally, column (5) reports the same results as Column (3), but including investment style dummies. All regressions include year dummies. Errors are clustered at the individual funds level and $t-statistics$ are reported in brackets below the estimated coefficients. One, two and three asterisks (*,**,***) denote significance at the 10%, 5% and 1% levels, respectively.

	(1) (a=0.5, N=31)	(2) (a=0.5, N=31)	(3) (a=0.5, N=23)	(4) (a=0.5, N=23)	(5) (a=0.5, N=23)
CONST.	0.000291*	0.000292*	0.000232**	0.000224**	0.000391***
	(1.67)	(1.78)	(2.24)	(2.18)	(2.63)
AFFDUMMY	0.001160***	0.000662***	0.000997***	0.001036***	0.000994***
	(3.55)	(3.77)	(4.28)	(4.45)	(4.27)
BANKDUMMY	-0.000573**	-0.000039	0.000002	-0.000007	0.000007
	(-2.13)	(-0.32)	(0.01)	(-0.04)	(0.04)
NETFLOW	0.000010*	0.000008	0.000006**		0.000007**
	(1.93)	(1.47)	(2.24)		(2.23)
NETFLOW (Off.)				0.000000	
				(0.08)	
LAGRET	-0.000592	0.001271**	0.000197	0.000213	0.000164
	(-0.67)	(1.96)	(1.02)	(1.09)	(0.67)
FEES	-0.000214***	-0.000040	-0.000139***	-0.000141***	-0.000110**
	(-2.70)	(-0.79)	(-2.99)	(-3.02)	(-2.45)
AUM (fund)	-0.000002	-0.000001	-0.000001	-0.000001	-0.000001
	(-1.47)	(-1.26)	(-1.13)	(-0.60)	(-1.14)
AUM (fund family)	0.000000*	-0.000000**	0.000000	0.000000	0.000000
	(1.70)	(-2.38)	(0.00)	(-0.32)	(-0.03)
Clustering	Fund	Fund	Fund	Fund	Fund
Year dummies	Yes	Yes	Yes	Yes	Yes
Style dummies	No	No	No	No	Yes
R^2	0.007	0.014	0.011	0.008	0.012
No. obs.	20,904	20,998	18,118	18,118	18,118

Figure 1: AUM and number of funds

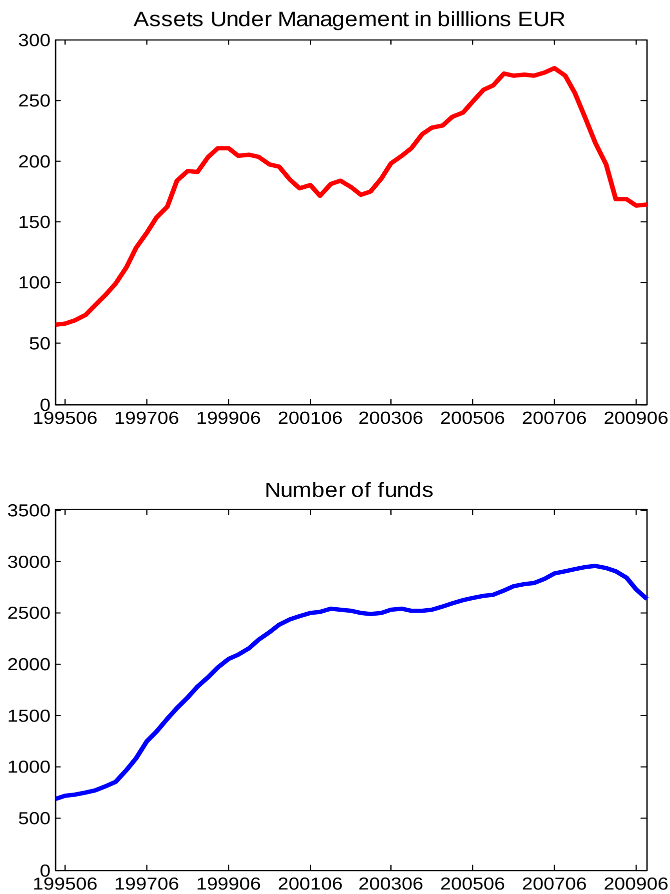

The figures plots the assets under management (AUM) and the number of funds for all the mutual funds registered with the Spanish SEC in the period $1995Q1$ to $2009Q2$.

Figure 2: Affiliated versus non-affiliated funds

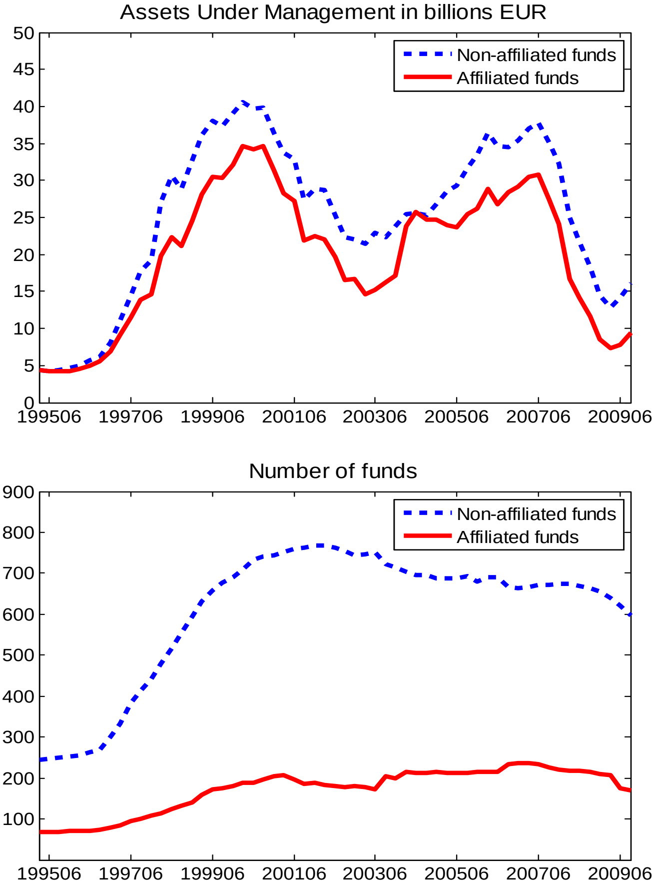

The figures plot the assets under management (AUM) and the number of funds in the affiliated and the non-affiliated fund categories. Affiliated funds are funds that belong to any of the 8 analyzed banks. Non-affiliated funds are those that either belong to non-banks or other banks. Only funds that traded with any of the analyzed banks at least once in the period $1995Q1$ to $2009Q3$ are included.

Bibliography

1. Ait-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset markets, *Journal of Finance* 53, 499-547.

2. Allen, Franklin and Douglas Gale, 1994, *Financial Innovation and Risk Sharing* (MIT Press, Cambridge, MA).

3. Amihud, Yakov and Clifford M. Hurvich, 2004, Predictive regressions: A reduced-bias estimation method, *Journal of Financial and Quantitative Analysis* 39, 813-841.

4. Anders, George, 1982, Options trading at expiration might influence prices of underlying stocks, studies indicate, *The Wall Street Journal*, April 15, 55.

5. Ang, Andrew and Geert Bekaert, 2007, Stock return predictability: Is it there?, *Review of Financial Studies* 20, 651-707.

6. Augen, Jeff, 2009, *Trading options at expiration: Strategies and models for winning the endgame* (Pearson Education, Upper Saddle River, NJ).

7. Avellaneda, Marco, and Michael D. Lipkin, 2003, A market-induced mechanism for stock pinning, *Quantitative Finance* 3, 417-425.

8. Bansal, Ravi and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

9. Bhattacharya, Uptal and Hazem Daouk, 2002, The world price of insider trading, *Journal of Finance* 57, 75-108.

10. Binsbergen, Jules H. van and Ralph Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance*, forthcoming.

11. Binsbergen, Jules H. van, Michael W. Brandt and Ralph Koijen, 2010, On the timing and pricing of dividends, *Unpublished manuscript.*

12. Bollerslev, Tim, George Tauchen and Hao Zhou, 2009. Expected stock returns and variance risk premia, *Review of Financial Studies 22*, 4463-4492.

13. Boudoukh, Jacob, Roni Michaely and Matthew Richardson, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877-915.

14. Boudoukh, Jacob, Matthew Richardson and Robert F., Whitelaw, 2008, The myth of long-horizon predictability, *Review of Financial Studies* 21, 1577-1605.

15. Bremer, Marc and Richard J. Sweeney, 1991, The reversal of large stock-price decreases, *Journal of Finance* 46, 747-754.

16. Bucks, Brian K., Arthur B. Kennickell, Traci L. Mach, and Kevin B. Moore, 2009, Changes in U.S. family finances from 2004 to 2007: Evidence from the survey of consumer finances, Federal Reserve Bulletin.

17. Campbell, John. Y., 1991, A variance decomposition for stock returns, *Economic Journal 101*, 157-179.

18. Campbell, John, Y. and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195-227.

19. Campbell, John.Y. and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509-1531.

20. Chevalier, Judith and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *The Journal of Political Economy* 105, 1167-1200.

21. Chiang, Chin-Han, 2010, Stock returns on option expiration dates, Unpublished manuscript.

22. Cochrane, John. H., 2005, *Asset pricing, revised ed.* Princeton, NJ: Princeton University Press.

23. Cochrane, John H., 2008a, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533-1575.

24. Cochrane, John H., 2008b, State-space vs. VAR models for stock returns, Unpublished manuscript.

25. Cox, Don R. and David R. Peterson, 1994, Stock returns following large one-day declines; Evidence on short-term reversals and longer-term performance, *Journal of Finance* 49, 255-267.

26. Durnev, Art and E. Han Kim, 2005, To steal or not to steal: Firm attributes, legal environment, and valuation, *Journal of Finance* 60, 1461-1493.

27. Ellis, Katrina, Roni Michaeli and Maureen O'Hara, 2000, When the underwriter is the market maker: An examination of trading in the IPO aftermarket, *Journal of Finance* 55, 1039-1074.

28. Fama, Eugene F., 1965, The behavior of stock-market prices, *Journal of Business* 38, 34-105.

29. Fama, Eugene F. and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3-27.

30. Favero, Carlo A., Arie Gozluklu and Andrea Tamoni, 2010, Demographic trends, the dividend-price ratio and the predictability of long-run stock market returns, *Journal of Financial and Quantitative Analysis*, forthcoming.

31. Ferreira, Miguel A. and Santa-Clara, Pedro, 2010, Forecasting stock market returns: The sum of the parts is more than the whole, *Journal of Financial Economics*, forthcoming.

32. Garleanu, Nicolae, Lasse H. Pedersen and Allen. M. Poteshman, 2009, Demand-based option pricing, *Review of Financial Studies* 22, 4259-4299.

33. Goyal, Amit and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455-1508.

34. Hansen, Lars P. and Robert J. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829-853.

35. Hodrick, Robert. J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357-386.

36. Investment Adviser Association, 2008, Evolution on revolution: A profile of the investment adviser profession, available at www.investmentadviser.org.

37. Jeannin, Marc, Giulia Iori, and David Samuel, 2007, Modeling stock pinning, *Physica* A 376, 467-479.

38. Jegadeesh, Narasimhan and Sheridan Titman, 1995a, Overreaction, delayed reaction, and contrarian profits, *Review of Financial Studies* 8, 973-993.

39. Jegadeesh, Narasimhan and Sheridan Titman, 1995b, Short-horizon return reversals and the bid-ask spread, *Journal of Financial Intermediation* 4, 116-132.

40. Jenkinson, Tim and Alexander Ljungqvist, 2001, *Going Public: the Theory and Evidence on How Companies Raise Equity Finance*, second edition (Oxford University Press, oxford, NY).

41. Johnson, Norman L., Samuel Kotz, and Adrienne W. Kemp, 1992. *Univariate discrete distributions*, 2nd ed. (John Wiley, New York, NY).

42. Kamara, Avraham and Thomas Miller Jr., 1995, Daily and intradaily tests of European put-call parity, *Journal of Financial and Quantitative Analysis* 30, 519-539.

43. Kolev, Gueorgui I., 2008, Forecasting aggregate stock returns using the number of initial public offerings as a predictor, *Economics Bulletin* 7, 1-8.

44. Krishnan, Hari, and Izzy. Nelken, 2001, The effect of stock pinning upon option prices, *RISK*, December 2001, S17-S20.

45. Lacerda, Filipe and Pedro Santa-Clara, 2010, Forecasting dividend growth to better predict returns, Unpublished manuscript.

46. Lakonishok, Josef, Andrei Shleifer, Richard Thaler and Robert Vishny, 1991, Window dressing by pension fund managers, *A.E.R. Papers and Proceedings* 8, 227-31.

47. La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer and Robert Vishny, 2000, Investor protection and corporate governance, *Journal of Financial Economics* 58, 3-27.

48. Lettau, Martin and Sydney Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815-849.

49. Lettau, Martin and Sydney Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics*, 76, 583-626.

50. Lettau, Martin and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607-1652.

51. Lewellen, Jonathan W., 2004, Predicting returns with financial ratios, *Journal of Financial Economics* 74, 209-235.

52. Lewellen, Katharina, 2006, Risk, reputation, and IPO price support, *Journal of Finance* 61, 613-653.

53. Lioui, Abraham, 2006, Black-Scholes-Merton revisited under stochastic dividend yields, *Journal of Futures Markets* 26, 703-732.

54. Lioui, Abraham and Jesper Rangvid, 2009, Stock return predictability in a monetary economy, Unpublished manuscript.

55. Marin, Jose M., and Jacques Olivier, 2008, The dog that did not bark: Insider trading and crashes, *Journal of Finance* 63, 2429-2476.

56. Martens, Martin and Dick van Dijk, 2007, Measuring volatility with the realized range, *Journal of Econometrics* 138, 181-207.

57. Massa, Massimo and Zahid Rehman, 2008, Information flows within financial conglomerates: Evidence from the banks-mutual funds relation, *Journal of Financial Economics* 89, 288-306.

58. Menzly, Lior, Tano Santos and Pietro Veronesi, 2004, Understanding predicatability, *Journal of Political Economy* 112, 1-47.

59. Mola, Simona and Massimo Guidolin, 2009, Affiliated mutual funds and analyst optimism, *Journal of Financial Economics* 93, 108-137.

60. Naranjo, Lorenzo, 2009, Implied interest rates in a market with frictions. Unpublished manuscript.

61. Nayak, Suhas, 2007, An equilibrium-based model of stock-pinning, *International Journal of Theoretical and Applied Finance* 10, 535-555.

62. Newey, Whitney K. and Kenneth D. West, 1987, A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

63. Ni, Sophie X., Neil D. Pearson and Allen. M. Poteshman, 2005, Stock price clustering on option expiration dates, *Journal of Financial Economics* 78, 49-87.

64. Parkinson, Michael, 1980, The extreme value method for estimating the variance of the rate of return, *Journal of Business* 53, 61–65.

65. Pearson, Neil D., Allen M. Poteshman and Joshua White, 2007, Does option trading have a pervasive impact on underlying stock prices?, Unpublished manuscript.

66. Prabhala, N. R., and Manju Puri, 1998, How does underwriter price support affect IPOs: Empirical evidence. Unpublished manuscript.

67. Ritter Jay and Donghang Zhang, 2007, Affiliated mutual funds and the allocation of initial public offerings, *Journal of Financial Economics* 86, 337-368.

68. Rytchkov, Oleg, 2008, Filtering out expected dividends and expected returns, Unpublished manuscript.

69. Ruud, Judith, 1993, Underwriter price support and the IPO underpricing puzzle, *Journal of Financial Economics* 34, 135-151.

70. Savickas, Robert and Arthur J. Wilson, 2003, On inferring the direction of option trades, *Journal of Financial and Quantitative Analysis* 38, 881 -902.

71. Schultz, Paul and Mir Zaman, 1994, Aftermarket support and underpricing of initial public offerings, *Journal of Financial Economics* 35, 199-219.

72. Schwarz, Adam L., Bonnie F. Van Ness, and Robert A. Van Ness, 2004, Clustering in the futures market: Evidence from the S&P 500 futures contract, *Journal of Futures Markets* 24, 413-428.

73. Sias, Richard W., Laura T. Starks and Sheridan Titman, 2001, The price impact of institutional trading, Unpublished manuscript.

74. Sirri, Erik and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589-1622.

75. Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-42.

76. Stoll, Hans R., and Robert E. Whaley, 1991, Expiration day effects: What has changed?, *Financial Analysts Journal* 47, 58-72.

77. Valkanov, Rossen, 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics* 68, 201-32.

78. Wurgler, Jeffrey, 2000, Financial markets and the allocation of capital, *Journal of Financial Economics* 58, 187-214.

79. Zhou, Hao, 2010, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, Unpublished manuscript.