

Option Implied Dividends

Jac. Kragt*

June 2, 2017

Abstract

I determine the valuation of future dividends for US companies as implied by option prices. This is the first paper in which the early exercise premium included in these prices is explicitly accounted for as part of finding these dividend valuations. From these implied dividend data, I build company-specific term structures of dividend growth relative to actual dividends. These term structures show substantial variation in slope over time as well as in the cross-section. Implied dividends predict actual dividends, particularly upward dividend changes. But if a dividend cut is correctly predicted by implied dividends, the average -2.3% stock price response to the announcement becomes negligible.

*j.c.kragt@uvt.nl, Tilburg University. I thank Joost Driessen, Frank de Jong, Rik Frehen and participants at the OptionMetrics 2016 conference for their helpful comments.

1 Introduction

The main component of a stock's value is its promise to future dividends. The present value identity for stock prices lays out the precise terms of the relationship, yet an adequate empirical connection between stocks and dividends is hard to establish. Shiller (1981) documents that the variation of dividends over time is far exceeded by the volatility of stock prices. This and other authors work with actual dividends as a proxy for future dividends. In view of these weak findings, could it be that actual dividends are not the best available representation for application in the present value of future dividends? I make the case to investigate the market valuation of future dividends for finding information about share prices.

Central to this argument is that the present value identity attributes value from the *discounted expectation* of dividends, and that the valuation of such expectations are more accurately measured than that future dividends are modeled. At the level of stock indices this is argued by several authors, summarized in Binsbergen and Koijen (2016). This literature circles around a particular product in financial markets: dividend derivatives. These products allow market participants to trade the dividends thrown off by the constituent companies of a stock index usually several years into the future. Their pricing thus provides an excellent gauge of future dividend valuations in the aggregate.

Addressing cross sectional asset pricing puzzles by means of dividend valuations requires access to priced dividends for individual companies. Unfortunately, such products are few and far between. Dividend futures for individual European stocks exist in small numbers since 2010, but they are scarcely traded and their liquidity is limited. In the US market, publicly traded dividend futures were introduced for the S&P 500 in 2016, but do not exist for individual companies. In order to find such pricing information, I propose to measure implied dividends for individual companies from the prices of stock options. This paper is the first to display such an approach and to apply its result in asset pricing.

The rationale to investigate option prices for this purpose is that they are a function of the assumptions made by market participants about the volatility of the underlying stock and about the dividends it will throw off. Options are the only instrument in financial markets which contain such forward-looking information about dividends for individual companies in meaningful numbers. They also exist for several expiry dates, which gives insight into a term structure of future dividends.

The case to imply dividend valuations paid by all constituent companies of the S&P 500 index from option prices has been made by van Binsbergen, Brandt and Koijen (2012). These authors apply the put-call parity (PCP) to index options for finding the implied dividends. This methodology presupposes that payoffs are the same for stock plus put

option and strike plus call option throughout the life of the options, which is true for options that are not exercisable before their expiry. For index options, the direct application of PCP is an appropriate methodology, since they are indeed such European style. Options of individual US stocks, however, are American style and exercisable before their expiry date. The embedded early exercise premium introduces a deviation in the PCP which, without proper adjustment, would disqualify it for the purpose of finding implied dividends directly from the prices of American options. I propose a methodology to circumvent this issue by modeling option prices and solving for implied volatility and implied dividends at the same time, using a binomial tree model. The core of this procedure is the minimization of the squared difference between the model prices and the market prices of put-call option pairs in a single step.

The procedure provides a database of implied dividends for each firm out to three quarters ahead of the observation date. The data set is as large as the set of companies on which stocks options are traded. A subset of companies have stock options expiring in excess of two years (LEAPS), with an according horizon for implied dividends¹.

The data set bears out that implied dividends vary substantially over time and in the cross-section. Nevertheless they track actual dividends well, despite that these observable data do not play a part in calculating implied dividends. Over time, they are on average valued lower than actual dividends, except for during bear markets. The average actual dividend yield is 56 bps per quarter. Expressed as an implied quarterly dividend yield, the average term structure shows a decline of a little over 1 bps per quarter. Implied dividends in the highest cross-sectional quartile are often substantially above actual dividends and in the lowest they are mostly below. For longer horizons this dispersion is smaller than for shorter horizons.

Implied dividends show that market participants anticipate changes in dividends. In fact, 54% to 57% of dividend increases are correctly predicted by implied dividends whereas dividend cuts are not predicted by implied dividends (close to 50 %). More generic tests also reveal that upward changes in dividends in particular are predicted by implied dividends. Future dividend raises are better predicted with higher sensitivity to implied dividends for longer horizons than dividend cuts.

These predictions are, however, obfuscated by the nature of dividend changes. Actual dividends are often left unchanged for an extended period of time by company management

¹As a check on the robustness of minimization procedure, I include stock loan fees as part of the return to shareholders. Such fees can be regarded as a source of income to be derived from owning a stock in addition to its dividend. By expanding the methodology to find implied dividends, stock loan fees are implied simultaneously but separately from implied dividends. Implied fees correspond to actual fees, while showing a negative slope in their term structure (see the Appendix).

pursuing a dividend smoothing policy and the prevalence of such no-changes in dividends is vast. An estimation methodology that takes account of this feature of the data orders them into a separate category. Actual dividend changes are categorized into upward small and large changes, negative changes and zero changes. An ordered Probit estimation then reveals the probabilities of a raise, a cut or a no-change occurring for a given implied dividend being higher or lower than the dividend that was last paid. Although significant, the material effect of implied dividends on a change in such probabilities is not large. The Probit results show that the marginal effect of implied dividends moving from below to above actual dividends predicting an increase amounts to at most 2%.

Just as implied dividends predict changes in actual dividends, the effect of an announcement of a dividend change on stock prices is also affected if implied dividends correctly predict the change. Larkin, Leary and Michaely (2016) note that stock prices typically fall around an announcement to cut dividends and by about twice as much as in case of a dividend raise. They argue that market participants are less surprised by a dividend change if a company has cut dividends before and they show that this reduces the stock price response. This mitigated element of surprise plays a part once implied dividends are included in testing the announcement effect as well. Although dividend cuts are not well-predicted by implied dividends, the three-day stock price response to a dividend cut is reduced substantially if it is correctly predicted by implied dividends. Any response reduction in case of dividend raises, however, is invisible.

This paper is organized as follows. The subsequent section sets out the procedure to imply dividend valuations from US stock options. This involves a discussion of the minimization procedure and the criteria to calculate the data and is followed by a description of the dividends as implied. An example of the data and the inclusion of stock loan fees in the minimization procedure concludes the section. In the subsequent section these data are deployed in several tests for their predictive power, both for actual dividends and the stock price response to announcements to change dividends. The final section contains the conclusions.

2 Implying dividends from option prices

The market valuation of dividends is found from the relationship between stock prices and futures. Pricing a future contract $F_{t,n}$ that matures at time n needs to take account of the time value of money r and for the fact that the buyer of a future does not receive the cash distributed by the asset – which is its dividend $P_{t,n}$ paid up to n in the case of a stock S_t

$$P_{t,n} = S_t - e^{-r_{t,n}} F_{t,n}. \quad (1)$$

As futures of individual stocks are not sufficiently available for obtaining an empirically meaningful data set, I use options instead. Put-call parity (PCP) is a model free relationship which enables to find dividends as implied by option prices. The intuition is that the combination of opposite positions in otherwise equal call and put options² creates the payoff of a future. This pricing relationship is given by the equality of the difference between the price of a European call option $c_{t,n}$ and put option $p_{t,n}$ to the discounted difference between option strike K and future price $F_{t,n}$

$$c_{t,n} - p_{t,n} = (F_{t,n} - K)e^{-r_{t,n}}. \quad (2)$$

Combining (1) and (2) describes the PCP as I apply it, where the price of the future is replaced by option prices:

$$P_{t,n} = S_t + p_{t,n} - c_{t,n} - Ke^{-r_{t,n}}. \quad (3)$$

This depiction is clear cut for European options as the value of dividends can be traced back directly to the variables on the right hand side, which are all observable. But American style options on US stocks allow for exercise earlier than the expiry date resulting in a strictly positive early exercise premium. There is no need to deviate from the principle of the PCP approach, but a correction for the premium that the early exercise presents is required to apply (3).

²Same underlying stock, strike and expiry date.

2.1 The model

The standard for pricing American options applies a binomial tree, introduced by Cox, Ross and Rubinstein (1979). The binomial CRR model builds up stock prices and option prices by up and down scenarios in constant time steps as the horizon extends into the future until the expiry date of the option. The put and call option prices in the CRR model are

$$p_{t,n} = f(\sigma_{t,n}, P_{t,n}, S_t, n, K, r_{t,n}), \quad (4)$$

$$c_{t,n} = f(\sigma_{t,n}, P_{t,n}, S_t, n, K, r_{t,n}), \quad (5)$$

in which $P_{t,n}$ is the present value at time t of all dividends to be paid between t and n as implied by the call and put option pair expiring at n , $r_{t,n}$ is the risk-free rate and $\sigma_{t,n}$ is the implied volatility of the stock with price S_t .

The build-up of the binomial tree requires to specify whether dividends are paid at discrete points in time, or continuous through time. Discretely timed dividends, as I apply the model, allows for the possibility of a payment shortly before expiry, which would trigger exercise at such a date in the life of the option. The CRR model thus takes into account the pricing consequences for options on the stock of such precipitous exercise.

If no assumptions are made about either implied volatility or implied dividends, then unique solutions cannot be found for each equation (4) and (5) independently from each other. At the same time, each of the two implied unknowns should be equal for calls and puts. The implied volatility refers to the volatility of the stock price, which is the same stock for the call and the put option. The implied dividend refers to the dividend implied to be paid by the company issuing the stock, which is also the same for both options. It follows that, with two equations and two unknowns, unique values can be found for implied dividend and implied volatility under this presumption for pairs of call and put options which have the same strike, expiry date and underlying stock characteristics.

There are concerns that the implied volatility in particular is not the same for calls and puts. The literature on the pricing of options and potential causes for mispricing in the Black-Scholes-Merton (1973) framework looks for trading demand and supply pressure as a potential cause. Bollen and Whaley (2004) argue that changes in implied volatility of index and stock options are related to net buying pressure measured by trades occurring away from the bid/ask midpoint. This causes the implied volatility function across moneyness of the options not to be flat as expected in the classic framework. Cremers and Weinbaum (2010) state that differences between call and put implied volatilities do not reflect pure arbitrage opportunities, rather these represent deviations from PCP, which can be viewed as

proxies for price pressures. The authors conclude that such deviations in option prices can lead stock prices by days. Gârleanu, Pedersen and Poteshman (2009) identify the relative demand for index and stock options by end-users. Option prices are priced away from the Black-Scholes-Merton framework as they cannot be hedged perfectly by intermediaries. The authors measure the expensiveness of options as their implied volatility relative to the expected volatility for the life of the option and establish a positive relation between option expensiveness and end-user demand.

In this paper concerns about short-lived price pressure are mitigated by the aggregation of the data. The implied dividend for an individual firm for a certain trading day is itself the product of several option pairs. Daily data are then averaged into a quarterly data set, which further reduces the impact of one-sided pressures as well as of noise.

There are other concerns with the exact application of the PCP. Kalay, Karakaş and Pant (2014) use deviations to establish the value of the voting premium. The difference of voting share prices from their non-voting synthetic cash flow equivalent using options measures the value of the right to vote as shareholders. They find the voting premium to be positive and increasing with the expiry of the options.

The two pricing equations for put and call options (4) and (5) each have the same two unknown implied dividend and implied volatility and the same observables. Cremers and Weinbaum (2010) as well as Gârleanu, Pedersen and Poteshman (2009) and Kalay, Karakaş and Pant (2014) employ in their analysis the implied volatilities delivered in the Ivy DB data set from OptionMetrics. As in Bollen and Whaley (2004), these are calculated under the constant dividend yield assumption (OptionMetrics, 2011). Harvey and Whaley (1992) state: "Since firms tend to pay stable quarterly dividends at regular periodic intervals during the calendar year, little uncertainty exists about the dividend parameters for short term index options.". This means that future dividends are proxied by the most recent dividend that actually has been paid. With that unobservable input to the model pinned down, it is possible to imply a value for the volatility of the underlying stock from the price of a single option.

From the market for dividend derivatives, however, it is established empirically that risk neutral dividend expectations are neither flat for any horizon nor constant over time (Binsbergen et al., 2013, and Kragt, de Jong and Driessen, 2016). The actual expectation about future dividends prevailing in the market will be different from its last known value for two key reasons. The first is that the present value of a dividend expected to be paid at time n depends on several ingredients:

$$P_{t,n} = PV_t(D_{t+n}) = D_t \exp(g_{t,n} - \theta_{t,n} - r_{t,n}). \quad (6)$$

Expected dividends $E_t(D_{t+n})$ grow at rate $g_{t,n}$ from currently paid dividends D_t and are discounted at the risk premium $\theta_{t,n}$ and the risk free rate $r_{t,n}$ for finding its present value $PV_t(D_{t+n})$. The assumption that the last known dividend will be continued at the same level beyond the observation date thus implicitly equates the objective expected growth rate $g_{t,n}$ to $\theta_{t,n} + r_{t,n}$ for all n .

Furthermore, instead of projecting dividends at a constant rate into the future, market participants are likely to change their expectations of future dividends over time based on their assessment of a company's willingness and ability to pay dividends. Such expectations may change whether the company publicly states to do so or not. It is the variation in the market valuation of dividends, as denoted in equation (6), that the following procedure is intended to capture.

I identify a call and put option pair that is the same for all characteristics, among which future dividends and volatility are unobservable. These two unknowns can be found, in principle, exactly given the option pricing equations (4) and (5). However, model implied prices $\hat{c}_{t,n}$ and $\hat{p}_{t,n}$ depend on building up the CRR binomial trees. They cannot be determined analytically by reverse engineering as this would involve starting from the back-end. This problem is resolved by equating model and observed prices respectively, while the option pricing equations for call and put are inverted at the same time.

In practice, the procedure starts by guessing values for the two unknown implied variables, calculating model prices and reiterating new guesses until convergence of model and observed prices. As both pricing equations require the same guessed value for the inputs, this is pursued by minimizing the quadratic difference between model prices and the observed end-of-day mid-prices of both a pair of call and put options:

$$F = (p_{t,n} - \hat{p}_{t,n})^2 + (c_{t,n} - \hat{c}_{t,n})^2, \quad (7)$$

where F is a function of the implied dividend and implied volatility of the modeled option prices. The minimization of F given the observed option prices results in a unique value for both implied dividends between t and n and for implied volatility for horizon n for each option pair in the data set.

The minimization method is based on two main assumptions: Markets are transparent and efficient, and end-of-day close option prices contain all economically relevant information relevant to market participants. Implied volatility and implied dividends are the same for call and put options with otherwise equal characteristics. As these implied values refer to a single stock, this is a statement of fact rather than an assumption.

2.2 Data and aggregation

The Ivy DB database provided by OptionMetrics contains all options traded on US stocks starting in 1996 and are included in the data set until August 2015. The minimization is performed under the condition that prices relate to companies with CRSP share codes 10 or 11, which restricts the set to regular companies only³. Overall, the data set contains option prices for 5,700 individual companies, which constitutes most of the stocks traded on the NYSE, AMEX and NASDAQ. The prices used are daily and are taken as the average of the bid and ask closing prices. In total, the data set includes a total of 337 million option prices. The daily stock prices S_t and current dividends D_t are also taken from the Ivy DB database. The risk free rate $r_{t,n}$ is proxied by the Eurodollar Libor spot rate, downloaded from Datastream.⁴

To construct a data set that is as large as possible, but from which irrelevant options are eliminated, I apply the following filters:

1. Options of which OptionMetrics does not provide a value for implied volatility are excluded. As OptionMetrics applies the constant dividend assumption, implied volatilities cannot be calculated for options which are far in-the-money or far out-of-the-money. Such options are less frequently traded, which is the main reason to exclude them.
2. Options expiring within 90 days of the observation date are excluded. Once a dividend is announced, there is little uncertainty about the magnitude of the payment being made. As some firms announce dividend payments farther in advance than others, the degree of risk embodied in implied dividends payable in the near future will vary among them. Only a small fraction of firms announce dividends beyond a horizon of 90 days, so restricting the data set to longer option maturities allows for a like-for-like comparison of implied dividends among firms.

Many stocks exist in the data set for just a part of the data period, they sometimes have options traded on them and not all companies pay dividends consistently. Consequently, it is insightful to regard the prevalence of the data per period⁵. Over the 79 quarters in the data period, there are 153,589 quartercompanies with options, or some 64% of all available stocks (see Panel A in Table 1). More often than not, a single stock will have options of the same expiry date with several strikes traded on it. Options with different strikes do not

³This excludes REITS, ETFs and several other underlying types.

⁴Gârleanu, Pedersen and Poteshman (2009), Cremers and Weinbaum (2010) also deploy Libor for this purpose.

⁵I work with quarters for reasons that are clarified later.

have the same implied volatilities, as the volatility surface is not flat but skewed. But as long as the expiry dates of option pairs are the same, the implied value for dividends remain equal to the market's value put on the dividends paid before expiry regardless of the option strike. Option pairs with the same expiry will thus produce the same implied dividends, but varying implied volatilities for varying strikes.

On average, for every observation date and expiry date combination, there are three option pairs with different strikes in the data set. To reduce measurement error, it is useful to find the implied dividends from all option pairs available. I calculate implied dividends from available options pairs for a given expiry date and underlying stock and take their median to determine the relevant value for each observation date. Alternative methods to find averages deliver materially the same results.⁶

The minimization is performed for dividends to be paid discretely, as this accounts for the early exercise premium in the CRR model. This requires identifying dates beyond the observation date at which dividends are assumed to be attributed to the stock. The relevant date to include in the pricing model is the ex-dividend date, which is the first business day following the date on which the dividend is attributed to the stock. The settlement date at which the dividend is actually paid into the account of the owner is not relevant. If the owner decides to sell the stock as of the ex-dividend date, but before the payment date, then the dividend will not be paid into the owner's account.

The latest date at which a dividend has actually been paid is known at any observation date and is used as a starting point. Most US companies pay dividends at a quarterly frequency (see Panel B in Table 1). The subsequent date at which the next dividend is assumed to be paid is set to the last ex-dividend date plus 91 calendar days, the following quarterly payment goes ex-dividend after 182 calendar days and so on. Some assumption about the exact dates cannot be avoided even if company statements about ex-dividend dates

⁶An alternative weighting scheme takes account of the impact of the bid ask spread on the precision of the estimates for the implied dividends. Bid ask spreads tend to be large for options that are traded infrequently, such as when their strikes are far in or out of the money. Large bid ask spreads involve a wide range of consistent mid prices, which increases the potential for measurement error. The larger the bid ask spread, the smaller the weight in the averaging of the dividend implied from among the option pairs

$$\hat{P} = \sum_{i=1}^n \frac{w_i \hat{P}_{K_i}}{\hat{P}_{K_i}} \quad (8)$$

in which the weight w_i for the option pair with strike K_i is set by

$$w_i = (0.05 + |BO_c| + |BO_p|)^{-1} \quad (9)$$

with $|BO_c|$ and $|BO_p|$ representing the bid ask spread of the call and put options respectively. A constant 0.05 is added to avoid inversion of zero. Although this approach appears more robust to measurement error, it delivers materially the same results and is not pursued.

were analyzed piecemeal. Many stocks have LEAPS⁷ traded on them, which are options with expiry dates extending out to over two and a half years (see Panel A in Table 1). In general, companies provide guidance about dividend dates of at most one year in the future. For these long dated options, an assumption about the date of going ex-dividend always has to be made.

Options with an expiry date of more than three months beyond the observation date can span more than a single implied dividend payment. The dollar value of each of them is assumed to be the same. Relaxing this assumption would increase the number of free variables in the minimization which prevents finding unique solutions in the minimization of an individual option pair with more than one dividend payments in (7). At the same time, at any observation date the present value of the first to be paid dividend is implied from the first to expire option during which life a dividend is expected to be paid. From the options with the next quarterly expiry date, both the first and the second future dividends present values are implied simultaneously. By deducting the first quarter implied value, the second dividend's implied value can be isolated, and by iteration the same applies to dividends further into the future.

Companies do not necessarily have to be current dividend payers to have dividends implied from their stock options. Often firms do not pay dividends for a protracted period of time, or never pay dividends at all. Sometimes they are current payers, but interrupt payment for a few years before they return as payers. Several hundred companies in the data set have paid dividends at an interval different than quarterly. And many among them switch payment frequency.

The choice made for aggregation deals with these issues by dividing the data into subsets of observation quarters of companies which pay dividends at a quarterly rate and of observation quarters companies which do not pay dividends. The quarterly frequency is chosen since this allows to establish a data set of implied dividends of which the horizon is constant. In practice, the expiry dates of regular stock options with expiries shorter than nine months are dispersed over a calendar year in a manner that is not the same for all companies. They tend to follow a pattern of expiries in January/April/July/October, February/May/August/November or March/June/September /December. Exceptions exist, particularly in the latter half of the data period. Also, LEAPS always expire in January of up to two years and eight months ahead. If a monthly frequency were pursued, then the horizon would reduce by a month twice for each month that passes and then jump ahead for two months. Aggregation of daily implied dividends to a monthly frequency thus produces a term structure of which the horizon is not the same across companies. Aggregating to

⁷Long-term Equity Anticipation Securities.

buckets of quarterly averages resolves this issue. Appendix I sets out the procedure for this aggregation.

2.3 Summary description of the minimization results

The results of the minimizations are pictured in dividend yields, that is, in present values of implied dividends relative to the share price at quarter t . All companies in the data set have regular options traded on their stock, while some companies also have LEAPS traded on their stock. The expiry of regular options extends out to no more than nine months, but the expiry of LEAPS can be as long as two years and eight months. As a result, growth rates beyond the third quarter are based on LEAPS and thus consists of a smaller sample of companies (see Panel B in Table 1). All data discussed are value weighted.

The term structure of dividend yields averaged over the data period for stocks without LEAPS (Figure 1) starts at 0.60% for currently paid dividends and drops to 0.53% for the average of the first, second and third quarter forward. Average implied dividends of stocks with LEAPS are at about the same level and show a noticeable decrease for quarters beyond the third. The implied dividend yield over the eight quarters following the current quarter is valued at no more than 0.45% to 0.48%, which is a fifth to a quarter less than the current dividend yield. These figures broadly match the levels and slope pattern for dividends based on option prices reported elsewhere (Binsbergen, Brandt and Koijen, 2012).⁸

Figures 2 and 3 illustrate implied dividend yields over time next to the current dividend yield for value weighted averages across all companies who pay dividends at a quarterly rate. Figure 2 shows the implied dividend yield of one quarter ahead and the sum of one and two quarters ahead, Figure 3 contains the implied dividend yield of one to three and one to eight quarters ahead. Several observations stand out. First of all, option prices appear to capture the expectation of dividends sensibly over time. The implied dividend yields follow the current dividend yield well, despite that they are calculated in the minimization process without reference to actual dividends. Second, implied dividends are below current dividends for longer dated horizons. This means that growth rate $g_{t,n}$ is smaller on average than the risky discount rate $\theta_{t,n} + r_{t,n}$ (see Equation (6)). Lastly, the level difference between current dividend yields and implied dividend yields decreases during the time span of the data set.

⁸These authors calculate dividend prices of the S&P 500 index based on PCP over the period 1996 to 2009. The levels they find are close despite differences in data period, smaller set of firms and S&P-weighting versus value-weighting. The slope of the term structure they report at an average of -5.15% for horizon 6 to 12 months ahead and -2.95% for 12 to 18 months ahead. The average slopes found here are a little less than twice as steep. The authors report a positive slope of 0.50% for 18 to 24 months ahead, whereas I find a decrease of -3.2% (figure 1). A slope from actual dividends to 6 months ahead implied dividends can not be calculated from the data they present.

In the latter half, the implied dividends are higher relative to current dividends for both short and long dated options. Lower interest rates can partly explain this structural shift.

The term structure emerges clearly in Figures 4 and 5, which portray sample average growth rates of implied dividends relative to actual dividends. Over time, there is substantial variation in dividend growth, with an upward trend as time passes. Implied dividend growth is negative throughout for longer horizons, but hovers around zero for horizons of one and two quarters. In 2008 a change has occurred taking dividend growth substantially higher. From their peak in 2008, current dividends measured as dividends per share fell by 10%. In part, the sudden increase in dividend growth is therefore a denominator effect.

Returning to the combined set of companies with and without LEAPS, current dividend yields show substantial variation. Breaking them down into deciles indicates a dividend yield of less than 0.2% per quarter for the lowest decile, the highest decile companies pay in excess of 1.1% dividend per quarter (Figure 6).

Growth rates of implied dividends vary substantially among individual companies as well. The shortest horizons give the biggest dispersion in growth rates (Figures 7 and 8). The breakpoint of the highest decile in first quarter growth is nearly 100%, while the breakpoint of the lowest decile is minus 50%. Options of companies in these deciles thus price a near doubling and halving respectively of dividends in the following quarter. Growth rates in consecutive quarters remain widely dispersed. The average growth rate after the third quarter ranges from nearly minus 40% for the lowest decile breakpoint, to over 60% and 30% growth for the highest decile breakpoints of three quarters and two years average dividend yields.

2.4 An example of a change in dividend policy and anticipatory pricing in options

The present value relationship assumes that the value of a stock is derived from expected dividends, but such valuation does not necessarily require that any dividend is currently paid. In fact, the expectation that dividends will be paid may lay far ahead in the future. Implied dividends give some insight into the length of that expectation horizon.

Often companies attain substantial market capitalization even though they do not pay dividends. A case in point is Apple Inc. (AAPL US) which attained a market capitalization in 2011 of US\$ 354 billion following seventeen years of not paying dividends. Its operating profitability reached US\$ 35 billion in fiscal year 2011.

The procedure to imply dividends from option prices works equally well for stocks whether dividends are being paid or not. The implied values found for such stocks are usually close to zero. Only when market participants anticipate dividend payments in the future will implied

dividends be significantly above zero.

Throughout the data set, there are 683 instances in which a company started to pay dividends after at least one year of no dividend payments. In these instances, there is some anticipation discernible from implied dividends ahead of time. Of these companies, the dividend valuation implied one quarter ahead amounted to a little over 30% of the newly installed dividend paid one quarter later. A quarter earlier, this had already reached 28% and three quarters before the payment the valuation reached 17%. For companies with LEAPS traded on their stock, dividend valuation going back even further could be obtained. However, there are only 19 instances of newly started dividend payments in the data set among such companies. In view of this small number, it is instructive to single out an example.

On 19 March 2012 Apple Inc. announced its intent to start paying dividends later in the year⁹. Figure 9 shows implied dividends for Apple priced into LEAPS expiring in January 2013. In the six months preceding the announcement, the implied dividend had already risen from zero to about \$ 0.50 per share. Market participants adjusted their expectations, anticipating at least some dividend payment until the start of the year 2013, which is clearly reflected in option prices before the announcement was made.

Next to the uncertainty over dividend initiation itself, there is uncertainty over its timing. Figure 10 shows the Apple dividends implied from an option series which expires in July 2012, which is before the first payment went ex-dividend in August 2012. It is clear from the implied dividends in this series too that the market anticipated payment of dividends during the first quarter of 2012 before the announcement on 19 March. But when Apple announced that the ex-dividend date would be in August, the implied dividend vanished from this expiry series. Market participants appeared to have realized that payment was to occur after its expiry in July.

Following the announcement, the option implied dividend per share per quarter until January 2013 rose from about \$0.50 first to about \$1.25 and then to \$2.00. The implied dividend remains lower than the announced level of \$2.65 because the exact number of discrete dividend payments deviates from the model assumptions at the time of the press release on the 19th of March.¹⁰

⁹Apple stated: "... the Company plans to initiate a quarterly dividend of \$2.65 per share sometime in the fourth quarter of its fiscal 2012, which begins on July 1, 2012."

¹⁰Apple paid two dividends of \$2.65 before the expiry of the January 2013 LEAPS. Since Apple didn't pay dividends up to that point, payments are modeled to occur 91 days, 182 and 273 days following the observation date. There are thus three payment dates over which to divide the present value of two anticipated dividend payments between the announcement on 19 March and 20 April, which is 273 days before the expiry of this LEAPS in January 2013. It is noteworthy, however, that after 20 April the implied dividend still doesn't reach the announced amount of \$ 2.65.

2.5 Stock loan fees and robustness

The PCP relationship contains all instruments relevant to an investor wishing to run a portfolio of options. However, market makers in particular use long and short positions in stocks as well to hedge their exposure to stock prices as a result of their option book. In cases where such market participants take short positions in the stock, they need to take into account the cost of doing so. D’Avolio (2002) documents that these stock loan fees amount to 25 basis points per annum on average. In his sample 91% of all stocks are not hard-to-borrow, and these can be borrowed at a running cost of 17 basis points.¹¹

If cash income from stocks not only appears as dividends, this should be incorporated in the PCP. It is straightforward to add loan fees f_t to the dividend $P_{t,n}$ paid by the company on its shares. The left-hand side of the PCP as depicted in equation (3) is thus expanded with fees as follows,

$$P_{t,n} + S_t(e^{f_{t,n}} - 1) = S_t + p_{t,n} - c_{t,n} - Ke^{-r_{t,n}}, \quad (10)$$

which, for small $f_{t,n}$, approximately equals

$$P_{t,n} = S_te^{-f_{t,n}} + p_{t,n} - c_{t,n} - Ke^{-r_{t,n}}. \quad (11)$$

Market data for stock loan fees are available and they can in principle be included in the minimization in (7). However, stock loan fees may be adjusted throughout the period that the stocks are on loan. Stock lenders and borrowers may terminate lending at any point in time (Avellaneda and Lipkin, 2009). As a consequence, the actual cost of borrowing a stock is not precisely captured by loan fee data. Moreover, fees in the PCP in equation (11) presume a cost for borrowing stocks throughout the life of the option. The typical period for borrowing stocks is measured in days. In the data sample, option expiries extend out to over two years. As a consequence, market data about fees that represent the actual cost anticipated for borrowing stocks until the expiry date of particular options do not exist. Including fees in the PCP would require unverifiable assumptions about their future value.

An alternative route to take account of stock loan fees is to imply them in a manner similar to implied dividends. The minimization in (7) is performed on two equations, allowing for two unknowns: implied volatility and implied dividend. To find another implied variable is feasible by expanding the number of option pairs in the minimization from one to two. The

¹¹These stocks are *General Collateral*, which are all loaned at the same fee since they are readily available and the fee only needs to compensate for the service of stock lending. Stocks that are hard-to-borrow are loaned at a higher fee, which makes up for the difference between the average fee of all stocks and the General Collateral fees documented by Avellaneda and Lipkin (2009).

two pairs are the same except for the strike, which increases the number of unknowns by a third variable: the implied volatility of the second option pair. The number of equations is increased from two to four, thus allowing for a fourth variable to be estimated as well: implied stock loan fees. Similar to dividends, the stock loan fee refers to the share and not the options. Therefore, the stock loan fee is the same for the modeled prices of both option pairs and can be implied as a fourth variable. The fact that they can be identified separately from dividends is due to the fact that stock loan fees are deemed continuous, whereas dividends are paid discretely. Appendix II contains a section which explains how the implied stock loan fee is included in the minimization procedure in more detail.

The implied stock loan fees produced by the minimization are shown in Figures 11 and 12. Throughout the data period, the fees move between 15 and 20 bppa up to two quarters ahead and slightly lower in the quarters beyond. Implied fees for different expiries move in conjunction with each other.

The levels of the implied fees come close to the market data for General Collateral fees of 17 bppa documented by D’Avolio (2002). The fact that implied fees move in a tight range reflects that investor opinions about future fees change little compared to current fees as time passes. This is confirmed by the term structure of implied stock loan fees in Figure 13. Its slope is negative but shallow. If average stock loan fees do not change much over time, investors may be motivated to make an assumption to this effect in the pricing of options.

Cross-sectional dispersion is shown in Figure 14. The breakpoint for the decile with smallest implied stock lending fees stands at about 12 bppa against 22 bppa for the decile with the largest fees. Such a difference is not substantial. Stocks that are hard-to-borrow can require fees for several percentage points per annum (Avellaneda and Lipkin, 2009) and the largest decile would be expected to have a higher breakpoint. However, such fees occur more often for small stocks (Avellaneda and Lipkin, 2009) and these are less likely to be present in the database of stock which have options traded on them. Moreover, in the minds of investors high fees do not necessarily remain high for a prolonged period. Implied fees represent their expectations of stock lending until expiry of the options, which may be led more by the average of fees than by the fee a particular stock that happens to be in demand.

As implied stock loan fees do not show large dispersion, their effect on the dispersion of implied dividends should be small. Implied dividends which are calculated while taking account of stock loan fees according to 11 indeed show very similar patterns and levels as those calculated without regarding fees. Although the minimization method to find implied dividends appears robust to the inclusion of fees, I chose to omit their mention for the remainder of this paper. The interpretation of implied dividends is thus to include stock loan fees throughout this paper.

3 The predictive power of implied dividends

An obvious question is whether implied dividends have predictive power for actual dividends. The opinion that investors express about the value of future dividends can be tested as an explanatory variable for the same dividends being paid later.

In the period 1996 to 2015, nearly 74,000 quarterly dividend payments were made by companies with stocks on which options were traded (Table 3). Most of the time companies leave dividend payments unchanged relative to the previous quarter, which confirms that a commonly accepted policy of dividend smoothing is widespread. Nonetheless, there are many instances of a change in dividend policy. Four types of policy changes are identified. In over 11,000 cases firms increase their dividends, where in more than 2,000 cases they decreased it. At only 66 occurrences, firms terminating dividends are rare, and the majority of them go out of business not long after. In total 683 times a company started paying dividends, which is defined as a dividend payment following at least four quarters of non-payment.

At a quarterly rate of 0.79%, the dividend yield is on average highest in the last quarter for companies that make a payment before they cease to do so. The group of companies decreasing dividends, reduce it from a dividend yield of 0.66% to 0.48% (the median \$ cut is -46%). When they increase payment, the change is more modest from 0.55% to 0.62% (the median \$ increase is 11%). Companies initiating dividend payments are the most frugal relative to their stock price and introduce payment at a yield of 0.44%. The biggest group of companies leaves their dividends unchanged at a dividend yield of 0.57%. This pattern confirms that companies are careful to increase their dividends, but once it happens they act more determined.

The example of implied dividends anticipating Apple's dividend initiation in 2012 suggests that dividends implied by option prices are forward looking not only in theory, but also in practice. An interesting question is whether this relationship holds up across companies and through time. In particular, I investigate whether implied dividends provide information about future dividends over and above lagged dividends. Since dividend changes are often zero and do not show a continuous distribution, an ordered Probit estimation is pursued in addition to a standard OLS. I conclude by considering the stock price response to announcements to change dividends in order to establish whether dividend changes that are correctly anticipated by implied dividends reduce this response.

3.1 Predictive OLS regressions

If market participants were to predict dividend changes correctly, then the relationship between actual and implied dividend changes would be linear, with a coefficient equal to one and zero intercept. The following regression equation serves as a first attempt to test the hypothesis:

$$d_{i,t+n} = \alpha_t + \beta_t p_{i,t,n} + \epsilon_{i,t+n} \quad (12)$$

The actual dividend change $d_{i,t+n}$ is defined as the relative change in the dividend paid n quarters following the observation quarter t for a given stock i :

$$d_{i,t+n} = \frac{D_{i,t+n} - D_{i,t}}{D_{i,t}}. \quad (13)$$

The implied dividend change $p_{i,t,n}$ is the dividend change for n quarters ahead relative to the observation quarter as implied by the data $P_{i,t,n}$:

$$p_{i,t,n} = \frac{P_{i,t,n} - D_{i,t}}{D_{i,t}}. \quad (14)$$

Implied dividends, at the same time, are valuations of future dividends and not objective expectations of future dividends (see equation (6)). The difference between the two is the dividend risk premium, which is not the subject of investigation here. Under the hypothesis that this risk premium is constant over time, it should surface as a non-zero intercept. However, if the dividend risk premium is time-varying, it may also show up in the coefficient of implied dividends. The appropriate hypothesis for testing the predictive content of implied dividends is that their relationship to actual dividends is close to linear with a coefficient close to one.

The implied dividend data are quite noisy. Figure 15 contains a scatter plot of actual and implied dividend changes two quarters ahead of the observation quarter. The plot shows that they are heavily scattered away from a hypothesized regression line of model (12), which is at least in part caused by measurement error at the stage of the calculations laid out in the previous section.

The plot also bears out that downward and upward dividend changes do not mirror each other. Under the model, downward changes should be concentrated in the southwest quadrant and upward changes in the northeast quadrant. For upward changes this appears reasonably clearly, but it is only barely visible for downward changes. As far as visible, downward changes align vertically close to zero implied dividend growth, while positive changes align horizontally slightly above zero actual dividend growth. Neither pattern fit

well to model (12).

To assess these patterns statistically in addition to the base model, upward and downward changes are separated and introduced as two different independent variables to an alternative model, as follows:

$$d_{i,t+n} = \alpha_t + \beta_{up,t} I_{p>0} \times p_{i,t,n} + \beta_{down,t} I_{p<0} \times p_{i,t,n} + \epsilon_{i,t+n}, \quad (15)$$

in which dummies $I_{p>0}$ and $I_{p<0}$ attain the value of 1 for a positive implied dividend change and negative respectively and 0 otherwise.

OLS regressions are run cross-sectionally for each quarter t of the data period of 1996 to 2015. This quarter-by-quarter approach allows for calculating standard errors of quarterly beta's, in a manner similar to Fama and MacBeth (1973). The predictive power of implied dividends is tested for n one to three quarters ahead, resulting in 77 to 79 regressions per n .

The estimated coefficients from the regressions of models (12) and (15) are displayed in Table 2. The averages of the coefficients of model (12) bear out that implied dividends do not line up strongly with actual dividends. Only 7 to 10 percent of a change in implied dividends translates into actual dividend changes. Even though these coefficients are small, they are significant, particularly for predictions two and three quarters ahead.

The alpha's in these regressions are close to but significantly larger than zero. There may be at least two explanations that implied dividend changes are smaller than actual dividend changes. Dividends increase throughout the data period, by on average 2.4 percent per quarter. Considering that only a fraction of such average change is picked up by implied dividends, the intercept reflects a substantial part of the average increase. A second explanation is that implied dividends are risk-neutral. If model (12) is true, then the alpha would reflect the risk premium and the time value of money as a result of which a positive value is expected.

The test results of the alternative model (15) confirms the visual inspection of the data in Figure 15. Applying up and down dummies to implied dividend changes shows that negative predicted changes attract an unexpected negative sign and are not significant. Only positive changes demonstrate statistically relevant predictive power for actual dividends¹². Their coefficients come out slightly stronger than in the base model, which is to be expected given the filtering of the weak power present in downward changes. Market participants only reflect upwards dividend changes in option implied dividends as dividend cuts remain largely unpredicted by implied dividends.

¹²The weakness of negative predictions is also reflected in an average R^2 of 0.35% for a model that only includes negative predictions (1.10% to 1.60% for model (15)).

3.2 Predictive ordered Probit

The scatter plot in Figure 15 depicts another particularity of the relationship between actual and implied dividends. Dividends are set by company management and are thus discretionary in nature. The plot shows that firms often choose a "round percentage" for a dividend change, for example, minus 50 percent¹³ or plus 50 percent. One in five dividend payments constitutes a change in the size of the payment, with increases occurring about five times as often as decreases (Table 3). At four out of five times, the most prevalent dividend change is therefore no change as also noted by Guttman, Kadan and Kandel (2010) and Baker, Mendel and Wurgler (2016). The last paid actual dividend is therefore highly relevant to the next payment due to the wide-spread policy of dividend smoothing.

While implied dividends may signal a change in future dividends, only when a change in actual dividends occurs will implied dividends be identified as relevant to future dividends. Moreover, any change in dividends tends to be large (Table 3). The average increase amounts to about 24 percent while the average decrease is -42 percent. If implied dividends predict a small change in actual dividends, it may often be too small to be effected by company management. Baker, Mendel and Wurgler (2016) note that this obscures a highly non-linear relationship where changes around zero occur much more frequently than larger movements. The high persistence in actual dividends thus presents a challenge to the linearity assumed in models (12) and (15).

Isolating the many zero-change occurrences in actual dividends in the data set requires a simple comparison of the number of instances in which a dividend change is correctly predicted by implied dividends. The north-east quadrant in Figure 15 contains correctly predicted increases, the south-west quadrant contains correctly predicted decreases and the other quadrants show incorrect predictions. A positive implied dividend, defined as an implied dividend value above its median, correctly predicts a future increase in dividends in 54 to 57% of the observed instances, depending on the horizon of the dividend implied. A decrease is predicted correctly by implied dividends below their median just about as often as it is predicted incorrectly (Table 4). This poor result for dividend cuts is matched by the weak explanatory power of implied decreases in the OLS estimates (Table 2).

When implied dividends are larger than actual dividends, there should be at least some predictive effect on the probability of actual dividends increasing in the future. The particular characteristics of the data set induces to order the actual dividends data set into categories. This is followed by Probit estimation which produces the probabilities of the dividends being changed as predicted by positive or negative implied dividends.

¹³Nearly one in four decreases is exactly minus 50 percent. Baker, Mendel and Wurgler (2016) discuss salience and round numbers in dividends and dividend changes in more detail.

Actual dividend changes d are ordered into four categories: a middle category that is identified as actual dividend growth close to zero, a decrease category and two increase categories for small and large increases. These categories are represented formally as follows:

$$y = \begin{cases} 1 & \text{if } d \leq -0.01 \\ 2 & \text{if } -0.01 < d \leq 0.01 \\ 3 & \text{if } 0.01 < d \leq b \\ 4 & \text{if } d > b \end{cases}$$

When actual dividends are cut, they tend to be cut drastically. Compared to cuts, dividend raises are smaller; the median of cuts is -45.71% and the median increase is 10.81% (Table 3). The small increases are often the product of a company policy to gradually raise dividends over time. Such dividend smoothing policies are hard to distinguish from zero dividend changes.

Actual dividend changes d are ordered into a category $y = 1$ for cuts, of which the upper boundary is set to -0.01 . Zero changes are categorized separately in $y = 2$.¹⁴ Small increases are allocated to the third category $y = 3$ and large increases are allocated to the fourth category $y = 4$. Carving up increases in y serves to account for the wide-spread policy of dividend smoothing, which leads to increases being both small and plentiful. Carving up the first category into one for small and large cuts is not material to the results due to the small number of small dividend cuts and because cuts are large more often than not. Boundary b between the third and the fourth category will be varied in order to exhibit the sensitivity of the parameter coefficients to it. Ordinal values y represent latent variable y^* which is a linear function of the model:

$$y_{i,t}^* = \beta_{up} I_{p>0} + \epsilon_{i,t}. \quad (16)$$

The up-dummy $I_{p>0}$ attains a value of 1 when implied dividends are larger than actual dividends and 0 otherwise. The probabilities of selecting one of the categories equal the probabilities of the latent variable falling into an estimated intercept (α_i) category. These probabilities then are equal to the standard cumulative distribution function Φ value of the

¹⁴The purpose of the boundaries that define $y = 2$ is solely to isolate the zero-change instances into a separate category. Setting these boundaries to small fractions other than but close to -0.01 and 0.01 is not material to the results.

estimated parameters:

$$P(y = 1) = P(y^* \leq \alpha_1 | p, \beta_{up}) = \Phi(\alpha_1 - I'_{p>0} \beta_{up})$$

$$P(y = 2) = P(\alpha_1 < y^* \leq \alpha_2 | p, \beta_{up}) = \Phi(\alpha_2 - I'_{p>0} \beta_{up}) - \Phi(\alpha_1 - I'_{p>0} \beta_{up})$$

$$P(y = 3) = P(\alpha_2 < y^* \leq \alpha_3 | p, \beta_{up}) = \Phi(\alpha_3 - I'_{p>0} \beta_{up}) - \Phi(\alpha_2 - I'_{p>0} \beta_{up})$$

$$P(y = 4) = P(\alpha_3 < y^* | p, \beta_{up}) = 1 - \Phi(\alpha_3 - I'_{p>0} \beta_{up})$$

Coefficient β_{up} indicates the impact of a unit change in implied dividends on the probability of one of the categories occurring. The model is estimated using maximum likelihood for each quarter in the data set. The coefficients are then averaged over all quarters and their t-statistics are calculated following the method of Fama & MacBeth (1973), similar to the t-statistics in the OLS regression in the previous section. Regressions are performed for dividends one, two and three quarters ahead of the observation quarter.

Results of the Probit estimation are shown in Table 5 for boundary value $b = 0.10$.¹⁵ For one quarter ahead the intercepts are further removed from the distribution mid-point than for two quarters ahead and, in turn, they are further removed than for three quarters ahead (Panel A). A change in implied dividends thus signals a larger likelihood of predicting a change in actual dividends the further ahead such change is implied. A potential reason is that nearby dividend changes may be known as they are announced sometimes several months before the ex-dividend date.

The positive implied dividend coefficient β_{up} has the expected positive sign. The inference is that a change in implied dividends from negative to positive reduces the likelihood of actual dividends falling in "cut" category 1, and increases the likelihood of falling in category 4, in which dividends rise by more than b percent. The coefficient of positive changes varies between 0.056 to 0.090 over the three quarters horizon, differing statistically from zero based on Fama & MacBeth t-statistics.

Are these coefficients material to the probabilities of actual dividends moving from one category to another? The marginal effects of the coefficients in this Probit estimation are defined as the change in the likelihood of falling into a particular category in response to implied dividends pointing to an increase instead of a cut. Translating this marginal effect to the discrete character of the dummy-variable is performed by differencing the probability

¹⁵Figures 16 and 17 show marginal effects for other values of b .

of a certain category of a dividend increase from a dividend cut:

$$\delta P(\alpha_{i-1} < y^* \leq \alpha_i | \beta_{up} I_{p>0}) = [\Phi(\alpha_i + \beta_{up} I_{p>0}) - \Phi(\alpha_{i-1} + \beta_{up} I_{p>0})] - [\Phi(\alpha_i) - \Phi(\alpha_{i-1})]$$

Panel B in Table 5 shows the probabilities of a category occurring in case of a negative implied dividend. Marginal effects on the likelihood of falling into the $y = 1$ (decreasing) or $y = 2$ (zero change) actual dividend category are negatively affected by implied dividend changes (Panel C in Table 5), as is determined by coefficient β_{up} . An implied dividend larger than actual dividend reduces the chance of falling into these categories and thus increases the chance of falling into the one of the two dividend increase categories $y = 3$ and $y = 4$.

The marginal effect of a positive implied dividend for an actual dividend change to fall in one of the positive categories is 2.1% to 2.7%. For dividend increases cut-off at 10%, the chances of falling into the increase $y = 4$ category above 10% is larger than to fall into the $y = 3$ small increase category. On average, a positive implied dividend hints at a larger increase in dividends rather than a small one.

Increases are accounted for in two separate categories to deal with dividend smoothing. This is a policy practiced by company management to change dividends only slowly in response to earnings changes. Based on survey evidence, Lintner (1956) observes that many managements believe that most stockholders prefer a reasonably stable dividend rate and that they put a premium on stability or gradual growth in the rate. Survey evidence in Brav et al. (2005) shows that managers feel strongly that the penalty for reducing dividends is substantially greater than the reward for increasing them. Evidenced by survey data, companies often pursue a policy to smooth dividend payments as long as conditions regarding stable earnings, among others, are fulfilled (Brav et al., 2005). Together with a desire to avoid dividend cuts, such a policy leads to increases in dividends being more frequent but also smaller than dividend cuts. It is worthwhile to find out whether smaller dividend increases, induced by regular policies, are indeed better predicted by implied dividends than larger ones that are more likely to reflect a one-off adjustment to the firm's structurally improved financial position. Cutting off the data set of dividend increases into two separate increase categories $y = 3$ and $y = 4$ serves this purpose.

The sensitivity of the marginal effect of a positive implied dividend to the value of the upper category boundary b is instructive (the results discussed thus far refer to $b = 0.10$ (Table 5)). The marginal effect relative to the probability of a dividend falling into the $y = i$ category for decreasing implied dividends provides insight into the importance of moving

from negative to positive implied dividends:

$$\frac{\delta P(\alpha_{i-1} < y^* \leq \alpha_i | \beta_{up} I_{p>0})}{P(\alpha_{i-1} < y^* \leq \alpha_i | \beta_{up} I_{p>0})} = \frac{[\Phi(\alpha_i + \beta_{up} I_{p>0}) - \Phi(\alpha_{i-1} + \beta_{up} I_{p>0})] - [\Phi(\alpha_i) - \Phi(\alpha_i - 1)]}{\Phi(\alpha_i) - \Phi(\alpha_{i-1})}.$$

Figures 16 and 17 shows this marginal effect conditional on boundary d for the two raise categories $y = 3$ and $y = 4$. A boundary slightly above 1% constructs a narrow category, which involves that a positive implied dividend has a limited effect on improving the probability of predicting an increase correctly to that small raise category. Larger boundaries increase the size of the category and thus the probability of a raise falling into it, but the probability of predicting the small raise category increases more than proportionally up to a boundary of approximately 25%¹⁶. Predicting the big raise category deteriorates already at a small boundary. Consequently, a move from negative to positive implied dividends improves the chances of predicting a raise smaller than 25%. Raises larger than that are predicted better by implied dividends only marginally. The policy of firms to raise dividends at a steady moderate pace rather than in large jumps appears to be reflected in market pricing.

3.3 Stock price response to dividend changes

The degree of predictability of future dividends may be picked up in the response of the stock price to a change in dividend policy. If an announcement to change dividends is expected, then it should be reflected in the stock price before the fact. Larkin, Leary and Michaely (2016) document that market reaction, measured as three-day stock return in absolute terms, to dividend reductions is more than double that of increases. They also find that institutional investors, mutual funds in particular, are more likely to hold dividend-smoothing stocks, but retail investors are less likely to do so. Dividend smoothing thus affects the composition of a firm's shareholders but, according to the authors, has little impact on its stock price.

The previous section discusses that negative implied dividends appear to have very little to say about actual dividends, while positive implied dividends do, and with the expected sign. Implied dividends thus indicate that investors do not see dividend cuts coming and may consequently be more surprised by them than in the case of increases. I investigate whether the difference in stock price responses between cuts and increases as well as the size of the response is connected with the expectations of future dividends that is part of the option implied valuation. I also test the impact to stock prices of dividend changes for being correctly or incorrectly predicted by implied dividends.

¹⁶The level difference between the proportional marginal effects of differing horizons shown in these figures is due to the differences found in the intercepts (Table 5).

To establish the data sample, I identify all dividend changes of companies paying dividends at a quarterly rate in the period of 1996 to 2015 with stocks that have options traded on them as before. Following the approach of Grullon et al. (2002)¹⁷, small changes due to rounding and recording of stock splits, as well as extreme observations, are eliminated by limiting the dividend announcements to those with absolute value of changes in quarterly common dividend per share between 12.5% and 500%. The sample is restricted to distribution events in which the declaration date is a non-missing trading date and there is no more than one dividend announcement made per event. For every dividend change, the three-day Cumulative Abnormal Return ($CAR_{i,t}$) is the sum of daily returns of the stock of the announcing firm i around the announcement ($[-1,+1]$ trading days) minus the CRSP value-weighted market return. The implied dividends $p_{i,t}$ are calculated as the median of the ten trading days preceding the dividend announcement date.

The stock price response is regressed on the actual dividend changes and implied dividend changes in various guises of the following equation:

$$CAR_{i,t} = \alpha + \beta_1 I_{d<0} + \beta_2 I_{d>0} \times d_{i,t} + \beta_3 I_{d<0} \times d_{i,t} + \beta_4 I_{p<0} + \beta_5 I_{d>0} \times I_{p>0} + \beta_6 I_{d>0} \times I_{p<0} + \beta_7 I_{d<0} \times I_{p>0} + \beta_8 I_{p<0} \times I_{p<0} + \epsilon_{i,t}, \quad (17)$$

from which parameters are included in the regression in several compositions. Results are presented in Table 6, showing only the two-quarter ahead implied dividends for ease of presentation. One and three quarter ahead implied dividends deliver broadly the same results, as usually different implied horizons jointly point to an increase or a cut.

Starting without implied dividends, the pattern found by Larkin, Leary and Michaely (2016) that dividends cuts weigh more on stock returns than dividends raises is repeated. The intercept shows a 0.7% three-day price return in response to dividend increases and a return of -1.6% for dividend cuts ($\alpha + \beta_1$, model (1)). Once the interaction of the size of the change is included, taking into account that the average cut is -0.42 , the average response to a cut is about the same ($\alpha + \beta_1 + \beta_3 \times \mu_d$, model (2)). This data sample also shows that the impact of the size of dividend cuts on stock prices β_3 is ten times as large as that of dividend increases β_2 ¹⁸.

Next I test whether implied dividends decrease the surprise impact of a dividend change on stock prices. If the implied dividend points to a dividend cut by means of a down-dummy $I_{p<0}$ then the positive surprise impact is indeed reduced to -0.4% ($\alpha + \beta_1 + \beta_4$, model (3)) versus -0.6% , model (2)). Statistically these results are not strong, the coefficient of the

¹⁷Their approach is followed by Larkin, Leary and Michaely (2016) as well.

¹⁸Larkin, Leary and Michaely (2016) report a ratio of four

down-dummy is not significant and the R^2 does not improve once the dummies are included.

Will stock prices respond differently to a dividend change that is correctly predicted by implied dividends from one that is not? A direct test of this question is pursued in model (4), which deploys four interaction terms. The regressors capture the four possible scenarios when a dividend change occurs: a correctly predicted increase or cut and an incorrectly predicted increase or cut. The results point to a very clear distinction between the response to the kind of change. A dividend raise is followed by a positive three-day stock return of $0.6 - 0.7\%$ whether correctly predicted by implied dividends or not. A cut, however, causes a substantial drop in stock prices (-2.6%) if it is not anticipated by implied dividends. If the cut is predicted, the response is only -0.2% . The coefficient for unpredicted cuts is strongly significant and subsumes most of the impact of a dividend cut on stock returns (β_1).

In summary, as a predictor of stock returns around dividend announcement, only positive implied dividends appear to reduce the impact (model (4)). Negative implied dividends do not bear on announcement returns in their own right. But once implied dividends are used to predict dividend changes, it is the incorrectly predicted cuts that move stock prices. Taking this conclusion together with the earlier results leads to the following interpretation on the element of surprise. Testing equation (15) shows that increased dividend payments are significantly predicted by implied dividends, whereas dividend cuts are not. It seems reasonable to presume that the expectations contained in implied dividends also prevail in the valuation of stocks; if market participants have reason to anticipate that a firm will raise dividends, then they will reflect this in the equilibrium price of its stock as well as in its dividends. Once the firm delivers on a higher dividend, the stock price may be unmoved by this fact alone. Since dividend cuts are not well predicted by implied dividends, a dividend cut usually is a surprise causing a stock's price to drop.

Larkin, Leary and Michaely (2016) find an explanation for the large negative response to dividend cuts in controlling for whether firms have cut at least once before. If they have, then the absolute impact on the stock price is about equal for an increase or a cut. A previous cut raises the awareness to the possibility of another cut in the future, reducing the element of surprise if a subsequent cut occurs. Their argument is thus similar to implied dividends predicting a dividend change: the announcement effect on stock prices reduces if there is reason to be less surprised about it.

4 Conclusion

There is a lot of information about the valuation of future dividends of individual firms to be gained from option prices. The extant literature on future dividends focuses on derivatives on stock indices. E.g. Binsbergen, Brandt and Koijen (2012) connect implied dividends to the term structure of risk premiums. But a cross-sectional approach to identifying future dividends has been lacking thus far. This study is the first to present a methodology to do so and applies the data found to their predictive power for dividends and stock returns.

A novel element in the methodology to imply dividends from option prices is that it clears the hurdle of dealing with the early exercise premium that is not present in index options but is an element of value in American style options of US stocks. The Cox, Ross and Rubinstein (1979) binomial tree takes account of this premium. The tree is constructed simultaneously for a pair of put and call options with otherwise equal characteristics by guessing the same values for implied dividends and implied volatilities for the put and call pricing models.

The implied data found do not equal actual dividend data. This should be a warning sign to the assumption often made in the model-pricing of options that future dividends are equal to actual dividends. On average, however, they are consistent with and slightly lower than actual dividend data, showing a modest negative term structure.

At the level of individual companies implied dividends are prone to measurement error. Nonetheless, they prove meaningful to predict actual future dividend changes, notably dividend increases. Dividend cuts are not well predicted. But if implied dividends do correctly predict a future dividend cut, the stock price response to the announcement thereof reported by Larkin, Leary and Michaely (2016) is reduced substantially.

In view of these results, there is good reason to pursue further research on implied dividends. The first is to improve on the methodology itself and to further investigate their relationship with actual dividends and stock prices. The second is to investigate whether implied dividends contain idiosyncratic risk that can be isolated from systemic risk, contradicting traditional CAPM reasoning. An obvious pursuit is to construct a factor based on implied dividends and establish whether portfolios of high implied dividends produce high or low stock returns.

References

- Avellaneda, Marco and Mike Lipkin, 2009, A dynamic model for hard-to-borrow stocks, *Working paper*.
- Baker, Malcolm, Brock Mendel and Jeffrey Wurgler, Dividends as reference points: a behavioral signaling approach, *The Review of Financial Studies*, 2016, 29, 697-738.
- Binsbergen, Jules H. van and Ralph Koijen, The term structure of returns: facts and theory, *Journal of Financial Economics*, 124 (1), 2017, 1-21.
- Binsbergen, Jules H. van, Wouter H. Hueskes, Ralph S. Koijen and Evert B. Vrugt, 2013, Equity yields, *Journal of Financial Economics* 110 (3), 503-519.
- Binsbergen, Jules H. van, Michael W. Brandt and Ralph Koijen, On the timing and pricing of dividends, *American Economic Review*, 102(4), 2012, 1596-1618.
- Bollen, Nicholas P.B. and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance*, 59, 711-753.
- Brav Alon, John R. Graham, Campbell R. Harvey and Roni Michaely, 2005, Payout policy in the 21st century. *Journal of Financial Economics*. 77 (3), 483-527.
- Cochrane, John H., 2011, Presidential Address: Discount rates, *The Journal of Finance*, 66, 4, 1047-88.
- Cox, John C., Stephen A. Ross and Mark Rubinstein, 1979, Option pricing: a simplified approach, *Journal of Financial Economics*, 7, 229-263.
- Cremers, Martijn and David Weinbaum, 2010, Deviations from put-call parity and stock return predictability, *Journal of Financial and Quantitative Analysis*, 45, 335-367.
- Gârleanu, Nicolae, Lasse H. Pedersen and Allen M. Potesman, 2009, Demand-based option pricing, *Review of Financial Studies*, .
- Grullon G., R. Michaely and R. Swaminathan, 2002, Are dividend changes a sign of firm maturity?, *Journal of Business*, 75, 3, 387-424.
- Kalay, Avner, Oğuzhan Karakaş and Shagun Pant, 2014, The market value of corporate votes; Theory and evidence from option prices, *Journal of Finance*, 69, 3, 1235-71.
- Kragt, Jac. C, Frank C.J.M. de Jong and Joost J.A.G. Driessen, 2016, The dividend term structure, *Working Paper*.
- Larkin, Yelena, Mark T. Leary and Roni Michaely, 2016, Do investors value dividend-smoothing stocks differently?, *Management Science*, 1, 1-23.
- Harvey, Campbell R. and Robert E. Whaley, 1992, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics*, 31, 43-73.
- Miller, Merton, and Franco Modigliani, 1961, Dividend policy, growth and the valuation of shares, *Journal of Business*, 34, 411-433.

OptionMetrics, 2011, *File and Data Reference Manual*, version 3.0.

Shiller, Robert, 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *The American Economic Review*, 71, 421-36.

Appendix I: Aggregation of daily implied dividend data into quarterly buckets

The daily dollar value of the implied dividend as found in the minimization of equation (7) is divided by the stock price of the day. This daily implied dividend yield is then averaged over the period from the first day following the third Friday in a given month until and including the third Friday of the next month. Constructing months this way guarantees that averages refer to the same option expiry date. Only if an implied dividend falls before an expiry date is it included in the option price. If it is paid after the expiry date it is not, even though it may still occur in the same calendar month. Hence the redefinition of months. The average is calculated after dropping the two most extreme values in each month.

The next step in the reduction of the data frequency is to bucket the values for each three months into a quarterly dividend yield. This involves taking an average again, but the horizon of the three implied dividends is not constant and it is therefore appropriate to interpret these as buckets instead of averages. Quarterly buckets are only taken if daily solutions for implied values exist for at least 75% of the days of the quarter. Equipped with quarterly dividend yields, a growth rate with a quasi-constant horizon is calculated, starting from a dividend actually paid in a given quarter to implied dividends in the ensuing quarters. A quarterly data set has the added benefit that it brings the frequency of the dividends implied in line with the dividend payment frequency that is the most prevalent among US companies. Lastly, I assume that any distortion in implied dividends due to inventory imbalances among market makers, asymmetric information or fixed cost is of constant mean over the daily implied values averaged to quarterly buckets.

There are two important arguments not to reduce the frequency to less than quarterly. First, the data period from 1996 to 2015 is only 20 years. Time series analysis would suffer from less variation due to fewer data points in case of an annual or semi-annual frequency. Second, most stocks have options maturing eight to nine months into the future, but many among them do not have longer dated LEAPS traded on them. A consistent pattern of implied dividends cannot be offered for annual frequency for these companies, although semi-annual would still be feasible.

Appendix II: Stock loan fees included in the minimization

In order to find implied stock loan fees, two pairs of call and put options are identified that have the same characteristics except for the strike instead of one pair of options. The modeled prices of all four options share the same implied dividend and implied stock loan fee. For each option pair with the same strike volatilities are implied separately from each other. This leads to four unknown variables to be implied: dividend, stock loan fees and the volatility of the first pair and the volatility of the second pair. F is a function of all four:

$$F = (p_{i,t,n} - \hat{p}_{i,t,n})^2 + (c_{i,t,n} - \hat{c}_{i,t,n})^2 + (p_{j,t,n} - \hat{p}_{j,t,n})^2 + (c_{j,t,n} - \hat{c}_{j,t,n})^2, \quad (18)$$

and is minimized as follows. Each model price is calculated starting by guessing values for the four unknown variables which serve as input to the CRR binomial tree. The four model prices are compared to their market prices in equation (18) after which new values are guessed until convergence of model and observed prices.

The procedure is performed for option pairs with strikes that are next to each other. For example, if there are n strikes available of options with the same expiry date, with strikes 1 to 6. The minimization is performed on option pairs with strikes 1 and 2, 2 and 3, and so on, until strikes 5 and 6. The implied values for dividends and stock loan fees represented as the output of the minimization procedure are the medians of the five values implied.

Table 1: **Option and LEAPS Distribution and Dividend Payment Frequencies**

Prevalence of options and dividend payment frequencies for US stocks traded from 1996 to August 2015. A quartercompany is a quarter of daily data in which at least 45 trading days of options are available for a given stock.

Panel A: Excluded from the data set are options of which OptionMetrics does not provide implied volatilities and with maturities shorter than three months. LEAPS are options with a maturity at start of at least two years.

Panel B: Payment frequency of dividends per quartercompany. The data set includes options of stocks with quarterly dividend frequencies only.

	Number of Quartercompanies
<i>Panel A: Prevalence of options</i>	
Stocks total	238,964
Stocks without options	85,375
Stocks with options	153,589
Stocks with LEAPS	39,389
Stocks without LEAPS	114,200
<i>Panel B: Dividend frequency of stocks with options</i>	
Any	153,589
Annual	813
Semi-annual	2,284
Quarterly	78,623
Monthly	531
None paid	71,338

Table 2: **Predictive regressions of actual dividend changes on implied dividend changes.**

Actual n -quarter dividend change is defined as the growth rate of a dividend paid n quarters ahead of the observation quarter t for a given company i : $d_{i,t+n} = \frac{D_{i,t+n} - D_{i,t}}{D_{i,t}}$. Implied dividend change is defined as the value in quarter t of the dividend implied by option prices for quarter $t+n$ relative to the actual dividend paid in quarter t : $p_{i,t,n} = \frac{P_{i,t,n} - D_{i,t}}{D_{i,t}}$. For each quarter in the data set the regression equation is estimated for one, two and three quarters ahead. The coefficient estimates are averaged over t and their t-stats are calculated by dividing the average coefficient values by their standard deviation, multiplied by the square root of the number of quarters. The t-statistics are reported in parentheses. The estimated equations are:

$$d_{i,t+n} = \alpha_t + \beta_t p_{i,t,n} + \epsilon_{i,t+n}$$

$$d_{i,t+n} = \alpha_t + \beta_{up,t} I_{d>0} \times p_{i,t,n} + \beta_{down,t} I_{d<0} \times p_{i,t,n} + \epsilon_{i,t+n}$$

	Number of quarters ahead of implied dividends					
	1		2		3	
<i>Intercept</i>	0.004 (0.97)	0.002 (0.33)	0.018 (4.95)	0.016 (4.27)	0.037 (6.26)	0.029 (5.04)
β_t	0.097 (2.38)		0.102 (6.68)		0.168 (5.88)	
$\beta_{up,t}$		0.100 (2.37)		0.105 (6.55)		0.184 (6.04)
$\beta_{down,t}$		-0.015 (-0.77)		-0.002 (-0.08)		-0.062 (-1.87)
R^2	0.009	0.011	0.011	0.014	0.012	0.016

Table 3: **Prevalence of dividend changes.**

Actual n -quarter dividend growth is defined as the growth rate of a dividend paid n quarters ahead of the observation quarter t for a given company i : $d_{i,t+n} = \frac{D_{i,t+n} - D_{i,t}}{D_{i,t}}$. The upper panel describes the number of quarterly dividend changes and their median and average size for five different types of dividend changes from 1996 to 2015 in US listed stocks that have options traded on them. The second panel shows the average dividend yield during the quarter before the quarter in which the dividend payment was made or terminated. The third panel shows the same for the quarter in which the dividend payment was made or terminated. The definitions for dividend changes made by Baker and Wurgler (2004, p 1134) partly coincide. Initiations as defined here equal new and list dividend payers, while Terminations equal new non-payers in Baker and Wurgler (2004).

<i>Panel A</i>	Increases	Decreases	Initiations	Terminations	Unchanged	Total
Number	11,564	2,357	683	66	59,211	73,881
Median change (%)	10.81	-45.71		-100	0	0
Average change (%)	24.16	-42.32		-100	0	2.46
<i>Panel B</i>	<i>Quarter before dividend change</i>					
Average Dividend Yield	0.55	0.66		0.79	0.57	0.57
StDev	0.37	0.54		0.60	0.40	0.40
<i>Panel C</i>	<i>Quarter of dividend change</i>					
Average Dividend Yield	0.62	0.48	0.44		0.57	0.57
StDev	0.38	0.43	0.39		0.41	0.41

Table 4: **Correctly and incorrectly predicted dividend changes.**

The table shows the number of instances in the data set of correctly and incorrectly predicted dividend changes (zero-changes are excluded) and the fraction of correct and incorrect predictions in brackets. Correct implied dividend predictions of non-zero changes in actual dividends are defined as follows: if an n -quarter ahead implied dividend at time t is higher than its median and the n -quarter ahead actual dividend of the corresponding firm is increased relative to the actual dividend at time t , then the implied dividend correctly predicts the increase. Correctly predicted dividend cuts are determined similarly but in the downward direction. Incorrectly predicted dividend changes occur when the implied dividend points to an increase while a decrease materializes and vice versa. Definitions of implied and actual dividend changes are as before.

	1 quarter ahead		2 quarters ahead		3 quarters ahead	
	Number	Fraction	Number	Fraction	Number	Fraction
Correct increases	4,807	0.577	8,790	0.573	11,118	0.545
Correct decreases	795	0.497	1,404	0.507	1,850	0.510
Incorrect increases	3,531	0.423	6,546	0.427	9,295	0.455
Incorrect decreases	805	0.503	1,365	0.493	1,779	0.490
Total	9,938		18,105		24,042	

Table 5: **Ordered Probit estimations of actual dividend changes on implied dividend changes.**

Ordinal values y represent latent variable y^* which is a linear function of the model: $y_{i,t}^* = \beta_{up,t} I_{d>0} + \epsilon_{i,t}$. The up-dummy $I_{d>0}$ attains a value of 1 for values of $p_{i,t,n}$ larger than zero, which means that the implied dividend is higher than the actual dividend paid in quarter t , and zero otherwise. Actual dividend changes d are ordered into categories as follows:

$$y = \begin{cases} 1 & \text{if } d \leq -0.01 \\ 2 & \text{if } -0.01 < d \leq 0.01 \\ 3 & \text{if } 0.01 < d \leq b \\ 4 & \text{if } d > b \end{cases}$$

Probabilities of falling into category $y = i$ are given by:

$$P(y = i) = P(\alpha_{i-1} < y^* \leq \alpha_i | \beta_{up} I_{d>0}) = \Phi(\alpha_i + \beta_{up} I_{d>0}) - \Phi(\alpha_{i-1} + \beta_{up} I_{d>0})$$

for $y = [1, 4]$, and where Φ is the cumulative normal distribution function and $\beta = [\beta_{up}, \beta_{down}]$. The coefficients indicate the impact of a unit change in implied dividends on the probability of one of the categories occurring.

The increase in the likelihood of falling into category i from a positive implied dividend is calculated by deducting the likelihood of a cut falling into a category i from the likelihood of an increase falling into that category: $\delta P(\alpha_{i-1} < y^* \leq \alpha_i | \beta_{up} I_{p>0}) = [\Phi(\alpha_i + \beta_{up} I_{p>0}) - \Phi(\alpha_{i-1} + \beta_{up} I_{p>0})] - [\Phi(\alpha_i) - \Phi(\alpha_{i-1})]$

The first panel describes the coefficients of the Probit estimation, t-statistics are reported in parentheses. The second panel contains the marginal effects of a move in implied dividends from negative to positive on the probability of falling into category $y = i$. Boundary b is set to 10%.

		Number of quarters ahead of implied dividends		
		1	2	3
<i>Panel A</i>	$\hat{\alpha}_{1,t}$	-1.855 (-73.4)	-1.596 (-55.09)	-1.442 (-45.91)
	$\hat{\alpha}_{2,t}$	1.077 (48.06)	0.630 (24.95)	0.297 (10.84)
	$\hat{\alpha}_{3,t}$	1.465 (54.18)	1.095 (33.2)	0.843 (21.09)
	β_{up}	0.090 (6.48)	0.080 (4.98)	0.056 (3.29)
<i>Panel B</i>		Probability of category y being chosen for negative implied dividends ($\beta_{up} = 0$)		
$d \leq -0.01$	$y = 1$	0.026	0.047	0.067
$-0.01 < d \leq 0.01$	$y = 2$	0.812	0.662	0.528
$0.01 < d \leq 0.10$	$y = 3$	0.077	0.136	0.189
$d > 0.10$	$y = 4$	0.085	0.155	0.216
<i>Panel C</i>		Marginal effect from a move in implied dividends from negative to positive		
$d \leq -0.01$	$y = 1$	-0.006	-0.008	-0.008
$-0.01 < d \leq 0.01$	$y = 2$	-0.015	-0.018	-0.014
$0.01 < d \leq 0.10$	$y = 3$	0.008	0.009	0.006
$d > 0.10$	$y = 4$	0.013	0.018	0.016

Table 6: **Stock price response to dividend announcements.**

The dependent variable $CAR_{i,t}$ is the stock price response to dividend change announcements, defined as the return of a stock on which a dividend change is announced during the three days around the announcement [-1,1]. Actual n -quarter dividend change is defined as the growth rate of a dividend paid n quarters ahead of the observation quarter t for a given company i : $d_{i,t+n} = \frac{D_{i,t+n} - D_{i,t}}{D_{i,t}}$. Implied dividend change is defined as the value in quarter t of the dividend implied by option prices for quarter $t+n$ relative to the actual dividend paid in quarter t : $p_{i,t,n} = \frac{P_{i,t,n} - D_{i,t}}{D_{i,t}}$. Dummies $I_{d>0}/I_{d<0}/I_{p>0}/I_{p<0}$ take on the value 1 when actual (D) or implied (P) dividends are increased/cut and are 0 otherwise. The regressions are performed by including either dummies or the interaction of the parameters with dummies. The t-statistics are reported in parentheses. The estimated equations are:

$$CAR_{i,t} = \alpha + \beta_1 I_{d<0} + \beta_2 I_{d>0} \times d_{i,t} + \beta_3 I_{d<0} \times d_{i,t} + \beta_4 I_{p<0} + \beta_5 I_{d>0} \times I_{p>0} + \beta_6 I_{d>0} \times I_{p<0} + \beta_7 I_{d<0} \times I_{p>0} + \beta_8 I_{p<0} \times I_{p<0} + \epsilon_{i,t}$$

	(1)	(2)	(3)	(4)
<i>Intercept</i>	0.007 (6.82)	0.006 (4.62)	0.005 (3.52)	
$I_{d<0}$	-0.023 (-11.38)	-0.007 (-1.35)	-0.006 (-1.25)	
$I_{d>0} \times d$		0.003 (1.74)	0.003 (1.71)	
$I_{d<0} \times d$		0.030 (3.49)	0.030 (3.49)	
$I_{p<0}$			0.002 (0.9)	
$I_{d>0} \times I_{p>0}$				0.006 (2.03)
$I_{d>0} \times I_{p<0}$				0.007 (4.08)
$I_{d<0} \times I_{p>0}$				-0.026 (-8.81)
$I_{d<0} \times I_{p<0}$				-0.002 (-0.52)
R^2	0.021	0.024	0.024	0.016

Figure 1: Dividend yields are averaged over the 1996-2015 data set for a given number of quarters ahead of the observation quarter for all firms (value-weighted) with and without LEAPS traded on their stocks.

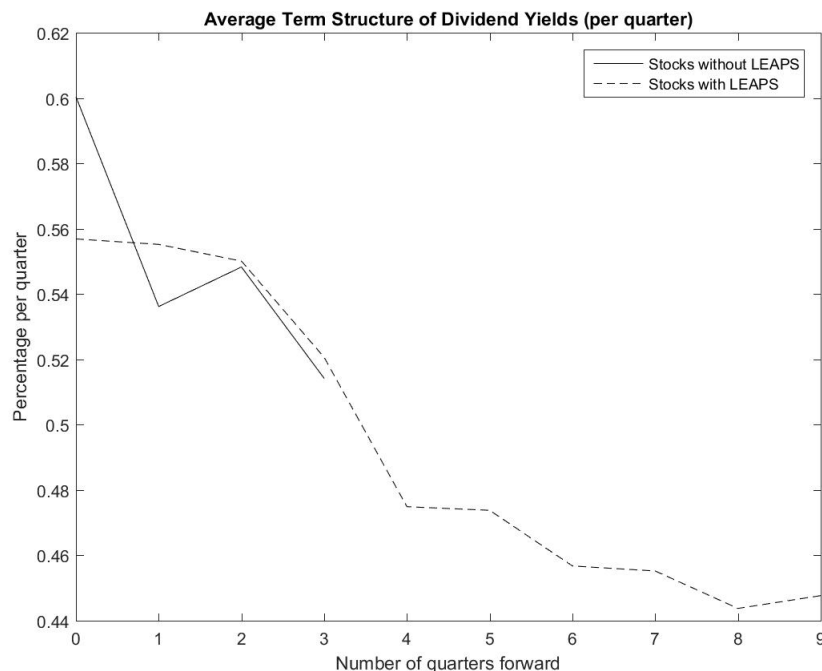


Figure 2: Dividend yields per quarter, current, 1 quarter ahead and 1 to 2 quarters ahead. Yields are value-weighted over all firms who pay dividends at a quarterly rate.

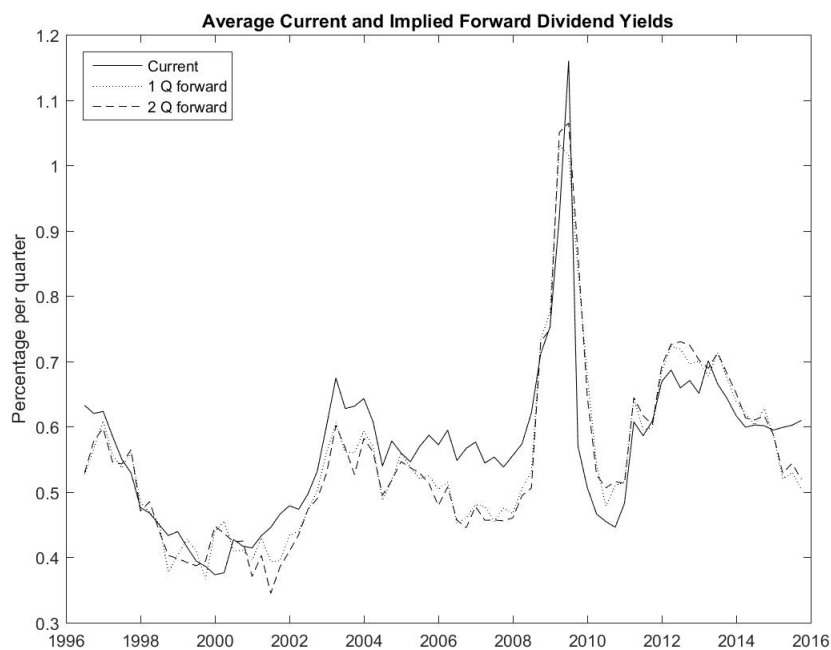


Figure 3: Dividend yields per quarter, current, 1 to 3 quarters ahead and 1 to 8 quarters ahead. Yields are value-weighted over all firms who pay dividends at a quarterly rate.

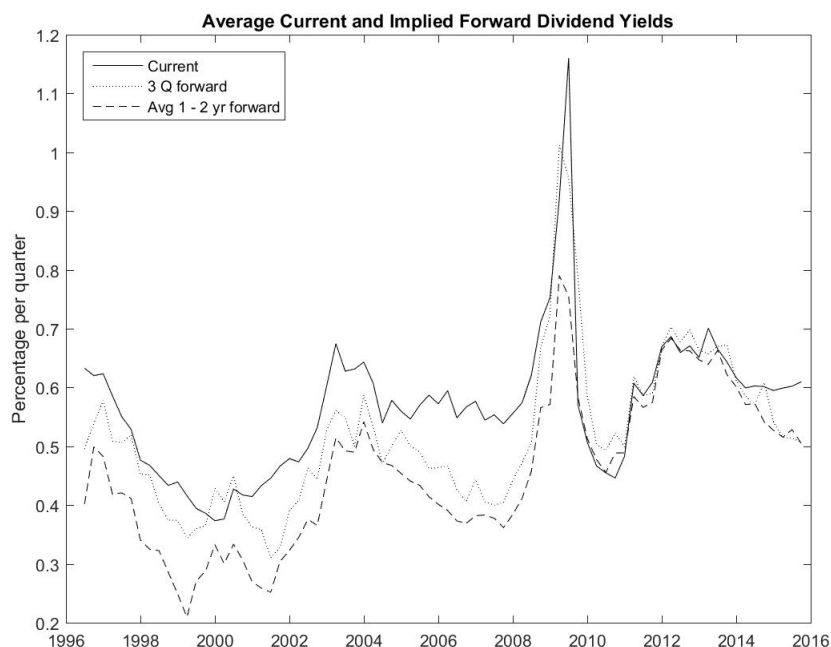


Figure 4: Growth rate for dividends 1 quarter ahead and 1 to 2 quarters ahead relative to current dividends. Growth rates are value-weighted over all firms who pay dividends at a quarterly rate.

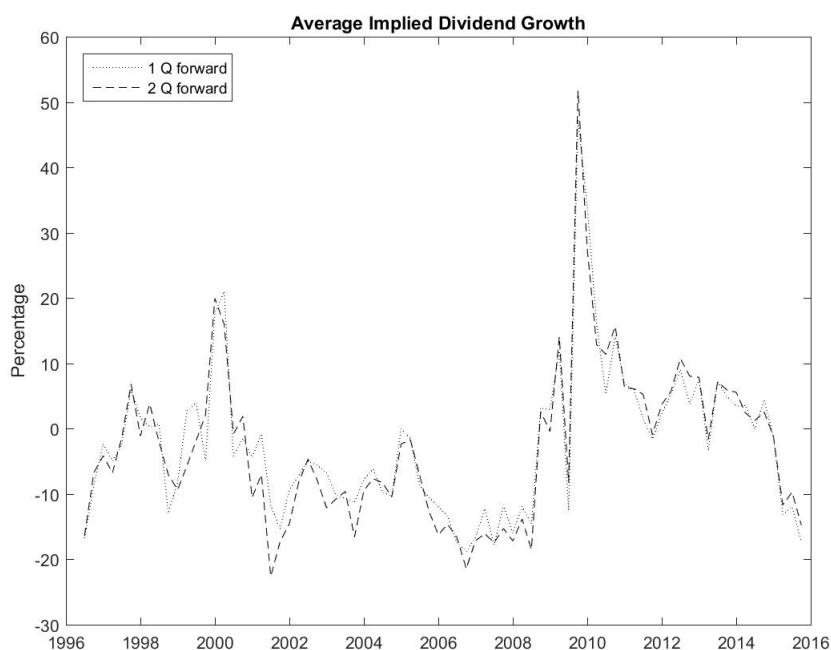


Figure 5: Growth rate for dividends 1 to 3 quarters ahead and 1 to 8 quarters ahead relative to current dividends. Growth rates are value-weighted over all firms who pay dividends at a quarterly rate.

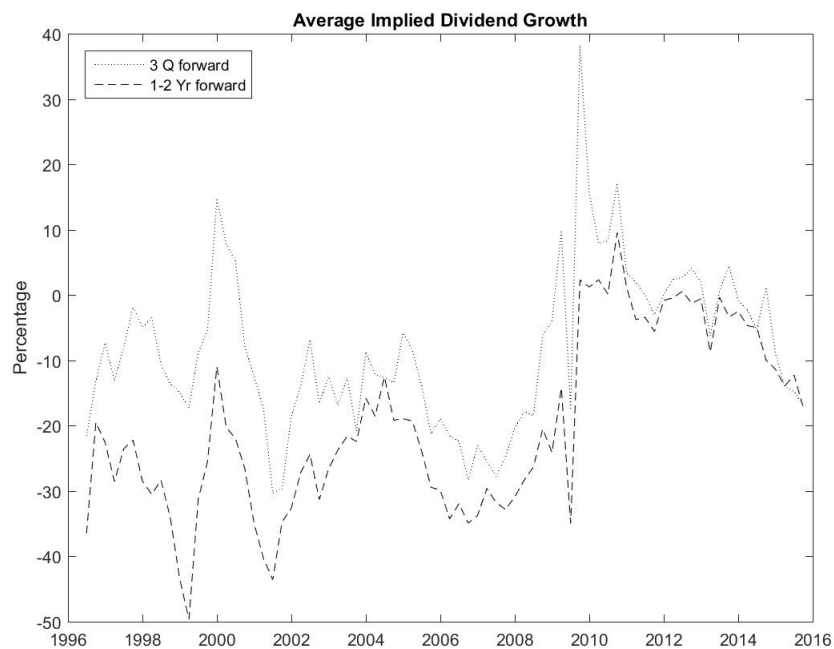


Figure 6: Decile breakpoints of current quarterly dividend yields for all quartercompanies.

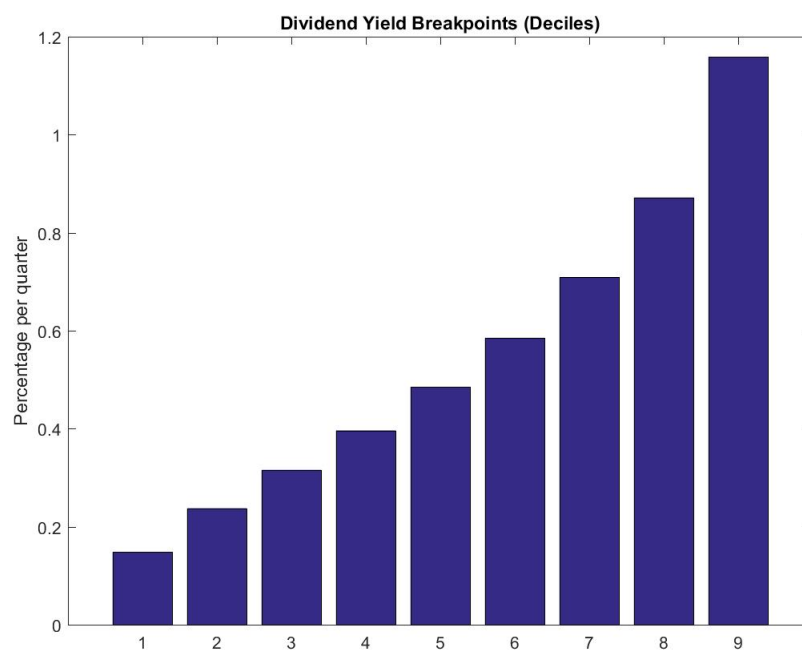


Figure 7: Decile breakpoints of implied dividend growth rates for 1 quarter ahead and 1 to 2 quarters ahead for all quartercompanies.

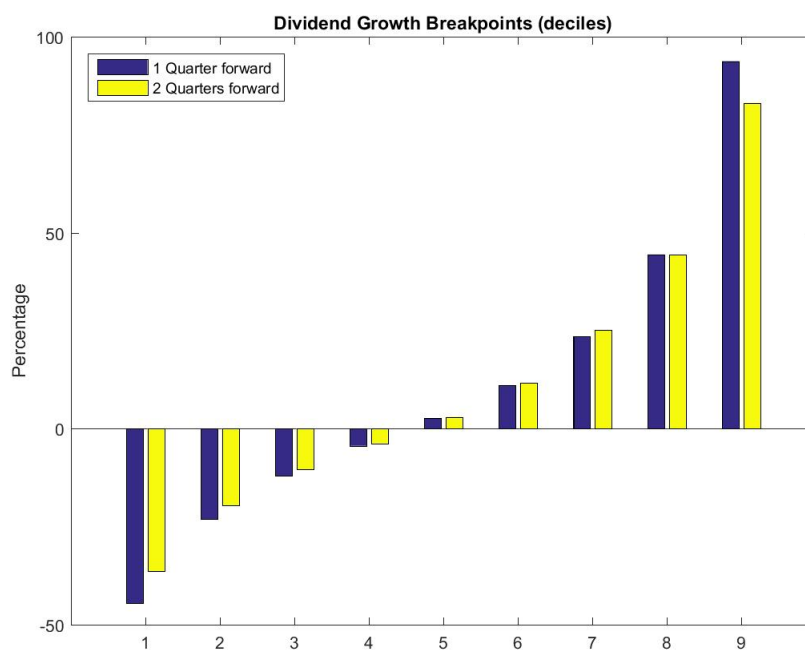


Figure 8: Decile breakpoints of implied dividend growth rates for 1 to 3 quarters ahead and 1 to 8 quarters ahead for all quartercompanies.

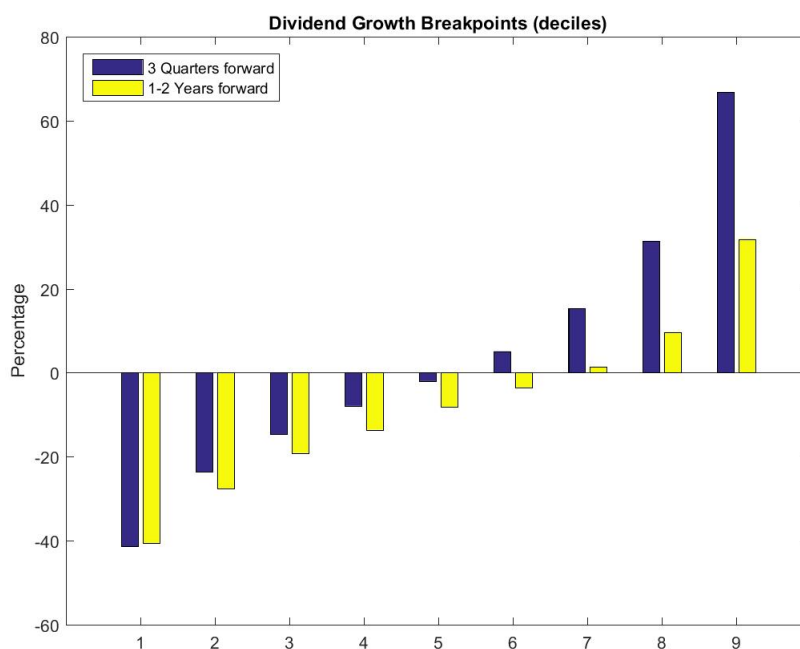


Figure 9: Apple Inc. dividends as implied by the LEAPS expiring in January 2013, normalized to a quarterly rate.

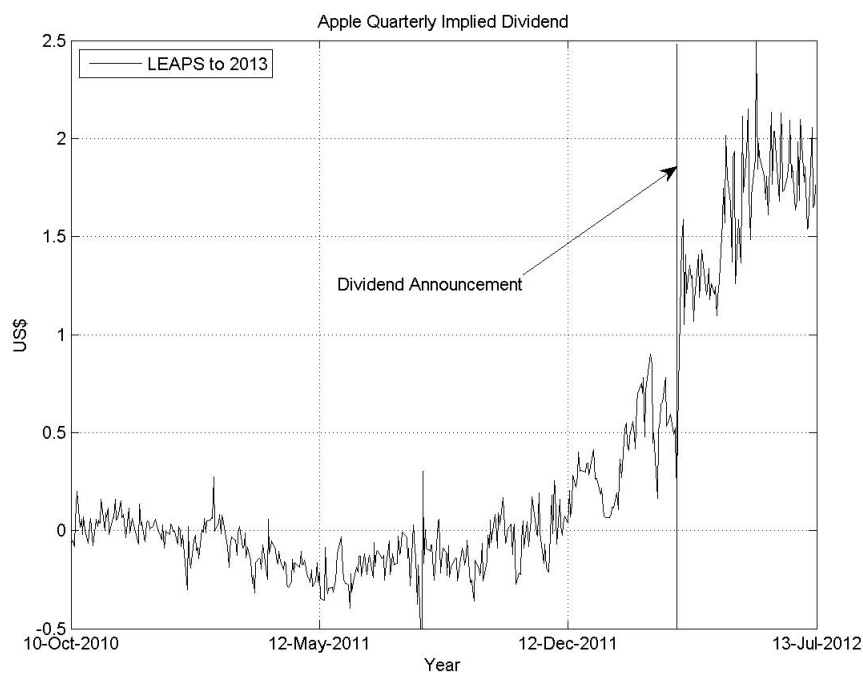


Figure 10: Apple Inc. dividends as implied by the stock option expiring in July 2012, normalized to a quarterly rate.

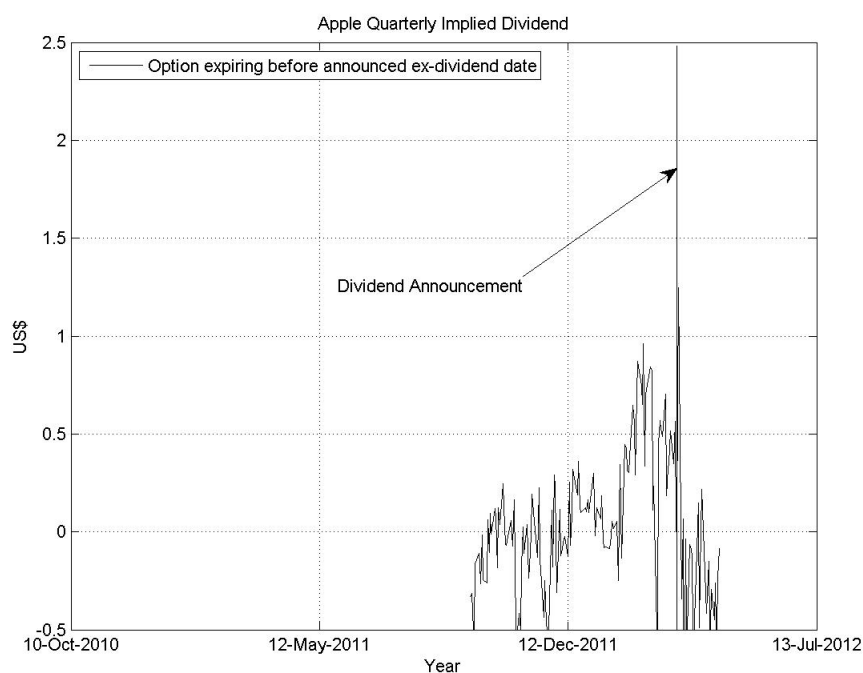


Figure 11: Implied stock loan fees 1 quarter ahead and 1 to 2 quarters ahead, normalized to bppa. Fees are value-weighted over all firms who pay dividends at a quarterly rate.

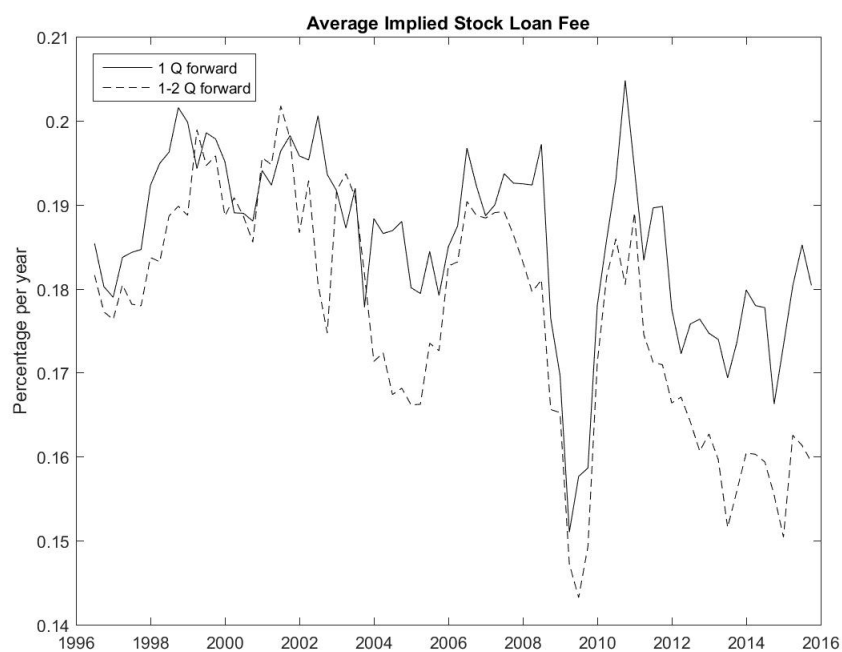


Figure 12: Implied stock loan fees 1 to 3 quarters ahead and 1 to 8 quarters ahead, normalized to bppa. Fees are value-weighted over all firms who pay dividends at a quarterly rate.

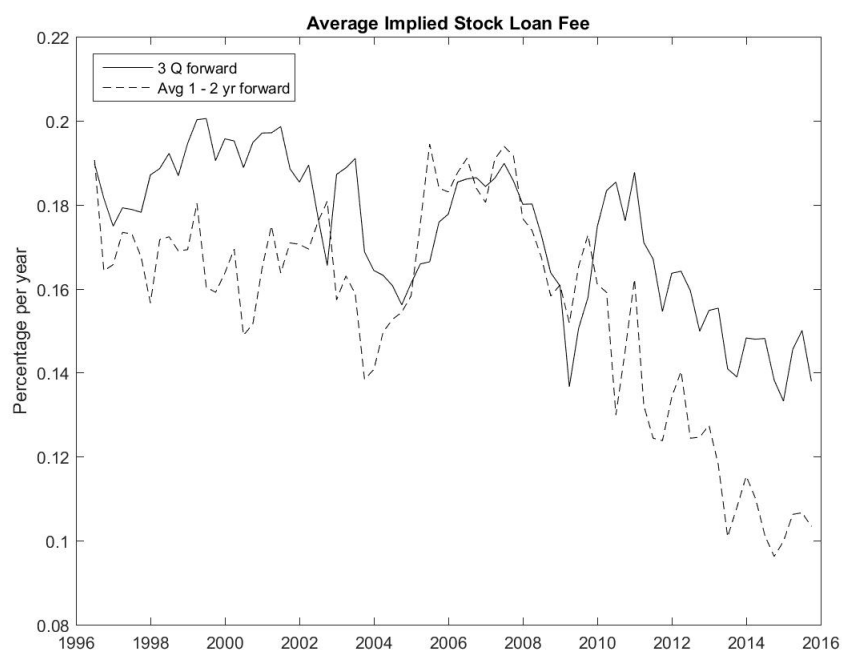


Figure 13: Stock loan fees are averaged over the 1996-2015 data set for a given number of quarters ahead of the observation quarter for all firms (value-weighted) with and without LEAPS traded on their stocks.

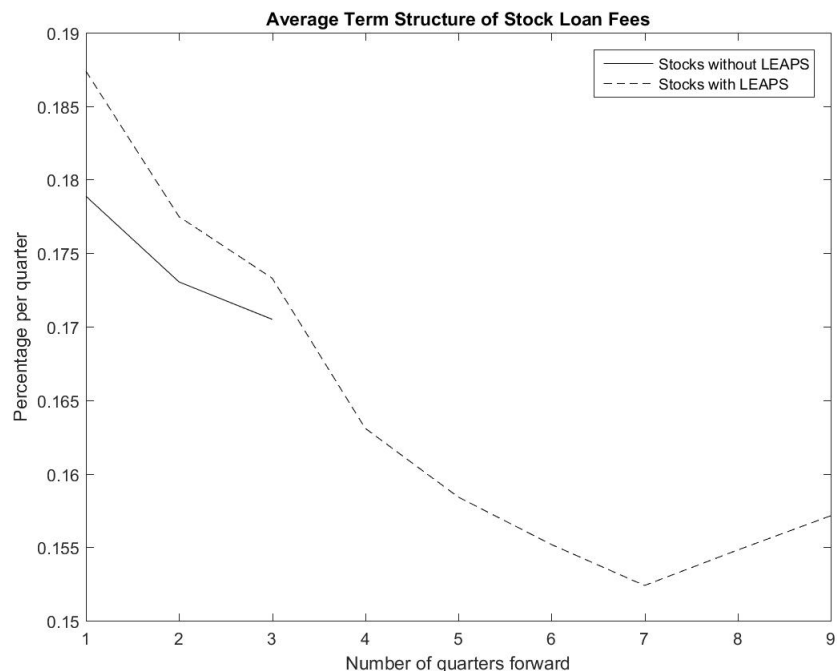


Figure 14: Decile breakpoints of implied stock loan fees for 1 quarter ahead, 1 to 2 quarters ahead and for 1 to 3 quarters ahead for all quartercompanies.

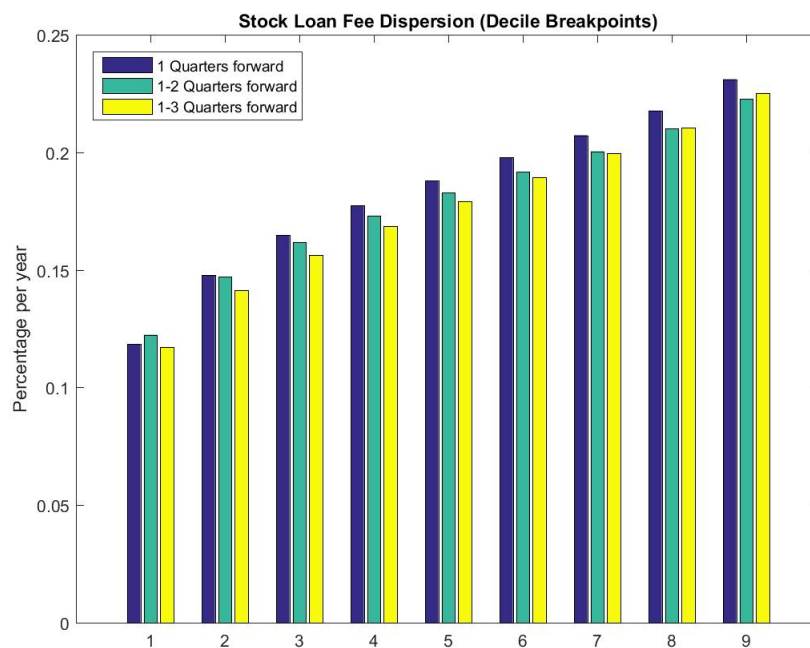


Figure 15: Contingency chart of actual dividends and implied dividends (2 quarters ahead). A dot indicates an actual change in dividends against a change as implied by the data $P_{i,t+n}$.

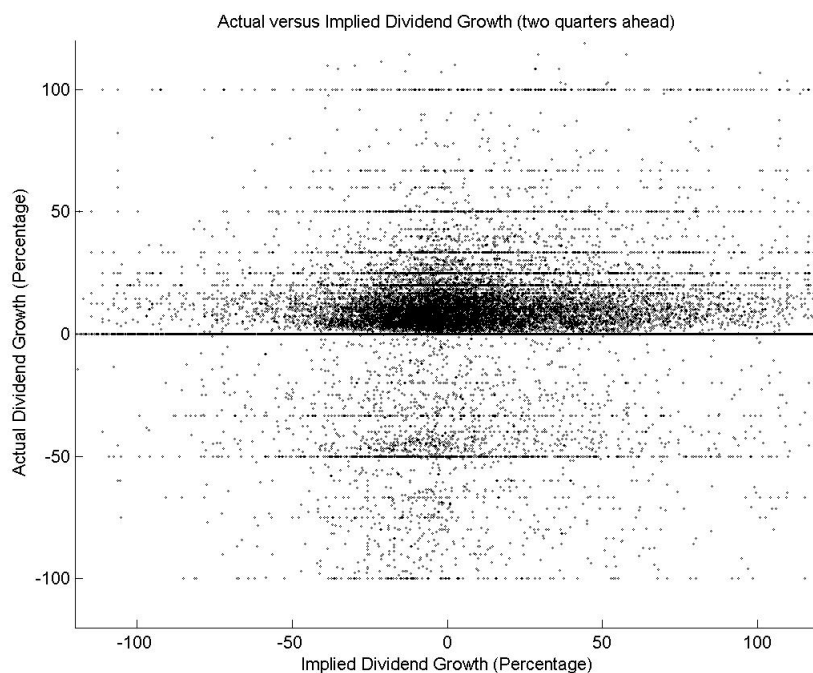


Figure 16: Increase in the likelihood of the small dividend raise d category $y = 3$, which is defined as $0.01 < d \leq b$, occurring due to a change in implied dividends moving from negative to positive. Marginal effects for 1, 2 and 3 quarters ahead.

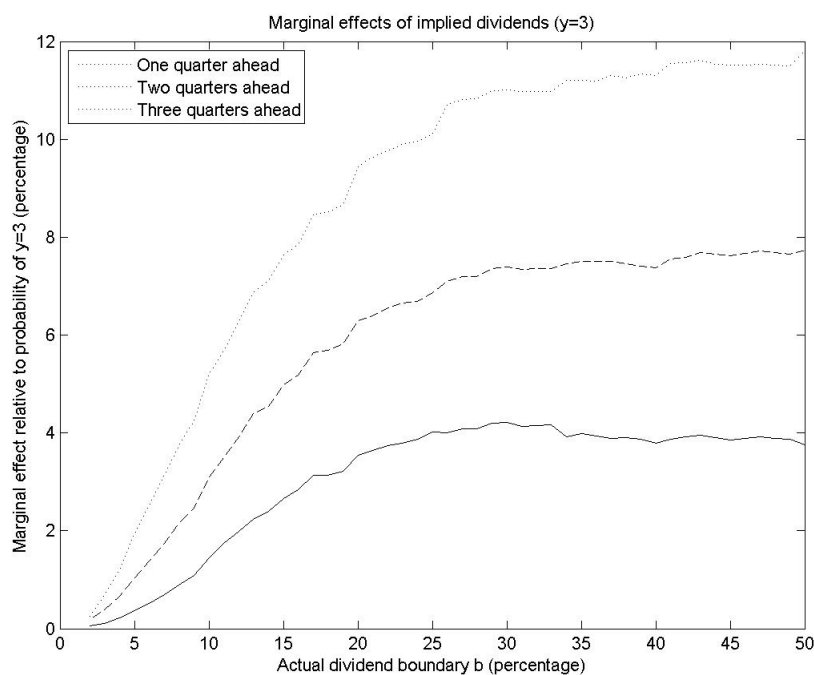


Figure 17: Increase in the likelihood of the small dividend raise d category $y = 4$, which is defined as $d > b$, occurring due to a change in implied dividends moving from negative to positive. Marginal effects for 1, 2 and 3 quarters ahead.

