Early exercise decision in American options with dividends, stochastic volatility and jumps

Antonio Cosma[*], Stefano Galluccio[**], Paola Pederzoli[***], and Olivier Scaillet[***]

[*]University of Luxembourg

[**]Incipit Capital, London

[***]University of Geneva and Swiss Finance Institute

Abstract

Using a fast numerical technique, we investigate a large database of investor suboptimal non-exercise of short maturity American call options on dividend-paying stocks listed on the Dow Jones. The correct modelling of the discrete dividend is essential for a correct calculation of the early exercise boundary as confirmed by theoretical insights. Pricing with stochastic volatility and jumps instead of the Black-Scholes-Merton benchmark cuts by a quarter the amount lost by investors through suboptimal exercise. The remaining three quarters are largely unexplained by transaction fees and may be interpreted as an opportunity cost for the investors to monitor optimal exercise.

We would like to thank the editor and the referee for constructive criticism and numerous suggestions which led to substantial improvements over the previous version. We also thank J. Detemple, D. Duffie, E. Gobet, J. Jackwerth, F. Moraux, D. Newton, A. Treccani and A. Valdesogo for valuable insight and help in addition to the participants at the 6th World Congress of the Bachelier Finance Society in Toronto, the 2011 European Econometric Society in Oslo, the 2013 Mathematical Finance Day in Montreal, the 2014 Mathematical and Statistical Methods for Actuarial Sciences and Finance in Salerno, the 2014 International Symposium on Differential Equations and Stochastic Analysis in Mathematical Finance in Sanya, the 7th General Advanced Mathematical Methods in Finance and Swissquote Conference in Lausanne, the 2015 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics in Cape Town, the 9th International Conference on Computational and Financial Econometrics in London, the 2015 International Conference on Computational Finance in London, SGF Conference 2016 in Zurich, the FERM 2016 in Guangzhou, the AFFI 2016 in Liège, the Stochmod16 in Louvain-la-Neuve, the 9th World Congress of the Bachelier Society in New York, the EFA 2016 in Oslo, and seminars at the University of Geneva and the University of Orléans. O. Scaillet received support from the Swiss NSF through the NCCR Finrisk. Paola Pederzoli acknowledges the financial support of the Swiss NSF (grant 100018 – 149307). Part of the research was made when she was visiting LSE. A previous version of this paper circulated under the title “Valuing American options using fast recursive projections”.

1. Introduction

Holders of short maturity American call options on dividend-paying stocks are known to miss exercising their options in an apparently suboptimal way (see e.g. Pool, Stoll, and Whaley (2008)). We investigate the suboptimal exercise by allowing alternative models for the underlying stocks dynamics. To do this, we compile data on 30 individual dividend-paying stocks listed on the Dow Jones, comprising a total of 101,295 series of short-term options amounting to approximately 9.5 million records. We show that, by taking into account stochastic volatility and jumps in the process of the underlying asset, we can explain up to 25% of the gain forgone due to non-optimal exercise decisions, as computed in Pool et al. (2008). Because financial frictions are a possible explanation of departure from the expected exercise behavior (e.g. Jensen and Pedersen (2016)), we also show that transaction costs cannot fully explain the non-exercise decisions.

Options written on single stocks are mainly of the American style. This means we can exercise them at any time before maturity. Economic theory gives a clear guidance on the early exercise decision: an investor should exercise the options if the proceeds from immediate exercise exceed the value of the contract when kept alive. The value of the underlying stock that makes the investor indifferent between holding onto the option or not, is the exercise boundary, or frontier. Such a clear normative dictate is nevertheless accompanied by a strikingly high level of suboptimal non-exercise. The figures reported in Pool et al. (2008), and confirmed in our data, show that, in given periods, investors hold onto between one third and one half of the contracts they should have exercised. This behavior can potentially be explained in terms of costs, whether explicit, as fees and other forms of financial frictions, or implied, as the opportunity cost of closely monitoring the option. Since the income in case of immediate exercise, that is the intrinsic value of the option, is readily known, the main determinant of the decision is the continuation value, that is the value of the option when it is not exercised. Our analysis focuses on the correct estimation of the latter. As already noticed in Pool et al. (2008) and in Barraclough and Whaley (2012), there are two ways of addressing this issue. One is the market-based approach. Market prices provide guidance in the exercise decision. If the market time value of the option, that is the difference between the market and intrinsic values, vanishes, then it is time for the investor to exercise. In practice, though, due to price discreteness and bid-asks spreads, the time value may not be exactly equal to zero, even when exercise is optimal. We prefer a model-based estimation. With this approach, we obtain an explicit figure for the continuation value, and whenever exercise would have been optimal, we can compute how much lower than the intrinsic value it is, thus quantifying the loss due to suboptimally holding onto the contract. This loss gives us an estimate of the opportunity cost or of the implied fees. Market prices are silent on this aspect, since the market time value of the option becoming zero only tells us that the continuation has reached from above the intrinsic value, not that it has crossed it, or, if it is the case, by how much. The model-based continuation value depends on the assumptions on the underlying stock process, and this leads us to check how the exercise boundary changes in different modelling environments. The crucial modelling assumption in this context is how to represent the dividend. The most recent theoretical literature on American options (Amin (1993), Adolfsson, Chiarella, Ziogas, and Ziveyi (2013)) models the dividend payment as a continuous dividend yield. We see in the following that when assuming a constant dividend yield, the introduction of jumps and stochastic volatility in the stock process may exacerbate, instead of partially explaining, the suboptimality figures.

In our empirical investigation, we show that by correctly modelling the dividend as a discrete payment, the boundary is higher under the Merton (1976) jump-diffusion and Heston (1993) stochastic volatility models than under the Black-Scholes model. In the subsequent discussion, we provide new theoretical insights on why this is the case. This allows us to revise the share of contracts that are suboptimally unexercised, since contracts that should be exercised under Black-Scholes, should be instead kept alive under the more sophisticated models. In our calibration, we price by fully taking into account the discrete nature of the dividend distributed by the underlying stocks and the American style of the call options, and we do so for different specifications of the stock dynamics. In the process, to our knowledge, we are the first to provide comprehensive descriptive statistics of the parameters driving the jumps and the stochastic volatility of the constituents of the Dow Jones Industrial Average Index (DJIA) traded in the period from January 1996 to December 2012. This feature is a peculiarity of our work, given that the standard empirical literature on options mainly focuses on European $S\&P500$ options with a dividend yield (Bakshi, Cao, and Chen (1997), Eraker, Johannes, and Polson (2003)). Broadie, Chernov, and Johannes (2007) and Broadie, Chernov, and Johannes (2009) approximate American prices with European ones, and show that transforming American options to European ones does not matter for calibration purposes when facing a continuous dividend yield since differences in early exercise premia are not so large in that case. This is not true with multiple discrete dividend payments, and we provide an example on how neglecting the discrete character of a dividend or its time of payment leads to an incorrect exercise decision. Overlooking this feature biases empirical findings on suboptimality.

To conclude this introduction, we point to the methodological contribution of our work. One of the reasons that has kept researchers from correctly modelling the dividend as a discrete process, is that the repeated option valuations needed to calibrate the different models for the underlying stock to such a large dataset (we have approximately 9.5 million records) are too time consuming with the existing numerical methods. To date, no empirical work on options written on dividend-paying stocks exists outside the Black-Scholes world[1]. We develop a new methodology to price American options, which, in the kind of pricing problems we face in our empirical study, is more than one order of magnitude faster than the techniques based on the discretization of the pricing partial differential equation, and four orders of magnitude faster than simulation-based techniques. Our technique belongs to the quadrature family of numerical routines (see Andricopoulos, Widdicks, Newton, and Duck (2007) and Chen, Härkönen, and Newton (2014), Fang and Oosterlee (2011)) and we refer to it as *Recursive Projections* to distinguish it from other variants. We are the first to characterize the convergence properties of a quadrature-based method in the presence of discrete dividends and with the underlying

[1]Broadie et al. (2007), state that "The computation time required for American options makes calibration to a very large set of options impractical." As reported earlier in the text, what Broadie et al. (2007) show in their Appendix A is that transforming American options into European ones does not matter for calibration purposes in their application, but we show in the following that it does make a difference in our study.

following a dynamics outside the Black-Scholes benchmark[2].

The paper is organised as follows. In Section 2, we give a numerical example of how introducing jumps and stochastic volatility in the stock process leads to a higher exercise boundary. This allows us to conjecture that by taking into account these features of the underlying process, the figures on suboptimal exercise should decrease. This is what we test, and find, in Section 3.2. Section 3.3 shows that transaction fees cannot reasonably explain the part of suboptimality in the exercise behavior that remains after taking into account jumps and stochastic volatility. In Section 3.4, we give a theoretical justification for the findings of Section 3.2. Section 4 provides some concluding remarks. Appendices A-C provide the detailed description of the pricing algorithms used in our empirical strategy, and gather further theoretical insights of the behavior of the exercise boundary under the different modelling environments[3].

2. Early exercise decision: a numerical example

In the example that follows, we show how the two most common assumptions made when investigating exercise decisions for American call options written on dividend-paying stocks, namely a continuous payment of the dividend, and a geometric Brownian motion for the underlying stock process, reinforce each other in introducing a downward bias in the computation of the exercise frontier. This negative bias increases in turn the number of contracts that appear to be irrationally kept alive. We design this example to provide an intuitive introduction to the impact of modelling choices on the exercise decision of contracts. A more formal treatment is left for the next section.

In our example, the call option has a remaining life of 6 months, that is the maturity is $T = 0.5$, and the underlying stock distributes a regular quarterly dividend d. The two dividend payments occur immediately, at $t_1 = 0$, and at $t_2 = 0.25$, corresponding to a time to maturity of 0.5 and 0.25, respectively. If we model correctly the dividend payment as an event that takes place at a precise point in time, then it is well known that it may be optimal to exercise the

[2]Our pricing library is publicly available on the website of the authors with a ready-to-use Matlab interface.
[3]The online supplementary materials gather additional technical details.

call option only immediately before the dividend payment (see for instance Hull (2011)). In this example, it means that we may exercise the option at times $t = 0$, and $t = 0.25$. It follows that the value of the call option at $t = 0.25$ is the maximum between the continuation value, that is the value of the option just after the dividend is paid if we do not exercise, and the intrinsic value, that is the proceeds from exercising immediately before the dividend payment. Formally, $C(S_t, T, K)$ is the value of the American option as a function of the underlying stock price S_t at time t, the strike K, and maturity T. Then, just before the stock goes ex dividend, that is an infinitesimal time interval ε before $t = 0.25$, it holds:

$$C(S_{0.25}, 0.5, K) = \max\{S_{0.25-\varepsilon} - K, C(S_{0.25} - d, 0.5, K)\}. \tag{1}$$

We follow the accepted convention that no jumps other than the drop in value of S_t due to the dividend payment occur at the dividend date, so that by continuity of the stock process, $S_{t-\varepsilon} = S_t$, and $S_t - d$ is the value of the stock immediately after the payment of the dividend[4]. The value $S^*_{0.25}$ such that $S^*_{0.25} - K = C(S^*_{0.25} - d, 0.5, K)$ is the exercise boundary at $t = 0.25$. For values $S_{0.25} \geq S^*_{0.25}$, it is optimal to exercise the option, while, for $S_{0.25} < S^*_{0.25}$, the optimal behavior is to keep the option until maturity. Since no further intermediate payment is foreseen before maturity, $C(S_{0.25} - d, 0.5, K)$ is the value of a European option:

$$C(S_{0.25} - d, 0.5, K) = \mathbb{E}[e^{-r(0.5-0.25)}(S_T - K)|S_t = S_{0.25} - d]. \tag{2}$$

The expectation in Equation (2) is taken with respect to the risk neutral transition density of the stock value from $t = 0.25$ to $t = 0.5$. For most of the stochastic processes chosen to model the stock price, Equation (2) has a closed or quasi-closed form solution. No arbitrage and inter-temporal consistency of asset prices dictate that at, $t = 0$, we have to compare the value of immediate exercise with the continuation value after the dividend payment, which now takes

[4]Whenever the dividend d does not appear in the argument of the call value, as in $C(S_t, K, T)$, we are considering the value of the call before the stock goes ex-dividend, that is after the maximum in Equation (1) is taken. If d appears explicitly, as in $C(S_t - d, T, K)$, we are referring to the continuation value of the call, after the dividend payment.

the more complex form:

$$C(S_0 - d, 0.5, K) = \mathbb{E}[e^{-r(0.25)}C(S_{0.25}, 0.5, K)|S_t = S_0 - d]. \tag{3}$$

This time the expectation is taken with respect to the risk neutral transition density of the stock value from $t = 0$ to $t = 0.25$. The value of the contract at $t = -\varepsilon$ is $\max\{S_{0-\varepsilon} - K, C(S_0 - d, 0.5, K)\}$, and the exercise boundary S_0^* is the value such that $S_0^* - K = C(S_0^* - d, 0.5, K)$. Contrary to Equation (2), the right hand side of Equation (3) in most cases does not have a simple-to-compute form. The Fast Recursive Projections, the numerical methodology we use in the empirical work of Section 3, takes fully into account the discrete nature of the dividend payment, and gives a consistent estimate of $C(S_0 - d, 0.5, K)$. We determine below the exercise boundaries S_t^* at $t = 0$ and $t = 0.25$ by computing the expectations in Equations (2) and (3) using the risk neutral transition densities implied by a pure diffusion model for the underlying stock (Black-Scholes), a stochastic volatility model (Heston) and a process allowing for jumps at discrete moments in time (Merton). A formal definition of the processes is left to Section 3.1. For the moment, we are only interested in giving a qualitative overview of how the exercise boundary changes as a consequence of the characterising features of the three process specifications.

We may be tempted to simplify the problem by approximating the discrete dividend payout process by a continuous dividend payment. In this modelling approach, there are no more a priori known dates when exercise could be optimal. Exercise can be optimal in any moment t between 0 and 0.5, whenever the stock price is deep enough in the money and higher than the critical value S_t^*, which is the exercise boundary for an American call with a continuous dividend yield. The boundary S_t^* is defined as the lowest value of S_t such that $S_t - K \geq C(S_t, T, K)$. In the continuous dividend approximation, there is therefore a value of S_t^* for each moment in time until maturity, whereas, in the discrete dividend case, we only had 2 possible exercise dates. The continuous dividend approach may not appear as a simplification since the comparison between the intrinsic value and the continuation value as in (3) has to be done at each point in time. In practice, the comparison takes place on a finite number of points between 0 and T and the most common numerical methods, such as trees and finite difference methods, adapt much more easily to a continuous dividend yield (that preserves the recombining property of these schemes) than to a process that is discontinuous at each dividend date. In Panel A of Figure 1, we plot the early exercise boundary for the Heston and Black-Scholes models for an American call option with a continuous dividend yield $r_d = 0.03$ (right graph) and with an equivalent quarterly discrete dividend $d = 1.38$ (left graph). We choose $d = 1.38$ to have an equivalent total annual dividend between the continuous dividend yield $r_d = 0.03$ and the discrete dividend case. Indeed, $1.38 = 0.03S^*/4$, where $S^* = 184$ is the critical stock price under the Black-Scholes model in the dividend yield case for maturity $T = 0.5$. The strike price is $K = 100$. We fix the parameters of the underlying processes at values commonly used in the studies of options written on single stocks (see Adolfsson et al. (2013))[5]. With a continuous dividend yield, the Heston early exercise boundary is always below the Black-Scholes boundary, whereas, with discrete dividend, we face the opposite. Indeed, in the only two points in time at which the option may be exercised, just before the dividend payments, the value of the exercise boundary is lower under the Black-Scholes model than under the Heston model. In comparing the exercise boundary under the Black-Scholes and the Merton models we slightly change the value of the contract parameters to make our discussion comparable with the one in Amin (1993). Maturity is still $T = 0.5$ and the dividends are distributed on $t = 0, 0.25$. The strike is now $K = 40$. In the Panel A of Figure 2, we plot the early exercise boundary for the Merton and Black-Scholes models for an American call option with a continuous dividend yield $r_d = 0.05$ (right graph) and in the case in which the stock pays an equivalent quarterly discrete dividend[6] $d = 1.125$ (left graph). For a discrete dividend, we observe a higher boundary in the Merton case than in the Black-Scholes case, which holds true for all maturities. For the continuous dividend case, for short maturities, the boundary is higher under the Merton model than under the Black-Scholes

[5]We use the following set of representative parameters: $T = 0.5$, $K = 100$, $r = 0.05$, $\sigma_0 = 0.2$, $\omega = 0.1$, $\sigma_{LT} = 0.3$, $\beta = 4$, and $\rho = -0.5$ as in Adolfsson et al. (2013). See Section 3.1 for the definition of the process parameters. For comparison, we follow Heston (1993), and we use the Black-Scholes model with a volatility parameter that matches the (square root of the) variance of the spot return over the life of the option in the Heston model.

[6]As before, we take $d = 1.125$ because $1.125 = 0.05S^*/4$, where $S^* = 90$ is the critical stock price with the dividend yield $r_d = 0.05$ for maturity $T = 0.5$. The process parameters for the Merton model are $r = 0.08$, $\gamma = 5$, $\sigma_M^2 = 0.05$, $\sigma_\psi^2 = 0.05$, $\mu_\psi = 0$. See Section 3.1 for the definition of the process parameters. We set the volatility parameter in the Black-Scholes model equal to the volatility of the underlying return over the life of the option in the Merton model.

model. For longer maturities, the boundary takes lower values in the Merton case than in the Black-Scholes case, and we observe a crossing of the early exercise boundary in the far right end of the right graph of Panel A, Figure 2. The exact point where the crossing between the exercise frontiers under Black-Scholes and Merton happens depends on the configuration of the process parameters. Nevertheless, this example shows how, for a choice of the parameters typically used in the literature on calls on dividend-paying stocks, the crossing takes place in the range of maturities that we use in the empirical investigation.

The example described in this section shows that by not correctly modelling the discrete nature of the dividend payment, and by not taking into account important features of the underlying process such as jumps and stochastic volatility, we introduce a bias in the determination of the exercise boundary[7]. When the discrete dividend is modelled correctly, the bias introduced by the simple Black-Scholes model is in the direction of lowering the value of the boundary. A decision of non-exercise that would appear not optimal under the Black-Scholes model, may appear much less so under a more realistic Heston or Merton model. This is why we expect that the figures of call options left suboptimally non exercised will decrease by adding jump and stochastic volatility in the process of the stock. This is what we test empirically, and justify theoretically, in the next section.

3. Empirics and theoretical insights

In this section we investigate the early exercise decision of call holders in light of the different values that the exercise boundary can take under distinct modelling assumptions for the underlying asset. Following the procedure suggested by Pool et al. (2008), we first check which contracts should be exercised by comparing the intrinsic value immediately before the dividend payment with the continuation value on the ex-dividend day. We quantify how much is economically lost in the case of a suboptimal non-exercise decision. This amount depends on the continuation value and is model-specific. We compare the results obtained under three

[7] In Appendix A we provide an additional example on how approximating American option prices by European option prices when the stock distributes a dividend biases downwards the estimation of the exercise boundary.

modelling environments that we formally define in the next paragraph, namely, Black-Scholes, Merton jump-diffusion, and Bates. We observe a lower degree of incorrect exercise decisions if we model the underlying security in the Merton or Bates framework, than if we restrict ourselves to the Black-Scholes dynamics. This suggests that investors incorporate features of the more sophisticated models when taking the exercise decisions. Finally, whenever we find evidence of a suboptimal non-exercise decision, we show that trading costs alone cannot justify the behavior of investors.

3.1. Data and modelling assumptions

The empirical analyses of this section are carried out using a large dataset of call options having a maturity of less than six months and written on the dividend-paying stocks that enter the Dow Jones Industrial Average Index (DJIA). The sample comprises daily observations between January 1996 and December 2012. The daily data on all option attributes, the stock price, and the dividend distribution details are from Optionmetrics. We obtain the daily data on the interest rates from the Treasury constant maturities of the H15 report of the Federal Reserve[8]. A total of 101,295 series of short-term options written on 30 stocks enter our database. The total number of records is approximately 9.5 million. Table 1 reports the number of quotes for each stock with a breakdown by maturity and moneyness. Our study focuses on the early exercise behavior of investors; hence, we focus on the in-the-money options, which are the category of options for which the number of quotes is the highest.

We now give formal definitions of the stochastic processes we use to model the evolution of the underlying stock. Our choice of environments follows the empirical findings of Bakshi et al. (1997), who suggest that jumps and stochastic volatility play a dominant role in pricing short-term options whereas modelling stochastic interest rates does not seem to significantly improve the pricing performance. We therefore choose a general model that allows for jumps in the mean of the process, and for time-varying variance. Bates (1996) was the first to suggest combining these features of the underlying process, and therefore, we refer to this process specification as

[8] We access the Optionmetrics and the H15 databases through the Wharton Research Data Services (WRDS) research platform.

the Bates model. Let $X_t = \log(S_t)$, then:

$$dX_t = (r - r_d - \gamma \upsilon - \frac{1}{2}\sigma_t^2)dt + \sigma_t dW_{1,t} + log(\psi)dq_t, \tag{4}$$
$$d\sigma_t^2 = \beta\big(\sigma_{LT}^2 - \sigma_t^2\big)dt + \omega\sqrt{\sigma_t^2} \cdot dW_{2,t}, \quad E(dW_{1,t} \cdot dW_{2,t}) = \rho dt.$$

In the specification (4), r is the risk-free interest rate, r_d is the continuous dividend yield paid by the asset, $dW_{1,t}$ and $dW_{2,t}$ two correlated Wiener processes, σ_t the time-varying volatility of stock returns. In the calibration, we impose $r_d = 0$, but we leave it explicitly in model (4) to give the general specification that we used in the numerical exercise of Section 2. The Poisson process, $q(t)$, is independent of $W_{i,t}, i = 1, 2$, and there is a probability γdt that a jump occurs in dt, and $1 - \gamma$ probability that no jump occurs. The parameter γ represents the mean number of jumps per unit of time. The random variable ψ is such that $\psi - 1$ describes the percentage change in the stock price if the Poisson event occurs, and $\upsilon = \mathbb{E}[\psi - 1]$ is the mean jump size. We further make the standard assumption (for instance, see Amin (1993); Bakshi et al. (1997)) that $\log(\psi) \sim N(\mu_\psi, \sigma_\psi^2)$. The parameter σ_{LT}^2 denotes the long term mean of the variance process, and ω is the volatility of volatility. If $\omega = 0$, the volatility is constant at the level σ_{LT}^2, and Equation (4) simplifies to the Merton (1976) jump-diffusion model. When referring to the Merton model, we will denote the (time-constant) volatility σ_M. If we simplify (4) further and set the jump intensity $\gamma = 0$, we recover the standard Black-Scholes model with no jumps. In the following, we will denote the Black-Scholes constant volatility with the symbol σ_{BS}. If instead $\gamma = 0$ but $\omega \neq 0$, Equation (4) is the Heston (1993) model, arguably the most popular process among the family of stochastic volatility models.

The main challenge in this exercise is to compute the continuation value at the ex-dividend date. As anticipated in Section 2, no closed form formula is available, and we need to rely on a fast and accurate numerical method. To this end, we develop a new pricing methodology, the Recursive Projections methodology, which consists in computing at specific moments in time the value of the contract at a fixed (i.e. time-homogenous) grid of values of the relevant state variables, that is the logarithm of the stock price and the stochastic variance. The Recursive Projections are related to quadrature models in that they allow to reconstruct la value of the contract only at the specific points in time that are relevant to our study, that is, maturity and ex-dividend dates. This characteristic is in contrast to finite different methods and trees, that need to compute the value function at close consecutive points in time, and allows us to avoid intermediate steps and to obtain fast pricing algorithms. The main feature that permits reconstructing the value function of the call option at a time t_1 from the value function function at time t_2, with t_2 and t_1 being arbitrarily distant from each other, is that we know the analytical form for the risk neutral transition density of the stock price and variance level from t_1 to t_2 (the so-called Green function) either directly, as in the Black-Scholes-Merton case, or its Fourier transform, as in the Heston-Bates case. Finally, the fact that the value of the call is computed at fixed grid point allows us to exploit some invariance properties of the transition densities that further increase the speed of the method. For the details of the methodology and the convergence properties, see Appendix B[9].

3.2. Estimating the cost of suboptimal non-exercise

The quantification of the suboptimality and the associated costs is computed in two steps: 1) At each day before the payment of the dividend the models of Black-Scholes, Merton and Bates are calibrated separately for each stock 2) By looking at the open interest at the end of the day, we calculate the money left on the table by investors who do not optimally exercise their options.

3.2.1. Calibration results

Table 2 reports the results of the calibration for the three modelling frameworks. We obtain the parameters through the minimization of the implied volatility mean squared error, as in Christoffersen and Jacobs (2004)[10]. The first line of Table 2 displays the average values of the parameters calibrated on our sample of single stocks, whereas the second line reports the average values that Bakshi et al. (1997) obtain for the same parameters from contracts written

[9]Section B of the online supplementary materials describes how the invariance properties of the transition densities simplifies the numerical implementation, while Section C makes an extensive comparison of the Recursive Projections with existing numerical methods.

[10]Section D of the online supplementary materials gives a detailed description of the data and the calibration procedure.

on the $S\&P500$ index. The parameters that rule the level of the volatility smile, namely, the Black-Scholes volatility σ_{BS}, the long term volatility σ_{LT}, and the spot volatility σ_0, are much higher in our single stock calibration than in the index calibration, which reflects the well known fact that an index is less volatile than its components. Indeed, in our sample, the average Black-Scholes volatility is 29%, σ_0, the instantaneous volatility the day of the calibration, is 28%, and the average long-term implied volatility is 32%, whereas for the index options, the same parameters take the values of 18.15%, 20%, and again 20%. The jump parameters in the Bates model show that jumps are on average less frequent in the single stock case than in the index case ($\gamma_{Stocks} = 0.5$ against $\gamma_{SP500} = 0.61$), but the amplitude and variability are higher for single stocks ($\mu_{\psi,Stocks} = -0.12$ and $\sigma_{\psi,Stocks} = 0.18$ respectively) than for the index ($\mu_{\psi,SP500} = -0.09$ and $\sigma_{\psi,SP500} = 0.14$). The remaining two parameters of the stochastic volatility component of the Bates model, the correlation parameter ρ and the volatility of volatility ω, have a specific impact on the shape of the implied volatility smile (Hagan, Kumar, Lesniewski, and Woodward (2002); West (2005)). A negative ρ implies a negatively sloped smile. The correlation parameter is in absolute value lower in the single stock case ($\rho_{SP500} = -0.52$ against $\rho_{Stocks} = -0.35$), meaning that the implied volatility smile for the index is more negatively sloped than for individual stocks. This finding is consistent with the findings of Bakshi, Kapadia, and Madan (2003) who describe the same relationship between the slopes of the index and of the individual stocks implied volatility smiles. Bollen and Whaley (2004) also find the same pattern, and explain it by relating the slope of the index smile to the buying pressure for index puts, with the demand for call options driving the shape of the smile of single stocks. The volatility of volatility ω determines the convexity of the implied volatility smile. The difference in the values taken by ω is striking. We find ω to be 75% for stock options, whereas Bakshi et al. (1997) find a much smaller value of 40% for short-term index options. This difference is due to the higher convexity of the implied volatility smiles of stock options versus that of index options, another feature also documented in Bollen and Whaley (2004). The parameters ρ and ω are related to the smile shape through the higher moments of the distributions of the returns of the underlying. A more negative ρ generates a more negatively skewed distribution of index returns with respect to stock returns, whereas a higher ω in the single stock returns leads to a higher kurtosis than in the index return distribution[11].

3.2.2. Suboptimality results

After having calibrated the models, we are able to compute the price $C(S_{t_h-1} - d, K, T)$ on the day previous to the dividend payment date t_h, by using the recursive projections. The price $C(S_{t_h-1} - d, K, T)$ is the continuation value of the option at date t_h, when the dividend d is distributed. By comparing it with the intrinsic value $S_{t_h-1} - K$, we can assess which options should be exercised on $t_h - 1$. If an option should be exercised (i.e., $C(S_{t_h-1} - d, K, T) \leq S_{t_h-1} - K$), then a positive open interest at the end of the day before ex-dividend ($OI_{t_h-1} > 0$) measures the failure of investors to exercise the option. In this case, we calculate the suboptimal non-exercise percentage as the following ratio:

$$NE_{\%} = \frac{OI_{t_h-1}}{OI_{t_h-2}}, \tag{5}$$

i.e., the number of contracts outstanding at the end of the day $t_h - 1$ to the total number of contracts outstanding at the end of day $t_h - 2$.[12] The total amount of money that is left on the table due to suboptimal non-exercise is given by the following formula:

$$TL = 100 \times OI_{t_h-1} \times [(S_{t_h-1} - K) - C(S_{t_h-1} - d, K, T)]. \tag{6}$$

The continuation value $C(S_{t_h-1} - d, K, T)$ depends on the model used for pricing; hence, the total loss due to suboptimal non-exercise (TL) is itself model-specific.

Table 3 presents the results on the suboptimal non-exercise behavior of investors. Table 3 clearly shows that the optimal early exercise decision depends on the model used for the stock price. Under the Black-Scholes model, approximately 9.5% of the outstanding contracts should

[11] In Section D of the online supplementary materials, we provide more detailed results, including a breakdown of the calibration by stock, and we show that the values of the calibrated parameters are homogeneous across stocks.

[12] The quantity defined in Equation (5) is actually an approximation of the actual non-exercise ratio, because it neglects a possible issue of new contracts on date $t_h - 1$. This event is unlikely; indeed Pool et al. (2008) test the approximation on a subsample of contracts for which they have the real exercise data. They conclude that the approximation (5) is a precise description of the actual exercise behavior of option investors.

be exercised, and the percentage decreases (approximately 7.5%) under the alternative models. This result is consistent with the numerical findings of Section 2, where we show that, in the case of discrete dividends, the early exercise boundary under the Black-Scholes model is lower compared to that implied by the Merton and Heston models. As a general rule, contracts that should be exercised under the Merton or Bates models should also be exercised under the Black-Scholes model. In our sample, we find some exceptions to this rule because, in Section 2, we choose the model parameters such that the total variance of the returns over the life of the option is the same in all models, whereas in real data, this condition may not hold. To give some examples, 4,680 options should be exercised under Black-Scholes but not under the Bates model, whereas the reverse is true only for 249 contracts. Similarly, we find that 2,872 options should be exercised under the Black-Scholes model but not under the Merton model, whereas the opposite occurs with only 53 options. The first important lesson we learn is that, by allowing for more sophisticated models than the Black-Scholes model, the number of contracts that should be optimally exercised decreases by almost 25%. The suboptimality figures are model-dependent and one may argue that they depend on the implemented calibration procedure. The comparison between our calibration results and the results of Bakshi et al. (1997) are reassuring in terms of the reliability of our calibration method. To justify the suboptimal behavior found in our sample entirely, we should obtain unreasonably high values for the jump and intensity parameters.

A second piece of evidence that stands out from Table 3 is that the percentage of investors who leave the options suboptimally non-exercised is higher under the Black-Scholes model than under the other models, 39% versus approximately 30%. We compute these percentages in accordance with Definition (5). If we restrict our attention to the 1,965 contracts in our sample that should be exercised under the Black-Scholes model but not under the Merton model or the Bates model, we find a striking 81% of no-exercise. These results may suggest that investors do not limit themselves to a Black-Scholes world when evaluating their options but rely on more sophisticated models that include jumps or stochastic volatility. Even if this evidence is a considerable step towards understanding the investor decision-making process, it does not fully solve the puzzle. Indeed, even in the Merton and Bates models, we still find a high percentage of suboptimal non-exercises, which leads to a global loss of approximately $130-140$ million dollars, down approximately 30% from the loss of 206 million dollars in the Black-Scholes model. The possible interpretation of this residual amount in terms of hidden fees, or as opportunity cost of closely monitoring the options, leads us to the discussion of Section 3.3.

As a further reliability check on our calibration procedure, we compare our results regarding the exercise decision with the ones obtained by Pool et al. (2008), limitedly to the Black-Scholes case. In that work, the authors apply the early exercise decision rule to all options series by using the Black-Scholes model with historical volatility and find that 53.1% of investors leave their options unexercised when instead they should have been exercised. Their data span over ten years (from 1996 to 2006) and to compare our results with theirs, we divide our sample into two subsamples, the first spanning the years 1996-2006 and the second spanning the years 2006-2012. For comparability, we calculate the average percentage of suboptimal non-exercise in the two subsamples and find that the percentage of suboptimal non-exercise under the Black-Scholes model is approximately 47% in the first subsample and 37% in the second. The decrease in the non-exercise behavior with time intimates that investors become more attentive in monitoring their investments. There is a small difference between our results (47%) and the 53.1% found in Pool et al. (2008). The explanation is most likely our focus on the constituents of the Dow Jones Industrial Average, whereas Pool et al. (2008) consider all option series. It is likely that, for large-cap companies, stock and option prices are monitored more closely than they are on average.

Throughout our empirical investigation, we choose a model-based approach to calculate the continuation value of the option $C(S_{t_h-1} - d, K, T)$. We could have also used a market-based approach where the continuation value is the market price of the option. The market-based approach checks whether the quantity $C_{MKT}(S_{t_h-1}, K, T) - (S_{t_h-1} - K)_+$ equals 0, where $C_{MKT}(S_{t_h-1}, K, T)$ is the observed market price at $t = t_h - 1$. As discussed in Pool et al. (2008) and in Barraclough and Whaley (2012), the market-based approach has shortcomings. The most important is that it does not make it possible to calculate the total loss due to suboptimal non-exercise, which we do in Equation (6). In addition, the bid-ask spread and the discreteness of the prices make it difficult to decide which C_{MKT} should be used. For all of these reasons, we follow Pool et al. (2008) and Barraclough and Whaley (2012) and use a model-based approach with Equation (5) to account for the actual exercise behavior of options investors. Barraclough and Whaley (2012) only use the market-based approach as a useful model-free test to confirm the presence of suboptimal non-exercise behavior. They find that the market-based approach gives a magnitude of suboptimality that is comparable to that implied by the model-based approach. This last piece of evidence is an additional argument against the possible objection that an incorrect calibration of the model parameters may be the source of the suboptimal exercise figures.

3.3. The role of fees

According to the recent literature on option prices (Jensen and Pedersen (2016), Christoffersen, Goyenko, Jacobs, and Karoui (2017b)), trading costs and financial frictions in general strongly affect both the option prices and the early exercise decision of American options. In this section, we investigate whether the suboptimal non-exercise behavior of investors is due to the trading costs that investors face when exercising their options.

Following Pool et al. (2008), we model the costs of exiting a long call position as a per share lump sum $\mathscr{F}$ that the investor must pay at the moment he decides to exercise. The specific value of $\mathscr{F}$ depends on how the exit is accomplished according to the different possible objectives of the investor. The most expensive value of the fee $\mathscr{F}$ is attained when the investor wants to exercise the option and reenter into the same call position. Pool et al. (2008) estimate an average value for the rollover costs $\mathscr{F}$ by using the commissions of the high-cost brokers, and they obtain a very conservative amount of $\mathscr{F} = 0.4446$ dollar per share. A detailed description of the components of the fee $\mathscr{F}$ can be found in Pool et al. (2008).

To understand the role of the fees in the early exercise decision, we perform two different empirical exercises. As a first check, we re-perform the exercise of Section 3.2.2 and compute the loss due to a suboptimal non-exercise decision, but this time using $C(S_{t_h-1} - d, K + \mathscr{F}, T)$ as the continuation value, and $(S_{t_h-1} - K - \mathscr{F})$ as the intrinsic value. The fee value is $\mathscr{F} = 0.4446$.

The fee $\mathscr{F}$ enters both in the exercise proceeds and in the continuation value. Indeed, at the moment of the exercise decision, the investor should decide whether to exercise and hence pay the exercise fee immediately or not exercise and postpone the payment of the exercise fee to a future date. Accordingly, the calculation of the total amount of money that is lost due to suboptimal non-exercise is given by the following formula:

$$TL_{\mathscr{F}} = 100 \times OI_{t_h-1} \times [(S_{t_h-1} - K - \mathscr{F}) - C(S_{t_h-1} - d, K + \mathscr{F}, T)]. \tag{7}$$

The second column in Table 3 shows the summary results including the fee. They are not very different from those obtained without considering the fee (first column of Table 3). We can conclude that the inclusion of trading costs does not change the big picture on the suboptimal non-exercise of investors, as outlined in the previous paragraph.

As a second empirical exercise, and in order to detect possible additional costs that are not taken into consideration in the fee $\mathscr{F}$, we calculate the value of the fee that would justify the non-exercise decision of investors. To do so, for each option for which $C(S_{t_h-1} - d, K, T) < (S_{t_h-1} - K)$, but that is not optimally exercised by some of the investors, we compute the value of the implied fee $\mathscr{F}$ that would justify the non-exercise decision. It amounts to numerically finding the zero of the following function:

$$f(\mathscr{F}) = C(S_{t_h-1} - d, K + \mathscr{F}, T) - (S_{t_h-1} - K - \mathscr{F}). \tag{8}$$

The results are reported in Table 4. The average implied fee is between 7 and 8 dollars per share, an incredibly high amount compared to the already conservative fee of 0.4446 dollar per share estimated by Pool et al. (2008). No realistic hidden fees can sum up to 7 dollars per share, and the trading costs of exiting a long call option position cannot fully justify the suboptimal non-exercise behavior of investors. We can interpret the difference between the implied fee of 8 dollars and the conservative fee of 0.4446 as an implied opportunity cost for the holder of the option to monitor the optimal exercise of the American option. Investors may not respond immediately to favorable stock price movements and may take some time before reacting and optimally exercising their option, which would be in line with the behavior on the early prepayment of mortgages documented by Stanton (1995).

3.4. Discussion of the findings on the early exercise decision

In Section 3.2.2, we have documented a decrease in the suboptimal exercise behavior when jumps and stochastic volatility are added to the stock process. The decrease is explained by an increase of the value of the early exercise boundary under the alternative modelling environments. We now extend the discussion on the early exercise boundary of Section 2, and provide theoretical insights that justify our empirical findings.

In the discussion of Section 2, we found that the early exercise boundary reacted differently when moving from Black-Scholes to Merton and Heston, depending on the modelling assumption of the dividend payment. The findings in the case of a continuous dividend do not provide an unambiguous prediction on the effect of the underlying process on the exercise boundary, and are in line with the ones of Adolfsson et al. (2013) for the stochastic volatility case, and of Amin (1993) for the jump-diffusion case. The former work documents a lower exercise boundary in the Heston case (for all maturities) with respect to the Black-Scholes case, while the latter shows evidence of a lower boundary in the Merton case than under Black-Scholes case for longer maturities, while the opposite is true for shorter maturities. Adolfsson et al. (2013) and Amin (1993) back up these findings using simulations, but do not explain why it is so. In Appendix C, we provide theoretical insights on these results.

The results of Section 2 that are more relevant to our empirical application are the ones in the case of a discrete dividend, and are, to our knowledge, entirely new. They consistently predict an increase of the early exercise boundary. In the following, we explain why it is the case. We first consider the stochastic volatility case. Assume there is only one discrete dividend to be paid. The continuation value of the call option immediately after the ex-dividend date is that of a European call with the remaining time to maturity. When the correlation $\rho \leq 0$, the price of European options for a deep in-the-money call, where early exercise could be optimal, is higher in the Heston case than in the Black-Scholes case (see Heston (1993), Hull and White (1987)). For instance, in the left graph of Panel B of Figure 1, for a time to maturity of 0.25, this would be the case in the range of stock prices of approximately 150. Even by taking into account the dividend drop in computing the continuation value, the ex-dividend stock price should remain in the region where the Heston price is higher. We can repeat the same argument for a number of discrete dividends sufficiently small (typically of the order of a couple percent) to prevent the stock price from falling in the price range where the call has more value under the Black-Scholes model. A higher continuation value under Heston model in turn implies a higher early exercise boundary.

We move on to considering the behavior of the early exercise boundary under the Merton model, that is, model (4) with the volatility of volatility $\omega = 0$ and σ_t set constant to σ_M. To interpret the graphs in Figure 2, we have to make an important distinction. For short maturity options, the jump component in the first line of Equation (4) dominates the diffusion component. As explained in Amin (1993) and Merton (1976), the result is higher prices for short maturities in-the-money call options under the Merton model than under the Black-Scholes model. We call this effect the jump effect. For longer maturities, the jump effect no longer dominates the diffusion component but instead creates an interplay that makes the jump-diffusion process observationally similar to a stochastic volatility process. For a discrete dividend, both the jump effect and the stochastic volatility effect, as previously discussed in the Heston case, predict a higher boundary in the Merton case than in the Black-Scholes case, which holds true for all maturities. The full Bates specification (4) must imply a higher exercise boundary than the Black-Scholes model in the discrete divided case, as it displays both jumps and stochastic volatility, and both features push the boundary upwards.

4. Concluding remarks

We investigate the exercise behaviour of investors in a large database of 101,295 series of short-term American call options (9.5 million prices). Pool et al. (2008) find that more than 40% of the investors fail to optimally exercise their contracts. We extend their analysis by including stochastic volatility and jumps to the process of the underlying stock. In order to deal with the large option database and the repeated calculations required for the calibration and pricing, we develop an option pricing technique which is at the same time fast, precise and which can handle both multidimensional dynamics and cash dividend distributions. By applying our technique to the dataset, we can explain up to 25% of the gain forgone due to suboptimal exercise decisions, as computed in Pool et al. (2008). This result confirms the insights we obtain from the theoretical developments of the paper. Indeed, we show that the exercise boundary is higher under the Merton and Heston models than under the Black-Scholes model if the dividend is discrete. This result underlines the importance of the correct modelling of the dividend distribution. We show that by modelling the dividend as a continuous yield instead of a discrete cash flow, the exercise frontier in the Bates model could have been lower instead of higher, and the suboptimal exercise behavior of Pool et al. (2008) would have been reinforced instead of mitigated. We further try to check whether we can explain the remaining part of the suboptimal behavior in terms of transaction costs (Jensen and Pedersen (2016)). We show that hidden transaction costs would need to be unrealistically large to explain the entire amount foregone by investors. This observation leads us to interpret the implied transaction fee as a monitoring cost.

As a final remark, the new pricing methodology (Recursive Projections) we developed is a powerful tool which allows new empirical research investigations. On the one side, American option pricing is only a particular case of stochastic optimal control problems. We can think of applying the recursive projection method to other problems, such as the optimal portfolio allocation involving complex and path-dependent financial assets or incentives contracts (Hodder and Jackwerth (2007)). Currently, these types of complex problems are solved by using Monte Carlo simulations (Detemple, Garcia, and Rindisbacher (2003)) and our method could offer a more efficient computational alternative. On the other side, our method facilitates the analysis based on individual stocks paying discrete dividends and sophisticated pricing models, like for instance the recent works of Kelly, Lustig, and Van Nieuwerburgh (2016b) and Christoffersen, Fournier, and Jacobs (2017a). We can also think, for example, of extending the recent studies of tail risk (Andersen, Fusari, and Todorov (2015, 2017)) and political uncertainty (Kelly, Pástor, and Veronesi (2016a)) to single name stock options. We can also use our method to efficiently compute the Greeks of the options, which are the risk metrics mostly used by option traders.

All these analyses are outside the scope of this paper and are left for future research.

References

Adolfsson, T., Chiarella, C., Ziogas, A., Ziveyi, J., 2013. Representation and numerical approximation of American option prices under Heston stochastic volatility dynamics. Research Paper 327, Quantitative finance research centre, University of Technology Sidney.

Amin, K. I., 1993. Jump diffusion option valuation in discrete time. The Journal of Finance 48, 1833–1863.

Andersen, T. G., Fusari, N., Todorov, V., 2015. The risk premia embedded in index options. Journal of Financial Economics 117, 558–584.

Andersen, T. G., Fusari, N., Todorov, V., 2017. Short-term market risks implied by weekly options. The Journal of Finance 72, 1335–1386.

Andricopoulos, A. D., Widdicks, M., Newton, D. P., Duck, P. W., 2007. Extending quadrature methods to value multi-asset and complex path dependent options. Journal of Financial Economics 83, 471–499.

Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. The Journal of Finance 52, 2003–2049.

Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options. Review of Financial Studies 16, 101–143.

Barraclough, K., Whaley, R. E., 2012. Early exercise of put options on stocks. The Journal of Finance 67, 1423–1456.

Bates, D. S., 1996. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. Review of Financial Studies 9, 69–107.

Bollen, N. P. B., Whaley, R. E., 2004. Does net buying pressure affect the shape of implied volatility functions? The Journal of Finance 59, 711–753.

Broadie, M., Chernov, M., Johannes, M., 2007. Model specification and risk premia: Evidence from futures options. The Journal of Finance 62, 1453–1490.

Broadie, M., Chernov, M., Johannes, M., 2009. Understanding index option returns. Review of Financial Studies 22, 4493–4529.

Chen, D., Härkönen, H. J., Newton, D. P., 2014. Advancing the universality of quadrature methods to any underlying process for option pricing. Journal of Financial Economics 114, 600–612.

Christoffersen, P., Fournier, M., Jacobs, K., 2017a. The factor structure in equity options. The Review of Financial Studies 31, 595–637.

Christoffersen, P., Goyenko, R., Jacobs, K., Karoui, M., 2017b. Illiquidity premia in the equity options market. The Review of Financial Studies 31, 811–851.

Christoffersen, P., Jacobs, K., 2004. The importance of the loss function in option valuation. Journal of Financial Economics 72, 291–318.

Cosma, A., Galluccio, S., Pederzoli, P., Scaillet, O., 2016. Early exercise decision in American options with dividends, stochastic volatility and jumps. Swiss Finance Institute Research Paper No. 16-73. Available at SSRN: https://ssrn.com/abstract=2883202 or http://dx.doi.org/10.2139/ssrn.2883202.

Detemple, J. B., Garcia, R., Rindisbacher, M., 2003. A Monte Carlo method for optimal portfolios. The Journal of Finance 58, 401–446.

Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. The Journal of Finance 58, 1269–1300.

Fang, F., Oosterlee, C. W., 2011. A fourier-based valuation method for bermudan and barrier options under heston's model. SIAM Journal on Financial Mathematics 2, 439–463.

Griebsch, S. A., 2013. The evaluation of European compound option prices under stochastic volatility using Fourier transform techniques. Review of Derivatives Research 16, 135–165.

Hagan, P. S., Kumar, D., Lesniewski, A. S., Woodward, D. E., 2002. Managing smile risk. WILMOTT Magazine pp. 84–108.

Haug, E. G., Haug, J., Lewis, A., 2003. Back to basics: a new approach to the discrete dividend problem. Wilmott magazine 9, 37–47.

Heston, S. L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.

Hodder, J. E., Jackwerth, J. C., 2007. Incentive contracts and hedge fund management. Journal of Financial and Quantitative Analysis 42, 811–826.

Hull, J., 2011. Options, Futures and Other Derivatives. Pearson Education, 8th ed.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. The Journal of Finance 42, 281–300.

Jamshidian, F., 1992. An analysis of American options. Review of Futures Markets 11, 72–80.

Jensen, M. V., Pedersen, L. H., 2016. Early option exercise: Never say never. Journal of Financial Economics 121, 278–299.

Kelly, B., Pástor, L., Veronesi, P., 2016a. The price of political uncertainty: Theory and evidence from the option market. The Journal of Finance 71, 2417–2480.

Kelly, B. T., Lustig, H., Van Nieuwerburgh, S., 2016b. Too-systemic-to-fail: What option markets imply about sector-wide government guarantees. American Economic Review 106, 1278–1319.

Kim, I., 1990. The analytic valuation of American options. Review of Financial Studies 3, 547–572.

Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125–144.

Pool, V. K., Stoll, H. R., Whaley, R. E., 2008. Failure to exercise call options: An anomaly and a trading game. Journal of Financial Markets 11, 1–35.

Stanton, R., 1995. Rational prepayment and the valuation of mortgage-backed securities. Review of Financial Studies 8, 677–708.

West, G., 2005. Calibration of the SABR model in illiquid markets. Applied Mathematical Finance 12, 371–385.

			Number of option quotes					
			maturity<60 days			maturity>60 days		
	Ticker	Stock	OTM	ATM	ITM	OTM	ATM	ITM
1	AA	ALCOA INC	42827	6379	41206	28788	9621	29371
2	AXP	AMERICAN EXPRESS CO	43360	10386	59656	27744	14941	50002
3	BAC	BANK OF AMERICA CO	53785	7767	57794	48365	15446	58196
4	BA	BOEING CO	41616	7385	52422	36022	13312	50357
5	CAT	CATERPILLAR INC DEL	40592	8651	55672	34438	15377	54361
6	CHV	CHEVRON CORPORATION	33444	5922	50092	28088	11238	51676
7	CSCO	CISCO SYS INC	58276	8194	63177	44295	14142	55598
8	KO	COCA COLA CO	34710	5055	44407	29054	9249	44421
9	DIS	DISNEY WALT CO	40138	7052	47571	31350	10730	40773
10	XOM	EXXON MOBIL CORP	39509	6427	55468	27944	9425	48375
11	GE	GENERAL ELECTRIC CO	49637	6734	58779	42696	13072	58291
12	HWP	HEWLETT PACKARD CO	54913	9182	53498	47707	17077	51029
13	HD	HOME DEPOT INC	42968	7052	56732	36047	13413	53921
14	INTC	INTEL CORP	55825	8666	61243	43796	14327	52367
15	IBM	INTER. BUS. MACHS	70249	9060	85569	45355	13330	62979
16	JNJ	JOHNSON & JOHNSON	34175	4551	47524	22974	7074	43322
17	MCD	MCDONALDS CORP	33730	5674	47969	27630	10145	48000
18	MRK	MERCK & CO INC	41559	7612	53181	31901	11486	47534
19	MSFT	MICROSOFT CORP	68396	8725	78393	50253	14163	65728
20	MMM	3M CO	36835	6737	46521	24402	10205	38750
21	JPM	MORGAN J P & CO INC	6991	3180	9940	4428	5874	10568
22	PFE	PFIZER INC	51100	6170	54751	47414	12109	52358
23	PG	PROCTER & GAMBLE CO	36971	5904	52782	25422	8974	46956
24	T	AT&T INC	42547	5542	52123	32851	9496	46524
25	TRV	TRAVELERS COMPANIES INC	21404	3803	27783	15069	5811	23646
26	UTX	UNITED TECHNOLOGIES CORP	34765	6366	48200	28973	11712	45580
27	UNH	UNITEDHEALTH GROUP INC	39924	8864	58885	32757	16724	56451
28	VZ	VERIZON COMMUNICATIONS INC	39642	6461	56527	34306	11100	54563
29	WMT	WAL-MART STORES INC	37172	5668	51904	31679	10429	52890
30	DD	DU PONT E I NEMOURS & CO	35262	7404	46302	24077	11097	41173

Table 1: Number of observations for in-the-money (ITM), at-the-money (ATM) and out-of-the-money (OTM) call option quotes for the stocks which are the constituents of the Dow Jones Industrial Average Index (DJIA). The data are further broken down by maturity. According to the classification of Bollen and Whaley (2004), a call option is considered OTM if its delta is less then 0.375, ATM if its delta ranges between 0.375 and 0.625 and ITM if its delta is above 0.625.

Underlying	BS	MRT				BTS							
	σ_{BS}	γ	σ_M	σ_ψ	μ_ψ	γ	σ_ψ	μ_ψ	ω	σ_{LT}	β	ρ	σ_0
All stocks	0.29	1.33	0.22	0.16	-0.12	0.50	0.18	-0.12	0.75	0.32	1.52	-0.35	0.28
SP500*	0.18	NA	NA	NA	NA	0.61	0.14	-0.09	0.4	0.2	3.93	-0.52	0.2

Table 2: Average values of the parameters of the models of Black-Scholes (BS), Merton (MRT) and Bates (BTS) calibrated at each day before the ex-dividend date on the options written on the dividend-paying stocks belonging to the Dow Jones Industrial Average Index (DJIA). In total we computed 1701 calibrations and the reported values are the averages across these calibrations.
The in-sample sum of squared error is on average equal to 0.26 for the Black-Scholes model, 0.20 for the Merton model, and 0.16 for the Bates model with stochastic volatility.

*Calibrated parameters of the SP500 dynamics are from Bakshi et al. (1997).

	Model	Without fee	With fee
Contracts outstanding		406 414 980	
Total market value		99 392 927 000	
Contracts that should be exercised	BS	38 527 586 (9.48%)	31 551 786 (7.76%)
	MRT	30 633 542 (7.54%)	25 340 009 (6.23%)
	BTS	30 486 666 (7.5%)	25 050 616 (6.16%)
Contracts that are left suboptimally non-exercised	BS	15 214 908 (39.49%)	11 077 913 (35.11%)
	MRT	9 404 406 (30.70%)	7 118 002 (28.09%)
	BTS	8 786 524 (28.82%)	6 702 799 (26.76%)
Money available due to exercise opportunity	BS	770 287 766 (0.77%)	647 088 372 (0.65%)
	MRT	555 922 918 (0.56%)	460 980 123 (0.46%)
	BTS	617 173 686 (0.62%)	509 713 158 (0.51%)
Total loss due to suboptimal non-exercise	BS	209 284 628 (27.17%) (0.21%)	172 424 967 (26.65%) (0.17%)
	MRT	133 130 786 (23.95%) (0.13%)	108 951 652 (23.63%) (0.11%)
	BTS	147 480 996 (23.9%) (0.15%)	123 615 413 (24.25%) (0.12%)

Table 3: Summary results of the total loss due to suboptimal non-exercise for the short-term call option series. The numbers are calculated for each series and each day before the ex-dividend date separately, and then pooled together.

The number of contracts outstanding is the total open interest of all contracts two days before the ex-dividend date. The contracts that should be exercised under a specific model, i.e. Black-Scholes (BS), Merton (MRT) and Bates (BTS), are the contracts outstanding for which the continuation value is lower than the exercise proceeds. The number of contracts that are left suboptimally non-exercised is the sum of the open interests one day before ex-dividend of the contracts that should have been exercised. We compute the other quantities in the table in the following way:

Total market value = Contracts outstanding$\times$Market price$\times$100,

Money available $= \max\{0, (S - K - \mathscr{F} - \text{Continuation value}) \times \text{Contracts outstanding} \times 100\}$,

Total loss $= \max\{0, (S - K - \mathscr{F} - \text{Continuation value}) \times \text{Open interest}_{t-1} \times 100\}$,

where $\mathscr{F}$ is the exercise fee. In the first column the results are computed considering $\mathscr{F} = 0$, while in the second column the results are computed considering the conservative fee of 0.44 dollar. The first percentage in parenthesis in the Total loss due to non-exercise is computed with respect to the money available due to exercise opportunities, while the second one is computed with respect to the total market value.

	Average rational implied fee			%Implied fee $>$ 0.4446		
Underlying	BS	MRT	BTS	BS	MRT	BTS
All stocks	7.54	7.27	7.23	94%	93%	93%

Table 4: The table reports the average implied fee per share which would explain the non-exercise behavior of investors in each model: Black-Scholes (BS), Merton (MRT) and Bates (BTS). The average implied fee is calculated for each option that should be exercised but which is not optimally exercised by some of the investors as the value of the trading costs $\mathscr{F}$ which makes the continuation value of the option equal to the early exercise proceeds: $C(S-d, K+\mathscr{F}, T) = (S-K-\mathscr{F})$. In the last three columns of the table we report the percentage of options for which the fee that would explain the suboptimal non-exercise behavior is higher than the conservative fee of 0.4446 dollar per share estimated by Pool et al. (2008).

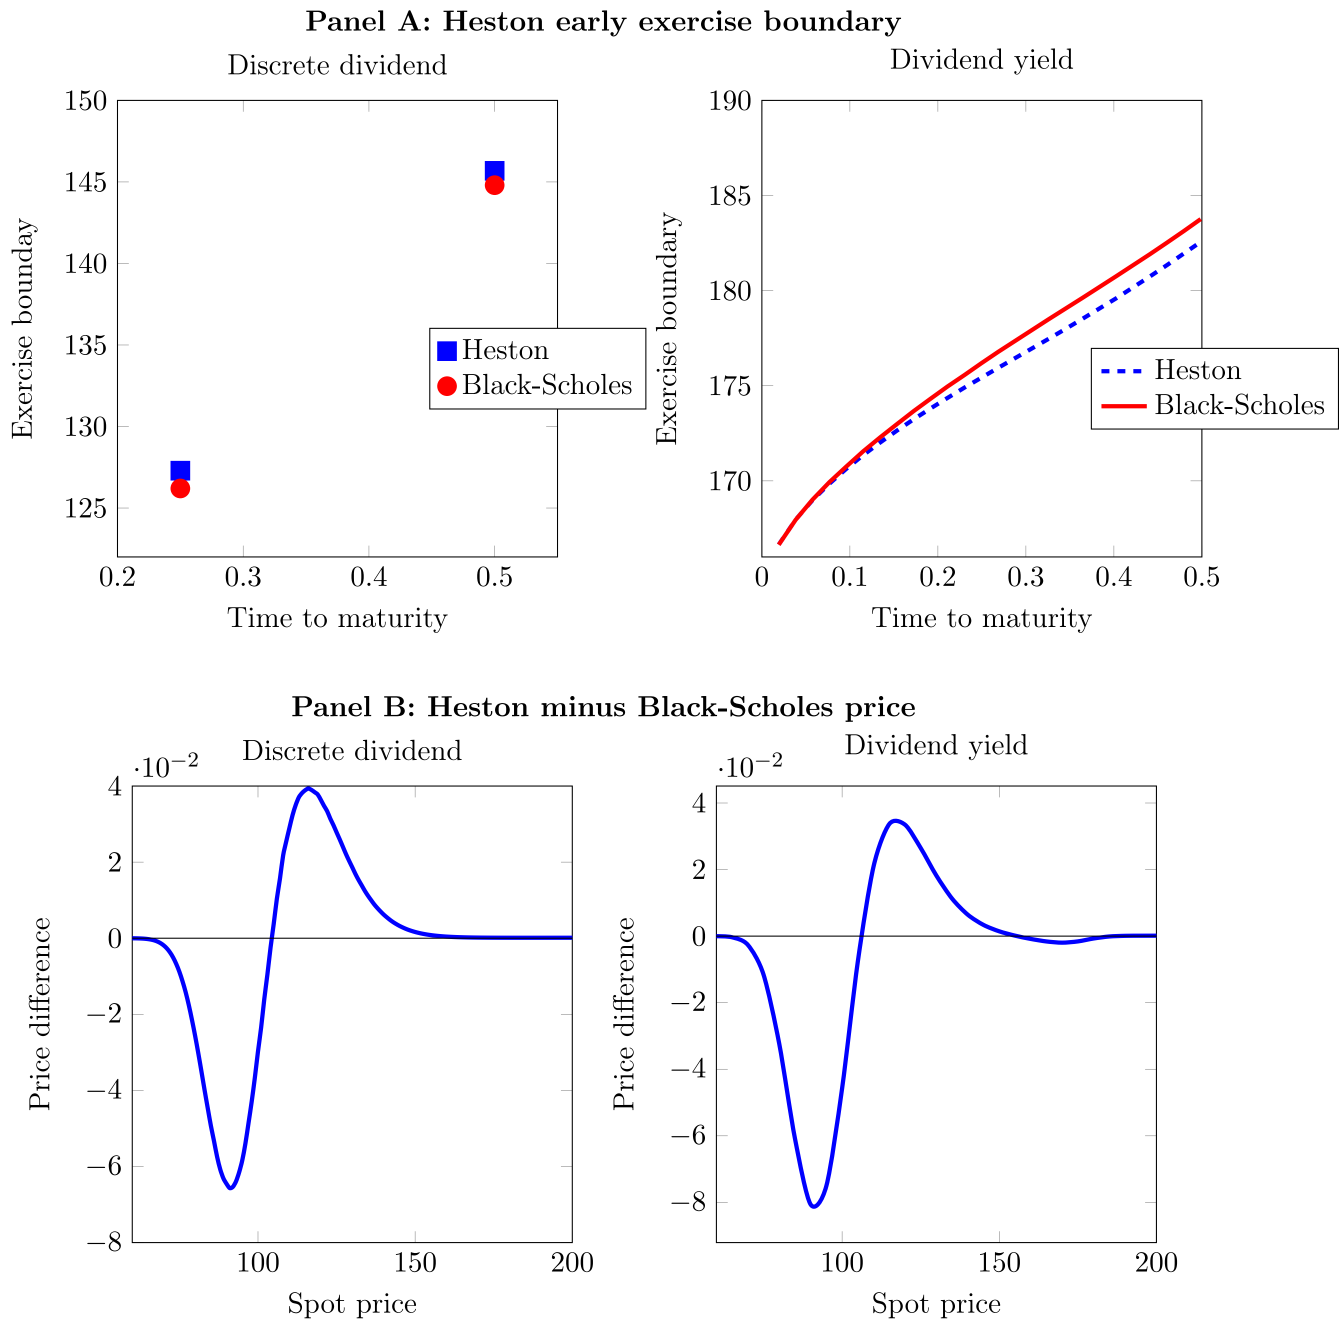

Fig. 1. Panel A. Comparison between the early exercise boundary in the Heston and Black-Scholes models of an American call with maturity 6 months, in the case in which the stock pays a dividend yield $r_d = 0.03$ (right) and in the case in which the stock distributes an equivalent quarterly discrete dividend of $d = 1.38$ (left). The remaining parameters are: $K = 100$, $r = 0.05$, $\sigma_0 = 0.2$, $\omega = 0.1$, $\sigma_{LT} = 0.3$, $\beta = 4$, $\rho = -0.5$. We set the volatility parameter in the Black-Scholes model equal to the volatility of the underlying return over the life of the option in the Heston model. Panel B. Heston minus Black-Scholes price of an American call with $T = 0.25$ for different values of S_0 in the case of discrete dividend (left) and continuous dividend yield (right). The remaining parameters are the same as in Panel A.

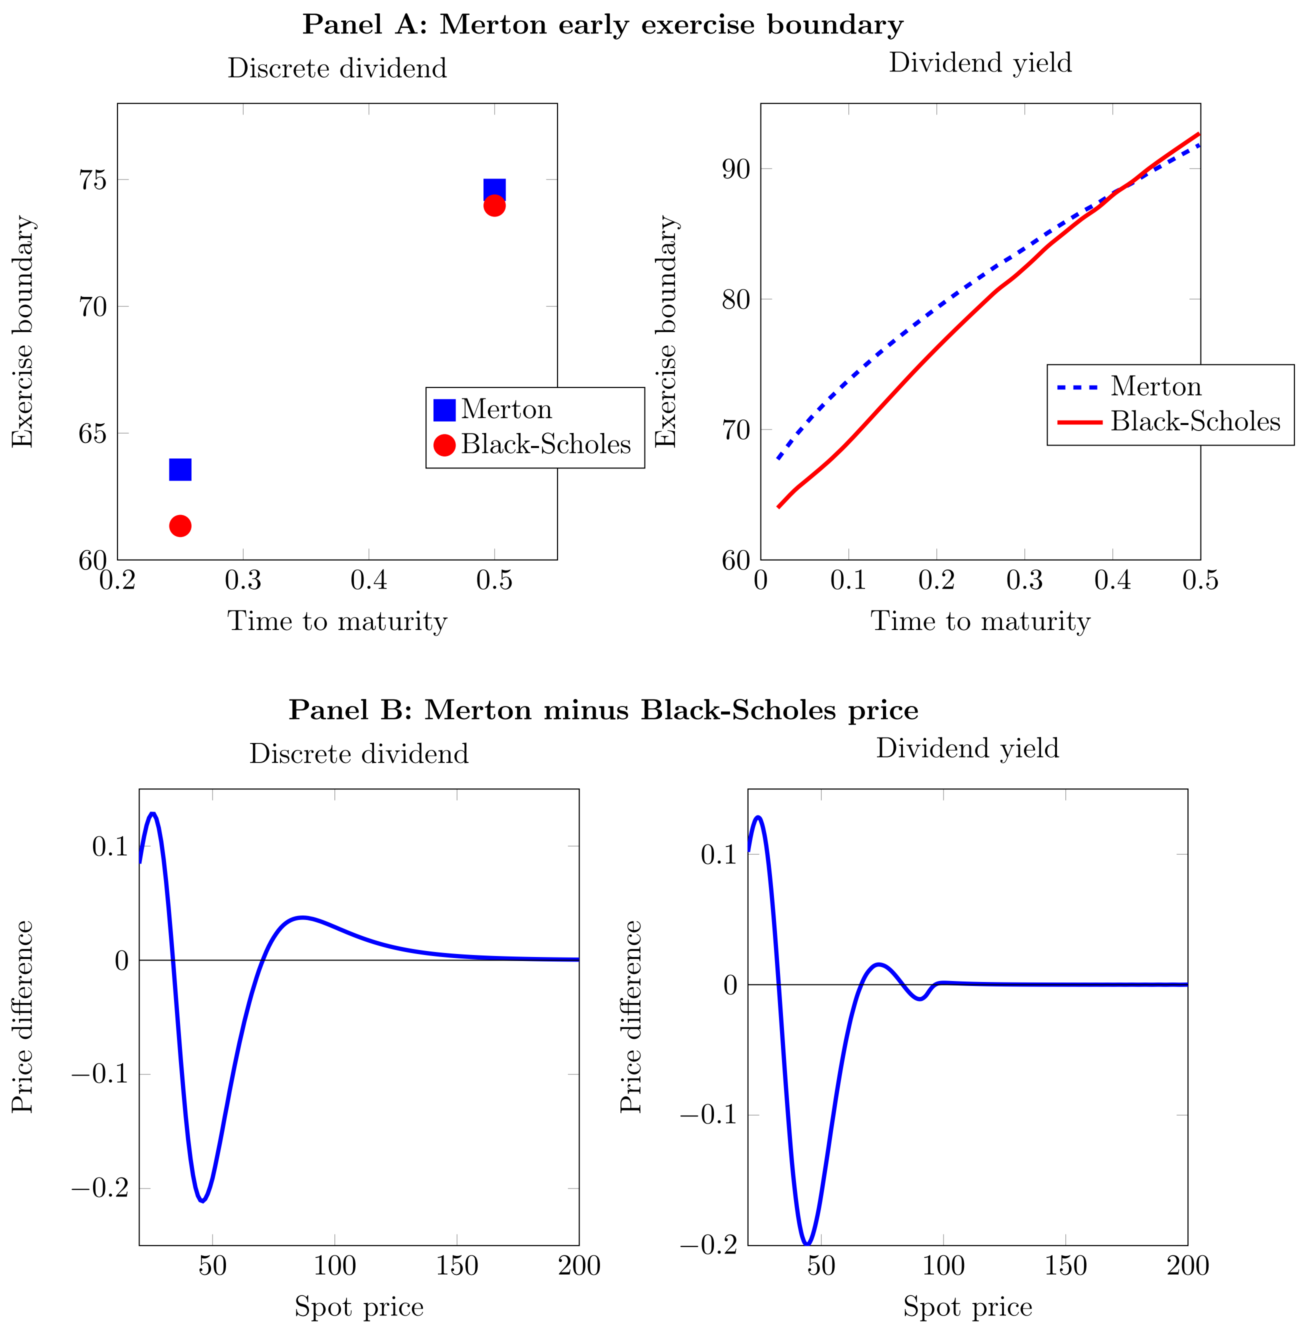

Fig. 2. Panel A. Comparison between the early exercise boundary in the Merton and Black-Scholes models of an American call with maturity 6 months, in the case in which the stock pays a dividend yield $r_d = 0.05$ (right) and in the case in which the stock distributes an equivalent quarterly discrete dividend of $d = 1.125$ (left). The other parameters are the following: $K = 40$, $T = 0.5$, $r = 0.08$, $\gamma = 5$, $\sigma_M^2 = 0.05$, $\sigma_\psi^2 = 0.05$, $\mu_\psi = 0$. We set the volatility parameter in the Black-Scholes model equal to the volatility of the underlying return over the life of the option in the Merton model. Panel B. Merton minus Black-Scholes price of an American call with the same parameters as those used in Panel A but different values of S_0 in the case of discrete dividend (left) and continuous dividend yield (right).

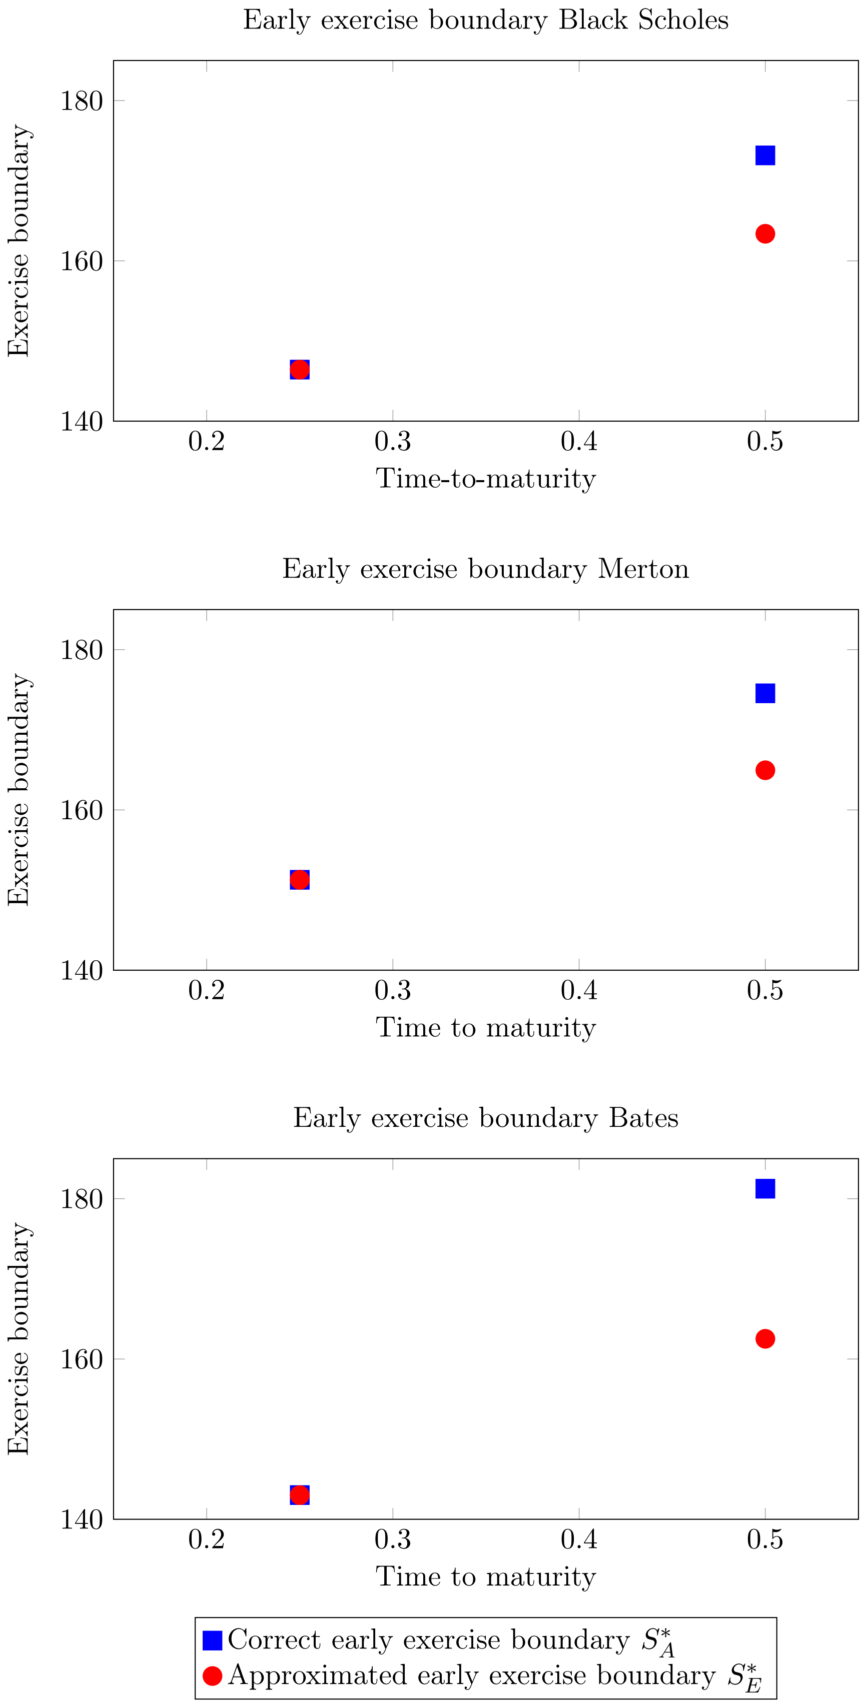

Fig. 3. Comparison between the true early exercise boundary S_A^* and the approximated early exercise boundary S_E^*. S_E^* is calculated by approximating the continuation value of the option with the price of a European option where the starting value of the stock is set equal to S_0 minus the present value of all future dividends.

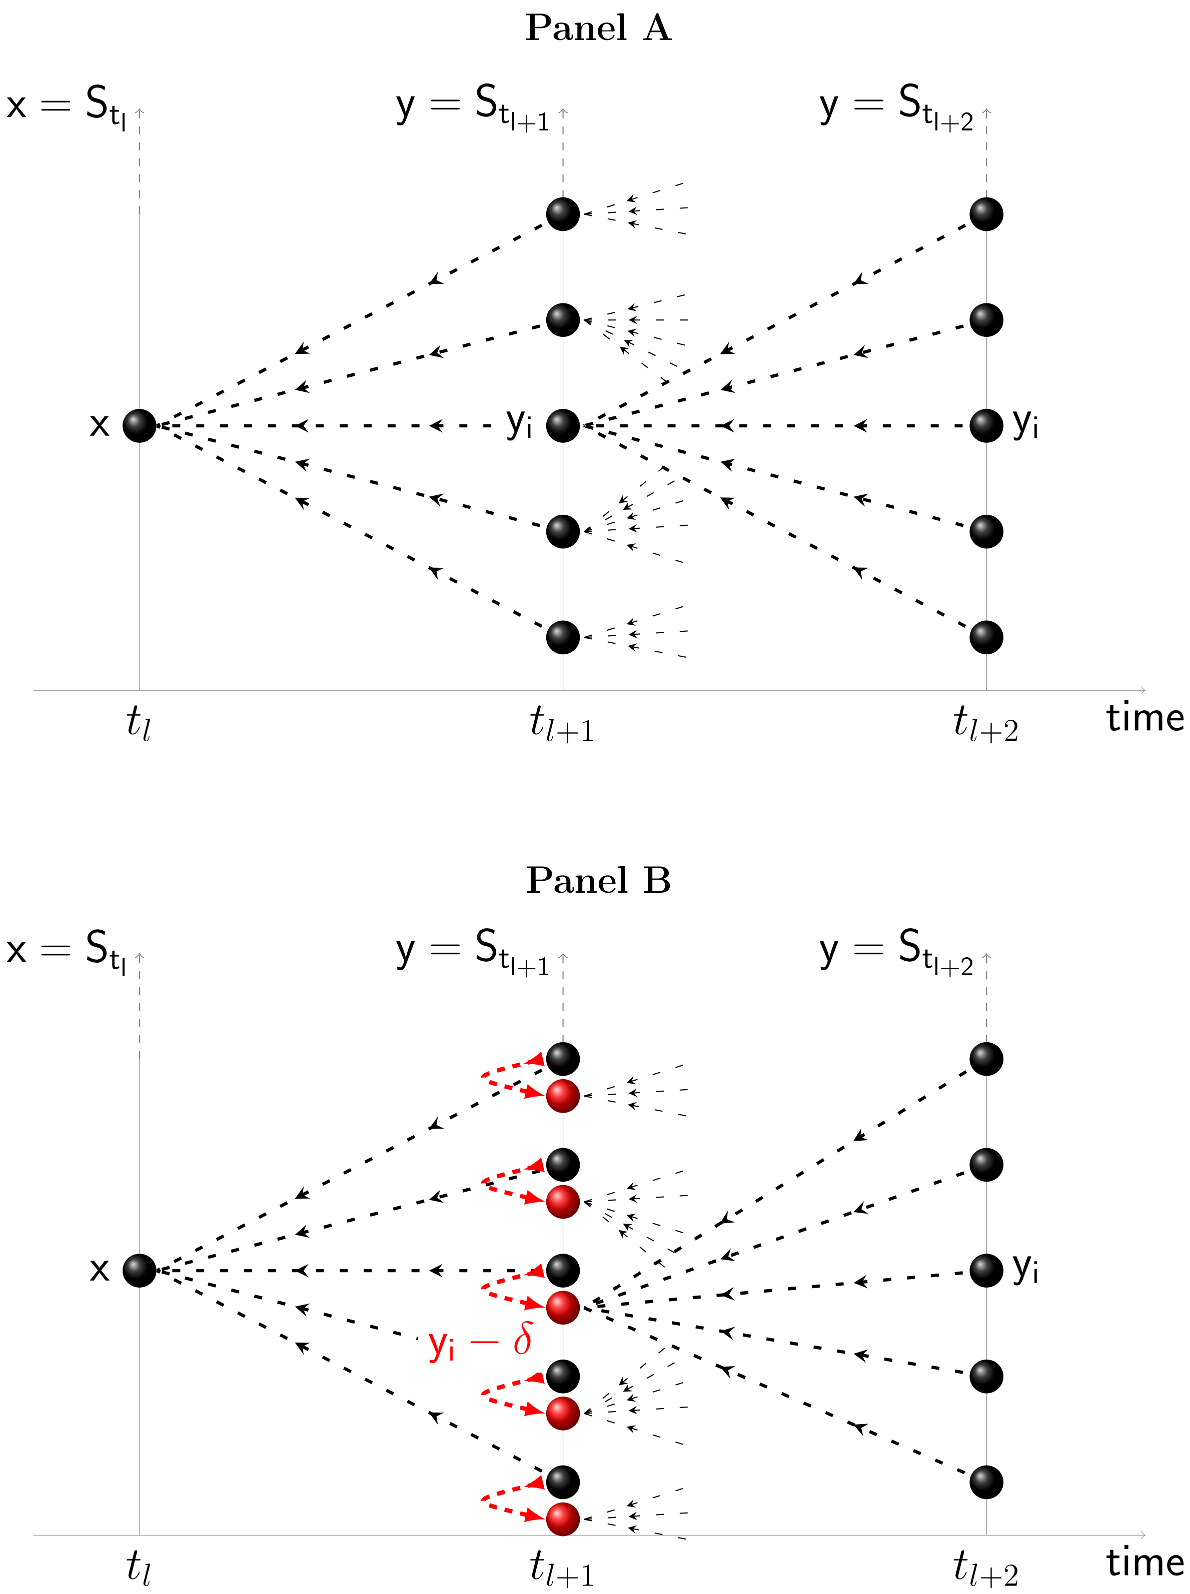

Fig. 4. Recursive scheme without dividends (Panel A) and with discrete dividends (Panel B). In Panel A, at date $t = t_{l+1}$, the intrinsic value $H(y_i, t_{l+1}) = \max\{y_i - K, 0\}$ is compared with the continuation value $C(y_i, T, K)$ computed at the same grid point y_i (black ball). In Panel B, at the ex-dividend date $t_h = t_{l+1}$, the intrinsic value $H(y_i, t_{l+1})$ at the grid point y_i (black ball) is compared with the continuation value $C(y_i - d, T, K)$ at $y_i - d$ (red ball).

Appendix A. Approximating American option prices by European prices

In this section, we provide further evidence on the importance of a correct modelling of the dividend as a discrete cash flow when computing the early exercise boundary.

The setup is the same as in Section 2 of the main text. There are two dividends to be paid over the remaining life of the option, which is 6 months. One is to be paid immediately ($t = 0$), the second in 3 months ($t = 0.25$). We compute the early exercise boundary by i) correctly taking into account that the dividends are discrete, and ii) by using the "escrow dividend" approximation, that is we price the contract at $t = 0$ as if it were a European option, with the continuation value after the first dividend payment not being computed at $S_0 - d$, but at $S_0 - d - d \cdot e^{-r(0.25)}$. We compute the early exercise boundary under the three models (Black-Scholes, Merton and Bates), with the following representative parameters: $\gamma = 5$, $\mu_\psi = 0$, $\sigma_\psi = 0.2$, $\sigma_{MRT} = 0.2$, σ_{BS}, $\sigma_0 = 0.2$, $\omega = 0.1$, $\sigma_{LT} = 0.3$, $\beta = 4$, $\rho = -0.5$, $r = 0.05$. Figure 3 displays the results of our numerical study. We call $S^*_{A,t}$ the early exercise boundary computed under assumption i), and $S^*_{E,t}$ the one under assumption ii). The two boundaries $S^*_{A,t}$ and $S^*_{E,t}$ coincide after the dividend at $t = 0.25$ is paid out, since there are no more intermediate cash flows before maturity, and the continuation value is in both cases the one of a European call with a time-to-maturity of 3 months. At $t = 0$, the two boundaries are different, with $S^*_{E,t} < S^*_{A,t}$. This means that an investor basing his exercise decision on the approximation $S^*_{E,t}$, may suboptimally decide to exercise the option at $t = 0$, if $S^*_{E,0} < S_0 < S^*_{A,0}$. He would then incur a loss given by $C(S_0 - d, 0) - (S_0 - K)$, where $C(S_0 - d, 0)$ is the correct price of an American option, computed at the stock value $S_0 - d$. The upper bound for the loss coming from the escrowed dividend approximation is reached exactly when $S_0 = S^*_{E,0}$. In this case, the following formulae give the maximum percentage loss at $t = 0$ under the three modelling environments considered and with the model parameters given above:

$$DL\%_{BS} = \frac{C_{BS}(S^*_{E,BS}) - (S^*_{E,BS} - K)}{C_{BS}(S^*_{E,BS})} = 0.63\%,$$
$$DL\%_{MRT} = \frac{C_{MRT}(S^*_{E,MRT}) - (S^*_{E,MRT} - K)}{C_{MRT}(S^*_{E,MRT})} = 0.48\%,$$
$$DL\%_{BTS} = \frac{C_{BTS}(S^*_{E,BTS}) - (S^*_{E,BTS} - K)}{C_{BTS}(S^*_{E,BTS})} = 1.14\%,$$

where we have suppressed for readability the dependence from time, and where the indexes BS, MRT, BTS mean that we have computed the relevant quantities under the Black-Scholes, Merton and Bates models, respectively. The investors can lose up to 1% of the market value of their option if they exercise according to the wrong boundary.

We finally check in our database if there are situations in which the stock reaches the value $S^*_{E,t}$ but not $S^*_{A,t}$ on the days before the ex-dividend dates. We find that this occurrence does happen, and not sporadically. For example, on May 10th, 2006, the Dupont stock closes at 45.71 dollars. The call option with $K = 30$ and $T = 0.45$ should not be exercised if the continuation value is computed correctly, but the option should be exercised if the continuation value is approximated with a European price. In this case, if an investor exercises his option wrongly, he will suffer a loss given by: $DL\%_{BS} = 0.07\%$, $DL\%_{MRT} = 2\%$, $DL\%_{BTS} = 0.06\%$. These are not small figures when compared to the bid-ask spreads.

Appendix B. Valuation by fast recursive projections

In this appendix, we detail the Recursive Projection method we developed to compute the option prices, and consequently the exercise boundaries, of our study. The main innovation of the methodology is that it can handle in a simple and fast way both multidimensional processes and discrete dividend distribution for the underlying asset.

Let us assume that we need the value of an option at some specific points in time, $(t_l)_{l=1,\cdots,L}$. We use the notation $V(x, t_l)$ for the value of the option at time t_l, and for a value of the underlying $S_{t_l} = x$. The value function $V(x, t_l)$ does not need to be the value of a call option. The Recursive Projections are a flexible tool that can value a wide range of path-dependent contracts (see Andricopoulos et al. (2007) for treatment of the kind of path-dependency that we can address). For readability, we do not include in the arguments of the value function the contract parameters T and K. At each time t_l, the holder of the option decides whether to exercise. Thus, the value of the option at time t_l is given by the maximum between the intrinsic value, i.e. the proceeds from exercising the option, and the continuation value of the option:

$$V(x, t_l) = \max\big\{H(x, t_l), \mathbb{E}\big[e^{-r(t_{l+1}-t_l)}V(S_{t_{l+1}}, t_{l+1})|S_{t_l} = x\big]\big\}. \tag{9}$$

$H(x, t_l)$ is a payoff function, as for instance $H(x, t_l) = \max\{x - K, 0\}$ in the case of a call option, or $H(x, t_l) = \mathbb{I}(x < K)$, $\mathbb{I}(\cdot)$ being an indicator function, for a digital put. At maturity, $t_L = T$, $V(x, T) = H(x, T)$. The main requirement to be able to apply the Recursive Projections, is that the discounted probability distribution (the so-called Green function), with respect to which the conditional expectation in (9) is taken, be known analytically, whether in direct form, or through its Fourier or Laplace transform. In the simple case when we have a direct analytical form for the Green function, as for instance in the Black-Sholes case, we can write the expectation in (9) as:

$$\mathbb{E}\big[e^{-r(t_{l+1}-t_l)}V(S_{t_{l+1}}, t_{l+1})|S_{t_l} = x\big] = \int G(x, t_l; y, t_{l+1})V(y, t_{l+1})dy. \tag{10}$$

The integral in Equation (10) is taken over the values $S_{t_{l+1}}$ that the underlying asset can take, and $G(x, t_l; y, t_{l+1})$ is the discounted transition density of the underlying asset form value x at t_l to the value y at time t_{l+1}. We now restrict our attention to the value of the derivative at a fixed grid of values of S_t: $(y_1, \ldots, y_N)$ with constant step $\Delta y = y_i - y_{i-1}$. The grid is the same for each value $t \in (t_l)_{l=1,\cdots,L}$. Then we can write the following equality in matrix form:

$$\mathbf{v}(t_l) \stackrel{\text{def}}{=} \max\{\boldsymbol{H}(t_l), \boldsymbol{G}(t_l; t_{l+1})\boldsymbol{V}(t_{l+1})\}, \tag{11}$$

where $\boldsymbol{V}(t_{l+1})$ is a $N \times 1$ vector with entries $\boldsymbol{V}(t_{l+1})_j = V(y_j, t_{l+1})$, $\boldsymbol{G}(t_l; t_{l+1})$ is a $N \times N$ matrix with entries $\boldsymbol{G}(t_l; t_{l+1})_{i,j} = G(y_i, t_l; y_j, t_{l+1})\Delta y$, and $\boldsymbol{H}(t_l)$ a $N \times 1$ vector with entries $\boldsymbol{H}(t_l)_j = H(y_j, t_l)$. It is easy to check that the multiplication of $\boldsymbol{V}(t_{l+1})$ by the i-th row of $\boldsymbol{G}(t_l; t_{l+1})$ is a discretization of the expectation on the right hand side of (9) for $x = y_i$. The vector $\mathbf{v}(t_l)$ is an approximation of the value of the contract $\boldsymbol{V}(t_l)$ at t_l, and can be fed again into Equation (11) to obtain $\mathbf{v}(t_{l-1})$:

$$\mathbf{v}(t_{l-1}) \stackrel{\text{def}}{=} \max\{\boldsymbol{H}(t_l), \boldsymbol{G}(t_{l-1}; t_l)\mathbf{v}(t_l)\}. \tag{12}$$

The main difference between Equations (11) and (12) is that in the former we assumed $\boldsymbol{V}(t_{l+1})$ known, as for instance at maturity where $\boldsymbol{V}(T) = \boldsymbol{H}(T)$, whereas in the latter, $\mathbf{v}(t_l)$ is an approximation of the real value. Equation (12) is at the heart of the Recursive Projections. It provides a recursive formula that relates the value of a contract at different points on a grid of values for the underlying, and at successive points in times. Equations (11) shows how the pricing of options that can be exercised at specific moments, as for instance Bermudan contracts, can be expressed as a series of matrix times vector operations. Furthermore if the time interval $\tau = t_{l+1} - t_l, l = 1, \ldots, L-1$, is constant and if the pricing operator enjoys a stationarity property (time translation invariance), then the matrix $\mathbf{G}(t_l; t_{l+1}) = \mathbf{G}(\tau)$ has constant entries, and the algorithm only involves one single computation of the matrix.

The methodology easily extends in the presence of discrete dividends paid on potential exercise dates. We only need to add the dividend δ to the continuation value in (12). Hence, in order to price an American option on a dividend-paying stock, Equation (12) must be modified by computing the state price density $G(x, t_l; y_j, t_{l+1})$ at the grid $(y_i - \delta(y_i))_{i=1,\ldots,N}$ for the conditioning value x, whenever t_l is dividend distribution date. The entries of the matrix $\mathbf{G}(t_l; t_{l+1})$ then become $\mathbf{G}_{ij} = G(y_i - \delta, t_l; y_j, t_{l+1})\Delta y$. Given the freedom in choosing where to sample G, $\delta(x)$ could be any function of x. If $\delta(x) = r_d\, x$, then we can accommodate for a proportional dividend. If $\delta(x) = d$, then we can accommodate for a discrete dividend amount d [13]. If $\delta(x) = 0$, then we return to the Bermudan option case. The value function $\mathbf{v}(t_l)$ still gives the value of the contract at the grid points $(y_1, \ldots, y_N)$; thus we can use its approximation $\mathbf{v}(t_l)$ as the input for the following step of the algorithm, and the recursive property of the algorithm is maintained. Figure 4 shows how the Recursive Projections easily accommodate dividend payments.

To carry out the analysis of Section 3.2.2, we need to be able to incorporate stochastic

[13] When we consider discrete dividends, we rule out arbitrage situations where the dividend is too large with respect to the stock price (see Haug, Haug, and Lewis (2003) for a discussion).

volatility and jumps in the description of the process. In the class of stochastic volatility models, there are two state variables, the underlying asset S_{t_l} and the variance $\sigma^2_{t_l}$. One of the most popular model belonging to this class is the Heston model. In the Heston model, we know the analytical form of the Fourier transform[14] of the bivariate state price density $G_2(S_{t_l}, \sigma^2_{t_l}, t_l; y, w, t_{l+1})$, which describes the discounted transition probability density from the asset level S_{t_l} and variance level $\sigma^2_{t_l}$ at time t_l to the asset level $y = S_{t_{l+1}}$ and variance level $w = \sigma^2_{t_{l+1}}$ at time t_{l+1}. Let its Fourier transform be $\hat{G}_2(S_{t_l}, \sigma^2_{t_l}, t_l; \lambda, \kappa, t_{l+1})$, so that $G_2(S_{t_l}, \sigma^2_{t_l}, t_l; y, w, t_{l+1}) = \frac{1}{4\pi^2} \iint d\lambda d\kappa e^{-\iota(\lambda y + \kappa w)} \hat{G}_2(S_{t_l}, \sigma^2_{t_l}, t_l; \lambda, \kappa, t_{l+1})$, where ι is the imaginary unit. There are different ways in which we can obtain a numerical approximation of G_2. One straightforward possibility is to consider the function $\Gamma_2(y_i, w_p, t_l; y_j, w_q, t_{l+1})$ obtained by numerically inverting $\hat{G}_2$ by using a Fast Fourier Transform (FFT)[15], and that gives the approximated discounted transition probability density from the asset level $S_{t_l} = y_i$ and variance level $\sigma^2_{t_l} = w_p$ at time t_l to the asset level $y_j = S_{t_{l+1}}$ and variance level $w_q = \sigma^2_{t_{l+1}}$ at time t_{l+1}, where $(w_1, \ldots, w_W)$ is fixed, equispaced of values taken by the stochastic variance σ^2_t, with step Δw.

In the stochastic volatility framework, the recursive pricing relation for an option that can be exercised on a discrete set of dates consists in moving backwards in time as in Equation (12) with:

$$V(x, \xi, t_l) = \max\big\{H(x, t_l), \mathbb{E}\big[e^{-r(t_{l+1}-t_l)} V(S_{t_{l+1}}, \sigma^2_{t_{l+1}}, t_{l+1}) \big| S_{t_l} = x, \sigma^2_{t_l} = \xi\big]\big\}. \tag{13}$$

Thus, the recursive projections in the Heston model are the discrete counterpart of (13). We give in the following proposition the convergence properties of the algorithm.

Proposition 1. *Let $H(y, w, T)$ be such that $|H(y, w, T) - H(y', w', T)| < C\underline{\Delta}$ for a positive constant C, and for $|y - y'| < \Delta y$, $|w - w'| < \Delta w$ and $\underline{\Delta} = \sqrt{(\Delta y)^2 + (\Delta w)^2}$. Let $v_{ip}(t_l)$ be*

[14] For an explicit formula, see Griebsch (2013).

[15] In Cosma, Galluccio, Pederzoli, and Scaillet (2016) we describe a more general approach to obtain an approximation of the G_2 function in terms of functional projections. The function $\Gamma_2(y_i, w_p, t_l; y_j, w_q, t_{l+1})$ is an approximation of $G_2(S_{t_l}, \sigma^2_{t_l}, t_l; y, w, t_{l+1})$ since it is obtained by a discrete transform, in place of the continuous transform we would need to recover the true G_2 transition density.

defined for a set of dates $(t_l)_{l=1,\dots,L}$, with $t_L = T$, as follows:

$$v_{ip}(t_l) = \max\Bigg\{H(y_i, w_p, t_l), \sum_{\substack{j=-\infty \\ q=1}}^{\infty} \Gamma_2(y_i, w_p, t_l; y_j, w_q, t_{l+1}) H(y_j, w_q, t_{l+1}) \sqrt{\Delta y \Delta w}\Bigg\}, \text{ for } l = L-1, \tag{14}$$

$$v_{ip}(t_l) = \max\left\{H(y_i, w_p, t_l), \sum_{\substack{j=-\infty \\ q=1}}^{\infty} \Gamma_2(y_i, w_p, t_l; y_j, w_q, t_{l+1}) v_{jq}(t_{l+1}) \sqrt{\Delta y \Delta w}\right\}, \qquad \text{for } l = 1, \dots, L-2. \tag{15}$$

Then, for each t_l in $(t_1, \dots, t_{L-1})$, the approximated values $v_{ip}(t_l)$ defined in (14) *and* (15) *converge to the true value $V(y_i, w_p, t_l)$ with an approximation error of the order $O\big(\underline{\Delta}^2\big)$.*

In almost all applications $H(y, w, t_l)$ only depends on the value y taken by the underlying asset at t_l, and the computed price $v_{ip}(t_{l-1})$ depends on the stochastic variance only through the conditioning value $\sigma^2_{t_{l-1}} = w_p$.

Proof. See Section A of the online supplementary materials. □

Appendix C. Exercise boundary with a continuous dividend

In this section, we provide theoretical insights on why the early exercise boundary in the stochastic volatility models is lower than the boundary implied by the Black-Scholes model when the underlying stock distributes a continuous dividend yield. The discussion also applies to jump-diffusion processes for longer maturities, because in this case, as explained in Section 3.4, they are observationally similar to stochastic volatility processes.

Following Kim (1990) and Jamshidian (1992), we can decompose the value $V(S_t, t)$ of an American option into two components, namely, the European value $V^E(S_t, t)$ and the early exercise premium $V^A(S_t, t)$, such that:

$$\begin{aligned} V(S_t, t) &= V^E(S_t, t) + V^A(S_t, t) \qquad (16) \\ &= e^{-r(T-t)} \mathbb{E}\big[(S_T - K)_+ \big| S_t, \sigma_t^2\big] + \int_t^T e^{-r(s-t)} \mathbb{E}\big[(r_d S_s - rK)\mathbb{I}_{(S_s > S_s^*)} \big| S_t, \sigma_t^2\big] ds, \end{aligned}$$

where S_s^* is the early exercise boundary at time s and $\mathbb{I}_{(S_s>S_s^*)}$ equals one if, at time s, the stock is in the exercise region, otherwise zero. We can interpret $V^A(S_t, t)$ as a continuum of European call options with maturity $T - s$, strike price S_s^*, and payoff $r_d S_s - rK$. For each of these European options, we can apply the results of Table II in Hull and White (1987) who compare the values of European options under general stochastic volatility dynamics with the Black-Scholes price. Call values under the stochastic volatility assumption are lower when the contracts are at-the-money and $\rho \leq 0$. The continuum of contracts composing the $V^A(S_t, t)$ are at-the-money when $S_s = S_s^*$. As confirmed from our numerical simulations, the S_s^* values are distributed in the region immediately above $S = 150$, that is, exactly where the price of the American option under the Heston model is lower than that under the Black-Scholes model, so that the early exercise premium is lower under the former of the two modelling environments. This in turn explains the negative bump in the right graph of Panel B of Figure 1.