How Much is the Gap?
-
Efficient Jump Risk-Adjusted Valuation of Leveraged Certificates

Ally Quan Zhang*
Swiss Finance Institute
and
Institut für Banking & Finance
Universität Zürich

Matthias Thul†‡
Commerzbank AG
Equity Markets and Commodities

First Version: July 16, 2013
Last Update: December 18, 2016

Abstract

This paper develops a novel and highly efficient numerical algorithm for the gap risk-adjusted valuation of leveraged certificates. The existing literature relies on Monte Carlo simulations, which are not fast enough to be used in a market making environment. This is because issuers need to compute thousands of price updates per second. By valuing leveraged certificates as multi-window barrier options, we explicitly model random jumps that occur at known times, such as between the exchange closing and re-opening. Our algorithm combines the one-day transition probability with Simpson's numerical integration rule. This yields a backward induction scheme which requires a significantly coarser spatial and time grid than finite difference methods. We confirm its robustness and accuracy through Monte Carlo simulations.

Keywords: leveraged certificates, barrier options, overnight gap, risk management

JEL Classification: G13, C63
MS Classification (2010): 35K20, 60G51

*(Corresponding Author) Address: Ally Quan Zhang, Institut für Banking & Finance, Plattenstr. 32, 8032 Zürich, ZH, Switzerland. e-mail: quan.zhang@bf.uzh.ch.

†Address: Matthias Thul, Commerzbank AG, Equity Derivatives Flow Trading, Equity Markets and Commodities, Mainzer Landstr. 153, 60327 Frankfurt am Main, e-mail: matthias.thul@gmail.com. The opinions expressed in this article are solely those of the authors and do not reflect any views by Commerzbank AG.

‡We thank David Colwell, David Feldman, Felix Kübler, Detlef Mages, participants at the 9th World Congress of the Bachelier Finance Society as well as one anonymous reviewer for their helpful comments and suggestions. Ally Quan Zhang acknowledges financial support from the Swiss Finance Institute.

1 Introduction

Leveraged certificates are knock-out barrier options with an at-the-money or in-the-money barrier. They have an almost constant delta when in-the-money and are thus often marketed as offering a risk exposure similar to that of a forward or futures contract. However, they are in fact limited liability assets with their premium payable upfront. Thus, they don't subject the buyer to margin calls in case of an adverse movement of the underlying asset. In contrast to forward contracts, the buyer is exposed to the seller's default risk but not vice versa. As their name indicates, they usually offer a significant gearing on the capital invested.

Leveraged certificates are commonly issued by financial institutions and sold to institutional and retail investors. There exist highly liquid on-exchange and over-the-counter markets for these products. According to Deutscher Derivate Verband (2015), as of December 2015, there were 355,632 leveraged certificates issued in the German on-exchange markets alone with an open interest of MEUR 717. They accounted for an average of daily turnover of KEUR 55,788 from 11,870 transactions during 2015.

From the issuer's perspective, confining the buyer's maximum loss to the initial premium comes at the cost of increased market risk. While forward and futures contracts can be hedged statically, leveraged certificates generally require a dynamic hedge that needs to be unwound upon a knock-out event. In any pure diffusion model with continuous trading, the corresponding underlying asset price is always equal to the barrier and there is no gap risk. The prices derived from such models almost always undervalue the leveraged certificate when the actual market dynamics exhibit discontinuities.

Despite the significant impact that the gap risk has on the prices of leveraged certificates, to the best of our knowledge, there exists no viable approach for its real-world valuation so far. This paper is the first to propose a numerical algorithm for the gap risk-adjusted valuation of leveraged certificates, which is sufficiently fast to be adopted in a market making environment. Our valuation framework accounts for random jumps that occur at known times and are normally distributed in logarithmic returns. While we focus on jumps that are induced by the discontinuity of exchange trading hours, our approach also applies to scheduled economic events such as central bank monetary policy decisions and company earnings releases, both intraday and overnight. We do not explicitly incorporate these jumps in the underlying asset price process but instead take the equivalent approach of modeling the leveraged certificate to have multiple barrier windows. The time durations for which the barrier is inactive are generally not set according to the actual lengths of the overnight periods or the economic events. Instead they are chosen such as to match the desired jump standard deviation.

Boes et al. (2007) find that the overnight jump component represents a significant proportion of the overall variance. Accurately valuing the corresponding gap risk in leveraged certificates is especially important since they are typically traded during extended hours. That is, they are quoted in the over-the-counter market or on listed derivatives exchanges even when the exchange of the underlying asset is closed. During these times, they are priced using the market makers' estimate of the underlying asset price. Since this shadow spot is not directly observable, the leveraged certificates cannot knock-out even when it breaches the barrier. Just like Boes et al. (2007), we model logarithmic overnight returns to be normally distributed.

A large strand of the existing literature on leveraged certificates focuses on empirically analyzing the price setting behavior, predominantly in the German market. Muck (2006) and Wilkens and Stoimenov (2007) find a significant and consistent overpricing across issuers. Their theoretical reference prices or super-hedging boundaries, however, are obtained under the assumption of continuous trading. Subsequent literature shows that a significant portion of the apparent markups can be attributed to the overnight gap risk.

Entrop et al. (2009) and Rossetto and van Bommel (2009) both recognize the need to value the gap risk in leveraged certificates associated with an overnight trading halt. While Entrop et al. (2009) impose normally distributed close-to-open jumps, Rossetto and van Bommel (2009) sample from the historical overnight return distribution. Both papers value open-end contracts through Monte Carlo simulations. They conclude that overpricing is still significant, even after accounting for gap risk. While straightforward to implement, their low computational speed renders Monte Carlo methods unsuitable for market making purposes. Large issuers simultaneously quote tens of thousands of leveraged certificates and their prices need to be continuously recomputed to avoid arbitrage. Our proposed approach resolves this problem.

Baller et al. (2016) don't impose a valuation model for the gap risk premium but jointly estimate its magnitude together with the actual markups that are controlled by the issuers. The authors argue that overnight gap risk is mainly driven by the short term volatility and use hourly dummies to capture its intraday dynamics. They find that the priced overnight gap risk ceteris paribus increases during the trading day and for leveraged certificates closer to the knock-out. This is in rough agreement with our model-based analysis. However, the authors only consider the time-of-day and barrier distance as determinants for the gap risk premium but not the number of trading days to maturity. Consequently, they likely overestimate the time decay in the actual margins set by the issuers.

Our key contribution to the literature is to propose a novel grid-based numerical algorithm for the gap risk-adjusted valuation of leveraged certificates. This algorithm is the first to be sufficiently fast for real-time pricing. Our main idea is to construct a backward induction scheme which exploits the availability of closed-form solutions to the one-day transition probability. This allows us to take time steps which are significantly larger than in standard finite difference methods. To further improve the convergence and stability, we introduce a control variate correction. Our algorithm yields pricing errors in the order of magnitude of 0.01% of the underlying asset price with as few as 50 spatial grid points. Its single-threaded pricing times are in the low single digit milliseconds on commodity hardware. A caching technique can be used to further reduce re-computation times. To confirm the robustness of our algorithm, we employ Monte Carlo simulations with a wide range of randomly generated contractual and market parameters. While our method is tailored to value leveraged certificates, the underlying approach can be adopted to handle other complex barrier structures. Finally, we show how our framework can also be extended to general jump events with a predictable time of occurrence as well as discrete barriers.

2 Continuous Trading

This section formally defines leveraged certificates and discusses their valuation within the Black and Scholes (1973) framework. We maintain the assumptions of (i) frictionless markets and (ii) the underlying asset following a constant coefficient geometric Brownian motion throughout this paper. The key additional premise made in this section is that continuous trading is possible. That is, we value contingent claims as if the market for the primary securities never closes. This assumption is dropped in Section 3, where we explicitly model the two-phase structure of exchange opening hours. The valuation functions derived in this section then serve as a benchmark. Like most analytical results in this paper, we obtain them by applying the method of images to higher-order power binaries; see Buchen (2001) and Skipper and Buchen (2003). We refer to Appendices A through C for a summary of the key results.

2.1 Model Setup

Let $W^* = \{W_t^* : t \in \mathbb{R}_{+,0}\}$ be a one-dimensional standard Brownian motion on a complete filtered probability space $(\Omega, \mathfrak{F}, \mathbb{F}, \mathbb{P}^*)$. As we are solely interested in pricing, we interpret $\mathbb{P}^*$ to be the risk-neutral probability measure corresponding to the bank account numéraire. The filtration $\mathbb{F} = (\mathfrak{F}_t)_{t \in \mathbb{R}_{+,0}}$ is the $\mathbb{P}^*$-augmentation of the natural filtration induced by the process W^*. The frictionless market consists of two primary assets. The first is a bank account $B = \{B_t : t \in \mathbb{R}_{+,0}\}$ with non-random dynamics

$$\mathrm{d}B_t = rB_t\mathrm{d}t,$$

where the risk-free interest rate $r \in \mathbb{R}$ is a constant and $B_0 = 1$. There further exists a single risky limited-liability spot asset $S = \{S_t : t \in \mathbb{R}_{+,0}\}$ with geometric Brownian motion dynamics

$$\mathrm{d}S_t = (r - \delta)S_t\mathrm{d}t + \sigma S_t\mathrm{d}W_t^*$$

and initial value $S_0 \in \mathbb{R}_+$. The continuous dividend yield $\delta \in \mathbb{R}$ and the diffusion coefficient $\sigma \in \mathbb{R}_+$ are constants. An extension to time-dependent parameters is straightforward. We deliberately exclude this generalization as it does not add significant insights while at the same time complicating the notation.

2.2 Contract Definition

A bull (bear) leveraged certificate is essentially a down-and-out call (up-and-out put) option with an at- or in-the-money barrier. It pays a non-negative rebate upon the first hitting time of the barrier with an amount equal to the then-current intrinsic value.

Definition 2.1 (Leveraged Certificate).

Let $T \in \mathbb{R}_+$ be the option maturity. We define the first hitting time of S to the barrier $B \in \mathbb{R}_+$ as

$$\nu = \inf \{t \in \mathbb{R}_{+,0} : \phi S_t \leq \phi B\},$$

where $\phi \in \{-1, +1\}$ indicates a bear or bull certificate. The terminal payoff of the leveraged certificate is given by

$$V_T = \phi (S_T - K) 1 \{\nu > T\} + e^{r(T-\nu)} (\phi S_\nu - \phi K)^+ 1 \{\nu \leq T\},$$

where $K \in \mathbb{R}_+$ is the strike price satisfying $\phi(B - K) \geq 0$. $\triangle$

Throughout this paper, we follow the convention to set $\nu = \infty$ on the set where S never breaches B. Note that in the present setup, there is no need to floor the rebate at zero, since $S_\nu = B$ whenever $\nu < \infty$. However, in the two-phase structure of Section 3, this no longer generally holds.

We remark that the payoff profile discussed here is a slightly simplified abstraction. In particular, the contractual definition of the rebate amount greatly varies. It is negotiated in the over-the-counter market or fixed by the issuing institution. The reference price used to calculate the intrinsic value is often not exactly equal to the underlying asset price at the first hitting time. Instead it is linked to the path over a short time interval immediately following the knock-out event.

From the seller's point of view, the prime reason for introducing an in-the-money barrier is to reallocate parts of the risk of underlying asset price jumps to the buyer. The majority of listed contracts however has an at-the-money barrier and fully exposes the issuer to the gap risk. That is, while the buyer receives a zero payoff in any case, the profit and loss of the seller depends on the underlying asset price at which she can unwind her delta hedge position.

In the current framework however, markets are frictionless, trading is continuous and the underlying asset follows a pure diffusion process. Thus, it is always possible to close out the complete delta hedge position exactly on the barrier. Consequently, there exists no gap risk.

2.3 Valuation

We apply the method of images to obtain the valuation function for leveraged certificates when trading is continuous. While these results are not novel, they serve two purposes later in this paper. First, they constitute the terminal components for the multiple window barrier options considered in Section 3. Second, they are used as the benchmark prices that we compare the gap risk adjusted values with. The non-negative difference between these two valuation functions is then interpreted as the gap risk premium.

Proposition 2.1 (Valuation of Leveraged Certificates).
The value of a leveraged certificate is given by

$$\begin{aligned}\tilde{V}(S,\tau) &= \underbrace{\phi\left(\mathcal{A}_B^{\phi}(S,\tau) - K\mathcal{B}_B^{\phi}(S,\tau) - \overset{B}{\mathcal{I}}\left\{\mathcal{A}_B^{\phi}(S,\tau) - K\mathcal{B}_B^{\phi}(S,\tau)\right\}\right)}_{\textit{terminal payoff value}} \\ &\quad + \underbrace{\phi(B-K)\mathcal{R}_B^{\phi}(S,\tau)}_{\textit{rebate value}}.\end{aligned}$$

Here, $\overset{B}{\mathcal{I}}$ is the image operator; see Definition A.3. $\mathcal{A}_\xi^s$, $\mathcal{B}_\xi^s$ and $\mathcal{R}_B^\psi$ are the valuation functions of first-order asset and bond binaries as well as a pay-at-hit rebate, respectively; see Definitions B.2 and C.2.

Proof Let

$$\tilde{V}(S,\tau) = \tilde{V}^1(S,\tau) + \tilde{V}^2(S,\tau),$$

where $\tilde{V}^1(S,\tau)$ and $\tilde{V}^2(S,\tau)$ correspond to the valuation functions of the terminal payoff and the rebate, respectively. $\tilde{V}^1(S,\tau)$ satisfies an initial boundary value problem (IBVP) with barrier B, flag ϕ and initial condition

$$\begin{aligned}f(S) &= \phi(S-K)1\{\phi S > \phi B\} \\ &= \phi\left(\mathcal{A}_B^{\phi}(S,0) - K\mathcal{B}_B^{\phi}(S,0)\right);\end{aligned}$$

see Definitions A.2 and B.2. The expression for $\tilde{V}^1(S,\tau)$ then follows by applying the method of images in Proposition A.1. $\tilde{V}^2(S,\tau)$ is the valuation function of $\phi(B-K)$ units of a pay-at-hit rebate; see Proposition C.2. □

Note, that all valuation functions for barrier options are conditional on no prior crossing of the barrier. For brevity, we do not make this explicit. Lemmata 2.1 and 2.2 give two useful auxiliary results.

Lemma 2.1 (Special Case Valuation of Leveraged Certificates).
When both the risk-free interest rate and the dividend yield are zero, then the value of the leveraged certificate is equal to its intrinsic value at all times, i.e.

$$\tilde{V}(S,\tau) = \phi(S - K).$$

This can be shown by explicitly setting $r = \delta = 0$ in the valuation function in Proposition 2.1. It is however more instructive to consider the corresponding semi-static replication strategy since the latter applies to all pure diffusion processes. In case of a bull (bear) leveraged certificate, it corresponds to buying (selling) one stock and borrowing (lending) the strike price. This position is unwound either upon a knock-out or the option maturity, whatever event comes first. Since the underlying asset has continuous sample paths, its price upon the first hitting time is always equal to the barrier. Consequently, the then-current value of the replication portfolio is equal to the rebate.

Lemma 2.2 (Deep-in-the-Money Value of Leveraged Certificates).
We have

$$\lim_{\phi \ln(S/K) \to \infty} \tilde{V}(S,\tau) = \phi\left(Se^{-\delta\tau} - Ke^{-r\tau}\right).$$

3 Valuation under Discontinuous Trading

In this section, the assumptions of frictionless markets and the underlying asset dynamics following a constant coefficient geometric Brownian motion are retained. However, we now explicitly account for the predictable discontinuities in the underlying asset price, which are exogenously imposed by the limited opening hours of real world exchanges.

3.1 Two-Phase Model Setup

We explicitly model each trading day to consist of two successive phases. Let T_i^{O} and T_i^{C} be the opening and closing times of the underlying asset's market on the i-th calendar day, respectively. For convenience, and without loss of generality, the current day corresponds to $i = 0$ and we define i-th calendar day to start at the opening T_i^{O} and to end at the next day's opening T_{i+1}^{O}. During the intraday period $\left[T_i^O, T_i^C\right)$, the market is open and the barrier is active. It is followed by an overnight period $\left[T_i^C, T_{i+1}^O\right)$, during which the market is closed and the barrier is not monitored. To simplify notation, we henceforth assume that the length of each intraday and overnight period is the same. Possible generalizations and practical considerations are discussed in Section 7.

Consider a leveraged certificate that expires at the closing time of the n-th trading day T_n^{C}. Since the underlying asset is only traded during the intraday periods, the barrier is monitored on the intervals $\left[t, T_0^{\mathrm{C}}\right)$ and $\left[T_i^{\mathrm{O}}, T_i^{\mathrm{C}}\right)$ for $i = \{1, 2, \ldots, n\}$. Consequently, the leveraged certificate can be viewed as a multi-window barrier option with a number of windows equal to the number of intraday periods until maturity. We define the set of monitoring intervals as

$$\mathcal{T}^{\mathrm{O}} = \bigcup_{i=0}^{\infty} \left[T_i^{\mathrm{O}}, T_i^{\mathrm{C}}\right).$$

The first hitting time of S to the barrier B during an intraday period is then given by

$$\nu = \inf\left\{t \in \mathcal{T}^{\mathrm{O}} : \phi S_t \leq \phi B\right\}.$$

We obtain the same expression for the terminal payoff as in Section 2.2 but now with respect to the new definition of the first hitting time.

3.2 Spatial Grid Construction

The key idea behind our numerical algorithm is to discretize the spatial domain through $m \in \mathbb{N}$ spot grid points $\{S_j\}_{j=1}^m$ and find an approximation for the one-day backward induction operator. That is, we aim at mapping the set of prices $\left\{V\left(S_j, T_{i+1}^{\mathrm{O}}\right)\right\}_{j=1}^{m}$ on the market opening of the $(i+1)$-st trading day into the corresponding set of prices $\left\{V\left(S_j, T_i^{\mathrm{O}}\right)\right\}_{j=1}^{m}$ on the market opening of the previous trading day. We suggest to employ numerical integration and exploit that the conditional transition probabilities between the node points $\left(S_j, T_i^{\mathrm{O}}\right)$ and $\left(S_k, T_{i+1}^{\mathrm{O}}\right)$ for $j, k \in \{1, 2, \ldots, m\}$ can be computed in closed-form. Since their expressions only involve a one-dimensional standard normal cumulative distribution function (CDF), analytical approximations up to machine precision are readily available.

The spatial grid is constructed through

$$S_j = Bg^{\phi}(j) \qquad \text{for } j \in \{1, 2, \ldots, m\},$$

where the function $g : \mathbb{R}_{+,0} \to \mathbb{R}_+$ is defined by

$$g^{\phi}(x) = \exp\left\{\phi\kappa_{\max}\left(\frac{x}{m}\right)^{\xi}\right\}.$$

Here, $\kappa_{\max} \in \mathbb{R}_+$ is the maximum logarithmic moneyness of the spot grid such that $g^{\phi}(m) = e^{\phi\kappa_{\max}}$. It can, for example, be set to a fixed number of standard deviations of annual logarithmic returns and we find that $\kappa_{\max} = \sigma$ works well in practice. The parameter $\xi \geq 1$ is a stretching factor that determines by how much the logarithmic spacing in the spot grid is increasing as we move away from the barrier. The special case $\xi = 1$ corresponds to an equally spaced logarithmic spot grid. Figure 1 gives an example to illustrate this construction for a bull leveraged certificate and two different values of ξ. As will be discussed later, the variability of the gap risk premium is highest near the barrier. An equidistant grid in the logarithmic spot price is either not sufficiently dense in this region or requires too many grid points. Through numerical experiments, we find that a stretching factor of $\xi = 1.5$ yields good results.

Insert Figure 1 about here.

3.3 Backward Induction Step

Let $\hat{\tau}^{\mathrm{I}} = T_i^{\mathrm{C}} - T_i^{\mathrm{O}}$ and $\hat{\tau}^{\mathrm{D}} = T_{i+1}^{\mathrm{O}} - T_i^{\mathrm{O}}$ be the length of an intraday period and a trading day, respectively, and define the survival domain by $\mathcal{D} = \{S \in \mathbb{R}_+ : \phi S > \phi B\}$. We start by giving an analytical expression for the valuation function $V\left(S, T_i^{\mathrm{O}}\right)$ on the market opening of the i-th trading day in terms of its values on the market opening of the following trading day, given that $i < n$. This expression can be decomposed into one term, $C\left(S, T_i^{\mathrm{O}}\right)$, which corresponds to the continuation value at T_{i+1}^{O} as well as two rebate related terms. The first, $R^{\mathrm{A}}(S)$, corresponds to a knock-out at the market opening T_{i+1}^{O} and the second, $R^{\mathrm{B}}(S)$, to a knock-out during the intraday period $\left[T_i^{\mathrm{O}}, T_i^{\mathrm{C}}\right)$; see also Figure 2.

Insert Figure 2 about here.

Lemma 3.1 (Interday Backward Induction).
For $i < n$, the solution for $V\left(S, T_i^O\right)$ is given by

$$\begin{aligned} V\left(S, T_i^O\right) &= \underbrace{e^{-r\hat{\tau}^D} \int_{\mathcal{D}} V\left(y, T_{i+1}^O\right) f_1(y|S)\mathrm{d}y}_{C\left(S, T_i^O\right) \sim \text{ continuation value}} + \underbrace{e^{-r\hat{\tau}^D} \phi \int_K^B (y-K) f_1(y|S)\mathrm{d}y}_{R^A(S) \sim \text{ value of rebate at } T_{i+1}^O} \\ &\quad + \underbrace{\phi(B-K)\int_0^{\hat{\tau}^I} e^{-ry} f_2(y|S)\mathrm{d}y}_{R^B(S) \sim \text{ value of intraday rebate}}, \end{aligned}$$

where

$$\begin{aligned} f_1(y|x)\mathrm{d}y &= \mathbb{P}^*\left\{ S_{T_{i+1}^O} \in \mathrm{d}y \,\cap\, \nu \geq T_{i+1}^O \middle| \nu > T_i^O,\ S_{T_i^O} = x \right\}, \\ f_2(y|x)\mathrm{d}y &= \mathbb{P}^*\left\{ \nu \in T_i^O + \mathrm{d}y \middle| \nu > T_i^O,\ S_{T_i^O} = x \right\}. \end{aligned}$$

Here, $f_1(y|x)\mathrm{d}y$ is the probability of an underlying asset path starting at x on the market opening T_i^{O} of the i-th trading day and ending in dy on the market opening T_{i+1}^{O} of the following trading day without breaching the barrier during the intraday period $\left[T_i^{\mathrm{O}}, T_i^{\mathrm{C}}\right)$. We note that $f_1(y|x)$ is not a probability density function (PDF).

Lemma 3.2 (Conditional Transition Probability Function).
The conditional transition probability function is given by

$$f_1(y|x) = e^{r\hat{\tau}^D}\left(\frac{\partial \mathcal{B}_{By}^{\phi-}}{\partial y}\left(x, \hat{\tau}^I, \hat{\tau}^D\right) - \overset{B}{\mathcal{I}}\left\{ \frac{\partial \mathcal{B}_{By}^{\phi-}}{\partial y}\left(x, \hat{\tau}^I, \hat{\tau}^D\right) \right\} \right),$$

where

$$\frac{\partial \mathcal{B}^{s_1 s_2}_{\xi_1 \xi_2}}{\partial \xi_2}(S,\tau_1,\tau_2) = -e^{-r\tau_2}\frac{s_2}{\xi_2\sigma\sqrt{\tau_2}}\mathcal{N}'(\alpha_{0,2})\,\mathcal{N}(\hat{\alpha})$$

and

$$\begin{aligned}\hat{\alpha} &= \frac{\alpha_{0,1}-\rho_{1,2}\alpha_{0,2}}{\sqrt{1-(\rho_{1,2})^2}},\\ \rho_{1,2} &= s_1 s_2\sqrt{\tau_1/\tau_2}.\end{aligned}$$

All other notation is as in Appendix B; see in particular Definition B.2.

Proof Define

$$F_1(y|x) = \int_0^y f_1(z|x)\mathrm{d}z.$$

We recognize this as the undiscounted value of a bond binary that pays off one unit of cash at the market opening T^{O}_{i+1} of the following trading day if the then-current underlying asset price is below y and the barrier was not breached during the intraday period $\left[T^{\mathrm{O}}_i, T^{\mathrm{C}}_i\right)$. The solution is given by

$$F_1(y|x) = e^{r\hat{\tau}^{\mathrm{D}}}\left(\mathcal{B}^{\phi-}_{By}\left(x,\hat{\tau}^{\mathrm{I}},\hat{\tau}^{\mathrm{D}}\right) - \overset{B}{\mathcal{I}}\left\{\mathcal{B}^{\phi-}_{By}\left(x,\hat{\tau}^{\mathrm{I}},\hat{\tau}^{\mathrm{D}}\right)\right\}\right).$$

The partial derivative corresponds to a special case of Proposition B.4. □

While we numerically approximate the continuation value in the following section, the expressions for the two rebate related terms are sufficiently low-dimensional to be used in their analytical form.

Lemma 3.3 (Rebate Valuation).
The values of the two rebate components in Lemma 3.1 are

$$\begin{aligned} R^A(S) &:= \underbrace{e^{-r\hat{\tau}^D}\phi\int_K^B (y-K) f_1(y|S)\mathrm{d}y}_{\text{value of rebate at } T^O_{i+1}} \\ &= \phi\Bigg(\mathcal{Q}^{\phi\phi}_{BK}\left(S,\hat{\tau}^I,\hat{\tau}^D\right) - \overset{B}{\mathcal{I}}\left\{\mathcal{Q}^{\phi\phi}_{BK}\left(S,\hat{\tau}^I,\hat{\tau}^D\right)\right\} - \mathcal{Q}^{\phi\phi}_{BB}\left(S,\hat{\tau}^I,\hat{\tau}^D\right) \\ &\quad + \overset{B}{\mathcal{I}}\left\{\mathcal{Q}^{\phi\phi}_{BK}\left(S,\hat{\tau}^I,\hat{\tau}^D\right)\right\} - (B-K)\mathcal{B}^{\phi\phi}_{BB}\left(S,\hat{\tau}^I,\hat{\tau}^D\right)\Bigg)\end{aligned}$$

and

$$
\begin{aligned}
R^B(S) &= \underbrace{\phi(B-K)\int_0^{\hat{\tau}^I} e^{-ry} f_2(y|S)\mathrm{d}y}_{\textit{value of intraday rebate}} \\
&= \phi(B-K)\left\{\mathcal{U}_B^{\phi}(S) - B^{-\beta(-\phi)}\left({}^{\beta(-\phi)}\mathcal{P}_B^{\phi}\left(S,\hat{\tau}^I\right) - \overset{B}{\mathcal{I}}\left\{{}^{\beta(-\phi)}\mathcal{P}_B^{\phi}\left(S,\hat{\tau}^I\right)\right\}\right)\right\}.
\end{aligned}
$$

Here, $\mathcal{P}_\xi^s$ *and* $\mathcal{U}_B^\psi$ *are the valuation functions of a first-order power binary and a perpetual binary, respectively; see Definitions B.1 and C.1. All other notation is as in Appendix B; see in particular Definition B.2.*

3.4 Approximation of the Continuation Value

We proceed to propose an approximation scheme for the continuation value component. We first split up the domain of integration into two parts.

$$
\begin{aligned}
C\left(S,T_i^{\mathrm{O}}\right) &= \underbrace{e^{-r\hat{\tau}^{\mathrm{D}}}\int_{\mathcal{D}} V\left(y,T_{i+1}^{\mathrm{O}}\right) f_1(y|S)\mathrm{d}y}_{\text{continuation value}} \\
&= e^{-r\hat{\tau}^{\mathrm{D}}}\left(\int_{\mathcal{D}\backslash\mathcal{D}_\infty} V\left(y,T_{i+1}^{\mathrm{O}}\right) f_1(y|S)\mathrm{d}y + \int_{\mathcal{D}_\infty} V\left(y,T_{i+1}^{\mathrm{O}}\right) f_1(y|S)\mathrm{d}y\right).
\end{aligned}
$$

$\mathcal{D}\backslash\mathcal{D}_\infty$ corresponds to the spatial domain covered by the spot grid and $\mathcal{D}_\infty$ is the far-field domain. Our analysis thus far has been exact. We could directly approximate the first term through Gaussian quadrature

$$
e^{-r\hat{\tau}^{\mathrm{D}}}\int_{\mathcal{D}\backslash\mathcal{D}_\infty} V\left(y,T_{i+1}^{\mathrm{O}}\right) f_1(y|S)\mathrm{d}y \approx e^{-r\hat{\tau}^{\mathrm{D}}}\sum_{j=1}^{m} w_j V\left(S_j,T_{i+1}^{\mathrm{O}}\right) f_1\left(S_j|S\right),
$$

where $\boldsymbol{w}$ is an m-dimensional vector of integration weights. Since the spot grid is in general not equally spaced, even in its logarithms, we could use the trapezoidal rule but not many higher-order rules. Instead, we first transform the domain of integration by applying the change of variables $z = \left(g^\phi\right)^{-1}(y/B)$ such that $\mathrm{d}y = h(z)\mathrm{d}z$, where

$$
\begin{aligned}
\left(g^\phi\right)^{-1}(x) &= \left(\frac{1}{\phi\kappa_{\max}}\ln(x)\right)^{1/\xi} m, \\
h(z) &= \frac{\phi\kappa_{\max}\xi}{m}\left(\frac{z}{m}\right)^{\xi-1} Bg^\phi(z).
\end{aligned}
$$

We then get

$$
\begin{aligned}
& e^{-r\hat{\tau}^{\mathrm{D}}} \int_{\mathcal{D}\setminus\mathcal{D}_\infty} V\left(y, T_{i+1}^{\mathrm{O}}\right) f(y|S)\mathrm{d}y \\
= \; & e^{-r\hat{\tau}^{\mathrm{D}}} \int_0^m V\big(Bg^\phi(z), T_{i+1}^{\mathrm{O}}\big) f\big(Bg^\phi(z)\big|S\big) h(z)\mathrm{d}z \\
\approx \; & e^{-r\hat{\tau}^{\mathrm{D}}} \sum_{j=1}^m w_j V\left(S_j, T_{i+1}^{\mathrm{O}}\right) f\left(S_j | \, S\right) h(j).
\end{aligned}
$$

Here we used that by construction, $\left(g^\phi\right)^{-1}(B/B) = 0$ and $\left(g^\phi\right)^{-1}(S_m/B) = m$. In case of a bull leveraged certificate, we thus transform the domain of integration from $[B, S_m]$ to $[0, m]$. We approximate this integral on the equally spaced grid $\{0, 1, \ldots, m\}$. This allows for an application of Simpson's rule, which corresponds to setting

$$
w_j = \begin{cases} 1/3 & \text{if } j = m \\ 2/3 & \text{if } j \text{ is even} \\ 4/3 & \text{if } j \text{ is odd} \end{cases} \cdot
$$

Since $h(0) = 0$, the weight w_0 can be directly set to zero. We choose m to be an even number so that the total number of intervals is even as well.

Now let $W(S, t)$ be the value of the otherwise identical barrier option with continuous monitoring. Let $\hat{\tau}_i^{\mathrm{O}} = T_n^{\mathrm{C}} - T_i^{\mathrm{O}}$ be the time-to-maturity at the opening of the i-th trading day. We approximate the second term through

$$
\begin{aligned}
& e^{-r\hat{\tau}^{\mathrm{D}}} \int_{\mathcal{D}_\infty} V\left(y, T_{i+1}^{\mathrm{O}}\right) f_1(y|S)\mathrm{d}y \\
\approx \; & e^{-r\hat{\tau}^{\mathrm{D}}} \int_{\mathcal{D}_\infty} \left(W\left(y, T_{i+1}^{\mathrm{O}}\right) + \chi_{i+1}\right) f_1(y|S)\mathrm{d}y \\
\approx \; & e^{-r\hat{\tau}^{\mathrm{D}}} \int_{\mathcal{D}_\infty} \left(\phi e^{-\delta\hat{\tau}_{i+1}^{\mathrm{O}}} y - \phi e^{-r\hat{\tau}_{i+1}^{\mathrm{O}}} K + \chi_{i+1}\right) f_1(y|S)\mathrm{d}y.
\end{aligned}
$$

The first step corresponds to approximating the value of the multi-window barrier option with that of a continuously monitored barrier option plus some constant χ_{i+1}. We comment on possible choices for χ_{i+1} below. The second step uses the asymptotic deep in-the-money value derived in Lemma 2.2. Consequently, the value of the second term is given by

$$
\begin{aligned}
\dots \quad = \quad & \phi e^{-\delta\hat{\tau}^{\mathrm{O}}_{i+1}} \left(\mathcal{A}^{\phi\phi}_{BS_m}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) - \overset{B}{\mathcal{I}} \left\{ \mathcal{A}^{\phi\phi}_{BS_m}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) \right\} \right) \\
& + \left(\chi_{i+1} - \phi e^{-r\hat{\tau}^{\mathrm{O}}_{i+1}} K \right) \left(\mathcal{B}^{\phi\phi}_{BS_m}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) - \overset{B}{\mathcal{I}} \left\{ \mathcal{B}^{\phi\phi}_{BS_m}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) \right\} \right) \\
=: \quad & w^{\mathrm{A}}_{i+1} A(S) + w^{\mathrm{B}}_{i+1} B(S),
\end{aligned}
$$

where

$$
w^{\mathrm{A}}_{i+1} = \phi e^{-\delta\hat{\tau}^{\mathrm{O}}_{i+1}}, \qquad w^{\mathrm{B}}_{i+1} = \chi_{i+1} - \phi e^{-r\hat{\tau}^{\mathrm{O}}_{i+1}} K.
$$

Here, the definitions of $A(S)$ and $B(S)$ are implicitly clear. The constant χ_{i+1} controls how the gap risk premium is extrapolated for deep in-the-money scenarios. For $\chi_{i+1} = 0$ we assume that the gap risk premium is zero beyond the maximum moneyness. Alternatively, we can set $\chi_{i+1} = V\left(S_m, T^{\mathrm{O}}_{i+1}\right) - W\left(S_m, T^{\mathrm{O}}_{i+1}\right)$ such that the gap risk premium stays constant beyond the maximum moneyness. While the former option tends to under-price the gap risk, the latter over-prices it. The degree of under- or over-pricing is increasing with the time-to-maturity. One could conceive more elaborate schemes, for example with an exponentially decaying gap risk premium. However, through numerical experiments, we find that choosing a sufficiently high value for $\kappa_{\max}$ and setting $\chi_{i+1} = 0$ works well.

3.5 Matrix Representation and Control Variate Stabilization

Let $\boldsymbol{V}\left(T^{\mathrm{O}}_i\right)$ be the $m \times 1$ column vector with elements $V\left(S_j, T^{\mathrm{O}}_i\right)$. We now formulate the backward induction algorithm defined above in matrix notation to obtain a recursive expression for $\boldsymbol{V}\left(T^{\mathrm{O}}_i\right)$ in terms of $\boldsymbol{V}\left(T^{\mathrm{O}}_{i+1}\right)$. Let $\boldsymbol{P}$ be the $m \times m$ backward induction matrix with elements

$$
P_{j,k} = e^{-r\hat{\tau}^{\mathrm{D}}} w_k f_1\left(S_k | S_j\right).
$$

Further, define $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{R}^{\mathrm{A}}$ and $\boldsymbol{R}^{\mathrm{B}}$ to be the $m \times 1$ column vectors with elements $A\left(S_j\right)$, $B\left(S_j\right)$, $R^{\mathrm{A}}\left(S_j\right)$ and $R^{\mathrm{B}}\left(S_j\right)$, respectively. Then

$$
\boldsymbol{V}\left(T^{\mathrm{O}}_i\right) \approx \boldsymbol{P} \cdot \boldsymbol{V}\left(T^{\mathrm{O}}_{i+1}\right) + w^{A}_{i+1}\boldsymbol{A} + w^{B}_{i+1}\boldsymbol{B} + \boldsymbol{R}^{\mathrm{A}} + \boldsymbol{R}^{\mathrm{B}}.
$$

We remark that the vector $\boldsymbol{V}\left(T^{\mathrm{O}}_1\right)$ of leveraged certificate prices on the market opening of the next trading day needs to be only computed once during any trading day and can then be cached. This requires that the other market parameters such as risk-free interest rate, dividend yield and volatility stay unchanged.

To further stabilize the numerical scheme, we recommend to employ the price of an early ending barrier option as a control variate. The latter has a plain vanilla terminal payoff at the market opening T_{i+1}^{O} of the following trading day if the barrier B is not breached during the intraday period $\left[T_i^{\mathrm{O}}, T_i^{\mathrm{C}}\right)$. Its time T_i^{O} price is given by

$$E(S) = \phi \left(\mathcal{Q}_{BK}^{\phi\phi}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) - \overset{B}{\mathcal{I}} \left\{ \mathcal{Q}_{BK}^{\phi\phi}\left(S, \hat{\tau}^{\mathrm{I}}, \hat{\tau}^{\mathrm{D}}\right) \right\} \right).$$

Define $\boldsymbol{E}$ to be the $m \times 1$ column vector with elements $E\left(S_j\right)$. Its numerical approximation is given by

$$\boldsymbol{E} \approx \boldsymbol{P} \cdot \boldsymbol{D} + \phi\left(\boldsymbol{A} - K\boldsymbol{B}\right),$$

where the elements of the $m \times 1$ column vector $\boldsymbol{D}$ correspond to the terminal payoff of the early ending barrier option, that is $D\left(S_j\right) = \left(\phi S_j - \phi K\right)^{+}$. We define the $m \times 1$ column vector of correction terms $\boldsymbol{X}$ element-wise as the ratio of the analytical and numerical prices of the control variate. The control variate backward induction scheme is then given by

$$\boldsymbol{V}\left(T_i^{\mathrm{O}}\right) \approx \left(\boldsymbol{P} \cdot \boldsymbol{V}\left(T_{i+1}^{\mathrm{O}}\right)\right) \circ \boldsymbol{X} + w_{i+1}^{A}\boldsymbol{A} + w_{i+1}^{B}\boldsymbol{B} + \boldsymbol{R}^{\mathrm{A}} + \boldsymbol{R}^{\mathrm{B}},$$

where $\circ$ denotes the element-wise Schur product.

3.6 Final Step

The algorithm discussed so far yields a price grid $\boldsymbol{V}\left(T_1^{\mathrm{O}}\right)$ corresponding to the market opening of the following trading day. We could again adopt the above approach of numerically integrating over the prices $\boldsymbol{V}\left(T_1^{\mathrm{O}}\right)$ multiplied by the (conditional) transition probability function. This is prone to instabilities in $V(S,t)$ and its spatial derivatives as t approaches T_1^{O} unless a very high number of grid points m is chosen. Doubling m however nearly quadruples the computational complexity as the dominating factor is the construction of the backward induction matrix $\boldsymbol{P}$.

Instead, we suggest an alternative scheme that yields very stable values and spatial derivatives while only marginally increasing the computational complexity. Our approach is independent of the particular valuation problem analyzed in this paper and is generally applicable to grid-based numerical methods.

Assumption 3.1 (Polynomial Approximation).
Assume that the valuation function $V\left(S, T_1^{\mathrm{O}}\right)$ within the survival domain can be well approximated by a piecewise polynomial in the underlying asset price, that is

$$V\left(S, T_1^{\mathrm{O}}\right) \approx \sum_{j=0}^{m-1}\left(\sum_{k=0}^{d} \alpha_{j,k} S^k\right) 1\left\{S \in [S_j, S_{j+1})\right\} + \left(\sum_{k=0}^{d} \alpha_{m,k} S^k\right) 1\left\{\phi S \geq \phi S_m\right\},$$

where $d \in \mathbb{N}_0$ is the polynomial degree.

Potential candidates are, for example, a piecewise linear interpolator or cubic (smoothing) splines. When evaluating the above expression, we need to distinguish two cases depending on the current market phase.

Lemma 3.4 (Continuation Value).
The continuation value for the final step is given by

$$C(S,t) = \sum_{j=0}^{m}\sum_{k=0}^{d} \beta_{j,k} \begin{cases} {}^{k}\mathcal{P}_{S_j}^{\phi}\left(S, \tau_1^{O}\right) & \text{if } t \in \left[T_0^{C}, T_1^{O}\right) \\ \left({}^{k}\mathcal{P}_{BS_j}^{\phi\phi}\left(S, \tau_0^{C}, \tau_1^{O}\right) - \overset{B}{\mathcal{I}}\left\{{}^{k}\mathcal{P}_{BS_j}^{\phi\phi}\left(S, \tau_0^{C}, \tau_1^{O}\right)\right\}\right) & \text{otherwise} \end{cases}.$$

where

$$\beta_{j,k} = \begin{cases} \alpha_{0,k} & \text{if } j = 0 \\ \alpha_{j,k} - \alpha_{j-1,k} & \text{otherwise} \end{cases}.$$

Proof When $t \in \left[T_0^{\mathrm{C}}, T_1^{\mathrm{O}}\right)$, then the market is already closed and the barrier is not monitored until the market opening of the next trading day. The continuation value is given by

$$\begin{aligned} C(S,t) &= e^{-r\tau_1^{\mathrm{O}}} \int_{\mathcal{D}} V\left(S, T_1^{\mathrm{O}}\right) g(y|S)\mathrm{d}y \\ &= e^{-r\tau_1^{\mathrm{O}}} \left(\sum_{j=0}^{m-1} \int_{S_j}^{S_{j+1}} \left(\sum_{k=0}^{d} \alpha_{j,k} y^k\right) g(y|S)\mathrm{d}y \right. \\ &\qquad \left. + \phi \int_{S_m}^{\phi\infty} \left(\sum_{k=0}^{d} \alpha_{m,k} y^k\right) g(y|S)\mathrm{d}y \right). \end{aligned}$$

where

$$g(y|x)\mathrm{d}y = \mathbb{P}^*\left\{ S_{T_1}^{\mathrm{O}} \in \mathrm{d}y \middle| S_t = x \right\}$$

is the transition probability function. Each of the inner integrals evaluates to

$$e^{-r\tau_1^{\mathrm{O}}} \int_{S_j}^{S_{j+1}} \left(\sum_{k=0}^{d} \alpha_{j,k} y^k \right) g(y|S)\mathrm{d}y = \sum_{k=0}^{d} \alpha_{j,k} \left({}^k\mathcal{P}_{S_j}^{\phi}\left(S, \tau_1^{\mathrm{O}}\right) - {}^k\mathcal{P}_{S_{j+1}}^{\phi}\left(S, \tau_1^{\mathrm{O}}\right) \right).$$

The weights $\beta_{j,k}$ on the power binary with strike S_j and exponent k are found by collecting terms. The proof for the case when $t \in \left[T_0^{\mathrm{O}}, T_0^{\mathrm{C}}\right)$ is fully analogous. □

The delta and vega can be computed analytically from this piecewise polynomial approximation using Propositions B.2 and B.3. We can use the same approach to obtain closed-form expressions for other Greeks.

4 Accuracy and Performance

We validate the numerical approximation by comparing its prices to those obtained through Monte Carlo simulations. The common contractual terms are a strike and barrier at $K = B = 100.00$ with the maturity on January 1, 2015. We value leveraged certificates without a rebate as these have the highest exposure to the overnight gap risk. The market opens at 9:00 AM and closes at 5:30 PM. The intraday and overnight periods have an equal time weight of 50%. We generate 10,000 random uniformly distributed samples for the remaining pricing relevant parameters. This ensures that a diverse set of test cases is employed. The pricing instant lies between midnight of June 1, 2014 and midnight of January 1, 2015. The logarithmic moneyness $\phi \ln(S_0/B)$ is in the interval $(0\%, 10\%]$ when the market is open and in $(-5\%, 10\%]$ when the market is closed. The risk-free interest rate and dividend yield are in the interval $[0\%, 10\%]$ and the volatility is sampled from $(0\%, 100\%]$. Bull and bear certificates are generated with equal probability. Test cases are only accepted if the unconditional probability of breaching the continuous barrier is in the interval $[1\%, 99\%]$. Each Monte Carlo simulation is based on 1,000,000 total paths including antithetics. In this and the next section, results for the grid-based scheme are based on $m = 100$ grid points, a scaling factor of $\xi = 1.5$, a maximum moneyness of $\kappa_{\max} = \sigma$ and piecewise linear interpolation for the final step, unless indicated otherwise.

Figure 3 shows the distribution of the price differences between the Monte Carlo simulation and the grid-based scheme, normalized to the respective number of Monte Carlo standard errors. The corresponding descriptive statistics are given in Table 1. Visually, the distribution appears fairly normal, except for a sharp peak around zero. The latter can be attributed to samples that have a very low hitting probability and consequently carry a negligible gap risk premium. Once these observations are removed, the resulting histogram resembles an almost perfect bell curve. This confirms the robustness of the grid-based scheme since the Monte Carlo estimates themselves have normally distributed errors.

Insert Figure 3 about here.

Insert Table 1 about here.

The left pane of Figure 4 illustrates the convergence of the numerical algorithm for different numbers of spatial grid points. We observe that already 48 points yield an average absolute error in the order of magnitude of only 0.01% of the underlying asset price. Using the control variate significantly improves the accuracy when a very small grid is used and there is no noticeable difference for 100 points or more. For $m \geq 20$, we observe that the absolute error roughly decreases by a factor of four whenever we double the number of grid points, thus indicating a quadratic convergence rate. This is confirmed by the black dash-dotted line in the left pane of Figure 4, which corresponds to the function $1/m^2$. When the leveraged certificate is deep in-the-money, the convergence rate is roughly cubic.

Insert Figure 4 about here.

The low absolute error for a relatively small number of spatial grid points is the main advantage of the proposed mixed scheme over standard finite difference methods. Since the transition probabilities from the market opening of one trading day to the next are computed in closed-form, the only source of discretization errors is the coarseness of the spatial grid. In contrast, finite difference schemes additionally demand for a very fine time grid to properly capture the discontinuity of the barrier. This significantly increases the number of matrix multiplications in the backward induction which dominates the computation time for a relatively long time-to-maturity and/or a large spatial grid.

The right pane of Figure 4 shows the convergence as a function of the computation time. The measurements are taken for a single threaded implementation on an Intel Core i7-4870HQ CPU clocked at 2.50 GHz. The algorithm is generally very fast, with 96 grid points taking on average 1.78 (1.66) ms with (without) the control variate correction for a leveraged certificate with three months to maturity. The black dash-dotted line in the right pane of Figure 4 corresponds to the function $10^{-4}/t$, thus indicating a linear convergence rate. This is in accordance with our previous analysis given that the computational complexity of constructing the backward induction matrix is $\mathcal{O}\left(m^2\right)$.

As remarked in Section 3.5, the vector $\boldsymbol{V}\left(T_1^{\mathrm{O}}\right)$ can often be cached between two calls to the valuation function as long as only the asset price and/or time change. The average cached pricing time for the above example is 0.20 ms.

5 Dynamics of the Gap Risk Premium

We define the gap risk premium in leveraged certificates to be the difference between the prices of the otherwise identical multi-window and continuous barrier options. For contracts without a rebate, it can be shown to be always non-negative. In this section, we discuss the dynamics of the gap risk premium as a function of its major driving factors. Throughout, we consider a bull leveraged certificate with a strike and barrier at $K = B = 100.00$. The market setup is as in Section 4. The diffusion coefficient is $\sigma = 20\%$ and we assume a zero risk-free interest rate and dividend yield $r = \delta = 0\%$.

To gain some intuition for the gap risk premium, we first consider the special case of a leveraged certificate with a single overnight period to maturity. Figure 5 plots its value as a function of the asset price and for different time points during $\left[T_{n-1}^{\mathrm{O}}, T_n^{\mathrm{O}}\right)$. Since the risk-free interest rate and dividend yield are zero, the value of the leveraged certificate during the market opening hours of the expiry day, and in particular at the market opening T_n^{O}, is equal to its intrinsic value. Consequently, its value during the overnight period $\left[T_{n-1}^{\mathrm{C}}, T_n^{\mathrm{O}}\right)$ is equal to that of an otherwise identical plain vanilla option maturing at T_n^{O}. The gap risk premium corresponds to the time value of this plain vanilla option; see the curves corresponding to 8:00 PM of T_{n-1} and 8:00 AM of the maturity day T_n. For any fixed time $t \in \left[T_{n-1}^{\mathrm{C}}, T_n^{\mathrm{O}}\right)$, the gap risk premium is the highest for asset prices equal to the barrier. For any fixed asset price, it is increasing in the time-to-maturity $T_n^{\mathrm{O}} - t$.

Insert Figure 5 about here.

During the market opening hours $\left[T_{n-1}^{\mathrm{O}}, T_{n-1}^{\mathrm{C}}\right)$, the leveraged certificate can be valued as an early ending barrier option with barrier end in T_{n-1}^{C} and maturing at T_n^{O}. The gap risk premium is zero for asset prices at or below the barrier. Above the barrier, it is first increasing and then converges to zero in the limit. As the trading day progresses, the maximum gap risk premium increases and the asset price in which the maximum is attained shifts further towards the barrier. This becomes clear when interpreting the early ending barrier option as a compound option, i.e., a knock-out barrier option whose payoff is at T_{n-1}^{C} is a plain vanilla option with maturity in T_n^{O}. Two opposing forces determine its value as the asset price approaches the barrier. On the one hand, the probability of breaching the barrier during the current trading day increases, thus lowering the price. On the other hand, the expected time value of the vanilla option payoff at T_{n-1}^{C} also increases, conditional on no prior knock-out. For asset prices that are close to (far away from) the barrier, relative to the remaining time during which it is active, the survival probability increases faster (slower) than the expected time value decreases. Thus, the gap risk premium increases (decreases).

Figure 6 shows the corresponding results for a time-to-maturity of three months, keeping all other parameters as in Figure 5. The gap risk premium stays nearly unchanged during the overnight period and for asset prices below the barrier. In this region, it is highly likely that the leveraged certificate is knocked-out either immediately on the market opening T_1^{O} or during the following intraday period $\left[T_1^{\mathrm{O}}, T_1^{\mathrm{C}}\right)$. Consequently, the gap risk premium is mainly driven by the current overnight window, just as in the case with a single overnight period to maturity. The gap risk premium in the in-the-money region is significantly higher and decays much slower as the asset price increases. This is due to the increased number of remaining overnight windows. Conditional on no prior knock-out, during each additional day-to-maturity, there is a chance that the asset price breaches the then-inactive barrier during an overnight window, thus increasing the value of the embedded gap risk premium. This also explains why time variability of the gap risk premium is lower when the leveraged certificate is deep in-the-money. In this region, the probability of a gap event during the next overnight window is very small and the premium mainly stems from paths that cross the barrier during one of the later overnight windows.

Insert Figure 6 about here.

Figure 6 also confirms that the gap risk premium represents a substantial part of the overall value of a leveraged certificate. At an asset price of USD 100.50, for example, the value of the continuous barrier option is constant at USD 0.50 while that of the multi-window barrier option varies between USD 0.56 and USD 0.72, depending on the time-of-day. The highest impact can be observed at 5:30 PM and when the spot is equal to the barrier. The corresponding gap risk premium is USD 0.38 with an unadjusted price of USD 0.00.

Insert Figure 7 about here.

Accounting for the time-discontinuity of the barrier not only yields a gap risk adjusted value but also induces different sensitivities for dynamic hedging strategies. We consider again the previous example of an asset price of USD 100.50. The delta of the continuous barrier option is equal to 100% at all times, while that of the multi window barrier option varies between 53% and 129%. As confirmed by Figure 7, the holder of a long position in the leveraged certificate is generally gamma long when the market is closed or when the asset price is sufficiently far away from the barrier. She faces a short gamma singularity as the asset price approaches the barrier and when the time approaches the market closing.

Insert Figure 8 about here.

While we model volatility to be fixed or at most time-dependent for tractability reasons, real-world implied volatilities are not constant. Figure 8 shows that the value of the gap risk premium is generally increasing in volatility. During the overnight period, the gap risk vega is highest for asset prices near the barrier. Its intraday dynamics exhibit a similar pattern as that of the gap risk premium itself. For asset prices close to the barrier, a higher volatility increases the potential overnight gap but at the same time makes a prior knock-out more likely.

6 Generalizations and Extensions

6.1 Market Clock and Intraday Events

The rate at which uncertainty is resolved in the market is generally neither constant nor deterministic. Extending prior work by Clark (1973), Ane and Geman (2000) account for this by modeling the asset price to follow a time changed geometric Brownian. Using high-frequency transaction data, they estimate the empirical timescale under which short-term returns are approximately normally distributed. We follow a similar approach and incorporate the average rate of information flow within each trading day through a deterministic time change.

Let $\theta : \mathbb{R}_{+,0} \to \mathbb{R}_{+,0}$ be a non-decreasing and continuous function that maps the physical time into the so-called market time. Its first derivative can be interpreted as the speed at which the market time evolves relative to the physical time. In general, this approach allows us to choose market time durations $\theta(t_2) - \theta(t_1)$ for any physical time interval $[t_1, t_2]$ such as to be proportional to the corresponding total variance of the underlying asset. We employ the time change approach in Section 5, to assign weights to the intraday and overnight periods that differ from their relative lengths. Similarly, it can be used to differentiate the overnight gap between two consecutive trading days from weekends and exchange holidays.

Insert Figure 9 about here.

Figure 9 shows two different estimates for the intraday clock of the DAX futures contract. It compares them to the case where time passes at a constant rate, or equivalently, total variance accumulates linearly. Based on all trades in 2014 and 2015, we first aggregate the data to five minute intervals. We then estimate the relative weight of each interval as its average share of the total daily variance. Approximately 14% (33%) of the total daily variance can be attributed to the period where the futures (spot) market is closed overnight. We observe that the intraday clock runs at an accelerated speed after the German spot market opening around 9:00 AM as well during the US spot market opening around 3:30 PM. It moves slower outside the German spot market trading times as well as around lunch time. The intraday clock between 8:00 AM and 22:00 PM can alternatively be estimated based on the transaction volume or the number of trades. We find that these two measures yield almost identical results and Figure 9 depicts the results for the volume-based market clock.

The time change approach can be further extended to handle predictable jump events during the market opening times. These can be modeled as very short physical time periods during which the barrier is inactive and market time passes at a very high rate. Monetary policy announcements, for example, are commonly followed by a very short period of extremely high volatility and relatively low liquidity. Stop loss orders placed in the market beforehand are likely to be filled at levels significantly beyond the barrier, should they be triggered.

6.2 Discrete Barriers

Our analysis thus far covers leveraged certificates with a constant barrier. A common variation has a discrete barrier which is strictly in-the-money and active only upon the set of market closing times

$$\mathcal{T}^{\mathrm{C}} = \bigcup_{i=0}^{\infty} T_i^{\mathrm{C}}.$$

In order to limit the issuer's gap risk exposure, these products additionally have a continuously monitored barrier equal to the strike which is only active during the intraday windows. The relevant hitting time is thus given by

$$\nu \quad = \quad \inf \left\{ t \in \mathcal{T}^{\mathrm{O}} : \phi S_t \leq \phi K, \; t \in \mathcal{T}^{\mathrm{C}} : \phi S_t \leq \phi B \right\}.$$

The previously defined payoff function still applies to discrete barriers. We can continue to employ the numerical algorithm in Section 3 with adjusted expressions for the initial step, transition probability function, rebate and final step. All details can be found in Appendix D.

Figure 10 depicts the corresponding gap risk premium as a function of the asset price. The in-the-money gap risk premium is generally lower than that of the otherwise identical leveraged certificate with a continuous barrier equal to the strike. The contract design reduces the likelihood of a path that crosses the barrier during an overnight window, since the asset price first has to close strictly above the barrier. Once this scenario occurs, however, the gap risk premium is dominated by the imminent knock-out upon the market opening against the strike level. It then closely resembles that of a standard leveraged certificate.

Insert Figure 10 about here.

7 Conclusion

We show that the overnight gap risk premium typically constitutes a non-negligible fair price component of leveraged certificates. The drastic computational demand driven by the large number of such barrier options in the market and their high price sensitivity necessitates and efficient valuation algorithm. Hence, standard approaches based on Monte Carlo simulations are infeasible. For this purpose, our paper develops a bespoke grid-based backward induction scheme that is both highly accurate and extremely fast. It can be extended to accommodate general predictable jump discontinuities as well as discrete barriers. Moreover, many insights carry over to the valuation of other complex barrier options.

References

Ane, Thierry and Hélyette Geman (2000) "Order Flow, Transaction Clock, and Normality of Asset Returns," *Journal of Finance*, Vol. 55, No. 5, pp. 2259–2284.

Baller, Stefanie, Oliver Entrop, Michael McKenzie, and Marco Wilkens (2016) "Market

Makers' Optimal Price-Setting Policy for Exchange-Traded Certificates," *Journal of Banking & Finance*, Vol. 71, No. 10, pp. 206–226.

Black, Fischer and Myron S. Scholes (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, Vol. 81, No. 3, pp. 637–654.

Boes, Mark-Jan, Feike C. Drost, and Bas J. M. Werker (2007) "The Impact of Overnight Periods on Option Pricing," *Journal of Financial and Quantitative Analysis*, Vol. 42, No. 2, pp. 517–534.

Buchen, Peter W. (2001) "Image Options and the Road to Barriers," *Risk Magazine*, Vol. 14, No. 9, pp. 127–130.

——— (2004) "The Pricing of Dual-Expiry Exotics," *Quantitative Finance*, Vol. 4, No. 1, pp. 101–108.

Clark, Peter K. (1973) "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, Vol. 41, No. 1, pp. 135–155.

Deutscher Derivate Verband (2015) "Marktvolumen von Derivativen Wertpapieren," available at http://www.deutscherderivateverband.de.

Entrop, Oliver, Hendrik Scholz, and Marco Wilkens (2009) "The Price-Settings Behavior of Banks: An Analysis of Open-End Leverage Certificates on the German Market," *Journal of Banking & Finance*, Vol. 33, No. 5, pp. 874–882.

Greene, William H. (2008) *Econometric Analysis*: Pearson Prentice Hall, 6th edition.

Ingersoll, Jonathan E. Jr. (2000) "Digital Contracts: Simple Tools for Pricing Complex Derivatives," *Journal of Business*, Vol. 73, No. 1, pp. 67–88.

Muck, Matthias (2006) "Where Should You Buy Your Options? The Pricing of Exchange-Traded Certificates and OTC Derivatives in Germany," *Journal of Derivatives*, Vol. 14, No. 1, pp. 82–96.

Rossetto, Silvia and Jos van Bommel (2009) "Endless Leverage Certificates," *Journal of Banking & Finance*, Vol. 33, No. 8, pp. 1543–1553.

Skipper, Max and Peter W. Buchen (2003) "The Quintessential Option Pricing Formula," Working Paper, School of Mathematics and Statistics, University of Sydney.

Wilkens, Sascha and Pavel A. Stoimenov (2007) "The Pricing of Leverage Products: An Empirical Investigation of the German Market for 'Long' and 'Short' Stock Index Certificates," *Journal of Banking & Finance*, Vol. 31, No. 3, pp. 735–750.

A Method of Images

The method of images for the heat transfer equation provides a simple solution to many initial boundary value problems. Analogously to the reflection principle for the paths of Brownian motions, it exploits symmetry relationships in the Kolmogorov forward equation satisfied by its transition PDF. The central idea is to extend a problem defined on a semi-infinite domain into a related problem on an infinite domain. The latter can then be solved though standard techniques by a convolution of its initial condition with the corresponding Green's function. The infinite domain problem is constructed such as to satisfy the same initial condition on the original domain and for its solution to match the prescribed value on the boundary. The method of images often significantly simplifies the solution of boundary value problems that are hard to handle through probabilistic methods. This is in particular true when dealing with high-dimensional problems such as the multiple barrier windows in this paper.

Here, we re-state without proof the necessary results needed in this paper. We refer to Buchen (2001) for an introduction in the context of standard barrier options. Buchen (2004) applies the method of images to the pricing of dual expiry exotics. We start by giving three central definitions that are based on Buchen (2001).

Definition A.1 (Black and Scholes (1973) Forward Operator).
The one-dimensional Black and Scholes (1973) forward operator $\mathcal{L}$ is defined as

$$\mathcal{L}\{\tilde{V}\} = -\frac{\partial \tilde{V}}{\partial \tau} + (r-\delta)S\frac{\partial \tilde{V}}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 \tilde{V}}{\partial S^2} - r\tilde{V}. \qquad \triangle$$

Definition A.2 (Black and Scholes (1973) Initial Boundary Value Problem).
The function $\tilde{V}(S,\tau)$ satisfies an IBVP with boundary level $B \in \mathbb{R}_+$ and initial condition $f : \mathbb{R}_+ \to \mathbb{R}$ if

$$\begin{aligned} \mathcal{L}\left\{\tilde{V}\right\}(S,\tau) &= 0 \qquad \text{for } (S,\tau) \in \mathcal{D} \times \mathbb{R}_+, \\ \tilde{V}(S,0) &= f(S), \\ \tilde{V}(B,\tau) &= 0 \qquad \text{for } \tau \in \mathbb{R}_{+,0}, \end{aligned}$$

where the active domain $\mathcal{D} \subseteq \mathbb{R}_+$ is given by

$$\mathcal{D} = \{S \in \mathbb{R}_+ : \psi S > \psi B\}$$

and $\psi \in \{-1, +1\}$ indicates an upper or lower boundary. $\triangle$

Definition A.3 (Image Operator).
Let $\tilde{V}(S, \tau)$ be a solution to the Black and Scholes (1973) PDE. The image of $\tilde{V}(S, \tau)$ relative to the barrier $B \in \mathbb{R}_+$ is given by

$$\overset{B}{\mathcal{I}} \left\{ \tilde{V}(S, \tau) \right\} = \left(\frac{S}{B} \right)^{2\alpha} \tilde{V} \left(\frac{B^2}{S}, \tau \right),$$

where

$$\alpha = \frac{1}{2} - \frac{r - \delta}{\sigma^2}. \qquad \triangle$$

We refer to Buchen (2001) for the general properties of the image operator. The following proposition provides the main result for the valuation of barrier options within the Black and Scholes (1973) framework.

Proposition A.1 (Method of Images for the Black and Scholes (1973) PDE).
Let $\tilde{V}(S, \tau)$ satisfy a Black and Scholes (1973) IBVP. Define $\tilde{V}_B(S, \tau)$ to be the solution to the corresponding full-range problem

$$\begin{aligned} \mathcal{L}\left\{ \tilde{V}_B \right\}(S, \tau) &= 0 \qquad \text{for } (S, \tau) \in \mathbb{R}_+^2, \\ \tilde{V}_B(S, 0) &= f(x) 1\{\psi S > \psi B\}. \end{aligned}$$

Then $\tilde{V}(S, \tau)$ is given by

$$\tilde{V}(S, \tau) = \tilde{V}_B(S, \tau) - \overset{B}{\mathcal{I}} \left\{ \tilde{V}_B(S, \tau) \right\}.$$

The main steps in the proof of this result are to first apply a coordinate transformation that reduces the Black and Scholes (1973) PDE to the one-dimensional heat transfer equation such that the transformed boundary is still a constant. The initial boundary value problem for the heat transfer equation can then be easily solved using the method of images. Reversing the change of variables in the solution yields the result in Proposition A.1.

The following three Lemmata provide the major Greeks for images of valuation functions. These results all follow from a careful differentiation and the proofs are omitted for brevity.

Lemma A.1 (Sensitivities of the Image of a Valuation Function).
Let $\tilde{V}(S,\tau)$ be some valuation function. The major partial derivatives of the image of $\tilde{V}(S,\tau)$ are given by

$$\begin{aligned}
\frac{\partial}{\partial S}\overset{B}{\mathcal{I}}\left\{\tilde{V}(S,\tau)\right\} &= \left(\frac{S}{B}\right)^{2\alpha}\left(\frac{2\alpha}{S}\tilde{V}\left(\frac{B^2}{S},\tau\right)-\frac{B^2}{S^2}\frac{\partial\tilde{V}}{\partial S}\left(\frac{B^2}{S},\tau\right)\right)\\
\frac{\partial^2}{\partial S^2}\overset{B}{\mathcal{I}}\left\{\tilde{V}(S,\tau)\right\} &= \left(\frac{S}{B}\right)^{2\alpha}\left(\frac{2\alpha(2\alpha-1)}{S^2}\tilde{V}\left(\frac{B^2}{S},\tau\right)\right.\\
&\quad\left.+\frac{2(1-2\alpha)B^2}{S^3}\frac{\partial\tilde{V}}{\partial S}\left(\frac{B^2}{S},\tau\right)+\frac{B^4}{S^4}\frac{\partial^2\tilde{V}}{\partial S^2}\left(\frac{B^2}{S},\tau\right)\right),\\
\frac{\partial}{\partial \tau}\overset{B}{\mathcal{I}}\left\{\tilde{V}(S,\tau)\right\} &= \left(\frac{S}{B}\right)^{2\alpha}\frac{\partial\tilde{V}}{\partial\tau}\left(\frac{B^2}{S},\tau\right),\\
\frac{\partial}{\partial \sigma}\overset{B}{\mathcal{I}}\left\{\tilde{V}(S,\tau)\right\} &= \left(\frac{S}{B}\right)^{2\alpha}\left(\frac{2(1-2\alpha)}{\sigma}\ln\left(\frac{S}{B}\right)\frac{\partial\tilde{V}}{\partial\sigma}\left(\frac{B^2}{S},\tau\right)\right).
\end{aligned}$$

B Higher-Order Binary and $\mathcal{Q}$ Options

As shown in Buchen (2001), the method of images yields a solution for the prices of barrier options in terms of a linear combination of cash-or-nothing (bond) and asset-or-nothing (asset) binaries and their respective images. Similarly, Ingersoll (2000) demonstrates that a wide variety of exotic payoffs can be priced in terms of elementary digital options. Skipper and Buchen (2003) generalize much of the previous literature on digital option pricing by introducing the so-called $\mathbb{M}$-binary. These multi-asset and multi-period exotic binaries can be used to price most rainbow options, whose terminal payoff depends on the asset prices at a discrete set of monitoring points.

The basic building blocks for the prices of the barrier options considered in this paper are higher-order power binaries. Our notation is very similar to the one used by Skipper and Buchen (2003).

Definition B.1 (n-th Order Power Binaries).
Let $\boldsymbol{s} = (s_1, s_2, \ldots, s_n)'$ be an n-dimensional vector of indicators with values $s_i \in \{-1, +1\}$, let $\boldsymbol{T} = (T_1, T_2, \ldots, T_n)' \in \mathbb{R}^n_+$ be an n-dimensional vector of maturity dates such that $T_i < T_j$ for $i < j$ and let $\boldsymbol{\xi} = (\xi_1, \xi_2, \ldots, \xi_n)' \in \mathbb{R}^n_+$ be an n-dimensional vector of strike prices. The n-th order power binary ${}^{\eta}\mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}$ with exponent $\eta \in \mathbb{R}$ has the time T_n payoff

$$ {}^{\eta}\mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}} \left(\boldsymbol{S_T}, 0 \right) = S^{\eta}_{T_n} 1_n \left\{ \mathrm{diag}(\boldsymbol{s}) \boldsymbol{S_T} > \mathrm{diag}(\boldsymbol{s}) \boldsymbol{\xi} \right\}. $$

Here, $\boldsymbol{S_T} = (S_{T_1}, S_{T_2}, \ldots, S_{T_n})' \in \mathbb{R}^n_+$ is the n-dimensional vector of underlying asset prices at the maturity dates and $\mathrm{diag}(\boldsymbol{s})$ is the $n \times n$ diagonal matrix created from the vector $\boldsymbol{s}$. The n-dimensional indicator function 1_n is defined element-wise as

$$ 1_n \left\{ \boldsymbol{x} > \boldsymbol{a} \right\} = \prod_{i=1}^{n} 1 \left\{ x_i > a_i \right\}. \qquad \triangle $$

Thus, an n-th order power binary has a payoff equal to the η-th power of the asset price at the terminal maturity date T_n, if and only if $s_i S_{T_i} > s_i \xi_i$ for all $i \in \{1, 2, \ldots, n\}$. Again, each $s_i \in \{-1, +1\}$ serves as an indicator for a less-than or greater-than inequality, respectively. Higher-order binaries can also be regarded as compound options. A second-order binary, for example, is a binary option to receive a first-order binary at time T_1 given that $s_1 S_{T_1} > s_1 \xi_1$. In general, for $n > 1$ we have at the first maturity date

$$ {}^{\eta}\mathcal{P}^{s_1 s_2 \ldots s_n}_{\xi_1 \xi_2 \ldots \xi_n} \left(S, 0, \tau_2, \ldots, \tau_n \right) = {}^{\eta}\mathcal{P}^{s_2 s_3 \ldots s_n}_{\xi_2 \xi_3 \ldots \xi_n} \left(S, \tau_2, \tau_3, \ldots, \tau_n \right) 1 \left\{ s_1 S > s_1 \xi_1 \right\}. $$

To keep the notation less verbose, we usually suppress the explicit dependence of the power binary price on all relevant times to maturity but instead simply write ${}^{\eta}\mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(S, \tau)$ in what follows.

Proposition B.1 (Valuation of n-th Order Power Binaries).
The time $0 \le t \le T_1$ value of the n-th order power binary is given by

$$ {}^{\eta}\mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(S, \tau) = F_{\eta} \mathcal{N}_n \left(\boldsymbol{\alpha}_{\eta}; \boldsymbol{C} \right), $$

where

$$ F_{\eta} = S^{\eta} \exp \left\{ \left((\eta - 1) \left(r + \frac{1}{2} \eta \sigma^2 \right) - \eta \delta \right) \tau_n \right\}. $$

$\mathcal{N}_n(\boldsymbol{x};\boldsymbol{C})$ is the n-variate standard normal CDF evaluated at n-dimensional vector $\boldsymbol{x}$ and with correlation matrix $\boldsymbol{C}$. $\boldsymbol{\alpha}_\eta = \mathrm{diag}(\boldsymbol{s})\boldsymbol{d}_\eta$ and the n-dimensional vector $\boldsymbol{d}_\eta = (d_{\eta,1}, d_{\eta,2}, \ldots, d_{\eta,n})' \in \mathbb{R}^n$ is defined as

$$d_{\eta,i} = \frac{\ln(S/\xi_i) + \left(r - \delta + \left(\eta - \frac{1}{2}\right)\sigma^2\right)\tau_i}{\sigma\sqrt{\tau_i}},$$

where $\tau_i = T_i - t$ is the i-th time-to-maturity and $\boldsymbol{C} \in \mathbb{R}^n \times \mathbb{R}^n$ is a symmetric positive definite correlation matrix given by

$$\boldsymbol{C}_{i,j} = \begin{cases} 1 & \text{if } i = j \\ s_i s_j \sqrt{\tau_i/\tau_j} & \text{if } i < j \\ s_i s_j \sqrt{\tau_j/\tau_i} & \text{otherwise} \end{cases}.$$

Proof This result is a special case of the generalized $\mathbb{M}$ binary valuation equation given in Theorem 1 in Skipper and Buchen (2003), pp. 10–12, when there is only a single underlying asset. Section 5.2.(b) of their paper gives an example. □

An n-th order bond (asset) binary corresponds to the special case $\eta = 0$ ($\eta = 1$) and has a payoff of one unit of cash (one asset). Following Buchen (2004), we further define higher-order $\mathcal{Q}$ options to have a payoff equal to that of a plain vanilla call option with strike price ξ_n at the terminal maturity date T_n conditional on $s_i S_{T_i} > s_i \xi_i$ for all $i \in \{1, 2, \ldots, n\}$. Introducing these special cases often helps making the notation more clear and compact.

Definition B.2 ($\boldsymbol{n}$-th Order Bond Binaries, Asset Binaries and $\boldsymbol{\mathcal{Q}}$ Options). Let $\boldsymbol{s}$, $\boldsymbol{\xi}$ and $\boldsymbol{T}$ be as in Definition B.1. The n-th order bond binary $\mathcal{B}^{\boldsymbol{s}}_{\boldsymbol{\xi}}$, asset binary $\mathcal{A}^{\boldsymbol{s}}_{\boldsymbol{\xi}}$ and $\mathcal{Q}$ option $\mathcal{Q}^{\boldsymbol{s}}_{\boldsymbol{\xi}}$ have the time T_n payoff

$$\begin{aligned} \mathcal{B}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(\boldsymbol{S_T}, 0) &= 1_n\left\{\mathrm{diag}(\boldsymbol{s})\boldsymbol{S_T} > \mathrm{diag}(\boldsymbol{s})\boldsymbol{\xi}\right\}, \\ \mathcal{A}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(\boldsymbol{S_T}, 0) &= S_{T_n} 1_n\left\{\mathrm{diag}(\boldsymbol{s})\boldsymbol{S_T} > \mathrm{diag}(\boldsymbol{s})\boldsymbol{\xi}\right\}. \\ \mathcal{Q}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(\boldsymbol{S_T}, 0) &= \mathcal{A}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(\boldsymbol{S_T}, 0) - \xi_n \mathcal{B}^{\boldsymbol{s}}_{\boldsymbol{\xi}}(\boldsymbol{S_T}, 0). \qquad \triangle \end{aligned}$$

Care has to be taken when multiplying the image of a power binary with an indicator. Lemma B.1 provides the necessary result and is used repeatedly. It immediately follows from Definition A.3.

Lemma B.1 (Product of the Image of a Power Binary and an Indicator). *Let ${}^{\eta}\mathcal{P}^{s_2}_{\xi_2}(S,\tau)$ be the valuation function of a first-order power binary. Then the following relationship holds*

$$\overset{B}{\mathcal{I}}\left\{{}^{\eta}\mathcal{P}^{s_2}_{\xi_2}(S,\tau)\right\}1\{s_1S > s_1\xi_1\} = \overset{B}{\mathcal{I}}\left\{{}^{\eta}\mathcal{P}^{-s_1s_2}_{\hat{\xi}_1\xi_2}(S,0,\tau)\right\},$$

where

$$\hat{\xi}_1 = \frac{B^2}{\xi_1}.$$

The property continues to hold for higher-order binaries.

The following Propositions derive some first-order sensitivities of power binary prices with respect to market and product static parameters. The delta and vega are important to hedge against changes in the underlying asset price and volatility, respectively. The strike sensitivity is required for computing transition densities in multi-period settings; see Section 3.3. These results represent novel contributions.

Proposition B.2 (Delta of n-th Order Power Binaries). *Let $\boldsymbol{\rho}_i$ be the i-th column of the matrix $\boldsymbol{C}$ and define*

$$\begin{aligned}
\hat{\boldsymbol{V}}_i &= \boldsymbol{C}^{-i} - \boldsymbol{\rho}_i^{-i}\left(\boldsymbol{\rho}_i^{-i}\right)', \\
\hat{\boldsymbol{D}}_i^2 &= \operatorname{diag}\left(\boldsymbol{V}_i\right), \\
\hat{\boldsymbol{\alpha}}_{\eta,i} &= \left(\hat{\boldsymbol{D}}_i\right)^{-1}\left(\boldsymbol{\alpha}_\eta^{-i} - \boldsymbol{\rho}_i^{-i}\alpha_{\eta,i}\right), \\
\hat{\boldsymbol{C}}_i &= \left(\hat{\boldsymbol{D}}_i\right)^{-1}\hat{\boldsymbol{V}}_i\left(\hat{\boldsymbol{D}}_i\right)^{-1},
\end{aligned}$$

where we use the notation $\boldsymbol{x}^{-i}$ ($\boldsymbol{X}^{-i}$) to indicate the i-th element (the i-th row and column) is removed from the vector $\boldsymbol{x}$ (the matrix $\boldsymbol{X}$). The time $0 \le t \le T_1$ first-order asset price sensitivity of the n-th order power binary is given by

$$\frac{\partial\, {}^{\eta}\mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}}{\partial S}(S,t) = \frac{F_\eta}{S}\left(\mathcal{N}_n\left(\boldsymbol{\alpha}_\eta;\boldsymbol{C}\right) + \sum_{i=1}^{n}\frac{s_i}{\sigma\sqrt{\tau_i}}\mathcal{N}'\left(\alpha_{\eta,i}\right)\mathcal{N}_{n-1}\left(\hat{\boldsymbol{\alpha}}_{\eta,i};\hat{\boldsymbol{C}}_i\right)\right).$$

Here, $\mathcal{N}'(x)$ is the univariate standard normal PDF and we define $\mathcal{N}_0(\cdot,\cdot) \coloneqq 1$.

Proof We have

$$\frac{\partial^\eta \mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}}{\partial S}(S,t) = \frac{F_\eta}{S}\left(\mathcal{N}_n\left(\boldsymbol{\alpha}_\eta; \boldsymbol{C}\right) + S\frac{\partial}{\partial S}\int_{-\infty}^{\boldsymbol{\alpha}_\eta} \mathcal{N}_n'(\boldsymbol{x}, \boldsymbol{C})\mathrm{d}\boldsymbol{x}\right),$$

where we used the integral representation of the n-variate standard normal CDF. Each of the n upper limits of integration $\alpha_{\eta,i}$ is a function of S. Thus, by a repeated application of the Leibniz rule we get

$$\frac{\partial}{\partial S}\int_{-\infty}^{\boldsymbol{\alpha}_\eta} \mathcal{N}_n'(\boldsymbol{x}, \boldsymbol{C})\mathrm{d}\boldsymbol{x} = \sum_{i=1}^{n} \frac{\partial \alpha_{\eta,i}}{\partial S}\mathcal{N}'\left(\alpha_{\eta,i}\right)\int_{-\infty}^{\boldsymbol{\alpha}_\eta^{-i}} \mathcal{N}_{n-1}'\big(\boldsymbol{x}^{-i}; \hat{\boldsymbol{\mu}}_i, \hat{\boldsymbol{V}}_i\big)\mathrm{d}\boldsymbol{x}^{-i}.$$

Here, conditional on $x_i = \alpha_{\eta,i}$, the vector $\boldsymbol{x}^{-i}$ has a multivariate normal distribution with mean vector $\hat{\boldsymbol{\mu}}_i = \boldsymbol{\rho}_i^{-i}\alpha_{\eta,i}$ and covariance matrix $\hat{\boldsymbol{V}}_i$ as given in Proposition B.2; see e.g. Theorem B.7 in Greene (2008). Consequently, $\big(\hat{\boldsymbol{D}}_i\big)^{-1}\left(\boldsymbol{x}^{-i} - \hat{\boldsymbol{\mu}}_i\right)$ has a multivariate standard normal distribution with correlation matrix $\hat{\boldsymbol{C}}_i$ as given in Proposition B.2. The result follows when substituting for the partial derivatives of $\alpha_{\eta,i}$

$$\frac{\partial \alpha_{\eta,i}}{\partial S} = \frac{s_i}{S\sigma\sqrt{\tau_i}}.$$

□

Proposition B.3 (Vega of *n*-th Order Power Binaries).

The time $0 \le t \le T_1$ first-order volatility sensitivities of the n-th order binary is given by

$$\frac{\partial^\eta \mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}}{\partial \sigma}(S,t) = -F_\eta \sum_{i=1}^{n} \frac{s_i d_{1-\eta,i}}{\sigma}\mathcal{N}'\left(\alpha_{\eta,i}\right)\mathcal{N}_{n-1}\big(\hat{\boldsymbol{\alpha}}_{\eta,i}; \hat{\boldsymbol{C}}_i\big).$$

Proof The proof is fully analogous to that of Proposition B.2 and thus omitted. We use that

$$\frac{\partial \alpha_{\eta,i}}{\partial \sigma} = -\frac{\alpha_{1-\eta,i}}{\sigma}.$$

□

Proposition B.4 (Dual Delta of *n*-th Order Power Binaries).

The time $0 \le t \le T_1$ first-order sensitivity of the n-th order power binary with respect to the j-th strike is given by

$$\frac{\partial^\eta \mathcal{P}^{\boldsymbol{s}}_{\boldsymbol{\xi}}}{\partial \xi_j} = -F_\eta \frac{s_j}{\xi_j\sigma\sqrt{\tau_j}}\mathcal{N}'\left(\alpha_{\eta,i}\right)\mathcal{N}_{n-1}\big(\hat{\boldsymbol{\alpha}}_{\eta,i}; \hat{\boldsymbol{C}}_i\big)$$

Proof The proof is fully analogous to that of Proposition B.2 and thus omitted. We use that

$$\frac{\partial \alpha_{\eta,j}}{\partial \xi_j} = -\frac{s_j}{\xi_j \sigma \sqrt{\tau_j}}.$$

□

C Pay-at-Hit Rebates

This section provides the valuation function for fixed rebates that are paid immediately upon the first hitting time of a barrier. This result is well-known and is usually obtained by an explicit integration over the first hitting time PDF. We take a different, though equivalent, approach that unveils the underlying structure in terms of perpetual and power binaries as well as their respective images. It is this decomposition, together with the valuation formula for higher-order power binaries in Proposition B.1, that allows us to value complex contingency structures involving pay-at-hit rebates.

Definition C.1 (Perpetual Binaries).
A perpetual binary $\mathcal{U}_B^{\psi}$ has a unit payoff upon the first hitting time ν of S to the barrier $B \in \mathbb{R}_+$ defined as

$$\nu = \inf \{t \in \mathbb{R}_{+,0} : \psi S_t \leq \psi B\}. \qquad \triangle$$

The valuation function for perpetual binaries is well-known. We thus state the following proposition without proof.

Proposition C.1 (Valuation of Perpetual Binaries).
The value of the perpetual binary is given by

$$\mathcal{U}_B^{\psi}(S) = \left(\frac{S}{B}\right)^{\beta(-\psi)},$$

where

$$\begin{aligned} \beta(\psi) &= \alpha + \psi\sqrt{\lambda + \alpha^2}, \\ \lambda &= \frac{2r}{\sigma^2} \end{aligned}$$

and α is as in Definition A.3.

We note that $\mathcal{U}_B^\psi(S,\tau)$ is a power of S. Thus, Proposition B.1 can be used to value higher-order options on perpetual binaries.

Definition C.2 (Pay-at-Hit Rebate).
Let ν as in Definition C.1. A pay-at-hit rebate with maturity $T \in \mathbb{R}_+$ has the payoff

$$\mathcal{R}_B^\psi\,(S,T) = e^{r(T-\nu)}1\{\nu \leq T\}. \qquad \triangle$$

Proposition C.2 (Valuation of Pay-at-Hit Rebates).
The valuation function of the pay-at-hit rebate is given by

$$\mathcal{R}_B^\psi(S,\tau) = \mathcal{U}_B^\psi(S) - B^{-\beta(-\psi)}\left({}^{\beta(-\psi)}\mathcal{P}_B^\psi(S,\tau) - \overset{B}{\mathcal{I}}\left\{ {}^{\beta(-\psi)}\mathcal{P}_B^\psi(S,\tau)\right\}\right),$$

where $\beta(\psi)$ is as in Proposition C.1.

Proof We first note that the solution to the pay-at-hit rebate valuation problem can be decomposed as

$$\mathcal{R}_B^\psi(S,\tau) = \mathcal{U}_B^\psi(S) - \tilde{V}^1(S,\tau).$$

Here, $\tilde{V}^1(S,\tau)$ satisfies an IBVP with barrier B, flag ψ and initial condition $f(S) = \mathcal{U}_B^\psi(S)$; see Definition A.2. That is, a pay-at-hit rebate corresponds to initially taking a long position in the otherwise identical perpetual claim and closing it out at the maturity, given that the barrier has not been breached. We use Definition B.1 to express the valuation function for the corresponding full-range problem $\tilde{V}_B^1(S,\tau)$ in terms of a power binary whose value is given by Proposition B.1. The valuation function for $\tilde{V}^1(S,\tau)$ then follows from the method of images; see Proposition A.1. $\square$

D Discrete Barriers

This section provides the necessary results for valuing discrete barriers using the numerical algorithm in Section 3. The survival domain is now defined as $\mathcal{D} = \{S \in \mathbb{R}_+ : \phi S > \phi K\}$.

Lemma D.1 (Intraday Valuation on the Maturity Date II).

The valuation function $V\left(S, T_n^O\right)$ *is given by*

$$V\left(S, T_n^O\right) = \phi\left(\mathcal{A}_B^{\phi}(S,\tau) - K\mathcal{B}_B(S,\tau) - \overset{K}{\mathcal{I}}\left\{\mathcal{A}_B^{\phi}(S,\tau) - K\mathcal{B}_B(S,\tau)\right\}\right).$$

Lemma D.2 (Interday Backward Induction II).

For $i < n$*, the solution for* $V\left(S, T_i^O\right)$ *is given by*

$$V\big(S, T_i^O\big) = \underbrace{e^{-r\hat{\tau}^D}\int_{\mathcal{D}} V\left(y, T_{i+1}^O\right) f_1(y|S)\mathrm{d}y}_{C\left(S,T_i^O\right) \sim \text{ continuation value}} + \underbrace{e^{-r\hat{\tau}^I}\phi\int_K^B (y-K) f_2(y|S)\mathrm{d}y}_{R(S) \sim \text{ value of the rebate at } T_i^C},$$

where

$$\begin{aligned} f_1(y|x)\mathrm{d}y &= \mathbb{P}^*\left\{ S_{T_{i+1}^O} \in \mathrm{d}y \,\cap\, \nu \geq T_{i+1}^O \middle| \nu > T_i^O,\ S_{T_i^O} = x\right\}, \\ f_2(y|x)\mathrm{d}y &= \mathbb{P}^*\left\{ S_{T_i^C} \in \mathrm{d}y \,\cap\, \nu \geq T_i^C \middle| \nu > T_i^O,\ S_{T_i^O} = x\right\}. \end{aligned}$$

Lemma D.3 (Conditional Transition Probability Function II).

The conditional transition probability function is given by

$$f_1(y|x) = e^{r\hat{\tau}^D}\left(\frac{\partial \mathcal{B}_{By}^{\phi-}}{\partial y}\left(x, \hat{\tau}^I, \hat{\tau}^D\right) - \overset{K}{\mathcal{I}}\left\{\frac{\partial \mathcal{B}_{By}^{\phi-}}{\partial y}\left(x, \hat{\tau}^I, \hat{\tau}^D\right)\right\}\right).$$

where the partial derivatives are defined as in Lemma 3.2.

Lemma D.4 (Rebate Valuation II).

The value of the rebate component in Lemma D.2 is

$$\begin{aligned} R(S) &:= e^{-r\hat{\tau}^I}\phi\int_K^B (y-K) f_2(y|S)\mathrm{d}y \\ &= \phi\bigg(\mathcal{Q}_K^{\phi}\left(S,\hat{\tau}^I\right) - \overset{K}{\mathcal{I}}\left\{\mathcal{Q}_K^{\phi}\left(S,\hat{\tau}^I\right)\right\} - \mathcal{Q}_B^{\phi}\left(S,\hat{\tau}^I\right) + \overset{K}{\mathcal{I}}\left\{\mathcal{Q}_B^{\phi}\left(S,\hat{\tau}^I\right)\right\} \\ &\quad -(B-K)\mathcal{B}_B^{\phi}\left(S,\hat{\tau}^I\right)\bigg). \end{aligned}$$

Lemma D.5 (Continuation Value II).

The continuation value for the final step is given by

$$C(S,t) = \sum_{j=0}^{m}\sum_{k=0}^{d}\beta_{j,k}\begin{cases} {}^{k}\mathcal{P}^{\phi}_{S_j}\left(S,\tau_1^O\right) & \text{if } t \in \left[T_0^C, T_1^O\right) \\ \left({}^{k}\mathcal{P}^{\phi\phi}_{BS_j}\left(S,\tau_0^C,\tau_1^O\right) - \overset{K}{\mathcal{I}}\left\{{}^{k}\mathcal{P}^{\phi\phi}_{BS_j}\left(S,\tau_0^C,\tau_1^O\right)\right\}\right) & \text{otherwise} \end{cases}.$$

where the coefficients $\beta_{j,k}$ *are as in Lemma 3.4.*