Quanto Skew

Peter Jäckel*

First version: 25th July 2009
This version: 4th June 2016

Abstract

We assess the effect of an implied volatility skew for an FX rate on quanto forwards and quanto options of an asset that itself is subject to an implied volatility skew using a simplistic double displaced diffusion model.

1 Introduction

Quanto contracts are financial derivatives whose payout currency differs from the natural denomination of their underlying financial observable. Their purpose is to provide exposure to the performance of the observable without exposure to a currency conversion risk. Simple examples include forward contracts or options on US or JPY government bonds, commodity futures, or equity indices written such that the numerical return on the underlying is paid out directly in a different currency, without the application of any FX conversion factor. This is equivalent to the concept of entering into the equivalent domestic derivative, e.g., a forward contract or option on a US asset written in USD, with a guarantee that whatever the payout is, it will be converted at the time of payment at an FX rate that is agreed upon at inception of the deal. Quanto contracts of vanilla nature are traded over-the-counter in significant size. They either are requested directly by end-investors, or are used as hedges for components of multi-asset investment strategies and other non-vanilla derivatives.

Despite their relative importance, surprisingly little research has been focussed on quanto derivative pricing. Whilst practitioners have gone to extreme lengths in their efforts to design, implement, and calibrate models for underlying assets in order to accomodate, or even explain, the observable implied volatility skew, when it comes to quanto derivatives, even vanilla contracts tend to be valued with simple adjustments on top of what may otherwise be a rather sophisticated stochastic processes for the underlying asset.

*OTC Analytics

Key words and phrases. quanto, skew, double displaced diffusion model.

2 Common quanto adjustments

The most commonly used quanto correction used in practice is without doubt to modify the expected value of the asset when quantoed, and otherwise reuse whatever stochastic model may be in favour for the underlying asset. In the case of correlated geometric Brownian motion for the underlying observable and the FX rate, it is well known that one can, by an argument of change of measure, justify the use of a simple adjustment for all quanto derivatives. This is because simply replacing the forward F of the underlying by

$$F' = F\mathrm{e}^{\hat{c}}, \quad \text{using } \hat{c} = \hat{\sigma}_S \rho \hat{\sigma}_Q T \,, \tag{1}$$

with $\hat{\sigma}_S$ being the domestic at-the-money implied volatility of the asset, ρ being the process correlation (i.e., correlation of increments, usually estimated from time series) between asset and FX rate (in terms of value of one investor currency unit expressed in the asset's domestic currency), and $\hat{\sigma}_Q$ being the at-the-money implied volatility of the FX rate, takes care of the effect of quantoing for *all derivatives* expiring at T on this asset. Possibly primarily for reasons of convenience, the quanto adjustment (1) has, however, also been deployed as a simple forward correction irrespective of what implied volatility skew may be observable for the underlying asset, and for the FX rate.

When skew parametrisations are used for domestic options on the asset in the shape of some functional form as in

$$\hat{\sigma}(K) = f(F, K, T, \lambda_1, \lambda_2, \ldots) \tag{2}$$

where $f()$ could be any functional form, e.g., SABR [HKL02], or a volatility implied from a model expressed in terms of an underlying martingale observable such as Heston [Hes93], CEV [CR76], displaced diffusion [Rub83], etc., then, in practice, the convential approach tends to be to retain all of the parameters $\lambda_1, \lambda_2, \ldots$, exactly as they are used in the domestic currency. Regarding the effective forward, there are two common schools of thought:-

a) Determine an effective volatility coefficient $\hat{\sigma}$ exactly as in (2) using the forward as if the option was entirely domestic. Subsequently, price the quanto option with Black's formula replacing the forward F by F' as given in (1). This approach, in essence, transfers the domestic skew unaltered to quanto options, whence we refer to it as the *Domestic-Forward-ATM-Quanto* method, or *DFAQ* for short. The term "ATM-Quanto" refers to the fact that the (conventionally approximated) quanto forward F' used in the final pricing stage is approximated as the domestic forward that is adjusted using at-the-money implied volatilities.

b) Determine the effective volatility coefficient directly with the quanto-adjusted effective forward F' as given in (1). Subsequently, price the quanto option with Black's formula using the same adjusted forward F'. In this article, we refer to this approach as the *Quanto-Forward-ATM-Quanto*, or *QFAQ* for short.

In this article, we attempt to shed some light on the question how well these simplistic quanto adjustments perform when we go beyond the Black-Scholes setting, when even one of the simplest of all self-consistent skew-generating models is used. In our investigation, we use an extremely light-weight model for the skew in the asset and the FX rate in aid of attaining closed form solutions for vanilla quanto options which we use for comparison. Whilst a very simple model is clearly not enough to permit a general statement about the quanto effect for all models, we believe it suffices to give a first order indication as to the quality of conventional quanto adjustments. This entails a closer look at the exact quanto forward correction, which turns out to be model dependent. Further, we attempt to give an indication of the *quanto effect on the model skew parameters*. Whilst the analysis presented here is, admittedly, predicated on very simple modelling ideas, it highlights that, in general, *it is necessary to adjust the full implied volatility skew when an asset is quantoed into another currency.*

3 Exact quanto valuation

Assume a financial observable S is denominated in currency X. Assume that an investor, whose natural home currency is Y, wishes to participate in the asset's performance between today and some future time horizon T, but wishes not to be exposed to any currency risk. This investor may be interested in what is known as a *quanto option* that pays $(\theta \cdot (S_T - K))_+$ directly in currency Y, with $\theta = +1$ for calls and $\theta = -1$ for puts.

Denote Q as the value of one currency unit of Y expressed in currency X, i.e., as the quotient Y/X. To the seller of a quanto option, the net present value of the contract, expressed in units of currency X (e.g., for reasons of availability of options to hedge the exposure to S), is

$$\tilde{V}^{\mathrm{X}}(t) = P(t,T) \cdot \mathsf{E}^{\mathrm{X}}[(\theta(S_T - K))_+ \cdot Q_T] \tag{3}$$

with $P(t,T)$ representing today's value of a zero coupon bond paying at T, and the expectation being taken in an X-denominated T-forward measure. Naturally, we can translate an X-denominated net present value into a Y-denominated net present value by spot FX conversion:

$$\tilde{V}^{\mathrm{Y}}(t) = \frac{\tilde{V}^{\mathrm{X}}(t)}{Q_t} = P(t,T) \cdot \mathsf{E}^{\mathrm{X}}\left[(\theta(S_T - K))_+ \cdot \frac{Q_T}{Q_t}\right] . \tag{4}$$

Equations (3) and (4) are model-independent.

4 Displaced Diffusion

The displaced diffusion model [Rub83] is a convenient device to generate a skewed implied volatility profile when explicit control over the curvature, i.e., the actual smile, is of secondary importance. It allows for vanilla option valuation formulae in terms of the Black-Scholes pricing formula with adjusted input parameters which makes it exceptionally easy and efficient in practical applications. Its main drawbacks are that the curvature of the implied volatility profile it generates is implicitly determined by its skew, and that the domain of the distribution it generates for the underlying financial observable does not begin at zero, as for the lognormal distribution, but at an offset, which is negative when the skew is negative. In other words, when calibrated to a typical equity or interest rate skew (at the money), it permits for the underlying observable to attain negative values. This is clearly a deficiency in a variety of circumstances. Nevertheless, it has become a favourite of many researchers and practitioners whenever an assessment of the impact of a mere skew of implied volatilities on a specific issue of financial engineering is required.

In the simple case of constant parameters, the displaced diffusion model can be summarised as the unique martingale process

$$S(t) = S(0) \cdot \left(\mathrm{e}^{-\frac{1}{2}\sigma_S^2\beta_S^2 t + \sigma_S\beta_S W(t)} - (1-\beta_S)\right) \Big/ \beta_S \tag{5}$$

for a financial observable $S(t)$ with $W(t)$ being a standard Wiener process. European vanilla option prices can be calculated by a simple transformation of variables and are given by

$$V_{\pm 1}(S,K,\sigma,\beta,T) = B_{\pm 1}(\tfrac{1}{\beta}S, K + \tfrac{(1-\beta)}{\beta}S, \sigma\beta, T) \tag{6}$$

with the Black-Scholes formula

$$B_\theta(S,K,\hat\sigma,T) = \theta\cdot\Big[\, S\cdot\Phi\left(\theta\cdot(\tfrac{x}{\hat\varsigma}+\tfrac{\hat\varsigma}{2})\right) - K\cdot\Phi\left(\theta\cdot(\tfrac{x}{\hat\varsigma}-\tfrac{\hat\varsigma}{2})\right)\Big] \tag{7}$$

$$x = \ln(S/K) \tag{8}$$

$$\hat\varsigma = \hat\sigma\cdot\sqrt{T} \tag{9}$$

where we have omitted all discounting. At the money, the pricing formula (6) becomes

$$V_{+1}(S,K,\sigma,\beta,T)|_{K=S} = S\cdot\left(2\Phi\Big(\tfrac{\sigma\beta\sqrt{T}}{2}\Big)-1\right)\tfrac{1}{\beta} \tag{10}$$

for call options. We can compare this with the Black-Scholes call option price at the money given by

$$B_{+1}(S,K,\hat\sigma(K),T)|_{K=S} = S\cdot\left(2\Phi\Big(\tfrac{\hat\sigma_{\rm atm}\sqrt{T}}{2}\Big)-1\right) \tag{11}$$

with $\hat\sigma_{\rm atm}=\hat\sigma(S)$. This allows for an explicit solution for the at-the-money Black implied volatility

$$\hat\sigma_{\rm atm} = \tfrac{2}{\sqrt{T}}\Phi^{-1}\left(\tfrac{1}{\beta}\Phi\left(\tfrac{\sigma\beta\sqrt{T}}{2}\right)-\tfrac{(1-\beta)}{2\beta}\right)\ . \tag{12}$$

For calibration purposes, the inverse solution is also of practical use:

$$\sigma = \tfrac{2}{\beta\sqrt{T}}\Phi^{-1}\left(\beta\Phi\left(\tfrac{\hat\sigma_{\rm atm}\sqrt{T}}{2}\right)+\tfrac{(1-\beta)}{2}\right)\ . \tag{13}$$

Of further interest is that we can derive from this and

$$\tfrac{\rm d}{{\rm d}K}B_{+1}(S,K,\hat\sigma(K),T) = \tfrac{\rm d}{{\rm d}K}V_{+1}(S,K,\sigma,\beta,T) \tag{14}$$

the at-the-money Black implied volatility skew solely in terms of β and the at-the-money Black implied volatility itself:

$$\tfrac{\rm d}{{\rm d}K}\hat\sigma(K)\big|_{K=S} = \tfrac{(\beta-1)}{2}\tfrac{\sqrt{2\pi}}{S\sqrt{T}}\left[2\Phi\Big(\tfrac{\hat\sigma\sqrt{T}}{2}\Big)-1\right]{\rm e}^{\frac{\hat\sigma^2 T}{8}} \tag{15}$$

$$= \tfrac{(\beta-1)}{2}\tfrac{\hat\sigma}{S}\left(1+\tfrac{\hat\sigma^2T}{12}+\tfrac{\hat\sigma^4T^2}{240}+\cdots\right)\ , \tag{16}$$

where $\hat\sigma = \hat\sigma_{\rm atm}$ on the right hand side. Equations (13) and (15) enable us to calibrate a displaced diffusion model to a given at-the-money volatility and skew with great ease.

5 Displaced Diffusion Quanto Skew

We assume a displaced diffusion model for both the underlying asset and the foreign exchange rate process in the domestic martingale measure of the asset (currency X) as discussed in section 3. We note that this, in general, prohibits us from explicitly changing measure to the investor's domestic measure (currency Y) since the reciprocal foreign exchange rate $1/Q$ is not a measurable process as Q can attain zero and even become negative when $\beta<1$. In this setting, the undiscounted quanto option price (4) is given by

$$\tilde V^{\rm Y}_{\pm1}(S_0,K,\sigma_S,\beta_S,\sigma_Q,\beta_Q,\rho,T)$$

$$= {\sf E}^{\rm X}\left[(\pm(S_T-K))_+\cdot\frac{Q_T}{Q_t}\right] \tag{17}$$

$$= {\sf E}^{\rm X}\left[\left(\pm\Big(\tfrac{S_0}{\beta_S}{\rm e}^{-\frac{\sigma_S^2\beta_S^2T}{2}+\sigma_S\beta_SW_S(T)}-\tfrac{1-\beta_S}{\beta_S}S_0-K\Big)\right)_+ \cdot\left({\rm e}^{-\frac{\sigma_Q^2\beta_Q^2T}{2}+\sigma_Q\beta_QW_Q(T)}-(1-\beta_Q)\right)\Big/\beta_Q\right] \tag{18}$$

$$= \tfrac{1}{\beta_Q}B_{\pm1}\Big(\tfrac{1}{\beta_S}S_0{\rm e}^{\tilde c},K+\tfrac{(1-\beta_S)}{\beta_S}S_0,\sigma_S\beta_S,T\Big) - \tfrac{(1-\beta_Q)}{\beta_Q}B_{\pm1}\Big(\tfrac{1}{\beta_S}S_0,K+\tfrac{(1-\beta_S)}{\beta_S}S_0,\sigma_S\beta_S,T\Big) \tag{19}$$

with

$$\tilde c = \rho\sigma_S\beta_S\sigma_Q\beta_QT \tag{20}$$

wherein ρ represents the correlation of $W_S(T)$ and $W_Q(T)$. The par strike for a quanto forward contract can equally be computed:

$$\tilde F = S_0\cdot\left[1+\left({\rm e}^{\tilde c}-1\right)/(\beta_S\beta_Q)\right]\ . \tag{21}$$

Armed with this information, we can attempt to find *matched quanto displaced diffusion parameters* whose purpose it is to enable the (approximate) valuation of Y-denominated options on the underlying asset without the need for a combined asset-FX model. For standard geometric Brownian motion, we do of course know how to do this change of measure exactly — the volatility information of S in the Y measure is the same as in the X measure, but the par forward changes to the quanto forward. For other models, the change of measure results in more complicated deformations of the quantoed asset distribution. In other words, not just the forward changes, but the *entire implied volatility profile* changes, too. Here, we try to approximate the *quanto skew* with the same form of parametrisation as the domestic skew: the quanto forward $\tilde F$ which is already given in equation (21), a quanto displaced diffusion volatility $\tilde\sigma$, and a quanto skew parameter $\tilde\beta$. The hope is that, given the right choice of the quanto skew parameters $\tilde\sigma$ and $\tilde\beta$, we can use the vanilla displaced diffusion formula (6) to approximate quanto options, i.e.,

$$V_{\pm1}(\tilde F,K,\tilde\sigma,\tilde\beta,T)\approx\tilde V^{\rm Y}_{\pm1}(S_0,K,\sigma_S,\beta_S,\sigma_Q,\beta_Q,\rho,T)\,. \tag{22}$$

For this purpose, we match the call option price and the skew at the quanto forward. This will ensure that at least for options struck near the quanto forward, we will have good agreement. We shall see later how far the approximation can be used. For this, in addition to formula (19), we need

$$\tfrac{\rm d}{{\rm d}K}\tilde V^{\rm Y}_{+1}(S_0,K,\sigma_S,\beta_S,\sigma_Q,\beta_Q,\rho,T) = \tfrac{1-\beta_Q}{\beta_Q}\Phi\left(\tfrac{\xi}{\varsigma}-\tfrac{\varsigma}{2}\right)-\tfrac{1}{\beta_Q}\Phi\left(\tfrac{\xi+\tilde c}{\varsigma}-\tfrac{\varsigma}{2}\right) \tag{23}$$

with

$$\xi = -\ln\left(1-\beta_S+\beta_S\tfrac{K}{S_0}\right) \tag{24}$$

$$\varsigma = \sigma_S\beta_S\sqrt{T}\ . \tag{25}$$

Using the abbreviations

$$\tilde{v} := \tilde{V}^{\mathrm{Y}}_{+1}(S_0, K, \sigma_S, \beta_S, \sigma_Q, \beta_Q, \rho, T)\Big|_{K=\tilde{F}} \tag{26}$$

$$\tilde{\kappa} := \tfrac{\mathrm{d}}{\mathrm{d}K}\tilde{V}^{\mathrm{Y}}_{+1}(S_0, K, \sigma_S, \beta_S, \sigma_Q, \beta_Q, \rho, T)\Big|_{K=\tilde{F}} \tag{27}$$

we can now express our objective as solving

$$\tilde{v} = \frac{\tilde{F}}{\tilde{\beta}}\left[2\Phi\left(\tfrac{1}{2}\tilde{\sigma}\tilde{\beta}\sqrt{T}\right) - 1\right] \tag{28}$$

$$\tilde{\kappa} = -\Phi\left(-\tfrac{1}{2}\tilde{\sigma}\tilde{\beta}\sqrt{T}\right) \tag{29}$$

for $\tilde{\beta}$ and $\tilde{\sigma}$. The solution is:-

$$\begin{aligned}\tilde{\beta} &= (1+2\tilde{\kappa})\tilde{F}/\tilde{v}\\ \tilde{\sigma} &= 2\;\Phi^{-1}\Big(\tfrac{1}{2}(1+\tilde{\beta}\tilde{v}/\tilde{F})\Big)\Big/\,(\tilde{\beta}\sqrt{T})\;.\end{aligned} \tag{30}$$

6 Examples

We now give some examples for the quality of the quanto skew approximation (30) for the double displaced diffusion setup. Before we look at the quality of our approximation, we start off by highlighting a fact that is unfortunately far too easily overlooked when dealing with quanto skew parametrisations and approximations.

6.1 No FX skew

This case is encompassed by our double displaced diffusion framework when $\beta_Q = 1$. It is well known that under these circumstances it is indeed possible to change measure to the investor currency Y since the FX rate is lognormally distributed and cannot attain zero (or even become negative). Given that this is arguably the simplest quanto framework that goes beyond the skew-free pure Black-Scholes model for the underlying asset and the FX rate, one might hope that industry-practice conventional quanto adjustments hold for this case exactly. We show in figure 1

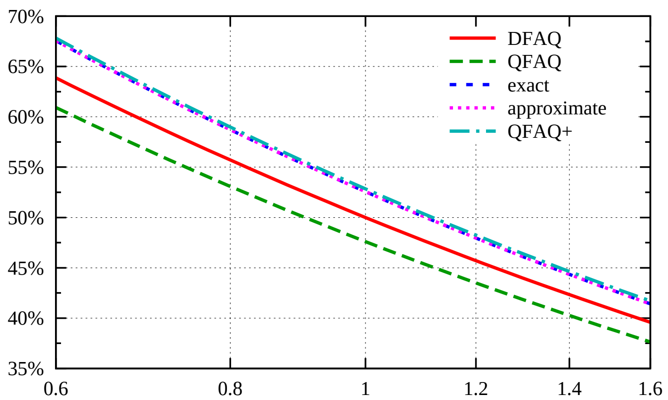

Figure 1: Implied volatilities as a function of strike K for $S_0 = 1$, $T = 2$, $\beta_S = \frac{1}{16}$, $\sigma_S = 48.98\%$, $\beta_Q = 1$, $\sigma_Q = 20\%$, and $\rho = -50\%$.

the following implied volatility curves in comparison:-

DFAQ — for this, we use directly the simple displaced diffusion formula (6). Note that this is by construction identical to the domestic skew.

QFAQ — this is the approach of replacing the forward and reusing the domestic parameters, i.e., $V(F', K, \sigma_S, \beta_S, T)$, as explained in section 2.

exact — this is formula (19).

approximate — this is the approximation $V(\tilde{F}, K, \tilde{\sigma}, \tilde{\beta}, T)$ with (21) and (30).

QFAQ+ — this is $V(F', K, \sigma_S \frac{S_0}{F'}, \beta_S, T)$.

All data are expressed in terms of implied volatilities using the respectively exact forward as reference. This is to mean that what is shown for any of the price curves is the unique solution of

$$v_\theta = B_\theta\left(F, K, \hat{\sigma}_\theta(v_\theta, F, K), T\right) \tag{31}$$

for $\hat{\sigma}$ (which thereby implicitly becomes a function of v_θ, F, and K), with v_θ being the respective price (or approximation) and F the par forward strike that is consistent with the respective price formula. Specifically, we have:-

$$\begin{array}{ll} \textit{DFAQ} & \text{— } F = F'^{1} \\ \textit{QFAQ} & \text{— } F = F' \\ \textit{exact} & \text{— } F = \tilde{F} \\ \textit{approximate} & \text{— } F = \tilde{F} \\ \textit{QFAQ+} & \text{— } F = F' \end{array}$$

Since we are using the exact forward as reference for implied volatilities, call-put parity is preserved, and it is mathematically irrelevant whether we use $\theta = +1$ or $\theta = -1$ (though we may have chosen to always compute values on out-of-the money options in order to minimize numerical truncation). Note that the odd choice for σ_S for the data in figure 1 was made such that the domestic skew is calibrated to $\hat{\sigma}_S = 50\%$ at $K = S_0$ using equation (13). As a final note of form before we discuss the poignant details of the figure, we mention that we are not precisely comparing like for like in the figure since we are using different reference forwards for the calculation of implied volatilities. With the exception of the *DFAQ* curve, the difference is small since, for this example, $F' = 0.904837$ and $\tilde{F} = 0.902336$ are very close.

The most striking effect that can be seen in figure 1 is that whilst the implied volatility skew of the *exact* quanto curve appears to be shifted up, or to the right, whichever way one prefers to see it, the implied volatility skew of the *QFAQ* quanto adjustment appears to be shifted down (or left). The discrepancy for this, admittedly, or perhaps arguably[2], exaggerated

[1] Corrected on 2016-06-04 from S_0 to F', which makes this consistent with the earlier statement "*Note that this is by construction identical to the domestic skew*".

[2] A level of 50% for a two year horizon may have been considered exaggerated for equities in the pre-2008 era, but since then much higher levels of implied volatility have been observed, if only temporarily.

example is about 5% in terms of implied volatility, which is clearly not negligible. The explanation for this significant error in the conventional *QFAQ* approach lies with the fact that we are dealing with a pricing framework that belongs to the family of local volatility models. When we use the QFAQ quanto adjustment, we tend to focus on forward prices, and too easily forget that we are effectively dealing with an approximation for a (possibly stochastic) drift introduced by the quanto term. Simply absorbing the drift in an (approximately) adjusted forward, which is then used in a formula of local volatility type has the side effect that we inadvertently change the absolute level of volatility for the process of the underlying. To lowest order, one can see this if we compare the initial absolute volatility level of the actual asset process in measure Y, which is $\sigma_S \cdot S_0$, whereas the QFAQ adjustment implies that the asset process has initial absolute volatility close to $\sigma_S F'$. As evidence for this effect, we included in figure 1 the curve marked as *QFAQ+* where the used process volatility has been adjusted by the factor $\frac{S_0}{F'}$ to compensate for the otherwise incurred loss of volatility. As we can see in figure 1, this first order adjustment does indeed get volatility levels almost right. Unfortunately, this simplistic compensation only works for moderate maturities and choices of $|\beta_S| \ll 1$, i.e., when the model is nearly Bachelier. As a side note, we mention that it can easily be shown that in this case of $\beta_Q = 1$ the *approximate* pricing curve is identical to the *exact* solution, which is of course a desirable feature to have for the approximation, and this is reflected in the data.

Since, as we mentioned, we are not exactly comparing like for like in figure 1, we also show a comparison in terms of actual option price in figure 2. The data

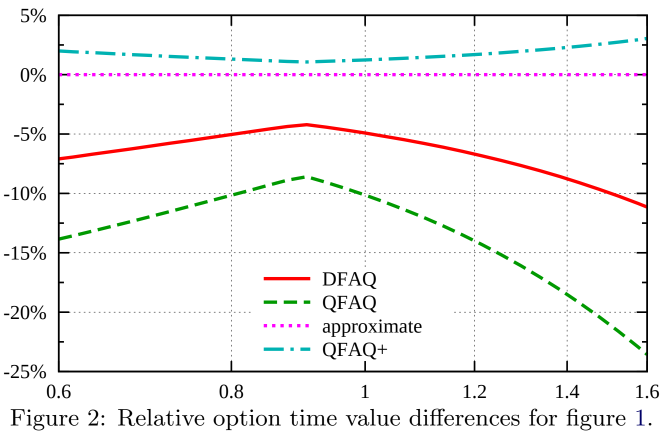

Figure 2: Relative option time value differences for figure 1.

in the figure are relative differences of the time value of the option, as priced with the respective methodology, when compared with the exact price. For instance, the curve marked "DFAQ" corresponds to

$$\frac{v_\theta^{\text{DFAQ}} - \big(\theta\cdot(\tilde{F}-K)\big)_+}{v_\theta^{\text{exact}} - \big(\theta\cdot(\tilde{F}-K)\big)_+} - 1$$

with the actual price v_θ^{DFAQ} computed as explained earlier in this section. Note that the apparent kink in the relative option time value difference curves is no error: it indicates the location of the exact quanto forward $\tilde{F}$. We can see in this figure that even in this very simple case of a pure local volatility smile that is consistent with an analytically tractable model, alas, both of the conventional quanto adjustment approaches discussed here give prices that, near the effective money, are about 5% different from the exact price. We can also see that the QFAQ+ approach appears to eliminate most of the error incurred due to the mis-estimation of initial absolute volatility that is inherent in the QFAQ approach. We hasten to add that this mis-estimation is particularly strong in the case of a setting that is consistent with a local volatility approach. When the underlying model ideas that give rise to the domestic smile are closer to a floating smile (also known as *sticky delta*) approach, the QFAQ methodology may be more natural than the DFAQ approach, though, this remains to be properly investigated. Our final observation regarding figure 2 is that the approximate method shows zero relative difference to the exact price, which is of course consistent with this method being identical to the exact pricing method for $\beta_Q = 1$ as was mentioned earlier. Since we consider the alternative representation of the data in terms of relative option time value errors a useful complement to the depiction of the resulting implied volatilities, we will from here on accompany all diagrams by this second view without further discussion in the text.

The effect we wanted to highlight in this section is thus: when dealing with quanto skew transformations, there is an intrinsic risk that by focussing on the quanto adjustment, which is often done in terms of an adjusted forward, we forget that the implied volatility parametrisation we use, which may be intended to be consistent with a model of local volatility style, incurs absolute volatility changes if we simply adjust the forward. The displaced diffusion model/parametrisation used in this article is not the only one that will have this problem. It is in fact common across the whole local volatility family: CEV, SABR, and the fully fledged non-parametric local volatility approach [Dup94]. Even when there is no FX skew, vanilla quanto option pricing deserves attention being paid to the choice of underlying model.

6.2 Co-inclining skew

In this case, the implied volatility skew and the FX skew lean in the same direction. We show in figure 3 the implied volatility curves for $\beta_Q = \frac{1}{16}$, with otherwise the same parameters as in figure 1. As we can see, the overall picture remains the same. In fig-

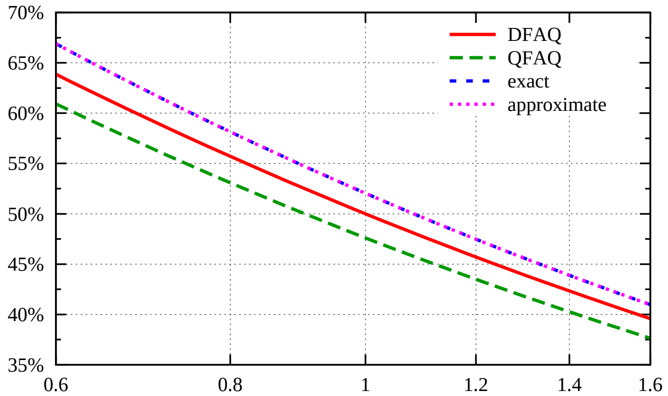

Figure 3: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{1}{16}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 48.98\%$, $\beta_Q = \frac{1}{16}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.95\%$, and $\rho = -50\%$. $F' = 0.9048$ and $\tilde{F} = 0.9024$.

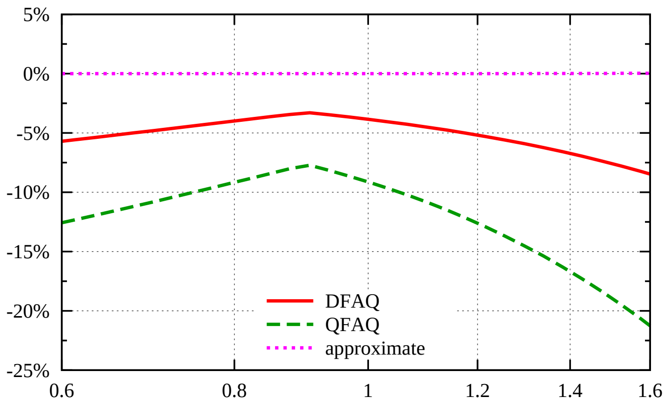

Figure 4: Relative option time value differences for figure 3.

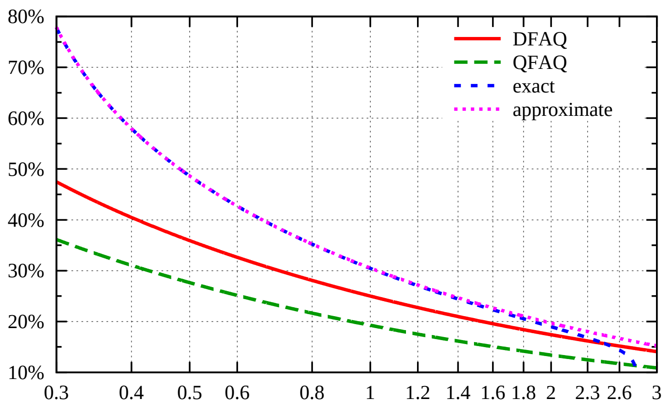

Figure 5: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{1}{16}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 23.76\%$, $\beta_Q = \frac{1}{16}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.95\%$, and $\rho = -50\%$. $F' = 0.6065$ and $\tilde{F} = 0.5405$.

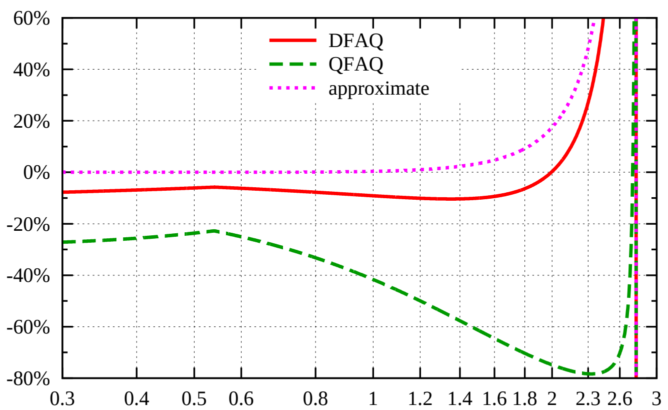

Figure 6: Relative option time value differences for figure 5.

ure 5, we show the extreme example where we have increased maturity to 20 years, though we reduced domestic at-the-money volatilities to 25%. We can see that the discrepancies widen significantly, and that our approximation (30) starts to break down for very high strikes $\gtrsim 2.6$, though, but agrees very well with the exact solution otherwise. It should be mentioned, however, that this breakdown is largely caused by the fact that the underlying model, due to its allowing for negative FX rates when $\beta_Q < 1$, in this extreme scenario, attains negative call option prices for $\gtrsim 2.77$. Thus, arguably, the failure of the approximation, which is designed to disallow negative option prices, is of no concern greater than the fundamental shortcoming of the chosen simplistic model to allow for negative option prices in extreme scenarios.

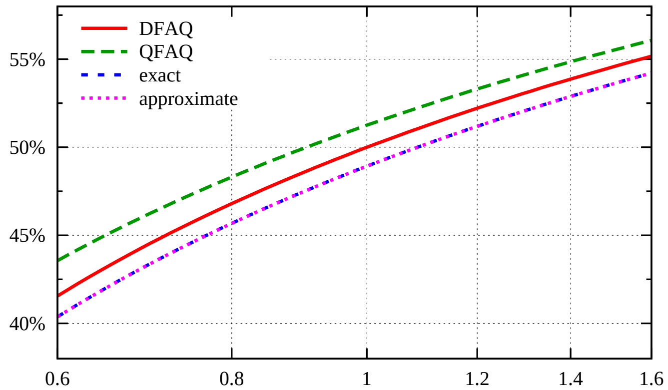

Figure 7: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 51.42\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.08\%$, and $\rho = -50\%$. $F' = 0.9048$ and $\tilde{F} = 0.9079$.

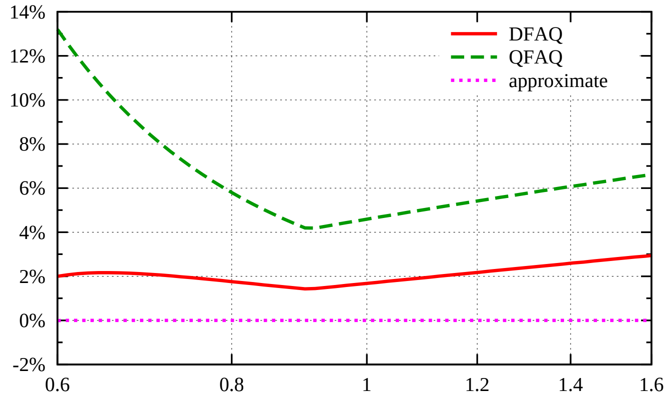

Figure 8: Relative option time value differences for figure 7.

In figures 7 and 9, we show the example of positive skew for the underlying asset and the FX rate with $\beta_S = \beta_Q = \frac{3}{2}$, $T = 2$, and $\rho = -50\%$ and $\rho = 50\%$, respectively. This is followed by a long dated example with positive skew in figure 11.

6.3 Contra-inclining skew

In this case, the implied volatility skew and the FX skew lean in opposite directions.

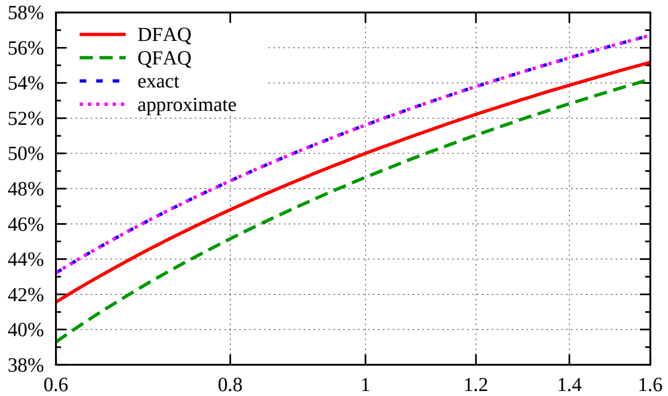

Figure 9: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 51.42\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.08\%$, and $\rho = 50\%$. $F' = 1.1052$ and $\tilde{F} = 1.1163$.

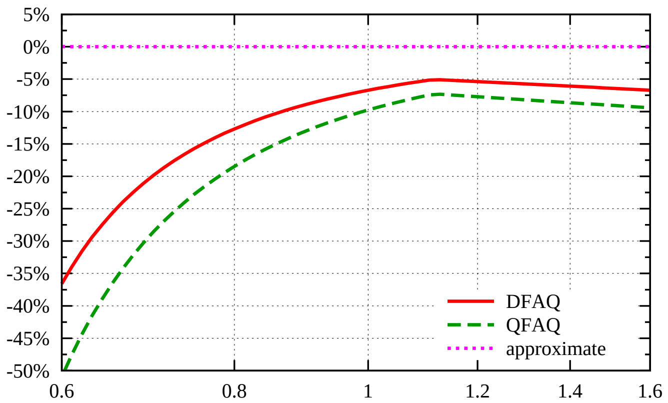

Figure 10: Relative option time value differences for figure 9.

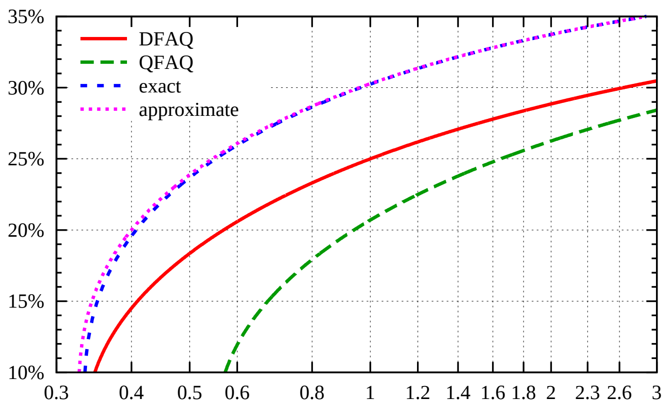

Figure 11: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 27.05\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.96\%$, and $\rho = 50\%$. $F' = 1.6487$ and $\tilde{F} = 2.1471$.

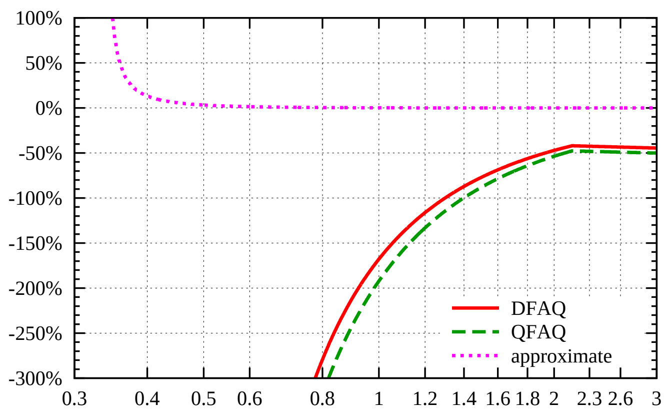

Figure 12: Relative option time value differences for figure 11.

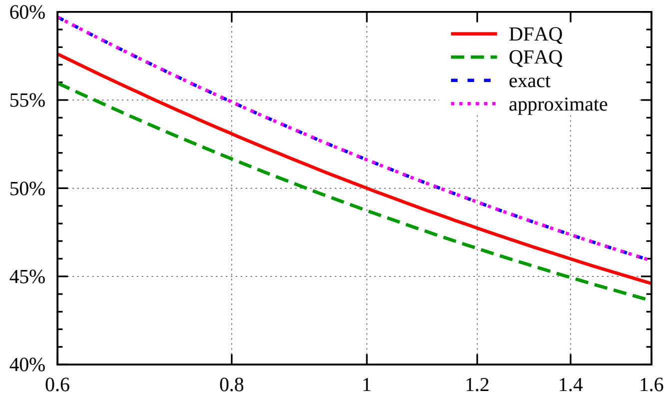

Figure 13: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{1}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 49.22\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.08\%$, and $\rho = -50\%$. $F' = 0.9048$ and $\tilde{F} = 0.9047$.

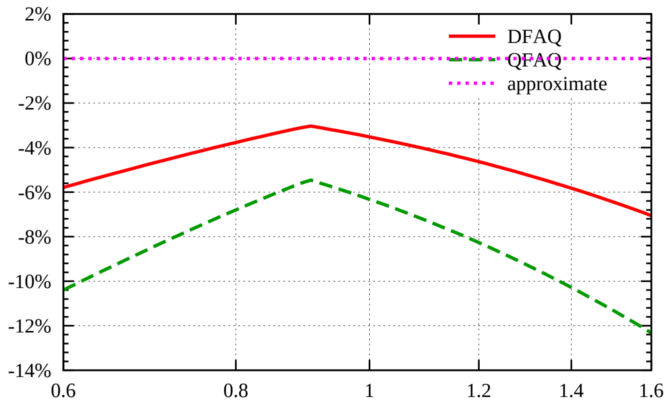

Figure 14: Relative option time value differences for figure 13.

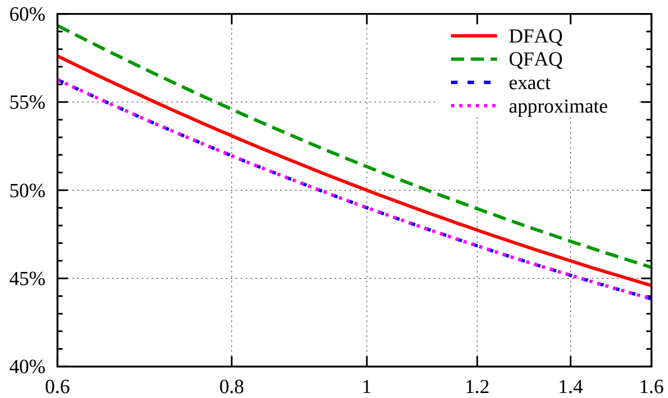

Figure 15: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{1}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 49.22\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.08\%$, and $\rho = 50\%$. $F' = 1.1052$ and $\tilde{F} = 1.1026$.

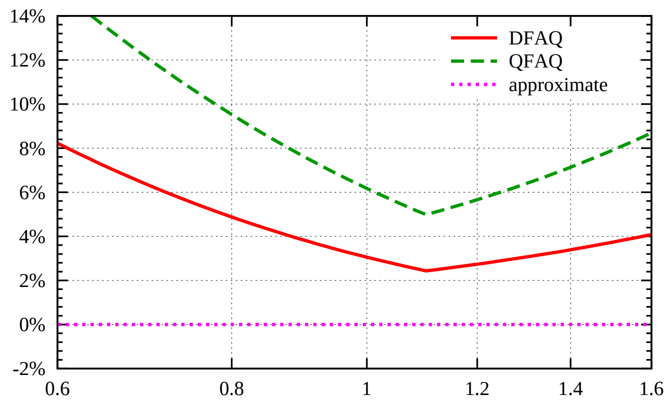

Figure 16: Relative option time value differences for figure 15.

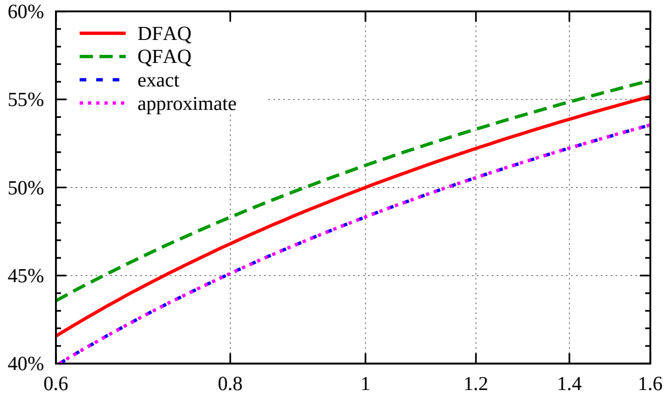

Figure 17: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 51.42\%$, $\beta_Q = \frac{1}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.95\%$, and $\rho = -50\%$. $F' = 0.9048$ and $\tilde{F} = 0.9047$.

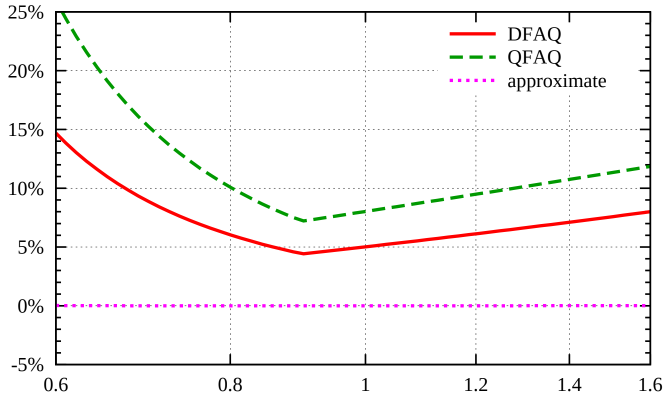

Figure 18: Relative option time value differences for figure 17.

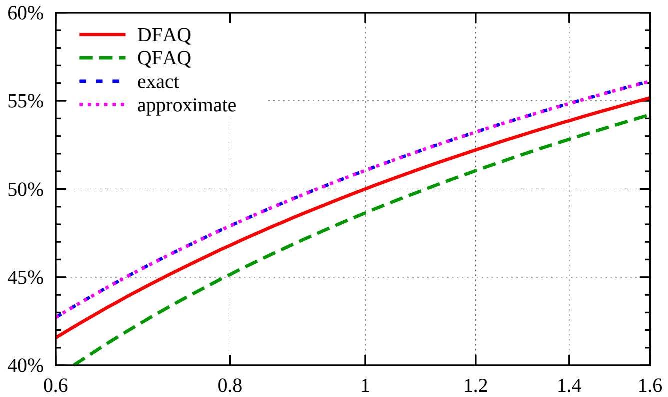

Figure 19: Implied volatilities for $S_0 = 1$, $T = 2$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 50\% \to \sigma_S = 51.42\%$, $\beta_Q = \frac{1}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.95\%$, and $\rho = 50\%$. $F' = 1.1052$ and $\tilde{F} = 1.1066$.

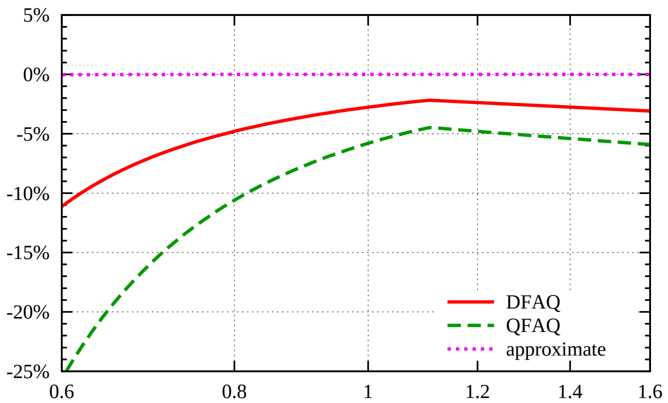

Figure 20: Relative option time value differences for figure 19.

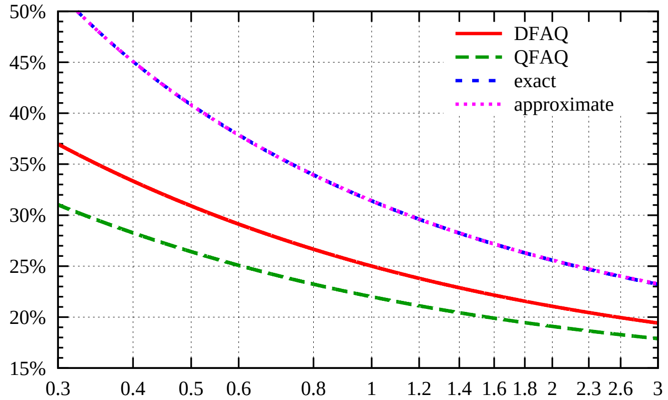

Figure 21: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{1}{2}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 24.04\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.96\%$, and $\rho = -50\%$. $F' = 0.6065$ and $\tilde{F} = 0.5804$.

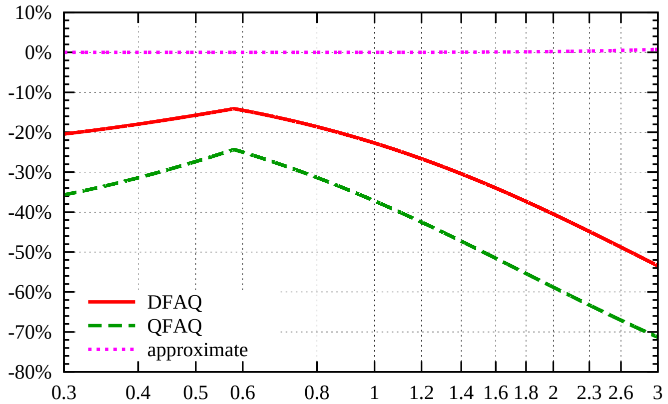

Figure 22: Relative option time value differences for figure 21.

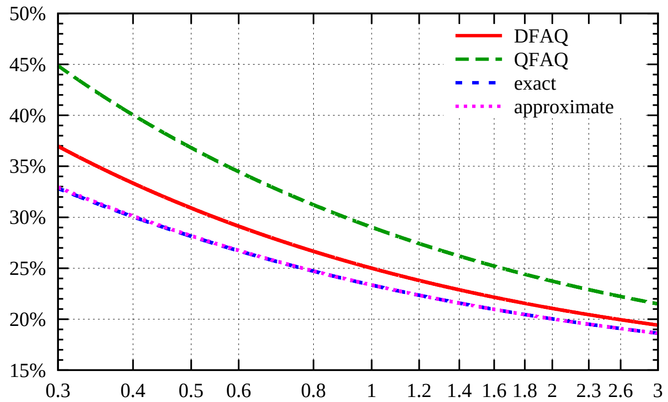

Figure 23: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{1}{2}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 24.04\%$, $\beta_Q = \frac{3}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 20.96\%$, and $\rho = 50\%$. $F' = 1.6487$ and $\tilde{F} = 1.6123$.

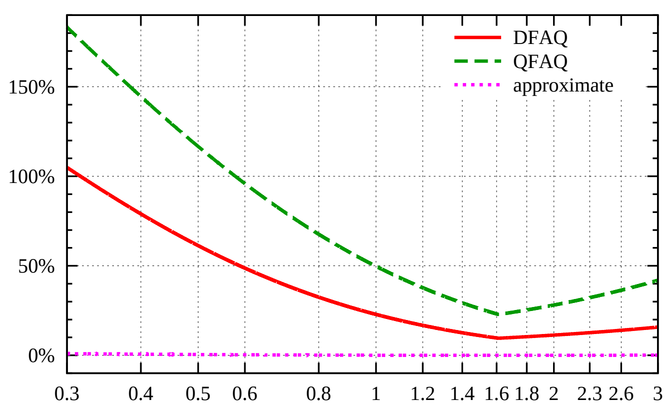

Figure 24: Relative option time value differences for figure 23.

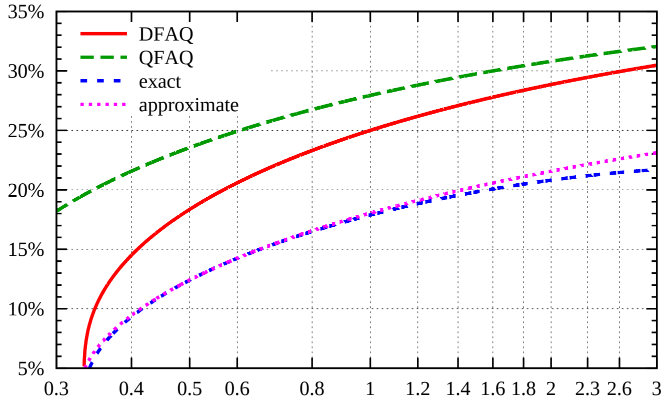

Figure 25: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 27.05\%$, $\beta_Q = \frac{1}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.51\%$, and $\rho = -50\%$. $F' = 0.6065$ and $\tilde{F} = 0.5642$.

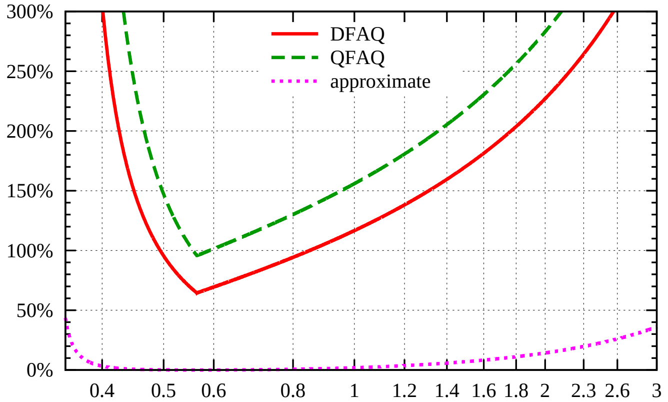

Figure 26: Relative option time value differences for figure 25.

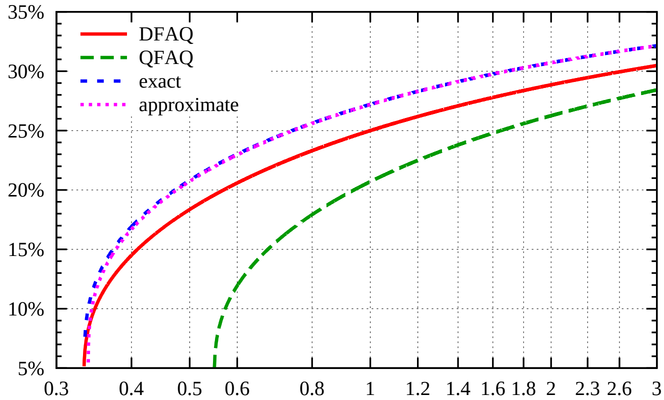

Figure 27: Implied volatilities for $S_0 = 1$, $T = 20$, $\beta_S = \frac{3}{2}$, $\hat{\sigma}_S = 25\% \to \sigma_S = 27.05\%$, $\beta_Q = \frac{1}{2}$, $\hat{\sigma}_Q = 20\% \to \sigma_Q = 19.51\%$, and $\rho = 50\%$. $F' = 1.6487$ and $\tilde{F} = 1.6474$.

Figures 13 to 27, which are for various contra-inclining combinations of $\beta_S < 1$ and $\beta_Q > 1$, and $\beta_S > 1$ and $\beta_Q < 1$, are explained in their respective captions.

7 Conclusion

We have shown that the comparatively simple financial derivative contract known as a "quanto" option requires model-dependent considerations when both the underlying asset and the converting FX rate's option markets exhibit implied volatility skews. This is

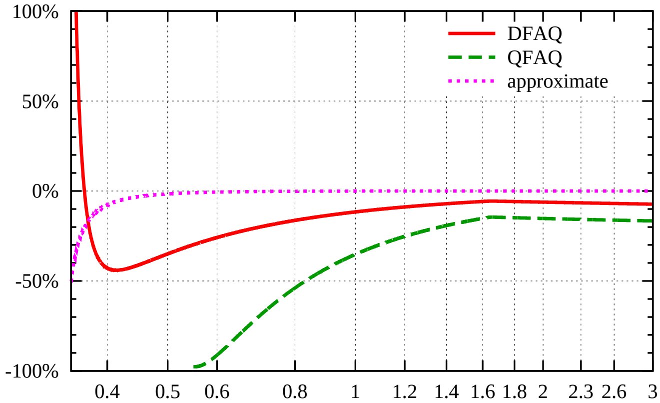

Figure 28: Relative option time value differences for figure 27.

in contrast to the most commonly used practices for the valuation of vanilla quanto options. For the specific example of a double displaced diffusion model, we have given the explicit pricing formula, and derived an analytical approximation that permits to represent the quanto option implied volatility skew in the same parametric setting as the domestic skew, within this model. We have given numerical examples for a variety of parameter combinations that appear to be in support of the practical viability of the approximation. Notwithstanding the fact that the analysis employed a double displaced diffusion model, we emphasize that the main result of the investigation is that quanto options, for anything but short-dated contracts, warrant explicit modelling for accurate pricing and consistent risk managament with respect to the underlying vanilla markets.

References

[CR76] J. C. Cox and S. A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, March 1976.

[Dup94] B. Dupire. Pricing with a Smile. *Risk*, 7(1):18–20, 1994.

[Hes93] S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6:327–343, 1993.

[HKL02] P. Hagan, D. Kumar, and A. S. Lesniewski. Managing Smile Risk. *Wilmott Magazine*, pages 84–108, September 2002.

[Rub83] M. Rubinstein. Displaced diffusion option pricing. *Journal of Finance*, 38:213–217, March 1983.