

Chapter 17

Foreign Exchange, Quantos and Inflation

Three markets that look unrelated on a trading floor are one piece of mathematics. An exchange rate is the price of one numeraire in units of another; a quanto is a payoff settled in the wrong numeraire; and inflation is foreign exchange with the real economy as the foreign country, the consumer price index as the exchange rate and index-linked bonds as the foreign bonds. One adjustment, $-\rho\sigma_1\sigma_2$, does the work in all three, and we derive it once and then spend the chapter recognising it. It ends by explaining a correction on year-on-year inflation swaps that is worth a dozen basis points at the long end and is invisible to anyone who has not made the identification.

17.1 An Exchange Rate Is a Ratio of Numeraires

Chapter 6 defined a numeraire as a strictly positive tradable that other prices are quoted against, and showed that each numeraire carries its own measure in which prices divided by it are martingales. Foreign exchange is what happens when there are two of them and they belong to different countries.

Let X_t be the spot exchange rate, quoted as units of domestic currency per unit of foreign. Let B_t^d and B_t^f be the two savings accounts, growing at the domestic and foreign short rates r_d and r_f .

Calculation 17.1 (The forward exchange rate). A domestic investor can hold the foreign savings account, but only by first converting into it. What they hold, valued in their own currency, is $X_t B_t^f$ — and this is a domestic tradable, so it must have drift r_d under the domestic risk neutral measure. Since B^f grows deterministically at r_f ,

$$\frac{d(X_t B_t^f)}{X_t B_t^f} = \frac{dX_t}{X_t} + r_f dt \stackrel{!}{=} r_d dt + \sigma_X dW_t^d, \quad (17.1)$$

which forces

$$\frac{dX_t}{X_t} = (r_d - r_f) dt + \sigma_X dW_t^d. \quad (17.2)$$

The exchange rate drifts at the interest rate differential. Integrating, the forward rate is

$$F_X(0, T) = X_0 \frac{P_f(0, T)}{P_d(0, T)} = X_0 e^{(r_d - r_f)T}. \quad (17.3)$$

Remark (What this says in words). If foreign interest rates are higher than domestic, the forward exchange rate is below the spot: the currency you earn more interest in is expected — under the pricing measure — to weaken by exactly the extra interest. It has to be, or borrowing domestically, converting, lending abroad and selling the proceeds forward would be a riskless profit. This is covered interest parity, and it is not an economic prediction. It is a no-arbitrage statement, enforced by anyone with access to both money markets.

That this is not an economic prediction deserves emphasis, because the historical record disagrees with it loudly: high-yielding currencies have not on average weakened by their interest differential, which is the entire basis of the carry trade. Chapter 6's distinction applies exactly here. (17.3) is a statement about $\mathbb{Q}$; the carry trade is a bet about $\mathbb{P}$; and the gap between them is a risk premium, not an arbitrage.

Remark (The symmetry, and what breaks it). Nothing above singled out the domestic currency. The foreign investor sees the exchange rate $1/X_t$ and, by the same argument in reverse, finds it drifts at $r_f - r_d$ under *their* risk neutral measure. Both are right. The two measures differ, and $1/X$ under the foreign measure is not the reciprocal of X under the domestic one in any naive sense — by Itô's lemma the reciprocal of a lognormal picks up a $+\sigma_X^2$ in its drift, and that term is exactly the difference between the two measures.

This is Siegel's paradox, and it is the first sign that being careless about which currency one is standing in is expensive. The rest of the chapter is about the cases where it is expensive by an amount someone quotes.

Calculation 17.2 (Options on the exchange rate). Since (17.2) is a geometric Brownian motion with drift $r_d - r_f$, everything in chapter 5 applies with r_f playing the part of a dividend yield: the foreign currency pays interest to whoever holds it, exactly as a stock pays dividends. The price of a call on the exchange rate is

$$C = P_d(0, T) [F_X N(d_1) - K N(d_2)], \quad d_{1,2} = \frac{\ln(F_X/K) \pm \frac{1}{2}\sigma_X^2 T}{\sigma_X \sqrt{T}}, \quad (17.4)$$

which is Black-76 on the forward (17.3). Under its own name this is the Garman-Kohlhagen formula; it is not a new formula, so that everything already established about the Black world — the smile of chapter 9, the dynamics of chapter 10 — transfers without rederivation.

17.2 The Quanto Adjustment

Now the case where the two currencies genuinely interfere.

Definition 17.3 (Quanto). A quanto is a derivative whose payoff is computed from a foreign asset but paid in domestic currency at a fixed exchange rate.

A typical example: a note paying the return on a foreign equity index, in dollars, with no currency exposure. The buyer wants the index and not the currency, and the quanto gives them exactly that. Someone has to price the separation.

Theorem 17.4 (The quanto drift adjustment). *Let S_t be a foreign asset with volatility σ_S , let X_t be the exchange rate (domestic per foreign) with volatility σ_X , and let ρ be their correlation. Then under the domestic risk neutral measure,*

$$\frac{dS_t}{S_t} = (r_f - \rho\sigma_S\sigma_X)dt + \sigma_S dW_t, \quad (17.5)$$

and the quanto forward of S to time T is the ordinary forward multiplied by $e^{-\rho\sigma_S\sigma_X T}$.

Proof. S is a foreign tradable, so under the foreign risk neutral measure $\mathbb{Q}_f$ it drifts at r_f . To value a payoff settled at home we need its drift under $\mathbb{Q}_d$, so we need the change of measure between them, and the Radon-Nikodym derivative is fixed by the requirement that the converted asset $X_t S_t$ be a domestic tradable.

Apply Itô's product rule to XS and impose drift r_d under $\mathbb{Q}_d$. Writing μ_S for the unknown drift of S under $\mathbb{Q}_d$,

$$\frac{d(XS)}{XS} = \frac{dX}{X} + \frac{dS}{S} + \frac{dX}{X} \frac{dS}{S} = (r_d - r_f)dt + \mu_S dt + \rho\sigma_S\sigma_X dt + (\dots)dW,$$

using (17.2) for the first term. Setting the total drift to r_d gives $\mu_S = r_f - \rho\sigma_S\sigma_X$, which is (17.5). Integrating a drift over $[0, T]$ multiplies the forward by its exponential. $\square$

Remark (Why the cross term is the whole story). The proof is three lines and every one of them is the Itô cross term $\frac{dX}{X} \frac{dS}{S}$.

Converting a foreign asset into domestic currency means multiplying two random quantities. If they tend to move together, the product moves more than either — you get more foreign currency exactly when each unit of it is worth more — and that extra covariance is worth something. A domestic investor holding the unhedged foreign asset receives it. A quanto buyer has been given the asset with the currency stripped out, so they do not receive it, and the price has to be lower by exactly that amount. Equation (17.5) is that statement with the accounting done.

Example 17.1 (The size of it). A foreign index with 25% volatility, a currency with 10% volatility, correlation 0.5. The adjustment is

$$-\rho\sigma_S\sigma_X = -0.5 \times 0.25 \times 0.10 = -1.25\% \text{ per year},$$

so the one year quanto forward is 98.76 where the plain forward is 100.

A hundred and twenty five basis points a year is easy to dismiss on a short trade and impossible to dismiss on a long one, because it is a drift and compounds. At five years the forward is six percent below the plain one. This is the reason the adjustment is quoted as a drift rather than as a price: the number that matters is the rate, and the maturity does the rest.

Exercise. A quanto *option*, rather than a forward, on the same asset. Show that its price is the ordinary Black formula with the forward replaced by the quanto forward, and that the volatility is unchanged. (Hint: (17.5) changed the drift and not the diffusion coefficient.) Then explain why this means a quanto option's implied volatility is the same as the plain option's, even though its price is different — and why that makes the adjustment easy to hide in a trading system that only stores implied volatilities.




Figure 17.1: The quanto forward as a fraction of the plain forward, at a 25% asset volatility and a 10% currency volatility. It is monotone through zero and almost linear in the correlation, since the adjustment enters as an exponent and the exponent is small; and it fans out with maturity, because the adjustment is a drift. Nothing is wrong at $\rho = 0$: a payoff settled in a currency that does not move with it costs nothing to redenominate.

Drawn by `quanto.forward`.

17.3 Inflation Is Foreign Exchange

An economy has two ways of measuring value. Nominal terms count currency. Real terms count goods — a basket, whose currency price is the consumer price index I_t . Then:

Foreign exchange	Inflation
domestic currency	nominal economy
foreign currency	real economy
exchange rate X_t	price index I_t
foreign short rate r_f	real short rate r
foreign bond $P_f(t, T)$	index-linked bond $P_r(t, T)$
domestic short rate r_d	nominal short rate n

The index I_t is the price of one unit of the real good in nominal currency, which is precisely what an exchange rate is. And an index-linked bond, which pays a principal scaled by realised inflation, is a bond denominated in goods: it pays one basket at maturity, converted at whatever the index then is. It is the foreign bond.

This identification is the Jarrow-Yildirim model. Everything in the first two sections now transfers with no new work.

Calculation 17.5 (The zero coupon inflation swap). A zero coupon inflation swap pays $I_T/I_0 - (1+b)^T$ at T , and b is quoted as the breakeven inflation rate. Valuing it is (17.3) under a different name: $I_t P_r(t, T)$ is a nominal tradable, so the forward index is

$$\mathbb{E}^T[I_T] = I_0 \frac{P_r(0, T)}{P_n(0, T)}, \quad (17.6)$$

and the breakeven follows,

$$(1+b)^T = \frac{P_r(0, T)}{P_n(0, T)}. \quad (17.7)$$

Remark (There is no model in this). Equation (17.7) contains no volatility, no correlation and no model. The breakeven is a ratio of two observable discount curves, and the trade that enforces it is static: buy the index-linked bond, sell the nominal bond, and hold both to maturity. It is the Fisher equation.

This matters because the breakeven is routinely described as “the market’s inflation expectation”, and it is not one. It is the expectation under the nominal T -forward measure, which differs from the real-world expectation by an inflation risk premium — investors will pay to be protected against inflation, so the breakeven sits above what they actually expect. Chapter 6’s $\mathbb{P}$ against $\mathbb{Q}$ distinction, for the third time.

Example 17.2 (A breakeven). Ten year nominal rates at 3.0% and ten year real rates at 0.5%. The breakeven is

$$\frac{P_r(0, 10)}{P_n(0, 10)} = e^{(0.030-0.005) \times 10} = e^{0.25},$$

so $(1+b)^{10} = e^{0.25}$ and $b = e^{0.025} - 1 = 2.53\%$. The difference of the two continuously compounded rates is 2.5%; expressed as an annually compounded rate it is 2.53%. The three basis point gap is a compounding convention.

17.4 Where the Model Comes Back: Year-on-Year Swaps

The zero coupon swap needed no model. The year-on-year swap, which pays $I_{T_i}/I_{T_{i-1}} - 1$ on each of a series of dates, needs one.

Calculation 17.6 (Why the naive answer is wrong). The obvious guess is that a year-on-year payment is the ratio of two zero coupon forwards. It is not, and conditioning shows where it breaks. Writing $S = T_{i-1}$ and $T = T_i$, and taking the expectation under the nominal T -forward measure in which the payment is valued,

$$\begin{aligned}\mathbb{E}^T\left[\frac{I_T}{I_S}\right] &= \mathbb{E}^T\left[\frac{1}{I_S}\mathbb{E}^T[I_T \mid \mathcal{F}_S]\right] \\ &= \mathbb{E}^T\left[\frac{1}{I_S} \cdot I_S \frac{P_r(S, T)}{P_n(S, T)}\right] \\ &= \mathbb{E}^T\left[\frac{P_r(S, T)}{P_n(S, T)}\right],\end{aligned}\tag{17.8}$$

where the second line is (17.6) applied at time S instead of time 0, which is legitimate because $I_t P_r(t, T)/P_n(t, T)$ is a T -forward martingale.

The index has cancelled out entirely. What is left is the expectation of a *future* bond price ratio — and unlike a forward, that is not pinned by any static portfolio. Bond prices at S depend on where rates are at S , and (17.8) therefore depends on how they get there. A model is unavoidable.

Remark (The general shape of this). The reader who has been through chapters 7 and 8 has met this twice already. A futures contract differs from a forward because it is settled under the risk neutral rather than the forward measure. A CMS payment differs from a swap rate because it is paid on the wrong date. Each time, an instrument that looks like a forward turns out to reference a rate under a measure in which that rate is not a martingale, and the gap is a convexity correction.

(17.8) is the same phenomenon and will produce the same kind of answer. What is new here is that part of the correction will turn out to be a quanto adjustment, which the other two cases have no analogue of, because they involve only one currency.

We now do the calculation. Take both short rates to be Hull-White as in chapter 8, with mean reversions a_n, a_r and volatilities σ_n, σ_r , the index lognormal with volatility σ_I , and correlations ρ_{nr} between the two rates and ρ_{rI} between the real rate and the index.

Calculation 17.7 (The real rate under the nominal measure). First, the piece that has no counterpart in an ordinary rates model. The real short rate is a foreign short rate, so its drift under the nominal measure is not the drift its own model was written with.

The argument is Theorem 17.4 again. $I_t P_r(t, T)$ must be a nominal tradable, and expanding by Itô's product rule as before,

$$\frac{d(IP_r)}{IP_r} = \underbrace{(n - r)dt}_{\text{index}} + \underbrace{\mu_{P_r}dt + (\cdots)dW}_{\text{real bond}} + \underbrace{\rho_{rI}\sigma_I(-B_r(t, T)\sigma_r)dt}_{\text{cross term}} \stackrel{!}{=} n dt + (\cdots)dW,$$

where $-B_r(t, T)\sigma_r$ is the real bond's volatility in the Hull-White model, $B_r(t, T) = (1 - e^{-a_r(T-t)})/a_r$. Solving gives $\mu_{P_r} = r + \rho_{rI}\sigma_I\sigma_r B_r$, which is the real bond drifting faster than r under the nominal measure. Translated back to the short rate, this is

$$dr_t = (\theta_r(t) - a_r r_t - \rho_{rI}\sigma_r\sigma_I)dt + \sigma_r dW_t^{\mathbb{Q}_n}.\tag{17.9}$$

The extra term is exactly (17.5), with the real rate in the place of the foreign asset and the price index in the place of the exchange rate.

Calculation 17.8 (The convexity). With both rates Gaussian, (17.8) is an expectation of an exponential of a normal and can be done in closed form. Writing y^n, y^r for each rate's deviation from its own initial forward curve and B_n, B_r for the two Hull-White factors over $[S, T]$, the bond reconstitution formula of chapter 8 gives

$$\frac{P_r(S, T)}{P_n(S, T)} = \underbrace{\frac{P_r(0, T) P_n(0, S)}{P_r(0, S) P_n(0, T)}}_{\text{the naive answer}} \cdot e^{D_r - D_n} \cdot e^{-B_r y_S^r + B_n y_S^n},$$

with D_i the deterministic terms. Taking the expectation of the lognormal and using that $D_i = -\frac{1}{2} B_i^2 V_i$ exactly cancels half of each variance, the whole correction collapses to

$$C = \exp \left(-B_r m_r + B_n m_n + B_n^2 V_n - B_r B_n \text{Cov}(y_S^n, y_S^r) \right), \quad (17.10)$$

where m_r, m_n are the two rates' means at S under the nominal T -forward measure and V_n is the nominal variance. The year-on-year payment is the naive ratio of forwards multiplied by C .

Remark (Reading (17.10)). The formula is opaque and the content is not. Every term is a mean or a covariance under the nominal T -forward measure, and the drifts that produce those means are three, listed in [Inflation::real_mean](#):

- the Hull-White term each model carries to stay fitted to its own initial curve, which would be there in a single-currency world and contributes nothing interesting;
- $-\rho_{rI} \sigma_r \sigma_I$, the quanto adjustment (17.9), present only because the real economy is a foreign one;
- $-\rho_{nr} \sigma_n \sigma_r B_n(u, T)$, from changing to the nominal forward measure, which is the ordinary convexity of chapters 7 and 8.

So the correction has two parents. Setting $\rho_{rI} = 0$ — or $\sigma_I = 0$ — deletes the quanto term and leaves the rest untouched, and that is a sharp enough statement to check numerically, which [quant/src/crosscurrency.rs](#) does.

Example 17.3 (Is it worth anything). At the calibration in the figure — a 1.0% nominal volatility, 0.7% real, 1.2% index volatility, $\rho_{nr} = 0.6$, $\rho_{rI} = -0.3$ — and against a flat 2.53% breakeven:

Period starts	Total correction	Quanto part
1y	-0.3 bp	-0.3 bp
5y	-1.8 bp	-1.1 bp
10y	-4.2 bp	-2.0 bp
20y	-8.8 bp	-3.2 bp
30y	-12.4 bp	-3.9 bp

A basis point is nothing at the front end and a dozen of them at thirty years is not: on a long dated year-on-year book it is the difference between a position that makes money and one that does not.




Figure 17.2: The correction a one year inflation payment needs, against the zero coupon curve, as a function of when the year starts. It grows with the start date because every term carries the accumulated effect of the measure changes up to that point, and the two components grow differently: the forward measure term is quadratic-ish in the start while the quanto term is closer to linear, so the quanto share falls from most of the correction at the front end to about a third of it at thirty years. Drag the correlation to zero and only one curve moves.

Drawn by `Inflation::yoy_convexity` and `Inflation::real_mean`.

Remark (Trusting this). Equation (17.10) is exactly the kind of result that is easy to get wrong and hard to notice being wrong: a sign error inside it changes a number by a few basis points and nothing visibly breaks. So the implementation carries a second, independent route to the same answer — a Monte Carlo that steps both rates forward under (17.9) and averages the bond price ratio it finds, sharing none of the algebra above. The two agree across a range of correlations and volatilities to well inside the simulation error.

17.5 What Carries Over

Two things carry over from this chapter beyond the formulas.

The first is that a change of numeraire is a change of country, and that the cost of moving between them is always the same object: the covariance between the thing being valued and the exchange rate connecting the two numeraires. Written as $-\rho\sigma_S\sigma_X$ it looks like a quanto convention. It is not; it is the Itô cross term, and it appears wherever a quantity is valued in a measure other than the one it was defined in.

The second is that recognising two markets as the same market is worth more than any individual result derived here. The inflation desk and the foreign exchange desk sit on different floors and quote in different units, and (17.9) is invisible from inside either one. It becomes obvious — almost forced — the moment the index is seen as an exchange rate. Chapter 18 does the same thing again, for correlation.

References

- Garman, M. B., & Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of International Money and Finance*, 2(3), 231–237.
- Jarrow, R., & Yildirim, Y. (2003). Pricing treasury inflation protected securities and related derivatives using an HJM model. *Journal of Financial and Quantitative Analysis*, 38(2), 337–358.
- Brigo, D., & Mercurio, F. (2006). *Interest Rate Models — Theory and Practice*, 2nd ed. Springer.