How to value a coco

Converting default risk into conversion risk provides a method for valuing contingent convertibles, according to Patrick Cheridito and Zhikai Xu

Since the financial crisis of 2007-2009, contingent convertible bonds (cocos) have been issued by a number of financial institutions, including Lloyds Banking Group, Rabobank, Credit Suisse, Bank of Cyprus, Australia and New Zealand Banking Group, UBS, Zurcher Kantonalbank and Macquarie.

A typical coco starts life as a standard corporate bond promising to make regular coupon payments and redeem the principal amount at maturity. But if before maturity a specified trigger event occurs, the coco converts into something else. A general coco is characterised by its principal amount F; maturity T, coupon dates (t_i, $i = 1,\ldots,n$) and annualised coupon rate c; and its trigger mechanism and conversion mechanism.

Most existing cocos are triggered if some accounting measure of capital adequacy falls below a critical level, and they convert into a fixed number G of the issuing firm's stock shares. Some have an additional regulatory trigger or an option for the holders to convert the coco early. Usually, a coco offers an attractive coupon rate as a compensation for the possibility of conversion. But to fully understand it, all relevant sources of risk have to be taken into account.

Three relevant sources of risk

A coco is a typical hybrid product in that it depends on various sources of risk. A standard coco is exposed to interest rate risk, conversion risk and equity risk. To be useful for pricing and hedging, a coco model should include all these types of risk. A general model that can be specified in different ways, such as, for example, a structural or reduced form model, is developed in Cheridito and Xu (2013a). A reduced form model with deterministic conversion intensity is studied in Cheridito and Xu (2013b). To illustrate the most important points we here use a simplified version of the latter model.

We suppose that there exists a risk-neutral probability measure and denote by:

τ – the conversion time

$Q(t)$ – the risk-neutral probability of $\tau > t$

$q(t)$ – the risk-neutral density of τ

$Z(t)$ – the price of a risk-free zero coupon bond paying \$1 at time t

S_t – the price of the issuing firm's stock at time t.

The value of a coco can be decomposed into three parts: the value of future coupon payments, the value of the redemption of the principal in case the coco does not convert and the value of a possible conversion. If the trigger event is assumed to be independent of future interest rates, and accrued interest payments are neglected, the first two components can be written as:

(a) $c\sum_{i=1}^{n}(t_i - t_{i-1})Z(t_i)Q(t_i)$

(b) $FZ(T)Q(T)$

The independence assumption is not unreasonable for an accounting trigger since a firm's capital adequacy is not directly related to the prevailing risk-free rate. But it can be relaxed by introducing forward measures. Usually, a coco also pays accrued interest if conversion occurs between two coupon dates. This can easily be covered. But it only has a minor influence on the total value of a coco.

The third value component is worth more attention, since it is specific to cocos. If conversion occurs before T, each share of coco immediately converts into G stock shares. So the value of conversion is

$$GE^{Q}[\tilde{S}_\tau 1_{\{\tau\le T\}}],$$

where E^Q denotes the expectation under the risk-neutral probability measure, $\tilde{S}_\tau$ is the discounted stock price at time τ and $1_{\{\tau\le T\}}$ is the indicator function of the event $\tau\le T$. If the stock pays no dividends, its discounted price is a martingale with respect to the risk-neutral probability measure, and it follows that the conversion value can be expressed as

(c) $GS_0(1-Q^*(T))$,

where $Q^*(T)$ is the probability of $\tau>T$ with respect to the measure obtained from Q by tilting it with $\tilde{S}_\tau/S_0$. The total value of a coco is the sum of (a), (b) and (c). Dividend payments can easily be covered by adjusting the formula for (c). But, like accrued interest payments, this does not affect the overall price of a coco much.

Model calibration

For a coco model to be applicable in practice it has to be calibrated to market quotes of related instruments. S_0 can be observed in the market, and $Z(t)$ can (up to interpolation) be deduced from the term structure of risk-free yields at time 0. It remains to specify $Q(t)$ and $Q^*(T)$.

Since conversion risk is related to default risk, $Q(t)$

can be estimated from market quotes of CDS spreads. $Q^*(T)$ depends on the specification of the discounted stock price. In a reduced form model, conversion comes as a surprise. Therefore, prices of the coco, stock and CDS contracts are expected to jump by a quantity we call jump-to-conversion (JTC). The JTC of a CDS can easily be calculated (see below). But the JTC of the coco depends on how one models the JTC of the issuing firm's stock. In reality it will depend on the specifics of the coco contract as well as the situation the firm will find itself in upon conversion.

There are two competing balance sheet effects if a coco converts. On the one hand, the liabilities of future coupon payments and redemption of the principal disappear. On the other, the value of stock shares is diluted since they increase in number. In addition, the mere fact that a coco is triggered can have informational effects on the stock price. As a benchmark case, let us here assume that the stock price is continuous. Then, with a deterministic conversion intensity, $Q^*(T)$ is equal to $Q(T)$.

CDS contracts exist with large maturities and are liquidly traded. For the model to be able to price them, one has to add scenarios in which the issuing firm defaults. Let us introduce the following:

θ – the default time
$P(t)$ – the risk-neutral probability of $\theta > t$
$p(t)$ – the risk-neutral density of θ
s_j, $i = 1,\dots,m$ – coupon dates of a CDS
δ – annualised coupon rate of the CDS
R – recovery rate in case of default.

If default is assumed to happen independently of interest rates, and accrued interest payments are neglected, the model prices a protection buyer position in the CDS as $PL - CL$, where

$$PL = (1-R)\int_0^{s_m} Z(t)p(t)\,dt\,,$$

is the value of the protection leg and

$$CL = \delta\sum_{i=1}^{m}(s_i - s_{i-1})\,Z(s_i)P(s_i)\,,$$

the value of the coupon leg. The conversion time τ and default time θ can jointly be modelled by augmenting the credit risk framework of O'Kane and Turnbull (2003) with a second random time arriving with a deterministic intensity. More precisely, let N_t be a counting process with deterministic jump intensity λ_t; that is, N_t starts from 0, is piece-wise constant, and the probability that it jumps up by 1 in a small time interval $[t,t+h]$ is approximately equal to λ_th. Now assume the trigger event τ happens at the first jump time of N_t and default θ at the second jump time. Then default cannot occur before the coco has been triggered, and the densities and q and p are given by

$$q(t) = \lambda_t e^{-\int_0^t \lambda_s\,ds} \quad\text{and}\quad p(t) = \lambda_t \int_0^t \lambda_s\,ds\, e^{-\int_0^t \lambda_s\,ds}\,.$$

If different CDS contracts are liquidly traded, one can choose the function λ_t so that they are priced correctly by the model. For instance, if one assumes λ_t to be constant between the tenors of the CDS contracts, it can uniquely be inferred from CDS spreads with the method of O'Kane and Turnbull (2003). This gives the risk-neutral density $q(t)$ and, therefore, the price of the coco. However, the approach needs the recovery rate R as input. Either it is estimated from time series data, or it is chosen such that the model prices the coco at its market value. The second method yields a coco-implied recovery rate.

Hedging

A coco with an accounting trigger can be hedged with the stock of the issuing firm, CDS contracts and interest rate swaps. To hedge a short position one tries to reproduce the sensitivities of the coco to the underlying risk factors. A long position is hedged with the opposite strategy. At time 0, equity risk is offset by buying $G(1-Q^*(T))$, stock shares.

Under the assumption that the stock price does not jump, the JTC of the coco at time 0 is $G\,S_0 - C_0$ – where C_0 is the coco price at time 0. If conversion were to happen immediately, the default time θ becomes the next jump time of N_t. So its density changes from $p(t)$ to $q(t)$, and the value of a protection buyer position in the CDS jumps to

$$(1-R)\int_0^{s_m} Z(t)q(t)\,dt - \delta\sum_{i=1}^{m}(s_i - s_{i-1})\,Z(s_i)Q(s_i)\,.$$

Hence the coco's JTC can be neutralised with an appropriate CDS investment. Finally, one immunises the portfolio against movements of the risk-free yield curve by investing in fixed income products. Since the sensitivities change over time, the hedging portfolio has to be updated continuously. To keep it self-financed, one balances it out by investing excess funds in a money market account.

Conclusion

This article discusses the pricing, calibration and hedging of cocos in a reduced form model. Many of the simplifying assumptions can be removed. For instance, it is easy to add dividends and accrued interest payments. One can also allow the stock price to jump at conversion or let the conversion intensity be stochastic.

Patrick Cheridito is a member of the program in applied and computational mathematics at Princeton.

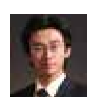

Zhikai Xu is a PhD candidate in operations research and financial engineering at Princeton University.