

Unlocking the language of structured securities



“The limits of my language,” wrote the philosopher Ludwig Wittgenstein, “mean the limits of my world.”

The languages of today’s complex financial markets often consist not simply of words and numbers but also of technical systems. The credit crisis has shown the importance of their powers – and limits.

Although few outsiders have heard of it, the single most important language of mortgage-backed securities and similar products is a system called Intex. It includes a computer language for defining deals’ intricate cash flow rules, a graphics-based tool for designing deals, and a truly remarkable computerised “library” of the parameters of the underlying asset pools and the cash flow rules of more than 20,000 deals. Intex is not cheap – one user told me his bank pays about \$1.5m a year for it – and it has competitors such as Bloomberg, but it is essential for all serious participants in structured securities.

Intex’s power as a language is to make instruments such as mortgage-backed securities mentally tractable. I confess I’ve always found them daunting. The rules governing a deal can occupy hundreds of pages of impenetrable legal prose, and the economic value of the deal’s tranches depends on three complex characteristics of the underlying mortgage pool: the rate at which borrowers prepay (redeem their mortgages early), their propensity to

default, and the loss severity (the proportion of the debt that cannot be recovered if a borrower defaults).

In July a friendly banker showed me Intex in action. He chose a particular mortgage-backed security, entered its price and a figure for each of prepayment speed, default rate, and loss severity. In less than 30 seconds, back came not just the yield of the security, but the month-by-month future interest payments and principal repayments, including whether and when shortfalls and losses would be incurred. The psychological effect was striking: for the first time, I felt I could understand mortgage-backed securities.

Of course, my new-found confidence was spurious. The reliability of Intex’s output depends entirely on the validity of the user’s assumptions about prepayment, default and severity. Nevertheless, it is interesting to speculate whether some of the pre-crisis vogue for mortgage-backed securities resulted from having a system that enabled neophytes such as myself to feel they understood them. Certainly, like any language, Intex aided communication. If you were planning a mortgage-backed deal, you could construct an Intex file, make it available to potential investors, and use it to discuss the deal’s features, modify those features, and gauge investors’ interest.

The limits of the language came when mortgage-backed securities were repackaged into collateralised debt obligations (CDOs), complex debt securities based on pools of other assets. You could still run Intex, first for each of the securities and then for the CDO, but it could be a slow process. Often, CDOs included not just mortgage-backed securities, but tranches of other CDOs, each maybe incorporating further CDOs. This multiplied enormously the number of underlying mortgage pools, causing a single valuation run to take hours. (On occasion, each of a pair of CDOs would buy a tranche of the other, creating a “loop” that slowed analysis). Sometimes, users did little more than one run using the prepayment, default and severity rates judged most likely. Those (such as the rating agencies) that needed to do

more nearly all took a fatal shortcut. Instead of analysing CDOs from the bottom (the underlying pools of mortgages) up, they shifted to a different mathematical language, which treated a CDO’s components (mortgage-backed securities and tranches of other CDOs), in effect, as if they were corporate bonds, with their properties inferred from their ratings. This often led to serious underestimation, especially by rating agencies, of correlation among these components.

The one bank I’ve found that did analyse CDOs based on mortgages from the bottom up was Goldman Sachs. It developed its own analytical techniques, and used a large “computer farm” in New Jersey to spread the analysis over multiple machines, so keeping the time each run took tolerable. Goldman’s Abacus CDOs – one of which was the flashpoint of the SEC’s recent investigations

– were apparently analysed this way.

A bottom-up analysis was – and still is – expensive. It requires clever quantitative analysts, multiple computers and software developers able to “parallelise” a program so it runs efficiently on many machines at once. Nevertheless, if going beyond the limits of existing languages in this way helped Goldman take the crucial late-2006 decision to liquidate or hedge its positions in mortgage-backed securities (enabling it to survive the crisis almost unscathed), it was well worth it. I hope its competitors have learned the lesson: a limited language means a dangerously limited world.

Professor Donald MacKenzie teaches sociology at The University of Edinburgh

The limits of the language came when mortgage-backed securities were repackaged into collateralised debt obligations