THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

‘A device for being able to book P&L’

Citation for published version:

MacKenzie, D & Spears, T 2014, "A device for being able to book P&L: The organizational embedding of the Gaussian copula", *Social Studies of Science*, vol. 44, no. 3, pp. 418-440.
<https://doi.org/10.1177/0306312713517158>

Digital Object Identifier (DOI):

[10.1177/0306312713517158](https://doi.org/10.1177/0306312713517158)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Social Studies of Science

Publisher Rights Statement:

The final version of this paper has been published in Social Studies of Science, 44/3, June/2014 by SAGE Publications Ltd, All rights reserved. © Donald MacKenzie and Taylor Spears, 2014. It is available at:
<https://journals.sagepub.com/doi/10.1177/0306312713517158>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

‘A device for being able to book P&L’: The organizational embedding of the Gaussian copula

Donald MacKenzie and Taylor Spears

School of Social & Political Science, University of Edinburgh, Edinburgh, Scotland

Abstract

This paper, the second of two articles on the Gaussian copula family of models, discusses the attitude of ‘quants’ (modellers) to these models, showing that, contrary to some accounts, those quants were not ‘model dopes’ who uncritically accepted the outputs of the models. Although sometimes highly critical of Gaussian copulas – even ‘othering’ them as not really being models – they nevertheless nearly all kept using them, an outcome we explain with reference to the embedding of these models in inter- and intra-organizational processes: communication, risk control and especially the setting of bonuses. The article also examines the role of Gaussian copula models in the 2007-08 global crisis and in a 2005 episode known as ‘the correlation crisis’. We end with the speculation that all widely-used derivatives models (and indeed the evaluation culture in which they are embedded) help to generate inter-organizational co-ordination, and all that is special in this respect about the Gaussian copula is that its status as ‘other’ makes this role evident.

Keywords

Gaussian copula, financial modelling, investment banking, finance, performativity

Corresponding author:

Donald MacKenzie, School of Social & Political Science, University of Edinburgh, Chrystal Macmillan Building, Edinburgh, EH8 9LD, Scotland.
emails: frances.j.burgess@ed.ac.uk

In a companion article (MacKenzie and Spears, 2014) we examine the development of Gaussian copula models, used in finance to model Collateralized Debt Obligations or CDOs, which are securities based on pools of assets such as corporate bonds. In this article, we discuss how Gaussian copula models became embedded in organizational practices in one of the two main contexts in which they are used, investment banking. The other main context of use, credit rating agencies, is discussed in MacKenzie (2011). Here, we examine the role the Gaussian copula played in two financial crises: the little-known 2005 episode that participants called ‘the correlation crisis’, and the wider credit crisis that erupted in summer 2007 and led to the near-collapse of the global banking system in autumn 2008. We end with a discussion of the ‘counter-performative’ role of financial models, in which their empirical accuracy is undermined through their practical use in market processes.

Our article bears on three themes found in the small but growing literature known as the social studies of finance, in which perspectives from disciplines such as anthropology, sociology and science and technology studies are applied to financial markets. The first theme is the attitude taken by participants to models. In media discussion, market participants are often portrayed as unthinkingly accepting the outputs of a model, or as ‘model dopes’, following Garfinkel’s (1967) ‘cultural dope’. However, it is far from clear that model dopes exist; research in the social studies of finance has failed to find empirical evidence of them. Mars (1998; see Svetlova, 2012) shows how securities analysts’ judgements of the value of shares are not driven by spreadsheet models; rather, they adjust the inputs into these models to fit their ‘feel’ for the ‘story’ about the corporation in question. Svetlova (2009, 2012) finds similar flexibility in how models are used; they are ‘creative resources’ rather than rules that unambiguously determine action. Beunza and Stark (2012: 413) find the traders they study to

be ‘intelligent, creative, thoughtful and independently minded’, fully conscious that the models they use could be wrong. Indeed, traders employ models ‘to gain cognitive distance’, practising ‘reflexive modelling’, in which they use models to infer others’ beliefs from patterns of prices and compare those beliefs with their own (Beunza and Stark, 2012: 411).

Those we interviewed were also not model dopes. Indeed, even in interviews conducted before the credit crisis, we found considerable hostility to Gaussian copula models. They were considered to be flawed models, possibly not even worthy of the term ‘model’. Such criticism was voiced even by those who had made important technical contributions to the development of Gaussian copula models, and who were still using them. This apparent paradox – sophisticated, sceptical participants continuing to employ models they disliked, even when alternatives were available – was important methodologically to our research, because it prompted interviewees to tell us why they felt compelled to act in this way, and in so doing provided us with our first clues to the organizational embedding of Gaussian copula models in investment banking.

The second theme in the social studies of finance literature upon which we build concerns what Muniesa et al. (2011: 1189) call ‘the description of financial objects’. This is ‘the problem of constructing robust, flexible, portable, and mutually compatible depictions of complex, multisided, and often ambiguous financial objects (products, trades, marketplaces)’ (p. 1189). This theme is explored in greater ethnographic depth in Lépinay’s (2011) participant-observation study of a leading investment bank. Such banks are complex organizations with multiple parts, including prestigious ‘front office’ activities (such as sales, trading, and modelling by the ‘front office’ quants who support trading); ‘middle office’ functions (including accounting and risk control); and ‘back office’ tasks such as trade

processing, clearing and settlement. ‘Quants’ is the vernacular, in the world of finance, for modellers.

As Lépinay (2011) emphasizes, the different parts of an investment bank often employ different ‘languages’ to capture the characteristics of financial objects most relevant to them. However, complete heterogeneity is not attractive to banks. If nothing else, it might require slow, expensive manual recoding of the characteristics of a product or trade at each stage of its processing. As we discuss below, the choice of models by a bank’s front-office traders and quants is constrained by the models employed by the bank’s accounting and risk control divisions. If the front-office quants use a model that diverges too far from that used by accountants, they imperil the granting of ‘Day 1 P&L’, in which the present value of the anticipated future income of a trade is credited to the trader at the time at which the trade is entered into (‘P&L’ is the acronym of profit and loss). At the same time, if the front-office traders and quants use a model that diverges too radically from that used by a bank’s risk controllers, their capacity to trade is threatened. As noted by Lépinay (2011), and as we discuss in our companion article (MacKenzie and Spears, 2014), there is a strong emphasis in the derivatives departments of banks on hedging, but whether or not a trade is seen as ‘properly hedged’ depends on the model used to calculate hedging ratios. If a trader’s hedges differ too much from those calculated by risk controllers, then the latter are likely to view the trader’s positions as unduly risky.

Intra-organizational matters such as these articulate with inter-organizational issues. A model used by multiple organizations – as the Gaussian copula was – has at least three advantages. First, the very fact that a model is used widely can make it a good predictor of price movements, a point to which we return shortly. Second, a widespread model can be a

medium of communication among organizations. Third, those who are not specialists in modelling – accountants, for example, often are not – are more likely to regard a widely used model favourably as opposed to one that appears idiosyncratic.

Organizational issues are prominent in the history of the Gaussian copula. While the emphasis in our companion article is on culture, this paper focuses, in part, on how ‘organization’ can trump ‘culture’. It is a tale of models that were widely judged inadequate, in particular by the standards of the ‘locally’ hegemonic evaluation culture – no-arbitrage modelling (MacKenzie and Spears, 2014). Nevertheless, the models were and still are used, because the organizational costs of abandoning them are too high. This does not, however, imply that evaluation cultures are unimportant. The costs of abandoning Gaussian copula models were and are in good part to do with the patterns of co-ordinated behaviour that had arisen around the models. Any meaningful concept of ‘culture’, we posit, must view it as a form of and a resource for co-ordinated action,¹ and this, we suggest, is the case for evaluation cultures in finance; precisely *because* such cultures cross-cut organizations, they facilitate communication and explicit or implicit co-ordination amongst organizations.

The third theme we develop from the social studies of finance literature pertains to the ‘performative’ aspect of models, which is understood as the way in which the models’ use alters, or even brings into being, the phenomena they model (Callon, 1998). The performativity of the Gaussian copula differs from that of the canonical Black-Scholes options model (MacKenzie, 2006). For roughly a decade, from the mid 1970s to the 1987 stock market crash, the practical use of Black-Scholes – especially by traders in performing arbitrage, the low risk or riskless exploitation of discrepancies in patterns of prices – had strongly performative effects, shifting patterns of prices towards the postulates of the model.

In effect, traders using Black-Scholes saw discrepancies between that model and patterns of option prices as profit opportunities, and their exploitation of those opportunities caused the discrepancies to diminish.

In contrast, there *always* was a systematic discrepancy between the Gaussian copula and patterns of market prices, and the main mechanism of the model's performativity did not involve arbitrage. As we discuss in more detail below, the discrepancy, known as the 'correlation skew', was seen as evidence of inadequacy of the Gaussian copula, and not as an arbitrage opportunity from which one could reliably profit. Black-Scholes could have a performative effect even when only a minority of traders were using it for arbitrage and most traders either had not heard of it or disagreed with it. The performative effects of the Gaussian copula, in contrast, involved implicit consensus. Because of the correlation skew, traders and quants using Gaussian copula models had to adjust them so that they fitted patterns of market prices. There was a dominant way of doing this (the 'base correlation' approach described in our companion article), and those employing it chose similar values for the models' parameters.² This implicit inter-organizational co-ordination meant that (so long as the consensus did not change) those models could be relied upon to continue to be a good fit to patterns of market prices, thus reinforcing their canonical role. That process broke down during the credit crisis, but how it broke down differed from how Black-Scholes eventually broke down; the mechanism of the counter-performativity of the Gaussian copula differed from that of Black-Scholes.

Our account of Gaussian copula models employs two main data sources: documents, especially the technical literature and specialist trade press, and 114 predominantly oral-history interviews. Twenty-nine interviews conducted with quants are particularly important.

This paper has several sections. First, we examine participants' criticisms of Gaussian copula models, before discussing how such models became organizationally embedded in the investment banking. The second section examines the role of Gaussian copula models in communication and the third outlines their role in the remuneration of traders and in risk control. The fourth and fifth sections explore the models' role in the correlation crisis and credit crisis. We conclude with a sixth section that considers the interplay between the organizationally embedded uses of Gaussian copula models, the counter-performativity of these and other models, and the economic crisis.

Criticisms of the Gaussian copula

We did not intend our research to focus on participants' attitudes to the validity of Gaussian copula models. Our initial focus was how 'correlation' had come to be reified and rendered 'tradable' in 'index' markets (MacKenzie and Spears, 2014). Nevertheless, in five of the eight interviews we conducted with quants prior to the credit crisis, interviewees expressed their views on the adequacy of Gaussian copula models. Because of the risk of participants inflecting their views of Gaussian copula models with hindsight of their role in the economic crises, we focus our analysis on the five pre-crisis interviews, discussing the more recent interviews more briefly.³

The closest to an explicit defence of the Gaussian copula was voiced by the most junior of the five, a young quant working for a hedge fund. 'You can't beat the Gaussian distribution in terms of its flexibility ... analytical tractability and ... computational efficiency', he said, also noting the Gaussian copula's role, discussed below, in facilitating communication between organizations: 'I think there will always be a place for the Gaussian copula.' Even he, though, acknowledged that '[t]he Gaussian copula doesn't have fat tails'.

In other words, it tends to underestimate the frequency of extreme events. He also said that, because of the ‘correlation skew’ we discuss below, ‘you run into trouble fitting the market [spreads]’. (Market participants normally characterize CDO tranches not by their prices but by their ‘spreads’, which are the increments they offer over the interest-rate benchmark, LIBOR, London Interbank Offered Rate.) For these reasons, he was also experimenting with using a copula function embodying the fatter-tailed t distribution.⁴

The four other quants we interviewed all expressed negative views of the Gaussian copula. One had developed and was a strong proponent of an alternative model. Unsurprisingly, his criticisms of the Gaussian copula were extensive. Asked by us to explain his comment that it was ‘unsatisfactory’, he joked: ‘I shall begin; we shall see if we run out of tape.’ More surprising were the comments of the three other quants, two of whom had made important technical contributions to the Gaussian copula family of models, and the third of whom was responsible for a lesser but still significant development. They too all expressed dissatisfaction, and two of them were just as outspoken in their criticisms as the proponent of the alternative model.

One of the common themes that emerged in the criticism of the Gaussian copula model was expressed by all of the interviewees, including the junior quant quoted above, as well as in the specialist trade press of the period (e.g., Hagger, 2006; Marmery, 2005). This criticism concerns what one has to do to get the standard version of the Gaussian copula to ‘fit the market’, or, in other words, what one has to do to replicate the spreads offered by the different tranches of the same CDO. If the Gaussian copula was correct, one should have been able to use the same correlation figure for each tranche, since that particular correlation was an intrinsic feature of the pool of assets underlying the CDO. But one never could. In particular, one always had to use a

higher correlation for the highest tranche (sometimes called ‘super-senior’) than for the intermediate, ‘mezzanine’ tranches. Instead of the ‘flat’ correlation structure that there should have been, there always was what participants called a ‘correlation skew’. The perfect fit of the standard Gaussian copula base correlation model to the market, referred to in our companion article, thus had what was widely seen as an arbitrary aspect; it could be achieved only by doing something incompatible with the ‘flat correlation’ ontology of the model. The existence of the correlation skew was perfectly explicable,⁵ but it could be modelled by Gaussian copulas only by modifying them in what was generally perceived as an ad hoc fashion. None of our research participants viewed the skew as a discrepancy in market prices that they could exploit, and thus reduce or eliminate, by arbitrage. Instead, they saw the skew’s existence as evidence of a flawed model. Thus arbitrage did not form the basis for any performativity of the model.

The second thread in criticism of the Gaussian copula, which is related to the above point about arbitrage, is more private, but forcibly expressed in three of the five pre-crisis interviews. As discussed in our companion article, in the culture of no-arbitrage modelling, dominant in the derivatives departments of investment banks, there was a clear prescription for how to model a derivative. This consisted of finding a ‘replicating portfolio’, which is a portfolio of more basic assets that, whatever happened to the price of those assets, would offer the same return as the derivative. (The portfolio would need continuous adjustment as those prices moved, but in a no-arbitrage model the adjustments are self-financing; once the portfolio is created, they can be made without further net expenditure.) Such a replicating portfolio could then be used as a recipe for hedging the derivative’s risks and, it provided an ‘objective’ price for the derivative. Its price must equal the cost of the replicating portfolio, because if it does not there is an opportunity for arbitrage, for riskless profit, and that opportunity cannot persist: traders will simply keep buying

whichever is cheaper – the derivative or the replicating portfolio – and selling the dearer until the price of the derivative is equal to the cost of the replicating portfolio.

The Gaussian copula was not a model of this kind. As discussed in our companion article, it was more heterogeneous in its inspiration, and the prices or spreads it generates are not imposed ‘objectively’ by arbitrage. As just noted, market participants saw the correlation skew as a defect in the model, not an arbitrage opportunity.⁶ Indeed, one quant interviewed pre-crisis denied that the Gaussian copula was worthy of the term ‘model’: ‘It’s fundamentally flawed. People refer to it as not a model but an elaborate interpolation, and I agree with that: that’s what it is.’ Another interviewee who had also contributed to the Gaussian copula model said in January 2007 that the Gaussian copula was not like the Black-Scholes option model. With Black-Scholes, ‘the price is something that is derived from a hedging strategy’. In contrast, the Gaussian copula was ‘a pricing model [that] gives everyone a consensus to all sort of use the same model, put in roughly the same inputs, and therefore everyone kind of agrees on the same price’. The Gaussian copula base correlation model, discussed in our companion article, ‘became performative’, said a third interviewee in October 2006, ‘in that the act of me going out and saying “This is a great valuation tool” ... meant ... everyone said “We’ll use [it].” Once everyone was using it, you have to use it as well’, because it then becomes a good guide to prices.

The above objections to the Gaussian copula – the arbitrary volatility skew, and the fact that it was not a ‘proper’ no-arbitrage model – continued to be voiced by the quants we interviewed after the crisis. For instance, one later interviewee, who made important contributions to Gaussian copula model, echoed our earlier interviewee’s denial that it counted as a model: ‘the nice thing is that it fits the market exactly ... The bad thing is it’s not a model ... [Y]ou’re not computing values of things as expectations under some well-defined measure’ (in

the probability-theory sense of ‘measure’; see MacKenzie and Spears, 2014). Said another: ‘Copulas are generally an early doodling activity in an area ... a simple trick ... perceived as a hack.’ Nevertheless, the fact that criticisms such as these – and others⁷ – were also expressed *prior* to the credit crisis raises the crucial question: why was a model that was widely seen as flawed still used? Until around 2005, part of the answer was the absence of competitors seen as adequate for modelling CDOs. As an interviewee put it, ‘the dominant solution [Gaussian copula base correlation] ... was unsatisfactory for a number of reasons [such as those outlined above] fairly well understood by everybody. But at the time there was no viable alternative. ... So we were on base correlation and grumbling.’⁸ However, in the latter part of the decade a number of alternatives emerged that were more attractive from the viewpoint of the culture of no-arbitrage modelling, such as the ‘gamma process’ model developed by the quant Martin Baxter (2007). But other than at one bank (a relatively small participant in the CDO market), which employed Baxter’s model, Gaussian copula base correlation remained in use. Furthermore, with the exception of the ‘tweak’ discussed below, this model remains dominant. Why? The answer, we posit, is the embedding of the Gaussian copula in intra- and inter-organizational processes in investment banking.

Talking with models

One form of this embedding was the role of the Gaussian copula as a medium of communication between people working for different banks or hedge funds. Like many derivatives, CDOs are complex products. Two different CDOs can be hard to compare, and it can be hard to judge whether the spreads offered by the tranches of the one are more or less attractive than those offered by the tranches of the latter.

Two decades earlier, the options market had faced the same problem of the lack of easy comparability of the prices of two different options. In response, market participants gradually adopted the practice of ‘talking with models’, especially with the canonical Black-Scholes model. It was used not only to price an option, but also to work out the level of volatility of the price of the underlying asset consistent with a given option’s price. (Other things being equal, the higher the volatility of the underlying asset the higher the price of the option.) ‘Implied volatility’ was calculated by running the Black-Scholes model ‘backwards’.⁹ Doing so allowed two different options with different features – for example, an option with a three-month maturity to buy IBM stock at \$240, and one with a six-month maturity to sell IBM stock at \$200 – to be compared on a single underlying parameter. ‘Implied volatility’ was invoked frequently when participants in options markets talked, even when they negotiated a price. Two traders haggling over the price of an option could talk to each other not in dollars but in implied volatilities. For example, if one trader was offering to buy an option with an implied volatility of 20 percent, and another was offering to sell it at 24 percent, they would perhaps split the difference at 22 percent. Indeed, this form of communication became sufficiently widespread that dealers’ quotations in options frequently take the form not of dollar prices but of implied volatility levels.

With many investment-bank participants in the CDO market having experience of trading and/or modelling options, it is unsurprising that a similar communicative practice emerged around CDOs. Participants ran Gaussian copula models ‘backwards’ to extract ‘implied correlation’ (the correlation level consistent with the ‘spread’ offered by a CDO tranche). To do so necessitated a considerable simplification; the correlations of all pairs of corporations or other debt issuers in the CDO’s ‘pool’ had to be assumed to be identical. Nevertheless, ‘implied correlation’ became a standard feature of how participants talked about CDOs, and this practice became an important form of the embedding of the Gaussian copula. If two traders from two

different banks or hedge funds were to talk successfully using ‘implied correlation’, they both had to be using CDO models that were sufficiently similar for the correlations ‘backed out’ from each to be comparable. Otherwise, as an interviewee put it, ‘Like two people speaking two different languages, they can’t really have a conversation.’ Only the Gaussian copula was used widely enough to serve as the necessary Esperanto. Whatever models different traders might privately prefer, ‘we communicate using the numbers implied by the Gaussian distribution’, this interviewee told us.

This use of the Gaussian copula for communication did not, however, become as deep as the equivalent use of the Black-Scholes model in options. The reason lies in the material implementation of the two models. Black-Scholes had an analytical solution: a formula for the price of an option that was an ordinary, explicit mathematical expression (MacKenzie, 2006: 264, equation 2). The Gaussian copula did not have a strictly analytical solution, except in the special case found by Vasicek for the large homogeneous pool (see MacKenzie and Spears, 2014). This exception aside, the Gaussian copula was at best semi-analytical: its solution involved computerized numerical methods, and there were choices to be made in how to implement those methods. As one interviewee said: ‘There is your [numerical] integration routine. Do you use a trapezium rule? Do you use Gauss[ian] quadrature? There are all sorts of nuances.’ And as another interviewee put it: ‘What is a single-factor Gaussian copula? ... The implementation is absolutely key. All it [the model] says is, integrate under here. How you choose to integrate under this function is still open to [different] implementations. So, yeah, everything will be slightly different.’ Even the standard Gaussian copula base correlation model was in material reality multiple; different implementations of it could yield somewhat different results.

In the case of Black-Scholes, two traders could agree a deal ‘priced’ as a level of implied volatility, and both their models would then output effectively the same dollar price. With the Gaussian copula, however, two traders could agree on a correlation level, but even if they were using what was in abstract ‘the same model’, its different implementations would often produce spreads that differed by small but economically consequential amounts, stymieing the consummation of the deal. As a quant told us in January 2007:

... everyone has agreed on this model [Gaussian copula base correlation], but ... let’s say you take two [implementations] built by two different quants. You put in the same correlations and you might find your CDO price is quite different. ... So if you had 100 basis points [one percentage point] implied spread on a CDO tranche, you might find that two different models would [output] 99 to 101, and [the difference] could even be more than that in certain places. So when people were initially quoting correlation, they found that it didn’t translate into being tradable, because it still didn’t allow them to pin down the price enough.

In 2004, the J.P. Morgan team who, as described in our companion article, were successfully pushing the idea of base correlation also tried to tackle this problem of different implementations head on. They sought to persuade others in the market for standardized indexed tranches all to use Vasicek’s large homogeneous pool model, with its simple analytical solution, as the way to move between correlation levels and prices:

We went out with ... a large-pool model, ‘cos I was hoping it was going to be [like] Black-Scholes ... my hope was, you could almost have it as a quoting mechanism, right, if everyone had the same model and they all agreed on the

same model it didn't matter whether it was a good model or not. ... [W]e could give someone the spreadsheet with it [the large-pool model] in. So, here you are, there's no add-ins [additional algorithms such as numerical integration] or stuff, ... it's just standard sums that you can look into, understand how it works and run it again and again and again. And we can give that [to market participants].

The effort did not succeed. J.P. Morgan's advocacy for the large homogeneous pool to be used as a convention for price quotation was misunderstood as advocacy for the internal use of the model, for example as a means of calculating 'deltas' (hedge ratios, as discussed in our companion article). The misunderstanding was perhaps wilful, because other global banks were seeking to contest J.P. Morgan's dominant position in the credit derivatives market. The effort to achieve communicative consensus around the large-pool model 'was fairly successful in Europe', said an interviewee, but 'not very successful in the U.S. where basically our, our sort of rival firms spun it as, "J.P. Morgan has got an inaccurate model."' Because the model assumed complete homogeneity of the assets in a CDO's pool, it implied exactly the same hedging ratio in respect to each asset, and plainly that was implausible. As another interviewee put it, '[market participants] all said, "deltas are rubbish", so they dropped the model.'

Because J.P. Morgan's effort did not succeed, the use of the Gaussian copula for purposes of communication never became as deeply entrenched as the equivalent use of the Black-Scholes model in the options market. As an interviewee said, 'because the standardized [large homogeneous pool] model failed, people had to drop correlation as a quotable' in the standard index tranche market. This process of ceasing to quote correlations was well underway when we conducted our first interviews in 2006. In the

case of more complex deals, however, the practice of agreeing on deals by agreeing on correlation levels continued because it provided a point of stability in *ad hoc* negotiations amongst sophisticated participants. For example, a manager of one of the leading hedge funds in this area told us:

... you can imagine that if you are having a negotiation with somebody, and you get to the end of the day, and you can say, 'I think we got a deal,' what is it that you have a deal on? ... What happens if, when you come in the next morning, spreads [on the underlying assets] are fifty basis points wider? ... what's the price? How can we agree that at 5 o'clock today we are going to make a fair adjustment based on how the market changes for when we get in tomorrow? Well, we can say, 'look, spreads are going to move, dispersion is going to move, let's just agree on what the implied correlation is'. We agree the implied correlation is 12 percent, you're done.

Remuneration and risk

A deeper form of organizational embedding of the Gaussian copula in investment banking occurred in the intertwined processes of determining traders' bonuses and assessing the riskiness of their trades. A critical issue was deciding when and how the anticipated future revenues from a trade should be 'booked' or recognized as profit in accounting terms. Our interviewees reported a universal desire among traders for future revenues from a credit derivatives deal (most of which last for between five and ten years) to be recognized as soon as the deal was done – as

‘Day 1 P&L’ – and so boost that year’s bonus as much as possible. (‘P&L’ is, as noted, profit and loss, the crucial determinant of traders’ bonuses.) ‘Let’s say ... you sell a deal for ... 100 and it’s really worth 95 [i.e. 95 percent of the sale price]’, said an interviewee. Another interviewee told us that in the early years of the credit derivatives market it was not unusual for traders to sell a deal ‘at par’ – 100 cents in the dollar – when their ‘bank[’s] system would have told them that this was worth about 70 cents’. A single trade ‘would make [\$]20 million of P&L’. Could the difference between price and value be booked immediately as Day 1 P&L, or would ‘you have to accrue that profit and you can only take, say it’s a ten-year deal, you can [only] take a tenth each year’? From the trader’s viewpoint, gradual accrual over five to ten years was deeply unattractive; many traders would likely leave the bank in question before five years were up; almost all would have done so before ten years.

Being able to ‘book’ the anticipated revenues from a credit derivatives deal as Day 1 P&L depended upon having a credible estimate of value, of how much the deal was ‘really worth’. Banks originally had considerable discretion concerning whether to book future revenues as Day 1 P&L, but Enron’s indiscriminate booking of Day 1 P&L from its energy-derivatives deals was thrown into the spotlight by its 2001 bankruptcy, and the issue began to attract the attention of regulators and auditors. In 2002, the Emerging Issues Task Force of the US Financial Accounting Standards Board (FASB) began to examine ‘whether unrealized gains or losses may be reported [i.e. recognized as profit] at inception of energy trading contracts’ (Emerging Issues Task Force, 2006: 3). With the Securities and Exchange Commission making clear that the underlying issue did not affect merely energy derivatives, with concern about the issue growing in Europe, and with the collapse in 2002 of Enron’s auditors, Arthur Andersen, which made the surviving auditing firms aware of just how big the dangers were, Day 1 P&L moved to centre stage.

The issue had two main aspects. The first was that the prices or mathematical parameters used in the calculation of P&L needed to be observable. At the start of the 2000s, it would have been hard to claim with credibility that the crucial parameter in Gaussian copula models, correlation, was observable; it could, at best, be estimated with difficulty. The method widely employed in the late 1990s was to use the easily observed correlation between two corporations' equity prices (share prices) as a proxy for the desired unobservable parameter, which was the correlation between the market values of their assets or of their survival times before default. Using equity prices, however, was too easily contested as a 'fudge'. As a textbook put it: 'There is no theoretical equality between equity correlation and default time correlation. ... [E]quity derived correlations have no theoretical justification' (Chaplin, 2005: 259-260). The issue was, interviewees reported to us, a major spur for the development of the standardized index tranche markets discussed in our companion article. Correlations 'backed out' from market prices in those markets using a Gaussian copula model were, in practice, agreed by auditors as having been 'observed' from the viewpoint of permitting the booking of revenues as Day 1 P&L, in part because that was the market-standard model. As an interviewee told us in May 2007, '[w]hen the [external] auditors or Finance [internal accountants and auditors] come in to look at our books we have to be market-standard.' Even at the one bank in which we discovered a radically different model being used instead of a copula, we were told that '[f]inance do look at [Gaussian] base correlations ... for reference'.

The second aspect of the issue was that future revenues could be treated as Day 1 P&L only if accountants and auditors could be persuaded that those revenues were reasonably certain. This meant that they had to view a deal as properly hedged, so that adverse price movements would not reduce or eliminate these future revenues. Traders and the quants supporting them

also needed to keep in mind the attitudes of risk controllers, who could constrain their capacity to trade or stop them trading. All derivatives traders in investment banking were, by the period discussed here, governed by risk-control procedures intended to disincentivize unhedged trading, and the models employed by risk-control departments were the basis by which those departments would judge if a trade was properly hedged. Hedging is a model-dependent activity. That is, calculation of the necessary hedge ratios requires a model of the movements of prices and spreads. Like accountants and auditors, risk controllers almost always used market-standard Gaussian copula models. If traders used such a model to determine hedge ratios, then their trading positions were thus likely to be judged properly hedged, and therefore both allowable from the viewpoint of risk control and predictable enough in their profitability to be eligible for Day 1 P&L.¹⁰

What would happen if you started to trade using a different model? Suppose tomorrow you ‘invented a fantastic model for pricing a CDO’ that was better than the Gaussian copula and closer to a no-arbitrage model. What could you then do with the new model, a quant asked us in November 2007. Could you ‘put on a massive position’ and make a huge profit? No, because ‘a really fantastic model ... is only going to be proved to be fantastic by the ability to go and hedge’ that position’s risks. Because others were still setting prices using the Gaussian copula, what were objectively the correct hedges (those implied by the superior new model) would appear to be wrong, and you could thus lose money, perhaps for years, and be vindicated only when long gone from the bank; ‘that’s what the depressing thing ... about being a quant is right now’:

[I]f I went to the people here [in his bank] and said, ‘we want to get this new model validated and use it in production [pricing and hedging]’ ... it would be a simple

point of ‘if your model is not fitting the market, sorry’. ... So you’ve got to put on an irrational hedge; that’s the only way to do it. Or you’ve got to face losing money. Painful.

The force of market-standard Gaussian models could be felt even without discovering what would happen if you tried to use a non-standard model for pricing and hedging. An important form of this force was via a service called Totem, administered by Markit, the leading data provider for credit derivatives. Each month, Totem sends trading desks a set of hypothetical CDOs to be priced. A front-office quant does the pricing, returns the result to Totem, and receives back the anonymized prices calculated by each trading desk using the service (unless the prices she or he has input are too far from the average, which raises suspicions that the quant was trying to manipulate the latter, in which case nothing is received back). Each bank’s accountants, auditors and risk controllers can thus use Totem results to assess the closeness of its quants’ and traders’ pricing to that of the rest of the market (most of the participants in which used and still use Gaussian copulas). ‘You do monthly submissions on [Totem], and as long as that is showing a happy result [prices close to the average of those submitted by other banks] then Finance will be pleased’, said an interviewee. That ‘happy result’ could of course most simply be achieved by using the market-standard model with parameter values similar to those others used.

The processes encouraging and on occasion even compelling use of the market-standard model had one particularly striking manifestation. J.P. Morgan, whose quants, as described in our companion article, developed the ‘base correlation’ version of the Gaussian copula that became the market standard, did not initially use it internally. It was only ‘a year later ... that we finally moved to base correlation as our valuation methodology. And the reason we did that is because everyone else did. ... [I]t becomes self-fulfilling: that’s what everyone uses, so that’s

how people assume [pricing is] going to work.’ As noted above, some participants’ pre-crisis awareness of this self-fulfilling aspect of the use of the Gaussian copula caused them disquiet, but was also a reason they felt they had to use it. As this interviewee said ‘... you need to know where the price is going to be tomorrow’, and to know that one had to use the model everyone else was using.

Certainly, their disquiet did not generally stop our interviewees using the Gaussian copula. The interviewee who proposed the above thought experiment concerning what would happen if he used a non-standard model, and who felt strong disquiet about the Gaussian copula, summed up why he had to keep using the latter. The most important role of a model in investment banking is as ‘a device for being able to book P&L’, he told us in this January 2007 interview:

[Y]ou can’t say, I have the most fantastic model ... I love this model and this model tells me I have made this much money so I want to book this much profit and pay my traders their bonuses. ... You can’t do that, you have ... to be able to say ... I have a hundred-name portfolio which I traded with a client and I’ve got [a] Gaussian copula base correlation [model] which is market-standard. I fit the model to the market. I then do all these tricks to price my product, and now it [the model] tells me that I’ve made x . [That] effectively allows me to do a ten-year trade and book P&L today ... without that people would be in serious trouble. All their traders would leave and go to competitors.

The crises of the Gaussian copula

The background to some of the unease expressed by our earliest interviewees was the May 2005 ‘correlation crisis’ that took place a year or so before our interviews and that was rooted in the popularity of synthetic single-tranche mezzanine CDOs. These were investment products (sold by investment banks to more minor banks and other institutional investors) that mimicked the risks and returns of buying ‘mezzanine’ (next-to-lowest) tranches of CDOs. Such tranches were attractive because they combined investment-grade credit ratings with healthy spreads (increments over LIBOR). By selling these single-tranche CDOs to their customers, investment banks thereby bought lots of ‘protection’ (quasi-insurance against default) on mezzanine tranches, which left them with a market exposure they did not want; if the cost of such protection fell sharply, they would suffer serious ‘mark-to-market’ losses as their trading positions were revalued to take account of price changes. Their desire to reduce this exposure created the possibility of what appeared to be a mutually beneficial trade between the investment banks (saying to themselves, as one interviewee put it, ‘we can’t have such a concentration of that risk’) and hedge funds, looking to make profits; salespeople at banks could say to their contacts in hedge funds ‘I could structure a trade like this, it’s great value, look at the [price] history.’¹¹

The way the trade worked was that investment banks offset the ‘protection’ they had bought by selling hedge funds ‘protection’ on the mezzanine tranches of standardized indices similar in their composition to single-tranche CDOs. The hedge funds then sold protection on the lowest tranches (the ‘equity’ tranches) of those indices, and the income they earned by doing so was greater than what they were spending on buying protection on mezzanine tranches from the investment banks. By choosing appropriate relative sizes of the mezzanine protection bought and equity protection sold, the result was a delta-neutral position (that is, a position hedged against improvements or deterioration in the perceived overall creditworthiness of the corporations whose debts underpinned the index in question) that would nevertheless make a consistent profit

for the hedge fund. Even some banks seem themselves to have been tempted into the trade. However, this form of trade exposed the banks to the possibility of correlation levels falling. In the terminology of the new field of correlation trading, taking part in such a trade makes one ‘long correlation’: you benefit if correlation rises, lose if it falls.. (High levels of correlation benefit those who have ‘insured’ – sold protection on – equity tranches because it makes outcomes more binary, as in the 0.99 case in figure 2 of our companion article. The chance of catastrophe sufficiently serious to hit even the most senior tranches increases, but the chance of little or no loss, and therefore an intact or almost intact equity tranche – and thus no claim on the insurance, or only a small claim – increases as well.) Put another way, the hedge funds were exposed to events that would provoke concern about idiosyncratic risks, or risks that affect just one corporation or a very small number of corporations; such risks endanger the sellers of protection on equity while leaving the situation of mezzanine tranches almost unchanged. (Equity is, as noted, the lowest tranche in a CDO, and thus the first to suffer losses, so the default of even a single corporation can affect the holders of the equity tranche. In contrast, several defaults need to take place before the mezzanine tranche suffers losses.)

Idiosyncratic risk was precisely what manifested itself on 5 May 2005, when Standard & Poor’s stripped General Motors and Ford of their investment-grade ratings, reducing GM to BB and Ford to BB+. It is a noteworthy event when a ratings agency reduces the obligations of the great mass-market car companies of the 20th century to ‘junk’. But it took place in generally benign economic conditions. It could be interpreted as an increase in a very specific risk. What appears then to have happened, interviewees told us, was that a particular large hedge fund (one interviewee named it, but it has been impossible to get confirmation of its identity) decided to unwind its position, which meant buying protection on equity tranches to cancel out its sales. The cost of ‘protection’ on those tranches thus increased, placing pressure on those who had similar

positions, who then also tried to unwind, further increasing the cost of protection on equity. In contrast, the cost of protection on mezzanine tranches fell (unwinding implied having to sell protection on those tranches). When Gaussian copula models were used to ‘back out’ correlation levels, that pattern of change in costs suggested that the correlation skew (explained in the second section of this article) had steepened sharply, hence the name ‘correlation crisis’.¹²

The result, said interviewees, was large losses for a number of hedge funds and some banks. The crisis attracted very little reporting, either at the time or subsequently, perhaps because of its complicated nature (and the absence of any spectacular bankruptcies). The *Financial Times* reporter Gillian Tett was one of the few to pick up the story, and her informants said that it involved undue faith in models: ‘People thought the models were almost infallible – the last few days have been a real shock’, one banker told her (Tett, 2005). Certainly, a naïve interpretation of the Gaussian copula model might have suggested that a position that was delta-neutral (as the trades central to the correlation crisis were intended to be) was thereby free of risk. However, when the first author suggested to another interviewee in January 2007 that the trade had been ‘model-driven’, he disagreed:

the press always wants to talk about these smart traders who were wrong because they believed in the models. I mean, no-one is that stupid that you put on a trade with a delta which is delta-neutral, I mean, no, you know that it’s only delta-neutral if nothing else changes.¹³

Even if not caused by a naïve interpretation of Gaussian copula models, the 2005 correlation crisis was certainly a temporary crisis for modelling practices. The steepening of the correlation skew during the crisis was sufficiently large that on some days market-standard

Gaussian copula base correlation models simply failed to calibrate: they could not find correlation levels that allowed them to match the spreads at which tranches were trading. An interviewee reported that this happened both to a particular model he had developed and more generally:

the [sharply reduced spreads on the] mezzanine tranche actually violated [the] lower bound that this stochastic correlation model was imposing. ... [The episode] was very upsetting to many people because their models simply stopped working. They couldn't match the market any more.

That calibration failure, however, was only temporary, and the 'correlation crisis' did not generate any major widespread change in the dominant practices of modelling. Far more persistent failures of models to calibrate were experienced in the second of the crises to afflict correlation modelling, the credit crisis that erupted in the summer of 2007. As the crisis deepened, the cost of protection on the apparently safest, super-senior tranches of the indices (as noted, these are widely traded standardized CDOs) rose to unprecedented levels, as fears of systemic collapse increased. Again, but much more frequently than in 2005, no correlation value at all could be found that enabled the spreads at which super-senior tranches were being quoted to be reproduced:

[Y]ou can derive some bounds on the value of the super-senior tranche [from the Gaussian copula model]. And those bounds were violated by the market. Spreads were too high for the super-senior tranches. You couldn't get there.

Hugely disruptive as failures to calibrate such as this are to the day-to-day work of pricing and hedging, there is nevertheless a sense in which the market-standard Gaussian copula base correlation model has survived even this crisis. It has been ‘tinkered with’, rather than discarded. Prior to the 2007-08 crisis, it was conventional to assume simply that if a corporation defaulted then the ‘recovery rate’ (the extent to which its creditors would get back what they were owed) would always be 40 per cent, a value that was roughly the historic average. More recently, however, that assumption has been discarded, and recovery rates have been modelled as stochastic. In particular, in the ‘one-factor’ Gaussian copula models discussed in our companion article, recovery rates have been made dependent upon the value of an underlying factor that can be interpreted as the ‘state of the economy’. It is assumed that in ‘bad’ states of the economy, recovery rates will be much lower than in ‘good’ states. Altering standard models in this way has made it possible for modelling to ‘work’ (to calibrate) most of the time, even in the very turbulent conditions of recent years. ‘Working’ has still not been universal – there have reportedly been particular days when even with this alteration standard models fail to calibrate (Brigo, Pallavicini and Torresetti, 2010: 104) – but the ‘fix’ has been good enough to keep the Gaussian copula dominant. In a situation in which the underlying markets have shrunk markedly, it has been judged better to ‘fix’ a model that was already understood by traders, accountants and risk-controllers than to suffer the financial, communicative and cognitive costs of moving to a radically different model.

‘The formula that killed Wall Street’?

What has just been discussed, however, is the (limited) effect of the credit crisis upon the Gaussian copula family of models. What, however, of the effect in the other direction? *Did* the Gaussian copula kill Wall Street, as Salmon (2009) suggests?

The market participants on which this article has focussed – the users of Gaussian copula models in the derivatives departments of investment banks – came under huge strain (including the calibration failures discussed in the last section), but their activities did not generate losses of sufficient magnitude to threaten the survival of their banks or of the financial system. Certainly, there were losses on the credit default swaps, the index tranches, and the CDOs (based on pools of corporate debt) with which those actors dealt, but those losses – while very big – were not catastrophic. As an interviewee said in July 2010: ‘Losses you hear around the place, “I lost a billion dollars” ... in normal times would be very notable’. A billion dollar loss, however, does not kill a global bank. The level of loss needed to do that (of the order of \$20-\$50 billion) did not come from the world discussed here: ‘the base corr guys [users of Gaussian copula base correlation models] are still standing... There were definitely bad days for everybody with the markets jerking around, and people felt the swings but I am not sure that there was anything in terms of an Armageddon for the models’.

Rather, the critical path by which the Gaussian copula *was* implicated in the credit crisis was via rating agencies, in particular in the rating not of ‘traditional’ CDOs based on pools of corporate debt, but of ‘ABS CDOs’, in which the underlying assets are asset-backed securities (ABSs), specifically mortgage-backed securities. We have discussed these and their role in the credit crisis elsewhere (MacKenzie, 2011). ABS CDOs were introduced somewhat later than corporate-debt CDOs, and originally were a small-scale business; only 3 percent of the CDOs issued in 1997-1999 were ABS CDOs (Newman et al., 2008: 34, exhibit 1). By the time ABS CDOs started to become large-scale (from 2001 onwards), the rating agencies already had in place an organizational division of labour. Both CDOs and ABSs fell within the remit of their structured finance departments, but those departments had separate groups rating CDOs, on the one hand, and ABSs on the other.

As discussed in MacKenzie (2011), the new ABS CDOs were therefore evaluated by the rating agencies in two temporally and organizationally separate steps. First, the underlying mortgage-backed securities or other ABSs were rated by the groups handling those securities, and then the overall CDO structure was rated by the CDO groups. Instead of considering ABS CDOs as radically different instruments that required an altogether new form of evaluation, the CDO groups simply made modest modifications to the techniques they already used to analyze CDOs whose pools consisted of corporate debt. From late 2001 onwards, those techniques increasingly involved the use of models in the Gaussian copula family, albeit usually one-period models analogous to CreditMetrics, *not* fully-fledged copulas of the kind introduced by David X. Li (MacKenzie and Spears, 2014). With little econometric data to draw upon (empirically estimating the correlation between ABSs is an even harder econometric problem than estimating correlations between corporations), the CDO groups employed largely judgment-based ABS correlation estimates, which were broadly similar in size to those they used for the analysis of corporate CDOs. When, for example, Standard & Poor's introduced its new one-period Gaussian copula system, *CDO Evaluator*, in November 2001 the same correlation (0.3) was used for the correlation between ABSs from the same sector (for example, ABSs based on subprime mortgages) as was used for the correlation between corporations in the same industry (Bergman, 2001).

The result of the assumption of only modest correlation was an extremely attractive opportunity for market participants to take ABSs of only modest credit quality (for example, the mezzanine tranches of subprime mortgage-backed securities with BBB ratings) and package them into CDOs with very large AAA tranches. Widespread exploitation of this opportunity had catastrophic consequences, both direct and indirect. A substantial proportion

of the gigantic losses that directly crippled global financial institutions were incurred on ABS CDOs. Citigroup lost \$34 billion on ABS CDOs, Merrill Lynch \$26 billion, UBS \$22 billion and AIG \$33 billion (see Benmelech and Dlugosz, 2009), and the avid demand of ABS CDOs for the mezzanine tranches of subprime mortgage-backed securities also had the indirect effect of side-lining the traditional buyers of such securities, who had typically scrutinized the underlying pools of mortgages with great care (Adelson and Jacob, 2008). ABS CDOs sat at the end of what market participants sometimes call an ‘assembly line’, in which subprime mortgages were bundled into ABSs, and then ABSs were bundled into ABS CDOs, with a view simply to achieving desirable ratings and with little effective concern for risks in the underlying assets that were not captured by those ratings.

In effect, market participants had ‘outsourced’ the analysis of ABS CDOs to the rating agencies. It was perfectly possible profitably to construct an ABS CDO without doing any correlation analysis of one’s own. All one had to do was to check that an intended structure would achieve the desired large AAA tranches, a task that was made easy by the fact that market participants could simply download Standard & Poor’s *CDO Evaluator* and its analogues at the other agencies. The first author vividly remembers a February 2009 interview in which he asked a senior figure at a firm that managed ABS CDOs what correlation model the firm had employed, only to be met with a blank stare: *no* model of its own had been used.

The major investment banks conducted some analysis of ABS CDOs beyond simply checking desired ratings, but, by the standards of the culture of no-arbitrage modelling, very little analysis took place in most cases. ABS CDOs often fell outside the remit of the derivatives departments of those banks. They were frequently constructed and analyzed by

other groups, such as those specializing in mortgage-backed securities. ‘The guys doing ABS had essentially different roles and different attitudes’, reported one interviewee. With the partial exception of Goldman Sachs, the modelling of ABS CDOs that was done did *not* take the form of no-arbitrage modelling.¹⁴ Rather, it involved either cashflow models of the underlying ABSs (with judgment-based estimates of likely mortgage default rates), whose outputs were then fed into a cashflow model of the CDO, or inferring the default probabilities of the ABSs from their ratings and using those probabilities in a Gaussian copula model of the CDO, in much the same way as the rating agencies modelled ABS CDOs. To those whose view was that the proper activity of a quant was no-arbitrage modelling, the catastrophic losses were thus on products that, in the words of one such quant, ‘were on the whole either less quanted or not quanted at all’.

An issue of ontology underlies judgements such as that made by the interviewee just quoted. As described in our companion article, no-arbitrage modelling extracts martingale or risk-neutral probabilities from patterns of market prices. With the partial exception of Goldman, this style of modelling – which is what the interviewee meant by ‘quanting’ – was, as far as we can discover, simply not applied to ABS CDOs. Rating agencies *did* model ABS CDOs, but rating agencies generally do not work with martingale probabilities; rather, they seek to estimate *actual* probabilities of default, and to do so they almost always use the historical records of defaults, not price patterns. In the case of subprime mortgage-backed securities, which dated only from the 1990s, such records encompassed a period of almost continuously rising house prices and only one relatively mild recession. Unfortunately, as we now know, when those benign conditions changed, such securities, and the mortgage borrowers on whom they were based, began to behave quite differently.

Conclusion

As the previous section has outlined, the Gaussian copula family of models *was* implicated in the processes that ‘killed Wall Street’. Salmon (2009), however, is quite wrong to focus on Li, the quant who, as discussed in our companion article, first introduced explicit use of copula functions. By the time of the crisis, the ratings agencies had moved only partially from the early one-period models to full fledged copula models of the kind introduced by Li, and the move was not central to the crisis. It was far less consequential than three other factors: the way in which, in their organizational structure, the rating agencies separated the analysis of ABSs from that of CDOs; the estimation of the probabilities of default on ABSs using data from a period of benign economic conditions; and the way in which the CDO groups in the agencies analyzed an ABS CDO in almost the same way as a CDO based on corporate debt, and in particular assumed that an ABS CDO would involve a level of correlation at most only modestly higher than that of a CDO based on corporate debt.

Nor would it be reasonable to blame the Gaussian copula family of models, in itself, for the crisis. These models did not have unitary, intrinsic effects. Rather, they had effects in combination with the organizational processes in which they were embedded. Gaussian copula models as employed by the rating agencies were quite different in their effects from Gaussian models employed in the derivatives departments of investment banks. Not only did the goals of modelling differ between rating agencies and investment banks, but also, as discussed in the previous section, the ontology. Moreover, the surrounding processes differed substantially. Governance, understood as risk control and the booking of profit, was certainly one aspect of the use of the Gaussian copula in investment banks. Ratings, however, were almost entirely about governance; many investment managers were forbidden (either by regulation or by organizational mandate) to buy anything other than investment-grade

securities, or the amounts of non-investment-grade purchases were strictly limited. In consequence, the ratings of such securities dictated the nature of the market for them.

The result of the embedding of Gaussian copula models in governance via ratings was the large-scale ‘gaming’ of them and of the other models employed by the ratings agencies. The crisis was caused not by ‘model dopes’, but by creative, resourceful, well-informed and reflexive actors quite consciously exploiting the role of models in governance. ‘[T]he whole market is rating-agency-driven at some level’, one of our earliest interviewees told us, a year before the crisis: ‘the game is ...to create ...tranches which are single-, double- or triple-A rated, and yield significantly more than a correspondingly rated [bond]’. This interviewee did not directly participate in the ‘game’ himself. His hedge fund profited only indirectly from the fact that, as he put it, ‘there are investors who are constrained by ratings.’ But other interviewees did. Two told us how they had employed optimization programs to find the highest-yielding pools of securities that would still make possible CDOs with sufficiently large AAA tranches; although they did not directly say this, the highest-yielding securities are those that market participants consider riskiest. Another interviewee described to us how his firm had developed, and sold to investment banks, a sophisticated software package designed to perform this perilous optimization.

Two dangers, however, attend these findings. First, our emphasis on knowledgeable, reflexive actors rather than model dopes could be read as a collapse into simplistic rational-actor, agency-theoretic explanations of the crisis. This would be quite the opposite of our intention. Culture and rationality are not opposed, even if rationality is construed as the pursuit of narrow self-interest. Even the most selfishly rational actor needs to calculate what is in his or her best interest, and that calculation of necessity partakes in the material cultures

of finance. Because such cultures differ, and because there is no a priori way to be entirely sure which practices are the most efficacious, even a fully reflexive, rational actor cannot stand wholly outside of finance's cultures of evaluation. Nor does the existence of these reflexive, rational actors diminish the co-ordinating role of models or other cultural resources. The way in which Gaussian copulas, a class of model that was often disliked, nevertheless helped achieve economically crucial outcomes (in particular the achievement of Day 1 P&L) shows that cultural resources can co-ordinate action even in the presence of widespread scepticism of their worth. One does not need to invoke cultural dopes to understand how cultural resources help produce co-ordinated action.

The second danger is that this article's findings will be read as an endorsement of no-arbitrage modelling, one of the hegemonic cultures of modern finance. Again, that is emphatically not our intention. Rather, we would note that there are multiple mechanisms of counter-performativity or, in other words, multiple ways in which the practical use of a model can undermine its empirical adequacy. One such mechanism that played a primary role in the credit crisis was the 'gaming' by market participants of the models (including the Gaussian copula) used by rating agencies. In essence, the gaming of models that assumed low default probabilities and low correlations helped bring about, in the way sketched in the previous section, huge levels of highly-correlated mortgage and ABS default.

There are, however, other mechanisms of counter-performativity. In particular, no-arbitrage models may be associated with a distinctive mechanism in which the hedging practices based on those models have effects on the market for the underlying assets that undermine the empirical adequacy of the assumptions about asset-price dynamics embedded in those models. The most obvious such case is the event that ended the period in which

patterns of option prices mirrored the Black-Scholes model relatively closely, the 1987 stock market crash. In this case, portfolio insurance (a form of hedging based on Black-Scholes) was at least to some degree implicated in violent price movements that were grotesquely unlikely on the geometric Brownian motion model underpinning Black-Scholes (MacKenzie, 2006). While they are not as well known as the 1987 crash, and seldom reported outside the specialist trade press (their details can be fiendishly complicated to outsiders), other examples of this mechanism exist.¹⁵ In such examples, what was (as far as we can tell) careful, diligent hedging based on ‘proper’ no-arbitrage models nevertheless caused substantial market disruption and serious losses, albeit closer to or lower than the \$1 billion scale of the ‘base correlation’ losses than to the \$20-\$30 billion of the ABS CDO losses.

It could be that here we have the beginnings of a typology of mechanisms of counter-performativity: models used for governance are undermined by being gamed; models used to hedge derivatives are undermined by the effects of that hedging on the market for the underlying asset.¹⁶ We end, however, with a speculation about the culture on which we have focused, no-arbitrage modelling. As this article has shown, the canonical Gaussian copula base correlation model played a *co-ordinating* role within and among investment banks. The use of the model helped to harmonize practices and prices and facilitate communication, and it provided a shared yardstick that enabled accountants and auditors to determine whether a valuation was correct and risk managers to assess whether a position was properly hedged. It therefore made possible Day 1 P&L, the up-front profit that is the essential lubricant of the trading of derivatives with maturity dates that stretch beyond traders’ likely working lives in their banks. This co-ordinating role of the Gaussian copula was visible to our interviewees – and therefore to us – precisely because they did not ‘naturalize’ the model; none believed that the Gaussian copula gave a faithful account of the economic world.

Perhaps a co-ordinating role is ever-present in shared models in finance, even those that are taken as capturing at least some aspects of the way the world is; perhaps this helps explain why investment banks – those apparently most capitalist of institutions – quite frequently give other market participants, free of charge, models in whose development they have invested much time and money. Perhaps the modelling of derivatives in investment banking always has an aspect of what one of our interviewees memorably called a ‘ballet’, in which highly-paid quants are needed not just to try to capture the way the world is, but also to achieve co-ordinated action. And – as Beunza and Stark (2012) have suggested in a different context – perhaps the seeds of disaster sometimes lie in that very achievement.

Acknowledgements

We are hugely grateful to our interviewees, to the referees and (especially) the editor of this journal, and to Daniel Beunza, Dominic O’Kane, Yuval Millo, Riccardo Rebonato and Felix Salmon for helpful comments. Bruce Worton kindly generated the graph in figures 2 of our companion article. All errors remain ours.

Funding

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013/ERC grant agreement no. 291733) and from the UK Economic and Social Research Council (RES-062-23-1958 and RES-598-25-0054). The pre-crisis pilot interviewing was supported by an earlier ESRC grant, RES-051-27-0062.

References

- Adelson M and Jacob D (2008) The sub-prime problem: Causes and lessons. *Journal of Structured Finance* 14(1): 12-17.
- Baxter M (2007) Gamma process dynamic modelling of credit. *Risk* 20(10): 98-101.
- Benmelech E and Dlugosz J (2009) The credit rating crisis. *NBER Macroeconomics Annual* 24: 161-207.
- Bergman S (2001) CDO Evaluator applies correlation and Monte Carlo simulation to determine portfolio quality. New York: Standard & Poor's, November 13. Available at: <http://www2.standardandpoors.com> (accessed 5 May 2006).
- Beunza D, Hardie I and MacKenzie D (2006) A price is a social thing: Towards a material sociology of arbitrage. *Organization Studies* 27: 721-745.
- Beunza D and Stark D (2012) From dissonance to resonance: Cognitive interdependence in quantitative finance. *Economy and Society* 41: 335-359.
- Brigo D, Pallavicini A and Torresetti R (2010) *Credit Models and the Crisis: A Journey into CDOs, Copulas, Correlations and Dynamic Models*. Chichester, WS: Wiley.
- Callon M (ed) (1998) *The Laws of the Markets*. Oxford: Blackwell.
- Cameron M (2013) Dashing the uridashi dream. *Risk* 26(3): 22-24.
- Chaplin G (2005) *Credit Derivatives: Risk Management, Trading and Investing*. Chichester, WS: Wiley.
- Credit Suisse First Boston (1997) *CreditRisk+: A Credit Risk Management Framework*. London: Credit Suisse First Boston. Available at: <http://www.csfb.com/institutional/research/assets/creditrisk.pdf> (accessed 4 June 2008).
- Emerging Issues Task Force (2006) *Issues Involved in Accounting for Derivative Contracts Held for Trading Purposes and Contracts Involved in Energy Trading and Risk Management Activities* (2/3). Available at: <http://www.fasb.org/st/index.shtml> (accessed 15 May 2008).
- Garfinkel H (1967) *Studies in Ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.
- Hagger E (2006) Traders look to play ABX series arbitrage. Creditflux Newsletter, October. Available at: <http://www.creditflux.com> (accessed 19 May 2008).
- Lépinay V (2011) *Codes of Finance: Engineering Derivatives in a Global Bank*. Princeton, NJ: Princeton University Press.
- MacKenzie D (1981) *Statistics in Britain, 1865-1930: The Social Construction of Scientific Knowledge*. Edinburgh: Edinburgh University Press.
- MacKenzie D (2006) *An Engine, not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.
- MacKenzie D (2011) The credit crisis as a problem in the sociology of knowledge. *American Journal of Sociology* 116: 1778-1841.
- MacKenzie D and Spears T (2014) 'The formula that killed Wall Street': The Gaussian copula and modelling practices in investment banking. *Social Studies of Science* 44.
- Marmery N (2005) The pricing puzzle. *US Credit* April: 42-44.
- Mars F (1998) *'Wir sind alle Seher': Die Praxis der Aktienanalyse*. PhD thesis, University of Bielefeld, Germany.
- Muniesa F, Chabert D, Ducrocq-Grondin M and Scott SV (2011) Back-office intricacy: The description of financial objects in an investment bank. *Industrial and Corporate Change* 20: 1189-1213.
- Nelken I (1999) *Implementing Credit Derivatives*. New York: McGraw-Hill.
- Newman D, Fabozzi FJ, Lucas DJ and Goodman LS (2008) Empirical evidence on CDO performance. *Journal of Fixed Income* 18(2): 32-40.

- Packer F and Wooldridge PD (2005) Overview: Repricing in credit markets. *BIS Quarterly Review* June: 1-13.
- Salmon F (2009) Recipe for disaster: The formula that killed Wall Street. *Wired*, 23 February. Available at: http://www.wired.com/print/techbiz/it/magazine/17-03/wp_quant, (accessed 25 February 2009).
- Svetlova E (2009) Theoretical models as creative resources in financial markets. In: Jansen SA, Schröter E and Stehr N (eds) *Rationalität der Kreativität? Multidiziplinäre Beiträge zur Analyse der Produktion, Organisation und Bildung von Kreativität* Wiesbaden: Springer, 121-135.
- Svetlova E (2012) On the performative power of financial markets. *Economy and Society* 41: 418- 434.
- Tett G (2005) Market faith goes out the window as the 'model monkeys' lose track of reality. *Financial Times* 20 May, 23.

Author biographies

Donald MacKenzie is a Professor of Sociology at the University of Edinburgh. His current research is on the sociology of financial markets, and he is focusing in particular on how market participants and technical systems evaluate (work out the economic worth of) financial securities. His books include *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance* (MIT Press, 1990) and *An Engine, not a Camera: How Financial Models Shape Markets* (MIT Press, 2006).

Taylor Spears is a Research Fellow in the Department of Sociology at the University of Edinburgh. His current research focuses on the community of derivatives quants and the development and social shaping of the financial models they build and use. He was previously a Research Fellow at the Science Policy Research Unit at the University of Sussex.

¹ Note that co-ordination does not necessarily imply harmony or the absence of competition. The most bitterly contested football match is still an example of co-ordinated action.

² The Black-Scholes model has one free parameter, volatility. The arbitrage that pushed patterns of option prices towards the postulates of Black-Scholes was 'spreading' (MacKenzie, 2006: 164-166), which did *not* involve the choice of a particular value of the volatility parameter. There was no full equivalent of 'spreading' with the Gaussian copula.

³ It is possible that even our earliest interviewees had been affected by the experience of the correlation crisis. See below for a brief discussion of this possibility.

⁴ On the history of this distribution, see MacKenzie (1981: 111-116).

⁵ The levels of correlation that fitted the 'spreads' of the lower tranches (the increments over benchmark interest rates that they offered) generated a spread on the most senior tranches that was far lower than the spread those tranches had to offer if investors were going to buy them. To get the model to generate the latter spread, it was necessary to input a higher level of correlation for the more senior tranches. As an interviewee put it to us, 'Maybe the model [with a flat correlation] says the super-senior tranche only pays [a spread of] three basis points [0.03 percent], but who the hell is going to read through the whole of the prospectus, figure out the risk, hire a lawyer to analyze the document, figure out how to book it, get a model approval, da da da, for something that only pays three basis points. They're saying, "Look, I'm really not going to get out bed for anything less than ten [basis] points." There is no science in that, it's just anything about ten sounds kind of good.'

⁶ We draw this stark contrast because it underpinned the critiques of the Gaussian copula we heard from our interviewees. It is in fact a simplistic contrast; arbitrage is a more complex matter than it suggests. See Beunza, Hardie and MacKenzie (2006), and also the discussion in our conclusion of counter-performativity.

⁷ For example, another criticism was that the Gaussian copula was essentially static. As an interviewee put it in February 2007: 'it has no [time] dynamics. Copulas are just a way of bolting together marginal distributions.'

⁸ Other models of likely losses in pools of assets were available, notably Credit Suisse's CreditRisk+ (Credit Suisse First Boston, 1997), and were used reasonably widely for modelling banks' credit risks, but seem to have been judged less suitable for modelling CDOs, especially in 'a trading situation' (Nelken, 1999: 237).

⁹ See Beunza and Stark (2012) for a discussion in a different context of 'backing out' a parameter. For more detail on the use of 'implied volatility' in options markets, see MacKenzie (2006: 168-169).

¹⁰ Participants acknowledge that the 'perfect' hedges of no-arbitrage modelling are not to be found in reality (apart from 'in Japanese gardens', as the traders' joke has it), and various 'reserves' are deducted from Day 1 P&L to try to take account of this. 'Reserving' is thus a crucial process, which unfortunately cannot be discussed here for reasons of space.

¹¹ The interviewee in question was talking more generally rather than about the specific trades discussed here.

¹² See the graphs of tranche spreads and the corresponding base correlation levels in early 2005 in Packer and Wooldridge (2005: 6).

¹³ Clearly, this assertion could involve hindsight, but it seems plausible, in that by 2005 'correlation trading' was well-established. It was predicated on the fact that correlation levels change, and if they do so a delta-neutral position can still suffer losses.

¹⁴ Goldman's modelling of ABS CDOs used estimates of default probabilities and correlations based on patterns of market prices, not, for example, the historical records of mortgage defaults used by other banks and by the rating agencies. Our hypothesis is that this may in part account for Goldman's decision to exit the subprime market (and indeed to 'short' it) as market conditions began to deteriorate late in 2006, a decision that made it possible for Goldman to survive the crisis almost unscathed financially. However, the post-crisis lawsuits faced by Goldman made it impossible for us to interview those involved.

¹⁵ There have been a number of market disruptions involving the hedging of a class of interest-rate derivatives known as constant maturity swaps (our attention was first drawn to these by an interviewee, who said one such episode in 2008 had caused 'chunky losses all around the City'). Another episode, in 2012, involved the hedging of *uridashi*, a form of option (heavily sold by investment banks to Japanese retail investors) that is linked to the Nikkei stock market index. *Risk* magazine reports total losses to the banks of up to \$500 million in this episode (Cameron, 2013).

¹⁶ We thank David Stark for pressing on us the importance of a more systematic understanding of counter-performativity. A third form (not found in the episodes discussed here) is what might one might call 'deliberate counter-performativity': the employment of a model that one knows overestimates the probability of 'bad' events, with a view to reducing the likelihood of those events (for an example, see MacKenzie, 2006: 209-210).