THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

'The formula that killed Wall Street'

Citation for published version:

MacKenzie, D & Spears, T 2014, 'The formula that killed Wall Street': The Gaussian copula and modelling practices in investment banking', *Social Studies of Science*, vol. 44, no. 3, pp. 393-417.
<https://doi.org/10.1177/0306312713517157>

Digital Object Identifier (DOI):

[10.1177/0306312713517157](https://doi.org/10.1177/0306312713517157)

Link:

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

Document Version:

Peer reviewed version

Published In:

Social Studies of Science

Publisher Rights Statement:

© MacKenzie, D., & Spears, T. (2014). The Formula That Killed Wall Street: The Gaussian Copula and Modelling Practices in Investment Banking. *Social Studies of Science*, 44(3), 393-417.
[10.1177/0306312713517157](https://doi.org/10.1177/0306312713517157)

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.

‘The formula that killed Wall Street’: The Gaussian copula and modelling practices in investment banking

Donald MacKenzie and Taylor Spears

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

Abstract

Drawing on documentary sources and 114 interviews with market participants, this and a companion article discuss the development and use in finance of the Gaussian copula family of models, which are employed to estimate the probability distribution of losses on a pool of loans or bonds, and which were centrally involved in the credit crisis. This article, which explores how and why the Gaussian copula family developed in the way it did, employs the concept of ‘evaluation culture’, a set of practices, preferences and beliefs concerning how to determine the economic value of financial instruments that is shared by members of multiple organizations. We identify an evaluation culture, dominant within the derivatives departments of investment banks, which we call the ‘culture of no-arbitrage modelling’, and explore its relation to the development of Gaussian copula models. The article suggests that two themes from the science and technology studies literature on models (modelling as ‘impure’ bricolage, and modelling as articulating with heterogeneous objectives and constraints) help elucidate the history of Gaussian copula models in finance.

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.

Email: frances.j.burgess@ed.ac.uk

The Formula That Killed Wall Street: that was how *Wired*'s editors introduced the Gaussian copula to the readers of a February 2009 article by journalist Felix Salmon. The model had 'devastated the global economy'. Its author, 'math wizard ... David X. Li ... won't be getting [a] Nobel [prize] anytime soon', wrote Salmon. 'Li's Gaussian copula formula will go down in history as instrumental in causing the unfathomable losses that brought the world financial system to its knees' (Salmon, 2009).

In this and a companion article published in this journal (MacKenzie and Spears, 2014) we examine the history of the Gaussian copula family of models, their embedding in organizational practices in finance, and their role in the global financial crisis. The current article presents a history of these models set in the context of a discussion of the dominant 'evaluation culture' – as we call it – of the modelling of financial derivatives, a culture that enjoys a degree of intellectual hegemony in modern investment banking. (A derivative is a contract or security the value of which depends on the price of an underlying asset or the level of an index or interest rate.)

Given how crucial mathematical models are to financial markets, surprisingly little research has been devoted to *how* financial models develop, which is our topic in this article; we return to other issues in the companion article. One theme in work on models by researchers on finance influenced by science and technology studies (STS) has been the 'performativity' of models: models are not simply representations of markets, but interventions in them and part of how markets are constructed (Callon, 1998, 2007). Models do indeed have effects, but – vital though that issue is – exclusive attention to their effects occludes attention to the processes that shape models and their development. Research on the history of financial modelling has seldom gone beyond 1970 when the canonical financial-

derivatives model, the Black-Scholes or Black-Scholes-Merton options pricing model, was constructed.¹ If the history of modelling in the decades since 1970 is treated in detail at all – and these are decades in which global financial markets have changed utterly – it is by practitioners (the best such work is Rebonato's [2004] history of interest-rate modelling).

Modelling more generally has, however, become a significant focus in STS and in philosophy of science. Much of the research on models in STS and philosophy of science addresses issues – for instance, whether modelling is a form of knowledge generation distinct from both theory and experiment (e.g., Dowling, 1999; Galison, 1997), or whether models are the crucial intermediaries between theory and reality (Cartwright, 1983) – that have no exact analogues in financial markets. Of course, the institutional contexts and purposes of modelling in finance and in science are different. The goal of most modelling in finance, after all, is to make money, not to contribute to knowledge. Experiment – the relationship of which to modelling in science has been an important topic for scholars (e.g., Morgan, 2005) – is much less prominent in finance. Finance does have its experiments (see Muniesa and Callon, 2007), but they are generally looser affairs. Furthermore, neither experiments nor experimental evidence played a part in the history discussed in this paper. Nor does 'theory' occupy the prominent place in finance that it does in many sciences; for many financial practitioners, 'theory' (option pricing theory, for example) simply *is* a collection of models, not something separate from models.

Nevertheless, there is much in the science-studies literature on models that can help frame research such as ours. A common finding is that building models is a creatively

¹ On the history of this model, see MacKenzie (2003) and Mehrling (2005); for an excellent sample of historical work on financial modelling, see Poitras and Jovanovic (2007).

‘messy’ process that draws on a heterogeneity of elements (e.g., Morgan and Morrison, 1999; Breslau and Yonay, 1999; Morgan, 2012). Boumans (1999), for example, argues that model construction in economics is ‘a trial and error process’, ‘like building a cake without a recipe’ from heterogeneous ingredients (pp. 95 and 67). In a one-word summary, model construction is bricolage (MacKenzie, 2003).

Other relevant themes from the science-studies literature on models emerge, for example, from Sismondo’s (2000) nuanced analysis of controversy surrounding Robert MacArthur and Edward O Wilson’s ‘island biogeography’ model, which posits a simple mathematical relationship between an island’s area and the number of species on it (MacArthur and Wilson, 1967). The model is not unitary, argues Sismondo (2000): ‘it has multiple uses and interpretations’. It can legitimately be viewed either as ‘true but quite abstract’ or as ‘interesting but false’. Its ‘success in representing nature’ thus depends on who is assessing that success, their professional identities, the kind of scientific work they engage in, and the role of models in that work. ‘[T]he success of IB [the island biogeography model] in representing nature depends in part upon [scientific] lifestyle and labor issues’ (Sismondo, 2000: 251-54). ‘Seeing models and simulations just in a space between theories and data’, Sismondo (1999) argues, ‘misses their articulation with other goals, resources, and constraints’ (p. 254). The articulation of Gaussian copula models with broader ‘goals, resources, and constraints’ is central to this paper and especially to our companion article (MacKenzie and Spears, 2014).

There is a particular tension that characterizes much of the history of Gaussian copula models. On the one hand, during the period on which we focus (from the late 1980s to the present), the modelling of financial derivatives was, as suggested above, characterized

increasingly by a dominant approach. On the other hand, there was a pressing need to evaluate a class of products known as Collateralized Debt Obligations (CDOs, explained below). The tension arose because CDOs could not initially be evaluated by models of the kind that were highly regarded in the dominant culture. The ensuing bricolage involved in the construction of Gaussian copula models thus took place in an uneasy interstitial space in which both practical demands and intellectual – and on occasion, perhaps even aesthetic – preferences played important roles.

The importance of intellectual preferences is part of what we want to highlight by emphasizing the presence here of an evaluation *culture*. We intend the term to signal a phenomenon similar to that captured by recent uses of ‘culture’ in science studies, in which the concept has been employed to express the pervasive finding that scientific practices (even within the same discipline at the same point in time) are not uniform: there are different ‘local scientific cultures’ (Barnes, Bloor and Henry, 1996), ‘experimental cultures’ (Rheinberger, 1997), ‘epistemic cultures’ (Knorr Cetina, 1999), ‘epistemological cultures’ (Keller, 2002) and ‘evidential cultures’ (Collins, 2004).²

Such differences in practices are found in finance too, as Smith (1999), for example, demonstrates in the case of the US stock market. An appropriate term for the more coherent and more distinct of clusters of practices that we focus on here is ‘evaluation cultures’ because evaluation – determining the economic worth and risks of financial instruments – is the activity at their core. An evaluation culture, as we use the term, is an at least partially shared set of practices, preferences, forms of linguistic or non-linguistic communication,

² Broadly similar invocations of ‘culture’ in economic life include the ‘cultures of economic calculation’ (Kalthoff, 2006) and ‘calculative cultures’ (Mikes, 2009).

meanings and beliefs, which perhaps includes an ontology or a distinctive set of assumptions about what ‘the economic world’ is made of, together with a mechanism of socialization into those practices and beliefs.³ Crucially, to count for us as an evaluation culture, such a set must go beyond the boundaries of any particular bank or other financial organization. Evaluation cultures ‘cross-cut’ organizations; as indicated very schematically in Figure 1, they are a different form of social patterning.⁴ Therein lies much of their importance; an evaluation culture can offer, for example, a route to career advancement complementary to internal promotions, as those who gain a good reputation with their counterparts in other financial institutions can (very) profitably move from one organization to another. The resultant circulation of personnel, along with industry meetings, training courses and other mechanisms, often make an evaluation culture’s practitioners personally known to each other, even in a large financial centre such as New York or London.⁵

Invoking ‘culture’ involves notorious potential pitfalls, of which perhaps the most serious is cultural essentialism. An evaluation culture is emphatically not an essentialist ‘package’ that is ‘coherent inside and different from what is elsewhere’ (Mol, 2002: 80). Cultures change, and the process of change is historically situated bricolage, not the unfolding of an internal logic. Cultures blend into each other, and borrow from each other. For example, the crucial step that separated the full-fledged Gaussian copula models of the

³ A strong form of evaluation culture may also convey an identity. In other words, it might involve not just what participants do but also whom they take themselves to be. That was not an issue we set out to investigate in this research, but our interviewees did occasionally employ formulations suggestive of identities, such as ‘you’re either a risk-neutral person [i.e. one who works with risk-neutral or martingale probabilities: see below] or you’re not.’

⁴ Figure 1 is of course not to be interpreted too literally. Neither cultures nor, indeed, organizations have clear boundaries; see the delightful parable in Hines (1988).

⁵ Beunza and Stark (2004) and Lépinay (2011) discover large differences in practices within the organizations they study, for example differences amongst trading ‘desks’ (subgroups). In our experience, intra-organizational differences of this kind are often manifestations of the intersection of evaluation cultures and organizations. For example, the practices of derivatives groups in a bank typically differ greatly from those of its ABS (asset-backed securities) desk, while the practices of derivatives groups often quite closely resemble those of derivatives groups in other banks.

2000s from earlier models ('one-period' models, as they are called) can be understood as an act of cultural borrowing: David X. Li, the modeller on whom Salmon focuses in his article for *Wired* (Salmon, 2009), drew upon a quite different mathematical tradition, that of actuarial science. Even the Black-Scholes model, which was paradigmatic for the evaluation culture on which we focus, was itself the result of bricolage (MacKenzie, 2003).

While the potential pitfall of cultural essentialism needs to be avoided, it should not stop us using a useful concept. 'Culture' suggests the richness and the importance to analysis of ideas and embodied practices, as opposed to the simplistic rational-choice terms in which financial markets are often cast. 'Culture' is implicitly plural (there are multiple cultures) and it highlights the role of even partially shared intellectual resources in communication and co-ordinated action, topics to be discussed in our companion article. Cultures, furthermore, are 'material cultures' shaped in part by the contingencies and affordances of physical artefacts and embodied practices. For example, it is important to our analysis that Gaussian copula models were not abstract ideas but programs running on computers that are physical machines consuming electricity and generating heat. Like Knuuttila and Voutilainen (2003), we take a 'material view' of models. For instance, as our companion article will show, the exigencies of the computational implementation of Gaussian copula models limited the extent to which they could be used in communication. The 'resources ... and constraints' (Sismondo, 1999: 254) with which modelling practices articulate thus include technological resources and physical constraints.

The history related in this paper is based on two sets of primary sources. The first is documentary sources, especially the extensive technical literature on Gaussian copulas. Although banks usually tried to keep new developments in modelling private for some time, secrecy

normally did not last long. As noted above, people move between organizations. ‘Quants’, as modellers are known in finance, are typically educated to the PhD level, and many retain something of an academic habitus and want their peers to know about their breakthroughs. There are also incentives of a quite different kind for banks to share models (MacKenzie and Spears, 2014). Major new developments of the Gaussian copula family thus became known relatively quickly, and the quants involved were often willing and able to publish detailed descriptions of them in specialist outlets, especially *Risk* magazine. Wider reporting in specialist trade magazines such as *Risk* forms another set of documentary sources.

Our second set of sources is 29 interviews with quants developing and/or using models such as Gaussian copulas or their rivals. These interviews are part of a larger corpus of 114 interviews with people involved (as traders, brokers, regulators, rating agency staff, etc.) in this modelling or in the market for the financial instruments – known generically as ‘credit derivatives’ – that Gaussian copulas were employed to model.⁶ The interviews took a broadly oral-history form: we led interviewees through their careers in modelling or in credit derivatives, with the goal of understanding the development of the market, the development and uses of models, and – crucially – interactions between the two. We use these interviews to construct a historical narrative of the development and changing uses of the Gaussian copula, crosschecking them against documentary sources. We also analyze the attitudes to the Gaussian copula expressed by interviewees (see the second section of MacKenzie and Spears, 2014). Because of the sensitivity of the topic, nearly all the interviews on which we draw have to remain anonymous. The exception is our 2001 interview with Oldrich Vasicek, the developer of the first Gaussian copula model in finance, conducted as part of an earlier project (MacKenzie, 2006).

⁶ 101 of these interviews were conducted by MacKenzie, 10 by Spears, and three by our colleague Iain Hardie. Other articles drawing on this wider corpus include MacKenzie (2011), which focuses on credit rating, and MacKenzie (2012), which focuses on the subprime credit-derivatives indices known as the ABX.

Given Vasicek's specific role, anonymity is impossible. We were unable to interview David X. Li, but were able to put questions to him by email. As with Vasicek, his role also makes anonymity impossible.

The main methodological difficulty we face is 'hindsight bias'. Credit derivatives were at the heart of the credit crisis that erupted in the summer of 2007, and although Salmon's argument needs to be qualified, the Gaussian copula was implicated. Those interviewed after the crisis may well be influenced by a desire to avoid blame. Fortunately, however, we had conducted a reasonable number of pilot interviews prior to the crisis; 29 of the 114 interviews, including eight of the 29 interviews with quants, took place before summer 2007. This is important to our analysis of criticisms by quants of the Gaussian copula family of models (MacKenzie and Spears, 2014).

The evaluation culture on which we focus, which we call the culture of no-arbitrage modelling, is outlined in the second section of this article. We discuss the way in which the culture organizes its activities in relation to an ontology of probabilities ('risk-neutral' or 'martingale' probabilities) invisible to others. We emphasize the culture's close connections to hedging practices in banks' derivatives departments (practices utterly central to their work: see Lépinay [2011]). The third and fourth sections turn to the history of the Gaussian copula family of models. We refer to a 'family' because the Gaussian copula is not a single, unitary model. It has been developed mathematically in different ways by different people in different contexts. Indeed, the modellers we discuss in the third section did not explicitly employ copula functions; only after such functions were introduced in this area by Li were earlier one-period models seen as Gaussian copulas. Section four shows how Li imported the idea of a copula function (an idea explained in that section and in Appendix 2) from actuarial

science. The section also sketches differences between the two most important organizational contexts in which Gaussian copulas were used: the credit rating agencies (Moody's, Standard & Poor's and Fitch) and the derivatives departments of investment banks, our focus here. The section ends by sketching the origins of the canonical (though still not entirely unitary) set of Gaussian models in investment banking: Gaussian copula base correlation models.

The culture of no-arbitrage modelling

As Kuhn (1970) emphasized, scientific cultures often coalesce around exemplary achievements. If the Nobel prize (the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel) can be used to gauge exemplary achievements, it is worth noting that it was awarded to two of the three individuals who developed the theory of options, Myron Scholes and Robert C. Merton, who, along with Fischer Black were associated with MIT and were early specialists in the nascent field of financial economics (Black and Scholes, 1973; Merton, 1973).

An option, one kind of derivative, is a contract or security that confers a right but not an obligation, for example to buy a set quantity of an underlying asset (such as a block of shares) at a fixed price (the so-called 'exercise price') at a given future time. One might expect that the price of an option should depend on expectations about whether the price of the underlying asset is going to rise or fall. On the Black-Scholes model, however, that is not so. The price of an option is determined by arbitrage, by the fact that the prices of two things that are worth the same – entitlements to identical cash flows – must be equal. If not there is an opportunity for riskless profit through arbitrage: one buys the cheaper thing and sells the dearer one, pocketing the difference.

In developing this ‘no-arbitrage’ model, Black, Scholes and Merton used what had by the early 1970s become the new specialism’s standard model of share-price movements: geometric Brownian motion. Brownian motion is the random movement of tiny particles, for example of dust and pollen, that results from collisions with the molecules of the gas or liquid in which they are suspended. The standard mathematical-physics model of this had been imported into finance. Given geometric Brownian motion and other simplifying assumptions (for example of a ‘frictionless’ market, in which both the underlying asset and riskless bonds can be bought or sold without incurring brokers’ fees or other transaction costs), Black, Scholes and Merton showed that it was possible to create a perfect hedge for an option, in other words a position in the underlying asset and in riskless bonds that, if adjusted appropriately, would have the same payoff as the option whatever the path followed by the price of the asset. Since the option and perfect hedge have identical payoffs, the price of the option must equal the cost of the hedge, or else there is an opportunity for arbitrage. This simple argument determines the price of the option, and nowhere in the formula for such a price is there any reference to investors’ attitudes to risk or beliefs about whether the price of the underlying asset was going to rise or fall.

The Black-Scholes model could have been taken as simply a surprising result about an unimportant security, because options were not central to finance in the early 1970s. Even as modelling of this kind was adopted in investment banking, it initially often was thought of as simply a cluster of loosely similar practices, sometimes called ‘the PDE approach’, because it typically involved finding a way of translating a problem in derivatives pricing into a partial differential equation akin to the canonical Black-Scholes equation.⁷ Gradually, however, ‘the PDE approach’ was supplanted by a more systematic conceptualization of no-arbitrage modelling based on the work of Stanford University applied mathematician and operations

⁷ The textbook that best exemplifies the PDE approach, especially in its early editions, is Hull (2000).

researcher, Michael Harrison, his economist colleague David Kreps and a former Stanford PhD student, Stanley Pliska. They proved the two propositions about arbitrage-free, ‘complete’ markets that have become known as the ‘fundamental theorems of asset pricing’, and in so doing introduced the idea, key to the ontology of no-arbitrage modelling, of ‘martingale probabilities’.⁸

Let us give a flavour of martingale probabilities by using ‘the parable of the bookmaker’, with which Martin Baxter and Andrew Rennie (quants at Nomura and Merrill Lynch, respectively) began an early textbook organized around martingale theory (Baxter and Rennie, 1996). Consider a race between two horses, and a bookmaker who knows the actual probabilities of each horse winning: 0.25 for the first horse and 0.75 for the second. The bookmaker could therefore set the odds on the first horse at ‘3-1 against’, and on the second at ‘3-1 on’. (Odds of ‘3-1 against’ mean that if a punter bets \$1 and wins, the bookmaker pays out \$3 plus the original stake. ‘3-1 on’ means that if a bet of \$3 is successful, the bookmaker pays \$1 plus the original stake. In this simplified parable, the adjustments to the odds necessary for the bookmaker to earn a profit are ignored.) Imagine, however, that ‘there is a degree of popular sentiment reflected in the bets made’, for example that \$5,000 has been bet on the first horse and \$10,000 on the second (Baxter and Rennie, 1996: 1). Over the long run, a bookmaker who knows the actual probabilities of each outcome and sets odds accordingly will break even, no matter how big the imbalance in money staked, but in any particular race he or she might lose heavily. However, a quite different strategy is available. The bookmaker can set odds not according to the actual

⁸ First, Harrison, Kreps and Pliska showed that a market is free of arbitrage opportunities if and only if there is an equivalent martingale measure, a way of assigning new, different probabilities (‘martingale’ probabilities) to the path followed by the price of an underlying asset such that the price of the asset (discounted back to the present at the riskless rate of interest) ‘drifts’ neither up nor down over time, and the price of the option or other ‘contingent claim’ on the asset is simply the expected value of its payoff under these probabilities, discounted back to the present. Second, that martingale measure is unique if and only if the market is complete, in other words if the securities that are traded ‘span’ all possible outcomes, allowing all contingent claims (contracts such as options whose payoffs depend on those outcomes) to be hedged with a self-financing replicating portfolio of the type introduced by Black, Scholes and Merton (Harrison and Kreps, 1979; Harrison and Pliska, 1981).

probabilities but according to the amounts bet on each horse; in this example, ‘2-1 against’ for the first horse, and ‘2-1 on’ for the second. Then, ‘[w]hichever horse wins, the bookmaker exactly breaks even’ (Baxter and Rennie, 1996: 1). As a probability theorist would put it, by adopting this strategy the bookmaker has changed ‘the measure’, replacing the actual probabilities of each outcome (a quarter and three-quarters) with probabilities that ensure no loss (a third and two-thirds). Those latter probabilities are the loose analogue of the ‘martingale’ probabilities introduced to finance by Harrison, Kreps and Pliska.

The diffusion from academia to banking of the martingale approach was pivotal to no-arbitrage pricing ceasing to be simply a cluster of mathematical practices and becoming in our terminology an evaluation culture. The shift in measure from actual probabilities to martingale probabilities (or ‘risk-neutral’ probabilities, as they are sometimes called) is common practice in the derivatives department of investment banks. The mathematics of derivatives pricing is then conducted not in the world of actual probabilities but in a world with a different ontology, the world of martingale probabilities which are simultaneously less real and more real than actual probabilities. Martingale probabilities are less real in the sense that they do not correspond to the actual probabilities of events, but are more real in the sense that (at least in finance) actual probabilities cannot be determined while martingale or risk-neutral probabilities can be calculated from current market prices. (Similarly, a bookmaker cannot actually know the true probabilities of the outcomes of a race, but can easily calculate how much punters have bet with him or her on each horse.) As an interviewee put it, martingale probabilities ‘have nothing to do with the past [they are not based on the statistical analysis of past events] or the future [they are not the actual probabilities of events] but are simply the recoding of ... prices.’

Working with martingale probabilities in a world in which prices change continually through time requires specialist training, because the underlying mathematics – ‘Brownian integrals’, or more generally stochastic calculus (the mathematics of random processes in continuous time) – is not part of standard university mathematics curricula. Socialization into the practices of no-arbitrage modelling was originally quite localized. At MIT, Robert C. Merton taught a notoriously mathematically demanding graduate course, described to us by two of his students. In the 1990s, however, such modelling was incorporated into textbooks (e.g. Baxter and Rennie, 1996) and into newly-created Masters courses in mathematical finance. ‘[T]here was an influx of people who were not scared of performing Brownian integrals and so on’, said an interviewee who worked in investment banking in the 1990s and 2000s. ‘I think it [no-arbitrage modelling using martingale probabilities] just generally became the de facto way of doing it [derivatives pricing].’

The adoption of no-arbitrage modelling was encouraged by a rough homology between it and financial practices in the derivatives departments of banks. The emphasis on hedging in the investment bank studied by Lépinay (2011) is consistent with our interviews; despite the widespread impression of reckless risk-taking that the crisis created, derivatives departments seek carefully to hedge their portfolios. Such hedging is incentivized by how traders are paid (see the discussion in our companion paper of ‘Day 1 P&L’ where ‘P&L’ is the acronym of profit and loss), and analyses of the exposure of derivatives portfolios to the risks of changing prices, interest rates, etc., are part of a daily routine that several interviewees described. The necessary modelling is very demanding computationally, even when grids of hundreds or thousands of interconnected computers are devoted to it. So the necessary risk-analysis programs are typically run overnight, while during the trading day no-arbitrage modelling is applied primarily in pricing (see Lépinay [2011] on ‘pricers’, which are software programs that run the necessary models). In

pricing, all the complication of no-arbitrage modelling and martingale probabilities reduces to a simple precept: ‘price is determined by hedging cost’ (McGinty et al., 2004: 20). It is the strategy of Black-Scholes modelling writ large: find a perfect hedge, a continuously-adjusted portfolio of more basic securities that will have the same payoff as the derivative, whatever happens to the price of the underlying asset or assets; use that portfolio to hedge the derivative; and use the cost of the hedge as a guide to the price of the derivative. In actual practice, of course, few if any hedges are actually perfect, and the price quoted to an external customer will be greater than that hedging cost, the difference generating the bank’s profit and the trader’s hoped-for Day 1 P&L.

The crucial role of a no-arbitrage model as a guide to hedging generates for traders a practical criterion of a good model. If they implement the hedges implied by the model, the profitability of the resultant trading position should be ‘flat’ (constant), indicating (e.g. to risk controllers) that the position is indeed hedged, that risks are fully controlled. P&L should not ‘swing too much’, said an interviewee: ‘That is what it is always about.’ Nevertheless, not all of the preferences of quants are purely pragmatic. An approach that can encompass the modelling of a huge range of complex derivatives but yet be boiled down to the two simple theorems formulated by Harrison, Kreps and Pliska fits well with the preferences for ‘elegance’ of many of those with advanced mathematical training. ‘The simplicity of it is alluring,’ said an interviewee. In the middle of their textbook, Baxter and Rennie, after recasting the derivation of the exemplary achievement, the Black-Scholes model, in the more general framework of martingale theory, paused:

with a respectable stochastic model for the stock [geometric Brownian motion], we can replicate any [derivative]. Not something we had any right to expect. ... Something subtle and beautiful really is going on under all the formalism. ... Before we push on, stop and admire the view. (Baxter and Rennie, 1996: 98)

By now, perhaps, the reader may feel our two articles are a badly telegraphed murder mystery: the culprit in the financial crisis is surely this strange, abstract culture, with its capacity to see things – martingale probabilities – of which others are unaware, and even to appreciate them as beautiful. Not so. The Gaussian copula family of financial models drew upon resources *from* that evaluation culture, but was never entirely *of* that culture.

The origins of the Gaussian copula

The first of what is now seen as the Gaussian copula family of models in finance was developed between 1987 and 1991 by Oldrich Vasicek, a probability theorist and refugee from the Soviet invasion of Czechoslovakia. He was hired in late 1968 by John McQuown, head of the Management Science Department of Wells Fargo in San Francisco. McQuown was a strong supporter of the new field of financial economics, hiring leading figures such as Black and Scholes as consultants, and financing conferences at which members of the bank's staff such as Vasicek were 'able to sit in and listen, wide-mouthed' (Vasicek, interview). Those conferences and his work for the bank introduced Vasicek to the Black-Scholes model and to Merton's use of stochastic calculus. In 1983, McQuown persuaded Vasicek to join him in a new venture, a firm called Diversified Corporate Loans. Banks' loan portfolios are often 'very ill-diversified', as Vasicek puts it – heavily concentrated in specific geographical regions or particular industries – and McQuown's idea was to enable banks to reduce these concentrations of risk by swapping 'loans that the bank has on its books for participation shares' in a much larger pool of loans originated by many banks (Vasicek, interview).

‘[I]t didn’t work’, says Vasicek – banks did not take up the idea. However, his modelling for it gave birth to what is to our knowledge the first Gaussian copula model in finance (although it was only retrospectively seen as a Gaussian copula). In order that the terms of the swap could be negotiated, it was necessary to model the risks both of default on a loan to one corporation and of multiple defaults in the bigger pool of loans. Financial economists, especially Merton (1974), had tackled the first problem, but not the second. It was immediately clear to Vasicek that defaults by different corporations could not plausibly be treated as statistically independent events. As he put it in an unpublished note, now in his private papers:

The changes in the value of assets among firms in the economy are correlated, that is, tend to move together. There are factors common to all firms, such as their dependence on [the] economy in general. Such common factors affect the asset values of all companies, and consequently the loss experience on all loans in the portfolio. (Vasicek, 1984: 9)

The task Vasicek set himself, therefore, was to model the value of a pool of loans to multiple corporations, taking account of the correlation between changes in the values of different firms’ assets. There was almost no direct empirical data to guide his modelling (even twenty years later the econometric estimation of the relevant correlations was still tricky; see MacKenzie [2011]). Given this absence of data, Vasicek simply reached for the standard model of asset-value fluctuations, geometric Brownian motion, imposing the mathematically most familiar correlation structure, that of a multivariate Gaussian distribution, the analogue for multiple variables of the familiar, bell-shaped, univariate normal distribution (see Appendix 1). Even with those simple choices, however, he could not find a general ‘analytical’ solution to his model, a solution that

would avoid recourse to computer simulation. Indeed, no analytical solution has subsequently been found.

Vasicek did, however, succeed in formulating an analytically solvable special case, that of a pool of a large number of equally sized loans, all of which fall due simultaneously, have the same probability of default and have the same correlation between the values of the assets of any pair of borrowers. This set of features leads to Vasicek's special case often being called the 'large homogeneous pool' model. He showed that as the number of loans increases the probability distribution of different levels of loss on the portfolio in a given, single time-period converges to equation 2 of Appendix 1 below. This equation generated Figure 2 of this paper.

The figure captures a crucial feature of all Gaussian copula models, which is the radical differences in the shape of the probability distribution of losses at different correlation levels. It is based on applying Vasicek's model to a large pool of homogeneous loans, each with a default probability of 0.02. (This corresponds roughly to a typical estimate of the probability of a firm with a low investment-grade rating such as BBB defaulting, with the time-period in a question being the coming five years.) The expected level of loss on the pool is in each case the same; it is just the probability of default on any individual loan, 0.02. If correlation is low (e.g., 0.1), the probability distribution of losses on the portfolio clusters reasonably tightly around this expected loss, while if correlation is higher the probability distribution 'spreads out' more; the probability of losing very little increases, but so does the probability of a loss markedly higher than 0.02. If correlation is very high indeed (e.g., 0.99), the probability distribution becomes bimodal ('twin-peaked'), with a palpable risk of almost complete loss; the entire pool is starting to behave like a single asset.

Vasicek's work was not published at the time. Modelling credit risk (the risk of borrowers defaulting) was critical to the business of Diversified Corporate Loans and KMV, which was Vasicek and McQuown's next, more successful firm. However, Vasicek's derivation of the large homogeneous pool model (some thirty lines of maths) did circulate privately. Li, for example, recalls seeing it in the form of a photocopy of a handwritten original, probably Vasicek's (email to MacKenzie, 24 May 2008). In particular, Vasicek's work was drawn upon by the developers of the first Gaussian copula model to be adopted at all widely in finance, J.P. Morgan's 1997 software system to measure credit risk, CreditMetrics (again, only in retrospect was this single-period model seen as a Gaussian copula).

The manager assigned to CreditMetrics originally wanted simply to hire Vasicek's firm, KMV, to produce the system for J.P. Morgan. However, negotiations 'with [KMV] were ponderous', said an interviewee. J.P. Morgan eventually produced CreditMetrics almost entirely in-house, but did 'engage ... KMV to work with us to build the correlation module.' CreditMetrics employed what was in effect the same overall mathematical framework as Vasicek had: fluctuations in the market value of each firm's assets were modelled as if driven by geometric Brownian motion, and interdependence among those fluctuations was modelled by imposing a multivariate Gaussian correlation structure. However, KMV had already had to abandon ('kicking and screaming' because 'they really loved the Vasicek closed form', as this interviewee put it) the radical simplifications of Vasicek's analytically solvable special case, which were judged too restrictive for practical use: who would believe a model in which all firms have the same probability of bankruptcy?

So both KMV and J.P. Morgan turned instead to computer simulation. CreditMetrics was implemented via a technique widely used in science, engineering and elsewhere: Monte Carlo

modelling (see Galison, 1997). Correlated, normally distributed (i.e. Gaussian) random numbers were used in CreditMetrics's software to generate a very large number of 'scenarios', and the corporate defaults in each of the scenarios were aggregated to form an estimate of the loss in that scenario, with the probability distribution of different levels of loss on the overall pool calculated by aggregating across all the scenarios (Gupton, Finger and Bhatia, 1997). It was far more demanding computationally than simply plugging numerical values into the equations expressing the analytical solution of Vasicek's special case – hundreds of thousands of Monte Carlo scenarios might be needed to achieve statistically stable estimates – but simpler to understand. CreditMetrics could be understood, at least in outline, by anyone who could grasp the idea of using correlated, normally distributed, random numbers to simulate statistical dependence among defaults by different corporations.

Broken hearts, corporate defaults and investment banks

Both Vasicek's special case – the large homogeneous pool – and CreditMetrics were, as noted, 'one-period' models: although the underlying stochastic processes took place in continuous time, all that was modelled was *whether* corporations defaulted within a single, given time period, and not *when* in that period they defaulted. It is at this point that the focus of Salmon's (2009) *Wired* article, David X. Li, enters the story. Li was brought up in rural China where his family lived because of the Cultural Revolution, and moved to Canada in the early 1990s for a Masters in Actuarial Science and a PhD in Statistics at the University of Waterloo. After a session (1994-5) teaching actuarial science and finance at the University of Manitoba, he worked as a quant, first at the Royal Bank of Canada and then Canadian Imperial Bank of Commerce, where he modelled 'credit derivatives' such as the CDOs discussed below.

‘I was aware of Vasicek[‘s] work’, Li told MacKenzie in an email message (24 May 2008). ‘I found that was one of the most beautiful math I had ever seen in practice. But that was a one period framework.’ The yields of a corporation’s bonds, or the prices of the new credit default swaps, could, however, be used to model the ‘survival time’ of an individual corporation (in other words, the time until it defaults). So, as Li put it, ‘the problem becomes how to specify a joint survival time distribution with marginal distribution [the probability distribution of the survival time of each individual corporation] given.’

To solve this problem, Li drew on a cultural resource not from financial economics but from actuarial science and ultimately mathematical statistics: copula functions. While at the University of Manitoba, Li had co-taught with the research actuary Jacques Carriere. Carriere was collaborating with Jed Frees of the University of Wisconsin and Frees’s doctoral student Emiliano Valdez on the problem of the valuation of joint annuities, in particular annuities in which payments would continue to be made to one spouse if the other died (email to MacKenzie from Frees, 23 January 2012). When pricing joint annuities, standard practice in insurance was simply to assume that the deaths of a wife and of a husband were statistically independent events. ‘With this assumption, the probability of joint survival is the product of the probability of survival of each life’ (Frees, Carriere and Valdez, 1996: 230). However, it was known empirically that the death of one spouse could increase the chances of death of the other, a phenomenon ‘often called the “broken heart” syndrome’ (Frees, Carriere and Valdez, 1996: 230).

To model broken-heart syndrome, Frees, Carriere and Valdez used copula functions, an approach developed in the 1950s by the Illinois Institute of Technology mathematician, Abe Sklar, which was attracting increasing attention in mathematical statistics in the 1990s. A copula function is a way of ‘coupling’ a set of marginal distribution functions (in the case of the

mortality of spouses, the function that specifies the probability that the wife will die at or before a given age, and the separate function that specifies the probability that the husband will die at or before another age). This coupling results in the ‘joint’ or ‘multivariate’ distribution function (see Appendix 2). Frees and his colleagues showed that taking into account the ‘correlation’ between the mortality of spouses by using an appropriate copula function reduced the value of a joint annuity by around 5 percent (Frees, Carriere and Valdez, 1996: 229).

Their work gave Li a way of linking his training in actuarial science and statistics to the practical problems of pricing CDOs and similar credit derivatives on which he was working. He treated a corporation’s default as analogous to a person’s death: the risks of different corporations defaulting were correlated, just as the mortality risks of spouses were, and copula functions could be used to model that correlation. This approach enabled Li to escape the limitation to a single period of Vasicek’s special case and CreditMetrics, while still retaining a connection to them. Viewed in the lens of Li’s work (Li, 1999, 2000), the model of correlation in them was a Gaussian copula, or a copula function that couples marginal distributions to form a multivariate normal distribution function. Although other copula functions were discussed by Li and by others who were also exploring the applicability of copula functions to insurance and finance (such as a group of academic mathematicians in Zürich with strong links to the financial sector),⁹ this connection to CreditMetrics – already a well-established, widely-used model – together with the simplicity and familiarity of the Gaussian (and the ease of implementing it: commercially available programs facilitated the use of correlated, normally-distributed random numbers in Monte Carlo simulation) meant that as others took up copula functions, the Gaussian copula had the single most salient role.

⁹ See Embrechts, McNeil and Straumann (1999).

Although it was also used to measure banks' overall credit risks, the most consequential modelling problem to which the Gaussian copula was applied was the evaluation of collateralized debt obligations (CDOs), a new class of securities becoming increasingly popular in the late 1990s and 2000s (see Figure 3). The firm, typically a large investment bank, creating a CDO would set up a legally-separate 'special-purpose vehicle' (a trust or special-purpose corporation) that would buy a pool of bonds or loans, raising the money to do so by selling investors interest-bearing securities that were claims on the cashflow generated from the pool. The lower 'tranches' of securities offered higher 'spreads' (interest payments were typically set as a given 'spread' or increment over a baseline interest rate, normally Libor, the London Interbank Offered Rate), but also greater risk of partial or complete loss on the investment if bonds or loans in the pool defaulted. For instance, the lowest tranche absorbed the first losses, and only once that tranche was 'wiped out' by these losses did they begin to affect the next-highest, mezzanine, tranche. In a typical CDO, if correlation amongst the bonds or loans in the pool was low, only the holders of the lowest tranche would be at substantial risk of losing some or all of their investment. If, however, correlation was very high (as in the 0.99 case in Figure 2), many of the bonds or loans might default, and losses could affect even the holders of the most senior tranche. So modelling correlation was the most crucial problem in CDO evaluation, and Gaussian copulas became – and still are (MacKenzie and Spears, 2014) – the canonical way of doing this.

As noted above, Li's work freed the Gaussian copula family from the restriction of earlier models to a single time period. It did not, however, free Gaussian models from the other chief limitation: in practical applications, no 'analytical' solution akin to that of Vasicek's large homogeneous pool could be found, so computationally intensive Monte Carlo simulation was needed. The consequences of this were very different in the two main contexts in which the

Gaussian copula was used to evaluate CDOs. In the credit rating agencies, the task was to assign a rating (BBB, AAA, etc.) to each tranche of a CDO by working out the probability of default on that tranche (or, in the case of Moody's, the expected loss on the tranche). For that task, a Monte Carlo Gaussian copula model akin to CreditMetrics was judged to be perfectly adequate. For example, when Standard & Poor's introduced such a model, *CDO Evaluator*, in November 2001, it reported that the simulation time needed to run 100,000 Monte Carlo scenarios on a PC was around two and a half minutes (Bergman, 2001). The several minutes it takes to run the scenarios were not salient; CDOs are complicated structures and assessing their legal and cash-flow aspects requires much more time. Nor was moving beyond one-period models to Gaussian copulas in Li's sense seen in the rating agencies as an urgent priority. Standard & Poor's made the move only with version 3.0 of *Evaluator*, released in December 2005, while Fitch simply kept using its one-period Gaussian model, *Vector*, analysing a multi-year CDO by running *Vector* for the first year and then again for the second year, and so on. (Moody's also seems to have stuck with one-period Monte Carlo formulations, although our interviews do not contain detailed information on practice at Moody's in this respect.)

The situation in the other main context, investment banking, was quite different. When CDOs first started to become a relatively large business, in the late 1990s, evaluating a CDO on a 'once and for all' basis, akin to practice at the rating agencies, was adequate (typically, the risks of all but the equity tranche were sold on to external parties), and CreditMetrics or similar one-period models were judged up to the job. In the early 2000s, however, new versions of CDOs became popular, of which the most important were 'synthetic' single-tranche CDOs. Instead of consisting of a special-purpose legal vehicle that bought a pool of debt instruments, these new CDOs were simply bilateral contracts between an investment bank and a client (such as a more minor bank or other institutional investor); such contracts mimicked the returns and risks of a

CDO tranche. They became popular because the CDO tranches in heaviest demand – mezzanine tranches – formed only a small part of the structure of a traditional ‘cash’ CDO of the kind shown in Figure 3, and were therefore in short supply.

For the client institution that bought a synthetic CDO tranche, the latter was a security that paid a decent rate of interest with a modest risk of default. For the investment bank that sold the tranche, it was a complex derivative. Because the bank did not own the pool of loans or bonds underpinning the contract (the pool was hypothetical, simply a way of calculating the interest payments the bank had to make to the client and the losses the latter might suffer), it had to find other means of hedging itself against losses throughout the contract’s lifetime. This involved the use of credit default swaps on each of the corporations whose debts made up the pool. In a credit default swap contract on a corporation – on Ford, for example – one financial institution pays set ‘insurance premiums’ to a second financial institution in return for the right, if Ford defaults on its loans or bonds, to hand over those loans or bonds to the second institution and receive their full face value. Hedging a synthetic CDO tranche using credit default swaps was roughly analogous to hedging an option with a position in the underlying shares, and the hedge ratios that were needed were christened ‘deltas’, the term already used in the options market.

The need to adjust credit-default-swap hedges of this sort, often daily, throughout the lifetime of a synthetic CDO – typically five, seven or ten years – meant that Gaussian copula models in investment banks had to satisfy demands quite different from those of the ‘one off’ analyses conducted by credit rating agencies. Single-period CDO models akin to CreditMetrics were not well suited to the calculation of deltas, so following Li’s work there was rapid, sustained interest in investment banking in full-fledged copula formulations. The need to

recalculate deltas and other risk parameters frequently meant that the computational demands of Monte Carlo simulation were a major problem for investment banks, not the minor one they were for rating agencies; extracting reliable estimates of a large set of partial derivatives such as deltas from a Monte Carlo copula model was vastly more time-consuming than using the model to rate a CDO tranche. In a situation in which the IT departments of many big banks were struggling to meet the computational demands of the overnight runs – ‘some days, everything is finished at 8 in the morning, some days it’s finished at midday because it had to be rerun’, an interviewee told MacKenzie in early 2007 – the huge added load of millions of Monte Carlo scenarios was unwelcome. The requisite computer runs sometimes even had to be done over weekends. An interviewee described one bank in which the Monte Carlo calculation of deltas took over forty hours. Such difficulties could be alleviated by distributing the necessary computations over grids of hundreds of interconnected computers, but there were often down-to-earth material and physical constraints on the size of grid: the finite capacity of computer room air-conditioning systems to cope with the resultant heat, and in some places (especially the City of London) constraints on electricity supply.

The innovative efforts of investment-bank quants were therefore focussed on developing what were christened ‘semi-analytical’ versions of the Gaussian and other copulas. These involved less radical simplifications than Vasicek’s model with its ‘analytical’ solution (equations 1 and 2 in Appendix 1), while being sufficiently tractable mathematically that Monte Carlo simulation was not needed. Much faster computational techniques such as numerical integration, Fourier transforms, and recursion sufficed. A commonly used simplification was introduced by, among others, Jon Gregory and Jean-Paul Laurent of the French bank BNP Paribas, first in a confidential May 2001 BNP working paper and then in Gregory and Laurent

(2003; see also Laurent and Gregory, 2005).¹⁰ The simplification was to assume that the correlations among the asset values or default times of the corporations in a CDO's pool all arose simply from their common dependence on one or more underlying factors. Most common of all was to assume a single underlying factor, which could be interpreted as 'the state of the economy'. The advantage of doing this was that given a particular value of the underlying factor, defaults by different corporations could then be treated as statistically independent events, simplifying the mathematics, avoiding Monte Carlo simulation and greatly reducing computation times.

'Factor reduction' (as this was sometimes called) and other techniques – such as, for example, a recursion algorithm introduced by the Bank of America quants Leif Andersen, Jakob Sidenius and Susanta Basu (Andersen, Sidenius and Basu, 2003) – made it possible for 'single-factor' Gaussian copulas and other copula models to run fast enough to be embedded in the hedging and risk management practices of investments banks' derivatives departments. Such techniques quickly became common knowledge amongst quants, and single-factor, semi-analytical Gaussian copula models became pervasive in investment banking. This style of modelling, and the associated hedging practices, helped make the creation and trading of CDOs culturally familiar to derivatives specialists. As one of them told us in January 2009, it 'had all the appearance of a derivative business ... [models with] parameters that they could look at and discuss, [which] had names they were familiar with like "correlation".' His choice of word – 'appearance' – is of course telling; as we discuss in our companion article, he and others felt it was mere appearance. Nevertheless, in his view, the resemblance 'did give people some comfort'.

¹⁰ Broadly analogous factor models were also discussed, for example, by Philipp Schönbucher (2001).

One final step was left in the creation of a de facto industry-standard modelling approach. It was triggered in 2003-04 by J.P Morgan and other investment banks collectively creating a set of tranches ‘index markets’, each of which traded what was in effect a standardized synthetic CDO with a specified underlying pool of corporate debt issuers (e.g., a given set of 125 investment-grade corporations domiciled in North America). Being standardized rather than negotiated ad hoc between a client institution and an investment bank, this CDO and its tranches could readily be bought and sold by multiple participants, and so had credible market prices. Indeed, the credibility of those prices was an important motivation for the creation of the index markets (MacKenzie and Spears, 2014). Those market prices provided a new way of determining correlation, the crucial parameter of a Gaussian copula model, turning it from a difficult econometric task to an easy modelling job. One could simply assume a common level of correlation among the corporations in the pool underlying a standardized index, and run a Gaussian copula model ‘backwards’, to discover the level of correlation consistent with market prices, that is, with the traded ‘spreads’ of the tranches. Indeed, correlation itself started to be reified. It was no longer just a parameter of a model, but something ‘correlation traders’ (as they started to be called) could trade in the new markets.

A last mathematical snag remained. In the case of many standardized index tranches, especially mezzanine tranches, running a Gaussian copula model backwards yielded not a single correlation value consistent with the ‘spread’ on the tranche, but two values. In other cases, the model would simply fail to calibrate: it could not reproduce market prices, and no correlation value would be generated. Problems of this kind could be avoided, a team at J.P Morgan argued, if correlation modelling shifted to what they called ‘base correlation’ (explained in Appendix 3).

Others in the CDO market quickly saw the advantages of doing so, and use of ‘base correlation’ rapidly became pervasive in investment banking.

That 2004-05 switch was the final stage in the construction of the canonical set of models, at least in investment banking: Gaussian copula base correlation models. Unlike in earlier years, when Gaussian copula models could be set against empirical data only partially and with difficulty, the new standard-index markets provided a ready empirical test of base-correlation models: the output of the latter could be compared directly to the traded ‘spreads’ in the index markets. At one level, the models passed this test unequivocally. A Gaussian copula base correlation model ‘fits the market exactly’, as an interviewee put it. Appropriately calibrated, the model could replicate precisely the market spreads. On other levels, though, we found deep disquiet – even in our pre-crisis interviews with quants – about the newly-emerged standard models, disquiet that forms the starting point of our companion article.

Conclusion

Cultures borrow; modelling is bricolage; modelling articulates not just with data (which has not played a large role in our story) but with ‘other goals, resources, and constraints’. With the exception of Li, with his background in actuarial science, all the contributors to the development of the Gaussian copula family of models whom we interviewed owed some degree of allegiance to what we have called the culture of no-arbitrage modelling. Indeed, as our companion article will describe, that allegiance was sufficiently strong that in private some of them distanced themselves markedly from the very family of models, Gaussian copulas, to which they had contributed. In their modelling *practices*, however, they had embraced productive heterogeneity. The Gaussian copula was not a no-arbitrage model, but

they adopted and developed it nonetheless. They were bricoleurs. For all their private qualms, they went for what ‘worked’, not for what was culturally homogeneous, ‘pure’ or ‘beautiful’.

What it was for a model to ‘work’ – to articulate successfully with organizational objectives and material constraints – was historically contingent and context-specific. For example, that running a Gaussian copula model backwards to infer an ‘implied correlation’ sometimes generated two values and sometimes none was not a major problem – not a salient failure to ‘work’ – until the emergence in 2003-04 of the standardized index tranche markets, which provided the data that made it easy repeatedly to run the model backward and provided the incentive to do so; if one saw oneself as trading ‘correlation’, then one needed continually to know what its level was. More generally, what it was for a model to ‘work’ was in part a matter of the rhythms of the working day, which differed radically in different organizational contexts. In rating agencies, with a CDO needing to be analyzed only once (or at most only a limited number of times during its lifetime), a delay of a couple of minutes while a computer performed a Monte Carlo simulation did not matter. In an investment bank, with thousands of CDO tranches needing to be reanalyzed every night, small delays of that sort could aggregate disastrously, with overnight runs not completed by the next morning and traders and risk controllers therefore not knowing where they stood. As noted above, the material constraints of heat generated and finite electricity supply limited the extent to which these problems could be circumvented by distributing the computational load over multiple computers. As Mirowski (2010) notes, markets can have computational limits.

Of course, in this area material constraint is never simply physical. A bank could – and our interview data suggested at least one, Goldman Sachs, did – decide to build, or hire space in, a data centre close to a major financial centre so as to be able to use an even bigger

grid of computers (in Goldman's case, we were told, the centre is in New Jersey). But most banks didn't – 'I would not have got the budget', an interviewee told us. 'There is no way I would have had a server farm' – the bank implicitly preferring simplification of models over much increased IT spending. Material constraints thus interact with organizational goals and resources, and, as our companion article will show, issues of organization created subtler incentives and constraints as well.

Despite such constraints, the technological, organizational and financial assemblages surrounding Gaussian copula models in investment banking *did* work. They were a major source of investment banking's rapidly growing profits from the late 1990s onwards. By 2004, as much as third of all investment-bank revenue in fixed-income (bonds and bond-like products) came from credit derivatives such as CDOs (Tett, 2005). These assemblages, indeed, had a certain local stability. Those who sought to change modelling practices often found that they could get 'no traction', as one of them put it to us (our companion article explores the reasons for this). Local stability, however, is not global stability. Many of the interviews we conducted in 2006 and early 2007 were suffused by a certain unease, on the part not just of the interviewer but at least sometimes also the interviewee; something seemed wrong, but one could not put one's finger on exactly what it was. In retrospect, that was because the central mechanism of the crisis was associated with simpler uses of the Gaussian copula than the sophisticated ones in investment banking we were investigating (MacKenzie and Spears, 2014). It is also impossible now to reread the transcripts of those early interviews without a certain chill, akin to what one of us once felt when reading archival documents from late-Edwardian England (MacKenzie, 1981). The world that generated those transcripts and those documents was in each case about to end in calamity. The role of the Gaussian copula in economic disaster is one of the topics to which we turn in our companion article.

Acknowledgements

These are to be found at the end of our companion article in this journal (MacKenzie and Spears, 2014).

Appendix 1: Vasicek's large homogeneous pool model

Vasicek applied to firms' asset values what had become the standard geometric Brownian motion model. Expressed as a stochastic differential equation,

$$dA_i = \mu_i A_i dt + \sigma_i A_i dz_i$$

where A_i is the value of the i th firm's assets, μ_i and σ_i are the drift rate and volatility of that value, and z_i is a Wiener process or Brownian motion, i.e. a random walk in continuous time in which the change over any finite time period is normally distributed with mean zero and variance equal to the length of the period, and changes in separate time periods are independent of each other.

Vasicek (1987, 1991) considers a portfolio of equally-sized loans to n such firms, with each loan falling due at the same time and each with the same probability of default p . Making the assumption that the correlation, ρ , between the values of the assets of any pair of firms was the same, Vasicek showed that in the limit in which n becomes very large, the distribution function of L , the proportion of the loans that suffer default, is

$$P[L \leq x] = N\left(\frac{\sqrt{1-\rho}N^{-1}(x) - N^{-1}(p)}{\sqrt{\rho}}\right) \dots (1)$$

where N is the distribution function of a standardized normal distribution with zero mean and unit standard deviation. The corresponding probability density function is:

$$f(x) = \sqrt{\frac{1-\rho}{\rho}} \exp\left(-\frac{1}{2\rho}(\sqrt{1-\rho}N^{-1}(x) - N^{-1}(p))^2 + \frac{1}{2}(N^{-1}(x))^2\right) \dots (2)$$

(Figure 2 shows graphs of this function with $p = 0.02$ and two different values of ρ .) Vasicek went on to show that the assumption of equally-sized loans was not necessary, and that this

limit result still held so long as $\sum_{i=1}^n w_i^2$ tended to zero as n became infinitely large, where w_i is

the proportion of the portfolio made up of loan i . ‘In other words, if the portfolio contains a sufficiently large number of loans without it being dominated by few loans much larger than the rest, the limiting distribution provides a good approximation for the portfolio loss’

(Vasicek, 2002: 160-1).

Appendix 2: ‘Broken heart’ syndrome and a bivariate copula function

Let X be the age at death of a woman and Y the age at death of her husband. In the notation of Frees, Carriere and Valdez (1996), let $H(x,y)$ be the joint distribution function of X and Y : i.e. $H(x,y)$ is the probability that the wife dies at or before age x , *and* that the husband dies at or before age y . Let $F_1(x)$ and $F_2(y)$ be the corresponding marginal distribution functions; e.g., $F_1(x)$ is the probability simply that the wife dies at or before age x .

A copula function C ‘couples’ (Frees, Carriere and Valdez, 1996: 236) F_1 and F_2 , the two marginal distributions, to form the joint distribution H . That is,

$$H(x,y) = C[F_1(x), F_2(y)]$$

If C , F_1 and F_2 are all known, then obviously H is known. What Sklar (1959) had shown was that a generalized version of the less obvious converse also held: ‘if H is known and if F_1 and F_2 are known and continuous, then C is uniquely determined’ (Frees, Carriere and Valdez, 1996: 236).

Appendix 3: Index tranches and base correlation

The credit indices that make ‘correlation trading’ possible are, in effect, a set of standardized, synthetic single-tranche CDOs. Consider, for instance, the DJ Tranchet TRAC-X Europe, set up by J.P Morgan and Morgan Stanley, the example of an index used in McGinty et al. (2004). Traders could buy and sell ‘protection’ against all losses caused by defaults or other ‘credit events’ suffered by the corporations whose debts were referenced by the index, or against specific levels of loss: 0-3 percent, 3-6 percent, 6-9 percent, 9-12 percent and 12-22 percent. Instead of running a Gaussian copula ‘backwards’ to work out the implied correlation (the ‘compound’ correlation, in the terminology of the J.P Morgan team) of each these tranches, the J.P Morgan team recommended inferring from the ‘spreads’ (costs of ‘protection’) on the tranches that were actually traded what the spreads would be on tranches of 0-6 percent, 0-9 percent, 0-12 percent and 0-22 percent. Running a Gaussian copula backwards on the traded 0-3 percent tranche and on these untraded tranches generates the ‘base correlations’ implied by the spreads in the index market.

Figure 1. Evaluation cultures and organizations: a schematic representation.

Figure 2. Probability distribution of losses on large portfolio of loans, each with default probability of 0.02, and identical pairwise asset correlations of 0.1 (upper graph) and 0.99 (lower graph).

Figure 3. A CDO (simplified and not to scale)

←----- capital investments by investors
 -----→ payments to investors

Investors in lower tranches receive payments only if funds remain after payments due to investors in more senior tranches are made. What is shown is a 'cash CDO'; in a 'synthetic CDO' the special purpose vehicle 'sells protection' on the assets via credit default swaps rather than buying them.

References

- Andersen L, Sidenius J and Basu S (2003) All your hedges in one basket. *Risk* 16(11): 67-72.
- Barnes B, Bloor D and Henry J (1996) *Scientific Knowledge: A Sociological Analysis*. Chicago, IL: Chicago University Press.
- Baxter M and Rennie A (1996) *Financial Calculus: An Introduction to Derivative Pricing*. Cambridge: Cambridge University Press.
- 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 and Stark D (2004) Tools of the trade: The socio-technology of arbitrage in a Wall Street trading room. *Industrial and Corporate Change* 13: 369-400.
- Black F and Scholes M (1973) The pricing of options and corporate liabilities. *Journal of Political Economy* 81: 637-654.
- Boumans M (1999) Built-in justification. In: Morrison MS and Morgan M (eds) *Models as Mediators*. Cambridge: Cambridge University Press, 66-96.
- Breslau, D and Yonay, Y (1999) Beyond metaphor: Mathematical models in economics as empirical research. *Science and Context* 12: 317-332.
- Callon M (ed) (1998) *The Laws of the Markets*. Oxford: Blackwell.
- Callon M (2007) What does it mean to say that economics is performative? In: MacKenzie D, Muniesa F and Siu L (eds) *Do Economists Make Markets? On the Performativity of Economics*. Princeton, NJ: Princeton University Press, 311-357.
- Cartwright N (1983) *How the Laws of Physics Lie*. Oxford: Clarendon.
- Collins H (2004) *Gravity's Shadow: The Search for Gravitational Waves*. Chicago, IL: Chicago University Press.
- Dowling D (1999) Experimenting on theories. *Science in Context* 12: 261-273.
- Embrechts P, McNeil A and Straumann D (1999) Correlation: Pitfalls and alternatives. *Risk* 12(5): 69-71.
- Frees EW, Carriere J and Valdez E (1996) Annuity valuation with dependent mortality. *Journal of Risk and Insurance* 63: 229-261.
- Galison P (1997) *Image and Logic: A Material Culture of Microphysics*. Chicago, IL: Chicago University Press.
- Gregory J and Laurent JP (2003) I will survive. *Risk* 16: 103-107.
- Gupton GM, Finger CC and Bhatia M (1997) CreditMetrics: Technical document. New York: JP Morgan.
- Harrison JM and Kreps DM (1979) Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20: 381-408.
- Harrison JM and Pliska SR (1981) Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications* 11: 215-260.
- Hines RD (1988) Financial accounting: In communicating reality, we construct reality. *Accounting, Organizations and Society* 13: 251-261.
- Hull JC (2000) *Options, Futures, & Other Derivatives*. Upper Saddle River, NJ: Prentice-Hall.
- Kalthoff H (2006) Cultures of economic calculation. In: Strulik T and Willke H (eds) *Towards a Cognitive Mode in Global Finance*. Frankfurt and New York: Campus, 156-179.

- Keller EF (2002) *Making Sense of Life: Explaining Biological Development with Models, Metaphors, and Machines*. Cambridge, MA: Harvard University Press.
- Knorr Cetina K (1999) *Epistemic Cultures: How the Sciences make Knowledge*. Cambridge, MA: Harvard University Press.
- Knuuttila T and Voutilainen A (2003) A parser as an epistemic artifact: A material view on models. *Philosophy of Science* 70: 1484-1495.
- Kuhn TS (1970) *The Structure of Scientific Revolutions* 2nd ed. Chicago, IL: Chicago University Press.
- Laurent JP and Gregory J (2005) Basket default swaps, CDOs and factor copulas. *Journal of Risk* 7(4): 103-122.
- Lépinay V (2011) *Codes of Finance: Engineering Derivatives in a Global Bank*. Princeton, NJ: Princeton University Press.
- Li DX (1999) The valuation of basket credit derivatives. *CreditMetrics Monitor* April: 34-50.
- Li DX (2000) On default correlation: A copula function approach. *Journal of Fixed Income* 9(4): 43-54.
- MacArthur RH and Wilson EO (1967) *The Theory of Island Biogeography*. 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 (2003) An equation and its worlds: Bricolage, exemplars, disunity and performativity in financial economics. *Social Studies of Science* 33: 831-868.
- 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 (2012) Knowledge production in financial markets: Credit default swaps, the ABX and the subprime crisis. *Economy and Society* 41: 335-359.
- MacKenzie D and Spears T (2014) 'A device for being able to book P&L': The organizational embedding of the Gaussian copula. *Social Studies of Science* 44.
- McGinty L, Beinstein E, Ahluwalia R and Watts M (2004) Credit correlation: A guide. London: JP Morgan. Available at: www.math.nyu.edu/~cousot/Teaching/IRCM/Lecture10/Base%20correlationJPM.pdf (accessed 24 June 2006).
- Mehrling P (2005) *Fischer Black and the Revolutionary Idea of Finance*. New York: Wiley.
- Merton RC (1973) Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4: 141-183.
- Merton RC (1974) On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29: 449-470.
- Mikes A (2009) Risk management and calculative cultures. *Management Accounting Research* 20: 18-40.
- Mirowski P (2010) Inherent vice: Minsky, markomata, and the tendency of markets to undermine themselves. *Journal of Institutional Economics* 6: 415-443.
- Mol A (2002) *The Body Multiple: Ontology in Medical Practice*. Durham, NC: Duke University Press.
- Morgan MS (2005) Experiments versus models: New phenomena, interference, and surprise. *Journal of Economic Methodology* 12: 317-329.
- Morgan MS (2012) *The World in the Model: How Economists Work and Think*. Cambridge: Cambridge University Press.
- Morgan MS and Morrison M (eds) (1999) *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press.

- Muniesa F and Callon M (2007) Economic experiments and the construction of markets. In: MacKenzie D, Muniesa F and Siu L (eds) *Do Economists Make Markets? On the Performativity of Economics*. Princeton, NJ: Princeton University Press, 163-189.
- Poitras G and Jovanovic F (eds) (2007) *Pioneers of Financial Economics*. Cheltenham, GL: Elgar.
- Rebonato R (2004) Interest-rate term-structure pricing models: A review. *Proceedings of the Royal Society of London A* 460: 667-728.
- Rheinberger HJ (1997) *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*. Stanford, CA: Stanford University Press.
- 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).
- Schönbucher PJ (2001) Factor models: Portfolio credit risks when defaults are correlated. *Journal of Risk Finance* 3: 45-56.
- Sismondo S (1999) Models, simulations, and their objects. *Science in Context* 12: 247-260.
- Sismondo S (2000) Island biogeography and the multiple domains of models. *Biology and Philosophy* 15: 239-258.
- Sklar A (1959) Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris* 8: 229-231.
- Smith CW (1999) *Success and Survival on Wall Street: Understanding the Mind of the Market*. Lanham, MD: Rowman & Littlefield.
- Tett G (2005) Market faith goes out the window as the 'model monkeys' lose track of reality. *Financial Times*. 20 May, 23.
- Vasicek O (1984) The philosophy of credit valuation. March 22. Privately circulated typescript. Personal papers, Oldrich Vasicek.
- Vasicek O (1987) Probability of loss on loan portfolio. February 12. Privately circulated addendum to Vasicek (1991).
- Vasicek O (1991) Limiting loan loss probability distribution. San Francisco: KMV Corporation, privately circulated. Available at: www.moodyskmv.com/research/whitepaper/Limiting_Loan_Loss_Probability_Distribution.pdf (accessed 5 May 2008).
- Vasicek O (2002) The distribution of loan portfolio value. *Risk* 15(12): 160-162.

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.