The Lemon-Squeezing Problem:

Analytical and Computational Limitations in CDO Evaluation

Iain Hardie* Donald MacKenzie

University of Edinburgh
Chrystal Macmillan Building
15A George Square
Edinburgh
EH8 9LD

* Corresponding author
E mail: Iain.Hardie@ed.ac.uk
Telephone: 0131-650-4249

Acknowledgements

We are extremely grateful to the referees and editor of this journal for helpful comments. The empirical work and analysis for this paper was supported by a grant from the UK Economic and Social Research Council (RES-062-23-1958), and its writing up was supported by the European Research Council (FP7/2007-2013/ERC grant agreement no. 291733).

Abstract

This article analyzes Collateralised Debt Obligations ('CDOs'), complex securities that were at the heart of the recent financial crisis. The difficulties of analyzing these securities are considered, and it is argued that the increasing complexity of CDOs that repackaged Mortgage-Backed Securities outpaced the returns available to investors, and therefore the resources available to pay for the analysis required to value the securities adequately within the timeframe available. CDOs therefore faced the problem of computational intractability. Such an outcome was, the article argues, inevitable in financial innovation that sought to create ever-higher returns from the fixed returns on a pool of assets. CDOs created what the article labels a Lemon-Squeezing Problem. Implications for regulatory responses to the crisis are briefly explored.

Keywords:

Financial markets, financial regulation, financial crisis, derivatives, market devices

The Lemon-Squeezing Problem:

Analytical and Computational Limitations in CDO Evaluation

‘If you can’t measure it, you can’t manage it’. Peter Drucker

Introduction

At the heart of the financial crisis, it is widely argued, lay Collateralized Debt Obligations (‘CDOs’), especially those whose underlying assets were Mortgage-Backed Securities (‘MBSs’). An understanding of these complicated financial instruments is central to understanding why difficulties in a relatively small financial market, sub-prime mortgages, became a systemic global financial crisis, and to determining appropriate regulatory responses. This article also examines CDOs in detail, but focuses on the problems of valuing these securities, to reach novel conclusions as to the underlying problem with the financial innovation that produced these ‘toxic assets’. It is shown that the nature of CDOs meant that it was impossible for the returns they offered to be sufficient to meet the costs of satisfactorily analyzing them. The successive ‘slicing and dicing’ of a finite cashflow from a pool of assets inevitably increases complexity and equally inevitably outruns the analytical computational capacity to complete timely evaluation that investors buying low return securities could justify paying. At the heart of the problem with CDOs was the fact that CDO evaluation involves not just a problem of asymmetric information (of ‘lemons’ in the terminology of the classic article by Akerlof 1970). It also involves what the article calls a ‘lemon-squeezing problem’: the inherent conflict between fixed returns (a finite quantity of ‘lemon juice’) and the growing complexity of CDOs, a complexity that outpaced the resources needed to understand it.

Market Devices and Market Complexity

Recognising the Cost of Market Devices

In making this argument, we contribute to two bodies of literature. The first is the literature on ‘market devices’, which is heavily influenced by the work of Michel Callon. Examples of market devices range from the mundane (supermarket shelving and trollies) to the esoteric (pricing systems in financial-derivatives markets) and from technological objects (such as stock tickers) to mathematical concepts (such as those deployed in financial economics). Market devices are ‘the material and discursive assemblages that intervene in the construction of markets’ (Muniesa, Millo and Callon 2007: 2); this collection edited by Muniesa and his colleagues provides a good range of case studies of them.

The work sparked by Callon and by the focus on the role of devices in economic life has been influential, not least in highlighting the importance of these often-overlooked influences on market outcomes: for detailed examples from finance, see MacKenzie (2006 and 2011); for examples from dealing with a physical commodity, see Çalışkan (2010). Nevertheless, the work has been the object of multiple critiques. It has been criticised, for example, for its ‘neglect of power and politics’ (Ertürk et al. 2013: 338; see e.g. Mirowski and Nik-Khah 2007). The critique that is most relevant here, however, is that the literature on market devices tends to ‘miss that which is precisely capitalist about capitalism: namely that the aim of any private enterprise is to generate profit, not construct a market’ (Ertürk et al. 2013: 339). Paradoxically, the burgeoning literature on the role of devices in economic life contains surprisingly little discussion of the economics of those devices, of how much they *cost*, relative to the revenue available for their acquisition and use. It is precisely that issue that is our focus here.

Complexity and Crisis in Finance

The second body of literature (also burgeoning, but rather more disparate than the first) to which we contribute concerns the role of complexity in finance, for example in the genesis of the credit crisis. Sometimes the notion of 'complexity' is used in the everyday sense; more rarely, the formal meaning of 'computational complexity' (touched on below) is deployed, as for example by Mirowski (2010). In both meanings of the term, it is clear that in the run-up to the credit crisis, the process of financial innovation – well characterized by Engelen and his colleagues as involving not a 'rationalist grand plan' but bricolage, 'the creative and resourceful use of materials at hand… to fashion new structures out of conjunctural events' (Engelen et al. 2011: 51) – led to a sharp increase in the complexity of financial instruments, and that this complexity (and the resultant opacity of those instruments and of the economic circuits in which they were implicated) was an important driver of the crisis.

Analytical Complexity and Computational Limitations

Issues of analytical complexity are clearly not confined to the social sciences; they are also, unsurprisingly, a central focus of computer science. One example is the 'Travelling Salesman problem', concerning the absolute and increasing computational difficulty of calculating the optimum journey between points. A recent solution involving 85,900 points justified a Princeton University Press monograph and required the equivalent of years of computing time. This is an 'NP-complete' (Nondeterministic Polynomial) problem: the difficulty of calculating a solution increases exponentially with additional points (Maymin 2011: 1). For the purposes here, a solution may be even more impractical if it must be found quickly with limited resources. With CDOs, these resources are the return on the securities in question, and they are intrinsically constrained by the initial pool of assets involved. The core of the self-undermining nature of the financial innovation involved in CDOs results from this constraint. Increasing complexity is required to produce ever-greater volumes of highly-rated securities from an underlying pool of assets whose returns cannot increase. (The assets underpinning a MBS or CDOs are debt instruments – mortgages, loans, etc. – and the returns from these are fixed by the underlying contracts.[1]) Fixed returns result in complexity outstripping calculative capacity. Financial innovators faced a 'Lemon-Squeezing Problem' which undermined financial stability.

The article therefore seeks to bring questions of computational complexity into discussions of financial markets and their regulation. At its most ambitious, the literature on computational complexity points out the incompatibility of computational and market efficiency (Maymin 2011), and looks to challenge neoclassical macroeconomics (Mirowski 2010). We join Hasanhodzic *et al*. (2009) and Arora *et al.* (2009) in seeing the issue as an extreme example of bounded rationality (Simon 1955). For Simon, our ability to act rationally is not limited by the availability of information, but rather by our limited ability to process that information because of limited computational capacity. In our analysis, the constraint on rationality is available resources relative to the computational capacity required to value securities.

Arora et al. (2009) see the issue in terms of informational asymmetries and Akerlof's lemons problem: 'designers of financial products can rely on computational intractability to disguise their information via suitable "cherry picking".' Such problems of information asymmetry certainly exist within the story of CDOs. Goldman Sachs' infamous Abacus transactions are a well-known example, and legal actions against Goldman seek to demonstrate the firm's information advantage (see, e.g., Lewis Baach 2011). We focus, however, on the 'intractability' or 'infeasibility' of calculation, within the available returns on the securities involved. Computational intractability, within available resources, is seen for a market actor with *all* relevant information (see Hasanhodzic et al. 2009, although their focus is not CDOs or derivatives), rather than with *almost all* relevant information (Arora et al. 2009). This is not a lemons problem, but a lemon-squeezing problem. Two striking features of the CDO market support a focus away from information asymmetries. First, banks that structured the securities and owned mortgage-originating companies made substantial losses on holding the AAA-rated tranches of CDO issues they themselves structured. Any superior information these banks possessed did not protect them. Second, the necessary information to value these securities was available, using information systems such as Intex. The problem was only partly that few used this system, remarkable as that is. The overriding problem was that using this system required uneconomic amounts of skilled inputs and impractical levels of computer capacity to complete analysis within the resources and time available.

Global Demand for Low Risk Securities and Blame for the Financial Crisis

The returns on CDOs, and the increasingly limited time available for considering a purchase, means, it is argued here, that they *must* be 'information-insensitive' securities, requiring limited analysis (see, e.g., Gorton 2010). The purchase of CDOs was part of the increased global demand for such debt securities in the years preceding the crisis. This fits with analysis that sees the crisis as more an issue of demand than supply. Global demand for low risk US$ securities in excess of available US Treasuries and Agency debt drove securitization (Bernanke et al. 2011; Caballero 2009). Such analyses draw us into the issue that dominates much discussion of the financial crisis, the question of who, if anyone, should be blamed.

This article is not interested in exonerating anyone, least of all the bankers and credit rating agencies involved in CDOs. Issues of complexity and the limits of knowledge as contributors to the crisis, and the implications for apportioning blame, are considered elsewhere (e.g., Bryan et al. 2012; Engelen et al. 2012). This article suggests fault lies with actors and structure, but (mainly) with actors who should have recognized the need for far greater analysis, and that the financial innovation involved in CDOs did not, and *could not*, allow returns that were sufficient for this analysis.

Research Methodology

Our analysis combines these two key issues: the economics of market devices and the complexity of financial instruments. 104 interviews were conducted with participants in the credit derivatives and asset-backed securities markets in the US and UK (including those who structured, invested in, quantitatively analysed and rated CDOs, based mainly in London and New York), on how participants evaluate financial instruments such as CDOs and MBSs. These were semi-structured interviews, which we recorded except on the rare occasions interviewees refused permission. Documentary sources are also analyzed, including the offering circulars (the information provided to prospective purchasers) for a range of MBSs and CDOs, and interview recordings, taken from the website of the United States' Financial Crisis Inquiry Commission (FCIC). The article investigates just what was involved in evaluating three classes of instrument:

1. A 'vanilla' MBS, the earliest and least complex securitization discussed here, for which the process of evaluation was reasonably tractable and economically feasible given the returns available.

2. A CDO of MBSs, in which tranches of MBSs form the 'pool' of assets underpinning the CDO. The huge increase in complexity is demonstrated relative to the first class of instrument, and the resultant sharp deterioration in the economics of evaluation. Evaluation of the kind that was feasible for vanilla MBSs became economically infeasible for CDOs of MBSs in the time available for market participants.
3. A CDO-squared, in which tranches of other CDOs form the asset pool. CDO-squareds involved yet another enormous increase in complexity. Increasingly, the constraints on evaluation were no longer simply economic; many CDO-squareds had features, which are discussed, that rendered evaluation effectively impossible, regardless of the time available.

Analysing Collateralised Debt Obligations

The Intex System

Our focus is on the single most important market device employed by participants, a software system called Intex. This was hugely important to the market in MBSs and CDOs, but has not yet attracted academic attention. Unlike the systems employed by ratings agencies (such as the Gaussian copula model: see MacKenzie 2011), Intex is simply a cash-flow model. Once the structural characteristics of the MBS or CDO are crystalized in the form of an Intex file (a task normally undertaken by those seeking to sell the tranches of the MBS or CDO), a prospective purchaser can choose a variety of scenarios, and use the Intex system to investigate the consequences for the tranches of the MBS or CDO. An MBS purchaser, for example, could choose a mortgage default rate and some other assumptions, and the Intex system will work out from the resultant cash flows whether or not a given tranche will default.

Although other similar systems were available, Intex seems to have been most widely used for tasks such as this. Even a simple vanilla MBS is a complex instrument (in the ordinary sense of the term ‘complex’), hard to get one’s head around: the offering circular is an incomplete account of the instrument, and one that cannot easily be processed for the purposes of economic analysis; the full legal documentation of a MBS or CDO runs to hundreds of daunting, unreadable pages. An Intex file, in contrast, captures the structure of a MBS or CDO in a way that is easily transferrable from the constructor’s computer system to the prospective purchaser’s. The purchaser can then try out the consequences of a variety of scenarios, decide whether or not to buy, or perhaps whether to request the constructor to change the instrument’s structure so that it was more attractive.

Other ways of evaluating MBSs and CDOs are also briefly touched on, but Intex is central to our argument. We ask what was involved in performing an Intex analysis of each of the above three instruments andhow long it might take. We also ask whether it could be paid for from the returns offered by the instruments. It is not claimed that quantitatively exact answers to those questions are provided – that would require data that we do not have, and may never have been collected by market participants – but the differences demonstrated between the economic and computational characteristics of the evaluation of the three classes of instrument are so large that the qualitative conclusions drawn are robust, and the implications significant.

After this introduction, the article next provides an overview of the structuring of CDOs, before considering the low returns available to those investing in and managing these securities. The need for timely analysis is then considered. The increasing complexity of the securitization process is then followed, starting with a 'vanilla' MBS. This is the standard structure whereby a pool of residential mortgages is packaged into a tradable financial asset, with various tranches offering different levels of creditworthiness and return. Such MBSs date from the 1970s, and are therefore a long-established part of financial markets. The article considers next the more recent innovation of CDOs that repackage certain of the tranches from MBSs (most commonly BBB rated) into further bonds with tranches of a similar range of creditworthiness and returns. These were then followed by 'CDO-squared', which performed the same repackaging on tranches of CDOs, again A or most often BBB rated. As this market became even more overheated, 'CDO-cubed' issues were also structured, repackaging tranches of CDO-squared, but we do not have access to an offering circular for such an issue. Throughout this analysis, the increased returns to investors from investment in the highest-yielding AAA-rated tranche of each structure are contrasted with the increasing complexity of those securities.

The presentation of a relatively straightforward linear process of financial innovation is of course highly simplified. Financial market actors, seeking to squeeze the lemon to create more salable securities from the capped returns from the underlying mortgages, were engaged in a process of far greater complication than we consider here. The focus here, however, is on the analytical tractability of these bonds, and so any additional complexity supports the argument.

Collateralized Debt Obligations: A Brief Introduction

A bank arranging a CDO sets up a Special Purpose Vehicle (SPV), a company whose sole purpose is to buy assets (collateral) that comprise other debt securities (see Bomfin 2005).

The assets considered here are Residential MBSs, which themselves buy and securitize residential mortgages. In return for a higher return, investors in the least senior tranches, accept the first losses from non-payment on the underlying assets. The most senior (including ‘super-senior’) tranches offer very low returns, a AAA rating and (supposedly) minimal risk of loss. Figure 1 sets out the structuring of a CDO which buys MBSs.

Figure 1: Packaging tranches of MBS into CDOs

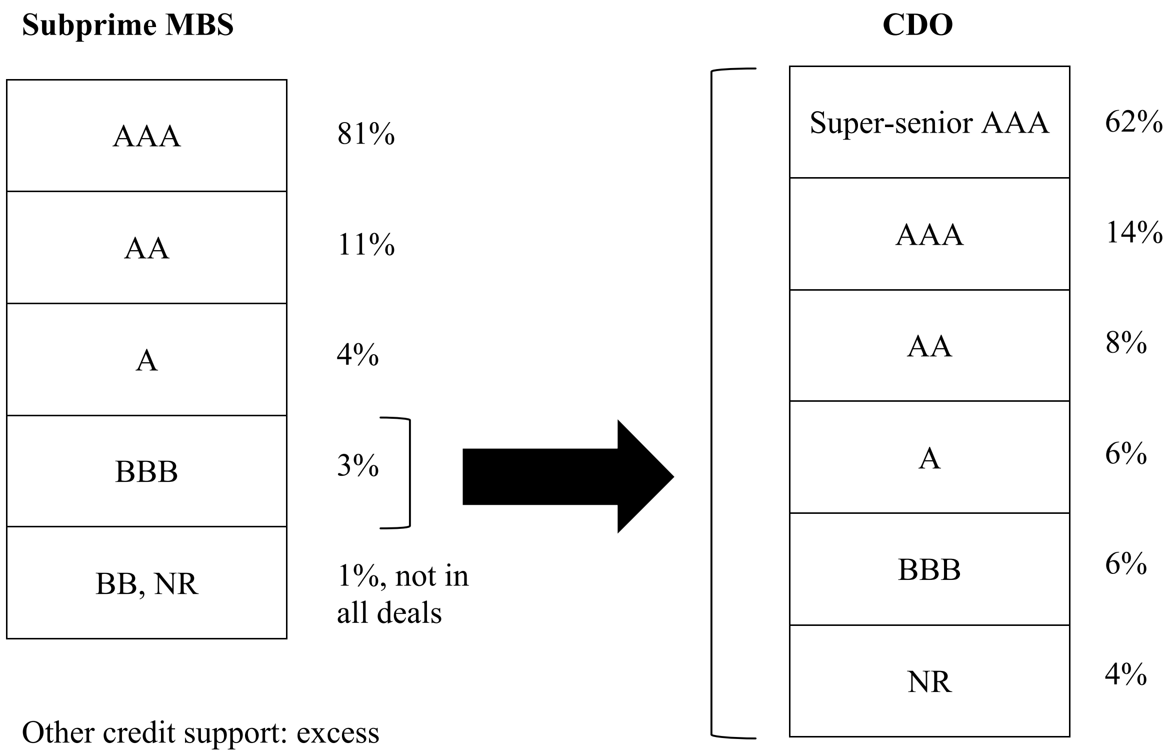

Other credit support: excess spread, over-collateralization

Other credit support: excess spread

Tranche sizes not to scale.
‘NR’ means ‘not rated’

Source: modified from Lucas (2007)

In the CDO market the less creditworthy tranches are further securitized into predominantly highly-rated, but a smaller volume of lowly-rated, tranches. The process was then repeated with these lowly-rated tranches. This is the lemon-squeezing whose inherent problems we explore. The highly-rated securities sold to investors offer higher returns at each stage of this chain of transactions, but the returns remain low relative to the increasing complexity, and therefore analytical difficulties, of the securities.

Low returns on highly rated CDO tranches

Greed and Low Returns

Popular discussion of, and political discourse regarding, the financial crisis have made much of financial market actors' greed (e.g., FCIC 2011; Madrick 2011; Mason 2009). It is therefore noteworthy how low the profit margins actually were in cash terms for those financial institutions bearing the ultimate risk: the Dublin unit responsible for the collapse of Landesbank Sachsen, which made investments in CDOs, made a total profit in 2005 of €44.2 million on investment of €8.3 billion (Kirchfeld and Simmons 2008), or 0.53 percent. Deutsche Bank, in its 2007 annual report, noted earnings of just €6 million on commitments supporting similar investments of €6.3 billion, Bank of New York Mellon US$3 million on commitments of US$3.2 billion, both returns of 0.10 percent (Acharya et al. 2011: 29; see also Arteta et al. 2008). Market participants therefore faced not only the computational difficulties of analysing these securities: it was also that the returns they offered were too low to allow for a full analysis.

Low Fees for Collateral Managers

Low returns for what should be complex analysis were not only a problem for final buyers; they also occur elsewhere in the chain of transactions involved in a CDO. Part of the complexity is that these financial structures generally use a collateral manager. These managers are not buying securities on their own behalf, but choose the underlying assets on behalf of the final investors (Engelen et al. 2011; Lewis 2010). Investors quite reasonably feel collateral managers should protect their interests, but management fees were not high for what should have been a complex role. In one manager's case, annual fees were said to be 0.09 – 0.17 percent of the transaction volume, with around 0.25 percent of additional performance fee (Shenn 2010). Another estimate is that fees were 0.45 - 0.75 percent (Salas and Hassler 2007). In equity fund management, these are fees closer to those charged by passive managers (investors that track the index rather than choosing individual stocks), not active managers, and they appear to have declined over time (Chau 2010). The total fees paid to CDO managers are large, at least US$1.5 billion for 2003-07 (FCIC 2011: 131). However, that represents only 0.11 – 0.23 percent of the total volume of CDOs (depending on the estimates of total volume used). CDO management was popular not because fees were high, but because managers had to do little analysis to earn them: 'The CDO manager, in practice, didn't do much of anything…"two guys and a Bloomberg terminal in New Jersey" was Wall Street shorthand for the typical CDO manager' (Lewis 2010: 141). Returns were maximized not by increased analysis of particular securities, but by not doing 'much of anything' on the largest volume of securities possible. Across the CDO industry, similar incentives resulted in huge volumes of insufficiently-analysed securities.

The need for timely analysis

Daily Valuation of Trading Books

Low returns had to pay for analysis that was not only highly complex, but also needed to be completed very quickly. Banks increased their return on equity, the return made for shareholders, before the crisis in part through higher leverage assisted by the favourable capital treatment of bonds held on trading books. The Federal Reserve's Norah Barger (2010) suggests that much of the buying of the most creditworthy and lowest yielding CDO tranches only occurred because of lower capital requirements on trading books, where banks held securities they were expected to buy or sell frequently.[2] Trading book assets must be revalued daily. They are also frequently financed, on a secured basis, overnight, in 'repo' markets widely seen as central to the financial crisis (Gorton and Metrick 2009), which also require daily (if not more frequent) valuation.[3] This introduces time pressure into calculations (see also Spears 2014).

Primary Market Purchases by CDO Managers

CDO managers also needed timely analysis. The competition amongst CDO managers to buy assets, combined with the incentives for the arranging banks to sell quickly, meant that the time available for analysis for the purchase of the 'mezzanine' (BBB rated) tranches on which the CDO market depended was extremely short. It also declined as the mania gathered pace, falling from a week to a day or even less. One CDO manager told us his firm was able to do the analysis using Intex 'in an hour or two'. Almost regardless of the amount spent on computer capacity, this was not sufficient time for a comprehensive analysis. The bespoke nature and small size of mezzanine tranches (only 3 percent of the original MBS in the typical structure shown in Figure 1 above) meant that they were illiquid in the secondary market, so CDO managers had little choice but to buy new issues. In an example of how the dynamics of a bubble can reduce the time available for analysis, all the tranches of a new MBS were often sold within less than four hours.

The Time Advantage of Short Sellers

The most common explanation for the success of short sellers – investors, such as hedge funds, that made trades in the expectation of prices falling – is that a lemons problem exists.

However, the success of short sellers fits well within our lemon-squeezing explanation. First, in contrast to the 'long only' buyer discussed elsewhere in this article, a hedge fund shorting an MBS or CDO is expecting substantial price falls and therefore substantial profit. The financial incentive to pay for analysis is therefore far higher. Second, the short seller does not face the time pressure of the MBS or CDO purchaser. The market opportunity is very unlikely to disappear within hours. Whatever the rights and wrongs of Goldman's Abacus transactions, it is clear that Paulson & Co. had the time they needed to select the assets they wished to short. Even so, it is noteworthy that the short sellers highlighted (e.g., by Lewis 2010) generally shorted MBSs, not CDOs, despite the fact that they would have made even greater profits from shorting CDOs. This may suggest that short sellers also faced difficulties in analyzing CDOs.

The increasing complexity of CDOs

In this section, we analyze three stages in the increased complexity of MBS and CDO securitization. We consider first a 'vanilla' MBS, the initial stage of mortgage securitization, and then a CDO, a structure which securitizes MBS. We end with consideration of a further stage in complexity, a 'CDO-squared'.

Stage 1: 'Vanilla' MBS

The analysis starts with a relatively straightforward, 'plain vanilla' MBS, discussed at length by the FCIC.[4] The approximately US$770 million deal was arranged by Citigroup in 2006; Citigroup Mortgage Loan Trust 2006-NC2 ('CMLT') bought 4507 subprime residential mortgages from across the United States, but with over 40 percent in California and Florida. 13 tranches, ranging in Standard & Poor's rating from AAA to BB+, were offered to investors globally. Despite its 318 page offering documents and subsequent problems (56.1

percent of loans were either delinquent or liquidated at the time of the FCIC analysis), analyzing this deal is not computationally complex. The returns on investing in the most creditworthy tranches are low. The highest yielding AAA tranche[5] gives a return of only 0.24 percent over LIBOR.

One of the interviewees demonstrated a single Intex 'run' for a specific tranche of such a MBS. Inputting the necessary data (he used simply fixed values of the inputs, not – as he normally would – inputs that varied through time) took around 20 seconds, and the Intex system then took around 30 seconds to calculate the future cash flows to an investor in the tranche. He was using the version of Intex available to internet subscribers, but he told the authors that the system's response time would be faster – a matter of seconds – using the version running on his bank's own servers. On the other hand, he would normally have to input full monthly curves of interest rates and default rates for the lifetime of the security. Nevertheless, he said that a single run could be completed in 'minutes'.

A single run would not count as adequate analysis of the MBS. This would require multiple runs, interviewees told us. A full Monte Carlo simulation[6] to get a 'rough price' would need 100,000 runs and 'a million simulations or more to get some decent risk measures'. The inputs for this simulation could of course be entirely automated, but with each run taking a few seconds of computation time, the simulation was a time consuming task. However, it could be speeded up, for example by distributing the task over multiple computers, and certainly was not entirely impossible. It involves assumptions about future events (for example, prepayments – early redemptions by the borrowers – of fixed rate mortgages are closely linked to US government bond yields), but so must all investment. Consideration of the 4507 underlying residential mortgages is essential to ensure the aggregated assumptions about this mortgage pool entered into Intex produce valid outputs. The need for accurate assumptions will add time to any analysis. Nor is Intex cheap; the bank at which our interviewee demonstrated Intex was paying US$1.5 million p.a. to use it. This cost is presumably a major factor in the low use of Intex amongst (particularly European) investors.

Nevertheless, the analysis of CMLT was, the article contends, feasible within the context of the resources available (the return on the securities), and could continue on an ongoing basis with the frequency required to hold these bonds on a bank trading book. This feasibility can only be a contention (and clearly initial purchasers are likely to have regretted their analysis, given subsequent events), but is based on a number of points. First, investment in the AAA rated tranches of MBS took place in large size by very large financial institutions. In the case of this security, the three publicly-sold tranches rated AAA, totaling US$580 million, sold to approximately 18 investors[7] (some investors appear more than once under slightly different names). While the list includes Memorial Hospital at Gulfport and Ursuline Society & Academy, an educational establishment, small investors like Italian or Norwegian municipalities (e.g., *New York Times*, 2 December 2007) represent a small part of the market. The MBS market was dominated by the largest investors who, for the number of mortgages in an MBS, could meet the cost of Intex. Furthermore, absolutely precise valuation (even if possible in all but a very small number of cases in finance) is not the issue, given the enormous losses subsequently made. Large financial institutions could still aggregate information across the various deals they held. To create sufficiently accurate outputs from Intex, it was not necessary to understand the detailed development of every mortgage, but to have a clearer picture of what was occurring with sub-prime mortgage lending in, say, California or Florida. This clearer picture could come with some analysis of the individual mortgages when the mortgages are of the number in an MBS. However, such mortgage-level analysis is still required with a CDO, but, even assisted by the use of Intex, is too time-consuming.

The simplistic approach here to comparing MBS with the subsequent classes is to consider the return per annum[8] from investing $100 million of borrowed money in the highest yielding AAA tranches of CMLT. Many of the investors in MBSs were banks[9] investing money borrowed from depositors or wholesale markets.[10] The London Interbank Offered Rate (LIBOR) represents a reasonable proxy for the cost of that borrowing. An investor therefore borrows funds at a cost of LIBOR, and uses those funds to invest in a security that gives a return of LIBOR plus 0.24 percent. The borrowing cost of LIBOR is matched by the LIBOR earned on the investment, leaving a fixed return of 0.24 percent, regardless of changes in LIBOR. This fixed return is US$240,000 on a US$100 million investment, or US$53.25 for each mortgage in the pool of underlying collateral. With a portfolio of US$1 billion (small by the standards of the large US and European banks), Intex or an alternative could easily be afforded. The overall argument in this paper does not rest on the reader's acceptance of this claim, but two points are central to the argument. First, the cashflow for the investor is fixed. The pool of mortgages created produces an income for the SPV. The investor has an amount of money to invest, in return for which it receives a portion of the cashflow from the mortgage pool. The bank arranging the deal can make changes to the structure, but the returns for investors do not materially change.[11] Adding more mortgages to the pool would increase the size of the overall transaction, but not the return to the investor, who is investing a fixed amount. The second point is that increasing the number of mortgages in the underlying pool increases the difficulties of analysis in line with the increased return from the additional mortgages. The complexity rise is, however, linear rather than, as in CDOs, exponential.

Stage 2: Collateralized Debt Obligation

The variety and complexity of CDOs of MBSs support the overall conclusions, but, in making the argument, CDOs can be sketched quite simply. CDOs have the same basic structure as the MBS discussed above. They buy assets to put into a collateral pool and issue tranches of securities secured on that collateral pool. The difference is that the assets in the collateral pool are the tranches of MBSs. Frequently, these tranches were the BBB-rated 'mezzanine' tranches towards the bottom of the seniority ranking of MBSs.

These CDOs, just as the MBS described above, take a pool of assets with a predetermined maximum cashflow (see Figure 1 above), and, thanks to the supposed diversification of the collateral pool and tranching, issue highly rated securities. They also inevitably increase the complexity of analysis of the securities. The buying of higher-yielding BBB MBS assets, in conjunction with the CDOs' tranching, allows some increased return to holders of the AAA tranches. In the case of Kleros Real Estate Finance I ('Kleros'),[12] the CDO analyzed here, the return for the highest yielding AAA tranche was 0.40 percent, approaching double the CMLT return for the same tranche. For a US$100 million investment, the annual return rises substantially to US$400,000. However, the increased complexity of analyzing Kleros is far bigger. The data necessary to give a precise number of the underlying assets in Kleros are not available, but a plausible estimate can be made, accurate in degree if not actual number.

Limitations on the Kleros collateral pool's investment in a single issuer's securities mean the pool contains a minimum of 88 MBS. The pool is also 5 percent Commercial MBSs, with the remainder Residential MBSs. A conservative assumption, therefore, based on the 4507 residential mortgages underlying the CMLT, is 4000 underlying mortgages in each MBS purchased by Kleros, or 352,000 underlying mortgages. The number of underlying mortgages has increased by over 7000 percent, the return to an investor in the highest yielding AAA-rated tranche has less than doubled. The amount of return available for each underlying mortgage (the proxy here for the resources for analysis) falls from $53.25 to $1.14.

Again, Intex was widely used to analyze CDOs based on MBS. That task was vastly more time consuming than analyzing a single MBS. First, appropriate inputs for the default rate, etc., must be chosen for each MBS: given the differences among them, choosing the same default rate for all of them was clearly inadequate, but judging the appropriate rates was a skilled, time-consuming human task. Then Intex had to be configured first to run a cash flow analysis for each underlying MBS, and second to feed the inputs from these analyses into the Intex model of the CDO. While, as noted above, a single run of Intex for a specific MBS tranche would take 'minutes', an interviewee told the authors that a single run of a CDO of MBSs, 'doing loan-level modeling for the underlying [MBS] bonds then applying it to [the] CDO… would have taken hours. So you might set [it] running on the evening then come back the next morning to look at the results [for a] single scenario'.

Given that, it is not surprising that interviews suggest that the analysis of CDOs using Intex typically took the form of only a handful of runs: a base case, incorporating the analysts' most likely scenario, and a small number of 'stressed' scenarios to get some sense of the consequences of adverse outcomes. No-one reported attempting anything approaching a full Monte Carlo simulation: with a single run taking several hours, a million runs was clearly infeasible, even if one were able to distribute the computational task over multiple machines. (In 2006-7, an investment-bank computer room might have upwards of a thousand machines in it, but 'parallelizing' a program is never fully efficient: running it on a thousand machines is not a thousand times as fast as running it on one, e.g. because the machines must communicate with each other.) Even if CDO investors had been allowed a week to make their decisions, and even deploying all available machines, the computation simply could not plausibly be completed in time.

This analysis also dramatically understates the challenges in calculating the value of CDOs. Many CDOs, including Kleros, employ a collateral manager (see above), which MBSs do not. The collateral manager's job involves – within constraints set out in the issue documentation, and, it has been claimed, subject to pressure from the bank arranging the CDO (see, e.g., Shenn 2010) – selling and buying assets in the collateral pool. Investors do not therefore face a constant collateral pool to analyze. A similar problem is that many deals have a 'ramp-up period' (though not Kleros). This involves a period – up to six months (Shivdasani and Wang 2009: 9) – after the investors have bought the CDO tranches when further assets are purchased (again within preset constraints) to increase the size of the collateral pool. Cashflows into the CDO from amortization, maturities, prepayment or sales are also reinvested by the manager. Although the investors know the broad characteristics of these assets, they do not know the specific assets. This is particularly important when, as in subprime mortgages, the quality of the underlying assets deteriorated over time. CDO investors in 2005, for example, when subprime mortgages were of generally higher quality, found that their CDO bought mortgages from 2006 and 2007, when quality had significantly deteriorated (Goodman et al. 2008: 286).[13]

Stage 3: CDO-Squared

The next stage in the development of the CDO market was the 'CDO squared', in which the immediate collateral assets are not MBS but other CDOs. To quantify the difficulties of analysis, we consider such a transaction, Timberwolf I.[14] 93 per cent of the collateral assets are Credit Default Swaps (CDS), so this is largely a 'synthetic CDO', but we ignore that here.[15] The offering circular for Timberwolf, unlike Kleros, lists the 56 CDO assets in the collateral pool. This is fewer than the ABS issues in Kleros. However, multiplying the 56 by 350,000 (the approximate number of underlying mortgages in Kleros) results in 19,600,000 mortgages. The increase in interest on the highest yielding AAA tranche is substantial, to 1.30 percent over LIBOR, but the return per underlying mortgage has fallen to US$0.07: seven cents per year. Although we have not found a 'CDO-cubed' – a CDO where the collateral assets are CDO-squareds – such issues were arranged. If the CDO-cubed had as its collateral 50 CDO-squared securities, the number of underlying mortgages rises to 98,000,000.

Sharp-eyed readers will notice that 98 million is more than the total number of mortgages outstanding in the United States. Even 19.6 million is approaching half the total. The lemon-squeezers faced another constraint: there were not enough mortgages outstanding to achieve the diversification on which the ratings should have depended. Two solutions were employed, but both increased analytical complexity. Different asset types were included in the collateral pool, as with the commercial MBSs in Kleros. This represents at least partial diversification, but increases the range of expertise needed for the assumptions underpinning analysis. The other solution is highly questionable in terms of diversification, and involved very significantly increased analytical complexity: 'looping', whereby CDOs invest in each other's asset pools. It became almost impossible even to understand the diversification of the CDOs. CDOs that included tranches of other CDOs in their asset pools rendered even a single Intex run hugely time consuming. One interviewee told us that in consequence he simply avoided such CDOs: 'I never did a … CDO with other CDOs in an underlying pool – never would have done – because I believe them…computationally intractable'. Most daunting of all were the cases, increasingly common in 2006-7, in which (often resulting from a 'scratch my back and I'll scratch yours' agreement between two CDO managers) CDO A would include a tranche of CDO B in its asset pool, and *vice versa*. The resultant computational 'loop' made it 'completely intractable' said this interviewee: even a single Intex run was now effectively impossible.

The Lemon-Squeezing Problem: Increasing Complexity Outpaces Increased Returns

Table 1 summarises the analytical issues and returns across the securities considered:

Table 1: Returns and Number of Underlying Mortgages, Various Securitization Structures

	MBS[16]	CDO[17]	CDO Squared[18]
Spread on Highest Yielding AAA Tranche	0.24%	0.40%	1.30%
Return on $100m Investment	US$240,000	US$400,000	US$1,300,000
Number of Underlying Mortgages	4507	352,000[19]	19,600,000[20]
Collateral Manager	No	Yes	Yes
Return per Underlying Asset	$53.25	$1.14	$0.07

Further anecdotal evidence strongly supports the conclusion that computational complexity constrained analytical capacity. Barger (2010) recounts how Citigroup in late 2007 told regulators that they had not been including AAA CDOs in calculations to determine capital against their trading book. They were told to calculate the requirement as soon as possible. Citigroup had to analyze individual mortgages (the process discussed above), and took 'several months'.[21] In addition, investment banks' valuation of the same financial product varied widely and valuation of different tranches within banks could be mutually inconsistent (Arora et al. 2009). In part, this results from 'clusters of [very different] evaluation practices' (MacKenzie 2011), but it also points to analytical difficulties. The long time that layered Intex models (i.e., analysis of CDOs that had purchased tranches of MBSs) take to complete their calculations discouraged (with hindsight essential) evaluation. Even with enormous computing capacity (see discussion of Goldman Sachs, which still employed a shortcut, below), it is not clear that the analysis could be fast enough for daily revaluation of all positions. The complaint by Basis Yield Alpha Fund against Goldman Sachs discusses a 'CDO valuation project', which used 'three different valuation methods to price all of its remaining CDO warehouse assets and unsold securities' (Lewis Baach 2011: 30);[22] i.e., *ex post* valuation of assets. The physical constraints involved were not just physical space and cost, but also of the air-conditioning capacity to deal with the heat from hundreds of computers.

Dealing with computational problems

Top-Down Analysis

Given the difficulties, outlined above, of using Intex to analyze a CDO from the 'bottom up' – i.e. by doing multiple Intex runs with different assumptions about the behavior of the underlying mortgage pools – it is unsurprising that market participants sought analytical 'shortcuts'. By far the most common was 'the bond method', or top-down analysis, involving no analysis of the underlying mortgage pools. Instead, the behavior of an MBS (or CDO in a

CDO-squared structure) tranche was simply inferred from its credit rating, and the CDO was analyzed as if it were simply a CDO made up of corporate bonds. That is how the rating agencies and many market participants analyzed CDOs (Mackenzie 2011). However, although default probabilities inferred from ratings had some credibility, the bond method also required inputting figures for the correlations among MBS tranches: 'There was never a good source of correlation numbers', reported an interviewee. As he told the authors, there was a widespread understanding amongst market participants that the correlation figures (of the order of 0.3) used by the rating agencies were far too low, but what was much less clear was how high a figure to use: 0.5, 0.7, maybe even 0.8. Even now, the correct correlation remains unclear.

Goldman Sachs developed perhaps the most intriguing shortcut. They employed, so an interviewee said, a hybrid of the 'bond method' (top-down) and bottom-up analysis. It was computationally very demanding – requiring a 'computer farm' in New Jersey, the authors were told – but tractable. Goldman, however, was as far as we can tell the only market participant to do things this way. Others either satisfied themselves with a relatively small number of Intex runs, or fell back on the bond method. Alternatively, many investors simply outsourced the credit analysis. A number of legal actions argue that CDO arrangers and/or managers had responsibility for the securities they sold, effectively challenging *caveat emptor* (e,g, Lewis Baarth 2011), or 'buyer beware'. There was clearly an agency problem in this market (Pozsar et al. 2010: 4), but it was an agency problem exacerbated by the low returns available to many agents relative to the analysis needed.

Problems of Outsourcing Analysis to Rating Agencies

The main outsourcing has been to the rating agencies, particularly for AAA-rated tranches, and criticism of their role is widespread. However, the key point to note here is that rating agencies face the same constraints as investors. The rating agencies cannot be exonerated for their role in the CDO market (and the reported failure of Moody's and possibly of other agencies even to subscribe to Intex is particularly noteworthy), but the extreme analytical challenge and time pressure of such volumes of initial ratings and ongoing monitoring was far beyond the challenge facing any individual bank or CDO manager. At the 2006 peak, Moody's was rating on average more than two CDOs every business day (FCIC 2011: 149).

None of the rating agencies did bottom-up analyses of CDOs: they all used variants of the bond method (fatally, with modest estimates of correlation). The closest to a bottom-up rating-agency analysis we found was a relatively small-scale experiment broadly similar to the hybrid method employed by Goldman Sachs (although as far as we can tell it was done entirely independently, with no interaction between the two teams). Correlations of around 0.8 were found, far above the 0.3 used in the bond method of CDO rating. It remained, however, just an experiment, with no influence on practice at the agency; the group conducting it did not have organizational responsibility for CDOs. It was 'a case of intellectual curiosity', one of the experimenters told us; CDO rating 'wasn't "under our watch" at the time'. The agencies did do significant analysis: Moody's, for example, included a matrix with 1000 scenarios (Adelson 2010), and Standard and Poor's Monte Carlo based modeling tool does 500,000 iterations (Alblescu 2010). However, these 500,000 iterations, which could be performed on a standard computer in less than a minute, did not 'drill down' to the underlying mortgages. Investors could not avoid the lemon-squeezing problem by outsourcing analysis to the rating agencies, as the agencies were not – and could not be – paid amounts sufficient for the analysis necessary.

Successful Outsourcing Before CDOs

When investors in, or those taking ultimate risk on, complex assets such as CDOs are not paid a sufficient return to allow for adequate analysis, reliance on others is always possible. In the market for MBSs before the emergence of CDOs, those others were investors in the mezzanine tranches of the structure (see Adelson and Jacob 2008). These specialist investors received relatively high returns for potentially significant risk. They therefore had the resources and incentive to complete meaningful credit analysis on the underlying mortgage portfolios. The unique position of these investors in the market also meant that they were given the necessary time to complete their analysis, which included examining the electronic records of the underlying mortgages. Essentially, investors in the more senior, low return tranches were dependent on the quality of this analysis. CDOs undermined this role, because CDOs purchased the mezzanine tranches at tighter spreads and with less time for analysis than traditional mezzanine investors were prepared to accept.

Conclusion

Space precludes a detailed discussion of how successful specific regulatory responses to the crisis have been in addressing the Lemon-Squeezing Problem. It is clear that increased capital requirements for trading books, especially for the securities we discuss here, leverage limits and greater scepticism about bank risk models are all potentially steps in the right direction. Complexity, in its more general sense, is both a recognised contributor to the crisis and an influence on regulatory responses, but these responses have addressed the issues discussed here largely accidently: the specific issue of the limits of computational capacity within finite resources has not been recognised. The result is that some responses – e.g., increasing competition in the credit rating industry – could arguably even make matters worse. The securitization industry is based on the use of increased complexity to squeeze ever greater returns from a finite income stream. That process must (and in only three steps did) reach the limits of computational capacity. Without regulatory constraint, this type of financial innovation will outpace increases in (affordable) computer power.

An important further question is obviously whether this inevitability is specific to CDOs, securitization generally, or is inherent in the bricolage (Engelen et al. 2011: 51) that characterises financial market innovation more broadly. It is necessary to remain cautious regarding that conclusion, but it is nevertheless important to consider, first, the inherent problem with securitization, and, second, the extent to which attempts by financial market actors to deal with this inherent problem can be seen as applying in other forms of financial innovation. The inherent problem with securitization is the inevitable conflict between fixed returns – the interest on the assets in any underlying portfolio – and complexity – the computational difficulties that arise in the tranching of securities that is the central innovation of the securitization process.

There are two basic ways to deal with these limits on computational capacity. The first is to 'outsource' full analysis to others, be they the arrangers of the CDOs or the credit rating agencies. This can only be effective if the arrangers or agencies are being paid fees sufficient for them to complete the necessary analysis. In a market constrained by the lemon-squeezing problem, this cannot happen. The second is to use assumptions, as inputs to computer models, to simplify the analytical process, rather than undertaking sufficient analysis for the assumptions to be valid, and/or considering numbers of scenarios to cover a sufficient range of assumptions. One question regarding the more general applicability of the problems in the innovation of CDOs is therefore the extent to which these two approaches might be taken in other areas of financial innovation. There are certainly reasons to suggest that they might. Heavy reliance on the rating agencies, not least thanks to regulation, pervades financial markets, and much of the market for Exchange Traded Funds involves investors relying on arrangers of highly complex structures. Assumptions, particularly regarding the applicability of past to future performance, are also central to financial markets, for example in the options market. Such assumptions were shown to be problematic with equity options in the 1987 stock market crash, but this did not prevent the assumption of no nationwide fall in US house prices causing financial disaster 20 years later.

The lemon-squeezing problem we highlight involves, at its heart, the increased complexity of trying to get ever more from a finite resource. Broad parallels can certainly be seen elsewhere in financial innovation, for example in developments in bank balance sheets in the years preceding the financial crisis. Banks increased their Return on Equity (the lemon juice) as their Return on Assets (the lemon) remained constant or even declined (Haldane 2009). This was achieved by financial innovation (aided by regulatory forbearance) which both watered down what counted as bank capital and restructured assets (for example, turning mortgages into MBSs or BBB CDO tranches into AAA) in ways that reduced risk-weighted assets and allowed higher leverage. The result was greater complexity: 'Banks appeared to have discovered a money machine, albeit one whose workings were sometimes impossible to understand' (Haldane 2009: 2). Equity investors temporarily received higher returns, but at the cost of banks they could not analyze. We argue in this article that financial innovation, in the specific case of CDOs, contains the seeds of its own destruction, because of the lemon-squeezing problem. Increasing complexity necessarily outpaces the resources to pay for the necessary analysis of this complexity. The extent to which this is inherent in financial innovation more widely is a question of considerable importance. The issues raised by computational complexity are therefore worthy of further study.

---

[1] Payments to MBS or CDO investors are fixed relative to LIBOR, and those investors in turn borrow at a cost also fixed relative to LIBOR. Therefore, interest rate rises increase the cashflow through the MBS or CDO, but do not generate extra income for investors.

[2] A further indication of their low returns.

[3] A transaction undertaken for only one day also obviously reduces the return very substantially. Lending US$100 million at, say, 3 per cent per annum overnight earns interest of slightly over $8000.

[4] See http://fcic.law.stanford.edu/.

[5] The highest yielding, least senior, AAA tranche is chosen in each case as the most conservative approach, given our hypothesis. Their overall conclusions, however, apply even more strongly to the more senior, even lower yielding tranches. In CMLT, two much larger and senior tranches offered LIBOR plus 0.04 percent and 0.16 percent. A further more senior tranche, yielding LIBOR plus 0.14 percent, was sold privately to Fannie Mae.

[6] Monte Carlo simulation involves defining the range of possible values of each parameter (in this case, mortgage default rate, recovery rate, recovery time, prepayment rate and interest rates); creating a 'scenario' by using random numbers to choose a set of values within those ranges; calculating the result (in this case, cash flow) for that scenario; repeating the exercise many times (i.e., generating multiple scenarios); and aggregating the results across scenarios.

[7] See http://fcic-static.law.stanford.edu/cdn_media/fcic-docs/2010-10-14_CMLTI_2006-NC2_20R_Senior Investors.pdf, accessed 9 February 2012.

[8] An arbitrary choice used for illustrative purposes only.

[9] See http://fcic.law.stanford.edu/.

[10] Hedge funds are similarly likely to be investing borrowed money.

[11] Faster repayment, for example, will increase cashflows, but not increase returns.

[12] See http://fcic-static.law.stanford.edu/.

[13] Particular features of CDOs' 'six-pack' structure also increase the sensitivity of any valuation to the underlying assumptions used, further increasing complexity.

[14] See http://fcic-static.law.stanford.edu/.

[15] A synthetic CDO sells CDS (i.e., is paid to take the risk of default on the reference securities underlying the CDS) rather than buying MBS or CDOs. Issuance volumes are therefore constrained only by the willingness of financial institutions to 'short' the reference MBS or CDOs, either for hedging or speculation. Returns and risks on synthetic CDOs, however, ultimately remain tied to the returns on the reference MBS or CDOs, which are securities of the kind discussed here, so the lemon-squeezing problem remains (for discussion of synthetic CDOs, see FCIC 2010, 142-6).

[16] Citigroup Mortgage Loan Trust 2006-NC2

[17] Kleros Real Estate Finance I

[18] Timberwolf I

[19] Authors'calculation based on maximum number if underlying securities under the single issuer concentration (prospectus, p.124), 88, multiplied by 4000, an estimation of the number of underlying mortgages in the chosen MBS. The lower number than 4507 is deliberately conservative, but allows for the fact that the collateral includes 'a substantial number' (but a minority) of Commercial MBS, which will have fewer underlying mortgages.

[20] 56 (number of underlying CDOs given in prospectus) multiplied by 350,000.

[21] Barger says that Bank of America was in a similar position.

[22] The case would presumably be strengthened Goldman were shown to know the true value of the assets throughout.

References

Acharya, V. V., P. Schnabl and G. Suarez. 2011. 'Securitization Without Risk Transfer.' Available at http://ssrn.com/abstract=1364525.

Adelson, M. 2010. Interview with the Federal Crisis Inquiry Commission, 22 October. Available at http://fcic.law.stanford.edu/resource/interviews.

Adelson, M. and D. Jacob. 2008. 'Risk Management Lessons from the Sub-prime Problem.' Adelson & Jacob Consulting LLC, 4 March. Available at http://www.securitization.net/pdf/Publications/Adelson_RiskMngmnt_Mar08.pdf.

Akerlof, G. A. 1970.'The Market for "Lemons": Quality Uncertainty and the Market Mechanism.' *Quarterly Journal of Economics* 84: 3: 488-500.

Alblescu, H. 2010. Interview with the Federal Crisis Inquiry Commission, 14 May. Available at http://fcic.law.stanford.edu/resource/interviews.

Arora, S., B. Barak, M. Brunnermeier and R. Ge. 2009. 'Computational Complexity and Information Asymmetry in Financial Products.' Working Paper, 19 October.

Arteta, C., M. Carey, R. Correa and J. Kotter. 2008. 'Revenge of the Steamroller.' 24 October. Available at http://apps.olin.wustl.edu/FIRS/PDF/2009/787.pdf.

Barger, N. 2010. Interview with the Federal Crisis Inquiry Commission, 31 March. Available at http://fcic.law.stanford.edu/resource/interviews.

Bernanke, B. S., C. C. Bertaut, L. Pounder DeMarco and S. Kamin. 2011. 'International Capital Flows and the Returns to Safe Assets in the United States, 2003-2007.' International Finance Discussion Papers 1014, Board of Governors of the Federal Reserve System.

Bomfim, A. N. 2005. *Understanding Credit Derivatives and Related Instruments.* San Diego and London: Elsevier.

Bryan D., R. Martin, J. Montgomerie and K. Williams. 2012. 'An important failure: knowledge limits and the financial crisis.' *Economy and Society* 41: 3: 299-315.

Caballero, R. 2009. 'The 'Other' Imbalance and the Financial Crisis.' MIT Department of Economics Working Paper No.09-32. Available at papers.ssrn.com/sol3/papers.cfm?abstract_id=1529764.

Çalışkan, K. 2010. *Market Threads*. Princeton and Oxford. Princeton University Press.

Chau, W. 2010. Interview with the Federal Crisis Inquiry Commission, 11 November. Available at http://fcic.law.stanford.edu/resource/interviews.

Engelen, E., I. Ertürk, J. Froud, S. Johal, A. Leaver, M. Moran, A. Nilsson and K. Williams. 2011. *After the Great Complacence.* Oxford: Oxford University Press.

Engelen, E., I. Ertürk, J. Froud, S. Johal, A. Leaver, M. Moran and K. Williams. 2012. 'Misrule of experts? The financial crisis as elite debacle.' *Economy and Society* 41: 3: 360-82.

Ertürk, I., J. Froud, S. Johal, A. Leaver and K. Williams. 2013. '(How) Do Devices Matter in Finance.' *Journal of Cultural Economy* 6: 3: 336-52.

Financial Crisis Inquiry Commission (FCIC). 2011. *The Financial Crisis Inquiry Report*. Available at http://fcic-static.law.stanford.edu/cdn_media/fcic-reports/fcic_report_full.pdf.

Goodman, L. S., S. Li, D. J. Lucas, T. A. Zimmerman and F. J. Fabozzi. 2008. *Subprime Mortgage Credit Derivatives.* Hoboken N.J.: John Wiley and Sons.

Gorton, G. 2010. *Slapped in the Face by the Invisible Hand.* Oxford and New York: Oxford University Press.

Gorton, G. and A. Metrick. 2010. 'Securitized Banking and the Run on Repo.' Available at http://ssrn.com/abstract=1440752.

Haldane, A. G. 2009. 'Small Lessons From a Big Crisis.' Remarks at the Federal Reserve Board of Chicago 45th Annual Conference 'Reforming Financial Regulation'', 8 May. Available at www.bankofengland.co.uk/publications/Documents/speeches/2009/speech397.pdf.

Hasanhodzic, J., A. W. Lo and E. Viola. 2009. 'A Computational View of Market Efficiency.' Working Paper, 31 August.

Kirchfeld, A. and J. Simmons. 2008. 'Lehman Toxic Debt Advice Led Leipzig Bank to Ruin via Dublin.' Bloomberg, 28 October. Available at http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aPxfNBXGvA4g.

Lewis, M. 2010. *The Big Short.* London: Allen Lane.

Lewis Baarth LLC. 2011. 'Basis Yield Alpha Fund (Master) v. The Goldman Sachs Group Inc.' Available at http://www.structuredfinancelitigation.com/files/2011/12/GoldmanAustralia.pdf.

MacKenzie D. 2006. *An Engine Not a Camera*. Cambridge MA and London: MIT Press.

MacKenzie, D. 2011. 'The Credit Crisis as a Problem in the Sociology of Knowledge.' *American Journal of Sociology* 116: 6: 1778-1841.

Madrick, J. 2011. *Age of Greed: The Triumph of Finance and the Decline of America.* New York: Alfred A. Knopf.

Mason, P. 2009. *Meltdown: The End of the Age of Greed.* London: Verso.

Maymin, P. Z. 2011. 'Markets are Efficient If and only If P = NP.' *Algorithmic Finance* 1: 1-11.

Milbradt, K. 2011. 'Level 3 assets: Booking profits, concealing losses.' 26 April. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1786985.

Mirowski, P. 2010. 'Inherent Vice: Minsky, Markomata, and the Tendency of Markets to Undermine Themselves.' *Journal of Institutional Economics* 6: 4: 415-43.

Mirowski, P. and Nik-Khah, E. (2007), 'Markets Made Flesh: Performativity, and a Problem in Science Studies, augmented with Consideration of the FCC Auctions.' Pp. 190-224 in *Do Economists Make Markets? On the Performativity of Economics*, edited by D. MacKenzie, Fabian M., and L. Siu. Princeton, NJ: Princeton University Press.

Muniesa, F., Y. Millo and M. Callon. 2007. 'An Introduction to Market Devices' In *Market Devices,* eds. M. Callon, Y. Millo and F. Muniesa. Oxford: Blackwell, 1-12.

New York Times. 2007. 'U.S. Credit Crisis Adds to Gloom in Norway.' 2 December. Available at www.nytimes.com/2007/12/02/world/europe/02norway.html?pagewanted=all.

Pozsar, Z., T. Adrian, A. Ashcraft and H. Boesky. 2010. 'Shadow Banking.' *Federal Reserve Bank of New York Staff Reports* no.458, July.

Salas, C. and D. Hassler. 2007. 'CDOs May Bring Subprime-like Bust for LBOs, Junk Debt.' Bloomberg News, 13 March. Available at www.bloomberg.com/apps/news?pid=newsarchive&sid=a0uS8xp4v2CE.

Senior Supervisors Group. 2009. 'Risk Management Lessons from the Global Banking Crisis of 2008.' October 21. Available at www.newyorkfed.org/newsevents/news/banking/2009/SSG_report.pdf.

Shenn, J. 2010. 'How Wing Chau Helped Neo Default in Merrill CDOs Under SEC View.' Bloomberg New, 9 May. Available at http://www.bloomberg.com/news/2010-05-10/how-wing-chau-helped-neo-default-in-matrix-of-merrill-cdos-under-sec-view.html.

Shivdasani, A. and Y. Wang. 2009. 'Did Structured Credit Fuel the LBO Boom?' Available at http://ssrn.com/abstract=1394421.

Simon, H. 1955. 'A Behavioral Model of Rational Choice.' *Quarterly Journal of Economics* 69: 99-118.

Spears, T. 2014. *Engineering Value, Engineering Risk: What Derivatives Quants Know and What Their Models Do*. University of Edinburgh PhD Thesis.