

Certification and quality signalling: regulatory reform for the ratings industry

Roger Bowden¹ and Peter Posch²

Abstract

In financial contexts such as credit rating and carbon trading, information provision has effectively become certification, highlighting the systemic importance of whether the information is fit for the intended purpose. System regulators can enhance the credibility of information signalling if agencies are offered a registration facility that sequesters part of their fee as a performance bond. Agencies can then signal the value of the rating by choosing whether or not to register and stake part of their fee. A real option analysis can be used to value the registration choice, and hence the amount that is spent by issuer on the incremental quality of research required. The framework can also inform margin and penalty setting by regulators; while an *ex ante* choice to register is a potential defence in professional liability litigation.

JEL reference numbers: G01, G18, G24, G28, G32.

Key words: Financial regulation; credit ratings; incentive alignment; real option; Bayesian decision theory.

¹ Kiwicap Research Ltd, Wellington New Zealand, and Institute of Finance, Faculty of Mathematics and Economics, Ulm University, Germany. Emails roger.bowden@kiwicap.co.nz; roger.bowden@uni-ulm.de .

² Institute of Finance, Faculty of Mathematics and Economics, Ulm University, Germany. Email peter.posch@uni-ulm.de.

1. Introduction

In the wake of the subprime crisis and ensuing credit crunch, the credit rating agencies have been singled out for regulatory attention and political attention. Complaints centred around conflicts of interest and adverse incentives arising because it was the issuer, not the investor, who paid for the rating; self selection biases from shopping around; market power and anticompetitive behaviour; general competence and a specific inability to handle the complexity of structured products; and whether the methodology that underpinned the ratings was fit for the intended purpose, or indeed any purpose at all beyond inflating the agency's own income numbers and stock price. These and cognate issues have been raised by a number of authors (e.g. Skreta & Veldcamp (2008), Liao et al (2009), Cobal *et al* (2009), Mukhopadhyay (2009), Bowden and Lorimer (2009), and White (2010)), in addition to ongoing official inquiries.

It is therefore unsurprising that many countries have evolved new legislation to govern the rating agencies. The US financial regulatory package of July 2010 imposes expert professional liability, bringing the agencies into line with accountants, lawyers and others involved in bond issues to the general public. Operative from June 2010, Treasury regulations in the UK require registration and certification of credit rating agencies and makes senior management personally liable for breaches. Australia has adopted similar provisions with the Australian Securities and Investments Commission (ASIC) as the administering body. Draft legislation from the European Commission envisages fines for proven conflicts of interest of up to 20% of annual income or turnover. A new official body, provisionally entitled the European Security Markets Authority (ESMA), will operate under the umbrella of the European System of Financial Supervisors. The new body will license and monitor European credit ratings agencies and will require a similar licensing regime to apply from offshore agencies for their debt products to be tradable in Europe. Also advocated is empowering the European export credit agency to issue ratings and somewhat oddly, the credit insurers, which would seem at odds with the conflict of interest ruling. Nor has it been ruled out that the ESMA will itself compete with the big three international agencies to rate debt issued in the EU markets. Japan and Hong Kong have indicated that they will fall into line with the registration regime proposed by the Europeans.

It remains to be seen whether such measures offer effective solutions, or are simply knee jerk reactions that are likely to be costly in bureaucracy and counterproductive in their effects. In the US the big three agencies (Moody's, Standard & Poor, and Fitch) are

reportedly already asking for their names to be removed from underwriting documents for new public issues, and the speculation has been that this will only encourage bypassing the public issue market for private placements, with adverse consequences such as high bid-ask spreads on the trading of debt.

A credit rating is technically just a piece of information, hopefully expert, about the likely capacity of the debt issuer to pay interest and repay the principal. The central problem leading up to the subprime crisis was that the agencies – and the big three in particular – had acquired the effective status of certifier for the issue. If an agency rated an MBS or ABS tranche as AAA, then that acquired the force of a sufficient statistic, in the sense that no further investor research was required for a prospective institutional placement. The implicit certification role was reinforced by official rulings in the Basle II capital adequacy regime (Standard Model), corresponding rulings from central banks, and Solvency II for the insurance industry. When the crunch did come, it was clear that the information provided by the agencies was in fact far from sufficient and the effective certification role misplaced. Yet the same problem of implicit or explicit certification crops up in other financial contexts, with potentially adverse fallouts for systemic stability. The carbon emissions market is an instance. It is, if anything, more extreme, for here the objects to be traded are certificates about carbon saving that are not much more than a consultant's opinion. The carbon consultant can assume the role of certifier by government license, much as is now proposed for the credit ratings agencies. But even assuming basic competence and probity as to conflict of interest, a wide range of estimates is possible for even quite standard industrial or agricultural activities, with the attendant possibility of major revisions down the track³.

The economic role of specialist agency information in enabling risk spreading, market making, and lowering the cost of capital is not something that should be handicapped by summary and possible ill considered legislation, let alone the deadweight cost of a bureaucratic superstructure to administer it. In all such cases, the problem arises whether prescriptive licensing and punitive liability provisions will be effective, or whether a combination of the carrot and stick might not be more effective.

The present paper suggests a facultative regime for the monitoring of credit rating agencies. With the approval of the debt issuer, an agency can elect two kinds of ratings, registered and unregistered. In the case of a registered rating, the agency has to deposit part of

³ Indeed, conventional green wisdoms such as 'planting forests is a net carbon sink' have progressively been revised to a net emitter of considerable proportions; Miner (2010) is a recent review. Such things are also a matter of scope.

the fee to be received from the issuer for the rating. The supervisory authority will also act as a stakeholder, refunding the fee plus interest if the track record of the registered rating is satisfactory within predefined limits, which will have attention to the purposes to which the typical investor will put the rating. A registered rating will by that token have extra credibility. In this way, the rating agency and debt issuer together can signal their confidence in the rating and the information provided.

Registration will likely cost more to the debt issuer. However, the issuer thereby gains a valuable real option (register or not) after having seen the rating result, with the signal value of the rating potentially reflected in a higher issue price. The real option aspect analysis can be set within the framework of Bayesian decision theory, with the payoff arising from the ex post value of sample information, interpretable here as the outcome of research to be performed by the ratings agency. The resulting real option value can be used to offset the cost of providing the extra rating quality, and the risk to the rating agency of losing part of the fee. Because the investor is in effect purchasing a signal as to quality, the registration facility mitigates the problem of misplaced agency, wherein it was the issuer alone who commissioned and paid for the rating.

The regulator can use the registration device to steer the rating in the direction of fitness for the intended purpose and to promote longer run objectives of achieving financial stability. It is argued in this respect that a crucial attribute has been missing from the existing body of credit rating practice, namely the maturity dimension. Special structures aside, credit rating for corporate issues is typically based on a ‘through the cycle’ or ‘over the cycle’ approach. Investors have in the past not been fully aware of this limitation. In addition, investors have their own maturity preferences. A BBB investment might be quite acceptable in a good state of the business cycle provided the intended holding maturity is short. Conversely, a given rating might be quite acceptable on an ‘over the cycle’ basis, but more exposed to temporary bad states of the business cycle. The time or maturity aspect can be summarised with the conditional expected time to default. Thus a given rating scale (AAA or Aaa downwards) can be put into alternative correspondences, one with probabilities of default and the other with the expected time to default. Both are useful ways of thinking about rating outcomes.

The scheme of the rest of the paper is as follows. Section 2 provides some background, noting that in the financial contexts of credit rating and carbon trading, information provision has effectively become certification. This highlights the social importance of whether the information is fit for the intended purpose. Section 3 elaborates in the context of credit rating

agencies, showing how the standard ‘through the cycle’ methodology needs to be augmented with the temporal dimension required by investors. Section 4 presents the core policy proposal: credit ratings can be registered, wherein part of the agency fee is sequestered contingent on performance, or unregistered. Section 5 turns to the demand for registration establishing a real options model for the value of registration, based in Bayesian decision theory. Section 6 contains some concluding remarks. The real options framework can inform margin and penalty setting by regulators; while an *ex ante* choice to register is a potential defence in professional liability litigation. A similar registration framework could also be used in the context of carbon trading.

2. Certification and credibility

On the face of it, credit rating and other applications in the finance industry are just a form of quality signalling in a situation where learning from experience about producer credibility would be very expensive for the buyers. From this point of view, it is quite in order for a firm trying to sell its debt issue to investors to hire credible third parties to signal the quality of the debt issue. The agency relationship is perfectly normal and quite explicit: the rating agencies are paid agents of the issuer. Concern that the debt issuer pays but the debt buyer is to benefit, is not necessarily a case of misplaced agency. The matter is rather similar to the distinction between real estate agents that can theoretically work for either seller or buyer, with different informational commissions reflecting the different reporting responsibilities.

What was unusual, at least prior to the subprime crisis, was that the credit rating agencies maintained their credibility when no formal revelatory mechanism existed to validate it. In an underwriting context⁴, for instance, investment banks signal the quality of the issue by establishing a price floor, charging the issuer a high fee – effectively an insurance premium – to absorb any risk to the underwriter from underpricing. Absent any such mechanism, the burden of maintaining credibility falls back on reputation and this in turn relies on track record. In the case of the credit rating agencies, their credibility might historically have resulted from the original arrangement where the issue buyers paid for the rating. Even if such memories persisted, which is problematical, then once the issue sellers were to meet the cost, things should have been viewed differently. Lewis (2010) and a series of US Senate hearings⁵ in 2010 document the abuses that resulted from managerial pressure on their line

⁴ Adriana *et al* (2009) discuss potential inefficiencies involved in this context.

⁵ US Congress Financial Inquiry Commission (FCIC). See also US Senate (2010) FCIC staff report.

research staff to hold on to the agency's market share. But the market as a whole was not aware of this at the time. Indeed the credibility of the major agencies was so strong, and their ratings so important for systemic stability, that it amounted to a form of certification.

At its most general level, certification carries a connotation that the body or object being certified is qualified to belong to some actual or implicit club (individuals or organisations) or a category of objects. To do that requires the certifier to have a body of qualifying information, even if it is as simple as the observed ability to swim 50 metres. For more complex social tasks, the boundary between the roles of certifier and the weaker role of information provider is sometimes blurred. For instance, credit rating agencies commonly claim their ratings are merely information about the financial ability of a firm to service its debt. But to the extent that the market requires a rating in order for the firm to be able to issue debt, then the information has tacitly acquired a stronger role, one of certification. Indeed, much of the edifice of special purpose vehicles prior to the subprime crisis was erected on the common understanding that a specific MBS or ABS structure or tranche could not be issued without a AAA (Aaa etc.) rating, or without being backed via CDS by a counterparty that possessed such a rating. In turn, this created huge client pressures on the rating agency to issue a good rating, or suffer the loss of business to a competing agency.

The nexus between information and certification becomes even stronger when official bodies of one kind or another stamp the rating as qualifying for some regulatory purpose or function. Thus under the Standard Model for the Basle II regulatory regime, asset credit ratings automatically qualify for designated capital ratios. Capital requirements vary from 0% (S&P equivalent AA- for sovereigns) to 150% for claims on corporates with an S&P rating below BB-, or on banks and sovereigns with an S&P rating below B- . Some central banks have similar rules for parallel regulatory regimes, where they have not signed up to the Basle regime.

In the different context of emissions trading schemes (ETS), private consulting firms prepare estimates of the carbon content or carbon saving of a specific activity. These become tradable certificates. The transformation from merely somebody's opinion⁶ to a tradable

⁶ Most financial markets trade on the basis of buyer and seller opinions, but about objects (shares, bonds etc.) that are unequivocally defined. In ETS markets, however, the tradable objects are themselves opinions. This is an endemic weakness of the market, for opinions are not only estimates of carbon content, but potential revisable in the future if better knowledge becomes available. To be sure, pork bellies and the like in commodity markets do have to be graded, so the distinction is not absolute. But a misclassified pork belly does not hang around for very long.

certificate is likewise achieved by government certification that the consultant is competent to prepare opinions that qualify as tradable objects in the market.

Where an actual or implied official recognition exists, two layers of certification are evidently involved in such applications, the primary layer being for the object, the secondary layer being for the certifier itself. The primary certifier effectively assumes an agency relationship with the secondary certifier. Thus credit rating firms become de facto agents of the Basle Committee or of the local central banks. Carbon consultants effectively become agents for the ETS market clearing house; or at one remove, of their respective governments or the IPCC.

A useful distinction can be made between weak or self referential certification and strong certification. In many contexts, certification is simply for the purposes of membership of an implied club. Once this is achieved, there is little or no further external impact or social consequences, beyond the money spent in achieving the certification; the prisoner's dilemma game provides an analogy⁷. Once the firm has achieved certification, the role of the certifier is ended, in the sense that subsequent product or service failures are not sheeted home to the certifier. In such cases, certifying criteria may well be incomplete in terms of the more embracing public good, but that is not the moral or legal responsibility of the certifier. In this sense, such certification is self referential.

'Strong certification', on the other hand, does require attention to the possibility that the rules themselves may be incomplete or even wrong in terms of some ulterior objective, and moreover there is a public interest involved in the outcome. Because it may be proved inadequate by subsequent events, certification of this kind is empirically testable against its stated objectives. Even more demanding, strong certification has to be robust against the possibility that its inadequacies may precipitate the very events it has been designed to avoid, a socially damaging form of adverse incentive where the dominant player plays a strategy that leads to a counter-productive reaction. Basle II certification has been singled out in this respect: so called 'regulatory arbitrage' induced banks to economise on scarce capital and take advantage of the high credit rating of structured credit tranches.

⁷ Business school certifications such as AACSB are arguably of this character. Their standards and rubrics are advertised as promoting good teaching, but ex post do not actually assure good teaching, indeed may detract from it because of the resources diverted to achieving accreditation. ISO 9001 in the commercial world operates off the same principle, in this case a set of procedural rules that the certifier thinks are conducive to delivery of a quality service.

In general, strong certification arises in contexts where there are significant social externalities if certification fails. Recriminations after the subprime crisis and subsequent international credit crunch laid blame at the doors of the credit ratings agencies. Complaints ranged from undue influence derived from close association with, or technical dependence on, those being rated (i.e. debt issuers); lax standards in pursuing or maintaining market share; technical inadequacies e.g. in estimating default probabilities of ordered credit tranches; and inappropriate models or modelling philosophy (see below). The problem was that information, as the agencies might have put it, had *de facto* become certification. Managers of synthetic CDS funds and other investor structures would take the credit raters' verdicts as sufficient statistics for their own investment decisions. An AAA rating, for example, absolved investment managers of any further responsibility to acquire further information about the asset or counterparty with whom they proposed to deal. Apparent endorsement of the expert status of the agencies came via their official certification status for purposes of Basle capital adequacy or equivalent acceptance by local financial regulators or central banks.

Thus a crucial issue in all this is whether the models and procedures used by the ratings agencies were fit for the intended purpose actually intended by investors and regulators, and what devices can be developed to enhance confidence that they are. In one aspect, this is matter of the technical expertise with which the implied tests were constructed and applied. But another concern arose because of a misunderstanding (as the agencies would claim) of the nature of the rating itself. The investors that relied on the rating had one understanding of the meaning of the rating; the agencies themselves had quite another. On this view, confusion as to purpose and meaning of credit ratings was an important cause of the subprime crisis. 'Fit for the intended purpose' therefore depended on an understanding as to purpose, quite as much or more than fitness. This particular aspect is explored in the next section.

3. Fit for what purpose? The time dimension

Ratings as constructed by the agencies consist of a scale intended to represent financial soundness. But does the implicit scale correspond to what the investors actually need? Any notion that there was a correspondence with probabilities of default took a knock in the aftermath of the subprime crisis, when the agencies acknowledged that an AAA corporate debt rating did not evidently entail the same probability as an AAA sovereign rating. But probability is a cardinal, not an ordinal measure. One cannot have two sets of probabilities, both adding up to unity. This would appear to leave the agency ratings, even when correctly

constructed in their own terms, as little more than a subjective ordinal rankings of apparent financial soundness relative to some peer group of investments. But even if a mapping to probabilities could be established, investors also need a time dimension for the ratings, for this affects the maturity choice for their investments. Does a BBB rating mean that there is a likelihood that the debt will default within a month; a year; or just at some time within the next ten years? In an ideal rating scale both correspondences should be established, one to estimated probabilities and the other to probabilities over stated periods of time, or equivalently, the expected time to default.

Credit ratings as constructed and sold by major agencies were, and usually still are, based on the ‘over the cycle’ or ‘through the cycle’ philosophy. With this approach, a credit rating is an estimated mapping from a set of financial variables or more particularly ratios (liquidity ratios, leverage ratios etc.) to probabilities of default or downgrade on the debt of the firm. The estimation sample space is a population of observed histories of similar firms over time, taken as long enough to encompass a range of business conditions, hence ‘over the cycle’. On the basis of estimated logits or probits (or similar scoring based methodology), a probability of default or downgrade is assigned to the debt being rated. The logits or probits themselves are conventionally estimated on the basis of pooled data. Each debt issue continues in the data pool until it is eliminated when and if it finally defaults. Up to that point, its associated financial ratios for each year enter equally into the data pool together with a no default outcome ($d = 0$) and the final default outcome ($d = 1$) at which it exits the pool. Possible selection bias in this data construction is econometrically problematic, but apart from this, little or no account is taken as to ambient business cycle conditions for each of the included years. Likewise in structured finance ratings, the sensitivity of subprime mortgages or commercial loans to a general business downturn, even for senior tranches, was revealed to be extreme (Coval *et al* 2009).

In the event, ratings constructed in one frame of reference, namely ‘over the cycle’, were evidently interpreted as valid for another, namely as referring to a probability of survival over the coming time period. Thus an AAA rating was tacitly accepted as denoting a very low probability of default in any stage of the business cycle and without reference to any developing economic or structural circumstances. Misunderstandings or communication problems aside, this raises the question as to whether the ‘over the cycle’ philosophy has much merit to begin with, i.e. just what information it can impart, and how the short run is embedded in the long run. If investors are concerned with the probability of survival, over

any given horizon, then the time dimension has to be explicitly introduced into the rating framework.

Define variables and functions as follows:

q = Scalar or vector metrics for financial strength or prudential reserves policy (often abbreviated to just ‘financial ratios’). A high q firm is sounder, more robust to bad states of the world and deserves a better credit rating. It is useful to think of q as an index, such as a probit or logit score for credit worthiness.

x = Scalar or vector indicator for the state of the business cycle and its effect on the firm’s operating environment; hence x_t at real time t is the current state of the cycle. Better states of the business cycle are represented by higher values of x .

τ = time to default.

$F(\tau; q, x_t) = F(\tau; q | x_t)$ = probability distribution function of time to default τ , given the current state of the business cycle and a designated value q for financial strength. Thus default could take place at $\tau = t + 1, t + 2, \dots$ or never ($\tau \rightarrow \infty$). The corresponding density is denoted $f(\tau; q, x_t)$.

Remark: Financial ratios could well vary over the cycle, so one could write $q_t = q(x_t, q^0)$, where q^0 is a state independent measure of financial quality, or the soundness of the firm’s policy in maintaining adequate prudential ratios or reserves. This can be absorbed within the above framework by indentifying q with q^0 , and this will be done in what follows.

Assume that $\frac{\partial F}{\partial q} < 0$; $\frac{\partial F}{\partial x_t} < 0$; $\frac{\partial^2 F}{\partial f \partial x_t} > 0$, uniformly in τ . Good financial ratios shift the

distribution of time to default to the right, likewise good states of the business cycle. To interpret the second derivative (interaction between q and x_t) suppose that x_t declines, i.e. $dx_t < 0$. Then for any given τ , the value of F should increase (which is bad). It will increase by less if q is higher, because this means firm has better financial ratios and can withstand adversity better.

An associated metric is the *conditional expected time to default* $\mu(q, x_t) = E[\tau | q, x_t]$. This may not exist for relevant distributions, since the time scale and hence right hand tail is inherently very long. An alternative is the *conditional median time to default*, $m(q, x_t)$.

The expected time to default (ETD) or median time to default (MTD) will be sensitive to states of the business cycle. For any given financial ratio vector q , the ETD or MTD will shorten in bad states. The extent to which this is true will depend upon the financial ratios q . Sound firms have high q and this enables them to withstand bad states better. Hence F stays reasonably stable though with some slight cyclical movement. On the hand, firms with poor q will exhibit a marked shortening of ETD or MTD in bad states of the business cycle. One implication is that a BBB firm has a higher probability of default over short horizons, especially in bad times. An AAA firm (~high q) maintains a lower probability of default over short horizons, even in bad states of the world. If An AAA firm defaults within a year this would be more indicative of an error in the credit rating.

The meaning of ‘over the cycle’ can be set within the above framework. Submerge or disregard the real time setting of the state of the business cycle, so that x_t becomes just x , which it is treated as an observation from a frequency function $\Psi(x)$ assembled by pooling data on x_t over all stages of the business cycle. More precisely, a long run stationary distribution is assumed to exist to any short run correlation in deviations from trend GDP and other relevant variables⁸.

Define the distribution function

$$(1) \quad \Phi(\tau; q) = \int_x F(\tau; q, x) d\Psi(x) = E_x[F(\tau; q, x)].$$

So $\Phi(\tau; q)$ is the unconditional distribution function of time to default, over all states of the business cycle. Moreover, the unconditional median time to default τ_m such that

$\Phi(\tau_m; q) = \frac{1}{2}$ is positively related to financial ratio quality:

$$(2) \quad \frac{\partial \tau_m}{\partial q} = -E_x \left[\frac{\partial F(\tau_m; q, x)}{\partial q} \right] / E_x [f(\tau_m; q, x)] > 0.$$

Expression (2) embodies the essence of the ‘over the cycle’ approach. On the average over the cycle, the expected time to default is greater the better the financial ratios. An AAA rating on the ‘over the cycle’ approach means that the ETD or MTD is very high, on the average over all states of the business cycle. A BBB rating (lower q) means that the ETD or MTD are on the average lower.

The trouble with the ‘over the cycle’ approach is that it ignores information about the current state of the world x , under which an apparently satisfactory ‘over the cycle’ rating

⁸ This means the distribution of $x_{t+\tau}$ should lose any dependence on x_t as $\tau \rightarrow \infty$, apart from trend.

could be unsatisfactory in a currently adverse state of the world. Thus with a given value of q , the unconditional MTD might be 10 years, suggesting quite a good credit rating for shorter dated issues. But in a bad current state a conditional MTD, i.e. one based on $F(\tau; q, x_t) = F(\tau; q | x_t)$, might shorten to just two years or so.

Figure 1 illustrates the distinction in terms of the median or expected time to default. Here there are two firms with identical business models and the same expected time to default when the latter is measured at the very best state of the world. One firm has very sound financial ratio policy q_H , and the other not, q_L . In normal states of the world (shaded area), there is little difference between the expected time to default for the two firms: q_L looks quite OK. But the difference widens dramatically in bad states (compare the two double headed arrows).




Figure 1: State dependence of median time to default

Using the unconditional MTD based on $\Phi(\tau; q)$ would mean that the given value of q is sufficient to withstand temporary stress, provided that the business cycle changes direction and improves, as is implicitly the case with the stationary distribution assumption. But it may not be sufficient to withstand a more prolonged stay in the left hand tail of the presumed stationary distribution, or to cope with a structural shift that might shift the stationary distribution itself. Using $F(\tau; q, x_t)$ would highlight the potential exposure to a longer than usual recession or a structural shift. Investors presented with both $\Phi(\tau; q)$ and $F(\tau; q, x_t)$ might take a quite different action that if presented with $\Phi(\tau; q)$ alone.

In other words, to base a rating solely on the unconditional ‘over the cycle’ approach is inappropriate and could even be dangerous. Nor does it materially help to solve the problem of fixed interest investment in real time, which involves choices as to maturity and their possible dependence on current states of the world. The discussion that follows develops mechanisms that signal a responsiveness and credibility in real time, rather than a notional ‘over the cycle’ or ‘on the average’ time.

4. Regulatory measures and aligned incentive structures

As noted in section 1, one approach to ensuring adequate certification outcomes is a public certification agency to make the ratings. The motivation behind such a suggestion arises from the systemic importance of this kind of certification, and its incorporation into regulatory rules for capital or liquidity adequacy. Public agencies, on this view, are less subject to related party pressures and more subject to scrutiny.

On the other hand, the public agency suggestion is problematic on a number of grounds. The competitive status with respect to the private agencies has not been elucidated, either as to how it would work, whether it would yield any better results, or whether it is even desirable to have public-private competition in such a function. The suggestion that follows does involve a public role in credit certification, but it centres on improving the revelatory function of private sector rating agencies and their judgements. It works by discriminating ratings in which the agencies have confidence from those in which they do not, creating in this sense a separating equilibrium observable by investors.

4.1 Surety margins

If credit rating agencies have any confidence in their ratings, then they should be prepared to place what is in effect a bet that a given rating will stand up within its own predefined limits. The present proposal is that a public supervisory agency should also run a margin system, in which the private agencies could choose to deposit a pre-defined portion of the fee received for making the rating. The fee is the performance margin. If a rating fails within a designated time period, e.g. because an AAA entity defaults within year, then some or all of the margin is retained by the public agency. Otherwise the margin is returned, along with interest earned on the sum deposited. Such an arrangement is analogous to a clearing house margin in derivatives exchanges; the ratings agency is tacitly trading its own reputation as an object of value.

Within this frame of reference there would be two kinds of rating, registered and unregistered. A *registered rating* is one for which surety margin has been paid. An *unregistered rating* is one for which no margin has been deposited. The surety margin is a performance bond. It is refundable if inconsistency arises between ratings prediction and subsequent history of the issue, with emphasis on ex post adverse outcomes. The surety margin is posted by the credit rating agency, with the consent of the firm issuing the debt to be rated. In many, if not most, instances the impetus to register a rating would arise with the issuer; section 5 elaborates on the economics involved in the demand for registration. From the administrative point of view, however, it is more workable to have the credit rating agency register an issue's rating.

A registered rating would by that token carry a connotation that the ratings agency is fairly sure of its ground, giving investors in the prospective debt issue a degree of comfort, of the kind that has evidently been lost in the wake of the subprime crisis. That being the case, rating agencies could charge more for a registered rating, defraying to some extent the risk of failure from unforeseeable states of the world, or the opportunity costs to shareholders of funds sequestered in posting margins with the registrar of the scheme. Natural limits on the amounts to be charged are disused in section 5.

An unregistered failure carries an implicit health warning to investors in that debt issue. The signal is that the implicit confidence limits for the rating are wider, the actual rating being regarded as the position indicator. The margin penalty structure is also of importance, i.e. the extent to which the margin is to be deducted in the event of some adverse outcome. The penalty should be responsive in progressive fashion to the degree of failure, just as with a progressive taxation scale, for this is a form of contingent taxation. This would provide an incentive for the agencies to periodically revise their ratings in the area of most sensitivity, namely for downgrades. Total and sudden default would carry a much greater penalty than getting there by stages. The latter approach entails less costly unpleasant surprises for investors and the system as a whole.

4.2 Margin setting

The setting of surety margins has to be responsive to the willingness of the agencies to pay them from time to time; there is not much point in having a registration option when nobody will ever use it. However before proceeding to discuss demand for registration, it is useful to canvass possible criteria to be used by the agency in setting comparative margins along dimensions such as ratings and issue maturities.

Surety margins should be responsive to dimensions such as the quality of the rating and the maturity of the debt instrument being rated. In this respect, ratings decisions should be guided by the information to be conveyed by the rating and by the costs and consequences of deficiencies. For just as in hypothesis testing, type II error is involved, where a rating is given that turns out to be incorrect, with costs to the users. Given the existence of such social costs, they should ideally be internalised in the surety margin structure.

Thus it should be expected that the higher the credit rating given to the debt issue, the higher the surety margin. A downgrade in a higher rating typically involves greater costs to investors and to the system as whole. This is not just matter of unsophisticated ‘mum and dad’ investors seeking safety; regulatory and related requirements are typically predicated on qualifying ratings. The Basle Standard model for bank capital adequacy has already been mentioned in this respect. As another instance, authorised trustee investment regulations commonly specify minimum allowable ratings, minimally investment grade (BBB), but often more stringent.

In addition, a high rating conveys to the user more certainty on the part of the rater as to the investment outcome. Conversely, a low rating conveys overtones of uncertainty more than does a high rated one. The strength or otherwise of the rating is in part a signal about the agnosticism or otherwise of the rating agency, calling for a greater degree of caution from the investor. In such an aspect one could think of the investor in an unseasoned issue as starting out with a Bayesian prior of BBB equivalent, i.e. that the proposed issue is at the margin of acceptability. If the rating agency now comes up with an AAA rating this creates a much more informative posterior, so the information content of the rating is the greater. Of course this could work both ways. An agency that come out with a BBB rating for US sovereign issues would be conveying some serious information relative to the customary investor prior of AAA. In this case, the type II error would be a more minor matter of opportunity cost in a decision not to invest, rather than the contingent cost of default.

The maturity of the proposed debt issue is also cognate. If a highly rated shorter dated issue collapses within a year, this could mean one of two things. The first is that the rater has missed some important warning signals and is therefore deficient in information on this count. Collapse of the structured investment vehicles during the subprime crisis is the obvious example, where the agencies failed to make the connection between information on deteriorating US consumer delinquency data and the home mortgages that made up the SIV portfolios. The second is that the rating itself has been of the ‘through the cycle’ variety. A

surety margin system that penalises early default would encourage the agencies to develop more state dependent procedures; or if this is not possible, to rate more conservatively in the chosen approach.

5. The demand for registration

The question of when and how much an agency is willing to pay for registration, and accept a contingent penalty, depends on how much the client (the debt issuer) is willing to pay in the form of an enhanced fee for a registrable rating. The latter question can be explored by combining Bayesian decision theory⁹ with the theory of real options.

To prepare a rating requires research by the rating agency. The outcome (i.e. the rating) is analogous to a sample estimate of the objectively correct rating, which could be thought of as embodying a one to one correspondence with investor welfare dimensions such as the probability of default or the expected time to default. In the first instance, the agency will be concerned with discovery of the quality of the issuing firm's financial ratios or management policies. However, the rating should ideally also encompass consequential metrics such as the probability of default $F(\tau_s; q, x_t)$ over selected maturities τ_s or else the median time to default $m(q, x_t)$, in each case conditional on the current state of the world.

For expositional purposes, the object to be estimated will be taken as some index of quality, which can be vector valued; q is the true value and \hat{q} is the agency's estimate on which the rating will be based, such as an estimated probit or logit score .

The more resources the agency devotes to the estimation, the closer is the estimate likely to be to the true or correct rating. But there will always be limitations of discoverability or the ability to assess future exposures and their consequences; so the posterior distribution never quite collapses on to the true value. A favourable outcome and the decision to register it endows the debt issue with incremental value. On the other hand, if the additional research is unfavourable the issuer can decide not to register it. The balance between the two creates a potential for real option value.

In addition to the positive outcome created by registering a favourable outcome, the real option value depends also upon how investors perceive the negative signalling effects of a decision not to register. The exposition that follows looks at two possibilities. In the first, investors treat a decision not to register as creating an uninformative outcome, just as though

⁹ Berger (1985) is a standard reference.

the additional research had never been done. In the second, investors interpret the decision not to register as indicative of a bad outcome.

5.1 Registration decision tree

Figure 2 depicts the decision tree under the first of the above alternatives. The agency starts with a Bayesian prior distribution for the true quality q , the basic or uninformative distribution U corresponding to a routine level of effort (the branch marked $c = 0$). It is assumed that registration will not be worthwhile on this branch¹⁰.

Alternatively, the agency can choose to put in a designated level of effort at cost $c > 0$, and prepare a rating using the information thus acquired. The issuer commissioning the rating sees the estimate \hat{q} and now has a choice: register, or not register. So there are two branches to the decision tree at this point:

(a) If the estimate is good (G), relative to the prior U, then the agency will ask the agency to register, signalling thereby to the investors that this is a better issue than they might have otherwise expected. The benefit is reflected in a high issue price for the debt, which is the payoff for this branch of the implied decision tree.

(b) If the rating estimate reveals little new information, then the issuer will ask the agency not to register it. Alternatively, the agency itself may choose not to register if it has less confidence in the numerical rating, mindful of the penalty (as in section IV) if the issue does default. This is the base case (U).

To illustrate, suppose that the true quality q has a Bayesian prior ‘base case’ distribution $p(q)$ marked as U in figure 2. An objective of the credit rating is to sharpen up on this distribution in the minds of investors, which can refer to changes in both location and scale (dispersion) dimensions; density G is the favourable possible outcome from spending the extra amount on research. If the research reveals no new information, then the result stays at the uninformative prior U. (The dotted density B can be ignored at this point.)

¹⁰ This branch could technically be subsumed by putting $c = 0$ in the upper branch; however it is kept for the sake of contrast.

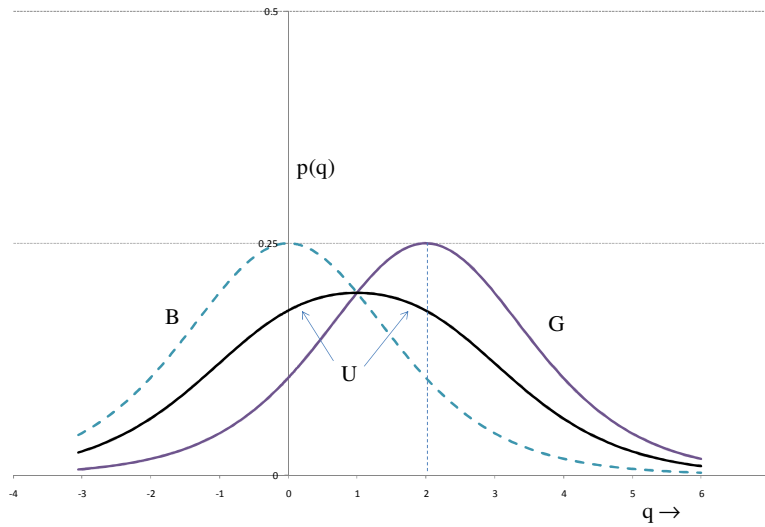


Figure 2: Prior and posterior densities for issue quality

The issuer's ability to choose how to proceed once the rating estimate is received amounts to a real option, with a valuable payoff under branch (a) above. The resulting decision tree is depicted in figure 3, with likely choices, as to register or not, highlighted as bold lines.

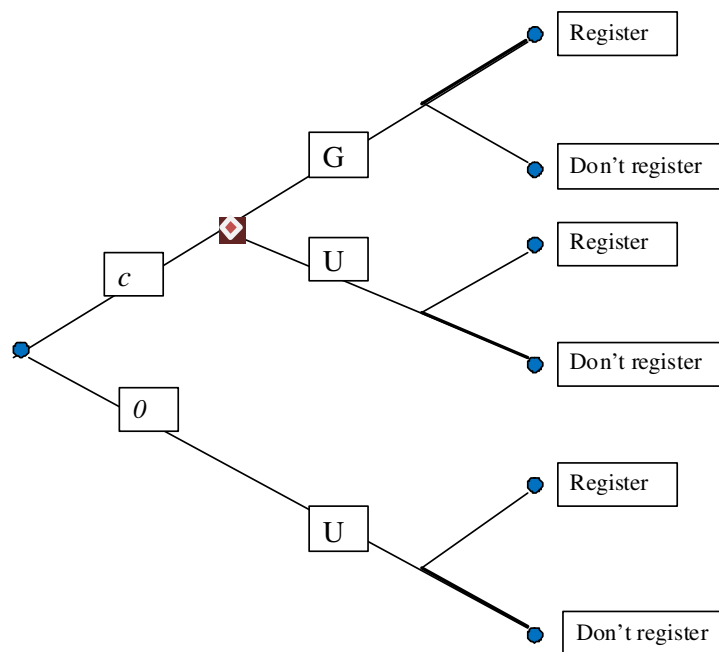


Figure 3: Decision tree for real option analysis

5.2 The expected value of rating information

The value of the implied real option sets the upper limit for the cost c to be paid to the credit rating agency and therefore the degree of effort to be put into the rating, the actual payment c to be determined by bargaining and by competition among agencies. In the language of Bayesian decision theory, the source of the option value corresponds to the expected value of sample information (EVSI); or in this context, the expected value of rating information (EVRI). The discussion that follows develops the connection.

The agency is to prepare an estimate \hat{q} with reliability in the form of a conditional density for the estimate given the value q , namely $p(\hat{q} | q; c)$, that depends also upon the effort and cost c devoted to the purpose. The resulting unconditional density for the estimate is denoted by

$$(3) \quad p(\hat{q}; c) = \int p(\hat{q} | q; c) p(q) dq.$$

Given the prior $p(q)$, the unconditional estimated density $p(\hat{q}; c)$ as in (3), and the reliability $p(\hat{q} | q; c)$, the Bayesian posterior density for the true quality q is given by:

$$(4) \quad p(q | \hat{q}; c) = \frac{p(\hat{q} | q; c) p(q)}{p(\hat{q}; c)}.$$

Suppose the agency's estimate indicates that density G in figure 3 is much more likely to obtain, i.e. is such that it favours a posterior $p(q | \hat{q}; c)$ close to G. Then the debt issuer will tell the agency to go ahead and register the rating, and the agency will likely concur, since it knows the reliability of its estimation procedure is high. The extra credibility that derives from registration will result in a higher issue price or proceeds for the newly rated debt. Alternatively, suppose that the posterior (4) is much closer to the uninformative outcome U. In this case there is no motivation on the part of the issuer to register and it can choose not to. Relative to the lower branch, it is worse off by the loss of the additional research fee, but seen *ex ante* it remains better off because the research fee itself has to factor in the possibility of a negative outcome.

More generally, the posterior value improvement arising from the agency's estimate is given by

$$(5a) \quad v(\hat{q}; c) = E_q[V; p(q | \hat{q}; c)] - E_q[V; p(q)].$$

If the quality q were known for certain then the issue would be worth $V(q)$. However investors do not know the precise quality, so the value of the issue is determined by their expected value over q . Thus in expression (5a), the term $E_q[V; p(q)]$ refers to the issue value arising from the uninformative null U , while the term $E_q[V; p(q|\hat{q};c)]$ refers to their expected value given that an estimate \hat{q} has been revealed. If the posterior distribution is revealed as close to U , then there is little or no value revealed to spending the extra sum c on research, in which case the value improvement $v(\hat{q};c)$ is zero or close to it. In this case, the issue will not be registered. On the other hand, if the posterior distribution is close to G , then $v(\hat{q};c)$ is positive. This will be signalled to investors by registering the rating.

The contingent payoff for the implied real option held by the debt issuer can be defined as
(5b) $v_+(\hat{q};c) = \max[v(\hat{q};c), 0]$.

Ex ante, the issuer does not know in advance the outcome \hat{q} , e.g. whether it will favour G or B in terms of the example. But taking the expectation over \hat{q} gives the expected option value of the information to be provided by the rating agency:

$$(5) \quad v_+^e(c) = EVRI = \int_{\hat{q}} v_+(\hat{q};c) p(\hat{q};c) d\hat{q}.$$

The quantity v_+^e is the expected value associated with the implied real option held by the issuer; in risk neutral world it would be its actual value. It will be diminished by the amount c to be spent on preparing the rating and will itself constitute an upper limit for that amount. The actual sum to be spent will hinge upon a bargaining game between the agency and the issuer as its client, in which environmental variables such as the extent of competition among ratings agencies will be important. In this respect, it should be noted that the amount spent c is not the only source of real option value. What is important is the amount of information conveyed to the investing public, which will include the signalling effect on the part of the agent in registering the rating.

5.3 Option value with unfavourable contingencies

Referring to figure 2, the base density U could itself be made up as a mixture of two densities, an unfavourable one marked as B (bad) and a favourable one G (good). The implied real option can have *ex ante* value even where the extra research could *ex post* reveal a bad outcome as in density B of figure 3. Choosing not to register the rating preserves uncertainty in the investor's mind, so the subjective density as seen by the investor is a fatter

version of B, i.e. one with higher uncertainty more akin to that of the original uninformative distribution U. If the underlying value curve V for the issue is convex upwards, then the prospective gains from registering a good outcome, as in G, outweigh the losses entailed in not registering a bad outcome B.

The real option to register the rating therefore provides a channel whereby the issuer can signal to the investors the quality of the rating. The agency cost becomes shared between the issuing firm and the investing public, for the latter end up indirectly paying some of the rating agency fee via a higher price for a registered issue. The registration scheme therefore indirectly addresses the public concern, noted in section I, about adverse incentives in existing payment arrangements.

6. Concluding remarks

1. As indicated in section 4.2, considerations of systemic stability are the primary consideration governing the setting of surety margins by the regulator. However, there has also to be sense of what the agencies can afford to pay, for otherwise the registration option will never be selected. In other words, demand is important as well as supply. The implied real option content could help to inform regulator decisions as to dimensions such as the proportion of rating fees to be withheld as surety margins, the validity period of the registration, and defined events which would induce partial or complete confiscation. Thus suppose the sequestration was set at 50% of the fee. An agency facing a contingent loss of 50% would have to ensure that the fee charged was high to begin with, which in turn means that the real option value to the client (the issuer) would have to be high. If the regulator wishes to encourage situations where the rating provides genuine and valuable information to the investing public, then a surety margin as high as 50% might make good economic sense. A similarly high fee could be charged for shorter dated registration maturities. In this respect, raters would have the option of specifying the registration period. This would assist investors in making it clear that there was a confidence band, as it were, around the maturity involved. If the agency chooses a registration period of 2 years for a ten year bond issue, that is a signal to investors that the survival of the issuer might be seen as contingent on the current state of the business cycle. It would, in any, case be unreasonable to expect a sequestration rate as high as 50% for the complete tenor of a ten year bond. If the registration is to cover the entire ten years than a reducing scale could be employed: e.g. 50% for the first year, 30% for the second, 10% for the ensuing years.

2. Registration is not seen as a substitute for professional liability, as it appears in recent legislation, which needs to be retained to cover misdemeanours such as aggravated conflict of interest. However, the willingness of the credit rating agency to propose or consent to registration of the subject rating might count as a defence or mitigation against a charge entailing professional liability. The act of registration, in other words, is a signal that the agency has prepared the rating in good faith, backed by the necessary research, and has done its best to ensure that the rating is fit for the intended purpose.

3. Carbon certification could in principle be explored within the same framework of registration. Suppose a firm wish to sell carbon credits for a new economic activity that it hopes will have a net carbon saving. To do that, it will need certification from a registered carbon certifier. The actual outcome of the measurement/estimation could be good (carbon saving), or it could even be bad (e.g. the new activity is actually a net emitter). The firm will know which once the certifier has reported back on his estimate. If it turns out to be good, then the firm will have a carbon certificate as a tradable product. If it turns out to be bad, then the firm will have the option of not proceeding (and the moral quandary as to whether to tell the government or IPCC if it is in fact negative). The *ex ante* expected real option value sets an upper limit on the cost to be paid for certification.

References

Adriani, F, L. Deidda and S. Sonderegger (2009) The role of financial intermediaries in securities issues: A theoretical analysis. MPRA Paper No. 16112, Munich University; <http://mpra.ub.uni-muenchen.de/16112/>.

Auriol E. and S.G.M. Schilizzi (2003) Quality signaling through certification: theory and an application to agricultural seed markets. Working paper, Agricultural and Resource Economics, University of Western Australia.

Berger, J. O. (1985) *Statistical decision theory and Bayesian analysis*, New York: Springer Verlag.

Bowden, R., and D. Lorimer (2009) Credit securitisation and system stability. Working paper, Institute of Finance, Ulm University and Victoria University of Wellington.

Coval, J., J. Jurek, and E. Stafford (2009) The economics of structured finance. *Journal of Economic Perspectives*, 23, 3-25.

Duffie, D., and H. Zhu (2009) Does a central clearing counterparty reduce counterparty risk? Working paper 2022, Stanford University School of Business.

Lewis, M. (2010) *The big short: inside the doomsday machine*. Penguin/Allen Lane.

Liao, H-H, Chen, T-K; Lu, C-W (2009) Bank credit risk and structural credit models: Agency and information asymmetry perspectives. *Journal of Banking & Finance*.33,8,1520-1530.

Miner, R. (2010) Impact of the global forestry industry on atmospheric greenhouse gases
FAO Forestry Paper # 2010.

Mukhopadhyay, B. (2009) Are Competitive Rating Agencies Efficient? *Journal of Emerging Market Finance*, 8(1), 67-85.

Skreta V., and L. Veldkamp (2009) Rating shopping and asset complexity: A theory of ratings inflation. *Journal of Monetary Economics* 56(5), 678-695.

US Senate (2010) Credit ratings and the financial crisis: Inquiry Commission Preliminary Staff Report June 2.

White, L.J. (2010) Markets: The credit rating agencies. *Journal of Economic Perspectives*, 24(2): 211–26.