Value of trust and trust in values

Theory, practice and possibilities for financial ratings

Pietro Battiston

March 1, 2010

"I am upset not because you lied to me, but because I can no longer believe in you."
Friedrich Wilhelm Nietzsche, Jenseits von Gut und Böse, *Chap. IV*

1 The framework

1.1 Introduction

No voluntary long-term relation of any kind between human beings is possible without them having *trust* in the fact that it will develop as intended. This is obviously true for what we may consider "more human" relationships between two person, as well as for the relation that an individual has with the State and with any single or collective entity.

This observation lies at the beginning of this analysis because it is its fundament and because the main "limitation" it recognizes - the statement doesn't apply to *coercion* mechanisms - fits perfectly in the framework on a study on *market*, considered to be *free*, in the sense that participants continually have to make choices, but even more in the sense that they choose to participate in the market itself.

Usually, trust is based on personal belief of some particular *qualities* of the counterpart. However, this is not strictly necessary: one may trust a counterpart just based on the assumption that it has no interest in cheating and avoiding respecting any agreement made. In some other cases, my trust may be based on the knowledge (or belief) that the counterpart has developed similar agreements - with me or others individuals - in the past and respected them, or more generally that other individuals have trust in it. Finally, there is the important, and slightly different from the ones reported so far, situation in which I may trust another entity because I know it is *obliged* to respect any agreement made (and I trust the influence of the body that is regulating or granting the agreement).

Traditionally, all this forms of trust have appeared in the economic field:

1. when two individuals decide to join efforts and capital in any economic business, they may each trust the other's desire to make the business run profitably;

2. when we buy in a shop an object - something of which we cannot immediately ascertain the quality by just looking at it - we may do it because we already bought here succesfully, or because other people did, or because we assume it is in the interest of the dealer to keep customers satisfied;

3. paying taxes is almost everywhere a *duty*, not a choice, but a citizen may find good to pay his taxes, and more explicitly could agree on a law that imposes a new tax, trusting in the fact that it will allow the State to supply a new/better service, of which the citizen itself will possibly benefit;

4. the concept is absolutely valid even outside from the "mainstream" economic models of gain/utility: I may want to gratuitously finance a particular nonprofit organization, trusting that it will use my money to apport some benefit to the society.

Notice that *qualities* of the counterpart we may base our trust on are often *moral* qualities - my money will not be stolen - but also that this is not the only truth - from someone who will manage my money, I also require specific skills, and my requirements depend upon the conditions my counterpart operates in. However, all the setting of this study will be based - not only for simplicity, but for adeherence to the reality we want to study - on a "yes/no" evaluation of the counterpart: not on the skills that allow it to perform better/worse, but on those that allow it to perform at least to a given level, which we (explicitly or implicitly) agree on.

1.2 Environment and motivations

What is common to all the examples given so far is that there is no *trust* without some kind of *personal knowledge* about the counterpart and the environment it operates in. While this may seem - and has been considered for a long time - perfectly natural, there is today a specific field where this is considered by many an unbearable limitation: *financial markets*. Those environments have traditionally gathered *experts* - of finance, or just of the specific kind of good/services they deal in - but more recently they became available to virtually anyone, thanks to developments of remote and telematic exchanges. This has led to the spread of the belief that size of market, and hence economic growth, can be enhanced by allowing everybody, even the so called *unsophisticated investors*, to invest its capital, in particular *savings*, in the financial market.

The aim of this presentation is not to establish if this belief is right or wrong. Suffice to observe that:

- it is widespread between policy makers, who usually foresee both the potential benefit of market in having more capital employed and of citizens having the possibility to receive part of market benefits,

- it is quite common among the numerous unsophisticated investors themselves - they wouldn't participate at all in the market if they didn't see in it an occasion to obtain some gain,

- it has good reasons to be in some sense the pipe dream of theoretic economists, since treating every member of the society as someone who freely participates in the market, with no constraints - practical or psychologic - would simplify enormously its economical analysis.

Hence, I'll assume our aim is to effectively enable any citizen to participate in the financial market, leaving aside the possibility that this may be a bad idea *per se*.

1.3 Knowledge and trust in what?

An observation is important: financial knowledge and know-how is *complex* in the very precise sense that it is *not linear*: knowing *a part* of it will not necessarily allow a player to gain a *correspondent part* of the benefits she would gain if she knew it all. This happens thanks to the high number of components and agents linked together, but even more because it is a *competition* where the field of action of strong players is basically not limited, and a kind of competition in which an opponent that is slightly better than me can possibly take all the benefits, leaving me with none. Statistic studies ([5]) show that a set of random stock investments has a good probability to reveal fruitful, at a level at least comparable with the aggregate stock index[1], though it is evident that the vast majority of investors base their choices on informations and reasonings (or even just instinct) - in other words, apparently the situation is not only *complicated*, but even easily *misleading*, because of wrong informations or wrong usage of them (behavioural economics are a field of studies that has evidenced many common distortionary effects, one of the most notable being the persistence of beliefs [6]).

2 A general view of trust administration

Before we start with a round-up of trust administration frameworks, let's state the requirements they are asked to fulfill:

1. the system should obviously allow (also unsophisticated) users to unbiasedly establish trust in trustful debt products and debt issuers, without requiring particular skills,

2. the system should not have any economic entry barrier for new users (not because of the recognition of some absolute individual right to trade, but simply because, as already stated, we are starting from the hypothesis that the more individuals enter the trading market, the better for everyone),

3. for exactly the same reason, the system should also avoid *psychologic* entry barriers: new users shouldn't be *afraid* of entering the market, and should not think that, for the simple fact of being newcomers, they have no chance of profiting from it. This requirement may seem too subjective, and in fact certainly the definition of "fear" is,

[1] While it may be tempting to think that an aggregate index should *by definition* represent the expected productivity of a random investment, the difference is exactly in the fact that price variations depend majorly on the "popularity" of some product, to which the random picking is immune.

but nevertheless the pragmatic reasons why it must be considered are evident, and the development of behavioural economic studies gives us scientific support in this direction (also, after the late crisis, it would be very ingenuous to consider *fear* as a component of financial market that an analists are free to not consider at all).

2.1 The current situation

Steming from the hypothesis that unsophisticated users have the right to be informed about credit worthiness, *Credit Rating Agencies*, organisms devoted to providing a rating of various types of debt obligations, are born and have grown since more than a century. CRAs are meant to provide public, independent and easily understandeable measurements of risk for investors, by classifying the obligations in different categories based on their risk of default.

However, in occasion of the economic crisis exploded in 2007/2008, many voices raised doubts about the reliability of those agencies, and those doubts revealed to be in the best of cases legitimate.

Obviously, if a CRA gives a wrong measure of a particular company or financial product, it doesn't automatically mean it is *not reliable*: a random/chaotic component is universally recognized in finance affairs, and what nobody can anticipate can't certainly be anticipated by CRAs neither; however, what struck were the evident cases of blatant errors, often motivated by huge conflicts of interests internal to the agencies, which get payed from the same entities whom reliability they must evaluate - entities that in many cases are also their customers for consultancy services. Ironically, the most common "mistake" of CRAs seems to be that same excessive persistence of beliefs so typical (also) among the unsofisticated users they should "defend".

Before continuing the analysis, an important theoretic problem must be raised: are Credit Rating Agencies meant to also evaluate *systemic risk*?

Asking this may seem like quibbling, since the detection by a Credit Rating Agency of some system risk (which, by definition, affects almost homogeneously market participants), would possibly convince it to lower a vast majority of its ratings, with no evident benefit for market players interested in *relative* ratings - is it better to invest in a product than in another?

However, this reasoning doesn't consider the base "limitation" this paper also started from: the *voluntariety* of participation. If in a particular period *all* ratings are below the historical average, one may choose to just *not enter* the financial market (at that time), or to leave it. This obviously raises a much more striking argument: deontological ethics are a very delicated matter for an entity to which objectivity is asked, but which choices may easily influence the matter of study itself, in a potentially catastrofic way (if a well-known agency lowered all ratings, this would easily have a bad impact on performance of the market, possibly spreading panic).[2]

[2] It is obviously ridiculous to assume that important CRAs think about those problems: far beyond just indulging with market "spontaneous deviations" such as interest misalignment, they have abused with creativity of their powers, as the case of Moody's blackmailing the German insurer Hannover Re, gratuitously rating its products and then finally downgrading them to the lowest rating, causing a loss of $175

One observation spares us the effort required to answer those doubts: at the moment, evaluation of systemic risk and of its possible influence on the trust which can be given to a certain product is technically far more difficult than just considering the characteristics of the product and of the company behind it: in other words, apparently *choosing to insert or not systemic risk among the evaluation parameters* can change only marginally the analysis of a product reliability, for technical reasons.

Anyway, we have an even stronger reason for not considering systemic risk in the main part of this exposition: while we will consider several possible alternative "trust administration frameworks", and the differences between them, none will have a significantly different predisposition to avoiding, or protecting the market from, systemic risk, with a possible exception which details will be given in the end of this paper. Again, this derives simply from the definition of "systemic risk": it affects market participants almost homogeneously, so it can't be tackled on an individual scale.

2.1.1 About CRAs competition

When in 2008 the economic crisis touched its most explosive period, one of the most striking facts seemed to be that the Rating Agencies market is at 90-95% owned by 3 participants:

1. Moody's,
2. Standard & Poors,
3. Fitch.

This may seem a major problem, because of the apparent lack of competition (the hypothesis that those few players own all the market because they formulate the most accurate ratings can be just be ruled out, *a posteriori* - more touching is the observation that, since most rated companies ask at least a couple of different ratings, there is not even much competition between those three, of which in fact the first two own 80% of the market, as underlined also by Sean Egans, Managing Director of the competitor Egan-Jones [3]). However, though it is probably reasonable that such a situation is per se *pathologic*, what is important to focus on is if the small number of (considerable) participants is not just a *symptom* of the "malady": in other words, couldn't this situation be the *natural result* of treating trust as if it was an ordinary good, more than the *cause* for the rating inaccuracies?

Two main causes may have brought this oligopolistic situation:

- the base prerequisite for a company that *sells trust* - which is, in the end, what Rating Agencies do - is to *be trusted*: trust in something is indeed a feeling that can change in time, but usually at a quite low rate (it is the principal subject of the persistency of beliefs we already mentioned), so a market in which bad agencies can be rapidly recognized and loose their clients is nothing more than a dream: low trust in an agency will only be established - by unsophisticated individuals - *after* those individuals spot

million dollars in a matter of hours, shows.[2]

some remarkable mistake or bias in the ratings (and even at this point, the asymmetry of information could not allow them to gather enough findings for an objective judgement, in particular considering that trusting in big Rating Agencies is a trend which implies a action-reaction circle - vicious or virtuous, it depends), and the direct clients - who are the ones who should be "graded" - will certainly not drop their too generous agencies spontaneously.

Notice that exposing publicly the methods used by those agencies to establish their ratings *could* in principle give an objective and immediate way to evaluate their accuracy... but certainly not to unsophisticated users, so this would just move power from the hands of the agencies to the hands of third parties in charge of judging - formally or informally - their methods: there is no reason why this new market of *"raters of raters"* - which would still deal in trust - should be free from the problems we are facing;[3]

- there is indeed a perverse mechanism in the fact that at the same time CRA ratings are used by private individuals who would like raters to be rather severe - to help them distinguish good deals from risky ones - and are part of regulations of the SEC and of the guidelines of Basel II accord (and are taken in consideration by other regulating institutions): final recipients of those regulations are tipically companies whose personal, vice versa, would prefer *relaxed* ratings, enabling them to freely choose to buy products (possibly risky - but they are investing someone else's money) with higher returns and still qualifying as part of their bank's net capital reserve (or allowing bond issuers to use a shortened prospectus, or to grant the strenght of an insurance company's reserves...).

  Ironically, of the two sets of individuals that are affected by CRA ratings, the most sophisticated ones are certainly mostly comprised in the latter, so it is really not so clear if an imprecise rating given by an agency will harm on the long term its "popularity" - here, the word "trust" would be at least ambiguous - and hence its attractiveness on debt issuers.

So, to resume, the theorethical framework underlying the current situation shows us a virtuous circle in which companies are interested in ratings of "fresh" and reliable rating agencies because they know that the public will give those ratings more importance in planning investments, and this implies rating agencies have interest in providing accurate results, and new rating agencies enter the market every time there is space for good agencies... but a high rate of leadership turnover implies a high speed of final recipients "trust updates", an hypothesis as far as possible from the "unsophisticatedness" - this is why, though for instance the difficulty of becoming a U.S. "Nationally Recognized Statistical Rating Organization" (a title that is required in several SEC policies referring to CRAs and that is currently recognized

[3] Though it may seem surprising, and an unneeded homage to fanatic competition theorists, that the USA Credit Rating Agencies Reform Act of 2006 explicitly prohibits the Securities and Exchanges Commission from regulating "methodologies by which any nationally recognized statistical rating organization determines credit ratings"!

to only 10 agencies) is well known, this entry barrier can hardly be seen as the real root of the lack, or bad functioning, of competition (in the end, *there are* 10 agencies, but only 3 of them have a dominant - and stable in time - position).[4]

Before leaving this argument, it may be interesting to mention a situation very different from the USA one, in a measure that could potentially reverse the reasonings just made: after Basel II, every regulator is invited to draw up a list of so called *External Credit Assessment Institutions*; following the guidelines provided on the 18th of January 2006 by the Committee of European Banking Supervisors, Bank of Italy currently provides a list of recognized ECAIs: this list is *very* limited[5], comprising only the three major US entities already mentioned, together with *Lince*, the only Italian one. In this case, it may be argued that the entry barrier *has indeed* an important distortionary effect; while it is probably true, I believe that a vast majority of the analyses and comparisons that will follow are still valid, even if the phenomenon of market closeness certainly influences the whole picture.

2.2 Are ratings necessary?

Theoretically speaking, one could assert there is no need for ratings, even in a market populated by unsophisticated users: the protection from frauds or just high risks can be obtained with the already available financial derivates and insurances. Premia that insurances and CDS issuers ask correspond, in a perfect competition market (which the derivatives one is certainly far more than the rating agencies') exactly to the price of gathering information about, and finally estimating reliably, the risk of default of a given product or debt issuer, plus the price of the risk taken itself. I do recognize some possible objections to this reasoning:

1. companies will try to act so that insurance premium on an issued product is kept low, in order to induce investors to buy it; to do so, they may keep informations that would give a bad image of the product secret;

2. insurance companies, or derivates issuers, should themselves be *trusted* as able to repay the possible losses;

3. products which are traded over-the-counter don't offer a "price menu", which is an essential feature for unsophisticated users to be interested in; certainly personal dealing of the premium is a prerogative of professional investors.

Though those are important points, the first can be answered by observing that secrecy of uncomfortable informations is *anyway* a problem, that is addressed, and must be, by regulators; on the contrary, institutional investors, such as insurances, certainly have the

[4] There is one point on which probably the bare existence of entry barriers does have a huge influence: the level of fees that rating agencies ask to their clients; it is natural to think that monopoly prices - Moody's had an astonishing operating margin of 54% in 1995 ([1]), and more in general the battle of other CRAs to become recognized NRSRO is a useful indicator of the advantages of the status - are a direct result of limited competition. However, the relation with the accuracy of ratings and the availability to final recipients is, as shown, at best much less direct.

[5] See http://www.bancaditalia.it/vigilanza/banche/ecai

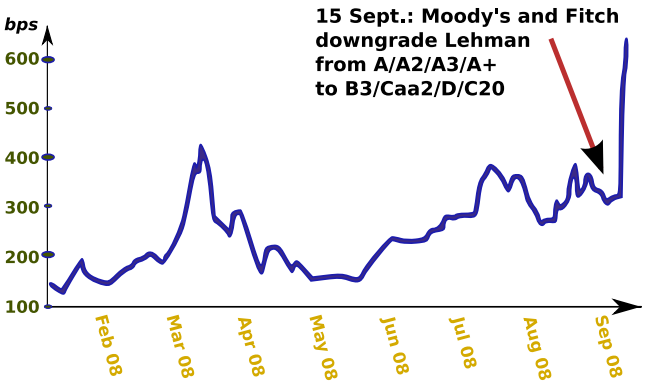

Figure 1: Lehman CDS spread in the months preceding its bankrupt, compared to Rating Agencies reactions.

possibility, more than small investors, to ascertain the perspectives of a product. About the second, as already said we are not considering *systemic* risk, and should not; on the other hand, an insuring company should be very resistent to particular market shocks, and ensuring this is true is another precise task of regulators.

The third point may be, still from the theoretical point of view, the most important: obliging insurances to be exchanged in official markets would represent a big and difficult regulatory intervention, and could greatly affect the market itself, since certainly secrecy is a key aspect of this market (the value of a case study is certainly much higher if its result is kept secret and revealed only to effective customers); however, notice that regulation of the derivates market is something that is already considered as a necessity by many regulators, and though it would certainly corrode profits of the specialists of the sector, it is exactly the price to pay to allow unsophisticated users to enter the game.

That said, it is quite probable that this model could simply never work in reality, simply because of mainly psychologic reasons: while typical small investors may already not be interested in sharing part of their profits with insurers, the real problem comes when we think to institutional investors and regulations such as SEC and Basel II recommandations, which view ratings as an official and constitutional part of market regulations, also between big investors. Requiring that all those institutional investors either respect insurances regulations or rely on external entities which respect them - and presumably impose high fees for their services - would bring a major upwheaval in the whole market, not only on a formal and theoretical point of view, but it would effectively disturb the freeness and economicity of the whole mechanism to a point probably not bearable, as well as distort the premia themselves (a company could choose to offer - possibly through a third party - CDS for its own bonds with a low prize, in order to fake diffused high trust in them).

2.3 Subscription-based rating agencies

"A critical distinction between Egan-Jones and its larger competitors in the credit rating industry is that its revenues are derived from the institutional investors who subscribe to its services, i.e., the business model which Moody's, S&P and Fitch followed during the era when they still enjoyed reputational capital." - [3]

The idea that those entities who are interested in ratings, and not the rated ones, should be *the direct customers* of rating agencies seems at first sight perfectly natural, for the very obvious reason that only in this case transparency would be an automatic result of the need to satisfy customers, by providing the most accurate ratings: in other terms, there would be an almost perfect alignment of interests between those making ratings and those using them.

Not by chance, the above is an excerpt of a testimony to the U.S. *House Committee on Oversight and Government Reform* by Sean Egan, Managing Director of the *Egan-Jones Rating Company*. This rating agency still wasn't mentioned in this paper because, though being a Nationally Recognized Statistical Rating Organization since December 21, 2007,[6] it seems not particularly important if we look at its size and its share of the market: what instead does make it unique, among the 10 NRSROs, is that it is based on a subscription model, where ratings informations are not made public, but instead delivered only to subscribers. This model of agency has many (mainly small) imitators around the world: in Italy, an example is Capp&CAPP (though its demand for EACI recognition from Bank of Italy was still not granted[7]).

The model is probably a winning one from the individual perspective, since it is the best an institutional investor can hope for; however, the limit of the approach is stated on the Egan-Jones website itself:

"Note: Egan-Jones Ratings products and services are exclusively for institutional investors. Regrettably, we have no products for individual investors."[8]

This is really regrettable if, as stated by many, such as the CFA Institute, this agency has been "often beating rivals S&P, Moody's, and Fitch in marking changes to perceived credit quality" ([4]).

But indeed, the subscription rating agency is necessarily a niche solution, not targeted at solving the needs of the large number of individual unsophisticated investors. Moreover, the fact that all ratings are private makes it difficult also to embed the solution in regulatory policies for institutional investors: Egan-Jones *is* recognized by the SEC, but as an alternative to traditional - often biased - rating agencies; even if it was established that only this model of rating agency was reliable, forcing all banks currently basing their solidity parameters on

[6] See http://www.egan-jones.com/publicdocs/Egan%20Jones%20Approval%20Order.pdf

[7] The usage, by Capp&CAPP, of the domain ecai.it for autopromotion is an example of cybersquatting in its purest form.

[8] From http://www.egan-jones.com/contact.aspx

freely available ratings to pay subscription fees would be a remarkable interference in the market; furthermore, a specific problem of non public ratings is that competition between the different agencies would be *severly hindered*, unless customers tend to subscribe to more than one, so they can compare; but currently, the low share of market those agencies own suggests the benefit/price ratio that potential customers see in them is too low for one subscription, let alone for several.

Finally, notice that large CRAs often receive informations from the monitored companies through private channels: while this is perfectly legal, SEC regulations established, until few days ago, that a CRA could only use such information if its ratings were made available to the public for free [9]; though this may seem an unneeded limitation binding, if we enlarge our view to the entire picture it is evident that this rule, which was abolished on September the 18th [10], was based on the rationale that officially enabling a secret - and certainly profitable - market of informations largely opens it to insider trading issues (which is what the first SEC directive we refer to, as well as similar regulations in other countries [7], aimed to address), and more generally risks, as any incentive toward obscurity, to tackle (even further) the reliability of "trust management".

On the whole, we can conclude that dealing in insurances on some products or on ratings about them is, though the distribution of risk and responsabilities completely changes, theoretically (which means, provided that there is no conflict of interest) equivalent on an aggregate basis: an agency gets a premium for its forecasting abilities, and has interest in keeping those abilities, and the yield of their usage, secret - which is exactly what regulators should avoid if they want unsophisticated users to take part in the market.

3 A new framework: the "debtrank" model.

3.1 Google and the Pagerank algorythm

At the beginning of the history of the World Wide Web, there existed a complete list of all sites, edited and kept up-to-date by Tim Berners-Lee, the creator of the WWW.[9] Through the nineties, however, the World Wide Web has grown enourmously, from the 130 web sites of 1993 to the almost 650.000 of 1997,[10] and far beyond. Hence, the necessity of specific tools to seek for some content arose, and the first web search engines were born: JumpStation, in 1993, and then in 1994 WebCrawler and the famous Lycos. Those search engines allowed the user to insert a query, containing some search terms, and to get back in return a list of pages containing them. Depending on the engines, the *affinity* of the page was established by looking at the frequency and position of words in it, or based on indexes manually created in precedence. Both methods were highly inefficient: the former because the criterion used to rank pages didn't take in consideration the *authority* or importance of a website, and considered only parameters, such as the frequency of words, which are a very vague (and

[9] Such a list, in a version back from 1992 listing all the 29 websites existing at the time, can be still accessed at http://www.w3.org/History/19921103-hypertext/hypertext/DataSources/WWW/Servers.html

[10] Estimations taken from http://www.mit.edu/people/mkgray/net/web-growth-summary.html

very easy to cheat, by providing specially crafted pages) proxy of affinity, the latter because of the huge amount of human work needed to create and keep updated the indexes.

Around 2000, Google began having a huge success because of a revolutionary algorythm used to automatically establish the importance of pages: the *Pagerank*[11]. It reflects the concept of "random surfer": it tries to establish the probability that a user of the web, who spends his existence clicking at random links of web pages, is on a given page at a given time. Practically speaking, if we call Q_p the set of all pages with links pointing at a given page p and $N(p)$ the number of external weblinks it contains, the Pagerank value (usually called just "pagerank") of p can be calculated as:[12]

$$\mathcal{R}(p) = \sum_{p' \in Q_p} \frac{\mathcal{R}(p')}{N(p')} \quad . \tag{1}$$

The definition is recursive, and hence can't be applied as it is; a feasible strategy to effectively calculate the pagerank of pages consists in starting by assigning a value of 1 to each one of them,[13] and then iterating on them all, by updating the value of each one accordingly.[14]

The result of this algorythm has the following nice properties:

1. every hyperlink pointing to another page works like a *vote* to this page: it is giving some importante to it,
2. the importance of my (page's) "vote" depends on my (page's) importance,
3. however, the more votes (hyperlinks pointing on other pages) I give, the less each vote is valued (so, there is no sense in filling every page of external weblinks only in order to raise importance of the targets).

Moreover, the properties 2 and 3 can be resumed in a more general one:
Given any set of pages P, there is no way to raise the aggregate value of pages in P just by creating, deleting or changing links between them.[15]

In other terms, there is no successful strategy, not even collective (links exchange), to fake importance inside a group: in the end, it can only derive from other sources.

[11] Named after its creator, Larry Page.

[12] For simplicity, we are omitting - or equivalently set to 0 - the so called *damping factor*, which is an important part of the original Pagerank algorythm ([8]), but which, in the hypothesis that the set of relations we are monitoring - hyperlinks, in this case - tends *per se* to stay up to date with real world changes, has no major influence on the *relative* importance to the different entities - pages, in this case - and has mainly a *normalizing* function.

[13] Or any constant, for that matter.

[14] The matricial version of this formula, where the set of all pagerank values is a vector iteratively multiplied by a matrix describing the graph of web links conveniently weighted, solves the problem of order inconcistencies - which page's pagerank shall I start updating? - simplifies notation and optimizes calculations.

[15] With the only exception - which would have anyway to be considered as a particular case - of pages with *no* external links.

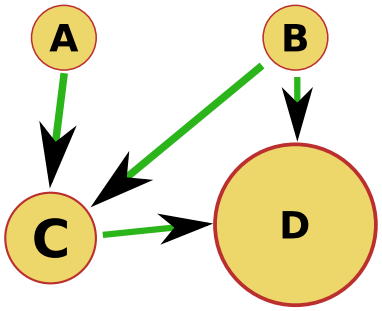

Figure 2: An example of Pagerank in a very small network: C, and D have the same number of incoming "*votes*", but D's ones have more weight, so D has more importance (the exact degree depends on the chosen dumping factor).

3.2 Back to finance

How does this digression apply to the subject of ratings?

Pagerank does nothing more than measuring the *confidence* about finding interesting informations in a given page. It does so by exploiting a very sparse and freely available information, collecting it in a usable form and interpreting every hyperlink as a declaration of *trust*. As its very simple description may suggest, it can also be of very general application: an interesting case in which it is already used is its implementation as a replacement for the traditional measures of scientific journals' *impact factor*[16] - basically any *directed graph* can be studied in this way.

It is hence straightforward to see its possible utility in finance: if every market player was able to declare its trust in a given entity - which can be another market player, an institutional investor, or even a "classic" rating agency - this would form a network of "votes"; then, it would be trivial (the calculations that every update of Google's database implies can't really be defined as "trivial", and they involve a very vast and advanced network of computers; however, they concern more than *1 trillion* pages[17], while the number of market players expressing their trust would, in the best of pipe dreams, reach the order of billions) to give a trust value to every player.

What is really not trivial is to *gather* information about individuals and entities' trust. It is not trivial, but neither too labour needing: all is needed would be a good *framework* for trust registration, where any market player, at any time, could register its "declaration of trust" towards another market player, which could be any type of entity. This framework should

[16] The implementation can be accessed http://www.eigenfactor.org

[17] Information reported on http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html ; there it is said that not all pages take part into the algorithm, but still, in the bottom of the page, one can read *"our distributed infrastructure allows applications to efficiently traverse a link graph with many trillions of connections"*.

certainly be totally virtual, possibly separated in several national agencies, tighted together by a particular protocol to which adherents conform, and it should be entirely managed by institutions, since its management, though very light and somewhat "passive" - it should only gather information, as the operation of implementing the algorithm could be then done by third parties - should guarantee absolute transparency and unbiasedness: in particular, every market player should be able to be represented - and to be represented only once[18] (this is the main reason why in the case of a network, a strict protocol should be implemented to ensure no "cross-site" cheating is possible).

3.2.1 How is finance different from the web

The functioning of a Pagerank-like financial trust administration system would (and *could*) differ under some aspect from the Google system, for several reasons, that should be taken into account:

- Google is meant to work in a hybrid environment, in the sense that a vast majority of the public only *uses* the system (by querying the search engine), a still big group of players *feeds* it (by providing sites, pages, or just content) and members of a small niche actively *study* it, try to understand what weaknesses of the system can be exploited, and in some way really *cheat* it. Instead, in the case of the financial market trust, the first two groups - those who *want to know* the trust attributed to some entity and those who *give* trust to other entities - would almost coincide: this would make the third group disappear, because what is supported by the mass of players is no more a cheat - it just influences the meaning that the mass will give to values - while what is supported by a minority will be, for the dynamics of the system (where an overevaluation will spontaneously catch attention), automatically resized by the mass action.

- The very important fact that *every evaluation will be voluntary* will give the possibility for a much more precise granularity of votes: mainly, instead of limiting the options to "vote/no vote", it will be possible - and highly recommended - to introduce *numeric* votes - for instance, -1, 0 and 1, or the range of integer numbers from -5 to +5, or even any decimal number an user may want to choose inside a given range. Obviously, the relation 1 should be modified in order to normalize votes in this extended version[19]

[18] In fact, if some big players were able to treach and get multiple "vote accounts", the benefit for them would be minimal, as new accounts would hold a single vote each, which probably would do almost no difference. What should be avoided, and what policy makers should battle against, would be only the possibility of mass account creation, or of mass istigation of small market players to double-account creation, which would anyway be macroscopic and hardly kept hidden actions.

[19] The most immediate, and probably most effective, way to do that would be to normalize votes so that sum of positive ones is 1 and sum of negatives is -1; then, to add another normalization constant that brings a contraction on negative votes, so that the overall sum is 1 (instead than 0); the special case in which a user gives only negative votes should be however treated differently - for instance, assume a node of the network always links itself. Notice that though it makes sense to face those technical problems, it doesn't mean public entities or in general regulators should bother about them: once the data (votes of users) would be public, there could be a spontaneous and free competition in finding the adaptation of Pagerank

- While behind Pagerank there is the idea that existing pages are almost static items,[20], and as time passes they basically only grow in number, the votes expressed by users could have a *date* attached, and they could be taken in less and less consideration as they age, or even expire after a given number of years, to reflect the fact that information is no more valid. Technically, this could be done on a vote basis and also on a user basis, where a market player that has entered no information in many months could have the weight of his votes diminished accordingly.[21]

3.2.2 Why should it work?

Trust is power, and power is money: there is no doubt about it, as many studies on reputational capital have shown, and as the high profits of Credit Rating Agencies also demonstrate.

In the case we are facing, if the described system of trust evaluation became vast and rich of information, the value, for any company, of having a high trust indicator would be considerable; hence we can imagine that also new players coming into the market would be kindly invited to show their support to market players they interact with, by signaling their trust in them.[22]

While this phenomenon would certainly be an engine of further growth for the system, an apparently distortionary effect can be easily forecasted: companies would easily put in place exchanges with small market players (which could take the form of an "offer" or even blackmail) of the form "register you trust in me and I'll offer you a special deal"; probably, many would accept such deals, and the trust value of the company would result as being inflated.

Indeed, this shows that the system of trust would *not* be the panacea to perfectly describe, with a single number, a random anonymous debt issuer in the world, since this is just *impossible*; and also suggests that the use we are making of the word "trust" has a significantly different and more abstract meaning than in common sense. But that said, consider the case of a user that receives such an offer: a particular debt offer which implies, among the conditions, that he must explicitly register (possibly faking it) his "trust" in the issuer. The user has two choices:

1. disgustingly refuse the offer (because of moral convictions, or just because he would require a higher price for him to disobey those moral convintions): at this point, the

which better suits the needs - possibly, different adaptations for different share of public: regulators may then limit themselves to endorsing a few adaptations well behaving for the general public, to which they may refer when the values given by this system are part of policies.

[20] Not in the sense that they cannot change, but that *"content is forever"*: what is written inside will be of interest forever, unless it is overridden by content providers - notice anyway that obsolete hyperlinks pointing to abandoned sites are a major incentive for the domains cybersquatting phenomenon.

[21] Technically, this means adding multiplier that decreases in time to all votes *to* the user considered, and simply diminishing accordingly also the constant that is taken as the starting value of any market player.

[22] Not very differently from what already happens in several online websites that host or manage exchanges or sales - such as the *feedback* mechanism in Ebay.

company will have lost a customer, and this customer will presumably give a *negative* vote to it;

2. accept the offer: in the end, by accepting, the user is still "trusting" the company, and hence the aggregate message passing, "trust this company", is *truthful*. There is really no *moral* meaning in the message, but just a transparent description of reality, which is exactly what the system aims to.

Having understood that the system will auto-incentivate itself, the problem is how to start the "chain reaction": certainly every big player would have interest to jump in a system of values on which policies are based; but certainly regulators *can't base policies* on a system in which big players still don't take part: there seems to be a chicken-or-the-egg problem. But luckily, there is a solution: the moment in which regulators would setup the framework could come much before the official entrusting of the system in policies, in particular considering that:

- the system itself is very cheap, and is anyway, even if no policy makes use of its indicators, a great gift to unsophisticated market players, that get a freely consultable tool (that they can optionally use *in parallel* with "classical" indicators, such as existing rating agencies' ratings),
- the system wouldn't necessarily be implemented worldwide at the same time: though doubling the number of independent systems would double the effort requested to players and lower the precision of each single trust network (but still, even small network could prove to be useful), single countries, or even single marketplaces, could already build a unified framework, that each new participant adopting the system could join at any time, until it becomes a single worldwide system;
- the novelty of the system would raise the interest of small but talented rating agencies or consulting companies; for them, it would be an occasion to get some publicity, in a market dominated by an oligopoly.

The last point may seem contradictory: in a system that is passed off as an *alternative* to Rating Agencies, *what is the place* for them?!

Rating agencies would simply be nodes of the network: the good ones, in which people trust more (even unsophisticated market player, which as time passes by would have a way to see the "percentage of guesses and of wrongs"), would be focal nodes, with very high trust values, and hence the power to distribute high values to other market players they entrust. They *could* be payed by those whom they rate, and even sell consultancies to them, but if their ratings were biased because of the conflict of interest with respect to their customers, they would loose the trust of users, and this would *immediately* tackle their power and prestige. It is hard to imagine that an untrustful company may pay a good rating agency for a rating, with the risk of getting a "bad grade" (and having payed it!); so, in the end, good rating agencies would give mainly good grades (this model will probably result in

something more similar to an unofficial *certification* than to a rating), but bad companies would completely loose the possibility to raise their trust rankings, and hence would still have an incentive to behave well. On the other hand, bad rating agencies, that nobody trusts, would simply have *no place* in this setup - and recall that "trusting" a rating agency would mean no more "being obliged by some regulation to buy debt of a given ranking", or "having the need to buy, just follow the flow and believe whom most of the market (seems to) believe"; it would mean "having a reason to explicitly express my trust".

4 Few words about rating and systemic risk

The hypothesis that influenced all this work was that the goal of rating agencies is "just" to help newcomers to enter the market, and that this is *good*.

But I evidenced another even more important limitation to the exposed arguments: the fact that none of the strategies reviewed and proposed could, at least directly, do much against systemic risk.

Now that the argument of trust has been dissected, it's time to try to make one step further, looking at the relation between it and systemic risk. The existence of a correlation between periods of low confidence and systemic crisis is clear and easily understandable, because the causal effect between the latter and the former is obvious; however what is interesting to study is the particularity of the reverse effect.

At first sight, a myopic approach to the problem, in the setup of the network of trust, might excess in optimism: one may think that, since values given by the pagerank are in the end totally relative - they have an ordinal, but absolutely not a cardinal meaning - they could be continuously normalized so that their average, or possibly their aggregate sum, or some particularly nice function of its distribution, is constant: market players would still see differences between poossible counterparts, but there would be no aggregate signal of "failure". What is very misleading in this way of thinking is that the whole system has no chances to work if it's not dominated by *total transparency*, including the publicity of all trust expressions, and once trust expressions are public, everyone can derive from them any desired estimate, including measures that would show, for instance, a declining trend in aggregate trust. That doesn't mean that particularly adapted measures couldn't be integrated in policies - for instance, to partially absorbe the effect of a crisis - but only that they can't be used to *hide* reality.

Systemic risk is a very complex argument, but its complexity can be easily traced to two (complex, in turn) main causes:

1. risk of unlucky external events (where for instance a raise in the price of fuel is considered as "external") that affect a vast share of the market players, directly or indirectly,
2. financial bubbles, which, when "exploding", have the perverse effect of *reversing* the

mispricings, putting market prices *below intrinsic values*,[23]

and obviously every possible combination of the two phenomena.

While trust between market players has hardly anything to do with the first point, it is an important component of the second phenomenon, since the underpricing, for instance, of shares is a consequence of the fall of trust in the issuer (and this can easily result in a vicious circle). More precisely, the word "consequence" most of the time comprises three effects - that, though hardly distinguishable in reality, where they sum up, we can theoretically consider as separate, because they have different backgrounds and final effects:

1. some individuals derive, from a set of new informations they receive, the conviction that some product was overvalued with respect to its intrinsec value, and hence they sell/lower the price for it;

2. others (not necessarily in possession of the same informations) feel that if the above individuals have sold/underpriced some product, it must be a sign that it was overvalued, and hence that the intrinsec value is lower: they will sell/lower the price in a similar way;

3. others may just see that, since the trust in the product, and hence the price of it, is going down, it is very probable that it will be, at least for a certain amount of time, *undervalued*, or at least that it will stay at a price lower than the usual: they will be the ones which will feel more the urgence to sell, even at a price (possibly known to be) under the intrinsec value.

Notice that, though the last "effect" is what was often described by the media, in troubled times, as "collective fear", this denotation may be misleading, since it is *not* (necessarily) an *irrational* behaviour: just "following the flow" may be motivated with very convincing mathematical arguments[24] (besides pure instinct), which simply *can not* be tackled on an individual base, because they are a result of the complexity of the system. Also notice that exactly the same flow of thinking is other times described as "speculation", when the "actor" is tipically a big, and presumably not "instintively afraid", institutional investor.

As already said, those effects cannot be entangled in reality and, as readers may guess, the word "others" used in the description of the latter two is a pure astraction, since *most* of the market participants will mix the three in their decisions making.

Now, we already observed that trust between individuals and entities, when motivated by *knowledge* and not just by possibly long a chain of trust relations, is a much more *stable* relation than what the simple *pricing* might be. But most importantly, *separating* the "market for trust" from the usual market can almost break the vicious circle (it would not avoid

[23] Which obviously doesn't mean that underpricing is *evident* during the burst, because if the intrinsec value could be easily and objectively established, bubbles would probably simply not occur.

[24] Tipically, some simple autocorrelations analyses of prices time series suffice.

bubble bursts - the first 2 effects listed above - but would be effective against the resulting undervaluation - the last element of the list): the same agents that are *afraid* that a share they own will loose price, and hence want to sell it, won't feel the same fear of declaring their trust in the issuer - they are, in principle, afraid that if they trust declarations are badly distributed, this can reduce other people's trust in themselves (and, as we have seen, trust is power), but evaluation of someone else's trust evaluation can be, and certainly *will be*, made on the long run.

In other words: *reputation capital* is by definition unattackable by speculation, so the market would be greatly helped if its value was clearly stated.

References

[1] Rating the rating agencies. *The Economist*, July 15, 1995. http://www.economist.com/finance/displaystory.cfm?story_id=9267952.

[2] Credit Raters' Power Leads to Abuses, Some Borrowers Say. *The Washington Post*, May 24, 2004. http://www.washingtonpost.com/wp-dyn/articles/A8032-2004Nov23.html.

[3] Testimony of Sean J. Egan, Managing Director Egan-Jones Rating Co., before the House Committee on Oversight and Government Reform, October 22, 2008. http://oversight.house.gov/documents/20081022102906.pdf.

[4] Jonathan Barnes. Solving the Credit Crisis. May-June 2009.

[5] L. Fisher. Outcomes for "Random" Investments in Common Stocks Listed on the New York Stock Exchange. *Journal of Business*, pages 149–161, 1965.

[6] Kwangmin J. Ko and Oliver Hansch. Persistence of Beliefs in an Investment Experiment. *SSRN eLibrary*, 2008.

[7] The Technical Committee of the International Organization of Securities Commissions. Report on the activaties of Credit Rating Agencies, September 2003. http://www.iosco.org/library/pubdocs/pdf/IOSCOPD153.pdf.

[8] Lawrence Page. Method for node ranking in a linked database, Jan 9, 1998. http://www.google.com/patents?id=cJUIAAAAEBAJ.

[9] Securities and Exchange Commission. Selective Disclosure and Insider Trading, October 2000. http://www.sec.gov/rules/final/33-7881.htm.

[10] Securities and Exchange Commission. Disclosure of Certain Money Market Fund Portfolio Holdings, September 18, 2009. http://www.sec.gov/rules/final/2009/ic-28903.pdf.