Information processing in complex networks

Quax, R.

Publication date
2013

Chapter 3

Information dissipation as an early-warning signal for the Lehman Brothers collapse in financial time series

Material from this chapter is submitted as a Research Article to Nature's Scientific Reports.

3.1 Introduction

A system consisting of coupled units can self-organize into a critical transition if a majority of the units suddenly and synchronously change state (143–145). For example, in sociology, the actions of a few can induce a collective tipping point of behavior of the larger society (76, 146–152). Epileptic seizures are characterized by the onset of synchronous activity of a large neuronal network (153–159). In financial markets the participants slowly build up an ever densifying web of mutual dependencies through investments and transactions to hedge risks, which can create unstable 'bubbles' (23, 160–163). Detecting the onset of critical transitions in these complex dynamical systems is difficult because we lack the mechanistic insight to create models with predictive power (72, 164–166).

A characteristic of self-organized critical transitions is that the network of interactions among the units leads to long-range correlations in the system, or in other words, every unit 'feels' the state of every other unit to some extent.

Here we measure this self-organized correlation in terms of the transmission of information among units. Shannon's information theory quantifies the number of bits that is needed to determine the state of a unit (i.e. Shannon entropy), as well as the fraction of these bits that is contributed by the state of any other unit (mutual information) (84). We use the information dissipation length (IDL) as a measure of the characteristic distance of the decay of mutual information in the system. As such it can be used to detect the onset of long-range correlations in the system that precede critical transitions.

We apply the IDL indicator to unique time series of interbank risk trading in the USD currency and find strong evidence that it indeed detects the onset of instability of the market several months before the Lehman Brothers bankruptcy. In contrast, we find that the critical slowing down indicator and other well-known early warning signals do not provide a clear warning. We repeated the same analyses on data in the EUR currency and find similar results, see Section B2 for details. Our results suggest that the Lehman Brothers bankruptcy was a self-organized critical transition and that the IDL could have served as an early-warning signal.

3.1.1 How information dissipation can lead to critical transitions

As a system's unit influences the state of another unit it transfers information (84) about its own state to the other unit (85, 86, 93, 94, 135, 167). For instance, each particle in an isolated gas 'knows' something about the momenta of neighboring particles due to the transfer of momentum during collisions. That is, the momentum of a particle is the result of its recent collisions with other particles. This information is in turn transferred to other particles in subsequent collisions, and so on. At each interaction the information is only partially transferred due to stochasticity and ambiguity (93–96), so information about the state of one particle can only reach a certain distance (IDL) before it is lost.

The IDL measures to what extent the state of one unit influences the states of other units. As the state of one unit depends on another unit, a fraction of the bits of information that determine its state becomes a reflection of the other unit's state. This creates a certain amount of mutual information among them. A unit can then influence other units in turn, propagating these 'transmitted' bits further into the network. This generates a decaying amount of mutual information between distant units that eventually settles at a constant. The higher the IDL of a system, the larger the distance over which a unit can influence other units, and the better the units are capable of a collective transition to a different state. Because of this we can measure the IDL of systems of coupled units and detect their propensity to a catastrophic change, even in the absence of a predictive model. See Sections B1 and B2 for a more detailed explanation and how it differs from existing indicators.

We measure the IDL of risk-trading among banks by calculating the IDL of the prices of interest-rate swaps (IRS) across maturities. The rationale is that the dependencies between banks are expected to be reflected in the dependencies of swap rates across maturities, as we explain next. Each financial institute is typically exposed to a significant amount of risk of changes in short-term and long-term interest rates, and buys corresponding IRSs to cancel out or 'hedge' these risks. If an institute has difficulties in financing its short-term interest rate hedges and consequently has a higher chance of default, then each long-term IRS that it holds becomes less valuable (and vice versa). The corresponding buyers of these long-term (short-term) IRSs must buy additional long-term (short-term) IRSs on the market to compensate, increasing the demand. An increased dependence between institutes can therefore lead to an increased dependence of the prices of IRSs of different maturities. A significant increase of this approximated IDL may indicate the onset of a critical event. We consider it to be a warning if a threshold of three times the long-term standard deviation is exceeded, which is a common criterion (168–170).

The IDL at time t is calculated as follows. The swap prices form a one-dimensional system because, for instance, a 3-year IRS logically consists of a 2-year IRS and a prediction of the value of a 1-year IRS that starts two years in prospect. That is, that the price of the i^{th} maturity depends on the price of a maturity $i-1$ and a (stochastic) prediction component. We therefore assume that the stochastic interaction between the IRS prices of maturities i and $i+1$ is equal for all i (171), which leads to an exponential decay of information between maturity 1 and i for increasing i (see Section 1.1.3). We fit the exponential decay $a+b\cdot(f^{(t)})^{1-i}$ to the measured Shannon information $I(r_1^{(t)} \mid r_i^{(t)})$ for $i=1,2,\ldots$, where a is the mutual information that all IRS rates have in common, b is the normalizing factor $I(r_1^{(t)} \mid r_1^{(t)})-a$, and $f^{(t)}$ is the rate of decay of the mutual information between the IRS rates across maturities. We define the IDL as the corresponding halftime $\log^{-1} f^{(t)} \cdot \log 1/2$. The mutual information $I(r_1^{(t)} \mid r_i^{(t)})$ is estimated by constructing a contingency table of the two vectors $r_1^{(t-w)},\ldots,r_1^{(t)}$ and $r_i^{(t-w)},\ldots,r_i^{(t)}$, which are the w most recently observed rates in the market at time t. To construct this table we divide the range of values of each vector into bins of constant size h such that two observed rates are considered equal if they fall into the same bin. Our results are robust against choosing the parameters w and h; see Section B3 for the analysis. The results in Figure 7 were produced with a window of $w=200$ trade days and $h=1/500$ percentage points.

Note that here we define IDL as the time it takes for information to reach a certain *fraction* of its maximum, whereas in the previous Chapter we defined the IDT to reach an arbitrary *constant*. This was done to apply the measure to the particular problem. In the previous Chapter the entropy of the system is time-invariant, and all nodes were identical. In the IRS data, on the other hand, the upper bound and lower bound of the cross-maturity mutual information curves change greatly over time. Possible causes include changing rules of dynamics in the market and a variable correlation with external indices (such as the house-price index). As a result, a constant amount of information has no clear and absolute meaning throughout time.

3.1.2 The data

The market of interest rate swaps (IRS) is the largest financial derivatives market today (172) with more than 504 thousand billion USD notional amounts outstanding, or almost 80% of the total market. The buyer of an IRS pays a fixed premium to the seller, while the seller pays the variable LIBOR interest rate to the buyer. In effect, the seller insures the buyer against unexpected fluctuations in LIBOR in return for the expected net value of the IRS. Swap prices can significantly influence the funding rates of financial institutions and therefore play a key role in the profit-and-loss and risk of financial institutions such as banks, insurance companies and pension funds.

Our data is provided by the ING Bank and consists of the daily prices of IRSs in the USD currency for the maturities of 1, 2, …, 10, 12, 15, 20, 25, and 30 years. The data spans more than twelve years from 04/29/1999 to 06/06/2011. The prices of IRSs are based on LIBOR, the average interbank interest rate at which banks lend money to each other. Our data correspond to IRSs with yearly fixed payments in exchange of quarterly variable payments because these swaps are the most liquidly traded across a large range of maturities.

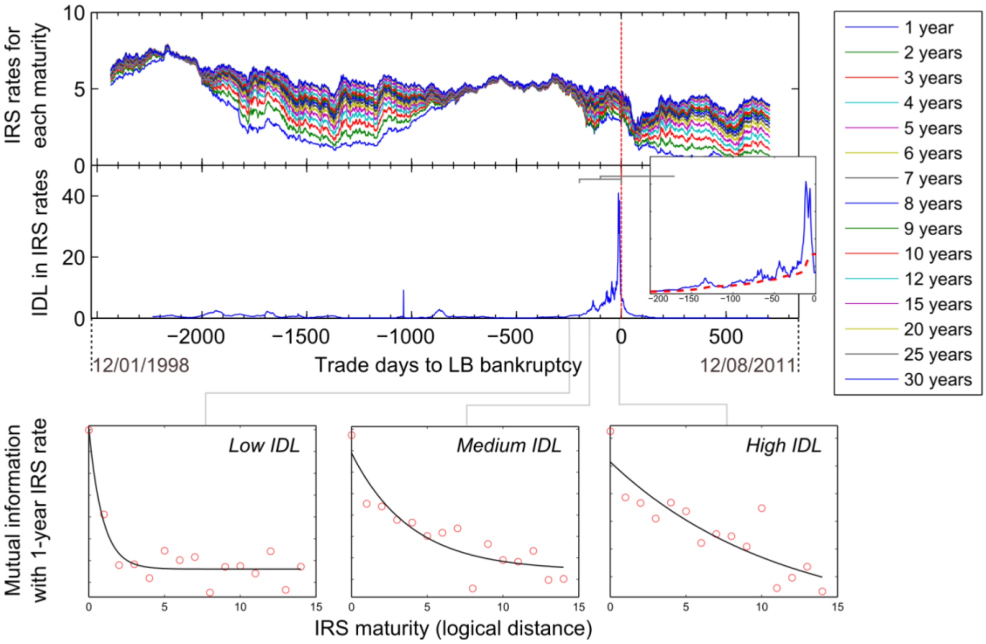

Figure 7: The original time series of the IRS rates for different maturities and the corresponding IDL indicator. Inset: the IDL indicator and a warning threshold during the 200 trade days preceding the LB bankruptcy. We set the warning threshold at three standard deviations above the mean IDL of a sliding window of 400 trade days. Bottom: the mutual information between the rates of the 1-year maturity IRS and all other maturities (red circles) at three different trade days in 2008. The fitted exponential decay is used to estimate the IDL of the IRS rates across maturities for that specific trade day.

3.2 Results and Discussion

3.2.1 Evidence of IDL as an early-warning signal

In Figure 7 we show the original time series of IRS rates with the corresponding values of IDL for the USD market. The day of the Lehman Brothers bankruptcy is preceded by a significant increase of IDL of one order of magnitude, a unique event in twelve years of risk-trading. This is consistent with our hypothesis that a self-organized transition requires that information about the state of a unit can travel a large distance through the system.

We also observe that the bankruptcy is followed by an abrupt drop of IDL and then a slow return to the long-term average. This phenomenon is consistent with interpreting a critical phenomenon as the release of built-up stress (143), similar to the way that an earth quake releases the built-up tension between tectonic plates. These two observations together suggest that the Lehman Brothers bankruptcy was a self-organized critical transition and that the IDL indicator is capable of detecting it. We verify experimentally that the IDL indicator detects correlations between subsequent maturities and is not prone to false alarms by computing the IDL for randomly generated time series with a known period of correlated time series; see Section B5 for details.

We find that the IDL indicator could have served as an early-warning signal for the Lehman Brothers bankruptcy. We define the earliest time at which a warning could be given as the point where the IDL increases beyond three times its long-term standard deviation (see the inset of Figure 7). In the USD market data we find that the earliest warning precedes the bankruptcy by 257 trade days and lasts for 146 days.

Due to the magnitude of the Lehman Brothers bankruptcy and the intimate relationship between the USD and EUR markets, the risk-trading in the EUR currency should have a coinciding peak of IDL. Therefore we repeated the analyses for the EUR risk-trading market over the same period, see Section B2. Indeed we find a strong coincident peak of IDL which could have served as an early-warning signal.

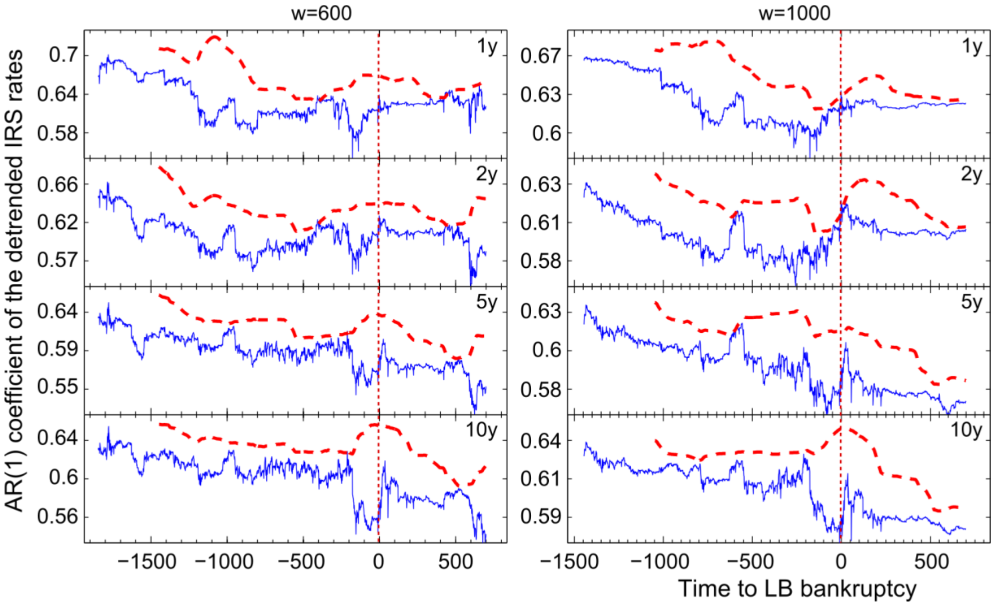

Figure 8: The solid blue line is the coefficient of the first-order autoregression of the detrended time series, which is a measure of critical slowing down. For certain types of critical transitions, this coefficient grows steadily leading up to the transition. The dashed red line is the warning threshold of three standard deviations above the mean of a sliding window of 400 trade days, as in Figure 7. The coefficient is computed of a sliding window of 1000 trade days which is detrended using a Gaussian smoothing kernel with a standard deviation of 5 trade days.

3.2.2 Comparison to critical slowing down and other indicators

The most well-known leading indicator of critical transitions is the increase of the autocorrelation of fluctuations of the system state (144, 173–176). The intuition is that if an unstable system is perturbed it returns more slowly to its natural state compared to a stable system. The more stable the system, the stronger the tendency to return to its natural state, so the more quickly it responds to transient perturbations.

We compute the first-order autoregression coefficient of the fluctuations of each maturity IRS time series for all possible window sizes and show the most representative results in Figure 8; see Sections B5.1 and B5.4 for more details. We find indeed signs of critical slowing down around the Lehman Brothers bankruptcy for certain window sizes. However, it is difficult to find parameter values that provide a sustained advance warning, that is, where the indicator crosses the warning threshold for more than a few days before the bankruptcy.

Another type of generic leading indicators used in the literature are the spatial correlation and spatial variance of the signals of the units of a system (145, 177–182). See Figure 9. In our data, the dimension of maturities can be taken as the 'spatial' dimension. We find however that in our time series they do not show a distinctive change of behavior around the time of the bankruptcy. More traditional indicators used for financial time series are the magnitude or spread of interest rates (183). However, Figure 7 and the bottom panel of Figure 9 show that neither measure provide a clear warning: a high and low-spread period occurred more than a year before the bankruptcy and was returning to normal at the time of the bankruptcy. We also find no warning in case the daily returns (relative differences) are used (see Section B4.2).

Lastly, the same swap with a different variable payment frequency (e.g., monthly, quarterly, semi-annually) were quoted at the same price in the market before 2007. During the recent crisis, a significant price difference across frequencies emerged (184). Although this has a major impact on the valuation and risk management of derivatives, this so-called 'basis' does not provide a clear early warning (see Section B4.3).

3.3 Perspectives

From an optimistic viewpoint, the IDL indicator may improve the stability of the financial derivatives market. Our observation that previously introduced leading indicators did not provide an early warning for the Lehman Brothers bankruptcy, and the crisis that followed, is consistent with the hypothesis that leading indicators lose their predictive power in financial markets (185). A plausible explanation is that an increase of a known leading indicator could be directly followed by preemptive policy by central banks (186), a change of behavior of the market participants, or both, until the indicator returns to its normal level. This would imply that the financial system is capable of avoiding the type of critical transitions for which it has leading indicators: it changes behavior as it approaches such a transition, while it remains vulnerable to other critical transitions for which it has no indicators. The fact that the IDL indicator provides an early warning signal suggests that it is capable of detecting a type of transition for which the financial system had no indicators at the time. Therefore, from this viewpoint the IDL indicator potentially makes the financial system more resilient because it improves its capability of avoiding catastrophic changes.

From a pessimistic viewpoint, on the other hand, the IDL indicator may actually decrease the stability of the financial system. Upon an increase of IDL, participants may respond in a manner that increases the IDL further, reinforcing the participants' response, and so on, propelling the financial system towards a crisis. This is a general dichotomy for all early warning indicators in finance (187). In the absence of a mechanistic model of the financial derivatives market it is difficult to predict the effect of a warning indicator.

Our results are a marked step forward in the analysis of complex dynamical systems. The IDL is a generic indicator that may apply to any self-organizing system of coupled units. For many such systems we lack the mechanistic insight necessary to build models with sufficient predictive power. Remarkably, we find evidence that the percolation of information can provide a tell-tale of self-organized critical phenomena even in the absence of a descriptive model. Although we study the financial derivatives market here, it seems reasonable to expect that it is true for a wide range of systems such as the forming of opinions in social networks (76, 147–152), the extinction of species in ecosystems (145, 174, 175, 179, 188–191), phase transitions and spontaneous magnetization in physics (177, 192–194), robustness in biological systems (195, 196), and self-organization of populations of cells (197) and even software components (198).

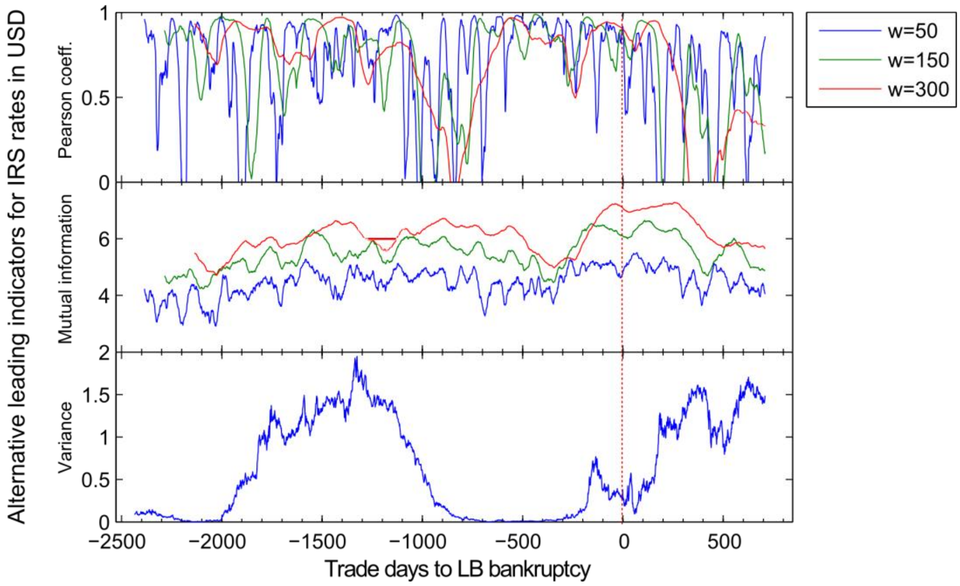

Figure 9: Alternative leading indicators for the IRS time series. We computed the average cross-maturity correlations for sliding window sizes of 50 days (blue line), 150 days (green line), and 300 days (red line) between the 1-year IRS and all other maturities. The variance at time t is computed of the rates of all 15 maturities at time t. Time point 0 on the horizontal axis corresponds to the day of the Lehman Brothers bankruptcy.